PSL UNIVERSITÉ PARIS

THÈSE DE DOCTORAT DE L'UNIVERSITÉ PSL

Préparée à l'Université Paris-Dauphine

Bootstrap percolation and kinetically constrained models: two-dimensional universality and beyond

Soutenue par
Ivailo HARTARSKY
Le 7 janvier 2022

École doctorale n°543
Ecole doctorale SDOSE

Spécialité
Mathématiques

Dauphine | PSL

Composition du jury :

Mme Oriane BLONDEL
Chargée de recherche CNRS,
Université Claude Bernard Lyon 1 — *Examinatrice*

M Hugo DUMINIL-COPIN
Professeur,
Institut des Hautes Etudes Scientifiques, Université de Genève — *Rapporteur*

M Aernout Coert Daniël VAN ENTER
Professeur émérite,
University of Groningen — *Examinateur*

Mme Alessandra FAGGIONATO
Professeure,
Università di Roma 'la Sapienza' — *Rapporteuse*

M Fabio MARTINELLI
Professeur,
Università degli Studi Roma Tre — *Examinateur*

M Justin SALEZ
Professeur,
Université Paris-Dauphine, PSL — *Président*

M Augusto TEIXEIRA
Professeur,
Instituto de Matemática Pura e Aplicada — *Examinateur*

Mme Cristina TONINELLI
Directrice de recherche CNRS,
Université Paris-Dauphine, PSL — *Directrice de thèse*

Contents

Acknowledgements

I list below a few of the people (and things) whom I owe thanks in relation to the present thesis. Naturally, I have to limit the list somehow, so I apologise in advance to those absent from it.

Professional Cristina Toninelli... for everything.

Fabio Martinelli for being like a second advisor to me.

The jury members and particularly the reviewers for accepting to handle the thesis, for their careful reading, interesting questions and thoughtful appreciation of the manuscript. Lyuben Lichev for numerous helpful comments on the presentation. Justin Salez for many useful remarks, particularly during the internal defence, and for his unceasing support over the years. Béatrice de Tilière and Antonin Chambolle for helpful information and suggestions.

Those who have de facto written parts of this thesis—Laure Marêché, Fabio, Tamás Mezei, Robert Morris, Réka Szabó and Cristina—but also my other collaborators for what they have taught me. A special mention for Cristina, Fabio and Laure for having been with me throughout the universality project. Beyond them, Aurelia Deshayes, Assaf Shapira, Réka and Erik Slivken for contributing to the friendly and cohesive team atmosphere. I present my excuses to Itai Benjamini, Giovanni Brigati, Barbara Dembin, Aurelia, Lyuben, Fabio, Réka, Vincent Tassion, Marius Tiba, Cristina and Ahad Zehmakan for eternally postponing projects, though I do intend to pay my debts... eventually.

Laëtitia Comminges and, more incidentally, Robin Ryder for teaching me to teach in very supportive conditions. LSO staff for making things work very smoothly without me having to understand them.

CEREMADE and particularly my officemates, Frank Alvarez Borges, Clément Berenfeld, Katharina Eichinger, Clément Moreau, Enguérand Petit for the nice working environment, albeit rarely visited. All Dauphine administrative and technical staff, Isabelle Bellier above all, for having helped me with some of the most difficult parts of the thesis. Enrolment was a large-scale collaborative project in which Yijun Wan's contribution is also warmly acknowledged. Similarly, Linda Arbane, Amirali Hannani, Valery and Vyara H. and Réka for help with the organisation of the defense.

The French state (via ENS via Dauphine) for funding the thesis and ERC MALIG for making life easier back when there was a life to make easier. EPFL, ETH, Roma Tre and Minu Balkanski foundation for hosting my research over short periods of time.

Personal My family for having been supportive of me during the thesis and long before it, naturally, Valery and Vyara above all. I specially thank Vyara for her headstrong interest in things beyond her understanding—an attitude I am yet to understand. She somehow regularly manages to teach me things she doesn't know.

My dearest friends in life outside mathematics, Nevena Eliseeva, Ivo Kortezov, Lyubo, Ganaël Roeland, Asiyat Shchamkhalova, Marius, Snezhana Veskova, for keeping me from doing mathematics more than advisable and... everything it is that friends do (and don't expect me to give a definition of that). Alongside them, my entire second family fathered by Minko Balkanski, by now too numerous to list. Especially, Plamen Bokov, Mihaela Dimitrova, Svilen Iskrov, Radoslav Kafov, Ivo, Lyubo, Anton Mladenov, Violeta Naydenova, Bogdan Stankov and Yanko Todorov for making organising summer schools and competitions a pleasure. This applies equally well to the Balkanski foundation non-pedagogical staff!

(Pre-)historic I insist on mentioning a few more people, without whom my career would never have gone in its present course. Starting with Veliko Kolev, the value of whose teachings I fully realised only recently, if ever; Minko, who enabled my coming to LLG (and France) and integrated me in the truly amazing second family mentioned above; Jean-François Logeais, who showed me that physics can be done properly; Giulio Biroli and Justin, who guided my first steps in research and probability; Raphaël Cerf, who directed me towards bootstrap percolation; Rob, who had the infinite patience needed to co-author my first mathematics paper, Chapter 10, (and I had much sharper views back then!); Dmitry Chelkak, who mentored me until passing me on to Cristina. And Cristina... for everything.

Résumé détaillé

Dans cette thèse nous étudions deux modèles issus de la physique statistique : la percolation bootstrap et les modèles cinétiquement contraints (KCM). Nous nous placerons principalement sur la grille $\mathbb{Z}^2$ avec condition initiale aléatoire, ce qui nous amènera à des considérations probabilistes et combinatoires. Néanmoins, il convient de signaler que ces modèles ont de nombreuses connections à d'autres domaines, aussi bien des mathématiques et de la physique que d'autres disciplines telles que l'informatique et les sciences sociales (voir Section 1.1).

1 Définition des modèles

1.1 Percolation bootstrap

Un modèle de percolation bootstrap est défini par un entier positif d (*dimension*) et une *famille de mise à jour* $\mathcal{U}$ qui est une famille finie non-vide de sous-ensembles finis non-vides de $\mathbb{Z}^d\setminus\{0\}$ appelés *règles de mise à jour*. Etant donné un ensemble $A \subset \mathbb{Z}^d$, on définit *l'action de bootstrap*

$$\mathscr{B}_\mathcal{U}(A) = A \cup \left\{x \in \mathbb{Z}^d : \exists U \in \mathcal{U}, x + U \subset A\right\}$$

avec $B + C = \{b + c : b \in B, c \in C\}$ et $b + C = \{b\} + C$ pour tous $B, C \subset \mathbb{R}^d$ et $b \in \mathbb{R}^d$. En mots, étant donné que l'ensemble A de *sites* est *infecté*, à la prochaine étape, on infecte de plus tout site tel qu'au moins une règle translatée en ce site est déjà entièrement infectée. Ce processus peut naturellement être vu comme un système dynamique en temps discret. Etant donné un ensemble de sites initialement infectés $A \subset \mathbb{Z}^d$, sa *clôture* est

$$[A]_\mathcal{U} = \bigcup_{n\in\mathbb{N}} \mathscr{B}^n_\mathcal{U}(A),$$

où $\mathscr{B}^n_\mathcal{U}$ signifie $\mathscr{B}_\mathcal{U}$ itéré n fois et $\mathbb{N} = \{0, 1, \ldots\}$ est l'ensemble des entiers naturels. Puisque $\mathcal{U}$ sera souvent fixé, nous l'omettons de toute notation, sauf si cela engendre de confusion. On dit que A est *stable* si $A = [A]$.

Pour donner un exemple, considérons le *modèle à r voisins*. Là on a $\mathscr{B}(A) = A \cup \{x \in \mathbb{Z}^d : |\{y \in A : y \sim x\}| \geqslant r\}$ avec $y \sim x$ si x et y sont voisins dans la structure de graph usuelle de $\mathbb{Z}^d$. Alternativement, on peut voir ceci comme $\mathscr{B}_{\mathcal{U}}(A)$ avec $\mathcal{U}$ la famille de $\binom{2d}{r}$ ensembles de r voisins de l'origine. C'est-à-dire si n'importe quels r voisins d'un site sont tous infectés, ce site le devient également.

Jusqu'alors le processus est entièrement déterministe. On introduit l'aléa seulement dans la condition initiale, en prenant $x \in A$ indépendamment pour chaque $x \in \mathbb{Z}^d$ avec probabilité q. On note la loi de A par μ_q et laisse q implicite lorsqu'il est clair du contexte. Habituellement en physique statistique, on peut à présent introduire la *probabilité critique*

$$q_{\mathrm{c}} = \inf\left\{q \in [0,1] : \mu_q\left([A] = \mathbb{Z}^d\right) > 0\right\}.$$

Puisque l'évènement ci-dessus est invariant par translation, l'ergodicité implique qu'en fait $\mu_q([A] = \mathbb{Z}^d) \in \{0,1\}$ pour tout $q \in [0,1]$. Tandis que pour $q < q_{\mathrm{c}}$ on peut étudier la géométrie de $\mathbb{Z}^d \backslash [A]$, dans le régime $q \geqslant q_{\mathrm{c}}$ on peut souhaiter être plus quantitatif. A cette fin on introduit le *temps d'infection* (de percolation bootstrap)

$$\tau_0^{\mathrm{BP}} = \min\{n \in \mathbb{N} : 0 \in \mathscr{B}^n(A)\} \in \mathbb{N} \cup \{\infty\},$$

qui est une variable aléatoire, en posant $\min \varnothing = \infty$ de manière usuelle.

1.2 KCM

Les KCM sont définis également par leurs dimension $d \geqslant 1$, famille de mise à jour $\mathcal{U}$ comme dans Section 1.1 et paramètre $q \in [0,1]$. Le $\mathcal{U}$-KCM est un processus de Markov à temps continu avec espace d'états $\Omega = \{0,1\}^{\mathbb{Z}^d}$, les zéros correspondant à des infections. On peut naturellement identifier toute configuration $\eta \in \Omega$ avec l'ensemble de ses infections (appelé A dans Section 1.1). La *contrainte en x* est donnée par

$$c_x(\eta) = \mathbb{1}_{\exists U \in \mathcal{U}, \eta_{x+U} = \mathbf{0}},$$

en écrivant η_X pour la restriction de $\eta \in \Omega$ à $X \subset \mathbb{Z}^d$ et $\mathbf{0}_X$ (resp. $\mathbf{1}_X$) pour la configuration entièrement infectée (resp. saine) sur X. On omet X de la notation s'il est clair du contexte. Si $X = \{x\}$ pour un $x \in \mathbb{Z}^d$, on écrit simplement x pour concision, de sorte que η_x est l'état du site $x \in \mathbb{Z}^d$ dans la configuration $\eta \in \Omega$. On écrit η^x pour la configuration obtenue en retournant l'état du site x dans la configuration η, c'est-à-dire $(\eta^x)_y = \eta_y$ pour tout $y \in \mathbb{Z}^d \backslash \{x\}$ et $(\eta^x)_x = 1 - \eta_x$.

On note $\mu_q = \mu$ la mesure produit de Bernoulli, telle que $\mu(\eta_x = 0) = q$, η étant une configuration aléatoire de loi μ. On note $\mu_X(f)$ l'espérance conditionnelle $\mu(f(\eta)|\eta_{\mathbb{Z}^d \backslash X})$ sur les états dans X pour tout ensemble fini $X \subset \mathbb{Z}^d$ et fonction réelle $f : \Omega \to \mathbb{R}$. Les variances par rapport à μ et μ_X sont notées Var et Var_X respectivement.

On peut maintenant définir le $\mathcal{U}$-KCM par l'action de son générateur sur les fonctions $f : \Omega \to \mathbb{R}$ dépendant seulement des états d'un nombre fini de sites

$$\begin{aligned}\mathcal{L}_{\mathcal{U}}(f)(\eta) &= \sum_{x\in\mathbb{Z}^d} c_x(\eta)\cdot(\mu_x(f)-f)(\eta)\\ &= \sum_{x\in\mathbb{Z}^d} c_x(\eta)\cdot(q\eta_x+(1-q)(1-\eta_x))\cdot(f(\eta^x)-f(\eta)),\end{aligned}$$

où μ_x est une abréviation de $\mu_{\{x\}}$ et de même pour Var_x. Ce processus peut également être défini par une construction graphique plus intuitive comme suit (voir [258] pour plus de contexte). Chaque site est muni d'un processus de Poisson standard (*horloge*) indépendant, dont les atomes sont appelés *sonneries*. Lorsque l'horloge en x sonne, on se donne de plus une variable aléatoire indépendante de loi (de Bernoulli) μ_x. Si la contrainte c_x est satisfaite, la mise à jour est dite *légale* et elle remplace l'état de x par celui de la variable. Dans le cas contraire (*mise à jour illégale*) la configuration reste inchangée après la sonnerie. Puisque le nombre de sonneries est localement fini et les contraintes ont une portée finie, ceci ne pose pas de problème [258]. Il est clair de chacune des deux définitions que la mesure μ est réversible et invariante pour le KCM, comme $c_x(\eta)$ ne dépend pas de η_x. Pour cette raison on appelle μ la *mesure d'équilibre*. Par contre, il faut remarquer que, par exemple, la mesure de Dirac sur la configuration entièrement saine $\mathbf{1}$ est invariante aussi.

La quantité d'intérêt principal est le temps (aléatoire) d'infection de l'origine

$$\tau_0 = \min\{t \geqslant 0 : \eta_0(t) = 0\},$$

$\eta(t)$ étant l'état du KCM en temps $t \in [0,\infty)$. On s'intéressera à $\mathbb{E}_\mu(\tau_0)$, l'espérance du temps d'infection pour le processus stationnaire (avec condition initiale de loi μ). Une quantité plus analytique, mais tout aussi importante, de la vitesse de la dynamique est son *temps de relaxation* T_{rel} défini comme l'inverse du trou spectral de $\mathcal{L}_{\mathcal{U}}$. Heureusement, on ne sera jamais amené à considérer le spectre d'opérateurs, grâce à la définition plus abordable

$$(T_{\mathrm{rel}})^{-1} = \inf_{f\not\equiv\mathrm{const}} \frac{\mathcal{D}_{\mathcal{U}}(f)}{\mathrm{Var}(f)}, \tag{1}$$

où $\mathcal{D}_{\mathcal{U}}$ est la forme de Dirichlet associée à $\mathcal{L}_{\mathcal{U}}$

$$\mathcal{D}_{\mathcal{U}}(f) = \sum_{x\in\mathbb{Z}^d} \mu(c_x\cdot\mathrm{Var}_x(f)) = -\mu(f\mathcal{L}_{\mathcal{U}}(f)).$$

2 Modèle à r voisins

Commençons par le modèle à r voisins provenant de [100, 155], qui est le plus classique à la fois en percolation bootstrap et KCM. Nous ne nous intéresserons qu'aux cas $r = 1$ et $d = r = 2$ (voir Section 1.4 pour les autres).

2.1 Modèle à un voisin

Le cas $r = 1$ étant trivial en percolation bootstrap, on se concentre à présent sur le KCM à un voisin appelé *FA-1f* en hommage à ses inventeurs, Fredrickson et Andersen [155]. Dans ce modèle lorsque $q \to 0$ le comportement typique d'une infection est le suivant. Avec taux q elle crée une autre infection à côté. Après, très rapidement l'une des deux infections disparait, résultant en un mouvement net de l'infection initiale correspondant à une marche aléatoire simple paresseuse. A des échelles de temps plus longues on peut également observer des branchements résultant de la création d'une troisième infection avant que la seconde ne soit détruite. En outre, deux infections effectuant leurs marches aléatoires qui arrivent à des sites voisins typiquement coalescent rapidement avant de pouvoir se déplacer. Ceci nous amène à introduire le modèle suivant au comportement identique, mais qui se prête bien plus facilement à l'analyse.

CBSEP

Soit $G = (V, E)$ un graphe connexe. Ses degrés minimal, maximal et moyen sont notés $d_{\min}, d_{\max}$ et d_{avg} respectivement. Le degré de $x \in V$ est noté d_x. Pour tout $\omega \in \Omega = \{0,1\}^V$ et sommet $x \in V$ on dit que x est *rempli/vide*, ou qu'il y a une *particule/trou* en x, si $\omega_x = 1/0$. On définit $\Omega_+ = \Omega \setminus \{0\}$ comme l'évènement qu'il existe au moins une particule. De même, pour une arête $e = \{x, y\} \in E$ on appelle $(\omega_x, \omega_y) \in \{0,1\}^{\{x,y\}}$ *l'état de* e dans ω et on écrit $E_e = \{\omega \in \Omega : \omega_x + \omega_y \neq 0\}$ pour l'évènement que e n'est pas vide.

Etant donné $p \in (0, 1)$, soit $\pi = \bigotimes_{x \in V} \pi_x$ la mesure produit de Bernoulli pour laquelle chaque sommet est rempli avec probabilité p et soit $\mu(\cdot) := \pi(\cdot|\Omega_+)$. Etant donné une arête $e = \{x, y\}$, on écrit $\pi_e := \pi_x \otimes \pi_y$ et $\lambda(p) := \pi(E_e) = p(2 - p)$.

Le processus d'exclusion simple avec branchement et coalescence (CBSEP) est une chaîne de Markov en temps continu sur Ω_+ pour laquelle l'état de chaque arête $e \in E$ telle que E_e arrive est remise à jour à taux 1 avec le mesure $\pi_e(\cdot|E_e)$. Ainsi, toute arête ayant exactement une particule déplace la particule à l'autre extrémité de l'arête (*mouvement d'échange*) à taux $(1 - p)/(2 - p)$ et crée une particule supplémentaire à son sommet vide (*mouvement branchant*) à taux $p/(2 - p)$. De plus, une arête contenant deux particules en tue une choisie uniformément (*mouvement coalescent*) à taux $2(1 - p)/(2 - p)$. Il est facile de voir que la chaîne est réversible par rapport à μ et ergodique sur Ω_+, comme elle peut rejoindre la configuration avec une particule dans chaque sommet. Si $c(\omega, \omega')$ est le taux de saut de ω à ω', la forme de Dirichlet $\mathcal{D}^{\mathrm{CBSEP}}(f)$ de la chaîne s'exprime comme

$$\mathcal{D}^{\mathrm{CBSEP}}(f) = \frac{1}{2}\sum_{\omega,\omega'}\mu(\omega)c(\omega,\omega')\left(f(\omega')-f(\omega)\right)^2 = \sum_{e\in E}\mu(\mathbb{1}_{E_e}\operatorname{Var}_e(f|E_e)).$$

Notons que les mouvements branchant et coalescent de CBSEP sont exactement ceux autorisés dans FA-1f. De surcroît, le mouvement d'échange pour l'arête $\{x,y\}$ de $(1,0)$ à $(0,1)$ peut être reconstruit, en utilisant deux mouvement successifs de FA-1f, le premier remplissant le trou en y et le second vidant x. Si l'on prend en compte aussi le taux de chaque mouvement, on obtient facilement la comparaison entre les formes de Dirichlet respectives : il existe une constante absolue $c > 0$ telle que pour tout $f : \Omega_+ \to \mathbb{R}$ on a

$$c^{-1}\mathcal{D}^{\mathrm{FA-1f}}(f) \leqslant \mathcal{D}^{\mathrm{CBSEP}}(f) \leqslant cd_{\max}p^{-1}\mathcal{D}^{\mathrm{FA-1f}}(f),$$

en posant le paramètre q de FA-1f égal au paramètre p de CBSEP. Dans notre application à FA-1f pour $p \to 0$ uniquement la borne supérieure, que nous croyons plus précise, comptera.

Malgré le fait que les deux modèles sont clairement étroitement liés, il convient de souligner que CBSEP a de nombreux avantages comparé à FA-1f, rendant son étude plus accessible. Particulièrement, CBSEP est *attractif* au sens qu'il existe un grand couplage (voir e.g. [256]) préservant l'ordre partiel sur Ω donné par $\omega \prec \omega'$ ssi $\omega_x \leqslant \omega'_x$ pour tout $x \in V$. De plus, on peut naturellement implanter dans CBSEP une marche aléatoire en temps continu $(W_t)_{t\geqslant 0}$ sur G telle que CBSEP a une particule en W_t pour tout $t \geqslant 0$. Cette propriété particulièrement fructueuse sera utilisée dans Section 3.5 et est difficile à reproduire pour FA-1f [59].

Dans Chapitre 3 nous établissons des bornes inférieures et supérieures sur le temps de relaxation et la constante de Sobolev logarithmique[1] de CBSEP sur des graphes finis arbitraires. Notre intérêt principal est pour le régime où $p = \Theta(1)/|V|$ et nos bornes sont exactes à corrections logarithmiques près pour de nombreux graphes usuels. Ces résultats entraînent des bornes supérieures correspondantes sur les temps de relaxation et la constante de Sobolev logarithmique du modèle FA-1f. A titre d'exemple, on obtient le résultat suivant, qui renforce des résultats de Pillai et Smith [299,300] sur le temps de mélange, tout en en fournissant une preuve plus directe.

Proposition 2.1. *Considérons FA-1f dans sa composante ergodique Ω_+ sur le tore $(\mathbb{Z}/n\mathbb{Z})^d$ avec sa structure de graphe usuelle et avec paramètre $q = \Theta(1/n^d)$. Sa constante de Sobolev logarithmique α satisfait*

$$\alpha^{-1} \leqslant O(1) \times \begin{cases} \log(1/q)/q^3 & d=1\\ \log^2(1/q)/q^2 & d=2\\ \log(1/q)/q^2 & d\geqslant 3.\end{cases}$$

[1] Cette constante est définie comme le trou spectral dans Eq. (1) avec $\operatorname{Var}(f)$ remplacé par l'entropie $\mu(f^2\log(f^2/\mu(f^2)))$.

Des bornes inférieures identiques à des facteurs logarithmiques près sont connues [88, 299], démontrant l'exactitude de ce résultat.

De plus, en utilisant Proposition 2.1, on établit des bornes similaires pour une généralisation de CBSEP avec $\{0,1\}$ remplacé par un espace d'états fini quelconque et sur le tore de taille arbitraire telle que $|V|q \to \infty$. Ces résultats démontrés dans Chapitre 3 et Section 5.B nous seront utiles par la suite pour l'étude du modèle FA-2f dans Chapitre 5.

2.2 Modèle à deux voisins

Percolation bootstrap à deux voisins

Passons à présent à la percolation bootstrap à deux voisins en deux dimensions, qui sera un prérequis pour l'étude de FA-2f.

Ce modèle a été très étudié, les premiers résultats rigoureux datant des années 80 [7, 358]. Les études les plus poussées [187, 190] ont établi

$$\exp\left(\frac{\pi^2 - O(\sqrt{q}\cdot\log^3(1/q))}{18q}\right) \leqslant \tau_0^{\mathrm{BP}} \leqslant \exp\left(\frac{\pi^2 - \Omega(\sqrt{q})}{18q}\right)$$

avec grande probabilité lorsque $q \to 0$. Nous renforçons la borne inférieure pour retrouver le second terme à un facteur multiplicatif borné près.

Théorème 2.2. *La percolation bootstrap à deux voisins en deux dimensions satisfait avec grande probabilité lorsque* $q \to 0$

$$\tau_0^{\mathrm{BP}} = \exp\left(\frac{\pi^2 - \Theta(\sqrt{q})}{18q}\right). \tag{2}$$

Ce théorème fait l'objet du Chapitre 10. En plus de raffinements des idées apportées par [7, 190, 225], nous sommes amenés à introduire un nombre d'avancées techniques aboutissant à une compréhension très fine de la croissance typique menant à la création d'une «gouttelette critique» (voir Section 1.4.3). Il convient de mentionner que Théorème 2.2 réfute des prédictions basées sur des simulations [337].

FA-2f

Considérons maintenant le KCM à deux voisins, FA-2f, toujours en deux dimensions. En raison de la difficulté de ce modèle la plupart des résultats sont restés heuristiques et seulement deux travaux [88, 272] ont fait un progrès rigoureux. Leur résultat est

$$\exp\left(\frac{\pi^2 - o(1)}{18q}\right) \leqslant \mathbb{E}_\mu(\tau_0) \leqslant \exp\left(\frac{O(\log^2(1/q))}{q}\right). \tag{3}$$

Malheureusement, ces bornes sont trop écartées pour discerner les bonnes prédictions non-rigoureuses. En effet, dans les derniers 35 ans de nombreuses conjectures conflictuelles sur l'asymptotique de $\mathbb{E}_\mu(\tau_0)$ se sont accumulées [86, 154, 157, 179, 180, 289, 306, 338, 345] et Eq. (3) ne permet pas de trancher cette controverse.

Notre résultat montre que seule la prédiction de [306, 345] est correcte. De plus, des conjectures plus fines sur le comportement à l'intérieur d'une gouttelette critique s'avèrent être justes dans [345], mais non pas dans [306].

Théorème 2.3. *Lorsque $q \to 0$, FA-2f à l'équilibre sur $\mathbb{Z}^2$ vérifie :*

$$\mathbb{E}_\mu(\tau_0) \geqslant \exp\left(\frac{\pi^2}{9q}\left(1-\sqrt{q}\cdot O(1)\right)\right), \tag{4}$$

$$\mathbb{E}_\mu(\tau_0) \leqslant \exp\left(\frac{\pi^2}{9q}\left(1+\sqrt{q}\cdot (\log(1/q))^{O(1)}\right)\right). \tag{5}$$

De plus, ces bornes valent aussi pour τ_0 avec grande probabilité.

Ceci constitue la première asymptotique exacte de $\log \mathbb{E}_\mu(\tau_0)$ dans toute la classe de KCM «critiques». Elle est établie dans Chapitre 5.

Remarque 2.4. Malgré les apparences, Théorème 2.3 n'est aucunement un corollaire du Théorème 2.2. Alors que la borne inférieure Eq. (4) découle assez directement d'Eq. (2), la preuve de la borne supérieure Eq. (5) est bien plus difficile. En particulier, elle demande d'intuiter un mécanisme efficace d'infection/guérison pour infecter l'origine sans analogue dans la percolation bootstrap à deux voisins, qui est monotone.

Heuristique L'intuition principale derrière Théorème 2.3 est que pour $q \to 0$ la relaxation à l'équilibre de FA-2f est dominée par le mouvement lent de groupements d'infection très improbables appelés *gouttelettes mobiles* ou simplement *gouttelettes*. Par analogie avec les *gouttelettes critiques* en percolation bootstrap, les gouttelettes mobiles ont une taille linéaire qui croît de manière polynomiale en q, i.e. elles vivent sur une échelle beaucoup plus petite que l'échelle métastable $e^{\Theta(1/q)}$ apparaissant en percolation bootstrap. L'une des considérations principales déterminant le choix d'échelle des gouttelettes mobiles est le fait que l'environnement d'infections typique autour de la gouttelette doit avec haute probabilité lui permettre de se déplacer sous la dynamique FA-2f en toute direction. Dans cette optique la contribution principale au temps d'infection de l'origine de FA-2f stationnaire devrait venir du temps qu'il faut pour que la gouttelette atteigne l'origine.

Afin de transformer cette intuition en preuve, on s'affronte à deux problèmes fondamentaux :

(1) une définition des gouttelettes mobiles précise, mais maniable ;

(2) un modèle efficace de leur évolution aléatoire «effective».

Dans [269, 272] et Chapitre 4 les gouttelettes mobiles (dites régions «super-bonnes» là) ont été définies de manière plutôt rigide comme des régions complètement infectées de forme et taille appropriée et leur mouvement a été modélisé comme un *FA-1f généralisé* sur $\mathbb{Z}$ [269, Section 3.1]. Dans ce dernier les gouttelettes sont librement créées et détruites avec les taux d'équilibre, mais *uniquement aux positions voisines d'une gouttelette déjà existante*.

Même si elle est puissante et robuste, cette approche ne peut donner accès à l'asymptotique *exacte* ni de (1), ni de (2) ci-dessus. En effet, une gouttelette mobile doit être autorisée à se déformer et se déplacer dans son voisinage comme une amibe, en réagençant ses infections, en utilisant les mouvements FA-2f. Ce mouvement «à l'amibe» entre positions voisines devra s'effectuer à une échelle de temps très inférieure à l'échelle globale nécessaire pour amener une gouttelette de très loin à l'origine. En particulier, il ne nécessite pas la création préalable d'une nouvelle gouttelette à partir de l'originelle pour détruire l'initiale seulement après (le mécanisme principal de la dynamique des gouttelettes sous le processus FA-1f généralisé).

Avec cette image à l'esprit, on propose une nouvelle solution à (1) et (2) ci-dessus qui donne accès à l'asymptotique exacte du temps d'infection. Concernant (1), notre raisonnement dans Section 5.2 se fait en deux étapes. D'abord on propose une définition multi-échelle sophistiquée des gouttelettes mobiles. En particulier, elle introduit un degré essentiel de *mollesse* dans la configuration microscopique d'infections.[2] La seconde étape est techniquement beaucoup plus difficile et consiste à développer les outils nécessaires pour l'analyse de la dynamique FA-2f à l'intérieur d'une gouttelette mobile. Notamment, nous démontrons deux propriétés clé (voir Propositions 5.2.7 et 5.2.9) :

(1.a) à l'ordre principal la probabilité ρ_{D} d'une gouttelette mobile est celle de percolation bootstrap :

$$\rho_{\mathrm{D}} \geqslant \exp\left(-\frac{\pi^2}{9q} - \frac{O(\log^2(1/q))}{\sqrt{q}}\right),$$

(1.b) le mouvement local «à l'amibe» des gouttelettes s'effectue à l'échelle de temps $\exp(O(\log(1/q)^3)/\sqrt{q})$ ce qui est sous-dominant par rapport à l'échelle de temps principale du problème et ne se manifeste que dans le second terme d'Eq. (5).

Propriété (1.a) découle relativement facilement de techniques de percolation bootstrap, tandis que la preuve de propriété (1.b), l'un des pas cruciaux de cette thèse, demande une quantité substantielle de nouvelles idées.

[2] Cette construction est inspirée par celle suggérée par P. Balister en 2017, qu'il conjecturait enlever les corrections logarithmiques redondantes dans Eq. (3) disponible à l'époque.

Alors que les propriétés (1.a) et (1.b) ci-dessus sont essentielles, elles ne suffisent pas à elles seules pour résoudre le problème (2) ci-dessus. En Section 5.4 nous proposons une modélisation (quoique seulement au niveau d'une inégalité de Poincaré, suffisante pour nos besoins) de l'évolution des gouttelettes mobiles par un CBSEP généralisé (mentionné dans Section 2.1), étudié dans Chapitre 3 à cette fin. Enfin, le fait que CBSEP relaxe à une échelle de temps proportionnelle à l'inverse de la densité des gouttelettes mobiles (modulo corrections logarithmiques – voir Proposition 5.4.1) donne l'asymptotique du Théorème 2.3. Nous tenons à insister sur le fait que la modélisation du mouvement à grande échelle des gouttelettes par un CBSEP généralisé plutôt qu'un FA-1f généralisé est une nouveauté absolue, y compris par rapport à la littérature de physique.

3 Universalité grossière

Lorsqu'on parle d'universalité pour les modèles de percolation bootstrap (resp. KCM) notre but ultime est de pouvoir répartir toutes les familles $\mathcal{U}$ en classes des sorte à ce que les représentants de chaque classe aient un comportement identique si l'on les regarde de loin. Ici par comportement identique on entendra notamment avoir des $\mu\left(\tau_0^{\mathrm{BP}}\right)$ (resp. $\mathbb{E}_\mu(\tau_0)$) de grandeur similaire lorsque $q \to 0$. Bien entendu, une telle classification est satisfaisante seulement si elle permet, étant donnée une famille de mise à jour $\mathcal{U}$, de pouvoir déterminer à quelle classe elle appartient seulement à partir de sa géométrie et combinatoire. Nous nous placerons exclusivement en deux dimensions, même si la quasi-totalité de nos arguments ne est pas restreinte à ce cas. Laissant l'histoire de l'universalité à Section 1.5.1, nous procédons directement à la définition des classes d'universalité grossière.

Soit $\|\cdot\|$ et $\langle\cdot,\cdot\rangle$ les norme et produit scalaire Euclidiens respectivement. Soit $S^1 = \{u \in \mathbb{R}^2 : \|u\| = 1\}$ le cercle unité que l'on identifie à $\mathbb{R}/2\pi\mathbb{Z}$ lorsque nécessaire. On appelle ses éléments *directions*. Les demi-plans ouvert et fermé de normale extérieure $u \in S^1$ et décalage $l \in \mathbb{R}$ sont

$$\mathbb{H}_u(l) = \left\{x \in \mathbb{R}^2 : \langle x, u\rangle < l\right\}, \qquad \overline{\mathbb{H}}_u(l) = \left\{x \in \mathbb{R}^2 : \langle x, u\rangle \leqslant l\right\}.$$

On omet l quand il est égal à 0. Sauf si cela engendre de confusion, on identifie tout sous-ensemble de $\mathbb{R}^2$, tel que $\mathbb{H}_u$, avec son intersection avec $\mathbb{Z}^2$.

Définition 3.1 (Direction stable)**.** Fixons une famille de mise à jour $\mathcal{U}$. Une direction $u \in S^1$ est *instable* s'il existe $U \in \mathcal{U}$ tel que $U \subset \mathbb{H}_u$ et *stable* sinon.

L'intérêt de cette définition vient du fait que $[\mathbb{H}_u]_{\mathcal{U}} = \mathbb{H}_u$ si u est stable (i.e. u est stable ssi $\mathbb{H}_u$ l'est) et $[\mathbb{H}_u]_{\mathcal{U}} = \mathbb{Z}^2$ si u est instable. On dit qu'une direction $u \in S^1$ est *rationnelle* si $\mathbb{R}u \cap \mathbb{Z}^2 \neq \varnothing$. Il n'est pas difficile de vérifier que l'ensemble de directions stables est une union finie d'intervalles fermés de S^1 aux extrémités rationnelles. Les extrémités d'intervalles de directions stables sont appelées *isolées* si l'intervalle est réduit à un point et *semi-isolées* sinon. Toute direction stable qui n'est ni isolée, ni semi-isolée est dite *fortement stable*. Pour des exemples illustratifs voir Fig. 1.2.

On est à présent en mesure de définir les classes d'universalité grossière.

Définition 3.2 (Partition d'universalité grossière)**.** Soit $\mathcal{C} = \{\mathbb{H}_u \cap S^1 : u \in S^1\}$ l'ensemble des demi-cercles ouverts de S^1. Une famille $\mathcal{U}$ est :

- *surcritique* s'il existe un demi-cercle ouvert $C \in \mathcal{C}$ dont toutes les directions $u \in C$ sont instables. Si de plus
  - il existe deux directions stables non opposées, $\mathcal{U}$ est *enraciné* ;
  - il n'existe pas deux directions stables non opposées, $\mathcal{U}$ est *déraciné*.
- *critique* si tout demi-cercle ouvert contient une direction stable et il existe un demi-cercle contenant un nombre fini de directions stables.
- *souscritique* si tout demi-cercle contient une infinité de directions stables. Elle est
  - *non triviale* s'il existe une direction instable ;
  - *triviale* si toutes les directions sont stables.

L'intérêt de ces classes devient apparent avec le résultat suivant résumé dans Tableau 1.2.

Théorème 3.3 (Universalité grossière [28, 74, 265, 267, 269])**.** *Soit $\mathcal{U}$ une famille de mise à jour bidimensionnelle. Si $\mathcal{U}$ est*

- *surcritique déraciné, alors $q_c = 0$, $\mu\left(\tau_0^{\mathrm{BP}}\right) = q^{-\Theta(1)}$ et $\mathbb{E}_\mu(\tau_0) = q^{-\Theta(1)}$.*
- *surcritique enraciné, alors $q_c = 0$, $\mu\left(\tau_0^{\mathrm{BP}}\right) = q^{-\Theta(1)}$ et $\mathbb{E}_\mu(\tau_0) = \exp(\Theta(\log^2(1/q)))$.*
- *critique, alors $q_c = 0$, $\mu\left(\tau_0^{\mathrm{BP}}\right) = \exp(q^{-\Theta(1)})$ et $\mathbb{E}_\mu(\tau_0) = \exp(q^{-\Theta(1)})$.*
- *souscritique non triviale, alors $q_c \in (0,1)$, $\mu\left(\tau_0^{\mathrm{BP}}\right) = \mathbb{E}_\mu(\tau_0) = \infty$ pour q assez petit.*
- *souscritique triviale, alors $q_c = 1$, $\mu\left(\tau_0^{\mathrm{BP}}\right) = \mathbb{E}_\mu(\tau_0) = \infty$ pour tout $q \in (0,1)$.*

Les mêmes asymptotiques s'appliquent à τ_0^{BP} et τ_0 avec grande probabilité et à T_{rel} lorsque $q \to 0$.

Nous reviendrons à des raffinements du résultat assez grossier sur les modèles critiques dans Section 4 et aux souscritiques dans Section 5.

4 Universalité raffinée des modèles critiques

4.1 Classes d'universalité raffinée

Pour pouvoir déterminer l'asymptotique de $\mu\left(\tau_0^{\mathrm{BP}}\right)$ et $\mathbb{E}_\mu(\tau_0)$ plus précisément que ce que nous fournit Théorème 3.3 pour les modèles critiques, on a besoin d'affiner notre partition et la notion de direction stable (voir Fig. 1.2 pour des exemples).

Définition 4.1 (Difficulté)**.** La *difficulté* $\alpha(u)$ de $u \in S^1$ est

- 0 si u est instable ;
- ∞ si u est stable, mais pas isolé ;
- $\min\{n : \exists Z \subset \mathbb{Z}^2, |Z| = n, |[\mathbb{H}_u \cup Z]_\mathcal{U} \backslash \mathbb{H}_u| = \infty\}$ sinon.

La *difficulté* de $\mathcal{U}$ est

$$\alpha = \alpha(\mathcal{U}) = \min_{C \in \mathcal{C}} \max_{u \in C} \alpha(u).$$

On dit qu'une direction $u \in S^1$ est *difficile* si $\alpha(u) > \alpha$.

Les difficultés affinent non seulement la notion de direction stable, mais aussi celle des classes d'universalité grossière. Plus précisément, il n'est pas difficile de vérifier qu'un modèle est surcritique ssi sa difficulté α est 0 ; critique ssi α est un entier strictement positif ; souscritique ssi $\alpha = \infty$.

Notons de plus que la tâche de déterminer les directions stables ou la classe d'universalité grossière d'une famille de mise à jour est facile, tandis que calculer les difficultés des directions stables ou la difficulté globale d'une famille critique est moins immédiat. Nous examinons cette question du point de vue de la complexité dans Chapitre 9, en démontrant qu'il est possible de calculer α en temps fini, étant donné $\mathcal{U}$, mais il est NP-difficile de le faire. Notre algorithme en temps fini repose sur des bornes quantitatives sur la distance que peut parcourir l'infection d'un ensemble initial de α infections ajouté à un demi-plan, si elle ne s'étend pas à l'infini. Le résultat de NP-difficulté, quant à lui, découle de l'immersion d'instances du problème classique de recouvrement d'ensemble. Malgré la difficulté pour déterminer les difficultés, on les considérera données par la suite.

Avec Définition 4.1 à notre disposition, on peut définir toutes les notions qui apparaitront dans la partition d'universalité raffinée.

Définition 4.2 (Types raffinés)**.** Une famille de mise à jour bidimensionnelle est

- *enracinée* s'il existe deux directions difficiles non opposées ;
- *déracinée* si elle n'est pas enracinée ;

- *déséquilibrée* s'il existe deux directions difficiles opposées ;
- *équilibrée* si elle n'est pas déséquilibrée, soit, il existe un demi-cercle *fermé* ne contenant aucune direction difficile.

On partitionne les familles équilibrées déracinées davantage en

- *semi-dirigées* s'il y a exactement une direction difficile ;
- *isotropes* s'il n'y a aucune direction difficile.

On sera amené à considérer de plus la distinction entre modèles à nombre fini ou infini de directions stables (soit, sans ou avec une direction fortement stable). Les derniers sont nécessairement enracinés, mais peuvent être équilibrés ou non. Ainsi on se retrouve avec une partition de toutes les familles de mise à jour critiques en deux dimensions en sept classes représentées en Fig. 1.2 que nous étudions par la suite.

4.2 Résultats d'universalité raffinée

L'universalité raffinée en percolation bootstrap étant connue de [70], on se focalise sur les KCM qui ont un comportement beaucoup plus riche et intriqué. En effet, nos résultats excluent des conjectures faites dans [269] sur l'asymptotique de $\mathbb{E}_\mu(\tau_0)$. Le résultat principal de la thèse résumé dans Tableau 1.3 s'énonce ainsi.

Théorème 4.3 (Universalité raffinée des KCM)**.** *Soit $\mathcal{U}$ une famille de mise à jour critique en deux dimensions avec difficulté α. Si $\mathcal{U}$ est*

(a) déséquilibré avec nombre infini de directions stables (donc enraciné), alors

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta\left((\log(1/q))^4\right)}{q^{2\alpha}}\right);$$

(b) équilibré avec nombre infini de directions stables (donc enraciné), alors

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta(1)}{q^{2\alpha}}\right);$$

(c) déséquilibré enraciné avec nombre fini de directions stables, alors

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta\left((\log(1/q))^3\right)}{q^{\alpha}}\right);$$

(d) *déséquilibré déraciné (donc avec nombre fini de directions stables), alors*

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta\left((\log(1/q))^2\right)}{q^\alpha}\right);$$

(e) *équilibré enraciné avec nombre fini de directions stables, alors*

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta\left(\log(1/q)\right)}{q^\alpha}\right);$$

(f) *semi-dirigé (donc équilibré déraciné avec nombre fini de directions stables), alors*

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta\left(\log\log(1/q)\right)}{q^\alpha}\right);$$

(g) *isotrope (donc équilibré déraciné avec nombre fini de directions stables), alors*

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta(1)}{q^\alpha}\right).$$

Il convient de souligner que, contrairement à ce qui était le cas en percolation bootstrap, ce résultat est l'état de l'art pour toute famille critique à l'exception de FA-2f, pour lequel l'unique résultat de précision supérieure (ou égale) est Théorème 2.3. Le seul KCM pour lequel le résultat fourni par Théorème 4.3 était connu précédemment est le celui de Duarte [267, 269].

La preuve du Théorème 4.3 a été effectuée en plusieurs étapes. Notamment, la borne inférieure à corrections logarithmiques près pour les modèles avec nombre fini de directions stables découle de [88] et résultats de percolation bootstrap [70]. En fait, il en est de même pour les bornes inférieures exactes du Théorème 4.3 pour les classes (d) et (g). La borne inférieure à corrections logarithmiques près pour les familles à nombre infini de directions stables est établie dans Chapitre 7 et donne de plus la borne inférieure exacte pour la classe (b). Les techniques en sont repris et développées davantage dans Chapitre 8 pour démontrer toutes les bornes inférieures du Théorème 4.3 dans un cadre unificateur.

Quant aux bornes supérieures, celles du Théorème 1.6.4 pour la classe (a) et à corrections logarithmiques près pour toutes les familles avec nombre infini de directions stables ont été obtenues dans [269]. La borne supérieure à corrections logarithmiques près pour les familles avec nombre fini de directions stables ainsi que la borne supérieure exacte pour la classe (c) sont prouvées dans Chapitre 4. Les bornes supérieures exactes restantes du Théorème 4.3 sont démontrées dans Chapitre 6, les classes (e) et (f) étant le défi le plus important.

Chacun des Chapitres 4 et 6-8 apporte de nombreuses nouvelles idées et techniques dont la présentation est laissée à Section 1.6 et aux chapitres correspondants. Nous ne signalons ci-dessous qu'une seule de ces avancées qui est particulièrement intéressante de manière indépendante et fait l'objet du Chapitre 2.

4.3 Dynamique microscopique et KCM inhomogènes

La démonstration de certaines des bornes supérieures du Théorème 4.3 dans Chapitre 6 fait apparaitre le problème suivant. On souhaite pouvoir remettre à l'équilibre l'état des sites sur la frontière d'une grande région infectée de forme polygonale. Vu de près, ceci revient à étudier la dynamique du KCM bidimensionnel d'origine restreint à un segment unidimensionnel avec une condition au bord dans le reste du plan. Quitte à identifier le segment à un sous-ensemble de $\mathbb{Z}$, ceci nous amène au cadre suivant (voir Section 2.1 pour les définitions formelles). Considérons un KCM

- sur un volume arbitraire $L \subset \mathbb{Z}$, $1 \leqslant |L| \leqslant \infty$, qui n'a pas besoin d'être un intervalle ;
- avec conditions au bord arbitraires dans $\mathbb{Z}$;
- conditionné à appartenir à une composante irréductible de l'espace d'états arbitrairement choisie ;
- avec espaces d'états finis arbitraires pour chaque site, qui peuvent différer d'un site à l'autre et n'ont pas besoin d'être uniformément bornés en taille ou probabilités des atomes, mais la probabilité d'être infecté est bornée inférieurement par $q > 0$;
- avec des familles de mise à jour qui peuvent varier d'un site à l'autre, pourvu qu'elles aient une portée uniformément bornée par $R < \infty$. Certains sites peuvent être complétement figés ou, inversement, libres de changer d'état sans contrainte.

Dans Chapitre 2 nous prouvons que dans ce cadre pour un certain $C_R > 0$ dépendant seulement de R

$$T_{\mathrm{rel}} \leqslant (2/q)^{C_R \log(\min(2/q,|L|))}.$$

Comme on l'a vu dans Théorème 3.3, cette borne est la meilleure possible pour tout KCM unidimensionnel homogène surcritique enraciné aux états binaires lorsque $q \to 0$. De manière surprenante, elle est également la meilleure possible pour certains KCM binaires homogènes surcritiques déracinés sur des intervalles, malgré le fait que sur $\mathbb{Z}$ leur temps de relaxation est seulement $q^{-\Theta(1)}$ d'après Théorème 3.3.

Notons que de tels KCM généraux possèdent souvent de nombreuses composantes irréductibles (il y en a toujours au moins deux sauf trivialités) et leur structure combinatoire peut être très intriquée. Elles se sont avérées difficiles à traiter à cause des dépendances de longue portée qu'elles induisent. Par conséquent, l'unique cas non trivial où le temps de relaxation sur une composante irréductible est contrôlé [59] (voir aussi [88,89]) est FA-1f sur un intervalle dans sa composante ergodique – la seule composante irréductible non triviale de ce KCM, donnée par $\Omega\backslash\{\mathbf{1}\}$.

Le lecteur peut consulter [321, 324, 326] pour des KCM inhomogènes, [88,89,342] pour des KCM homogènes avec diverses règles et conditions aux bord et [269] et Chapitre 3 pour des espaces d'état généraux. Néanmoins, il convient de souligner qu'aucune paire d'aspects parmi les suivants n'a été considérée simultanément par le passé : espaces d'état généraux, règles inhomogènes, composantes irréductibles. Formellement, comme on le verra, les domaines non-intervalles, les conditions au bord et les composantes irréductibles autres que l'ergodique peuvent être absorbés dans l'inhomogénéité des règles, mais de tels KCM arbitrairement inhomogènes n'ont jamais été considérés avant.

5 Modèles souscritiques

Pour finir, abordons les familles souscritiques, qui sont le moins bien comprises, à la fois en percolation bootstrap et en KCM. Pour simplifier le propos on exclue d'emblée les modèles souscritiques triviaux et dit simplement «souscritique» pour un modèle souscritique non trivial.

5.1 Transition de phase et percolation orientée

Avant de nous occuper des modèles souscritiques généraux, concentrons-nous sur le premier et le plus fondamental parmi eux. Il est défini par $\mathcal{U}^{\mathrm{OP}} = \{\{(1,0),(0,1)\}\}$ (voir Fig. 1.2b) et on l'appellera *percolation orientée* (OP) pour de raisons qui s'éclairciront par la suite, tandis que sa version KCM est connue sous le nom Nord-Est (NE). En percolation bootstrap OP a été considéré en premier dans [313] et NE – dans [307] immédiatement après. Néanmoins, OP était déjà très bien étudié en tant que modèle de percolation (voir [131, 258], ainsi que Chapitre 11).

L'équivalence entre percolation bootstrap avec la famille de mise à jour ci-dessus et OP est la suivante (observée dans [313, 315]). Un site $x \in \mathbb{Z}^2$ devient infecté en temps t ssi le chemin initialement sain le plus long, faisant des pas vers le haut ou vers la droite, commençant en x, a longueur t. En particulier, $\tau_0^{\mathrm{BP}} = \infty$ ssi l'origine appartient à un tel chemin infini sain. Ainsi, entre autres, il est bien connu que $0 < q_{\mathrm{c}}(\mathcal{U}^{\mathrm{OP}}) < 1$. Schonmann a noté aussi que ceci implique directement $q_{\mathrm{c}}(\mathcal{U}) < 1$ pour tout $\mathcal{U} \supset \mathcal{U}^{\mathrm{OP}}$.

Cette observation se généralise à tous les modèles souscritiques comme suit. Par définition toute famille souscritique (non triviale) $\mathcal{U}$ a une direction instable $u \in S^1$. Alors, il existe une règle de mise à jour $U \in \mathcal{U}$ contenue dans le demi-plan ouvert correspondant $\mathbb{H}_u$. On appelle la percolation bootstrap associée à la famille (à une règle) $\{U\}$ une *percolation orientée par sites généralisée* (GOSP) pour tout U contenu dans un demi-plan ouvert tel que U engendre $\mathbb{R}^2$ en tant qu'espace vectoriel (sinon $\mathcal{U}$ serait surcritique). Comme le nom le suggère, GOSP se comporte comme OP et respecte la même relation d'équivalence avec une représentation par percolation de chemins aux pas dans U. Il est alors facile de voir que $q_c < 1$ pour GOSP (par exemple par comparaison avec un processus de branchement) et donc aussi pour le modèle d'origine souscritique $\mathcal{U}$.

L'inégalité complémentaire, $q_c > 0$ pour les modèles souscritiques est sensiblement plus difficile et constitue la difficulté principale du Théorème 3.3 pour cette classe. Elle a été résolue dans [28] via une renormalisation multi-échelle assez technique.

Pour conclure, notons que pour les KCM souscritiques en général, essentiellement rien n'est connu de plus que ce qui relève de la percolation bootstrap. Plus précisémment l'intégralité des études se limite au modèle NE [88,108,246,355] très spécial. Dans le cadre général, les uniques résultats pour les KCM souscritiques sont ceux valides pour tout KCM (souscritique ou non) :

- La transition d'ergodicité/mélange du $\mathcal{U}$-KCM intervient à $q_c(\mathcal{U})$ de la $\mathcal{U}$-percolation bootstrap correspondante, c'est-à-dire la probabilité critique de $\mu(\tau_0^{\mathrm{BP}} = \infty)$. Ceci est démontré dans [88].
- La transition d'annulation du trou spectral du $\mathcal{U}$-KCM intervient à $\tilde{q}_c(\mathcal{U})$ de la $\mathcal{U}$-percolation bootstrap, c'est-à-dire la probabilité critique de décroissance exponentielle de $\mu(\tau_0^{\mathrm{BP}} > t)$ (voir Chapitre 12). Ceci est démontré dans Théorème 12.3.7.

En vue de ceci, on restreint notre attention par la suite à la percolation bootstrap avec famille souscritique non triviale.

5.2 De GOSP aux modèles de percolation bootstrap souscritiques généraux

GOSP

La discussion ci-dessus de la non trivialité de la transition de phase des modèles souscritiques fournit au moins deux raisons d'étudier GOSP en détail. D'abord, ce sont les modèles souscritiques les plus simples et donc un point de départ convenable. Deuxièmement, la compréhension de GOSP peut être répercutée vers les familles générales. De plus, GOSP est intéressante par elle-même en tant que modèle de percolation et automate cellulaire probabiliste (voir Section 11.2). Pour ces raisons, dans Chapitre 11, on étudie GOSP en dimension arbitraire $d \geqslant 2$, en nous concentrant sur la phase $q < q_c$, soit la *surcritique* en langage de percolation.

Les résultats du Chapitre 11 que nous énonçons informellement par la suite peuvent heureusement être vues comme des boîtes noires. Pour commencer, en employant le langage de percolation, on dit que $a \in \mathbb{Z}^2$ est connecté à $b \in \mathbb{Z}^2$ s'il existe un chemin sain de a à b avec pas dans la règle $U \subset \mathbb{H}_u$ qui définit notre GOSP pour un $u \in S^1$. Considérons GOSP restreint à $\overline{\mathbb{H}}_v$ pour un certain $v \in S^1$, c'est-à-dire, les chemins doivent être contenus dans $\overline{\mathbb{H}}_v$. On s'intéresse si 0 est connecté à l'infini avec probabilité positive en fonction de q et v. Il s'avère que l'ensemble de directions v telles que ceci arrive est un intervalle qui varie avec q d'une manière continue et strictement monotone pour $q \in [0, q_c)$. A $q \to q_c-$ l'intervalle converge vers un demi-cercle et pour $q \geqslant q_c$ il est vide. De plus, pour v hors de la clôture (topologique) de cet intervalle, la longueur du chemin le plus long depuis 0 (qui est p.s. fini par définition) a une queue qui décroît exponentiellement.

Approche directionnelle aux modèles souscritiques

Pour pouvoir faire un usage plus sophistiqué de GOSP que la simple comparaison $q_c(\mathcal{U}') \leqslant q_c(\mathcal{U})$ lorsque $\mathcal{U}' \supset \mathcal{U}$, on aura besoin d'une décomposition directionnelle de q_c (ou plutôt $\tilde{q}_c$, la probabilité de décroissance exponentielle de $\mu(\tau_0^{\mathrm{BP}} > t)$). A cette fin on introduit la notion suivante, dont la définition précise est laissée au Chapitre 12. La *densité critique* de $u \in S^1$ (pour une famille de mise à jour $\mathcal{U}$) est moralement

$$d_u \approx \inf \{q \in [0,1] : \mu(0 \notin [A \cup \mathbb{H}_u]_{\mathcal{U}}) = 0\}.$$

Le lecteur attentif aura remarqué que ceci est exactement la notion que nous venons de considérer pour GOSP. Ainsi, on considérera que les densités critiques de GOSP sont des fonctions du cercle bien comprises, certes non explicites. Il est clair aussi que $d_u = 0$ pour toute direction instable ou stable isolée u, donc cette notion est adaptée aux directions fortement stables.

Avec les densités critiques à notre disposition, le résultat central du Chapitre 12 s'énonce

$$\tilde{q}_c = \inf_{C \in \mathcal{C}'} \sup_{u \in C} d_u, \tag{6}$$

où $\mathcal{C}'$ est l'ensemble des demi-cercles fermés. Pour transformer Eq. (6) en une version raffinée de la comparaison basique $\tilde{q}_c(\mathcal{U}') \leqslant \tilde{q}_c(\mathcal{U})$ pour $\mathcal{U}' \supset \mathcal{U}$, il suffit d'observer que, de même, $d_u(\mathcal{U}') \leqslant d_u(\mathcal{U})$ sous la même condition. Ainsi,

$$\tilde{q}_c(\mathcal{U}) \leqslant \inf_{C \in \mathcal{C}'} \sup_{u \in C} \min_{U \in \mathcal{U}} d_u(\{U\}),$$

ce qui nous permet de transférer des bornes sur les densités critiques de GOSP à des modèles souscritiques arbitraires (les densités critiques des familles non surcritiques à une règle qui ne sont pas des GOSP sont identiquement 1). Dans Chapitre 12 nous illustrerons que ceci donne effectivement de meilleures bornes dans des situations génériques.

Dans le cadre des densités critiques nous retrouvons aussi des résultats connus sur la famille Spiral de [344], basés sur une application moins directe d'Eq. (6). Spiral est essentiellement l'unique modèle de percolation bootstrap souscritique autre que GOSP, qui est relativement bien compris, grâce à ses connections fortes à OP (voir aussi [235, 236, 346, 347] pour d'autres modèles étroitement liés). Il est particulièrement intéressant en vue de la discontinuité de sa transition de phase [344, 346] : il satisfait

$$\mu_{q_c}\left(\tau_0^{\mathrm{BP}} = \infty\right) > 0,$$

comme attendu pour la transition de blocage dans des milieux granulaires [320].

Nous laissons les résultats supplémentaires concernant la décroissance exponentielle (en particulier redémontrant des résultats de [315]), sensitivité au bruit et réponses à certaines questions de [28] au Chapitre 12.

6 Organisation

La thèse est structurée comme suit. Hormis l'introduction indispensable, chaque chapitre est basé sur un papier différent parmi [207, 209, 210, 213–216, 218–220, 222]. Les chapitres sont regroupés dans des parties relativement indépendantes appropriées pour des lecteurs de différentes communautés ou goût.

Partie I (Chapitres 2-6) contient des bornes supérieurs sur des KCM. Elle est appropriée pour les lecteurs compétents en chaînes de Markov et, plus spécifiquement, la dynamique de systèmes de particules en interaction.

Partie II (Chapitres 7-10) présente des bornes inférieures sur des KCM critiques (Chapitres 7 et 8) suivies de résultats de percolation bootstrap critique. Tous deux sont adéquats pour les lecteurs habiles en combinatoire (probabiliste).

Partie III (Chapitres 11 et 12) traite la percolation bootstrap souscritique, ce qui la rend appropriée pour les lecteurs expérimentés en percolation.

Le contenu de chaque chapitre est le suivant.

Chapitre 1 est une introduction générale à la thèse ainsi qu'à la percolation bootstrap et KCM. On s'y efforce d'être pédagogique et complet. Il inclue un survol de la littérature et contexte de la thèse, ainsi que de l'histoire des résultats que nous présentons et ceux étroitement liés. Cette introduction est supposée connue dans tous les autres chapitres.

Chapitre 2 [207] traite les KCM généraux inhomogènes arbitraires sur volumes finis ou infinis unidimensionnels conditionnés à une composante irréductible (voir Section 4.3). Nous établissons une borne sur leur temps de relaxation sans imposer de condition, en développant une approche de bissection révisée.

Chapitre 3 [218] coécrit avec Fabio Martinelli et Cristina Toninelli étudie CBSEP (voir Section 2.1) et une généralisation naturelle sur des graphes finis arbitraires. Nous démontrons des bornes sur les temps de relaxation, mélange et Sobolev logarithmique en toute généralité, qui sont souvent exactes à corrections logarithmiques près. On s'intéresse particulièrement à la limite où le nombre de sommets du graphe diverge comme l'inverse de la densité de particules à l'équilibre q et aux applications à FA-1f, en particulier retrouvant des résultats de Pillai–Smith [299, 300].

Chapitre 4 [216] coécrit avec Fabio Martinelli et Cristina Toninelli montre la borne supérieure du Théorème 4.3 à corrections logarithmiques près pour les modèles critiques à nombre fini de directions stables (voir Section 4).

Chapitre 5 [215] coécrit avec Fabio Martinelli et Cristina Toninelli montre Théorème 2.3 sur FA-2f sur $\mathbb{Z}^2$, ainsi établissant la première asymptotique exacte pour un KCM critique et tranchant les conjectures conflictuelles en physique (voir Section 2.2). La preuve repose de manière cruciale sur Chapitres 3 et 10 pour les bornes supérieure et inférieure respectivement et contient une version adaptée de Section 4.4.

Chapitre 6 [210] montre les bornes supérieures sur $\mathbb{E}_\mu(\tau_0)$ du Théorème 4.3 pour toutes les classes raffinées de familles bidimensionnelles critiques sauf les déséquilibrées au nombre infini de directions stables (a) traitées par [269] (voir Section 4). On se fie aux Chapitres 2 et 5.

Chapitre 7 [214] coécrit avec Laure Marêché et Cristina Toninelli montre la borne inférieure du Théorème 4.3 à corrections logarithmiques près pour les modèles critiques à nombre infini de directions stables (voir Section 4).

Chapitre 8 [213] coécrit avec Laure Marêché montre les bornes inférieures sur $\mathbb{E}_\mu(\tau_0)$ du Théorème 4.3 pour toutes les classes raffinées de familles bidimensionnelles critiques, en se basant sur Chapitre 7. Ainsi, avec Chapitre 6 nous achevons l'universalité raffinée des KCM critiques en deux dimensions (voir Section 4).

Chapitre 9 [219] coécrit avec Tamás Mezei examine le paramètre clé de l'universalité en deux dimensions à la fois pour la percolation bootstrap et KCM – la difficulté α – d'un point de vue computationnel (voir Section 4.1). Nous démontrons que la tâche de la déterminer, étant donné $\mathcal{U}$ est NP-difficile et fournissons un algorithme pour la trouver en temps fini.

Chapitre 10 [220] coécrit avec Robert Morris établit la borne inférieure du Théorème 2.2, déterminant l'ordre de grandeur du second terme de τ_0^{BP} pour la percolation bootstrap à deux voisins sur $\mathbb{Z}^2$ (voir Section 2.2). La preuve requiert une compréhension fine de la croissance typique d'une gouttelette critique et demande une quantité d'innovations techniques.

Chapitre 11 [222] coécrit avec Réka Szabó montre que GOSP se comporte comme la percolation orientée ordinaire dans sa phase surcritique (voir Section 5.2).

Chapitre 12 [209] fournit des résultats généraux sur les familles souscritiques (voir Section 5). Notamment, nous répondons à plusieurs questions posées dans [28]. La sortie du Chapitre 11 s'insère naturellement dans Chapitre 12 pour en étendre certains résultats à la généralité complète recherchée.

Chapter 1

Introduction

The present thesis concerns rigorous aspects of two classes of statistical mechanics models: bootstrap percolation and kinetically constrained models. Both models are located at the interface between statistical physics, probability and combinatorics, while the former is also studied from the viewpoint of computer and social sciences. Our work is probabilistic and combinatorial, mostly leaving out the other sides of the subject. The topic is at the border of several domains as diverse as interacting particle systems, probabilistic combinatorics, percolation, graph theory, cellular automata, computational complexity, glassy dynamics, social network phenomena and many more, so it is likely that the reader may come from one area or the other. Therefore, in the present chapter we provide a light and accessible introduction to the field assuming only knowledge of basic probability and giving heuristics rather than proofs. Our focus is on the two models on the square grid $\mathbb{Z}^2$ with random i.i.d. initial condition at low density and, more precisely, their universality.

The introduction is intended for linear reading and the readers should familiarise themselves with it before venturing to later chapters, which may be read in a more disordered and selective fashion. We provide a corresponding roadmap in Section 1.7, once we know what we are talking about, but before diving into details.

1.1 Motivation and background

In view of the volume and technicality of the manuscript, we owe the reader a furnished justification of our study. Our main motivation for the present work is the following universal proposition.

Proposition 1.1.1. *Everything is useful.*

Proof. This follows immediately from “C’est véritablement utile, puisque c’est joli,” as remarked by the little prince, in conjunction with the exclamation “C’est bien plus beau lorsque c’est inutile!” of Cyrano de Bergerac. □

The rest of Section 1.1 is devoted to an alternative, more cumbersome, less elegant, nonrigorous and chronologically posterior proof of a particular case of the above assertion in the setting of bootstrap percolation and kinetically constrained models provided for the sake of completeness. The mathematically predisposed reader, who would surely be satisfied with Proposition 1.1.1, is advised to skip ahead to Section 1.2.

1.1.1 Bootstrap percolation

Let us begin with the elder bootstrap percolation. To introduce the model informally, let us recall a folklore exercise. Consider an $n \times n$ bar of chocolate, certain of whose pieces are mouldy. Every day the edible pieces with at least 2 neighbouring mouldy pieces also perish. What is the smallest number of initially spoiled pieces needed to eventually spoil the entire bar?

The dynamics described above is what we call *2-neighbour bootstrap percolation*. It is more commonly phrased in terms of certain vertices of a graph G being initially infected, while healthy vertices become infected when they have at least r infected neighbours, and infections never heal. This results in the *r-neighbour bootstrap percolation* model on G.

The present thesis is only concerned with the typical infection time of, say, the origin in r-neighbour bootstrap percolation (and kinetically constrained models to be discussed in Section 1.1.2) and its generalisations with random initial set of infections on $G = \mathbb{Z}^d$ with its standard nearest neighbour graph structure. Before restricting ourselves to this setting, in which bootstrap percolation first emerged, let us discuss several other domains and settings in which it appears or is closely related to. These mentions are necessarily simplistic and likely abusive, taking into account that bootstrap percolation has long grown beyond the possibility of an exhaustive survey.

Other viewpoints on bootstrap percolation

Cellular automata—computer science A *cellular automaton* is a process which evolves in discrete time by updating simultaneously all the sites (vertices) of $\mathbb{Z}^d$ according to a rule, which depends on the current state of the process, but only on vertices within a certain distance of the site. The rule is assumed to be translation invariant. We further assume cellular automata to be *binary*, that is, each site has two possible states (healthy/infected). Thus, r-neighbour bootstrap percolation is a cellular automaton. The feature of bootstrap percolation that healing never occurs is referred to as *freezing* property. The fact that it inspects the total number of infections among the neighbours and that adding infections favours further propagation of infection then leads to calling it a *freezing totalistic monotone cellular automaton*. Alternatively, one may say *threshold* instead of *totalistic* or even *majority*, when r is exactly half of the number of neighbours, $2d$.

One viewpoint on cellular automata is the computational universality one. Namely, one asks whether it is possible to embed a Turing machine in them in order to perform arbitrary computations. Related questions are whether certain problems can be solved in finite time and, if so, in polynomial time. For instance, one may try to decide whether the state of the origin will ever change, given the initial state. We direct the reader to the recent thesis [262] on the subject and the references therein for such considerations. We also briefly examine this viewpoint in Chapter 9 for problems more related to our probabilistic perspective.

Cores—statistical physics In statistical physics and graph theory, a k-*core* is a subset of the vertices of a graph G which induces a graph of minimal degree at least k. In other words, each of the vertices has to have at least k neighbours in the subset. If G is d-regular, k-cores are equivalent to $(d-k+1)$-neighbour bootstrap percolation in the following sense. For any set of infected vertices, run the bootstrap dynamics until it becomes stationary. It is not hard to check that the resulting set of infections is the complement of a k-core and any k-core can be obtained this way. In particular, if $k=1$, we recover the hard-core model: all vertices become infected if and only if the initially healthy sites form an independent set.

The k-core model can be viewed as a constraint satisfaction problem and, as such, is related to spin glasses on one side and computational phase transitions in computer science on the other. Several questions become natural from this thermodynamics perspective. Firstly, one would like to know possible sizes of k-cores and, more generally, the number of such cores given their size, how to efficiently sample such configurations from a Boltzmann distribution, etc. Similarly, one may rather look at sets of initial infections which infect the entire graph (i.e. sets whose complement has no k-core) and investigate their thermodyniamics. We direct the reader to [198] and the references therein for this line of research and to its section II.D for more related models. Further see [302, 320] concerning the discontinuity of the phase transition of k-cores.

Propagation of contagion/opinion—computer and social sciences The bootstrap percolation dynamics and appropriate generalisations are used directly to model the propagation of infection, influence, opinions, etc. in networks. One is then usually interested in extremal questions such as what is the minimal size of a *dynamic monopoly*, i.e. a target set of customers (represented by vertices) one needs to convince (bribe, bombard with spam, etc.), so that their opinion can spread to the entire population or a large portion of it. How long does the spreading take? How can one determine such an optimal choice or approximate it algorithmically? We direct the reader to the survey [42].

Moving closer to our setting of interest, the same extremal questions may be asked for boxes or tori of $\mathbb{Z}^d$ as in the folklore exercise recalled above. Some notable references on the subject are [40, 45, 46, 201, 206, 278, 282, 304, 308].

Probabilistic viewpoint beyond lattices Moving still closer to our point of view, one may select the initial infected set randomly and ask whether the entire graph or large portion thereof becomes infected at a given 'density' of initial infections. This matter has been addressed on various graphs such as trees [39, 54, 72, 79, 150, 200, 322], random regular graphs [41, 233], (Gilbert–)Erdős–Rényi graphs [21, 234, 241, 351], hyperbolic lattices [311], Hamming tori [186, 328] and many more. Particularly interesting from the point of view of applications are graphs with 'real world' features such as sparsity, small diameter, heavy-tailed degree distributions, community structure, geometry etc. [1, 15–17, 78, 95, 147, 151, 164, 352, 363]. Yet closer to our setting are works on high dimensional hypercubes [31, 34, 35].

A somewhat different setting on lattices is the one of *polluted bootstrap percolation*. This may be seen as bootstrap percolation on the infinite supercritical cluster obtained from Bernoulli site percolation on the lattice. Alternatively, one may think of working on the entire lattice, but some sites are declared immune to infection from the start at random. Works in this setting include [165, 189, 191, 192].

Related models

Ising The (Lenz–)Ising model [228, 254] is perhaps the most studied model of magnetism together with its mean-field Curie–Weiss version. It so happens that bootstrap percolation was originally introduced [100, 303] precisely to model magnetic materials under appropriate conditions. It is therefore not surprising that the two are related. Firstly, bootstrap percolation was viewed in [7] as a first step towards understanding the metastability of the Glauber dynamics of the Ising model at (very) low temperature [77, 174, 268, 293]. Indeed, progress on bootstrap percolation did propagate to the Ising model [71, 99, 119, 120, 316, 317]. A further connection exists with the fixation of the zero-temperature dynamics of stochastic Ising [55, 149, 279] (also see [249]).

Graph bootstrap percolation Graph bootstrap percolation, also known as weak saturation, [67] sets out with two graphs G, H and a subset F of the edges of G viewed as initially infected. At each step one infects edges of G which complete an isomorphic copy of H. This model is often similarly flavoured to r-neighbour bootstrap percolation (corresponding to H being an $(r+1)$-star) and has sometimes been instrumental for the study of the latter, owing to the connections of the former to linear algebraic problems

[201, 282]. See [13, 36, 38, 73, 152, 239, 247, 273] for some combinatorial works and [22, 37, 44, 199] for probabilistic ones.

And more Let us also mention connections to the abelian sandpile [148, 280], two-way bootstrap percolation [368] (also see [314]), graph burning [75] and the recently introduced elliptic bootstrap percolation [14]. We ask the kind reader's forgiveness for other relatives having escaped our attention and refer to the references in [2, 4, 117, 359] for a few more.

1.1.2 Kinetically constrained models

Paradigmatic models

Although they were introduced independently, kinetically constrained models (KCM) are natural stochastic nonmonotone versions, canonically associated to any bootstrap percolation model. For the r-neighbour bootstrap percolation considered above the companion KCM is the *Fredrickson–Andersen r-spin facilitated model* (FA-rf) [155, 156], the parameter r being traditionally denoted j. The formal definitions can wait until Section 1.2.2, but let us give an informal one. Consider the lattice $\mathbb{Z}^d$ with some vertices initially infected. For each vertex we attempt to update its state at rate 1 (i.e. at random intervals of time with exponentially distributed mean 1 length). We are only allowed to do so if the site has at least j infected neighbours. If this constraint is satisfied, we flip a coin to decide what the new state will be, the coin yielding infection with probability q and health with probability $1-q$. Thus, infections can now heal and updates are no longer synchronous, but still have the same facilitation constraint as in bootstrap percolation. Another fundamental KCM is the *East model* [231] on $\mathbb{Z}$. There we may only update sites whose left neighbour is infected, but still do so at rate 1 w.r.t. a Bernoulli law with parameter q.

A model of glass

The FA-jf models were developed to model the glass transition, a brief introduction to which is in order. Yet, the domain of glassy physics is significantly more vast than bootstrap percolation and KCM together, so we do not even dream of providing a detailed account. We wholeheartedly recommend the very informative survey [23] including everything we relate, along with scores of references supporting it (see also [24, 255, 309, 341] for more emphasis on KCM). It is important to note that a great number of models of glasses have been proposed and physicists are nowhere near a consensus on which, if any, is 'correct.' KCM are but one of these theories with its virtues and sins, which we only partially account for.

A glass is obtained experimentally from a liquid by cooling it suddenly below its freezing temperature. Essentially every material can form a glass if asked gently, including SiO_2, yielding many of the everyday objects commonly called 'glass,' but by far not limited to it. Glasses are *rigid*, but not *solid*. The former means that they practically do not flow or do so extremely slowly and have mechanical properties of solid bodies. They are not solid, however, in the sense that they do not have the microscopic structure of solid state matter, namely a crystalline lattice. Instead, microscopically they are disordered and essentially indistinguishable from their fluid form. The 'transition' between liquid and glassy phases occurs very suddenly, with decreasing the temperature. It is uncertain whether one should speak of a proper thermodynamical phase transition at non-zero temperature, but it is clear that if a material takes a year for 1% of its molecules to move by a micron, only a mathematician would call it a liquid. As a matter of fact, glassy phenomenology goes well beyond the realm of the glasses defined above (usually called *structural glasses* to distinguish) and extends to various objects such as mesoscopic particles, grains, but also people (in a crowd) and more abstract entities such as solutions to constraint satisfaction problems (such as constructing a timetable or a thesis committee).

The idea behind the FA-jf models and KCM more generally is the following simple observation. If a person is in a dense crowd, so that there are people all around, he cannot move out of the 'cage' created by his neighbours. On the other hand, if somewhere in the crowd there is a little space, e.g. a person 'missing,' that is where movement can occur. From this perspective one thinks of infections as bits of free volume needed for anything to happen. In fact, from this microscopic perspective, it is more natural to consider conservative KCM (discussed below) to avoid people disappearing or popping up out of thin air. Indeed, FA-jf models are more appropriate from a coarse-grained vision of the material, some zones being sparser/better arranged than others, which can be viewed as excitations able to appear and disappear. The appeal of KCM comes from this intuitive explanation and their simplicity. Yet, from what we have said so far it is unclear whether they produce adequate glassy phenomenology known to occur in real materials.

If this thesis is being written, it is doubtlessly because this is the case. Indeed, KCM are capable of producing various emergent qualitative features (not hard-coded in the model's definition) characteristic of glasses. Such are the sharp divergence of relaxation time scales as the temperature (density of infections) decreases, dynamical heterogeneity (some regions move relatively fast, while others remain completely frozen for a long time), aging effects (the dynamics after a quench from high to low temperature becomes increasingly slower with time, retaining the 'age' of the system) and so on.

More convincingly yet, as we will see in detail, even one-dimensional KCM can exhibit two types of behaviour, which quantitatively match measurements in real glasses. Of course, it is delusive to think that the viscosity of any concrete glass is accurately predicted at a given temperature by the FA-1f model. However, some universal features are strikingly reproduced by KCM. More specifically, experimental studies have revealed two rough types of behaviours of (e.g. the relaxation times of) glasses called *strong* and *fragile*. Strong glasses, such as windows, feature Arrhenius behaviour, which may be explained by the existence of some local relaxation mechanism (think of swapping two molecules), which requires some fixed energy. Such mechanisms and quantitative divergence forms are exactly the ones associated to the simplest class of KCM we will call *supercritical unrooted*, whose simplest representative is none other than the FA-1f model.

Fragile glasses exhibit sharper super-Arrhenius divergences, as the (experimental) glass transition is approached, which appear to roughly follow a universal law. There are several simple expressions which give good fits for it and one of them is the one exhibited by another class of KCM called *supercritical rooted*, such as the East model. In contrast to supercritical unrooted models, for supercritical rooted ones many infections need to be created in order for macroscopically distant regions to interact. As we will see, for more difficult models even sharper divergences can occur, owing to 'cooperative' relaxation mechanisms requiring many particles to move collectively in a coordinated fashion. Depending on the choice of constraint, KCM can exhibit a finite-temperature dynamical phase transition to the freezing of a fraction of the system, as observed in jamming.

To be certain, all of this is quite remarkable to uncover in a model one can define in 3 lines. Yet, KCM are not 'the' widely accepted solution to the glass problem. The most straightforward objection is 'How do these constraints emerge from a first-principle microscopic description and what is the exact coarse-graining procedure giving rise to them?' (see [93, 94, 169, 353] for progress in this direction). Moreover, one of the most crucial features of KCM is the fact that their equilibrium measure is trivial, there is no static interaction between infections. This was done on purpose in order to see how much glassy behaviour one can recover only based on dynamical constraints devoid of static interaction and static phase transition. Instead, other theories argue that precisely static interactions are at the heart of glassy behaviour and, perhaps, a complex type of order should appear at the transition. Since such theories have their own supporting evidence, it is not surprising that KCM have little chance of capturing the full picture. Yet, even capturing a fragment of this complex problem is ample justification for the study of KCM, if any is needed.

Mathematical challenges

Even if the physical motivation of KCM is sound and the models are simple, that is not enough to make them easy to handle mathematically. Indeed, one of the main reasons for their special mathematical treatment is that they lack almost every nice feature one may hope for in an interacting particle system. The good properties of KCM are exhausted by the availability of a simple explicit reversible measure, the finite range of interactions and the dynamics being of Glauber type. The other side of the scale is more copious.

Non-attractiveness A spin system is called 'attractive' if infections favour infections. That is, one can couple the dynamical evolution starting with different initial conditions in such a way that if we add infections to an initial configuration, this will not lead to having fewer infections at a later time than if evolving from the original initial condition (see e.g. [258]). This crucial feature is lacking for KCM. Indeed, adding an infection may enable the healing of another site through an update which would otherwise violate the constraint.

The consequences of attractiveness in other interacting particle systems such as bootstrap percolation, the contact process, the voter and stochastic Ising models, to name a few, are numerous and heavily relied on. For instance, it is often sufficient to understand the behaviour of the system starting from the fully infected and fully healthy states and those evolve in a stochastically monotone way. In many cases one of the two states leads to a trivial dynamics or the two play symmetric roles, narrowing the study down to a single very simple initial condition. Furthermore, attractiveness is known to enable the use of censoring [298]—we may disregard certain preselected updates. This is particularly convenient to ensure that some parts of space have relaxed individually before having them interact with the rest of the system.

Degenerate rates The very definition of KCM imposes hard constraints on the dynamics. This means that we cannot modify the state of a given site at a finite cost. Moreover, the dynamics cannot relax locally (e.g. a fully healthy region cannot change on its own). This is an intrinsic difficulty related to the dynamical heterogeneity of KCM (and glasses): some regions are forced to remain inactive for long periods of time.

Consequently, another common tool, the (modified) logarithmic Sobolev constant, degenerates in infinite volume [88]. Indeed, mixing in finite volume, even with fully infected boundary condition is necessarily at least polynomially slow in the volume, since infections can only propagate at finite speed. In fact, it is not clear and not always true that relaxation occurs at all in KCM. Indeed, as already mentioned, some undergo a dynamical phase transition causing a portion of the volume to freeze at finite temperature and breaking ergodicity.

Cooperative dynamics In conservative dynamics as opposed to Glauber ones, it is common to need to move something in order to relax rather than relaxing locally throughout the volume and appropriate tools are not lacking. A problem in implementing them (other than non-attractiveness) is that it turns out to be hard to conveniently embed random walks even in the simple FA-1f model [59]. As we will see in Chapter 3, having random walks and attractiveness at our disposal for a model makes a real difference.

However, it is the cooperative behaviour of many KCM, which poses a real challenge. For them it is not sufficient to move a single microscopic object throughout the volume of the system, but rather necessary to create and move many infections in a coordinated fashion. This is quite uncommon and entails uncommonly sharply diverging relaxation times.

Putting these difficulties together makes even the study of the equilibrium dynamics of KCM particularly challenging. As a result, our task is three-fold. Firstly, we need to identify bottlenecks, which involve significant combinatorial difficulties (owing to the degenerate rates), even if the approach is standard in dynamics of interacting particle systems. Secondly, we need to develop an heuristic understanding of the intricate efficient relaxation mechanisms used by the dynamics (particularly due to the cooperative dynamics). Finally, we need to develop appropriate tools for translating the heuristics into mathematical results, where standard methods fail (because of non-attractiveness and the like).

Mathematical works beyond the standard setting

Before turning to our setting of interest for KCM, namely those on $\mathbb{Z}^d$ at equilibrium (with initial condition distributed according to the product Bernoulli measure with parameter q) in the low temperature limit $q \to 0$, let us review other mathematical viewpoints on KCM. Since they are somewhat less studied than bootstrap percolation, we may attempt to supply the reader with an exhaustive bibliography. Accounts of nonrigorous studies can be found in [24, 53, 173, 255, 309, 341]. We further recommend the recent monograph [342] as an excellent reference, particularly for mathematical works in the field.

Other graphs Appropriate versions of FA-jf and East models have been studied on graphs other than $\mathbb{Z}^d$, including hyperbolic lattices [311], trees [91, 271] and arbitrary graphs [89] (see also Chapter 3).

Out of equilibrium While we will be almost exclusively interested in equilibrium properties of KCM at low temperature, from the physics point of view it is very relevant to study these models subjected to a temperature quench. That is to start with an initial condition e.g. distributed according to a product measure with given infection density different from the invariant one, q. Exponential convergence to equilibrium is studied in [59, 92, 105, 283] and generalised in [263, 266]. The next step are results on the front of the East and FA-1f models in one dimension [57, 60] (e.g. with only negative integers initially infected). Finally, based on the above, cutoffs have been proved for East and FA-1f in one dimension [140, 162]. Unfortunately, all these results are either specific to the East model or only valid in the high temperature regime—a most undesirable restriction in models whose glassy behaviour is exhibited when approaching zero temperature. For the very special East model additional results out of equilibrium can be found in [104, 143, 145] and in [105, 106, 114] for its higher dimensional analogue. A KCM of a very different type, the North-East model, is studied out of equilibrium in [108].

Let us further mention that out-of-equilibrium results of a different kind [299, 300] concerning finite graphs whose size diverges jointly with temperature going to 0, are the only ones of use to us and will be discussed and derived in Section 1.4.2 and Chapter 3.

Finally, large deviations in trajectory space have been studied for East and FA-1f in [66], since physicists proposed viewing the behaviour of KCM as an ordinary static phase transition in the space of trajectories, driven by a 'dynamical activity' parameter [65, 171, 172, 230].

Diffusion One may wish to study the way a tagged particle diffuses within a KCM. For works on the subject see [58, 63]. This is all the more natural for conservative KCM discussed in the next paragraph [49, 64, 142, 343].

Related models

Conservative KCM As mentioned above, in some cases conservative versions of KCM may be more appropriate. The first and most classical such models are the Kob–Andersen ones [244], in which one is allowed to move a particle to an unoccupied neighbouring position, provided that a constraint similar to the one of FA-jf models is satisfied. More precisely, the particle (healthy site) should have at least j empty (infected) neighbours both before and after the transition. These models and their variants also known as kinetically constrained lattice gases have attracted significant attention [64, 90, 142, 270, 288, 323, 343, 345]. They turn out to be very closely related to non-conservative ones, yet none of the results presented in this thesis for non-conservative KCM has a conservative analogue as of now. Some other related kinetically constrained lattice gases can be found in [49, 61, 62, 76, 118, 139, 176–178, 232] and are currently quite active.

Interacting KCM Another natural modification of KCM is to introduce static interaction between particles. This may be achieved by updating each site w.r.t. a measure depening on the current state of other sites. This venue has only been explored in [89] for high-temperature Gibbs measures. One reason for very systematically considering non-interacting infections is that KCM were introduced precisely to investigate the effects of dynamical constraints in the absence of interesting static interactions. Besides, the non-interacting case is already mathematically very challenging and its product invariant measure is one of the very few useful properties available to work with.

Inhomogeneous KCM As in polluted bootstrap percolation, it is possible to introduce non-homogeneous versions of KCM. For instance, one may consider some sites in $\mathbb{Z}^d$ having FA-2f constraint, while others have FA-1f constraint. Such models are studied in [324, 326]. Generalisations thereof, albeit in one dimension, are studied in Chapter 2 as a tool for higher dimensional homogeneous KCM.

East-related models The East model being one of the simplest KCM with very useful additional features (namely it is oriented), it has been extensively studied and relates to other models of interest. For instance, a certain upper triangular matrix walk projects to an East process if one looks at the last column [163, 297]. A model of hierarchical coalescence was used for studying aging in East [144, 146]. Also, East being quite well understood, it has been possible to study random walks with East as random environment [26].

Plaquette models Plaquette models are also spin models with Glauber dynamics but, contrary to KCM have (local) static interactions instead of kinetic constraints. These were introduced to show that kinetic constraints can emerge from static interactions at low temperatures [169, 290]. Rigorous works on these can be found in [107, 109, 110].

And more Quantum versions of KCM are being studied lately as models for many-body localisation [170, 296], though mathematical treatment seems to be lacking for the moment. Finally, let us mention that some techniques developed in KCM are exported to other settings [52, 96].

To conclude, let us gather all PhD theses dedicated to KCM we are aware of [24, 56, 141, 255, 264, 321, 341], which the reader having reached this point may wish to consult.

1.2 Setup

Now that we are hopefully convinced of the importance of bootstrap percolation and KCM, let us formally define our models of interest, starting with the less technical bootstrap percolation. We leave the history of the emergence of $\mathcal{U}$-bootstrap percolation and KCM from the r-neighbour case to Section 1.5.3.

1.2.1 Bootstrap percolation

A bootstrap percolation model is defined by a positive integer d (*dimension*) and an *update family* $\mathcal{U}$ which is a finite non-empty family of finite non-empty subsets of $\mathbb{Z}^d\backslash\{0\}$ called *update rules*. Given a set $A \subset \mathbb{Z}^d$, we define the *bootstrap map*

$$\mathscr{B}_{\mathcal{U}}(A) = A \cup \left\{x \in \mathbb{Z}^d : \exists U \in \mathcal{U}, x + U \subset A\right\}$$

with $B+C = \{b+c : b \in B, c \in C\}$ and $b+C = \{b\}+C$ for any $B, C \subset \mathbb{R}^d$ and $b \in \mathbb{R}^d$. In words, given that A is infected, at the next step we additionally infect each site such that some rule translated at it is already fully infected. This process is naturally viewed as a dynamical system in discrete time. Given an initial set of infected sites $A \subset \mathbb{Z}^d$, its *closure* is

$$[A]_{\mathcal{U}} = \bigcup_{n\in\mathbb{N}} \mathscr{B}_{\mathcal{U}}^n(A),$$

where $\mathscr{B}_{\mathcal{U}}^n$ stands for $\mathscr{B}_{\mathcal{U}}$ iterated n times and $\mathbb{N} = \{0, 1, \dots\}$ is the set of non-negative integers. Since $\mathcal{U}$ will usually be fixed, we will omit it from all notation, unless confusion may arise. We say that A is *stable* if $A = [A]$.

Let us make a few observations. By definition bootstrap percolation is monotone in two distinct ways. Firstly, $A \subset \mathscr{B}(A)$ for any $A \subset \mathbb{Z}^d$, making it monotone in time. Secondly, the rules are monotone in the sense that $A \subset B$ implies $\mathscr{B}(A) \subset \mathscr{B}(B)$ for any $A, B \subset \mathbb{Z}^d$. In other words, they only ask for 'enough' infections, rather than 'exactly' some amount of infections. Up to these monotonicity properties and the assumptions that rules are translation invariant and finite range, the family of models defined by all $\mathcal{U}$ is as general as possible. Indeed, any infection condition can be readily recast in this language by writing it in its disjunctive normal form, which is only subject to the restriction not to contain negations (in other words, every increasing set can be written as the union of the increasing sets induced by its minimal elements). To give an example, let us take the r-neighbour model. Here we have $\mathscr{B}(A) = A \cup \{x \in \mathbb{Z}^d : |\{y \in A : y \sim x\}| \geqslant r\}$ with $y \sim x$ if x and y are neighbours in the usual graph structure on $\mathbb{Z}^d$. Yet, we may also view this as $\mathscr{B}_{\mathcal{U}}(A)$ with $\mathcal{U}$ the family of $\binom{2d}{r}$ sets of r neighbours of the origin. That is, if any r of the neighbours of a given site are all infected, it also becomes infected.

So far the process is purely deterministic. We introduce randomness only in the initial condition, by taking $x \in A$ independently for each $x \in \mathbb{Z}^d$ with probability q. We denote the law of A by μ_q and drop q whenever it is clear from the context. As usual in statistical physics, we can then introduce the *critical probability*

$$q_{\mathrm{c}} = \inf\left\{q \in [0,1] : \mu_q\left([A] = \mathbb{Z}^d\right) > 0\right\}. \tag{1.1}$$

Since the event above is translation invariant, ergodicity says that in fact $\mu_q([A] = \mathbb{Z}^d) \in \{0,1\}$ for all $q \in [0,1]$. While for $q < q_c$ one can study the geometry of $\mathbb{Z}^d \backslash [A]$, in the regime $q \geqslant q_c$ one may want to be more quantitative. For this purpose we introduce the (bootstrap percolation) *infection time*

$$\tau_0^{\mathrm{BP}} = \min\{n \in \mathbb{N} : 0 \in \mathscr{B}^n(A)\} \in \mathbb{N} \cup \{\infty\},$$

which is a random variable, setting $\min \varnothing = \infty$ as usual.

Let us note that in bootstrap percolation results are usually sought for τ_0^{BP} w.h.p. That is, one seeks to prove that for some real functions $a(q), b(q)$ it holds that

$$\lim_{q \to q_c+} \mu_q\left(\tau_0^{\mathrm{BP}} \in [a(q), b(q)]\right) = 1.$$

However, it is equally meaningful to (approximately) determine the mean, $\mu\left(\tau_0^{\mathrm{BP}}\right)$, or, perhaps, the rate of decay of the tail of the distribution of τ_0^{BP} for $q \geqslant q_c$. As we will see, it is quite systematically possible to enhance results holding w.h.p. to results in expectation, while decay rates are somewhat differently behaved. We will preferentially state results for $\mu\left(\tau_0^{\mathrm{BP}}\right)$ for the sake of consistency with KCM tradition.

Furthermore, results in bootstrap percolation are often stated in terms of a critical probability in finite volume, but we will prefer infection times in infinite volume again for compatibility with KCM. While there is no general result saying that determining finite volume critical probabilities and the distribution of τ_0^{BP} is equivalent, in practice all proofs available in the domain work for both quantities and it is a matter of personal taste whether both or just one is explicitly stated and/or proved.

1.2.2 KCM

KCM are similarly defined by their dimension $d \geqslant 1$, update family $\mathcal{U}$ as in Section 1.2.1 and parameter $q \in [0,1]$. The $\mathcal{U}$-KCM is a continuous time Markov process with state space $\Omega = \{0,1\}^{\mathbb{Z}^d}$, zeros corresponding to infections. We naturally identify any configuration $\eta \in \Omega$ with the set of its infections (called A in Section 1.2.1). The *constraint at* x is given by

$$c_x(\eta) = \mathbb{1}_{\exists U \in \mathcal{U}, \eta_{x+U} = \mathbf{0}}, \tag{1.2}$$

writing η_X for the restriction of $\eta \in \Omega$ to $X \subset \mathbb{Z}^d$ and $\mathbf{0}_X$ (resp. $\mathbf{1}_X$) for the fully infected (resp. healthy) configuration on X, omitting X when it is clear from the context. If $X = \{x\}$ is a singleton, we may write just x for brevity, so that η_x is the state of the site $x \in \mathbb{Z}^d$ in the configuration $\eta \in \Omega$. We further write η^x for the configuration obtained from η by flipping the state of site x, that is, $(\eta^x)_y = \eta_y$ for all $y \in \mathbb{Z}^d \backslash \{x\}$ and $(\eta^x)_x = 1 - \eta_x$.

We denote by $\mu_q = \mu$ the product Bernoulli measure, such that $\mu(\eta_x = 0) = q$, η denoting a random configuration with law μ. We further denote by $\mu_X(f)$ the conditional average $\mu(f(\eta)|\eta_{\mathbb{Z}^d\setminus X})$ over the states in X for any finite set $X \subset \mathbb{Z}^d$ and any real function $f : \Omega \to \mathbb{R}$. The variances w.r.t. μ and μ_X are denoted by Var and Var_X respectively.

We are then ready to define the $\mathcal{U}$-KCM by the action of its generator on functions $f : \Omega \to \mathbb{R}$ depending on the states of finitely many sites

$$\begin{aligned}\mathcal{L}_{\mathcal{U}}(f)(\eta) &= \sum_{x\in\mathbb{Z}^d} c_x(\eta)\cdot(\mu_x(f)-f)(\eta)\\ &= \sum_{x\in\mathbb{Z}^d} c_x(\eta)\cdot(q\eta_x+(1-q)(1-\eta_x))\cdot(f(\eta^x)-f(\eta)),\end{aligned}$$

where μ_x stands for $\mu_{\{x\}}$ and similarly for Var_x. This process can also be defined via a more intuitive graphical construction as follows (see [258] for background). Each site is equipped with an independent standard Poisson process (*clock*) whose atoms are called *clock rings*. When the clock at x rings, we are additionally given an independent random variable with (Bernoulli) law μ_x. If the constraint c_x is satisfied, we call the update *legal*, change the state of x to that variable and do nothing otherwise (*illegal update*). Since the number of clock rings is locally finite and the constraints have finite range, this poses no problems [258]. It is clear from either description that μ is a reversible invariant measure for the KCM, due to the fact that $c_x(\eta)$ does not depend on η_x. For this reason we call μ the *equilibrium measure*. It should be noted, however, that e.g. the Dirac measure on the healthy configuration $\mathbf{1}$ is also invariant.

The most intuitive quantity of interest is the (random) infection time of the origin

$$\tau_0 = \min\{t \geqslant 0 : \eta_0(t) = 0\},$$

denoting the state of the KCM at time $t \in [0,\infty)$ by $\eta(t)$. We will then be interested in $\mathbb{E}_\mu(\tau_0)$, the expected infection time for the stationary process (with initial state distributed according to μ). A more analytic but equally important measure of the speed of the dynamics is its *relaxation time* T_{rel} defined as the inverse of the spectral gap of $\mathcal{L}_{\mathcal{U}}$. Fortunately, we will never need to consider the spectrum of operators, thanks to the more manageable definition

$$(T_{\mathrm{rel}})^{-1} = \inf_{f\not\equiv\mathrm{const}} \frac{\mathcal{D}_{\mathcal{U}}(f)}{\mathrm{Var}(f)}, \tag{1.3}$$

where $\mathcal{D}_{\mathcal{U}}$ is the Dirichlet form associated to $\mathcal{L}_{\mathcal{U}}$

$$\mathcal{D}_{\mathcal{U}}(f) = \sum_{x\in\mathbb{Z}^d} \mu(c_x\cdot\mathrm{Var}_x(f)) = -\mu(f\mathcal{L}_{\mathcal{U}}(f)). \tag{1.4}$$

As for τ_0^{BP}, one may seek results for τ_0 w.h.p., in expectation (for $\mathbb{E}_\mu(\tau_0)$), in large deviations rate (the rate of decay of the tail of τ_0). For our purposes all three approaches, as well as T_{rel}, are essentially equivalent. Of course, they

(a) FA-1f/1-neighbour BP (b) East (c) FA-2f/2-neighbour BP

Figure 1.1 – The three paradigmatic update families in one dimension. The sites of each rule are represented by dots and the origin is marked by a cross. The arrows indicate in which direction infection can travel via the rule.

are not equivalent a priori and sometimes more work is needed to control all of them. Occasionally, in order to avoid some technicality, we do not pursue results for T_{rel}, but these can probably also be recovered with some more work. This is important, as we will often use T_{rel} for simple models, in which case we will make sure to prove corresponding bounds.

1.3 One dimension

As an instructive warm-up let us start by examining the one-dimensional case. For bootstrap percolation it is nearly trivial, while for KCM there are already interesting phenomena. In view of the results in two dimensions we are aiming for, the results presented in this section are out of fashion, but would have been quite interesting 5-10 years ago. Indeed, any one-dimensional model is readily embedded in two dimensions.

Let us begin by recalling three paradigmatic models—FA-1f, East and FA-2f (see Fig. 1.1). In FA-1f an infection is needed in at least one of the neighbours, in East we only look at the left neighbour, while FA-2f requires both neighbours to be infected in order to legally update the vertex. We call the corresponding bootstrap percolation models 1-neighbour, East and 2-neighbour respectively.

1.3.1 Bootstrap percolation

A complete solution to the bootstrap percolation problem is a simple exercise in all three cases.

Proposition 1.3.1. *In one dimension we have $\mu(\tau_0^{\mathrm{BP}}=0)=q$ and, conditionally on $\tau_0^{\mathrm{BP}}>0$,*

(1) τ_0^{BP} has the geometric distribution with parameter q and mean $1/q$ for the East bootstrap percolation.

(2) τ_0^{BP} has the geometric distribution with parameter $1-(1-q)^2=2q-q^2$ and mean $1/(2q-q^2)$ for the 1-neighbour bootstrap percolation.

(3) $\tau_0^{\mathrm{BP}}=1$ with probability q^2 and $\tau_0^{\mathrm{BP}}=\infty$ otherwise for 2-neighbour bootstrap percolation.

Proof. We treat each update family independently.

(1) The infection time equals the distance to the first initially infected site to the left of the origin. Indeed, each infection propagates one to the right at each step.

(2) The infection time equals the distance to the closest initially infected site (on either side). Indeed, on each step each infection propagates by one in both directions.

(3) Two consecutive healthy vertices remain healthy forever. Hence, if the origin is to become infected, either it or both its neighbours are initially infected. □

Remark 1.3.2. In order to dissipate a natural conjecture, we leave it as an exercise for the inquisitive reader to construct an update family $\mathcal{U}$ such that τ_0^{BP} has full support $\mathbb{N} \cup \{\infty\}$ and for the lazy or busy reader to verify that $\{\{-1,1,2\}\}$ is such a family.

Since such explicit results are beyond reach in general, unless otherwise stated, we consider the limit $q \to 0$ and use the following asymptotic notation. For any real functions $f(q)$ and $g(q)$ defined for $q > 0$ sufficiently small such that $g > 0$, we write

- $f = \Theta(g)$ when $cg(q) \leqslant f(q) \leqslant Cg(q)$,
- $f = \Omega(g)$ when $cg(q) \leqslant f(q)$,
- $f = O(g)$ when $|f(q)| \leqslant Cg(q)$

for some constants $0 < c \leqslant C < \infty$ and all $q > 0$ sufficiently small. Such implicit constants are allowed to depend on $\mathcal{U}$ but not on q. We further write $f = o(g)$ if $\frac{f(q)}{g(q)} \to 0$ and $f \sim g$ if $\frac{f(q)}{g(q)} \to 1$ as $q \to 0$. We use the symbol $\approx$ more vaguely for the purposes of heuristics in the present chapter.

Proposition 1.3.1 admits the following natural generalisation.

Definition 1.3.3. Let $\mathcal{U}$ be a one-dimensional update family and let

$$\zeta = \zeta(\mathcal{U}) = \min\{|A| : A \subset \mathbb{Z}, |[A]_{\mathcal{U}}| = \infty\}$$

be the minimum number of infections required to infect infinitely many sites. We say that $\mathcal{U}$ is *supercritical* if $\zeta < \infty$ and *trivial subcritical* otherwise.

Proposition 1.3.4 (Bootstrap percolation universality in one dimension)**.** *For any one-dimensional update family $\mathcal{U}$ one of the following alternatives holds.*

(1) $\mathcal{U}$ is supercritical and $\mu\left(\tau_0^{\mathrm{BP}}\right) = \Theta(q^{-\zeta})$.

(2) $\mathcal{U}$ is trivial subcritical and $\mu\left(\tau_0^{\mathrm{BP}}\right) = \infty$ for all $q \in (0,1)$.

Proof. Consider a large infected segment. Either it is able to extend in one direction (if the model is supercritical) or a large healthy segment will remain such forever (if the model is trivial subcritical), since it cannot be disrupted from either side. Thus, for trivial subcritical models $\mu(\tau_0^{\mathrm{BP}} = \infty) \geqslant (1-q)^C$ for C large enough.

Turning to supercritical models, using that $\mathcal{U}$ has finite range, it is not hard to see that any finite set of infections either infects finitely many sites or creates a periodic infection pattern to its left or right. This periodic pattern propagates linearly in time. Thus, the time needed to infect the origin is at most the distance to the closest set generating infinitely many infections on the correct side of the origin and with the correct remainder modulo the period. Then for small q this is dominated by sets of infections of minimal size, entailing the desired upper bound.

To see that $\mu\left(\tau_0^{\mathrm{BP}}\right) \geqslant \Omega(q^{-\zeta})$, observe that up to such distance from the origin with probability $\Omega(1)$ there are only sets of nearby infections generating finitely many infections. Since two such patches of infection do not interact until the infection they generate meets, one can prove uniform bounds on how far their infection can spread, concluding the proof (see Lemma 9.2.1 for more details). □

It is possible to refine the above result by observing that for small q the positions of sets of infections with infinite closure converge to a Poisson point process. Under suitable assumptions it then suffices to examine the number of such sets of size ζ modulo translation, governing the intensity of the process, and the speed at which a large infected interval expands in each direction. Such considerations were carried out in [181–184], but we will not pursue this direction further and consider bootstrap percolation in one dimension as completely resolved by Proposition 1.3.4.

However, it is worth mentioning that this probabilistic viewpoint completely disregards e.g. the computational complexity of determining ζ, given $\mathcal{U}$. We will return to closely related questions in Chapter 9.

1.3.2 KCM

For KCM the situation is significantly more intricate. Nevertheless, Definition 1.3.3 and Proposition 1.3.4 are still relevant, as they show that for trivial subcritical models $\mathbb{E}_\mu(\tau_0) = \infty$ for any $q < 1$. Indeed, if bootstrap percolation cannot infect the origin, neither can the corresponding KCM. At this point it is good to note that KCM with $q = 1$ degenerate into a continuous time version of bootstrap percolation, which can be treated identically to the discrete time one modulo technical details.

It then remains to determine the scaling of $\mathbb{E}_\mu(\tau_0)$ for supercritical models as $q \to 0$. Before doing this, let us consider the FA-1f and East models, which will turn out to be behave very differently (like a strong and a fragile glass respectively).

FA-1f

In order to gain some intuition, let us examine the typical evolution of the FA-1f dynamics on $\mathbb{Z}$ starting from a single infection at 0, when q is small. The only allowed move from the initial condition is infecting 1 or -1, which occurs at time of order $1/q$. W.l.o.g., assume we obtain infections at 0 and 1. Each one can disappear at rate $1-q$, much more than the rate of creating additional infections. Hence, w.h.p. we quickly either return to the original configuration or reach the one with only 1 infected. This sequence of moves nets a move of the central infection to the left or right with equal probability at rate roughly $1/q$. The result over larger lapses of time is a symmetric random walk slowed down by a factor q. The first deviation from this behaviour is observed w.h.p. at times of order $1/q^2$, at which point we may create a third infection before eliminating one of the two present when attempting to perform a move of the random walk. If, say, $-1,0,1$ are infected, several things may happen, all at rate of order 1. We may eliminate -1 or 1 and then reach a state with a single infection; or 0 may heal first, leading to two separated infections. In the latter case the two infections typically remain intact for a long time: a branching has occurred. When two (or more) infections are present, each one typically performs an independent random walk, until they come next to each other, at which point they typically coalesce.

With this picture in mind, let us begin our quantitative study of the model. For FA-1f, we have $\mathbb{E}_\mu(\tau_0) \geqslant (1-o(1))/q$, since w.h.p. the origin is not initially infected and the rate at which it becomes infected is either 0 or q, depending on the current configuration. Providing an upper bound is a bit harder. From our consideration of 1-neighbour bootstrap percolation, we recall that we should take into account the infection closest to the origin, which is typically at distance $1/q$ at equilibrium. The most naive idea would be to simply proceed as in bootstrap percolation and infect successively the sites from it to the origin. Let us see why this is not an efficient infection mechanism for FA-1f. As we are working with the stationary process, at all times the configuration is distributed according to μ. Therefore, witnessing a very unlikely event, such as roughly $1/q$ infections next to the origin, is not likely before a given time horizon T. More precisely, we would need to wait at least until time of order $q^{-1/q}$ to have a good chance to see this event.

Hence, we need to refine our infection mechanism to have fewer infections close to the origin at any given time. Fortunately, a simple solution is available, as suggested by the typical evolution of a lone infection. Starting from the infection closest to the origin, we infect its neighbour in the direction of the origin and then remove the original one and repeat until we reach the origin. The total effect of this is moving the infection by one site, but not creating anomalously many infections. This simple argument is formalised via cannonical paths (originating from [237], see [256, 310] for background) in [325] to show $\mathbb{E}_\mu(\tau_0) \leqslant O(1/q^3)$ (also see [88]). Indeed, starting from most configurations, we go through each configuration with two consecutive infections at positions $0 \leqslant x \leqslant 1/q$ and $x+1$. There are $1/q$ such positions and the corresponding configurations restricted to the $(1/q)$-neighbourhood of the origin have probability roughly $1/q^2$, leading to the bound claimed. We would like to keep the present chapter free from technical details, so we will not provide more details and appeal to the reader's intuitive understanding. Let us reassure the overly suspicious reader that they will have ample opportunity to savour formal proofs in the remaining chapters.

Since, heuristically, $1/q^3$ also corresponds to the time needed for an infection performing a random walk with steps with rate q, we would expect this upper bound to be sharp. This is indeed the case, leading to the following recent result, more difficult than the heuristics suggests.

Theorem 1.3.5 (FA-1f in one dimension [325])**.** $\mathbb{E}_\mu(\tau_0) = \Theta(1/q^3)$.

The same result holds for T_{rel} [88] and, using this, some physics predictions [48, 362] were ruled out.

East

We next examine the East model, which, contrary to FA-1f, was originally introduced precisely in one dimension [231] and this is the setting in which it is the most studied. Let us begin by considering the typical microscopic evolution starting from a single infection at 0. We first infect 1 at time of order $1/q$. However, instead of propagating further we typically hurry to undo this, since 1 heals at rate $1-q$ and 2 is infected at rate only q. After repeating this attempt roughly $1/q$ times, we eventually create infections at $0, 1, 2$ and the infection at 1 heals before the one at 2. This happens on time scale roughly $1/q^2$. Though 2 remains infected for quite some time, we have not made great progress. Indeed, even if we manage to infect 3, it will typically heal immediately and long before we manage to infect both 3 and 4, we tend to infect 1 again, leading to the healing of 2 and then 1. So we are quite likely to end up at our starting point at time only $1/q$ after we eventually managed to infect 2. Reaching 4 is then much harder and requires waiting for the above procedure to be repeated $1/q$ times until time $1/q^3$, at which point we may infect simultaneously $0, 2, 3, 4$ and be lucky enough for 3 to heal before 4 does. Ending up with $0, 2, 4$ infected, we have some chance of healing 2 before it destroys the infection at 4, thus ensuring that 4 stays infected for a time of order $1/q^2$, which might suffice to infect 6 or even 7. It should then be clear that the dynamics of the East model is much slower that FA-1f and features hierarchical back-and-forth motion, which only rarely manages to expel an infection far from the origin, often to be pulled back soon after, but sometimes becoming metastable for a long time.

With this image in mind, let us try to determine the typical timescales of East. Here the rough lower bound of $1/q$ above is still valid (and in fact so is $\Omega(1/q^3)$). However, the upper bound argument clearly breaks down, because we can no longer simply move an infection, leaving nothing behind. Should we then expect $\mathbb{E}_\mu(\tau_0) \approx q^{-1/q}$? The answer is negative, but we do have to create many infections. To be precise, it is a cute exercise to show that, starting from only one infection at 0 and using at most n additional infections at any single time, the East process can infect the site $2^n - 1$ (rather than the trivial n that we saw), but nothing further away from the origin. As the expression suggests, this can be achieved by an inductive procedure as follows. Use $n-1$ infections to place an infection at $2^{n-1}-1$. Then infect 2^{n-1} and remove all other infections. This is possible, since the dynamics is reversible. Now imagine 2^{n-1} is the origin and proceed in the same way with $n-1$ infections, starting from it. Hence, we only need about $\log_2 L$ infections to reach distance L and this is optimal, as stated below.

Definition 1.3.6 (Legal path)**.** Given an update family $\mathcal{U}$ in any dimension, a *legal path* γ is a finite sequence $\omega_{(0)}, \dots, \omega_{(k)}$ of configurations in Ω such that for each $i \in \{1,\dots,k\}$ there exists $v = v(\omega_{(i-1)}, \omega_{(i)}) \in \mathbb{Z}^d$, such that $\omega_{(i)} = (\omega_{(i-1)})^v$ and $c_v(\omega_{(i)}) = 1$. That is, successive configurations differ by a single legal spin flip.

Proposition 1.3.7 (Combinatorial bottleneck for East)**.** *Consider the East model on $\{1,\dots,M\}$ defined by fixing $\omega_0 = 0$ at all time. Then any legal path γ connecting the fully healthy configuration $\mathbf{1}$ to a configuration ω such that $\omega_M = 0$ goes through a configuration with at least $\lceil \log_2(M+1) \rceil$ infections.*

The above 'hierarchical' relaxation mechanism was noticed in [274, 329] and [112] examined the combinatorics (and in particular the number) of configurations reachable with a given number of infections more closely. From this one can easily obtain $\mathbb{E}_\mu(\tau_0) \geqslant q^{-c\log(1/q)}$ for c small enough as $q \to 0$. Indeed, by Proposition 1.3.7, with probability at least $1/2$ we need to witness exactly $\log(1/\sqrt{q})$ infections at distance at most $1/\sqrt{q}$ from the origin to its left to infect it. The number of such configurations (restricted to the $1/\sqrt{q}$ sites left of to the origin) is at most $(1/\sqrt{q})^{\log(1/\sqrt{q})}$ and thus subdominant w.r.t. their probability which is at most $q^{\log(1/\sqrt{q})}$. A more involved argument [8] was used to show that

$$\begin{aligned}\exp\left(\log^2(1/q)\,(1/(2\log 2) - o(1))\right) &\leqslant \mathbb{E}_\mu(\tau_0)\\ &\leqslant \exp\left(\log^2(1/q)\,(1/\log 2 + o(1))\right).\end{aligned}$$

The upper bound matched the physics conjecture [329] and was achieved by comparison with a different process and using canonical paths based on the iterative procedure above.

Surprisingly, it turned out that the lower bound was the sharp one [88]. The proof of this fact introduced the bisection technique, which we present next. The idea is to establish bounds on $\mathbb{E}_\mu(\tau_0)$ (rather T_{rel}, but the two are related, so that the final result holds for both) on finite segments by induction on their size. More precisely, one decomposes a segment essentially into its left and right halves (hence the name) and uses the following crucial two-block lemma to effectively bring an infection at the boundary of the right half, using the left half.

Lemma 1.3.8 (Two-block dynamics). *Let $(\mathbb{X},\pi)$ be the product of two finite probability spaces $(\mathbb{X}_1,\pi_1)$ and $(\mathbb{X}_2,\pi_2)$. Let $\mathrm{Var}_1(f) = \mathrm{Var}(f(X_1,X_2)|X_2)$ and similarly for $\mathrm{Var}_2(f)$. Fix an event $\mathcal{X}\subset\mathbb{X}_1$. Then for any $f:\mathbb{X}\to\mathbb{R}$*

$$\mathrm{Var}_\pi(f) \leqslant \frac{\mathbb{E}_\pi(\mathrm{Var}_1(f)+\mathbb{1}_{\mathcal{X}}\mathrm{Var}_2(f))}{1-\sqrt{1-\pi(\mathcal{X})}} \leqslant \frac{2}{\pi(\mathcal{X})}\mathbb{E}_\pi\left(\mathrm{Var}_1(f)+\mathbb{1}_{\mathcal{X}}\mathrm{Var}_2(f)\right).$$

A way to interpret this is as a Poincaré inequality (i.e. bound on the relaxation time) for a chain which updates X_1 at rate 1 and updates X_2 at rate 1, provided that $\mathcal{X}$ occurs. The original proof of [88] is simple but spectral, so we prefer to provide a probabilistic one which will be extended in Chapter 2.

Proof. Couple two copies of the above chain, by attempting the same updates in both (see [256] for background on Markov chains). For this, use a graphical representation as in Section 1.2.2 attempting updates at X_1 and X_2 with rate 1, but deeming those in X_2 illegal if $\mathcal{X}$ does not occur. The two chains clearly coalesce as soon as we update X_1 so that $\mathcal{X}$ occurs and then immediately update X_2. Consider (legal or illegal) updates on X_2 preceded by an update at X_1. Their number up to time T is $\lfloor N/2\rfloor$ with N a Poisson random variable with mean T. Each one succeeds in coupling the chains independently with probability $\pi(\mathcal{X})$. The result then follows directly from the elementary fact that $\mathbb{E}(\lambda^N)=e^{-T(1-\lambda)}$ for any $\lambda\in(0,\infty)$. □

Returning to the upper bound on T_{rel} of East, we apply Lemma 1.3.8 with X_1 and X_2 the states of the halves of our current segment and $\mathcal{X}$ requiring the presence of an infection in the left half close to the right one. Taking into account that $\pi(\mathcal{X})$ quickly becomes close to 1 when the volume grows beyond $1/q$, this can be turned into a proof [88] of

$$T_{\mathrm{rel}} = \exp\left(\log^2(1/q)(1/(2\log 2)+o(1))\right).$$

Since [88] more precise results have been established [104] (see also [145] for a review on the East model). We will not require even the precise constant

Model	Timescale	Supercritical Unrooted	Supercritical Rooted	Trivial subcritical
BP	$\mu\left(\tau_0^{\mathrm{BP}}\right)$	$\Theta\left(q^{-\zeta}\right)$		∞
KCM	$\mathbb{E}_\mu(\tau_0)$	$q^{-\Theta(1)}$	$e^{\Theta(\log^2(1/q))}$	∞

Table 1.1 – Summary of one-dimensional universality Proposition 1.3.4 and Theorem 1.3.10.

in the exponent above for our purposes, but the bisection technique and two-block lemma will prove essential. In addition, we should point out that they have been successfully applied beyond the realm of KCM [52, 96]. A more detailed exposition of this technique can be found in [342], while its origins lie in [268, Proposition 3.5].

Universality

The difference in the behaviour of FA-1f and East imposes a further ramification of Definition 1.3.3 introduced in [269].

Definition 1.3.9. A one-dimensional supercritical update family $\mathcal{U}$ is *unrooted* if there exists a finite set $A \subset \mathbb{Z}$ such that $[A] = \mathbb{Z}$ and *rooted* otherwise.

That is, for rooted families infection can propagate indefinitely in one of the two directions. The importance of this distinction is captured in the following result summarised together with Proposition 1.3.4 in Table 1.1.

Theorem 1.3.10 (KCM universality in one dimension [265, 267, 269])**.** *For any one-dimensional update family $\mathcal{U}$ one of the following alternatives holds.*

(1) $\mathcal{U}$ is supercritical unrooted and $\mathbb{E}_\mu(\tau_0) = q^{-\Theta(1)}$.

(2) $\mathcal{U}$ is supercritical rooted and $\mathbb{E}_\mu(\tau_0) = e^{\Theta(\log^2(1/q))}$.

(3) $\mathcal{U}$ is trivial subcritical and $\mathbb{E}_\mu(\tau_0) = \infty$ for all $q < 1$.

This result can be proved essentially following what we did for our three example models above (recall Fig. 1.1). For the upper bound of (1) it suffices to move a sufficiently long infected segment instead of a single infection and proceed as for FA-1f. Similarly, for the upper bound in (2) we replace single infections by sequences of infections in the bisection proof for East. These are due to Martinelli, Morris and Toninelli [269]. The lower bound in (2) was deduced by Martinelli, Marêché and Toninelli [267] from a combinatorial result of Marêché [265] generalising Proposition 1.3.7. Conveniently, for the reader, she did write out a proof specifically for the one-dimensional case. We will outline her argument directly in two dimensions in Section 1.5.2, as it is at the base of Chapter 8.

Inhomogeneity

Before closing our warm-up one-dimensional Section 1.3, let us present the result of an improved bisection technique we develop in Chapter 2 [207]. As explained above, our focus is on homogeneous (i.e. translation invariant) models, but the inhomogeneous one-dimensional setting will prove useful for two-dimensional homogeneous models in Section 6.A.1, which is our motivation for introducing them. Nevertheless, the result and especially its quite simple proof are of independent interest.

Our setting is the following (see Section 2.1 for the formal definitions). Consider KCM

- on an arbitrary volume $L \subset \mathbb{Z}$, $1 \leqslant |L| \leqslant \infty$, which need not be an interval;
- with arbitrary boundary conditions;
- conditioned to belong to an arbitrarily chosen irreducible component of the state space;
- with arbitrary on-site finite state spaces, which may vary from site to site and need not have uniformly bounded size or atom probabilities, but the probability of being infected is uniformly bounded from below by $q > 0$;
- with arbitrary update rules, which may vary from site to site, but have a range uniformly bounded by $R < \infty$. Some sites may be completely unconstrained or, inversely, frozen.

In Chapter 2 [207] we prove in this setting that for some $C_R > 0$ depending only on R

$$T_{\mathrm{rel}} \leqslant (2/q)^{C_R \log(\min(2/q,|L|))}.$$

As we saw in Theorem 1.3.10(2), this is sharp for all homogeneous supercritical rooted KCM as $q \to 0$. In addition, it may come as a surprise that this is also sharp for some homogeneous supercritical unrooted KCM on intervals with binary states, despite the fact that on $\mathbb{Z}$ their relaxation time is only $q^{-\Theta(1)}$ (recall Theorem 1.3.10(1)).

Let us note that for such general KCM there are usually many irreducible components (there are always at least two, save for trivialities) and their combinatorial structure can be very intricate. They have proved hard to deal with due to the long-range dependencies they introduce, like those present in conservative KCM. Consequently, the only nontrivial case in which the relaxation in an irreducible component is under control [59] (see also [88, 89]) is the FA-1f model on an interval in its so-called ergodic component—the only nontrivial component for this KCM (corresponding to $\Omega\setminus\{\mathbf{1}\}$).

We direct the reader to [321, 324, 326] for inhomogeneous KCM, to [88, 89, 342] for KCM with various homogeneous rules and boundary conditions and to [269] and Chapter 3 for general state spaces. Yet, we emphasise that no two among: general state spaces, inhomogeneous rules and irreducible components have featured simultaneously until present. Formally, as we will see, non-interval domains, boundary conditions and irreducible components other than the ergodic one can be absorbed in the inhomogeneity of the rules but such arbitrarily inhomogeneous KCM have not been considered previously.

1.4 r-neighbour bootstrap percolation and FA-jf

We may now turn our attention to the most classical, central and, perhaps, instructive of all bootstrap percolation models and KCM. Namely the r-neighbour bootstrap percolation and associated FA-rf in d dimensions with $2d \geqslant r \geqslant 1$ and $d \geqslant 2$. Recall that their constraint is to have at least r infected neighbours. These models were introduced in [100][1] and [155] respectively. The most interesting case is $d = r = 2$ and, more generally, $d \geqslant r \geqslant 2$ and that is where attention has been primarily directed. But before it, let us eliminate the remaining cases, which are similar to what we saw in one dimension.

1.4.1 High thresholds

The least relevant models are those with $2d \geqslant r > d$. We already saw an instance with 2-neighbour bootstrap and FA-2f in one dimension. The phenomenology in higher dimensions is no different—the models with $2d \geqslant r > d$ should still be called *trivial subcritical*, as with positive probability τ_0 and τ_0^{BP} are infinite for any value of $q < 1$. The reason is the same as in one dimension: a hypercube of healthy vertices remains such for all times.

It is worth mentioning that these parameter values are actually being investigated until present. The viewpoint is, of course, different, making them the only 'interesting' ones. Namely, one is interested in the presence of an infinite connected component in $\mathbb{Z}^d \setminus [A]$. To summarise in a few words, the main conclusion is that they behave exactly like ordinary percolation. We direct the reader to [82, 83, 102, 111, 275] for (nonrigorous) work on the topic (see also [14]) and disregard these models for the rest of our considerations.

1.4.2 FA-1f

1-neighbour bootstrap percolation with $d \geqslant 2$ is no harder than it is in one dimension—the infection time is equal to the distance to the closest initial infection. FA-1f on the other hand is both harder to analyse and more useful. Initial results on it were obtained in [88] focusing on the relaxation time. There it was proved that T_{rel} is q^{-2} up to a logarithmic factor in two

[1] Let us note that, as far as the closure is concerned, this model is defined earlier, for instance in [303].

dimensions, but in higher dimensions non-matching powers of q were given as lower and upper bounds. Lower bounds were subsequently improved in [325] to show that $T_{\mathrm{rel}} = \Theta(q^{-2})$ in $d \geqslant 3$. The latter paper also shows that $\mathbb{E}_\mu(\tau_0)$ is equal to T_{rel} up to a $\Theta(1)$ factor in all dimensions $d \neq 2$ and up to the unknown logarithmic factor in two dimensions.

Before discussing the ideas behind these results, let us mention that the value of the exponent in higher dimensions has a clear intuitive interpretation making the debate prior to [325] attractive for physicists. Indeed, several authors made predictions about this exponent based on various nonrigorous techniques including simulations, renormalisation group and field theory. We will not attempt to give a full list, but we may indicate that different false predictions were given in [330, 362] and the correct one as above in [229]. One heuristic suggests that, as in one dimension, the infection time of the origin is the time needed for a random walk moving at rate q to travel the typical distance between the origin and the closest infection. This would give $T_{\mathrm{rel}} \approx \mathbb{E}_\mu(\tau_0) \approx 1/q^{1+2/d}$ for $d \geqslant 2$. An alternative argument stipulates that we rather need to wait until the walk has had the time to visit a number of sites equal to the inverse of the infection density, giving the correct result $1/q^2$ for $d \geqslant 2$ (up to logarithmic corrections).

Turning to the idea behind this result, since the argument is more involved for $\mathbb{E}_\mu(\tau_0)$, we focus on T_{rel}. The lower bound of [325] consists in examining the number of connected clusters of infections truncated at distance $1/q$ from the origin as test function f in Eq. (1.3). Since infections are rare, they are mostly isolated and $\mathrm{Var}(f)$ scales like q^{1-d}. Moreover, the number of clusters changes by at most $2d$ after a flip and only changes if the flip occurs at a site with two or more adjacent infections. Thus, contributions to the Dirichlet form only come from transitions with three infections at or around a given vertex, yielding $\mathcal{D}(f) \approx q^{3-d}$. Hence, Eq. (1.3) gives $T_{\mathrm{rel}} \geqslant 1/q^2$.

Rather than explaining the upper bound's original proof from [88], we will take a simpler but less direct route by defining a closely related model of *coalescing and branching simple exclusion process* (CBSEP) and then deducing the result on FA-1f. Essentially, FA-1f is CBSEP's evil twin lacking nice properties, but behaving exactly the same way.

An auxiliary model: CBSEP

Let $G = (V, E)$ be a connected graph. Minimum, maximum, and average degrees in G are denoted by $d_{\min}, d_{\max}$ and d_{avg}, respectively. The degree of $x \in V$ is denoted by d_x. For any $\omega \in \Omega = \{0,1\}^V$ and any vertex $x \in V$ we say that x is *filled* (resp. *empty*), or that there is a *particle* (resp. *hole*) at x, if $\omega_x = 1$ (resp. 0). We define $\Omega_+ = \Omega \backslash \{\mathbf{0}\}$ to be the event that there exists at least one particle. Similarly, for any edge $e = \{x, y\} \in E$ we refer to $(\omega_x, \omega_y) \in \{0,1\}^{\{x,y\}}$ as the state of e in ω and write $E_e = \{\omega \in \Omega : \omega_x + \omega_y \neq$

$0\}$ for the event that e is not empty (its vertices are not both empty).

Given $p \in (0,1)$, let $\pi = \bigotimes_{x\in V} \pi_x$ be the product Bernoulli measure in which each vertex is filled with probability p and let $\mu(\cdot) := \pi(\cdot|\Omega_+)$ (if G is infinite, then simply $\mu = \pi$). Given an edge $e = \{x, y\}$, we write $\pi_e := \pi_x \otimes \pi_y$ and $\lambda(p) := \pi(E_e) = p(2-p)$.

CBSEP is a continuous time Markov chain on Ω_+ for which the state of any edge $e \in E$ such that E_e occurs is resampled with rate one w.r.t. $\pi_e(\cdot|E_e)$. Thus, any edge containing exactly one particle moves the particle to the opposite endpoint (the *SEP move*) with rate $(1-p)/(2-p)$ and creates an extra particle at its empty endpoint (the *branching move*) with rate $p/(2-p)$. Moreover, any edge containing two particles kills one of the two particles chosen uniformly (the *coalescing move*) with rate $2(1-p)/(2-p)$. The chain is readily seen to be reversible w.r.t. μ and ergodic on Ω_+, because it can reach the configuration with a particle at each vertex. If $c(\omega, \omega')$ denotes the jump rate from ω to ω', the Dirichlet form $\mathcal{D}^{\mathrm{CBSEP}}(f)$ of the chain has the expression

$$\mathcal{D}^{\mathrm{CBSEP}}(f) = \frac{1}{2}\sum_{\omega,\omega'} \mu(\omega)c(\omega,\omega')\left(f(\omega') - f(\omega)\right)^2 = \sum_{e\in E} \mu(\mathbb{1}_{E_e} \operatorname{Var}_e(f|E_e)). \tag{1.5}$$

Notice that the branching and coalescing moves of CBSEP are exactly the moves allowed in FA-1f. Moreover, the SEP move for the edge $\{x, y\}$ from $(1,0)$ to $(0,1)$ can be reconstructed using two consecutive FA-1f moves, the first one filling the hole at y and the second one emptying x. If we also take into account the rate for each move, we easily get the following comparison between the respective Dirichlet forms (see Eq. (1.4) and e.g. [256, Chapter 13.4]): there exists an absolute constant $c > 0$ such that for all $f : \Omega_+ \to \mathbb{R}$ it holds that

$$c^{-1}\mathcal{D}^{\mathrm{FA-1f}}(f) \leqslant \mathcal{D}^{\mathrm{CBSEP}}(f) \leqslant c d_{\max} p^{-1} \mathcal{D}^{\mathrm{FA-1f}}(f), \tag{1.6}$$

setting the parameter q of FA-1f equal to the parameter p of CBSEP. In our application to FA-1f for $p \to 0$ only the upper bound, which we believe to be sharper, counts.

Although the two models are clearly closely related, we would like to emphasise that CBSEP has many advantages over FA-1f, making its study simpler. Most notably, CBSEP is *attractive* in the sense that there exists a grand-coupling (see e.g. [256]) which preserves the partial order on Ω given by $\omega \prec \omega'$ iff $\omega_x \leqslant \omega'_x$ for all $x \in V$ (as it can be readily verified via the construction of Section 3.5.1). Furthermore, it is also natural to embed in CBSEP a continuous time random walk $(W_t)_{t\geqslant 0}$ on G such that CBSEP has a particle at W_t for all $t \geqslant 0$. The latter is a particularly fruitful feature, which we will use in Section 3.5, and which is challenging to reproduce for FA-1f [59].

In view of the lower bounds of [325] explained above, Eqs. (1.3) and (1.6), in order to upper bound $T_{\mathrm{rel}}^{\mathrm{FA-1f}}$ and recover the results of [88] it suffices to prove the following.

Proposition 1.4.1. *Let $G = \mathbb{Z}^d$. If $d = 2$, then $T_{\mathrm{rel}}^{\mathrm{CBSEP}} \leqslant O(\log(1/p)/p)$, while $T_{\mathrm{rel}}^{\mathrm{CBSEP}} \leqslant O(1/p)$ for $d \geqslant 3$.*

A proof is supplied in Chapter 3 and Section 5.B jointly with Fabio Martinelli and Cristina Toninelli [215, 218] (up to minor modifications). In fact, we will prove much more. Namely, we will treat CBSEP on arbitrary graphs, establishing often sharp bounds on T_{rel}, but, more importantly, also on its logarithmic Sobolev constant.[2] A corollary of such stronger results and Eq. (1.6) is control of the mixing and L^2-mixing times of FA-1f. This recovers, strengthens and generalises results of Pillai and Smith [299, 300] proved in a different and somewhat more involved way.

In addition, in Chapter 3 and Section 5.B we study a generalised version of CBSEP with with general state spaces per site instead of $\{0,1\}$. For this generalised model we establish appropriate mixing time bounds crucial for Chapter 5. This is our original motivation for considering CBSEP in the first place.

1.4.3 2-neighbour bootstrap percolation

Returning to r-neighbour bootstrap percolation and FA-rf in d dimensions, the remaining values, $d \geqslant r \geqslant 2$, are called *critical*. Since they are the most studied, we will need to review the literature on them. We begin with bootstrap percolation as usual, since it is a prerequisite for FA-rf. More specifically, we focus on $d = r = 2$, leaving higher dimensions to Section 1.4.5.

As already noted, the 2-neighbour bootstrap percolation originates from [100] (see also [245, 303]). Initially it was believed that $q_{\mathrm{c}} > 0$ based on simulations (see [3] and references therein) with estimated values in $(0.035, 0.17)$. However, it was proved soon after [358] that in fact $q_{\mathrm{c}} = 0$. This was the first manifestation of what would grow to be called the *bootstrap percolation paradox* we will keep returning to. To give it in a somewhat simplistic sentence, it refers to the observation that predictions on bootstrap percolation based on simulations always fail, no matter how advanced rigorous results they take into account. An early discussion of this paradox concerning the above can be found in [360], while subsequent reassessments include [116, 187].

Coarse threshold

The first quantitative statement in the domain of bootstrap percolation, which naturally laid its foundations is due to Aizenman and Lebowitz [7]

[2] This constant is defined like the spectral gap in Eq. (1.3) with $\mathrm{Var}(f)$ replaced by the entropy $\mu(f^2 \log(f^2/\mu(f^2)))$ (see Section 3.2.1).

(for nonrigorous precursors see [253]). They proved that

$$\tau_0^{\mathrm{BP}} = \exp(\Theta(1/q)) \tag{1.7}$$

w.h.p. We provide a sketch of the argument, as it introduces ingredients essential to us. The first thing to note about 2-neighbour bootstrap percolation is that the closure of any set of infections is the smallest (in terms of inclusion) collection of rectangles (with sides parallel to the axes of the lattice) at graph distance at least 3 from each other containing the infections. Thus, the closure of any set can be determined via the following *rectangles process.* We start off with a collection of rectangles consisting of each of the initial infections. At each step we merge two of them at graph distance 2 or less, replacing them by the smallest rectangle containing their union. Repeating this until the process becomes stationary yields the collection of rectangles in the closure. A corollary of this process is the following fundamental lemma.

Lemma 1.4.2 (Aizenman–Lebowitz [7])**.** *We say that a rectangle R is* internally filled *(by the set A of initial infections), if $[A \cap R] = R$. If R is internally filled, then for every $k \leqslant \operatorname{long}(R)$ there exists an internally filled rectangle $S \subset R$ such that $k \leqslant \operatorname{long}(S) \leqslant 2k$, where* $\operatorname{long}(R)$ *denotes the number of sites on the longer side of R.*

Clearly, $\tau_0^{\mathrm{BP}} < \exp(c/q)$ implies that the origin belongs to an internally filled rectangle with long side at most $\exp(c/q)$ with c to be chosen appropriately later. Then Lemma 1.4.2 shows that within distance $\exp(c/q)$ of the origin there should be an internally filled rectangle R of long side of our choice up to a factor 2. The right side length to choose, which we refer to as *critical scale*, is $1/q$. Observing that such an internally filled rectangle cannot contain two consecutive healthy rows/columns, we get

$$\mu([A \cap R] = R) \leqslant \left(1 - (1-q)^{2\operatorname{long}(R)}\right)^{\lfloor \operatorname{long}(R)/2 \rfloor} = \exp\left(-\Theta(1/q)\right),$$

concluding the proof that $\tau_0^{\mathrm{BP}} \geqslant \exp(\Omega(1/q))$ w.h.p. by the union bound on all possible positions of R, choosing c small enough.

A matching upper bound is guided by a similar idea (explaining the title 'Metastability effects in bootstrap percolation' of [7]). We first make sure to internally fill a square of (supercritical) side, say, q^{-3} and then this *critical droplet* is likely to grow and infect the entire grid at roughly linear speed. The internal filling can be directly forced starting from one infection and asking for it to find another one on its right and top side on each line as it progressively infects a growing square. This has probability

$$q \prod_{k=1}^{\infty} \left(1 - (1-q)^k\right)^2 \approx \exp\left(2 \int_0^{\infty} \log\left(1 - e^{-qx}\right) \mathrm{d}x\right) = \exp\left(-\Theta(1/q)\right) \tag{1.8}$$

and thus is likely to occur within distance $\exp(C/q)$ of the origin for C large enough. We may then ensure that with overwhelming probability every vertical or horizontal line of length q^{-3} at distance at most $\exp(C/q)$ from the origin contains an infection, so that the critical droplet does grow roughly linearly until it engulfs the origin after time $\exp(O(1/q))$, proving Eq. (1.7).

Let us now see how to enhance Eq. (1.7) to $\mu\left(\tau_0^{\mathrm{BP}}\right) = \exp(\Theta(1/q))$ in the way we will systematically handle such improvements. Notice that the above argument in fact gives the upper bound on τ_0^{BP} not only with probability $1 - o(1)$ but $1 - \exp(-c/q^2)$ for some $c > 0$. It therefore suffices to show a very mild bound on the tail of τ_0^{BP}:

$$\mu\left(\tau_0^{\mathrm{BP}} > t^2\right) \leqslant C\left(e^{-ctq^3} + e^{-ct^2\exp(-C/q)}\right) \tag{1.9}$$

for some $C, c > 0$. To see this, simply note that the origin becomes infected in time at most t^2 if it is contained in an internally filled square of size t. The latter event can be guaranteed by the presence of an internally filled square of size $1/q^3$ and two paths of boxes of size $1/q^3$ containing an infection in every line, such that the paths cross the square of size t from top to bottom and left to right respectively. Since each box contains the required infections with high probability, the existence of such paths from a positive fraction of the boxes contained in the square of size t is exponentially likely in tq^3 (by standard percolation results, see e.g. [196]). Resorting to the Harris inequality [202] (see Section 6.1.1), we obtain the desired Eq. (1.9) from Eq. (1.8). In fact, it is not hard to improve the above argument to obtain exponential decay rather than stretched exponential and improve the q^3 rate. Arguments along these lines will be provided in more detail in Chapters 4 and 5 in the KCM setting (and in [272] for the rest of the argument), but the reader may also want to consult [18, 19, 284, 315] for more precise results on such decay rates in bootstrap percolation.

Sharp threshold

Naturally, following Eq. (1.7) the question of the day became determining the implicit constant. This came about in a breakthrough of Holroyd 15 years later [225], proving that

$$\mu\left(\tau_0^{\mathrm{BP}}\right) = \exp\left(\frac{\pi^2 + o(1)}{18q}\right). \tag{1.10}$$

We will prove stronger lower and upper bounds in the sequel, so it is useful to give an idea of the proof, which introduced several crucial techniques commonly used thereafter. As in the Aizenman–Lebowitz result, the main difficulty is controlling the probability of a rectangle of size roughly $1/q$ being internally filled. More precisely, Eq. (1.10) follows once we show that for R

of size C/q for C large

$$\mu([A \cap R] = R) \approx e^{-\pi^2/(9q)}. \tag{1.11}$$

To be precise, a version of Eq. (1.10) w.h.p. follows and one can deduce Eq. (1.10), recalling the crude Eq. (1.9).

The lower bound in Eq. (1.11) is not hard. Once again we start from a single infection and make it infect progressively larger rectangles. However, it grows by an amount larger than 1 in each direction before switching to the other. Namely, the right choice is to grow in steps of $1/\sqrt{q}$. The use of this is that we do not need an infection on every line, but on every second line. This is the origin of the constant $\pi^2/9$: it arises like the integral in Eq. (1.8), but for a function corresponding to the lack of two consecutive rows/columns of healthy sites. If one thinks about the two-term recurrence relation this function should come from (we only need to remember if an infection was found on the previous line or the one before it), it is not surprising that it appears as the root of a certain quadratic equation. The reader interested in the links of this function and its integral with integer partitions may consult [84, 227]. Actually, the sketch above is not quite the way the result is proved in [225], but anticipates [187] and Chapter 5 discussed below.

The converse bound is much harder. Roughly speaking, the basic idea of the proof in [225] is that, given an internally filled rectangle R, we would like to associate with R a 'hierarchy:' a constant-size rooted unary-binary tree of sub-rectangles that 'encodes' the way in which the set $A \cap R$ grows to infect the rest of R. The leaves of this tree correspond to small internally filled rectangles ('seeds'), a vertex with two children corresponds to two (not too small) rectangles merging to form a larger rectangle, and a vertex with one child corresponds to a rectangle 'growing on its sides' to fill a slightly larger rectangle. Crucially, we would like all of these (increasing) events to occur disjointly, so that we can apply the van den Berg–Kesten inequality [357] (see Section 10.2.3) to bound the probability of their intersection. Since there are few possible such hierarchies (since their size is bounded and there are few choices for each sub-rectangle), and each is (roughly speaking) at least as as unlikely as a single 'seed' growing to fill R, one can deduce a sufficiently strong bound on the probability that R is internally filled. This translates the idea that the least costly way for infecting a large rectangle is to start from a small one and build up the infection little by little, rather than starting from several places. Indeed, one may consider a rectangle called 'pod' with dimensions essentially given by the sum of dimensions of all seeds in the hierarchy and show that the probability of the hierarchy occurring is at most the product of the costs of all seeds multiplied by the probability of the pod expanding gradually to fill R.

Threshold window

Equation (1.10) might as well have been the end of the story, had it not been a new manifestation of the bootstrap percolation paradox. Numerical estimates [2, 5, 289] of the constant $\pi^2/18$ above had yielded less than half the correct value. This naturally leads to the question of how fast the convergence in Eq. (1.10) is (i.e. quantify the $o(1)$) and how sharp the transition is. Let us start with the latter, addressed contemporaneously with Eq. (1.10). In [30] it was noticed that a general result on the sharpness of transitions of Boolean functions [158] directly applies to bootstrap percolation on the torus to show that the probability of the entire torus becoming infected has a transition much sharper than what is displayed in Eq. (1.10).[3] With some work, unfortunately specific to the 2-neighbour model due to Lemma 1.4.2, results could be transferred to the grid [30]. A window in terms of the size of the box instead of q was established in [187] similarly to what we do below for the transition window for the infection time.

Proposition 1.4.3 (Time window)**.** *Let* $\varepsilon \in (0, 1/2)$, $T_\varepsilon = \min\{t : \mu(\tau_0^{\mathrm{BP}} < t) \geqslant \varepsilon\}$ *and* $T_{1-\varepsilon} = \max\{t : \mu(\tau_0^{\mathrm{BP}} < t) \leqslant 1 - \varepsilon\}$. *For 2-neighbour bootstrap percolation in two dimensions for* $q > 0$ *small enough* $T_{1-\varepsilon}/T_\varepsilon \leqslant q^{-3}$.

Sketch proof. Observe that $\exp(c/q) < \tau_0^{\mathrm{BP}} \leqslant T_\varepsilon$ with probability at least $\varepsilon/2$ as $q \to 0$ for c small enough (not depending on ε) by Eq. (1.7). But this event implies (recall Lemma 1.4.2) the existence of a very large internally filled rectangle containing the origin fully infected in time at most $2T_\varepsilon$. This in turn implies that with probability at least $1 - \varepsilon/2$ at time $\exp(2c/q)$ there is an infected rectangle of size $\exp(c/q)$ at distance at most $C_\varepsilon T_\varepsilon$ from the origin for some C_ε large enough. Since with probability $1 - o(1)$ it grows at speed at least $1/q^3$ until reaching the origin, we are done. □

Remark 1.4.4. We would expect that for some positive constants $c_\varepsilon, C_\varepsilon$ we have $c_\varepsilon \leqslant T_{1-\varepsilon}/T_\varepsilon \leqslant C_\varepsilon$, as in [187], but this would require more work. Indeed, one possibility is to establish a limit shape theorem for the infection induced by a large infected rectangle (see [188] for a related model along these lines). Such matters have been addressed abundantly for supercritical models under additional restrictive conditions in [183–185]. A statement along these lines features in [184, Theorem 4.2], but the proof seems to have never appeared.

Speed of convergence

In view of Eq. (1.7), Proposition 1.4.3, and Remark 1.4.4, the transition window is tiny compared to its location, so this does not provide an explanation of the discrepancy of Eq. (1.10) with simulation results. For that reason we

[3] See [43] for related considerations of noise sensitivity.

quantify the error term in Eq. (1.10), again contradicting simulation predictions [337] (see also [190] for more) and showing that the convergence is very slow.

Theorem 1.4.5 (Second term)**.** *For 2-neighbour bootstrap percolation in two dimensions it holds*

$$\mu\left(\tau_0^{\mathrm{BP}}\right) = \exp\left(\frac{\pi^2 - \Theta(\sqrt{q})}{18q}\right).$$

The upper bound was established by Gravner and Holroyd [187] and is based on the mechanism presented for Eq. (1.10). Roughly speaking, the main difference, which is at the origin of the negative sign of the second term, is taking entropy into account. More precisely, rather than growing our squares in steps of $1/\sqrt{q}$, we allow the exact length of these increments to vary, while being of order $1/\sqrt{q}$. The entropy gained from this is sufficient to outweigh the energetic cost of deviating from a square shape.

The lower bound is significantly harder and is the subject of Chapter 10 joint with Robert Morris [220]. It was preceded by a bound with an additional polylogarithmic factor due to Gravner, Holroyd and Morris [190] (see also [84] for partial progress on decreasing the exponent of the redundant factor, without removing it). They required two additional ideas w.r.t. Holroyd's work [225]. Firstly, they needed their seeds to be much smaller (of size $1/\sqrt{q}$, rather than $o(1/q)$), and rectangles to grow geometrically (rather than linearly) as a function of their height in the tree (hierarchy). As a result, the number of possible hierarchies became very large (too large to use a naive union bound), and to deal with this issue they partitioned the family of hierarchies according to the number of 'big' seeds. We will use refinements of both of these ideas in Chapter 10.

In order to prove the lower bound of Theorem 1.4.5, we can only afford to lose a factor of $\exp\left(O(1/\sqrt{q})\right)$ (in the expected number of 'satisfied' hierarchies), and since our hierarchies will typically have height $\Theta(1/\sqrt{q})$, this means that we can only allow ourselves a constant number of choices at each step, unless we 'pay' for extra choices via some unlikely event occurring. Fortunately, this is intuitively possible: the only things that could prevent us from choosing the next rectangle in an almost unique way are: (a) the existence of a 'double gap' of consecutive healthy rows or columns blocking the growth of the critical droplet, or (b) the merging of two reasonably large internally filled rectangles. Our challenge will be to show that we gain enough from these events to compensate for the extra choices we are forced to make.

To do so, we will need to encode the existence of double gaps in our hierarchies, which causes two immediate problems: the events cease to be increasing, and cease to occur disjointly. To avoid these issues we only use the fact that the double gaps are healthy in a single path through the hierarchy (which we call the 'trunk'); outside the trunk we use increasing events defined on the complement of the double gaps. In Section 10.3 we will state (and sketch the proofs of) a pair of technical 'crossing' lemmas which provide sufficiently strong bounds on the probabilities of these events. We remark that we gain from the existence of these double gaps in two distinct ways: they force us to find either two infected sites close together, or one infected site in a relatively small region, and when the rectangle is very large they are themselves unlikely to exist.

Bounding the expected number of 'satisfied' hierarchies with height at most $O(1/\sqrt{q})$ will then be relatively straightforward; unfortunately, however, the height is not always so small. In Section 10.5 we will have to deal with various other types of hierarchy: those with too many vertices, with too many (or too large) seeds, and those whose growth deviates from the diagonal by a macroscopic amount (see Lemma 10.4.12). One additional innovation that we will need in order to deal with this last case is Lemma 10.4.16, which provides us with two 'pods,' instead of the single pod required by Holroyd.

1.4.4 FA-2f

Moving on to FA-2f (again in two dimensions), the story is much shorter. Indeed, the analogues of all the results for 2-neighbour bootstrap were not known before the present contribution—from the 1988 Aizenman–Lebowitz [7] one, Eq. (1.7), to Theorem 1.4.5 completed here. Our task is then to review the only two previous rigorous results [88, 272] and copious nonrigorous ones. The reader may have already noticed the abundant occurrence of references to the work of Cancrini, Martinelli, Roberto and Toninelli [88], which supplied the first rigorous results for various KCM and the most central model, FA-2f, is no exception.

Background

As for bootstrap percolation, the initial expectation was that FA-2f would exhibit a nontrivial transition [155]. We should emphasise that here and in the other works to be quoted below, predictions were made, taking into account bootstrap percolation results already available. In particular, a transition was expected despite its absence in bootstrap percolation [156]. This was quickly dissipated by physicists [157, 306], though rigorous results in this direction came only two decades later [88] (see also [87]) along with the bisection technique. Denoting the semigroup of the KCM by $(P_t)_{t\geqslant 0}$, the *ergodicity critical parameter* is defined as

$$q_{\mathrm{c}} = \inf\left\{q > 0 : \forall f \in L^2(\mu), \lim_{t\to\infty} P_t f = \mu(f)\right\}.$$

It was proved in [88] that for all $\mathcal{U}$-KCM this transition coincides with the corresponding $\mathcal{U}$-bootstrap percolation transition (Eq. (1.1)), which is why we still denote it q_{c}. It also coincides with the more standard ergodic theory definition: for $q > q_c$ the eigenvalue 0 of $\mathcal{L}_{\mathcal{U}}$ is simple and, therefore, by the ergodic theorem we also have

$$q_c = \inf\{q > 0 : \mathbb{P}_\mu(\tau_0 < \infty) = 1\}.$$

The same paper also discarded the possibility that for FA-2f (and many other models, but not all $\mathcal{U}$-KCM, as we will see) e.g. the tail $\mathbb{P}_\mu(\tau_0 \geqslant t)$ of the infection time would decay as a stretched exponential.[4] The pure exponential decay they established was quite unexpected as numerous non-rigorous works had exhibited evidence of stretching, though with various stretching exponents [11, 86, 153, 154, 156, 157, 179, 180, 205, 306] according to [87, 309]. The exponential decay of the above quantity follows rather easily, once it is established that $T_{\mathrm{rel}} < \infty$, though this had seemingly eluded physicists, who also had various predictions for the scaling of T_{rel} as $q \to 0$, as we will see.

The last results of [88] for FA-2f are the quantitative bounds on T_{rel}

$$\exp\left(\frac{\pi^2 - o(1)}{18q}\right) = \Omega\left(\mu\left(\tau_0^{\mathrm{BP}}\right)\right) \leqslant \mathbb{E}_\mu(\tau_0) \leqslant T_{\mathrm{rel}}/q \leqslant \exp\left(O\left(1/q^5\right)\right), \tag{1.12}$$

in particular establishing that it is finite. The first two inequalities hold for any $\mathcal{U}$-KCM and are not hard. The upper bound is both harder and not useful to us, so we do not discuss it further. Unfortunately, Eq. (1.12) does not give the correct scaling of $\mathbb{E}_\mu(\tau_0)$ and leaves discriminating between the conflicting expressions suggested by physicists [86, 154, 157, 179, 180, 289, 306, 338, 345] remained an open problem (e.g. [309] asked for settling this controversy). Progress in this direction was made recently by Martinelli and Toninelli [272], who improved the upper bound to

$$\exp\left(\frac{O(\log^2(1/q))}{q}\right), \tag{1.13}$$

much closer to the lower one, but still inconclusive. Indeed, by 2019, when [272] was published, several (different) predictions not only for the presence or absence of a logarithmic factor but also on the potential sharp constant, based on Eq. (1.10), had been accumulated in 35 years. The proof of [272] is again not very useful to us, so we do not discuss it.

Before settling the matter, let us explain the different predictions. The first one appeared in [289], where, based on numerical simulations, a faster than exponential divergence in $1/q$ was conjectured. The first to claim an exponential scaling $\exp(\Theta(1)/q)$ was Reiter [306]. He argued that the infection process of the origin is dominated by the motion of 'macro-defects,' i.e. rare regions having probability $\exp(-\Theta(1)/q)$ and size $q^{-\Theta(1)}$ that move

[4] They worked rather with the first time when the origin changes state rather than becomes infected, but this is unimportant.

at an exponentially small rate $\exp(-\Theta(1)/q)$. Later Biroli, Fisher and Toninelli [345] considerably refined the above picture. They argued that macro-defects should coincide with the critical droplets of 2n-BP having probability $\exp(-\pi^2/(9q))$ and that the time scale of the relaxation process inside a macro-defect should be $\exp(c/\sqrt{q})$, i.e. sub-dominant with respect to the inverse of their density, in sharp contrast with the prediction of [306]. Based on this and on the idea that macro-defects move diffusively, the relaxation time scale of FA-2f in $d = 2$ was conjectured to diverge as $\exp(\pi^2/(9q))$ [345, Section 6.3]. Yet, a different prediction was later made in [338] implying a different scaling of the form $\exp(2\pi^2/(9q))$.

Result

Our result proved in Chapter 5 jointly with Fabio Martinelli and Cristina Toninelli [215] shows that the scaling prediction of [306, 345] is correct, contrary to those of [289, 338]. Moreover, our result on the characteristic time scale of the relaxation process *inside* a macro-defect (see Proposition 5.2.9) agrees with the prediction of [345] and disproves the one of [306].

Theorem 1.4.6. *As $q \to 0$ the stationary FA-2f model on $\mathbb{Z}^2$ satisfies:*

$$\mathbb{E}_\mu(\tau_0) \geqslant \exp\left(\frac{\pi^2}{9q}\left(1 - \sqrt{q}\cdot O(1)\right)\right), \tag{1.14}$$

$$\mathbb{E}_\mu(\tau_0) \leqslant \exp\left(\frac{\pi^2}{9q}\left(1 + \sqrt{q}\cdot (\log(1/q))^{O(1)}\right)\right). \tag{1.15}$$

Moreover, these also hold for τ_0 w.h.p.

In particular, recalling Theorem 1.4.5, we have w.h.p.

$$\mathbb{E}_\mu(\tau_0) = \tau_0^{1+o(1)} = \left(\mu\left(\tau_0^{\mathrm{BP}}\right)\right)^{2+o(1)} = \left(\tau_0^{\mathrm{BP}}\right)^{2+o(1)}.$$

This is the first sharp asymptotics of $\log \mathbb{E}_\mu(\tau_0)$ within the whole class of 'critical' KCM. It is established in Chapter 5.

Remark 1.4.7. Despite the resemblance, Theorem 1.4.6 is by no means a corollary of Theorem 1.4.5. While the lower bound Eq. (1.14) does indeed follow rather easily from Theorem 1.4.5 together with an improvement of the 'automatic' lower bound $\mathbb{E}_\mu(\tau_0) \geqslant \Omega\left(\mu\left(\tau_0^{\mathrm{BP}}\right)\right)$ from Eq. (1.12), the proof of Eq. (1.15) is much more involved. In particular, it requires guessing an efficient infection/healing mechanism to infect the origin, which has no counterpart in the monotone 2-neighbour bootstrap percolation model.

Behind Theorem 1.4.6: high-level ideas

The main intuition behind Theorem 1.4.6 is that for $q \to 0$ the relaxation to equilibrium of the stationary FA-2f process is dominated by the slow motion of unusually unlikely patches of infection, dubbed *mobile droplets* or just *droplets*. In analogy with the *critical droplets* of bootstrap percolation, mobile droplets have a linear size which is polynomially increasing in q (with some arbitrariness), i.e. they live on a much smaller scale than the metastable length scale $e^{\Theta(1/q)}$ arising in bootstrap percolation. One of the main requirements dictating the choice of the scale of mobile droplets is the requirement that the typical infection environment around a droplet is w.h.p. such that the droplet is able to move under the FA-2f dynamics in any direction. Within this scenario the main contribution to the infection time of the origin for the stationary FA-2f process should come from the time it takes for a droplet to reach the origin.

In order to translate the above intuition into a mathematically rigorous proof, one is faced with two different fundamental problems:

(1) a precise, yet workable, definition of mobile droplets;

(2) an efficient model for their 'effective' random evolution.

In [269, 272] and Chapter 4 mobile droplets (dubbed 'super-good' regions there) have been defined rather rigidly as fully infected regions of suitable shape and size and their motion has been modelled as a *generalised FA-1f process* on $\mathbb{Z}$ [269, Section 3.1]. In the latter process mobile droplets are freely created or destroyed with the correct heat-bath equilibrium rates but *only at locations which are adjacent to an already existing droplet.*

While rather powerful and robust, this solution has no chance to give the *exact* asymptotics of either (1), or (2) above. Indeed, a mobile droplet should be allowed to deform itself and move to a nearby position like an amoeba, by rearranging its infection using the FA-2f moves. This 'amoeba motion' between nearby locations should occur on a time scale much smaller than the global time scale necessary to bring a droplet from far away to the origin. In particular, it should not require to first create a new droplet from the initial one and only later destroy the original one (the main mechanism of the droplet dynamics under the generalised FA-1f process).

With this in mind we offer a new solution to (1) and (2) above which indeed leads to determining the exact asymptotics of the infection time. Concerning (1), our treatment in Section 5.2 consists of two steps. We first propose a sophisticated multiscale definition of mobile droplets which, in particular, introduces a crucial degree of *softness* in their microscopic infection's configuration.[5] The second and much more technically involved step is developing the tools necessary to analyse the FA-2f dynamics inside a mobile droplet. In particular, we then prove two key features (see Propositions 5.2.7 and 5.2.9):

[5] This construction is inspired by one suggested by P. Balister in 2017, which he conjectured would remove the spurious log-corrections in Eq. (1.13) available at that time.

(1.a) to the leading order the probability ρ_{D} of mobile droplets matches Eq. (1.11):

$$\rho_{\mathrm{D}} \geqslant \exp\left(-\frac{\pi^2}{9q} - \frac{O(\log^2(1/q))}{\sqrt{q}}\right),$$

(1.b) the 'amoeba motion' of mobile droplets between nearby locations occurs on a time scale $\exp(O(\log(1/q)^3)/\sqrt{q})$ which is sub-leading w.r.t. the main time scale of the problem and only manifests in the second term of Eq. (1.15).

Property (1.a) follows rather easily essentially as explained for Eq. (1.11), while proving property (1.b), one of the pivotal steps of this thesis, requires a substantial amount of new ideas.

While properties (1.a) and (1.b) above are essential, they are not sufficient on their own for solving problem (2) above. In Section 5.4 we propose to model (admittedly only at the level of a Poincaré inequality, which however suffices for our purposes) the random evolution of mobile droplets as a generalised CBSEP process (recall Section 1.4.2), studied in Chapter 3 for this purpose. Finally, the fact that CBSEP relaxes on a time scale proportional to the inverse density of mobile droplets (modulo logarithmic corrections—see Proposition 5.4.1) yields the scaling of Theorem 1.4.6. We emphasise that modeling the large-scale motion of droplets by generalised CBSEP instead of a generalised FA-1f process is an absolute novelty, also with respect to the physics literature.

1.4.5 Higher dimensions

For completeness, let us also mention progress for the cases $d \geqslant r > 2$ and $d > r = 2$, which fall outside the scope of our work. To begin with, $d > r = 2$ bootstrap percolation and FA can usually be treated in exactly the same way as the $d = r = 2$ case discussed above (particularly because Lemma 1.4.2 is still correct and stable infected sets are boxes). This is particularly true for upper bounds, where the only difference is that we look for one infection in a $(d-1)$-dimensional slice instead of a row/column. In terms of lower bounds, deducing FA-2f ones from 2-neighbour bootstrap percolation ones works identically. Bootstrap percolation lower bounds offer more difficulty, in particular because a box now has several aspect ratios. As a result, upper bounds corresponding to the Gravner–Holroyd and lower ones on the level of Holroyd's are the state of the art [32, 354], though it seems plausible that a lower bound as in Theorem 1.4.5 might be accessible using the techniques already available in [32] and Chapter 10.

For $r > 2$ bootstrap percolation matters are more difficult, but the same level of precision as for $d > r = 2$ has been attained [32, 354]. The intuition to keep in mind (particularly relevant for upper bounds) is that on the side of a large infected cube we see essentially a $(r-1)$-neighbour bootstrap percolation process, albeit with certain sites further helped by slices neighbouring the side. The bounds are then unsurprisingly $r-1$ times iterated exponential in a power of q. Historically, the nontriviality of the transition was proved in [315], the iterated exponential upper bound in [360], the coarse threshold in [97, 98], the sharp threshold in [32, 33] (also see [226]) and finally a second order upper bound in [354].

For FA-jf the situation is quite favourable due to the low precision associated to an iterated exponential scaling (error terms are in the top exponent). Indeed, lower bounds follow from those in bootstrap percolation already using the trivial bound $\mathbb{E}_\mu(\tau_0) \geqslant \Omega\left(\mu\left(\tau_0^{\mathrm{BP}}\right)\right)$ from Eq. (1.12) without the need of the refinement of Chapter 5. Upper bounds are also much easier and paradoxically more precise than for FA-2f. Namely, in an upcoming work [217] we prove upper bounds of the exact same shape as those for bootstrap percolation, improving previous bounds from [88, 272]. Additionally, the proofs are simpler, as it is possible to use rigid droplets like [272] and for $r \geqslant 4$ it is even possible to move them very brutally without affecting the final result.

1.5 Rough universality in two dimensions

We already saw our first universality results in Section 1.3. Ideally, we would like to have similar results in all dimensions determining the scaling of $\mu\left(\tau_0^{\mathrm{BP}}\right)$ and $\mathbb{E}_\mu(\tau_0)$ as $q \to 0$ for all models in terms of the geometry and combinatorics of the rules. Since any model can be embedded in higher dimensions the classification necessarily becomes more and more ramified as the dimension increases. Consequently, we will recover our supercritical and trivial subcritical classes in all dimensions, along with various other classes absent in one dimension. As of now, only the two-dimensional case is treated and that is the one we will focus on. Nevertheless, some results in higher dimensions can be obtained with the same techniques. Indeed, on the bootstrap percolation side universality results have been announced as upcoming since the early 2010s and conjectures can be found in [70, 74, 280, 281]. For KCM, we expect that relatively little work will be needed to complete the rough universality picture once bootstrap percolation is dealt with. Of course, refined universality for critical models in dimension 3 and higher is currently well beyond fiction for bootstrap percolation, let alone KCM. Therefore, for the remainder of Chapter 1 we work in two dimensions.

We next define the rough universality classes. Let $\|\cdot\|$ and $\langle\cdot,\cdot\rangle$ denote the Euclidean norm and scalar product respectively. Let $S^1 = \{u \in \mathbb{R}^2 : \|u\| = 1\}$ be the unit circle identified with $\mathbb{R}/2\pi\mathbb{Z}$ when needed, whose elements we call *directions*. The open and closed half-planes with outer normal $u \in S^1$ and offset $l \in \mathbb{R}$ are

$$\mathbb{H}_u(l) = \{x \in \mathbb{R}^2 : \langle x, u\rangle < l\}, \qquad \overline{\mathbb{H}}_u(l) = \{x \in \mathbb{R}^2 : \langle x, u\rangle \leqslant l\}.$$

We omit l if it is equal to 0. Unless confusion arises, we will identify all subsets of $\mathbb{R}^2$, such as $\mathbb{H}_u$, with their intersection with $\mathbb{Z}^2$.

Definition 1.5.1 (Stable directions)**.** Fix an update family $\mathcal{U}$. A direction $u \in S^1$ is *unstable* if there exists $U \in \mathcal{U}$ such that $U \subset \mathbb{H}_u$ and *stable* otherwise.

The relevance of this definition and the terminology come from the fact that $[\mathbb{H}_u]_{\mathcal{U}} = \mathbb{H}_u$ if u is stable (i.e. u is stable iff $\mathbb{H}_u$ is) and $[\mathbb{H}_u]_{\mathcal{U}} = \mathbb{Z}^2$ if u is unstable. We say that a direction $u \in S^1$ is *rational* if $\mathbb{R}u \cap \mathbb{Z}^2 \neq \varnothing$. It is also not hard to check that the set of stable directions is a finite union of closed intervals of S^1 with rational endpoints. Endpoints of intervals of stable directions are called *isolated* if the interval is reduced to a point and *semi-isolated* otherwise. All stable directions which are neither isolated nor semi-isolated are called *strongly stable*. An illustration of the above in eleven examples can be found in Fig. 1.2, which we will explain progressively up to Section 1.6.1. Note that unstable directions can be determined by looking at each rule separately and taking the union.

We are ready for the rough universality classes consistent with Definitions 1.3.3 and 1.3.9. Indeed, one may consider stable/unstable directions also in one dimension, but, since there are only two meaningful directions, it would be somewhat less natural to define classes this way.

Definition 1.5.2 (Rough universality partition)**.** Let $\mathcal{C} = \{\mathbb{H}_u \cap S^1 : u \in S^1\}$ denote the set of open semicircles of S^1. An update family $\mathcal{U}$ is:

- *supercritical* if there exists an open semicircle $C \in \mathcal{C}$ whose directions $u \in C$ are all unstable. If additionally
  - there exist two non-opposite stable directions, $\mathcal{U}$ is *rooted*;
  - there do not exist two non-opposite stable directions, $\mathcal{U}$ is *unrooted*.
- *critical* if every open semicircle contains a stable direction and there exists a semicircle containing finitely many stable directions.
- *subcritical* if every semicircle contains infinitely many stable directions. It is
  - *nontrivial* if there exists an unstable direction;
  - *trivial* if all directions are stable.

Their relevance is clear from the following result generalising Proposition 1.3.4 and Theorem 1.3.10. It is summarised in Table 1.2.

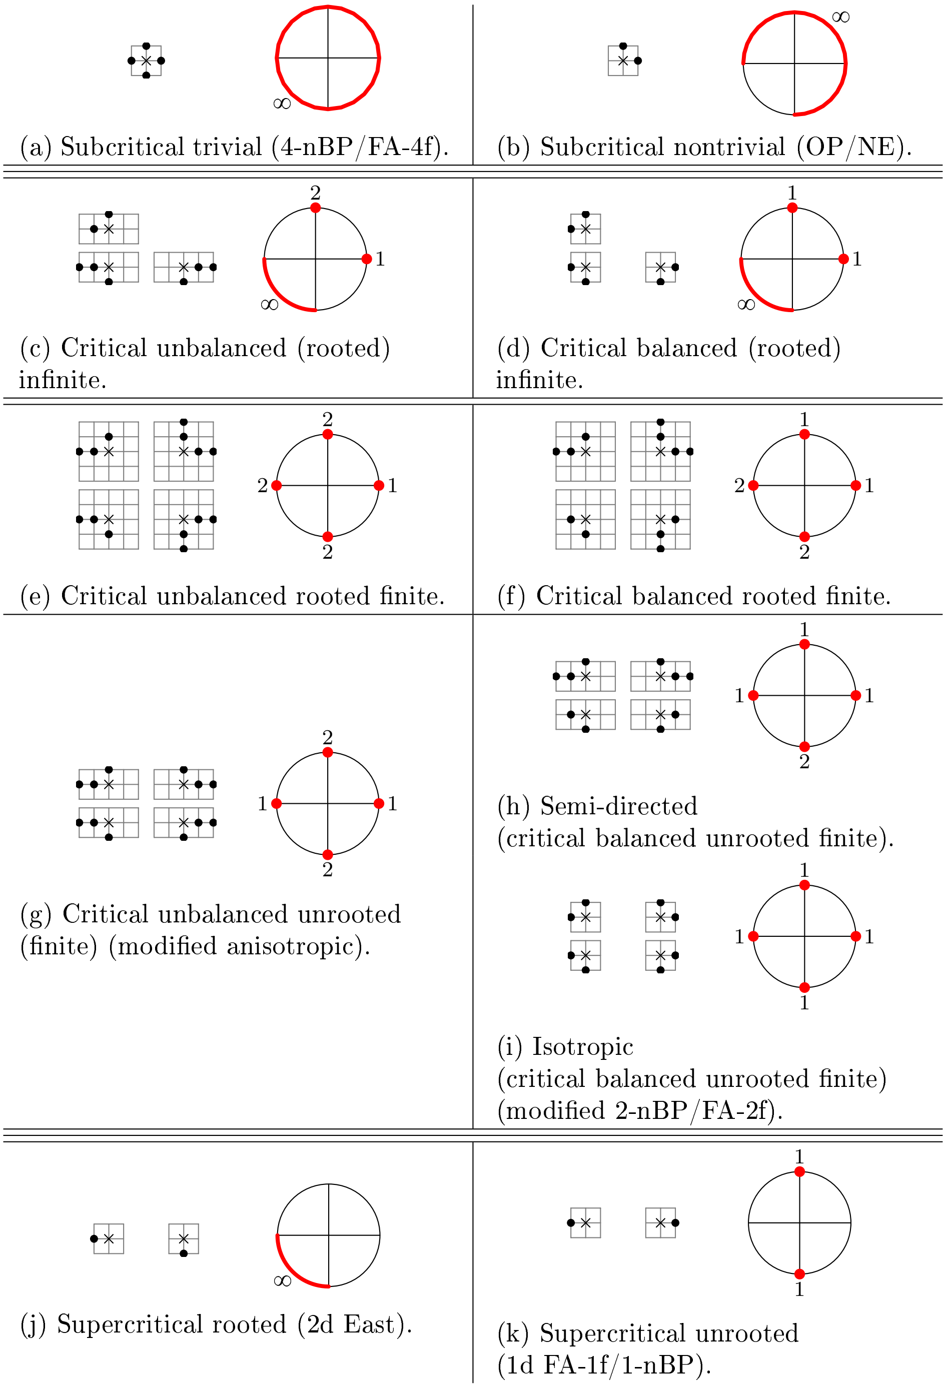

(a) Subcritical trivial (4-nBP/FA-4f).

(b) Subcritical nontrivial (OP/NE).

(c) Critical unbalanced (rooted) infinite.

(d) Critical balanced (rooted) infinite.

(e) Critical unbalanced rooted finite.

(f) Critical balanced rooted finite.

(g) Critical unbalanced unrooted (finite) (modified anisotropic).

(h) Semi-directed (critical balanced unrooted finite).

(i) Isotropic (critical balanced unrooted finite) (modified 2-nBP/FA-2f).

(j) Supercritical rooted (2d East).

(k) Supercritical unrooted (1d FA-1f/1-nBP).

Figure 1.2 – Representatives of universality classes of $\mathcal{U}$-KCM. Update rules are depicted on the left with 0 marked by a cross and the sites of the rule denoted by dots. Stable directions are thickened on the right with their difficulties indicated. Isolated stable directions are marked by dots. All critical models have difficulty $\alpha = 1$ witnessed by the right-hand open semicircle.

	Supercritical		Critical	Subcritical	
	Unrooted	Rooted		Nontrivial	Trivial
q_c	0	0	0	$\in (0,1)$	1
$\mu\left(\tau_0^{\mathrm{BP}}\right)$, τ_0^{BP}	$q^{-\Theta(1)}$	$q^{-\Theta(1)}$	$e^{q^{-\Theta(1)}}$	∞	∞
$\mathbb{E}_\mu(\tau_0)$, τ_0, T_{rel}	$q^{-\Theta(1)}$	$e^{\Theta(\log^2(1/q))}$	$e^{q^{-\Theta(1)}}$	∞	∞

Table 1.2 – Summary of Theorem 1.5.3. Critical parameter and characteristic timescale of rough universality classes as $q \to 0$. Cf. Table 1.1.

Theorem 1.5.3 (Rough universality [28, 74, 265, 267, 269])**.** *Let $\mathcal{U}$ be a two-dimensional update family. If $\mathcal{U}$ is*

- *supercritical unrooted, then $q_c = 0$, $\mu\left(\tau_0^{\mathrm{BP}}\right) = q^{-\Theta(1)}$ and $\mathbb{E}_\mu(\tau_0) = q^{-\Theta(1)}$.*
- *supercritical rooted, then $q_c = 0$, $\mu\left(\tau_0^{\mathrm{BP}}\right) = q^{-\Theta(1)}$ and $\mathbb{E}_\mu(\tau_0) = \exp(\Theta(\log^2(1/q)))$.*
- *critical, then $q_c = 0$, $\mu\left(\tau_0^{\mathrm{BP}}\right) = \exp(q^{-\Theta(1)})$ and $\mathbb{E}_\mu(\tau_0) = \exp(q^{-\Theta(1)})$.*
- *subcritical nontrivial, then $q_c \in (0,1)$, $\mu\left(\tau_0^{\mathrm{BP}}\right) = \mathbb{E}_\mu(\tau_0) = \infty$ for q small enough.*
- *subcritical trivial, then $q_c = 1$, $\mu\left(\tau_0^{\mathrm{BP}}\right) = \mathbb{E}_\mu(\tau_0) = \infty$ for all $q \in (0,1)$.*

The same asymptotics hold for τ_0^{BP} and τ_0 w.h.p. as $q \to 0$. The same asymptotics as for $\mathbb{E}_\mu(\tau_0)$ apply to T_{rel}, while

$$\limsup_{t\to\infty} \frac{\log \mu\left(\tau_0^{\mathrm{BP}} > t\right)}{t} \leqslant -q^{O(1)} \tag{1.16}$$

for all critical and supercritical families.

Remark 1.5.4. To be precise, in bootstrap percolation [74] only w.h.p. bounds were proved, while those in expectation and exponential decay rates are new. In fact, the corresponding qualitative statement is discussed in Chapter 12 and a few particular cases in [315]. It is interesting to note that for supercritical or critical models with two opposite stable directions it is easy to prove a lower bound for lim inf instead of lim sup matching Eq. (1.16) (consider a healthy strip perpendicular to those directions). However, for models without such directions we rather expect $\log \mu(\tau_0^{\mathrm{BP}} > t) \approx -t^2$, though we will not prove such upper bounds. For example, the 1-neighbour model clearly satisfies $\mu(\tau_0^{\mathrm{BP}} > t) = (1-q)^{2t^2+2t+1}$.

Sections 1.5.2 to 1.5.4 are dedicated to supercritical, critical and subcritical models respectively, also discussing the proofs of Theorem 1.5.3. However, before that, in Section 1.5.1, we outline the history of the universality setting.

1.5.1 History of universality

Bootstrap percolation

Initially, variants of the original r-neighbour model from [100] were introduced to explore the effect of rules on the behaviours, beyond simply varying the parameter r, e.g. [123]. Schonmann [313, 315] studied $\mathcal{U}$-bootstrap percolation with rules $U \in \mathcal{U}$ contained in the set of nearest neighbours of the origin. He appears to have been the first to deal with arbitrary $\mathcal{U}$-bootstrap percolation models in two and more dimensions in search of universality. Although nearest neighbour rules do provide representatives of many of the various classes we will be led to consider, the picture is quite degenerate. Moreover, his treatment is based to some extent on exhaustively treating models one by one (or comparing them when possible) and does not provide major insight on what general principles govern their behaviour. Nevertheless, Schonmann [315] did correctly discern the distinction between trivial subcritical and other models in general. He further contributed to the first insight on representatives of both the subcritical nontrivial class and critical 'unbalanced' one we will consider below [313, 360].

A further step was made by Gravner and Griffeath in the series of works we already mentioned [181–185]. They not only considered the $\mathcal{U}$-bootstrap percolation model with arbitrary range, but also introduced certain universality classes and some of the crucial notions we use in Section 1.6.1 [184]. An additional motivation [183] were well-known systematic phenomenological explorations of cellular automata, whose taxonomy is still used today (see [295, 366]), classifying them according to qualitative features observed. Unfortunately, Gravner and Griffeath focused mostly on update families defined by a neighbourhood and a threshold number of infections in that neighbourhood needed for infection and would impose symmetry and more technical conditions. Their emphasis was on the combinatorics of supercritical models, which were not the ones of central interest for physicists. They made little progress on critical ones and essentially excluded subcritical ones from their studies.

This general setting for critical models was revived by Duminil-Copin and Holroyd [125] in a preprint claiming sharp thresholds for a somewhat general class of models resembling the 2-neighbour one. To quote their abstract, 'This article represents a further step towards an understanding of universality of two dimensional bootstrap models.' and indeed this was the case. Nevertheless, their work only concerned a 'small' subclass of what we will call 'isotropic' critical models in the sequel (further generalisations remain open).

The first truly universal results were those of Bollobás, Smith and Uzzell [74], who insisted on not imposing any additional assumption, most notably symmetry, on their update families $\mathcal{U}$ (other than being composed of a finite number of finite rules). This was natural, since a number of specific models studied over the years [123, 184, 235, 236, 315, 346, 347] were neither supercritical, nor trivial subcritical, nor isotropic, often due to their lack of symmetry. Thus, [74] sought to unify all existing models in a single theorem.

Sadly, the results of [74], based on the proof scheme of [7], were much weaker than a coarse threshold in general. Namely, for each concerned model previously considered in the literature the very first results beyond $q_c = 0$, were systematically more precise than the universal one [7, 286, 360, 361].[6] Hence, the universal result was interesting mostly for its generality and its simple classification, significantly clarifying the now obsolete one of [184]. In fact the classification suggested in [74] into supercritical, critical and subcritical models did not become fully legitimate until Balister, Bollobás, Przykucki and Smith [28] supplied a proof that it indeed correctly identifies the distinct rough behaviours possible.

Universality was brought to an entirely different level by the work of Bollobás, Duminil-Copin, Morris and Smith [70], who proved a coarse threshold result for all critical models (see Section 1.6). This required subdividing the critical class in two refined classes, again clarifying considerations of Gravner and Griffeath. Thus, only models for which a sharp threshold or more is established, [69, 84, 128, 129, 225, 227] and Theorem 1.4.5, remain uncovered by their general result. While these are again essentially all common models, it should be noted that sharp thresholds are much more recent, harder and model-specific, leaving relatively little hope of a universal statement at that level. As noted above, the only step in this direction was made in [125], some of whose ideas were incorporated into [70] (as well as [32]).

KCM

$\mathcal{U}$-KCM were first considered by Cancrini, Martinelli, Roberto and Toninelli [88]. In addition to providing essentially the first rigorous results beyond the East model and several indispensable techniques such as the bisection one. Interestingly, though they were unaware of universality considerations of Gravner and Griffeath, they sought to treat representatives of as many different behaviours as possible (e.g. all those in Table 1.2). However, universality results were not sought until after all bootstrap percolation ones [28, 70, 74] were available. Indeed, we will see that a good understanding of those is a prerequisite for KCM universality.

The quest for establishing universality for KCM was taken up by Martinelli and Toninelli [272]. At the time subcritical models were quite automatically identified as a rough universality class for KCM by the results of [28, 74, 88], so the focus was on the remaining supercritical and critical models. Together with Morris [269] and with Marêché [265, 267] they did esta-

[6] We thank Aernout van Enter for pointing out that for the Duarte model both [286, 360] were aided by Schonmann.

blish an analogue of the rough result of [74], confirming the one-dimensional partition of supercritical models into rooted and unrooted (established precisely at this point together with its two-dimensional version).

1.5.2 Supercritical models

Supercritical models are dealt with quite similarly to Section 1.3 in one dimension. This is not surprising given that no further ramification arises in two (or more) dimensions. Additionally, the definition of supercritical families says precisely that there is a direction in which they can propagate (the midpoint of a semicircle of unstable directions), like in the one-dimensional case.

The lower bound $\tau_0^{\mathrm{BP}} \geqslant \Omega(1/\sqrt{q})$ w.h.p. is immediate, given that infections can only propagate at linear speed and typically none is present closer than that. Upper bounds for bootstrap percolation are also immediate—it suffices to observe that a conveniently shaped polygon of infections can reproduce itself translated in a rational direction which is the midpoint of an open semicircle of unstable directions (see [74] for details). This readily entails that τ_0^{BP} is upper bounded by a geometric random variable with mean $q^{-O(1)}$ and upper bounds w.h.p., expectation and Eq. (1.16) follow. As mentioned above, more advanced results related to limit shapes, precise scaling or combinatorial problems (e.g. determining the smallest number of infections able to create infinitely many) for supercritical models can be found in [181–185].

The proof of the upper bounds in Theorem 1.5.3 for KCM are also quite straightforward, given those of Theorem 1.3.10, which do require some work. Indeed, for unrooted models we can find two opposite open semicircles of unstable directions (recall that we view S^1 both as the unit circle and $\mathbb{R}/2\pi\mathbb{Z}$). Restricting our attention to a thick strip parallel to their midpoints, we may reason exactly as in one dimension (that is, as for the FA-1f model). The upper bound for supercritical rooted families can be proved in the same way, considering a strip parallel to the midpoint of a semicircle of unstable directions. See [269] for more details.

The lower bound for supercritical unrooted KCM follows from bootstrap percolation via the automatic bound $\mathbb{E}_\mu(\tau_0) \geqslant \Omega(\mu\left(\tau_0^{\mathrm{BP}}\right))$ from Eq. (1.12). Thus, we are left with proving that $\tau_0 \geqslant \exp(\Omega(\log^2(1/q)))$ w.h.p. for supercritical rooted KCM. Here again we can copy the proof for the East model sketched in Section 1.3.2 to see that only one difficulty remains—proving an analogue of the combinatorial bottleneck Proposition 1.3.7. This is the hardest and most novel part of Theorem 1.5.3 for supercritical families w.r.t. Theorem 1.3.10 and was proved by Marêché [265]. As we will need a substantially more sophisticated version of her argument in Chapter 8 to deal with critical KCM (see Section 1.6), we provide a sketch of her proof in the case of the two-dimensional East model depicted in Fig. 1.2j for concreteness.

We aim to show that in order for an infection to reach the center of a box R of size 3^n initially fully healthy it is necessary to visit a configuration with at least n infections simultaneously present in the box. This is achieved by an inductive argument that we describe next. It is enough to show that if R is initially fully healthy and we can only have strictly less than n infections at the same time in R, we cannot reach a configuration with all infections in the middle of R. Indeed, this implies that there are strictly less than $n-1$ infections at the same time in the middle of R, and one can make an induction on n. To show that we cannot reach a configuration with all infections in the middle of R, by reversibility we may instead prove that if we start with infections only in the middle of R, but we are never allowed to have n infections simultaneously in R, we cannot reach a configuration fully healthy in R. The idea is to ensure that for any legal path of the dynamics (recall Definition 1.3.6) starting with all infections in the middle of R in which we never have n infections at the same time in R, the following two conditions remain true at all times. Firstly, a buffer zone (see the shaded frame B in Fig. 8.1) with no infections remains intact. Secondly, there is always an infection in the internal region encircled by the buffer, so the dynamics cannot reach a configuration completely healthy in R.

In order to achieve that, we use a second induction, on the number of updates along the path already performed. We know that so far an infection remains trapped in the internal region encircled by the buffer, so we only have $n-1$ infections available for disrupting the buffer from the outside, which is impossible by induction on n. Therefore, it suffices to show that we may not disrupt the buffer from the inside either. By projecting the two-dimensional East model on each axis it is clear that no infection can enter the left and bottom parts of the buffer from the inside, and the projections of the lowest and leftmost particles in the region inside the buffer need to remain where they were initially. The right part of the buffer (and similarly the top one) cannot be reached from the inside, because at least one infection needs to remain as far left as the leftmost initial one was, so we only have $n-1$ infections with which to reach the right part of the buffer, which is impossible by induction on n. This completes the sketch of the result of [265] and, therefore, of Theorem 1.5.3 for supercritical models.

1.5.3 Critical models

For the moment we restrict our goal to sketching the proof of Theorem 1.5.3 for critical families, leaving stronger results to Section 1.6, of course, reusing some of the techniques, primarily those from bootstrap percolation. The intuition about critical families should be roughly the following, especially for bootstrap percolation. The behaviour is the same as for supercritical ones except that now we need infected droplets (polygons) of size $1/q^{\Theta(1)}$ rather than constant, able to propagate in the direction of the midpoint of a semicircle with finitely many stable directions. Thus, zooming the picture out, so as not to notice microscopic details and so that droplets look like finite-sized regions, we recover to a large extent the phenomenology of a corresponding supercritical model with all isolated stable directions removed. The expressions of $\mu\left(\tau_0^{\mathrm{BP}}\right)$ and $\mathbb{E}_\mu(\tau_0)$ are then intuitive, as the supercritical model to which the critical model morally renormalises has infections (representing polynomially sized infected regions) with density $q_{\text{eff}} = \exp(-1/q^{\Theta(1)})$.

The upper bound $\tau_0^{\mathrm{BP}} \leqslant \exp(q^{-O(1)})$ follows similarly to 2-neighbour bootstrap percolation discussed in Eq. (1.7). Starting from a large constant-sized infected polygon (with appropriately chosen sides), we can make it grow in a direction given by the midpoint of a semicircle with finitely many stable directions. However, since some of the directions are stable a little help is needed. Namely, to append a line of sites of $\mathbb{Z}^2$ to the side of the polygon perpendicular to an isolated stable direction, it suffices to find several consecutive infections along that line. However, notice that if both endpoints of the semicircle are semi-isolated stable directions, the growing droplet only grows in one direction and the sides to which we want to append lines remain constant, obstructing the direct application of Eq. (1.8). To remedy this, we note that every now and then a group of infections will be found allowing the extension of the side perpendicular to the last direction of the semicircle. Taking this into account, the proof for 2-neighbour goes through, since droplets now do grow to infinity with probability $\exp(-1/q^{O(1)})$ (see [74] for the details). It should be noted that, though not needed there, the mechanism of finding infections on the side of the growing droplet in order to make it grow a bit more than one-dimensionally is also applicable to supercritical models.

To prove Eq. (1.16), from which the bound in expectation also follows, one may consider boxes of size $q^{-\Theta(1)}$, elongated polynomially in q in the direction of growth, arranged in a brick wall fashion (see [269]). Then it suffices to find one which is entirely infected and such that a thick oriented path of boxes containing typical amounts of sets of consecutive infections sufficient to sustain the growth, emanates from this initial droplet. Any two such paths forming a wedge become infected and everything between them also does in roughly linear time.

This mechanism is the driving principle of the proof of the upper bound of Theorem 1.5.3 for critical KCM in [269, 272]. The proof was really completed in [269], but most of the effort there was directed to obtaining the correct power of q in the upper bound, which as we will see was mostly unsuccessful. Instead, the sketch we present is, perhaps, more in the spirit of the earlier [272]. The idea is to concentrate on a path as above starting from a fully infected box and propagate that box along the path. Each step of this movement can be performed as in bootstrap percolation. However, we cannot afford to infect the entire path, as that is too unlikely. Instead, we can proceed in an East-like fashion, leading to an upper bound of $\exp(O(\log^2(1/q_{\text{eff}})))$ with $q_{\text{eff}} = \exp(-1/q^{O(1)})$ the density of fully infected boxes.

Turning to lower bounds (for bootstrap percolation and thus for KCM, using the comparison), a generalisation of the rectangles process is needed. It was introduced in [74] under the name of *covering algorithm* and proceeds roughly as follows. Consider each initial infection and cover it by a large constant-sized polygon (with appropriate shape, as usual). Then merge any two polygons which are close to each other into the smallest polygon with sides parallel to the original ones containing both. The outcome of this process should then contain the closure of the original set of infections. To conclude two ingredients are needed. First, an analogue of the Aizenman–Lebowitz Lemma 1.4.2, based on the fact that the result of merging two polygons is not much bigger than the sum of the two for an appropriate notion of size (think of the diameter). The other ingredient, which also follows from this subadditivity, is that the number of infections required to obtain a large polygon in the output is at least linear in the size of the final polygon. With this at hand, we have both the necessity of the presence of a large covered droplet (produced by the covering algorithm) and a bound on its probability of occurrence, which conclude the proof by the union bound as for 2-neighbour bootstrap percolation. This proof scheme will be of use to us in Chapters 7 and 8.

1.5.4 Subcritical models

Let us rid ourselves also of trivial subcritical models, which, as the name indicates, are not of interest to us and are very easy to handle. Since all directions are stable, it is not hard to check that there exist finite polygons, which remain healthy forever if they are such initially. Hence, nothing new is gained w.r.t. the one-dimensional case.

Phase transition and oriented percolation

Nontrivial subcritical models, which we will call subcritical for short, assuming nontriviality unless otherwise stated, only arise in two and more dimensions. Before studying them in general, let us focus on the first and most fundamental example of subcritical model. It is defined by $\mathcal{U}^{\text{OP}} = \{\{(1,0),(0,1)\}\}$ (see Fig. 1.2b) and we will call it *oriented percolation* (OP) for reasons to become apparent, while its KCM version is known as the North-East model (NE). In bootstrap percolation OP was first considered in [313] and NE in [307] immediately after. Yet, OP was already very well studied as an ordinary percolation model (see [131, 258], as well as Chapter 11).

The equivalence between bootstrap percolation with the above update family and OP is the following (noticed in [313, 315]). A site $x \in \mathbb{Z}^2$ becomes infected at time t iff the longest path of initially healthy vertices going up or right at each step, starting at x, has length t. In particular, $\tau_0^{\mathrm{BP}} = \infty$ iff the origin belongs to such an infinite healthy path. In particular, among many other things, it is well known that $0 < q_c(\mathcal{U}^{\mathrm{OP}}) < 1$. Schonmann also observed that this immediately implies that $q_c(\mathcal{U}) < 1$ for any $\mathcal{U} \supset \mathcal{U}^{\mathrm{OP}}$.

This observation generalises to all subcritical models as follows. By definition any (nontrivial) subcritical family $\mathcal{U}$ has an unstable direction $u \in S^1$. Thus, there exists an update rule $U \in \mathcal{U}$ contained in the corresponding open half-plane $\mathbb{H}_u$. We call the bootstrap percolation model associated to the (one-rule) update family $\{U\}$ a *generalised oriented site percolation* model (GOSP) for any U contained in some open half-plane satisfying that U generates $\mathbb{R}^2$ as a vector space (otherwise the original model must be supercritical). As the name suggests, GOSP behaves like OP and respects the same equivalence with a percolation representation with paths with steps in U. It is then easy to see that $q_c < 1$ for GOSP (e.g. by comparison with a branching process) and consequently for the original subcritical family $\mathcal{U}$.

The converse inequality, $q_c > 0$ for subcritical models is significantly harder and constitutes the prime difficulty of Theorem 1.5.3 for this class. It was resolved in [28] via a fairly technical multiscale renormalisation scheme, which we will not describe, as we will not use any of it. It should be noted that here the importance of opposite stable directions is apparent (recall Remark 1.5.4 on critical and supercritical models). More precisely, what matters for subcritical models are strongly stable directions and, accordingly, it is important whether the model has two opposite strongly stable directions. If this is the case, the proof of $q_c > 0$ is simple. It suffices to partition the lattice into large rhombi, whose sides are close to being orthogonal to the two opposite strongly stable directions. If there is a bi-infinite oriented path of initially healthy rhombi, they do remain so at all times (see [28, 313]).

This argument reveals the relevance of strongly stable directions—they make slightly curved healthy regions stable. This will be of use to us also in Chapter 7. Indeed, the multiscale argument of [28] for models without opposite strongly stable directions proceeds by encapsulating infections in triangular contours with wiggling boundary, so as to be able to avoid other rare infections, preventing the use of straight lines.

Let us conclude the discussion by saying that for subcritical KCM essentially nothing is known beyond bootstrap percolation results. To give an exhaustive list of rigorous references on the subject, the reader may consult [355] for a combinatorial consideration in the spirit of [112], [246] superceded by [88] for ergodicity, positivity of the spectral gap and decay of correlations and [108] for mixing times. However, crucially, all these works only concern the simplest and very special NE model. Though it may be possible to generalise to GOSP, based on Chapter 11, treating subcritical KCM correspondingly is a very distant goal. In general, the only results are those valid for any KCM (subcritical or not):

- The ergodicity/mixing transition of the $\mathcal{U}$-KCM occurs at $q_{\mathrm{c}}(\mathcal{U})$ of the corresponding $\mathcal{U}$-bootstrap percolation, that is, the critical probability of $\mu(\tau_0^{\mathrm{BP}} = \infty)$. This is proved in [88].

- The non-zero spectral gap transition of the $\mathcal{U}$-KCM occurs at $\tilde{q}_{\mathrm{c}}(\mathcal{U})$ of the corresponding $\mathcal{U}$-bootstrap percolation, that is, the critical probability of exponential decay of $\mu(\tau_0^{\mathrm{BP}} > t)$ (see Chapter 12). We prove this in Theorem 12.3.7.

We therefore restrict our attention to (nontrivial subcritical) bootstrap percolation for the rest of Section 1.5.4.

From GOSP to general subcritical bootstrap percolation models

GOSP The above study of the nontriviality of the phase transition of subcritical models provides at least two reasons to study GOSP in detail. Firstly, they are the simplest of subcritical models and thus a good starting point. Secondly, the understanding of GOSP may be used to directly infer information on general families. Moreover, the models are interesting in their own right as percolation models, as well as probabilistic cellular automata (see Section 11.2). For these reasons, in Chapter 11, together with Réka Szabó [222], we study GOSP in arbitrary dimension $d \geqslant 2$, focusing on the phase $q < q_{\mathrm{c}}$, that is, the *supercritical* phase in percolation jargon.

The results from Chapter 11 relevant for the rest of our discussion of subcritical models can fortunately be taken as black boxes, allowing us to disregard most of the chapter if we are willing to accept them. Let us state them informally here. Firstly, assuming percolation language, we say that $a \in \mathbb{Z}^2$ is connected to $b \in \mathbb{Z}^2$ if there exists a healthy path from a to b with steps in the rule $U \subset \mathbb{H}_u$ defining our GOSP for some $u \in S^1$, which we take equal to $(0,1)$ for concreteness. Consider GOSP restricted to $\overline{\mathbb{H}}_v$ for some $v \in S^1$, that is, paths have to be contained in $\overline{\mathbb{H}}_v$. We are interested whether 0 is connected to infinity with positive probability as a function of q and v. It turns out that the set of directions v such that this happens is an interval varying with q in a continuous and strictly monotone fashion for $q \in [0, q_{\mathrm{c}})$. At $q \to q_{\mathrm{c}}-$ the interval converges to a half-circle and for $q \geqslant q_{\mathrm{c}}$ it is empty. Moreover, for v strictly outside of the (topological) closure of this interval, the length of the longest path from 0 (which is a.s. finite by definition) has an exponentially decaying tail.

Directional approach to subcritical models Before we can make better use of GOSP models than the bare comparison $q_{\mathrm{c}}(\mathcal{U}') \leqslant q_{\mathrm{c}}(\mathcal{U})$ when $\mathcal{U}' \supset \mathcal{U}$, we will need a directional decomposition of q_{c} (or rather $\tilde{q}_{\mathrm{c}}$, the critical probability of exponential decay of $\mu(\tau_0^{\mathrm{BP}} > t)$). For that purpose we introduce the following notion, whose precise definition is left to Chapter 12.

The *critical density* of $u \in S^1$ (for an update family $\mathcal{U}$) is morally

$$d_u \approx \inf\left\{q \in [0,1] : \mu(0 \notin [A \cup \mathbb{H}_u]_{\mathcal{U}}) = 0\right\},$$

where A still denotes the random set of infections with law μ. The vigilant reader has doubtlessly noticed that this is the exact same notion we just discussed for GOSP. Hence, we will consider GOSP critical densities as well-understood, though not explicit, functions on the circle. It is also clear that $d_u = 0$ for unstable directions and isolated stable ones, so this notion is tailored for strongly stable directions.

With critical densities at hand our central result of Chapter 12 reads

$$\tilde{q}_{\mathrm{c}} = \inf_{C \in \mathcal{C}'} \sup_{u \in C} d_u, \tag{1.17}$$

where $\mathcal{C}'$ denotes the set of closed semicircles. The proof is not hard and proceeds again by considering droplets growing in the middle direction of the semicircle above. In order to transform Eq. (1.17) into a refined version of the basic comparison $\tilde{q}_{\mathrm{c}}(\mathcal{U}') \leqslant \tilde{q}_{\mathrm{c}}(\mathcal{U})$ for $\mathcal{U}' \supset \mathcal{U}$, it suffices to observe that, likewise, $d_u(\mathcal{U}') \leqslant d_u(\mathcal{U})$ under the same condition. Hence,

$$\tilde{q}_{\mathrm{c}}(\mathcal{U}) \leqslant \inf_{C \in \mathcal{C}'} \sup_{u \in C} \min_{U \in \mathcal{U}} d_u(\{U\}),$$

allowing us to fruitfully transfer bounds on the critical densities of GOSP to arbitrary subcritical models (critical densities for non-supercritical one-rule families which are not GOSP are identically 1). In Chapter 12 we will illustrate that this indeed gives better bounds in generic situations.

We further recover known results about the Spiral family from [344], based on a less straightforward application of Eq. (1.17). Spiral is essentially the only subcritical bootstrap percolation model other than GOSP, which is relatively well understood, owing to its close relation to OP (see also [235, 236, 346, 347] for closely related models). It is is particularly interesting in view of the discontinuity of its phase transition[344, 346]: it satisfies

$$\mu_{q_{\mathrm{c}}}\left(\tau_0^{\mathrm{BP}} = \infty\right) > 0,$$

as expected for the jamming transition of granular systems [320].

We leave further results concerning exponential decay (in particular recovering results from [315]), noise sensitivity and answers to some questions of [28] to Chapter 12. To conclude our discussion of subcritical models, let us mention the following conjecture rectifying a question of [28], although the proper question whether it holds was already asked in [315].

Conjecture 1.5.5. *For all update families it holds that* $q_{\mathrm{c}} = \tilde{q}_{\mathrm{c}}$.

As we already discussed, this is true for all supercritical, critical, trivial subcritical models and GOSP, as well as Spiral, amounting to all models currently understood (see [208] for subsequent progress).

1.6 Refined universality for critical models in two dimensions

We may now turn to refined universality of critical bootstrap percolation models and KCM. Our ultimate goal are coarse thresholds—Aizenman–Lebowitz precision level for all of them.

1.6.1 Refined universality classes

To state the universality partition, we need another refinement of the notion of stable direction (see Fig. 1.2 for examples).

Definition 1.6.1 (Difficulty)**.** The *difficulty* $\alpha(u)$ of $u \in S^1$ is

- 0 if u is unstable;
- ∞ if u is stable, but not isolated;
- $\min\{n : \exists Z \subset \mathbb{Z}^2, |Z| = n, |[\mathbb{H}_u \cup Z]_{\mathcal{U}} \backslash \mathbb{H}_u| = \infty\}$ otherwise.

The *difficulty* of $\mathcal{U}$ is

$$\alpha = \alpha(\mathcal{U}) = \min_{C \in \mathcal{C}} \max_{u \in C} \alpha(u).$$

We say that a direction $u \in S^1$ is *hard* if $\alpha(u) > \alpha$.

The definition implies that there is always an open semi-circle $C \in \mathcal{C}$ containing no hard direction. It is interesting to note that the above definition may be viewed as an analogue of critical densities adapted to critical families, making α the analogue of $\tilde{q}_{\mathrm{c}}$ (or q_{c}, accepting Conjecture 1.5.5). Of course, in reality difficulties were introduced much earlier by Gravner and Griffeath (see e.g. [184]) and served as inspiration for critical densities of Chapter 12. Inversely, looking back at Definition 1.3.3, it is clear that difficulties should be interpreted as directional analogues of the parameter ζ appropriate for supercritical models.

Difficulties not only refine the definition of stable direction, but that of the rough universality classes. Namely, it is not hard to check that a model is supercritical iff its difficulty α is 0; critical iff α is a positive integer; subcritical iff $\alpha = \infty$. Consequently, for supercritical models directions are hard iff they are stable.

Further note that, while determining the stable directions or rough universality class of a family is an easy task, determining the difficulties of stable directions or the overall difficulty of the family is not. With Tamás Mezei [219] in Chapter 9, we examine this from a complexity viewpoint, showing that it is possible to determine α in finite time, given $\mathcal{U}$, but doing so is NP-hard. Providing a finite time algorithm relies on quantitative bounds on how far infection can travel from an initial set of α infections added to a half-plane, if it does not reach infinity, while the hardness result follows from embedding instances of the classical set cover problem. Regardless of the difficulty of determining difficulties, we will take them for granted in what follows.

With Definition 1.6.1 at our disposal, we can define all the notions which will appear in the refined universality partition. Although they may look arbitrary or zoological at first sight, we will see that they do correspond to different and sensible behaviours.

Definition 1.6.2 (Refined types)**.** A critical or supercritical two-dimensional update family is

- *rooted* if there exist two non-opposite hard directions;
- *unrooted* if it is not rooted;
- *unbalanced* if there exist two opposite hard directions;
- *balanced* if it is not unbalanced, that is, there exists a *closed* semicircle containing no hard direction.

We further partition balanced unrooted update families into

- *semi-directed* if there is exactly one hard direction;
- *isotropic* if there are no hard directions.

Of course, the reader may verify that rooted and unrooted supercritical models are (re)defined consistently with Definitions 1.3.9 and 1.5.2.

We will further consider the distinction between models with finite and infinite number of stable directions (that is, without or with a strongly stable direction). The latter are necessarily rooted but possibly balanced or unbalanced. Hence, we end up with a partition of all two-dimensional update families into the eleven classes represented in Fig. 1.2, seven of which are the critical ones we study below.

Let us remark that models in each class may have one axial symmetry, but non-subcritical models invariant under rotation by π are necessarily either isotropic or unbalanced unrooted (thus with finite number of stable directions), while invariance by rotation by $\pi/2$ implies isotropy.

Before stating the refined universality theorem, let us give a vague idea as to why Definition 1.6.2 may be relevant. Firstly, the distinction between balanced and unbalanced models was already apparent in the proof of Theorem 1.5.3. More precisely, we saw that we may be constrained to a one dimensional growth if the model is unbalanced, having to find rare sets along the side of the growing droplet, which tends to make it very elongated. The distinction between rooted and unrooted models was already apparent in one dimensional KCM and is related to the question whether or not we can move back and forth. This is also clearly relevant in view of the proof of Theorem 1.5.3, where the East-like motion we performed may be suboptimal for unrooted models. The presence or absence of strongly stable directions is also not absurd to take into account. That is because for those not even rare sets of infections can help us, making it 'impossible' for a droplet, however large or mobile, to move significantly in those directions. On the other hand, hard isolated stable directions are not 'impossible,' since rare sets of infections are sufficient for growing in those directions. Indeed, that was the procedure we used to prove rough bounds for critical models—we did not take difficulty into account, but only the fact that directions are not strongly stable. Sadly, we are unable to provide a simple and convincing explanation as to why semi-directed and isotropic models are very different but this will become clear with a bit more effort.

1.6.2 Refined universality results

Let us begin with the simpler case of bootstrap percolation, which already requires a substantial effort to prove.

Theorem 1.6.3 (Refined universality for bootstrap percolation [70]). *Let $\mathcal{U}$ be a two-dimensional critical update family with difficulty α. If $\mathcal{U}$ is*

- *unbalanced, then*
$$\mu\left(\tau_0^{\mathrm{BP}}\right) = \exp\left(\frac{\Theta(\log^2(1/q))}{q^\alpha}\right);$$
- *balanced, then*
$$\mu\left(\tau_0^{\mathrm{BP}}\right) = \exp\left(\frac{\Theta(1)}{q^\alpha}\right).$$

Once again, results in [70] were proved in terms of the probability of infecting a large torus, but can be recast into the above bounds on τ_0^{BP} w.h.p. and then improved to bounds in expectation, using rough decay rates as in Theorem 1.5.3. This theorem was relatively natural to conjecture, given that the relevance of difficulties and the balanced character of families was known to be important [184] and that several representatives of each class were studied and all shared these two scaling forms. Nevertheless, the proof and particularly the lower bound for unbalanced models was quite a challenge.

For KCM the picture was far less clear a priori. As testimony, we should point out that in the upper bounds provided in [269] were conjectured to be sharp up to logarithmic corrections. This was based on the intuition coming from a supercritical renormalisation point of view: if the model is unbalanced, then its droplets should be more costly as in bootstrap percolation; if the

	Infinite stable directions	Finite stable directions	
		Rooted	Unrooted
Unbalanced	(a) $2, 4, 0$	(c) $1, 3, 0$	(d) $1, 2, 0$
Balanced	(b) $2, 0, 0$	(e) $1, 1, 0$	(f) $1, 0, 1$ S.-dir. Iso. (g) $1, 0, 0$

Table 1.3 – Classification of critical $\mathcal{U}$-KCM with difficulty α from Theorem 1.6.4. For each class $\mathbb{E}_\mu(\tau_0) = \exp\left(\Theta(1)\left(\frac{1}{q^\alpha}\right)^\beta\left(\log\frac{1}{q}\right)^\gamma\left(\log\log\frac{1}{q}\right)^\delta\right)$ as $q \to 0$. The label of the class and the exponents β, γ, δ are indicated in that order. For comparison, $\mu\left(\tau_0^{\mathrm{BP}}\right) = \exp\left(\frac{\Theta(1)}{q^\alpha}\left(\log\frac{1}{q}\right)^{\gamma'}\right)$ with $\gamma' = 2$ for unbalanced $\mathcal{U}$ and $\gamma' = 0$ for balanced $\mathcal{U}$ by Theorem 1.6.3.

model is rooted, then it is obliged to move in an East-like way. This intuition was supported by the fact that their bounds did match automatic bootstrap percolation lower bounds in some cases and, more importantly, were proved to be sharp even with the logarithmic corrections for the specific case of the Duarte-KCM with $\mathcal{U} = \{\{(0,1),(-1,0)\},\{(0,-1),(-1,0)\},\{(0,1),(0,-1)\}\}$ in [267], establishing the first coarse threshold for a critical KCM. As it will become apparent in the following result, even the power of q given by the above heuristic was incorrect, making the distinction between finite and infinite number of stable directions unexpected. Including logarithmic corrections then uncovered yet another unpredicted and more intricate phenomenon leading to the following outcome summarised in Table 1.3.

Theorem 1.6.4 (Refined universality for KCM)**.** *Let $\mathcal{U}$ be a two-dimensional critical update family with difficulty α. If $\mathcal{U}$ is*

(a) unbalanced with infinite number of stable directions (so rooted), then

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta\left((\log(1/q))^4\right)}{q^{2\alpha}}\right);$$

(b) balanced with infinite number of stable directions (so rooted), then

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta(1)}{q^{2\alpha}}\right);$$

(c) unbalanced rooted with finite number of stable directions, then

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta\left((\log(1/q))^3\right)}{q^{\alpha}}\right);$$

(d) *unbalanced unrooted (so with finite number of stable directions), then*

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta\left((\log(1/q))^2\right)}{q^\alpha}\right);$$

(e) *balanced rooted with finite number of stable directions, then*

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta\left(\log(1/q)\right)}{q^\alpha}\right);$$

(f) *semi-directed (so balanced unrooted with finite number of stable directions), then*

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta\left(\log\log(1/q)\right)}{q^\alpha}\right);$$

(g) *isotropic (so balanced unrooted with finite number of stable directions), then*

$$\mathbb{E}_\mu(\tau_0) = \exp\left(\frac{\Theta(1)}{q^\alpha}\right).$$

We should emphasise that, at stark contrast with bootstrap percolation, this result is the state of the art for all critical update families except FA-2f, for which the only result of higher (or equal) precision is Theorem 1.4.6. The only KCM for which such a coarse threshold was known previously is the Duarte one [267, 269].

The proof of Theorem 1.6.4 could not be taken in one serving. Indeed, the following weaker statement was established as an intermediate step.

Corollary 1.6.5. *For two-dimensional critical update families*

$$\log \mathbb{E}_\mu(\tau_0) = \begin{cases} \left(\log \mu\left(\tau_0^{\mathrm{BP}}\right)\right)^{1+o(1)} & \textit{finite number of stable directions,} \\ \left(\log \mu\left(\tau_0^{\mathrm{BP}}\right)\right)^{2+o(1)} & \textit{infinite number of stable directions.} \end{cases}$$

The lower bound of Corollary 1.6.5 for models with finite number of stable directions follows directly from Theorem 1.6.3 and the automatic comparison $\mathbb{E}_\mu(\tau_0) \geqslant \Omega\left(\mu\left(\tau_0^{\mathrm{BP}}\right)\right)$ from Eq. (1.12). Actually, the same holds for the sharp lower bounds of Theorem 1.6.4 for classes (d) and (g). The remaining lower bound of Corollary 1.6.5 is proved in Chapter 7 jointly with Laure Marêché and Cristina Toninelli [214]. The proof incidentally giving the sharp scaling of Theorem 1.6.4 for class (b). The techniques there are developed further in Chapter 8 jointly with Laure Marêché [213] to prove all the lower bounds of Theorem 1.6.4 in a unified way.

Turning to upper bounds, the ones of Theorem 1.6.4 for class (a) and Corollary 1.6.5 for families with infinite number of stable directions were given in [269], the heuristic of which was already discussed in Section 1.5.3.

The remaining one in Corollary 1.6.5 along with Theorem 1.6.4(c) are proved in Chapter 4 jointly with Fabio Martinelli and Cristina Toninelli [216]. We prove the remaining upper bounds of Theorem 1.6.4 in Chapter 6, classes (e) and (f) being the most challenging.

It is worth mentioning that it is possible to optimise the proof of Theorem 1.6.4, so that some parts of Chapters 4 and 7 are not required (see Remark 6.0.2 and note that Section 7.4 is not used in Chapter 8). We have kept the original proofs in order to allow one to back out at the level of Corollary 1.6.5 and shun the fairly technical Chapters 6 and 8 if necessary.

1.6.3 Aspects of the proof for bootstrap percolation

Upper bounds

Let us begin with the upper bounds in Theorem 1.6.3, which are simpler and provide the intuition of the dominant infection mechanisms.

Balanced families For balanced models one proceeds very similarly to [7, 74], considering progressively growing infected droplets with polygonal shape and ensuring that, once they are large enough, they invade space roughly linearly. The latter assertion requires no new input w.r.t. Theorem 1.5.3, so we focus on the probability that a droplet of size $1/q^C$ for C large enough is internally filled. Our goal is to show that this probability is at least $\exp(-O(1)/q^{\alpha})$, where α is the difficulty of the balanced update family under consideration.

To obtain the desired bound, [70] proceeds as follows. Start from a large constant-sized infected droplet with a suitable polygonal shape as for Theorem 1.5.3 and demand the occurrence for each line of lattice sites we wish to append to the droplet a less unlikely event than the one in Section 1.5.3. Namely, we ask the presence of 'clusters' of at most α initial infections close to each other rather than some large constant. Fortunately, a few such clusters are enough to infect a line. Indeed, by Definition 1.6.1 for each direction u with $\alpha(u) \leqslant \alpha$ there exists a cluster of α infections, which generates infinitely many infections if placed next to an infected half-plane directed by u. As in one-dimension, it is not hard to see that the infections necessarily contain a periodic pattern along the line in at least one of the two directions. Hence, a bounded number of clusters, called a 'helping set' suffice to infect a half-line. This half-line in turn infects the entire line, as it contains arbitrarily large sequences of consecutive infections. Using this mechanism leads to the desired bound without additional effort w.r.t. Theorem 1.5.3.

Unbalanced families For unbalanced models we no longer need to pay any attention to the mechanism for infecting a 'critical' droplet. Namely, for such families a critical droplet may be taken to be a fully infected thick frame with polygonal shape. As we saw in Section 1.5.3, such a droplet occasionally finds sequences of infected sites allowing it to grow transversally. However, to reach them, it should be able to sustain its growth in the direction corresponding to the midpoint of a semicircle of difficulty at most α for a distance of order $1/q^C$. This can be achieved if the initial critical droplet has size $C' \log(1/q)/q^\alpha$ with C' sufficiently large compared to C. Indeed, for such sizes the probability that there is no helping set for a given line is roughly $1/q^{C'}$. Hence, by the time the first missing helping set is reached the droplet is already of size $1/q^C$ and helping sets are now exponentially likely to be present. Thus, the probability that a large droplet becomes infected by sites at distance at most $1/q^{C'}$ from it is essentially bounded by the probability of the initial frame being fully infected, leading to the desired $q^{C' \log(1/q)/q^\alpha}$.

Lower bounds

The lower bounds for Theorem 1.6.3 are the main difficulty, especially in the unbalanced case, for which we direct the reader to [70]. The simpler $\mu\left(\tau_0^{\mathrm{BP}}\right) \geqslant \exp(\Omega(1)/q^\alpha)$ bound, valid for all critical update families, can be proved essentially as in Section 1.5.3. More precisely, it suffices to 'cover' clusters of α sites close to each other rather than single infections. Since Definition 1.6.1 (with some work) guarantees that fewer than α sites cannot create arbitrarily large amounts of infection at the boundary of a droplet, covered droplets contain the closure of the initial infections. Hence, we may proceed identically to what was described in Section 1.5.3 to obtain the desired bound. We will see a more sophisticated version of this argument in Chapter 7.

1.6.4 KCM lower bounds: combinatorial bottleneck

Let us now discuss the proof of the lower bounds of Theorem 1.6.4. They are all treated in Chapters 7 and 8 in a systematic and comprehensive way, by showing that the main obstacle for the dynamics in all cases corresponds to a combinatorial bottleneck similar to the one of the two-dimensional East model (see Section 1.5.2 and Fig. 1.2j). For the sake of concreteness, we focus on the model from Fig. 1.2f, for which $\alpha = 1$.

Morally speaking, in this model the smallest mobile entity ('droplet') is an infected square of size roughly $1/q$. Indeed, typically on its right and top sides one can find an infection, which allows it to infect the column of sites on its right and the row of sites above it. However, it is essentially impossible for the infection to grow down or left, as this requires two consecutive infections and those are typically only available at distance $1/q^2$ from the droplet. We will only work in a region R of size $1/q^{7/4}$ around the origin, so, morally, such couples of infections are not available. Thus, the droplets follow the dynamics of the two-dimensional East model.

On a very high level we will proceed in the same way as for this supercritical rooted model in [265] (see Section 1.5.2). However, there are several obvious problems in making the above reasoning rigorous. Firstly, we said above that the smallest mobile entities, 'droplets,' were infected squares of size $1/q$, but the smallest mobile entities are actually more complicated. One needs to identify an event which says whether or not something is a droplet and this event should be deterministically necessary for infection to spread. Moreover, the event should be sufficiently unlikely, so that having many droplets at the same time has probability small enough to be a good bottleneck. It turns out that the notion of 'spanning' introduced in [70] for the proof of the lower bound of Theorem 1.6.3 for unbalanced models, following [97], is flexible enough for us. Roughly speaking (see Definition 8.1.1), a droplet is spanned if the infections present inside it are sufficient to infect a connected set touching all its sides. We call a droplet critical if it has size roughly $1/q$. It is known from [70] and obtained again in Section 8.A in a more adapted form that the probability of a specific critical droplet being spanned is roughly $\exp(-1/q)$. Unfortunately, given a configuration, spanned critical droplets may overlap and, in order to obtain good bounds on the probability of the configuration, one needs to consider disjointly occurring ones. We may then define the number of spanned critical droplets as the maximal number of disjointly occurring ones.

Having fixed these notions, we encounter a more significant issue: by changing their internal structure the (spanned critical) droplets may move a bit without creating another droplet. Worse yet, they are not really forbidden to move left or down, but simply are not likely to be able to do so wherever they want: it depends on the dynamic environment. Indeed, being able to move by a single step down is allowed by the presence of a couple of infections on the side of the droplet, which has probability only as small as q and is by far not something we can ensure never happens up to time $T \approx \exp(\log(1/q)/q)$.

In order to handle these problems, we introduce the crucial notion of crossing (not to be confused e.g. with the one of [70]). Consider a vertical strip S of width $1/q^{3/2}$ of our domain, R, which is a square of size $1/q^{7/4}$. Roughly speaking (see Definition 8.1.5 for a more precise statement), we say that S has a crossing if the following two events occur. Firstly, the infections in S together with the entire half-plane to the right of S are enough to infect a path from right to left in S (this is essentially the notion of crossing in [70]). Secondly, S does not contain a spanned critical droplet. Notice that these two events go opposite ways—the former is favoured by infections, while the latter is not. In Section 8.B we show that the probability of a crossing decays exponentially with the width of S at our scales.

Having established such a bound on the probability of crossings, we may safely assume (it happens with high probability) that they never occur until time T and this is the property we use to formalise the intuition that 'moving down or left is impossible.' More precisely, assume that initially the only critical droplet is on the right of S and S never has a crossing. Then, simply because the KCM dynamics can never infect more than what bootstrap percolation can, starting from the same initial condition, the droplet will not be able to reach the left side of S. Indeed, if it could, there would be a 'trail' of infectable sites from the right of S to its left, which would imply a crossing.

To conclude, let us emphasise a key aspect of the proof of upper bounds on spanning (and crossing, which are based on those for spanning), which we do not obtain like [70]. Indeed, in [70] one could not transfer analogous but easier bounds on covering, which were established there anyway, to spanning. That is because the covering algorithm in [70] lacks the key property of being closure-invariant: the collections of droplets associated to the closure of the initial infections being equal to the collection for the initial infections. Gaining an approximate version of this property is highly nontrivial, as in order not to overshoot in defining the droplets, one is forced to ignore small patches of infections (larger than the ones in [70]), which can possibly grow significantly when we take the closure for the bootstrap percolation process and especially so if they are close to a large infected droplet. In order to remedy this problem, we introduce a relatively intrinsic notion of 'crumb' (see Definition 7.3.1) such that its closure is still a crumb and does not differ too much from the original. A further advantage of our covering algorithm over the one of [70] is that it is somewhat canonical, with a well-defined unique output, which has particularly nice 'algebraic' description and properties (see Remark 7.3.10). Moreover, the closure property now allows us to directly transport probability bounds from covering to spanning.

1.6.5 KCM upper bounds: relaxation mechanisms

We finally attempt an heuristic explanation of Theorem 1.6.4 from the viewpoint of mechanisms, which is mostly related to upper bound proofs of Chapter 6. Instead of outlining the mechanism used by each class, we focus on techniques which are somewhat generic and then apply combinations thereof to each class. In figurative terms, we will develop several computer hardware components (three processors, four RAMs, etc.), give a general scheme of how to compose a generic computer out of generic components and, finally, assemble seven concrete computer configurations using the appropriate components for each, sometimes changing a single component from a machine to the other. Moreover, within each component type, different instances will be strictly comparable, so, at the assembly stage, we might simply choose the best possible component fitting with the requirements of model at hand. The purpose is twofold. Firstly, this enables us to highlight the robust tools we develop, which correspond to the components and how they are manufactured, as well as give a clean universal proof scheme which they are plugged into. Secondly, on the technical level, the modular structure will allow us to create each component only once and force us to make it as multipurpose as possible. Indeed, as it is clear from Table 1.4b, proceeding component-wise (row by row) is much easier than model-wise (column by column). We hope that the reader will be able to navigate through this more efficient albeit less straightforward procedure.

Our different components are called the *microscopic, internal, mesoscopic and global dynamics* and correspond to progressively increasing length scales on which we are able to relax, given a suitable infection configuration. As the notion of 'suitable,' which we call *super good* (SG), depends on the class and lower scale mechanisms used, we will mostly use it as a black box input extended progressively over scales in a recursive fashion. When we say that a convex polygonal region, called *droplet* and systematically equipped with a SG event (which makes the KCM inside the droplet ergodic), 'relaxes,' we mean that in a certain 'relaxation time' the dynamics restricted to the SG event and to this region 'mixes.' Formally, this translates to a constrained Poincaré inequality for the conditional measure, but this is unimportant for our discussion.

One should think of droplets as extremely unlikely objects, which are able to move within a (slightly) favourable environment. Indeed, at all stages of our treatment, we need to control the inverse probability of droplets (being SG) and their relaxation times, keeping them as small as feasible. Furthermore, due to their inductive definition, the favourable environment required for their movement should not be too costly. Indeed, that would result in the deterioration of the probability of larger scale droplets, as those incorporate the lower scale environment in their internal structure.

Scales

Microscopic dynamics is about modifying infections at the level of the lattice along the boundary of a droplet, while respecting the KCM constraint.

Internal dynamics is about relaxation on scales from the lattice level to the *internal scale* $\ell^{\mathrm{i}} = C^2 \log(1/q)/q^{\alpha}$, where C is a large constant depending on $\mathcal{U}$. This is the most delicate and novel step. Up to ℓ^{i} we account for the main contribution to the probability of droplets, which then saturates at a certain value ρ_{D}. Thus, it is important to only very occasionally ask for more than α infections to appear close to each other (helping sets). This means that up to the internal scale hard directions are practically impenetrable.

Mesoscopic dynamics is about relaxation on scales from ℓ^{i} to the *mesoscopic scale* $\ell^{\mathrm{m}} = 1/q^{C}$. As our droplets grow to the mesoscopic scale and past it, it becomes possible to require larger helping sets, which we call *W-helping sets*.

These allow droplets to move also in hard directions of finite difficulty, while nonisolated stable directions are still blocking.

Global dynamics is about relaxation on scales from ℓ^{m} to infinity. The extension to infinity being fairly standard (and not hard), one should rather focus on scales up to the *global scale* given by $\ell^{\mathrm{g}} = \exp(1/q^{3\alpha})$, which is notably much larger than all time scales we are aiming for.

Roughly speaking, on each of the last three scales, one should decide how to move a droplet of the lower scale in a domain on the larger scale.

For simplicity, in the remainder of Section 1.6.5, we assume that all update rules are contained in the axes of the lattice. This allows considering rectangular droplets (see Section 6.1.3). We further assume that all directions in the right-hand semicircle have difficulties at most α (under the above assumption only the four axis directions can be isolated or semi-isolated), while the down direction is hard, unless there are no hard directions (isotropic class), as in all critical examples in Fig. 1.2.

Microscopic dynamics

The microscopic dynamics is the only place where we actually deal with the KCM directly and is the same, regardless of the size of the droplet and the universality class. Roughly speaking, from the outside of the droplet, we may think of it as fully infected, since it is able to relax and, therefore, bring infections where they are needed. Thus, it is as though we are working on a single line of lattice sites, say column, next to an infected region. For an isolated (or semi-isolated) stable direction this induces a supercritical one-dimensional KCM on the column. Hence, provided a few suitable helping sets close to the column, we can apply results on one-dimensional inhomogeneous KCM from Chapter 2 to establish that the column can relax in time $\exp(O(\log(1/q))^2)$. Assuming we know how to relax on the droplet itself, this allows us to relax on a droplet with one column appended. However, applying this procedure recursively line by line is not efficient enough to be useful for extending droplets more significantly.

One-directional extensions

We next explain two fundamental techniques beyond the microscopic dynamics which we use to extend droplets on any scale in a single direction, say, right. Each of them can be viewed as a large scale version of either CBSEP or East in one dimension.

As mentioned above, our droplets have three aspects: geometry, SG event and relaxation mechanism bounding the relaxation time conditionally on the SG event. An extension takes as input a droplet with all its aspects and produces a larger (wider or taller), extended, one. While extending the geometry and SG event is a matter of definition and the relaxation mechanism is an heuristic image of the dynamics, bounds on the probability of the event and the relaxation time require a proof. This proof reflects the nature of the extension of the geometry and event, itself guided by the intuition of the underlying one-dimensional spin model and enabling the use of the proof technique for its relaxation time. We thus collectively refer to the procedure of extending the geometry, event and relaxation of a droplet as *extension*.

CBSEP-extension Recall CBSEP from Section 1.4.2. The relaxation time of this model on volume V is roughly at most $\min(V, 1/q)^2$ in one dimension and $\min(V, 1/q)$ in two and more dimensions, as we will see in Chapter 3.

For us particles will represent SG droplets, which we would like to move within a larger volume. However, as we would like them to be able to move possibly by an amount smaller than the size of the droplet, we need to generalise the model a bit. We equip each site of $\mathbb{Z}$ with a state space corresponding to the state of a column of the height of our droplet of interest in the original lattice $\mathbb{Z}^2$. Then the event 'there is a SG droplet' may occur on a group of ℓ sites (columns). The long range generalised CBSEP (which is actually a generalisation of what we call generalised CBSEP in Chapter 3), which we will call CBSEP by abuse in the present section, fixes some range $R > \ell$ and resamples groups of R consecutive sites if they contain a SG droplet, preserving this feature. Thus, one move of this process essentially delocalises the droplet within the range.

It is important to note (and this is crucial in Chapter 5) that CBSEP does not *have to* create a droplet in order to evolve. Indeed, conditionally on having a droplet within a certain domain, its position will be approximately uniform owing to the symmetric construction, so that, as long as it is able to move easily by one line both right and left, its position will quickly mix. It is for this initial step that we rely on the microscopic dynamics and helping sets. However, in order to achieve the displacement by one line we further need to be able to internally shuffle the SG event in an amoeba-like manner, so as to contract most of its internal structure in the direction we are moving to. Then, together with a suitable structure on the additional column granted by the microscopic dynamics, it becomes a droplet shifted by one step.

Below *CBSEP-extension* (Definition 6.2.4) refers to the procedure of extending a droplet's geometry, event and relaxation with CBSEP as underlying toy model. Geometry is simply extended both right and left, while the extended SG event requires the presence of the original SG droplet inside the extended one, in addition to helping sets throughout the rest of the extended droplet sufficient to catalyse the motion of the original droplet in both directions. The relaxation of the extended droplet via this mechanism is very swift. Indeed, the time needed to move the droplet is roughly a power of the volume times the inverse rate of the microscopic dynamics, which is itself fast, and the inverse rate of contraction, which is small, as we will discuss later. However, CBSEP-extensions can only be used for sufficiently symmetric update families. That is, the droplet needs to be able to move indifferently in both directions and its position should not be biased in one direction or the other.

East-extension Recall the East model from Section 1.3.2 and Fig. 1.1b with relaxation time $q^{-O(\log\min(L,1/q))}$ on a segment of length L. Its long range generalised version is defined similarly to the one of CBSEP. The only difference is that now $R > \ell$ consecutive columns are resampled together if there is a SG droplet on their extreme left. It is clear that this does not allow *moving* the droplet, but rather forces us to recreate a new droplet at a shifted position before we can progress. The associated *East-extension* (Definition 6.2.2) of a droplet corresponds to extending its geometry to the right, while the extended SG event requires that the original SG droplet is present in the leftmost position and helping sets are available in the rest of the extended droplet to allow its (long range generalised) East evolution.

The generalised East process goes back to [269], while the long range version is implicitly used in Chapter 4. However, both use a brutal strategy consisting of creating the new droplet from scratch. Instead, in Chapter 6 we will have to be much more careful, particularly in view of semi-directed models. Indeed, take ℓ large and $R = \ell + 5$. Then it is intuitively clear that the presence of the original leftmost droplet overlaps greatly with the occurrence of the shifted SG one we would like to craft. Hence, the idea is to take advantage of this and only pay the *conditional* probability of the droplet we are creating, given the presence of the original one.

This is not as easy as it sounds for several reasons. Firstly, we should make the SG structure soft enough, as in Chapter 5 and in contrast with [269] and Chapter 4, so that small shifts do not change it much. Secondly, we need to actually have a quantitative estimate of the conditional probability of a complicated multi-scale event, given its translated version, which necessarily does not quite respect the same multi-scale geometry. To make matters worse, we do not have at our disposal a very sharp estimate of the probability of SG events (contrary to what is the case in Chapter 5), so directly computing the ratio of two rough estimates would yield a very poor bound on the conditional probability. In fact, this problem is also present when contracting an amoeba in the CBSEP-extension—we need to evaluate the probability of a contracted version of the amoeba conditionally on the original amoeba being present.

We deal with these issues in Section 6.B. We establish that, as intuition may suggest, to create a droplet shifted by $R - \ell$, given the original one, we roughly only need to pay the probability of a droplet on scale $R-\ell$ rather than ℓ, which provides a substantial gain. Hence, the time necessary for an East-extension of a droplet to relax is essentially the product of the inverse probabilities of a droplet on scales of the form 2^m up to the extension length.

Internal dynamics

The internal dynamics is where most of our work will go. This is not surprising, as the probability of SG events will saturate at its final value ρ_{D} ($\exp(-\Theta(1)/q^\alpha)$ for balanced models and $\exp(\Theta(-\log^2(1/q))/q^\alpha)$ for unbalanced ones) at the internal scale. These are the values familiar from bootstrap percolation.

Unbalanced internal dynamics Let us begin with the simplest case of unbalanced models. If $\mathcal{U}$ has unbalanced with infinite number of stable directions (class (a)), droplets in [269] on the internal scale consist of several consecutive infected columns, so that no relaxation is needed (the SG event is a singleton). The columns have size ℓ^{i}, which justifies the value of ρ_{D}.

If $\mathcal{U}$ is unbalanced with finite number of stable directions (classes (c) and (d)), droplets on the internal scale are fully infected square frames of thickness $O(1)$ and size ℓ^{i}, which gives a similar value of ρ_{D}. In order to relax inside the frame, one can infect several columns next to the frame (inside it) and move them throughout the area enclosed in the frame with the help of the frame. This can be done similarly to a CBSEP-extension, by infecting the next column and removing the previous one (see Fig. 4.7). The time necessary for this relaxation is easily seen to be $\rho_{\mathrm{D}}^{O(1)}$ (the cost for creating the infected columns).

CBSEP internal dynamics If $\mathcal{U}$ is isotropic (class (g)), up to the conditioning problems of Section 6.B described above, we need only minor adaptations of the strategy of Chapter 5 for FA-2f. Droplets on the internal scale will have an internal structure as obtained by iterating Fig. 6.5a (see also Fig. 5.2). Our droplets will be extended little by little alternating between the horizontal and vertical directions, so that their size is multiplied essentially by a constant at each extension. Thus, roughly $\log(1/q)$ extensions are required to reach ℓ^{i}. As isotropic models do not have any hard directions, we can move in all directions and thus the symmetry required for CBSEP-extensions is granted. Hence, this mechanism leads to a very fast relaxation of droplets in time $\exp(q^{-o(1)})$.

Remark 1.6.6. The vigilant reader may have noticed that CBSEP requires an actual symmetry, while for a general isotropic model we only know that there are no hard directions. We circumvent this issue by artificially symmetrising our droplets and events, asking for helping sets in directions which do not need any and asking for the symmetric of the helping set of the opposite direction. Although these are totally useless for the dynamics, they are important to ensure that the positions of droplets are indeed uniform rather than suffering from a drift towards an 'easier' non-hard direction.

East internal dynamics The most challenging case is the balanced non-isotropic one (classes (b), (e) and (f)). Indeed, the hard direction prevents us from using CBSEP-extensions. To be precise, for semi-directed models (class (f)) it is possible to perform CBSEP-extensions horizontally, but the gain is insignificant, so we treat all balanced non-isotropic models identically up to the internal scale.

We still extend droplets, starting from a microscopic one, by a constant factor alternating between the horizontal and vertical directions. However, in contrast with the isotropic case, extensions are done in an oriented fashion, so that the original microscopic droplet remains anchored at the corner of larger ones (see Fig. 6.3b). Thus, we may apply East-extensions on each step and obtain that the cost is given by the product of conditional probabilities for East extensions over all scales and shifts of the form 2^n. In total, a droplet of size 2^n needs to be paid once per scale larger than 2^n. A careful computation shows that only droplets larger than $q^{-\alpha}$ provide the dominant contribution and those all have probability essentially ρ_{D}. Thus, the total cost would be $\rho_{\mathrm{D}}^{\Theta(\log\log(1/q))^2}$, since there are $\log\log(1/q)$ scales from $q^{-\alpha}$ to ℓ^{i}, as they increase exponentially. This is unfortunately a bit too rough for the semi-directed class. However, the solution is simple—it suffices to introduce scales growing double-exponentially above $q^{-\alpha}$, so that the product becomes dominated by its last term, giving the final cost $\rho_{\mathrm{D}}^{\Theta(\log\log(1/q))}$.

Mesoscopic dynamics

For the mesoscopic dynamics we are given as input a SG event for droplets on scale ℓ^{i} and a bound on their relaxation time.

CBSEP mesoscopic dynamics If $\mathcal{U}$ is unrooted (classes (d), (f) and (g)), recall that the hard directions (if any) are vertical. Then we can perform a horizontal CBSEP-extension directly from ℓ^{i} to ℓ^{m}, since $\ell^{\mathrm{i}} \approx \log(1/q)/q^{\alpha}$ makes it likely for helping sets to appear along all segments of length ℓ^{i} until we reach scale $\ell^{\mathrm{m}} \approx q^{-C}$. The resulting droplet is very wide, but short (see Fig. 6.6). However, this is enough for us to be able to perform a vertical CBSEP-extension, requiring W-helping sets, since they are now likely to be found. Again, CBSEP dynamics being very efficient, its cost is negligible.

East mesoscopic dynamics If $\mathcal{U}$ is rooted (classes (a)-(c) and (e)), then CBSEP-extensions are still inaccessible. We may instead East-extend rightwards from ℓ^{i} to ℓ^{m} in a single step. Once again we obtain a very wide but short droplet. If the model has finite number of stable directions or it is balanced (classes (b), (c) and (e)) we may perform an East-extension upwards, since W-helping sets are now likely to be found as in the CBSEP mesoscopic dynamics. This leads to a droplet of size ℓ^{m} in time $\rho_{\mathrm{D}}^{\Theta(\log(1/q))}$. If the family has finite number of stable directions (classes (c) and (e)), we may afterwards CBSEP-extend (vertically or horizontally) as we did for unrooted models above. Note that for balanced models with infinite number of stable directions we can only move rightwards and upwards, while for unbalanced families with infinite number of stable directions we can only move rightwards and the following additional mechanism is needed.

Stair mesoscopic dynamics For unbalanced families with infinite number of stable directions (class (a)) the following *stair mesoscopic dynamics* was introduced in [269]. Recall that in this case the internal droplet is simply a few infected columns. While moving the droplet right via an East motion, we pick up W-helping sets above or below the droplet (as in Section 1.5.3). These sets allow us to make all droplets to their right shifted up or down by one row. Hence, we manage to create a copy of the droplet far to its right but also slightly shifted up or down (see [269, Fig. 6]). Repeating this (with many steps in our staircase) in a two-dimensional East-like motion, we can now relax on a mesoscopic droplet with horizontal dimension much larger than ℓ^{m} but still polynomial in $1/q$ and vertical dimension ℓ^{m} in time $\rho_{\mathrm{D}}^{\Theta(\log(1/q))}$.

Global dynamics

The global dynamics receives as input a SG event for a droplet on scale ℓ^{m} with probability roughly ρ_{D} and a bound on its relaxation time, as provided by the mesoscopic dynamics. Its goal is to move such a droplet efficiently to the origin from its typical initial position at distance roughly $\rho_{\mathrm{D}}^{-1/2}$.

CBSEP global dynamics If $\mathcal{U}$ has a finite number of stable directions (classes (c)-(g)), since the mesoscopic droplet is large enough, it can perform a CBSEP motion in a typical environment. Therefore, the cost of this mechanism is given by the relaxation time of CBSEP on large volumes with density of droplets given by ρ_{D}. Performing this strategy carefully and using the two-dimensional CBSEP, this yields roughly $1/\rho_{\mathrm{D}}$. Preferring a two-dimensional over a one-dimensional CBSEP strategy is not of particular importance for Theorem 1.6.4, since we only know $\log \rho_{\mathrm{D}}$ up to a constant factor. However, this is crucial in Chapter 5 for Theorem 1.4.6.

East global dynamics If $\mathcal{U}$ has infinite number of stable directions (classes (a) and (b)), the strategy is identical to the CBSEP global dynamics, but

Global		Mesoscopic		Internal		
CBSEP	East	CBSEP	East, Stair	CBSEP	East	Unbal.
$\rho_{\mathrm{D}}^{-1+o(1)}$	$\rho_{\mathrm{D}}^{-O(\log(1/\rho_{\mathrm{D}}))}$	$e^{q^{-o(1)}}$	$\rho_{\mathrm{D}}^{-O(\log(1/q))}$	$e^{q^{-o(1)}}$	$\rho_{\mathrm{D}}^{-O(\log\log(1/q))}$	$\rho_{\mathrm{D}}^{-O(1)}$

(a) The relaxation time cost associated to each choice of dynamics mechanism on each scale in terms of the probability of a droplet ρ_{D}.

	(a)	(b)	(c)	(d)	(e)	(f)	(g)
Global	East*	East*	CBSEP	CBSEP*	CBSEP	CBSEP	CBSEP*
Mesoscopic	Stair	East	East*	CBSEP	East*	CBSEP	CBSEP
Internal	—	East	Unbal.	Unbal.*	East	East*	CBSEP

(b) The fastest mechanism available to each refined universality class of Theorem 1.6.4 on each scale. The * indicates a leading contribution for the class.

Table 1.4 – Summary of the mechanisms and their costs. The microscopic one common to all classes and with negligible cost is not shown.

employs an East dynamics. Now the cost becomes the relaxation time of an East model with density of infections ρ_{D}, which yields $\exp(O(\log(1/\rho_{\mathrm{D}}))^2)$.

Assembling the components

To conclude, in Table 1.4a we provide a summary of the mechanisms for each scale and their cost to the relaxation time. The results are expressed in terms of the probability of a droplet ρ_{D}, which equals $\exp(-O(\log(1/q))^2/q^\alpha)$ for unbalanced models and $\exp(-O(1)/q^\alpha)$ for balanced ones. The final bound on $\mathbb{E}_\mu(\tau_0)$ for each class then corresponds to the product of the costs of the mechanism employed at each scale. To complement this, in Table 1.4b we indicate the fastest mechanism available for each class on each scale and further indicate which one gives the dominant contribution to the final result appearing in Theorem 1.6.4, once the bill is footed.

Finally, let us avert the reader that, for the sake of concision, the proofs in Chapter 6 do not systematically implement the optimal strategy for each class as indicated in Table 1.4b if that does not deteriorate the final result. Similarly, when that is unimportant, we may give weaker bounds than the ones in Table 1.4a.

1.7 Organisation

The remainder of the thesis is structured as follows, illustrated in Fig. 1.3. Each chapter is based on a different paper or preprint among [207, 209, 210, 213–216, 218–220, 222]. The reader will be pleased to learn that they are already completely familiar with the introductory material of [207, 210, 213, 215, 219, 220] and partially for the remaining works [209, 214, 216, 218, 222]. Each chapter is self-contained enough to be read independently of all others

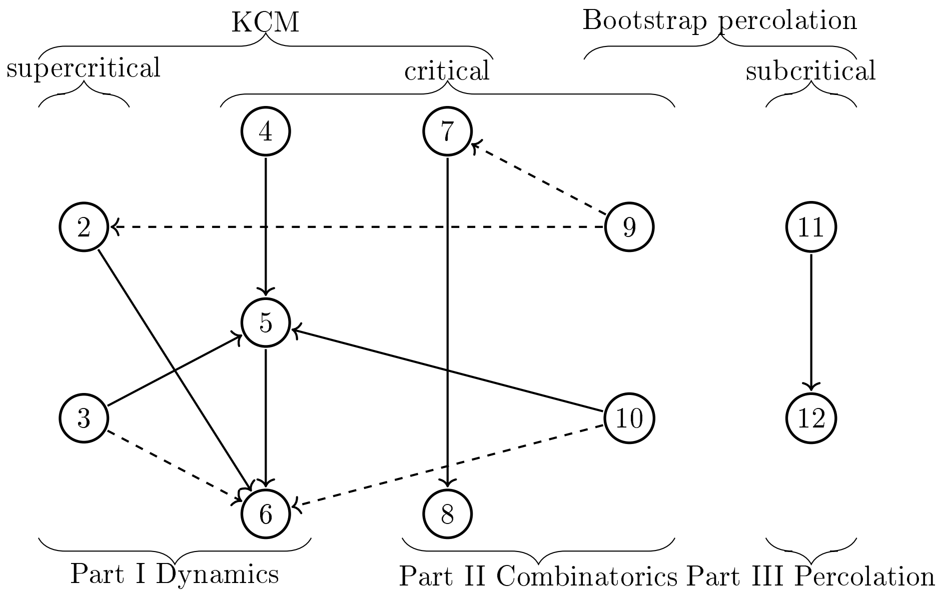

Figure 1.3 – Organisation of the thesis. Vertices represent chapters. Solid arrows indicate that some results are imported, while dashed ones indicate less explicit influence or inspiration.

(except the present Chapter 1, familiarity with which is assumed throughout). We have not altered the published or pre-published originals for the sake of verifiability, with the exception of removing parts already covered in Chapter 1, acknowledgements and open problems solved by subsequent chapters. In particular, we have left the notation of each chapter as in the corresponding original, so any notation defined there primes over Chapter 1, but in most cases the notation is precisely the one we have seen already.

Though they are self-contained, many dependencies are present between chapters, as results from other chapters are occasionally used (see Fig. 1.3). Chapters are presented from supercritical through critical to subcritical models, covering KCM before addressing bootstrap percolation. Naturally, the reader may choose a different order and even subset. Indeed, the graph in Fig. 1.3 is divided into relatively independent parts suitable for readers with different backgrounds or tastes.

Part I (Chapters 2 to 6) contains upper bounds for KCM. It is appropriate for readers proficient in Markov chains and, more specifically, dynamics of interacting particle systems.

Part II (Chapters 7 to 10) presents lower bounds for critical KCM (Chapters 7 and 8) followed by critical bootstrap percolation results. Both are appropriate for readers adept at (probabilistic) combinatorics.

Part III (Chapters 11 and 12) treats subcritical bootstrap percolation models, making it appropriate for readers versed in percolation theory.

We regard Chapters 6, 8 and 12 as the core of each part in decreasing level of importance and difficulty. We view the rest as prerequisites, though all of them are of independent interest and some specialists would doubtlessly find Chapters 5 and 10 of greater importance, Chapter 11 of greater interest and Chapters 2 and 3 of greater use.

We should acknowledge that Chapters 10 and 12 (i.e. [209, 220]) were essentially completed during the author's Master degree, the latter being the author's Master thesis. They are included, since they are tightly linked to the remainder of the study. Instead, we do not include works unrelated to our main subject, completed both prior to the beginning of the thesis [115, 206, 211] and during its course [212].

The content of each chapter is as follows.

Chapter 2 [207] treats arbitrary general inhomogeneous KCM on finite or infinite one-dimensional volumes conditioned to an irreducible component, as discussed in Section 1.3.2. We establish an East-like bound on their relaxation times without imposing any condition via a revised bisection approach.

Chapter 3 [218] joint with Fabio Martinelli and Cristina Toninelli studies CBSEP discussed in Section 1.4.2 on arbitrary finite graphs and a natural generalisation thereof. We provide bounds on relaxation, mixing and logarithmic Sobolev times in full generality, often sharp up to logarithmic corrections. We are particularly interested in the limit of the number of vertices of the graph diverging as the inverse of the equilibrium particle density p and implications for FA-1f in particular recovering results of Pillai–Smith [299, 300].

Chapter 4 [216] joint with Fabio Martinelli and Cristina Toninelli proves the upper bound on $\mathbb{E}_\mu(\tau_0)$ in Corollary 1.6.5 for critical models with finite number of stable directions.

Chapter 5 [215] joint with Fabio Martinelli and Cristina Toninelli proves Theorem 1.4.6 on FA-2f on $\mathbb{Z}^2$, thus establishing the first sharp threshold for a critical KCM and settling conflicting conjectures from physics, as explained in Section 1.4.4. The proof crucially relies on input from Chapters 3 and 10 for the upper and lower bounds respectively and includes an adapted version of Section 4.4.

Chapter 6 [210] proves the upper bounds on $\mathbb{E}_\mu(\tau_0)$ in Theorem 1.6.4 for all critical two-dimensional refined universality classes except unbalanced families with infinitely many stable directions (a) handled by [269]. The mechanisms involved were duly outlined in Section 1.6.5. We rely on Chapters 2 and 5.

Chapter 7 [214] joint with Laure Marêché and Cristina Toninelli proves the lower bound on $\mathbb{E}_\mu(\tau_0)$ in Corollary 1.6.5 for critical models with infinite number of stable directions.

Chapter 8 [213] joint with Laure Marêché proves the lower bounds on $\mathbb{E}_\mu(\tau_0)$ in Theorem 1.6.4 for all critical two-dimensional refined universality classes, relying on Chapter 7. Thus, together with Chapter 6 we complete the (refined) universality of critical two-dimensional KCM. The ideas were detailed in Section 1.6.4.

Chapter 9 [219] joint with Tamás Mezei examines the key parameter of two-dimensional universality of both bootstrap percolation and KCM—the difficulty α—from a computational perspective. As mentioned in Section 1.6.1, we show that determining it, given $\mathcal{U}$, is NP-hard and exhibit an algorithm for finding it in finite time.

Chapter 10 [220] joint with Robert Morris establishes the lower bound in Theorem 1.4.5, determining the order of magnitude of the second term of $\mu\left(\tau_0^{\mathrm{BP}}\right)$ for 2-neighbour bootstrap percolation on $\mathbb{Z}^2$ discussed in Section 1.4.3. The proof requires a very precise understanding of the typical growth of a critical droplet, and involves a number of technical innovations.

Chapter 11 [222] joint with Réka Szabó shows that GOSP behaves like ordinary oriented percolation in its supercritical phase, as noted in Section 1.5.4.

Chapter 12 [209] provides general results on subcritical families, as mentioned in Section 1.5.4. In particular, we answer several questions posed in [28]. The output of Chapter 11 naturally plugs into Chapter 12 to extend some of its results to the desired complete generality.

Part I

Dynamics

Chapter 2

Bisection for KCM revisited

This chapter is based on [207]. Recall Section 1.3.2.

2.1 Formal statement

For all *sites* $x \in \mathbb{Z}$ fix a finite positive probability space $(\mathcal{S}_x, \pi_x)$ called *state space* and $\mathcal{I}_x \subset \mathcal{S}_x$ satisfying $\pi_x(\mathcal{I}_x) \geqslant q > 0$. We say that $x \in \mathbb{Z}$ is *infected* when the event $\mathcal{I}_x$ occurs and *healthy* otherwise. Thus, we refer to $q = \inf_{x\in\mathbb{Z}} \pi_x(\mathcal{I}_x)$ as the *infection probability*. The *volume* $L \subset \mathbb{Z}$ is a finite or infinite set. Consider the corresponding product space $\mathcal{S}_L = \prod_{x\in L} \mathcal{S}_x$ and measure $\pi_L = \bigotimes_{x\in L} \pi_x$. We will usually denote elements of $\mathcal{S}_L$ (*configurations*) by η, ω, ξ, etc. and corresponding restrictions to any $X \subset L$ by η_X and η_x when $X = \{x\}$. A *boundary condition* is any $\omega \in \mathcal{S}_{\mathbb{Z}\setminus L}$ or an appropriate restriction, when some of the states of ω are unimportant. Given two configurations $\eta_L \in \mathcal{S}_L$ and $\eta'_{L'} \in \mathcal{S}_{L'}$ for volumes L, L' with $L \cap L' = \varnothing$, we denote by $\eta_L \cdot \eta'_{L'} \in \mathcal{S}_{L\cup L'}$ the configuration equal to η_x if $x \in L$ and to η'_x if $x \in L'$.

For all $x \in L$ we fix *an update family* $\mathcal{U}_x$ that is a finite family of finite subsets of $\mathbb{Z}\setminus\{x\}$. Its elements are called *update rules*. We assume that there exists a *range* $R \in [1, \infty)$ such that for all $x \in L$, $U \in \mathcal{U}_x$ and $y \in U$ we have $|x - y| \leqslant R$. For $x \in L$ we say that the *constraint at x is satisfied* if

$$c_x^\omega(\eta) = \mathbb{1}_{\exists U\in\mathcal{U}_x, \forall y\in U, (\eta\cdot\omega)_y\in\mathcal{I}_y}$$

equals 1. In words, we require that for at least one of the rules all its sites are infected, taking into account the boundary condition. The transitions allowed for the KCM are those changing the state of a single site whose constraint is satisfied (before and, equivalently, after the transition, since rules for x do not contain x). In these terms, $\mathcal{U}_x = \varnothing$ corresponds to a site unable to update under any circumstances, while $\mathcal{U}_x \ni \varnothing$ corresponds to a site whose constraint is always satisfied. The transitions define an oriented graph with vertex set $\mathcal{S}_L$ and symmetric edge set (containing the reverses of its edges). We call its connected components *irreducible components* of the KCM and view them as events. Given an irreducible component $\mathcal{C} \subset \mathcal{S}_L$, we set $\mu_L = \pi_L(\cdot|\mathcal{C})$. We further write $\mu_X = \mu_L(\cdot|\eta_{L\setminus X})$, $\mu_x = \mu_{\{x\}}$ for $x \in \mathbb{Z}$ and $X \subset \mathbb{Z}$ and denote by Var_X and Var_x the corresponding variances.

The *general KCM* defined by L, $\mathcal{S}_x$, π_x, $\mathcal{I}_x$, ω, $\mathcal{U}_x$ and $\mathcal{C}$ is the continuous time Markov process with generator and Dirichlet form acting on functions $f : \mathcal{C} \to \mathbb{R}$ depending on the states of finitely many sites given by

$$\begin{aligned}\mathcal{L}_L(f)(\eta) &= \sum_{x\in L} c_x^\omega(\eta) \cdot (\mu_x(f(\eta)) - f(\eta)),\\ \mathcal{D}_L(f) &= \sum_{x\in L} \mu_L\left(c_x^\omega \cdot \mathrm{Var}_x(f)\right)\end{aligned}$$

respectively. In other words, this is the continuous time Markov process which resamples the state of each site at rate 1 w.r.t. μ_x, provided its constraint is satisfied. It is useful to note that when $c_x^\omega = 1$, we have $\mu_x = \pi_x$. For the existence of such infinite-volume processes see [258] and for basic background refer to [88, 89]. It is also not hard to check that π_L and, therefore, μ_L is a reversible invariant measure for the process. Finally,

$$(T_{\mathrm{rel}})^{-1} = \inf_{f\neq\mathrm{const.}} \frac{\mathcal{D}_L(f)}{\mathrm{Var}_L(f)} \in [0,1]$$

is the *spectral gap* of $\mathcal{L}_L$ or inverse *relaxation time.*

Theorem 2.1.1. *There exists an absolute constant $C > 0$ such that for any range $R \in [1,\infty)$, infection probability $q \in (0,1]$, volume L and general KCM with these parameters it holds that*

$$T_{\mathrm{rel}} \leqslant (2/q)^{CR^2 \min(\log|L|, R\log(2/q))}. \tag{2.1}$$

Remark 2.1.2. Equation (2.1) and its proof apply to general KCM on a circle $\mathbb{Z}/n\mathbb{Z}$ (uniformly on n). For trees of maximum degree Δ and diameter D we can only retrieve that for some C depending on Δ and R,

$$T_{\mathrm{rel}} \leqslant (2/q)^{C\log D}.$$

As mentioned in Section 1.3.2, Theorem 2.1.1 is sharp not only for all homogeneous rooted supercritical models, but also for some unrooted ones. Indeed, an unrooted KCM in finite volume may lack clusters of infections mobile in both directions, but only be able to create them, using ones mobile in a single direction. Such is the case of the homogeneous $\{\{-2\},\{1,2\}\}$-KCM on $L = \{1,\dots,2n\}$ with healthy boundary condition, only 1 and $|L|$ infected initially (so that it is in its "ergodic component," able to infect the entire volume). As usual, a test function showing that $T_{\mathrm{rel}} \geqslant \exp(\Omega(\log^2(1/q)))$ for $|L| = 1/q \to \infty$ is the indicator of configurations reachable from the initial state above without creating $\log(1/q)/10$ infections simultaneously. The phenomenon is not related to the lack of symmetry—a similar reasoning applies to the $\{\{-9,-8,-6\},\{-7,-6,-4\},\{-6,-5,-3\},\{3,5,6\},\{4,6,7\},\{6,8,9\}\}$-KCM on $L=\{1,\dots,6n+3\}$ with the ergodic initial condition $\{1,2,4,6n-1,6n,6n+3\}$.

2.2 Proof

Let us begin with a straightforward but important corollary of reversibility.

Observation 2.2.1. The irreducible component of a general KCM naturally identifies with the set of sites which can be eventually updated, together with the state of all remaining sites. We call the set of the sites that can be updated in L *closure*[1] and denote it by $\{\eta\}_L^\omega \subset L$. We denote the state of the remaining sites by $\eta^0 := \eta_{L\setminus\{\eta\}_L^\omega}$ and refer to it as *initial condition.*

Since sites in $L\setminus\{\eta\}_L^\omega$ can never be updated, we may remove them from L and replace ω by $\omega\cdot\eta_{L\setminus\{\eta\}_L^\omega}$. With this reduction, we may assume that $\{\eta\}_L^\omega = L$ for the original general KCM and omit this condition. Further note that we may absorb any boundary condition in the inhomogeneous update rules by removing infected sites in ω from update rules and removing update rules containing non-infected sites in ω. Thus, we may further assume that our initial general KCM is defined so that its rules do not depend on the boundary condition and therefore discard ω. Moreover, once the boundary condition is irrelevant, we may replace L by an interval of length $|L|$, if L is finite, or $\mathbb{N}$ or $\mathbb{Z}$ if L is infinite in one or two directions. Finally, approximating L by large finite segments if it is infinite (see [88, Section 2] and [258, Chapter 4]), we may assume $|L|<\infty$. Note that during the proof we will consider smaller domains and will then specify the closure and boundary condition.

Henceforth, we fix a general KCM subject to the above simplifications specified by its volume $L=\{1,\dots,|L|\}$, state spaces $(\mathcal{S}_x,\pi_x)$, infection events $\mathcal{I}_x$, and update families $\mathcal{U}_x$. We will prove Theorem 2.1.1 by induction on $|L|$. The induction step is provided by the following two-block result, which is the core of the argument.

Proposition 2.2.2. *Let* $L_1=\{1,\dots,\ell\}$ *and* $L_2=\{\ell-\Delta+1,\dots,|L|\}$ *with* $\ell\in[1,|L|]$ *and* $\Delta\in[0,\ell]$. *Then*

$$\operatorname{Var}_L(f) \leqslant \gamma(\Delta)\sum_{i\in\{1,2\}}\mu_L\left(\operatorname{Var}_{L_i}\left(f|\{\eta\}_{L_i}^{\eta_{L\setminus L_i}},\eta^0_{L_i}\right)\right), \tag{2.2}$$

[1] Note that this is smaller than the bootstrap percolation closure denoted $[\cdot]$ in Chapter 1, which will not be used in the present chapter to avoid confusion.

setting for some absolute constant $C > 0$

$$\gamma(\Delta) = \begin{cases} 1 + \exp\left(-\Delta q^{CR} / \left(CR^2\right)\right) & \Delta \geqslant CR^2/q^{CR} \\ (2/q)^{CR^2} & \textit{otherwise.} \end{cases}$$

Remark 2.2.3. This statement can be viewed as a Poincaré inequality for a Markov process with two symmetric moves performed at rate 1. We update the state η_{L_1} from the measure π_{L_1} conditioned on the irreducible component of the current state in L_1. Crucially, the closure is taken only inside L_1, without infecting sites in L_2 and going back to L_1, but using η_{L_2} as a (frozen) boundary condition. In particular, the variance in Eq. (2.2) is not $\mathrm{Var}_{L_i}(f)$.

Before proving Proposition 2.2.2, let briefly recall how to deduce Theorem 2.1.1, referring to [88, Theorem 6.1] for more details. We average Eq. (2.2) over $N \approx |L|^{1/3}$ choices of ℓ so that the $L_1 \cap L_2$ for different choices are disjoint and $\Delta \approx |L|^{1/3}$ is fixed. All ℓ are chosen so that $\ell - |L|/2 \in [-N\Delta/2, N\Delta/2]$. Denoting by Γ_l the maximum of T_{rel} over all general KCM of range (at most) R and infection probability (at least) q on volume with cardinal at most l, this yields the recurrence relation

$$\Gamma_{|L|} \leqslant (1 + 1/N)\gamma(\Delta)\Gamma_{|L|/2+N\Delta}.$$

Iterating this, we derive the desired Eq. (2.1).

Thus, our task is to prove Proposition 2.2.2, for which we need the following.

Claim 2.2.4. Let Λ be a volume. Then for any irreducible component $\mathcal{C} = (\{\eta\}_\Lambda^\omega, \eta^0)$, under $\pi_\Lambda(\cdot|\mathcal{C})$ the infections in the closure $(\mathbb{1}_{\mathcal{I}_x})_{x \in \{\eta\}_\Lambda^\omega}$ stochastically dominate i.i.d. Bernoulli variables with parameter q.

Proof. Indeed, for any x in the closure, conditionally on $\eta_{\Lambda\setminus\{x\}}$ and $\mathcal{C}$, either $\mathcal{I}_x$ has to occur to obtain the correct closure or η_x has the law π_x. □

Sketch of the easier case of Proposition 2.2.2. As a warm-up, let us sketch the proof of Eq. (2.2) with $\gamma(\Delta) = (2/q)^{CR^2}$, which is valid for all values of Δ.

We aim to couple two copies η and η' of the chain in Remark 2.2.3, so that they meet with appropriate rate. To do this, we require that the following sequence of events all occur in both chains uninterrupted by any other updates. Each chain is updated on L_1 to a state such that the R sites in $L_1 \setminus L_2$ closest to L_2 (if $|L_1 \setminus L_2| < R$, take all sites in $L_1 \setminus L_2$) which are in the closure $\{\eta\}_{L_1}^{\eta_{L \setminus L_1}}$ of the current state in L_1 are infected. Then do the same in L_2, infecting all possible sites at distance at most R from L_1 in $L_2 \setminus L_1$. Repeat this couple of operations $R + 1$ times. The configurations provided to η and η' so far are chosen independently, but updates occur at the same times for both. Next update L_1 in both η and η' to the same configuration still with infections next to L_2 as above and finally update L_2 in both chains to the same configuration, forcing them to meet.

In order for this to work, we need two ingredients. Firstly, we need to check that the rate at which this sequence of updates occurs is at least $(q/2)^{CR^2}$, which is clear from Claim 2.2.4. Secondly, we need to check that this is a valid coupling, namely that in the last two steps the two chains are indeed resampled from the same distribution. For this it suffices to see that after $R+1$ repetitions of the alternating updates in L_1 and L_2, necessarily the R sites in $L_2 \backslash L_1$ closest to L_1 are all infected. This is not surprising, since each time we provide the best possible boundary condition and so the sequence of these boundary conditions is nondecreasing.

Therefore, it remains to see that after a couple of updates as above either the boundary condition is already fully infected or it increases strictly. Assume the last R sites in $L_1 \backslash L_2$ remain unchanged after updating L_2 and then L_1 as above. Then none of the remaining non-infected sites could be updated at all, since even the best boundary condition L_2 can provide does not allow L_1 to infect them. Since it was assumed that $\{\eta\}_L = L$, this implies that all R sites are infected, as desired. □

Note that the above is sufficient to obtain Theorem 2.1.1 for $|L| \leqslant (2/q)^{CR}$.

Proof of the harder case of Proposition 2.2.2. We consider two copies of the process from Remark 2.2.3 denoted $(\eta(t))_{t\geqslant 0}$, $(\eta'(t))_{t\geqslant 0}$. It is well known [256] that it suffices to couple them so the probability that they do not meet before time T is at most $Ce^{-T/\gamma(\Delta)}$ for any T large enough. Observe that whenever several successive updates are performed at L_1 (and similarly for L_2), the final result is preserved if we discard all but the last update, since the dynamics of Remark 2.2.3 is of Glauber type. Hence, we may consider a discrete time chain with the same state space which updates L_1 at odd steps and L_2 at even ones (so the update from time 0 to time 1 is in L_1). Conditionally on the number of alternating updates N up to time T, after removing redundant ones as indicated above, the two chains η and η' meet if their discrete time versions do. We denote the latter by ω and ω'.

We assume that $\Delta \geqslant CR^2/q^{CR}$, the alternative being treated in a similar but simpler way as sketched above. We call any set $B \subset L_1 \cap L_2$ of $2R+1$ consecutive sites a *block* and say it is *infected* if $\mathcal{I}_x$ occurs for all $x \in B$.

Claim 2.2.5. Fix $\theta \in \mathcal{S}_L$ such that $\{\theta\}_L = L$ and an infected block $B = x + \{1, \dots, 2R+1\}$. Then $\{\theta\}^{\theta_B}_{\{1,\dots,x\}} = \{1, \dots, x\}$.

Proof. This follows immediately from the fact that the closure is increasing in the set of infections (since constraints are), since an infected block is the maximal possible boundary condition. □

Let us denote by $M = \ell - \lfloor \Delta/2 \rfloor + \{-R, \dots, R\}$ the middle block. Our coupling of ω and ω' is the following for integer $t \geqslant 0$.

- The two chains evolve independently between $2t$ and $2t+2$, unless

$$\{\omega(2t)\}_{L_1}^{(\omega(2t))_{L\setminus L_1}} \cap \{\omega'(2t)\}_{L_1}^{(\omega'(2t))_{L\setminus L_1}} \supset M. \tag{2.3}$$

- If Eq. (2.3) occurs, we first sample a two independent configurations $\xi, \xi' \in \mathcal{S}_{L_1}$ with the laws of $(\omega(2t+1))_{L_1}$ and $(\omega'(2t+1))_{L_1}$, given $\omega(2t)$ and $\omega'(2t)$. Let $x + \{-R, \dots, R\}$ be the rightmost block which is infected in both ξ and ξ', if it exists. We set $\omega(2t+1) = \xi \cdot (\omega(2t))_{L\setminus L_1}$ and

$$\omega'(2t+1) = \xi_{\{1,\dots,x\}} \cdot \xi'_{\{x+1,\dots,\ell\}} \cdot (\omega'(2t))_{L\setminus L_1}$$

  and sample $\omega(2t+2) = \omega'(2t+2)$ with their (common) law given the state at time $2t+1$. If no such block exists, $(\omega(2t+1))_{L_1} = \xi$ and $(\omega(2t+1))_{L_1} = \xi'$ and the two evolve independently between $2t+1$ and $2t+2$.

This is a legitimate Markov coupling of the homogeneous chains $(\omega(2t))_{t\geqslant 0}$ and $(\omega'(2t))_{t\geqslant 0}$. Indeed, by Claim 2.2.5, conditionally on $x + \{-R, \dots, R\}$ being the rightmost infected block, $\xi_{\{1,\dots,x\}}$ and $\xi'_{\{1,\dots,x\}}$ are identically distributed. Let us define $X(t) = \left|M \cap \{\omega(2t)\}_{L_1}^{(\omega(2t))_{L\setminus L_1}}\right|$ and similarly for ω'. Equation (2.3) then reads $X(t) = X'(t) = 2R+1$. We will lower bound $\min(X(t), X'(t))$ by the discrete time Markov chain $Y(t)$ on $\{0, \dots, 2R+2\}$ started at 0, which increments by 1 with probability

$$\left(1 - \left(1 - q^{4R+2}\right)^{\Delta/(4R+3)}\right)^4 \tag{2.4}$$

and jumps to 0 otherwise with an absorbing state $2R+2$ assigned only if ω and ω' have already met. We call a transition of Y to 0 a *failure.*

Lemma 2.2.6. *For all $t \geqslant 0$, we have*

$$\mathbb{P}(Y \textit{ is not absorbed}) \geqslant \mathbb{P}(\omega \textit{ and } \omega' \textit{ have not met}).$$

Proof. It suffices to prove that if Eq. (2.3) holds, ω and ω' meet in two steps at least with the probability in Eq. (2.4), while if Eq. (2.3) fails, at least with the probability in Eq. (2.4) each of X and X' not equal to $2R+1$ increases.

Assume that $X(t) = X'(t) = 2R+1$. By Claim 2.2.5 (note that if $X(t) = 2R+1$, then M can be infected inside L_1) we have $\{(\omega(2t))_{L_1}\}_{L_1}^{(\omega(2t))_{L\setminus L_1}} \supset \{1, \dots, \ell - \lceil \Delta/2 \rceil\}$. Recalling Claim 2.2.4 and the fact that the configurations ξ and ξ' are chosen independently, we obtain that the probability that $\omega(2t+2) \neq \omega'(2t+2)$ is at most $(1-q^{4R+2})^{\Delta/(4R+3)} \leqslant$ (2.4), since $\Delta \geqslant CR^2/q^{CR}$.

Next assume w.l.o.g. that $X(t) < 2R+1$. Then $\omega(2t+2)$ and $\omega'(2t+2)$ are independent conditionally on $\omega(2t), \omega'(2t)$, so it suffices to establish that

$$\mathbb{P}\left(X(t+1) > X(t) | \omega(2t) = \eta\right) \geqslant \left(1 - \left(1 - q^{2R+1}\right)^{\Delta/(4R+3)}\right)^2 \tag{2.5}$$

for any η compatible with $X(t) < 2R+1$. Consider the event $\mathcal{E}$ that in $\omega(2t+1)$ for at least one block B to the left of M all sites in $B \cap \{\omega(2t)\}_{L_1}^{(\omega(2t))_{L \setminus L_1}}$ are infected and likewise for $\omega(2t+2)$, a block B' to the right of M and $B' \cap \{\omega(2t+1)\}_{L_2}^{(\omega(2t+1))_{L\setminus L_2}}$. By Claim 2.2.4, $\mathbb{P}(\mathcal{E}|\omega(2t) = \eta)$ is bounded by the r.h.s. of Eq. (2.5). Thus, Lemma 2.2.7 below concludes the proof of Lemma 2.2.6. □

Lemma 2.2.7. *In the above setting $\mathcal{E}$ implies $X(t+1) > X(t)$.*

Proof. Fix blocks $B = x+\{-R, \dots, R\}$ and $B' = x'+\{-R, \dots, R\}$ witnessing the occurrence of $\mathcal{E}$ and denote $\theta = (\omega(2t+1))_{L_1}$, $\zeta = (\omega(2t+1))_{L\setminus L_1}$, $\theta' = (\omega(2t+2))_{L_2}$ and $\zeta' = (\omega(2t+2))_{L\setminus L_2}$ for shortness.

We know that $\{\theta\}_{L_1}^{\zeta} \cap B$ is infected. Therefore,

$$\begin{aligned} \{\theta\}_{\{x,\dots,\ell\}}^{\theta_{\{x-R,\dots,x-1\}}\cdot\zeta} &= \{\theta\}_{L_1}^{\zeta} \cap \{x, \dots, \ell\}, \\ \{\theta\}_{\{1,\dots,x\}}^{\theta_{\{x+1,\dots,x+R\}}} &= \{\theta\}_{L_1}^{\zeta} \cap \{1, \dots, x\}, \end{aligned} \tag{2.6}$$

by Claim 2.2.5 applied to the general KCM restricted to L_1 after performing the reductions from the beginning of Section 2.2. Consequently,

$$\begin{aligned} \mathcal{M} := M \cap \{\omega(2t)\}_{L_1}^{\zeta} &= M \cap \{\theta\}_{L_1}^{\zeta} = M \cap \{\theta\}_{\{x,\dots,\ell\}}^{\theta_{\{x-R,\dots,x-1\}}\cdot\zeta} \\ &\subset M \cap \{\theta_{L_1 \cap L_2} \cdot \zeta\}_{L_2}^{\zeta'} = M \cap \{\omega(2t+1)\}_{L_2}^{\zeta'}. \end{aligned} \tag{2.7}$$

Using the analogous relation for the second transition, we obtain $X(t+1) \geqslant X(t)$ and equality holds iff Eq. (2.7) and its analogue are equalities.

Assume that $X(t+1) = X(t)$. Then, for an augmented configuration $\bar{\omega}$ equal to $\omega(2t+1)$ with additionally all sites in $\mathcal{M}$ infected, neither update can modify states in $M \setminus \mathcal{M}$. Thus, for $\bar{\omega}$ the block M simultaneously has the isolation property Eq. (2.6) of B and its analogue for B'. Hence,

$$\{\bar{\omega}\}_L = \{\bar{\omega}\}_{\{1,\dots,\ell-\lfloor\Delta/2\rfloor-R-1\}}^{\bar{\omega}_M} \cup \mathcal{M} \cup \{\bar{\omega}\}_{\{\ell-\lfloor\Delta/2\rfloor+R+1,\dots,|L|\}}^{\bar{\omega}_M},$$

since the update rules of each site in M cannot look both to the left of M and to its right. Recalling that $\{\bar{\omega}\}_L \supset \{\omega(2t+1)\}_L = L$, we get $\mathcal{M} = M$ yielding the desired contradiction, since $X(t) = |\mathcal{M}|$. □

Returning to the proof of Proposition 2.2.2, clearly, in order for Y not to be absorbed at least one failure must occur in every $2R+2$ steps. Hence, the probability that η and η' have not met by time $T \geqslant 2$ is at most

$$e^{-T} \sum_{n=0}^{\infty} \frac{T^n}{n!} \left(1 - \left(1 - \left(1 - q^{4R+2}\right)^{\Delta/(4R+3)}\right)^{8R+8}\right)^{\lfloor n/(4R+4) \rfloor} \leqslant e^{-T} T^{9R} + \exp\left(-T\left(1 - A^{1/(9R)}\right)\right),$$

since N has the Poisson distribution with parameter T, setting

$$A = 1 - \left(1 - \left(1 - q^{4R+2}\right)^{\Delta/(4R+3)}\right)^{8R+8} \leqslant (8R+8) \exp\left(\frac{-\Delta q^{4R+2}}{4R+3}\right). \quad \square$$

Chapter 3

Coalescing and branching simple symmetric exclusion process

This chapter is based on joint work with Fabio Martinelli and Cristina Toninelli [218].

3.1 Introduction

In this chapter we study a coalescing and branching simple symmetric exclusion process (CBSEP) on a general finite graph $G = (V, E)$. The model was first introduced by Schwartz [319] in 1977 (also see Harris [204]) as follows. Consider a system of particles performing independent continuous time random walks on the vertex set of a (finite or infinite) graph G by jumping along each edge with rate 1, which coalesce when they meet (a particle jumping on top of another one is destroyed) and which branch with rate $\beta > 0$ by creating an additional particle at an empty neighbouring vertex. The process is readily seen to be reversible w.r.t. the Bernoulli(p)-product measure with $p = \frac{\beta}{1+\beta}$. Initially the model was introduced in order to study the biased voter model [319] (also known as Williams–Bjerknes tumour growth model [365]), which turns out to be its dual additive interacting particle system [194].[1] A further duality in between the two processes in the Sudbury–Lloyd sense [332] has been established since then, which shows that the law of CBSEP at a fixed time can be obtained as a p-thinning of the biased voter model (see [334, Exercise 3.6]). When $\beta = 0$ this model reduces to coalescing random walks, additive dual to the standard voter model, which have both been extensively studied (see e.g. [257, 258]).

[1]In fact, biased voter and Williams–Bjerknes models slightly differ on non-regular graphs. For such graphs CBSEP is the additive dual of the former.

When the graph is the d-dimensional hypercubic lattice, the first results were obtained by Bramson and Griffeath [80, 81]. In particular, they showed that the law of CBSEP converges weakly to its unique invariant measure starting from any non-empty set of particles and for any dimension d. Moreover, building on their work, Durrett and Griffeath [133] proved a shape theorem for this process, which easily implies that CBSEP on the discrete torus of side length $L \to \infty$ exhibits mixing time cutoff (but without any control on the critical window). In the case of the regular tree a complete convergence result is due to Louidor, Tessler, and Vandenberg-Rodes [260]. In the particular setting of $\mathbb{Z}$ a key observation is that the rightmost (or leftmost) particle performs a biased random walk with explicit constant drift (see Griffeath [194]). For more advanced results see e.g. the work by Sun and Swart [333].

While the main focus of the above-mentioned works was the long-time behaviour of the process on *infinite* graphs, our interest will concentrate instead on the mixing time for *finite* graphs. We determine the logarithmic Sobolev constant and relaxation time of the model quite precisely on a wide spectrum of relatively sparse finite graphs and for values of the branching rate β which are $o(1)$ as $|V| \to \infty$ (see Theorem 3.3.1 and Corollary 3.3.2). For instance, our results imply that for transitive bounded degree graphs the inverse of the logarithmic Sobolev constant and relaxation time when $\beta = 1/|V|$ are, up to a logarithmic correction, equal to the cover time of the graph. We will then use these results to strengthen and extend the findings of Pillai and Smith [299, 300] on the mixing time for the FA-1f kinetically constrained model in the same regime (see Corollary 3.3.3). Motivated by a different application to the kinetically constrained models FA-jf with $j > 1$ (see Chapter 5), we then investigate a version of the model in which the single vertex state space $\{0, 1\}$ is replaced by an arbitrary finite set and we bound its mixing time (see Theorem 3.3.5).

3.1.1 The CBSEP and g-CBSEP models

Recall CBSEP and the associated notation from Section 1.4.2, which can be seen to be equivalent to the coalescing and branching random walks described above up to a global time-rescaling. In the sequel we will always assume for simplicity that p is bounded away from 1 (e.g. $p \leqslant 1/2$).

We will also consider a *generalised* version of CBSEP, in the sequel g-CBSEP, defined as follows. We are given a graph G as above together with a probability space (S, ρ), where S is a finite set and ρ a probability measure on S. We still write $\rho = \bigotimes_{x \in V} \rho_x$ for the product probability on $\Omega^{(g)} := S^V$. In the *state space* S, we are given a bipartition $S_1 \sqcup S_0 = S$, and we write $p := \rho(S_1) \in (0, 1)$. We define the *projection* $\varphi : \Omega^{(g)} \to \Omega = \{0, 1\}^V$ by $\varphi(\omega) = (\mathbb{1}_{\{\omega_x \in S_1\}})_{x \in V}$ and we let $\Omega_+^{(g)} = \{\omega \in \Omega : \sum_x (\varphi(\omega))_x \geqslant 1\} =$

$\varphi^{-1}(\Omega_+)$. For any edge $e = \{x, y\} \in E$ we also let $E_e^{(g)}$ be the event that there exists a particle at x or at y for $\varphi(\omega)$. In g-CBSEP every edge $e = \{x, y\}$ such that $E_e^{(g)}$ is satisfied is resampled with rate 1 w.r.t. $\rho_x \otimes \rho_y(\cdot|E_e^{(g)})$. A key property is that its projection chain onto the variables $\varphi(\omega)$ coincides with CBSEP on G with parameter p. As with CBSEP, the g-CBSEP is reversible w.r.t. $\rho_+ = \rho(\cdot|\Omega_+^{(g)})$ and ergodic on $\Omega_+^{(g)}$.

3.1.2 The FA-1f KCM

Of particular relevance for us are the beautiful works of Pillai and Smith [299, 300] on FA-1f that we present next. For any positive integers d and L, set $n = L^d$, and let $\mathbb{Z}_L = \{0, 1, \ldots, L-1\}$ be the set of remainders modulo L. The d-dimensional discrete torus with n vertices, $\mathbb{T}_n^d$ in the sequel, is the set $\mathbb{Z}_L^d$ endowed with the graph structure inherited from $\mathbb{Z}^d$. For the *discrete time* version of FA-1f on $\mathbb{T}_n^d$ with $p = c/n$ [299, 300] provide a rather precise bound for the (total variation) mixing time $T_{\text{mix}}^{\text{FA}}$. Translated into the continuous time setting described above, their results read

$$\begin{aligned} C^{-1}n^2 &\leqslant T_{\text{mix}}^{\text{FA}} \leqslant Cn^2\log^{14}(n) \quad d = 2 \\ C^{-1}n^2 &\leqslant T_{\text{mix}}^{\text{FA}} \leqslant Cn^2\log(n) \qquad d \geqslant 3, \end{aligned} \tag{3.1}$$

where $C > 0$ may depend on d but not on n.

Remark 3.1.1. In [299, Section 2] it was argued that $T_{\text{mix}}^{\text{FA}}$ should be lower bounded by the time necessary to get two well-separated particles starting from one. By reversibility, and since to move an isolated particle by one step, we should first create a particle at a neighbouring site at rate p, this time should correspond $p^{-1}T_{\text{meet}}^{\text{rw}}$, where $T_{\text{meet}}^{\text{rw}}$ is the meeting time of two independent continuous time random walks on $\mathbb{T}_n^2$ with independent uniformly distributed starting points. In particular, since in two dimensions $T_{\text{meet}}^{\text{rw}} = \Theta(n\log(n))$, in [300, Remark 1.1] it was conjectured that $T_{\text{mix}}^{\text{FA}} = \Omega(p^{-1}n\log(n)) = \Omega(n^2\log(n))$ in the regime $p = \Theta(1/n)$ and this was recently confirmed by Shapira [325, Theorem 1.2]. As it will be apparent in the proof of Theorem 3.3.1(d), this heuristics together with the attractiveness of CBSEP will allow us to prove a lower bound on the logarithmic Sobolev constant and relaxation time of CBSEP on a general graph.

3.2 Preliminaries

In order to state our results we need first to recall some classical material on *mixing times* for finite Markov chains (see e.g. [223, 310]) and on the *resistance distance* on finite graphs (see [101, 339], [256, Ch. 9] and [261, Ch. 2]).

3.2.1 Mixing times and logarithmic Sobolev constant

Given a finite state space Ω and a uniformly positive probability measure μ on Ω, let $(\omega(t))_{t\geqslant 0}$ be a continuous time ergodic Markov chain on Ω reversible w.r.t. μ, and write $P^t_\omega(\omega') = \mathbb{P}(\omega(t) = \omega'|\omega(0) = \omega)$. Let also $h^t_\omega(\cdot) = P^t_\omega(\cdot)/\mu(\cdot)$ be the relative density of the law $P^t_\omega(\cdot)$ w.r.t. μ. The total variation mixing time of the chain, $T_{\rm mix}$, is defined as

$$T_{\rm mix} = \inf\left\{t > 0 : \max_{\omega\in\Omega_+} \|P^t_\omega(\cdot) - \mu(\cdot)\|_{\rm TV} \leqslant 1/(2e)\right\},$$

where $\|\cdot\|_{\rm TV}$ denotes the *total variation distance* defined as

$$\|P^t_\omega(\cdot) - \mu(\cdot)\|_{\rm TV} = \frac{1}{2}\|h^t_\omega(\cdot) - 1\|_1,$$

where $\|g\|^\alpha_\alpha = \mu(|g|^\alpha), \alpha \geqslant 1$. The ℓ^2-mixing time T_2 or, more generally, the ℓ^q-mixing times T_q, $q \geqslant 1$, are defined by

$$T_q = \inf\left\{t > 0 : \max_{\omega\in\Omega_+} \|h^t_\omega(\cdot) - 1\|_q \leqslant 1/e\right\}.$$

Clearly $T_{\rm mix} \leqslant T_q$ for $q > 1$ and it is known that for all $1 < q \leqslant \infty$ the ℓ^q-convergence profile is determined entirely by that for $q = 2$ (see e.g. [310, Lemma 2.4.6]). Moreover, (see e.g. [310, Corollary 2.2.7],

$$\frac{1}{2}\alpha^{-1} \leqslant T_2 \leqslant \alpha^{-1}(1 + \frac{1}{4}\log\log(1/\mu_*)), \tag{3.2}$$

where $\mu_* = \min_{\omega\in\Omega_+} \mu(\omega)$ and α is the *logarithmic Sobolev constant* defined as the inverse of the best constant C in the logarithmic Sobolev inequality valid for any $f : \Omega_+ \to \mathbb{R}$

$$\mathrm{Ent}(f^2) := \mu(f^2 \log(f^2/\mu(f^2))) \leqslant C\mathcal{D}(f). \tag{3.3}$$

Recalling the definition Eq. (1.3) of $T_{\rm rel}$, it is not difficult to prove that $T_{\rm rel} \leqslant T_{\rm mix}$ and that (see e.g. [122, Corollary 2.11])

$$2T_{\rm rel} \leqslant \alpha^{-1} \leqslant (2 + \log(1/\mu_*)) \times T_{\rm rel}. \tag{3.4}$$

Notation warning. In the sequel, unless otherwise indicated, all the quantities introduced above will not carry any additional label when referring to CBSEP. On the contrary, the same quantities referring to other chains, e.g. the FA-1f KCM or g-CBSEP, will always carry an appropriate superscript.

3.2.2 Resistance distance

Given a finite connected simple graph $G = (V, E)$, let $\vec{E}$ denote the set of *ordered pairs* of vertices forming an edge of E. For $\vec{e} = (u, v) \in \vec{E}$ we set $-\vec{e} = (v, u)$. Given an anti-symmetric function θ on $\vec{E}$ (that is $\theta(\vec{e}) = -\theta(-\vec{e})$) and two vertices x, y we say that θ is a *unit flow from x to y* iff $\sum_{v:(u,v)\in\vec{E}} \theta((u, v)) = 0$ for all $u \notin \{x, y\}$ and $\sum_{v:(x,v)\in\vec{E}} \theta((x, v)) = 1$. The *energy of the flow θ* is the quantity $\mathcal{E}(\theta) = \frac{1}{2}\sum_{\vec{e}\in\vec{E}} \theta(\vec{e})^2$ and we set

$$\mathcal{R}_{x,y} = \inf\{\mathcal{E}(\theta) : \theta \text{ is a unit flow from } x \text{ to } y\}. \tag{3.5}$$

The Thomson principle [340] states that the infimum in (3.5) is actually attained at a unique unit flow.

The quantity $\mathcal{R}_{x,y}$ can be interpreted as the effective resistance in the electrical network obtained by replacing the vertices of G with nodes and the edges with unit resistances. In graph theory it is sometimes referred to as the *resistance distance*. It is also connected to the behaviour of the simple random walk on G via the formula $2|E|\mathcal{R}_{x,y} = C_{x,y}$, where $C_{x,y}$ is the expected commute time between x and y. Furthermore, if we let $T_{\text{rel}}^{\text{rw}}$ be the relaxation time of the random walk, the bound $\max_{x,y} \mathcal{R}_{x,y} \leqslant c\sqrt{T_{\text{rel}}^{\text{rw}}}/d_{\min}$ holds [292, Corollary 1.1] where $c > 0$ is a universal constant (see also [9, Corollary 6.21] for regular graphs). Finally, by taking the shortest path between x, y and the flow θ which assigns unit flow to each edge of the path, $\mathcal{R}_{x,y} \leqslant d(x, y)$, where $d(\cdot, \cdot)$ is the graph distance, with equality iff x, y are linked by a unique path. In the sequel and for notation convenience we shall write $\bar{\mathcal{R}}_y$ for the spatial average $\frac{1}{n}\sum_x \mathcal{R}_{x,y}$.

Remark 3.2.1. For later use we present bounds on $\max_y \bar{\mathcal{R}}_y$ for certain special graphs. If G is the d-hypercube ($n := |V| = 2^d$) it follows from [301] that $\bar{\mathcal{R}}_y = \Theta(1/\log n)$ for all $y \in V$. If instead G is the regular b-ary tree with $b \geqslant 2$ then $\max_y \bar{\mathcal{R}}_y = \Theta(\log n)$. If G is a uniform random d-regular graph with $n \to \infty$, and d independent of n, then w.h.p. $T_{\text{rel}}^{\text{rw}} = O(1)$ [85,159], and therefore w.h.p. $\bar{\mathcal{R}}_y = \Theta(1)$ for all $y \in V$. Finally, if G is the discrete d-dimensional torus $\mathbb{T}_n^d \subset \mathbb{Z}^d$ with n vertices, then, as $n \to \infty$ and d is fixed, it follows from [261, Proposition 2.15] that

$$\max_y \bar{\mathcal{R}}_y = \Theta(1) \times \begin{cases} n & \text{if } d = 1, \\ \log(n) & \text{if } d = 2, \\ 1 & \text{if } d \geqslant 3. \end{cases}$$

3.3 Main results

Our first theorem establishes upper and lower bounds for the inverse of the logarithmic Sobolev constant, α^{-1}, and relaxation time, T_{rel}, of CBSEP in the general setting described in the introduction.

Let $T^{\text{rw}}_{\text{meet}}$ denote the expected meeting time for two *continuous time* random walks jumping along each edge at rate 1 and started from two uniformly chosen vertices of G. We refer the reader to [9, 240] for the close connections between $T^{\text{rw}}_{\text{meet}}$ and $\mathcal{R}_{x,y}$. Let also $T^{\text{rw}}_{\text{mix}}$ denote the mixing time of the *discrete time* lazy simple random walk on G (i.e. staying at its position with probability 1/2).

Theorem 3.3.1. *Let $p_n \in (0,1)$ and consider CBSEP with parameter p_n on a sequence of graphs $G = G_n = (V_n, E_n)$ with $|V_n| = n$, maximum degree $d_{\max} = d_{\max}(n)$, minimum degree $d_{\min} = d_{\min}(n)$, and average degree $d_{\text{avg}} = d_{\text{avg}}(n)$.*

(a) If $p_n = \Omega(1)$, then

$$\alpha^{-1} \leqslant O(n) \tag{3.6}$$
$$T_{\text{rel}} \leqslant O(1). \tag{3.7}$$

(b) If $p_n \to 0$, then for some absolute constant $c > 0$

$$\alpha^{-1} \leqslant c \max\left(\frac{d_{\text{avg}} d^2_{\max}}{d^2_{\min}} T^{\text{rw}}_{\text{mix}} \log(n), \left(\max_y \bar{\mathcal{R}}_y\right) n|\log(p_n)|\right) \tag{3.8}$$
$$T_{\text{rel}} \leqslant cn \max_y \bar{\mathcal{R}}_y. \tag{3.9}$$

(c) There exists an absolute constant $c > 0$ such that for all $p_n \in (0,1)$

$$\alpha^{-1} \geqslant \frac{cn}{d_{\text{avg}}} \tag{3.10}$$
$$T_{\text{rel}} \geqslant \frac{1 - \mu(\sum_x \omega_x = 1)}{p d_{\text{avg}}}. \tag{3.11}$$

(d) If $p_n = O(1/n)$, then we have the stronger bound

$$\alpha^{-1} \geqslant T^{\text{rw}}_{\text{meet}} \Omega(1 + |\log(np_n)|) \tag{3.12}$$
$$T_{\text{rel}} \geqslant T^{\text{rw}}_{\text{meet}} \Omega(1). \tag{3.13}$$

For the reader's convenience, and in view of our application to the FA-1f KCM, we detail the above bounds for the graphs discussed in Remark 3.2.1 when $p_n = \Theta(1/n)$.

Corollary 3.3.2. *In the setting of Theorem 3.3.1 assume that $p_n = \Theta(1/n)$. Then:*

(1) hypercube:

$$\Theta\left(\frac{n}{\log n}\right) = 2T_{\text{rel}} \leqslant \alpha^{-1} \leqslant O(n),$$

(2) *regular b-ary tree, $b \geqslant 2$ independent of n:*

$$\Theta(n\log(n)) = 2T_{\mathrm{rel}} \leqslant \alpha^{-1} \leqslant O(n\log^2(n)),$$

(3) *uniform random d-regular graph, d independent of n: w.h.p.*

$$\Theta(n) = 2T_{\mathrm{rel}} \leqslant \alpha^{-1} \leqslant O(n\log(n)),$$

(4) *discrete torus $\mathbb{T}_n^d$ with d independent of n:*

$$\alpha^{-1} \leqslant O(1) \times \begin{cases} n^2\log(n) & d=1,\\ n\log^2(n) & d=2,\\ n\log(n) & d\geqslant 3,\end{cases}$$

and

$$\alpha^{-1} \geqslant 2T_{\mathrm{rel}} = \Theta(1) \times \begin{cases} n^2 & d=1,\\ n\log(n) & d=2,\\ n & d\geqslant 3.\end{cases}$$

The corollary follows immediately from Theorem 3.3.1(b) and (d) together with Remark 3.2.1, the well-known results on $T^{\mathrm{rw}}_{\mathrm{mix}}$ for each graph and the fact that (see [10, 101, 240]) for the graphs in Remark 3.2.1 it holds that

$$T^{\mathrm{rw}}_{\mathrm{meet}} = \Theta(n) \max_y \bar{\mathcal{R}}_y.$$

Indeed, the upper bounds on $T^{\mathrm{rw}}_{\mathrm{mix}}$ are only needed to see that for these graphs the maximum in the r.h.s. of (3.8) is achieved by the second term. Using Corollary 3.3.2 together with (1.6) and (3.2), we immediately get the following consequences for the FA-1f KCM to be compared with the r.h.s. of (3.1).

Corollary 3.3.3. *Consider the FA-1f KCM on $G = \mathbb{T}_n^d$ with parameter $p_n = \Theta(1/n)$ and let $T^{\mathrm{FA}}_{\mathrm{mix}}$ and T^{FA}_2 be its mixing time and ℓ^2-mixing time respectively. Then*

$$\begin{aligned} T^{\mathrm{FA}}_{\mathrm{mix}} \leqslant T^{\mathrm{FA}}_2 \leqslant \left(\alpha^{\mathrm{FA}}\right)^{-1}\log(n) \leqslant cn\log(n)\alpha^{-1} \\ \leqslant O(1) \times \begin{cases} n^3\log^2(n) & d=1\\ n^2\log^3(n) & d=2\\ n^2\log^2(n) & d\geqslant 3.\end{cases}\end{aligned} \tag{3.14}$$

Remark 3.3.4. Our results in $d \geqslant 2$, besides being more directly proved than in [299, 300], hold in the stronger logarithmic Sobolev sense, and extend to other graphs, e.g. all the graphs discussed in Corollary 3.3.2. Furthermore, contrary to the approach followed in [299, 300], our methods can be easily adapted to cover other regimes of p_n. For $d = 1$ the above upper bound on $T^{\mathrm{FA}}_{\mathrm{mix}}$ can be proved to also be sharp up to logarithmic corrections, using the technique discussed in [88, Section 6.2].

Our second theorem concerns the total variation mixing time of the generalised model, g-CBSEP.

Let τ_{cov} denote the cover time of the simple random walk on G (see e.g. [256, Ch. 11] and also [101] for a close connection between the average cover time and the resistance distance), and let

$$T_{\text{cov}}^{\text{rw}} = \inf\left\{t > 0 : \max_{x \in V} \mathbb{P}_x(\tau_{\text{cov}} > t) \leqslant 1/e\right\}.$$

Theorem 3.3.5. *Consider g-CBSEP on a finite connected graph G of minimum degree $d_{\min}$ with parameter $p = \rho(S_1)$ and let T_{mix} be the mixing time of CBSEP on G with parameter p. Then there exists a universal constant $c > 0$ such that*

$$T_{\text{mix}} \leqslant T_{\text{mix}}^{g\text{-}CBSEP} \leqslant c(T_{\text{mix}} + T_{\text{cov}}^{\text{rw}}/d_{\min}).$$

The main reason to bound the total variation mixing time of g-CBSEP, instead of the ℓ^q-mixing times as for CBSEP, is that the scaling of the logarithmic Sobolev constant for g-CBSEP is very different from that of the CBSEP, as the following example shows.

Example 3.3.6. Let $G = \mathbb{T}_n^2$, $p_n = 1/n$, $S = \{0, 1, 2\}$, and $\rho(1) = p, \rho(0) = \rho(2) = (1-p)/2$. Then,

$$\left(\alpha^{g\text{-CBSEP}}\right)^{-1} = n^{3/2+o(1)}. \tag{3.15}$$

In the same setting Corollary 3.3.2 gives $\alpha^{-1} = O(n \log^2(n))$. To prove (3.15) it is enough to take as test function in the logarithmic Sobolev inequality for g-CBSEP the indicator that a vertical strip of width $\lfloor\sqrt{n}/2\rfloor$ of the torus is in state 0.

3.4 CBSEP—Proof of Theorem 3.3.1

For this section we work with CBSEP in the setting of Theorem 3.3.1 and abbreviate $p = p_n$. In the sequel c shall denote an absolute constant whose value may change from line to line.

3.4.1 Upper bounds—Proof of Theorem 3.3.1(a) and (b)

Let us first prove the easy upper bound Theorem 3.3.1(a), assuming that $p = \Omega(1)$. We know from [89, Theorem 6.4] that $T_{\text{rel}}^{\text{FA}} = O(1)$. Recalling (1.6) and the definition of the relaxation time, we get that for CBSEP $T_{\text{rel}} = O(1)$, yielding (3.7). By (3.4) this gives $\alpha^{-1} = O(n)$ and concludes the proof of Theorem 3.3.1(a).

The rest of this section is dedicated to the proof of the main upper bound—Theorem 3.3.1(b). The starting point is the following decomposition of the entropy of any $f : \Omega_+ \to \mathbb{R}$

$$\mathrm{Ent}(f^2) = \mu\left(\mathrm{Ent}(f^2|N)\right) + \mathrm{Ent}\left(\mu(f^2|N)\right), \tag{3.16}$$

where $N(\omega) = \sum_{x\in V} \omega_x$ is the number of particles and $\mathrm{Ent}(f^2|N)$ is the entropy of f^2 w.r.t. the conditional measure $\mu(\cdot|N)$ (see (3.3)). The first term in the r.h.s. above is bounded from above using the logarithmic Sobolev constant of the SEP on G with a fixed number of particles.

Proposition 3.4.1. *There exists an absolute constant $c > 0$ such that*

$$\mu\left(\mathrm{Ent}(f^2|N)\right) \leqslant c\log(n)\frac{d_{\mathrm{avg}}d_{\max}^2}{d_{\min}^2}T_{\mathrm{mix}}^{\mathrm{rw}}\mathcal{D}(f),$$

Proof. Let

$$\mathcal{D}_G^{\mathrm{SEP}}(f) = \frac{1}{2}\sum_{e\in E}\mu\left((f(\omega^e) - f(\omega))^2\right),$$

where ω^e is the configuration obtained from ω by swapping the states at the endpoints of the edge e, denote the Dirichlet form of the symmetric simple exclusion process on G. Similarly let

$$\mathcal{D}_{K_n}^{\mathrm{BL}}(f) = \frac{1}{2n}\sum_{e\in E(K_n)}\mu\left((f(\omega^e) - f(\omega))^2\right)$$

be the Dirichlet form of the Bernoulli-Laplace process on the complete graph K_n. The main result of [12, Theorem 1] implies that[2]

$$\mathcal{D}_{K_n}^{\mathrm{BL}}(f) \leqslant c\frac{2|E|}{n}\frac{d_{\max}^2}{d_{\min}^2}T_{\mathrm{mix}}^{\mathrm{rw}}\mathcal{D}_G^{\mathrm{SEP}}(f).$$

Using $2|E| = \sum_x d_x$ we get, in particular, that

$$\mathcal{D}_{K_n}^{\mathrm{BL}}(f) \leqslant c\frac{d_{\mathrm{avg}}d_{\max}^2}{d_{\min}^2}T_{\mathrm{mix}}^{\mathrm{rw}}\mathcal{D}_G^{\mathrm{SEP}}(f).$$

On other hand, the logarithmic Sobolev constant of the Bernoulli-Laplace process on K_n with $k \in \{1,\dots,n-1\}$ particles is bounded by $c\log n$ uniformly in k [252, Theorem 5]. Hence,

$$\mu\left(\mathrm{Ent}(f^2|N)\right) \leqslant c\log(n)\mathcal{D}_{K_n}^{\mathrm{BL}}(f) \leqslant c\log(n)\frac{d_{\mathrm{avg}}d_{\max}^2}{d_{\min}^2}T_{\mathrm{mix}}^{\mathrm{rw}}\mathcal{D}_G^{\mathrm{SEP}}(f).$$

The proposition then follows using $p \leqslant 1/2$ and

$$\mathcal{D}_G^{\mathrm{SEP}}(f) \leqslant \frac{(2-p)}{(1-p)}\mathcal{D}(f). \qquad \square$$

[2]Actually the comparison result proved in [12] is much stronger, since it concerns weighted exchange processes on G and on K_n.

We now examine the second term $\mathrm{Ent}(\mu(f^2|N))$ in the r.h.s. of (3.16). Let

$$g(k) := \mu\left(f^2|N = k\right)^{1/2}$$

for $k \geqslant 1$, so that $\mathrm{Ent}(\mu(f^2|N)) = \mathrm{Ent}_\gamma(g^2)$, where γ is the probability law of N on $\{1, \dots, n\}$. Clearly, γ is $\mathrm{Bin}(n, p)$ conditioned to be positive, so that for any $2 \leqslant k \leqslant n$ we have

$$\gamma(k)(1-p)k = \gamma(k-1)p(n-k+1). \tag{3.17}$$

Proposition 3.4.2. *There exists an absolute constant $c > 0$ such that*

$$\mathrm{Ent}_\gamma(g^2) \leqslant c\log(1/p) \times p\sum_{y\in V}\mu\left(\left[f(\omega^y) - f(\omega)\right]^2(1-\omega_y)\right).$$

where we recall that ω^y denotes the configuration ω flipped at y.

Proof. The proof starts with a logarithmic Sobolev inequality for γ w.r.t. a suitably chosen reversible death and birth process on $\{1, \dots, n\}$.

Lemma 3.4.3. *There exists an absolute constant $c > 0$ such that for any non-negative function $g : \{1, \dots, n\} \to \mathbb{R}$*

$$\mathrm{Ent}_\gamma(g^2) \leqslant c\log(1/p) \times \sum_{k=2}^{n}\gamma(k)k[g(k) - g(k-1)]^2.$$

Leaving the tedious proof to Section 3.A, we move on to bounding the r.h.s. above for the special choice $g = \mu(f^2|N)^{1/2}$.

Claim 3.4.4. For any $2 \leqslant k \leqslant n$ we have

$$(g(k) - g(k-1))^2 \leqslant \frac{A_k^2}{g^2(k-1) + g^2(k)},$$

where

$$A_k = \frac{1}{n-k+1}\sum_{y\in V}\mu\left((1-\omega_y)\left[f^2(\omega) - f^2(\omega^y)\right]|N = k-1\right).$$

Proof. We first observe that

$$[g(k) - g(k-1)]^2 = \frac{[g^2(k) - g^2(k-1)]^2}{[g(k) + g(k-1)]^2} \leqslant \frac{[g^2(k) - g^2(k-1)]^2}{g^2(k) + g^2(k-1)}. \tag{3.18}$$

Next we write

$$\begin{aligned} g^2(k-1) &= \sum_{\omega:\ N(\omega)=k-1}\frac{\mu(\omega)}{\gamma(k-1)}f^2(\omega) \\ &= \frac{1}{n-k+1}\frac{1}{\gamma(k-1)}\sum_{y\in V}\sum_{\omega:\ N(\omega)=k-1}\mu(\omega)(1-\omega_y)f^2(\omega). \end{aligned}$$

With the change of variable $\eta = \omega^y$ we get that the r.h.s. above is equal to

$$\frac{1}{n-k+1}\frac{1}{\gamma(k-1)}\sum_{y\in V}\sum_{\omega:\ N(\omega)=k-1}\mu(\omega)(1-\omega_y)\left[f^2(\omega)-f^2(\omega^y)\right] \\ + \frac{\gamma(k)(1-p)}{p(n-k+1)\gamma(k-1)}\sum_{y\in V}\sum_{\eta:\ N(\eta)=k}\frac{\mu(\eta)}{\gamma(k)}\eta_y f^2(\eta),$$

the second line being equal to $g^2(k)$ by (3.17). In conclusion $g^2(k-1) = g^2(k) + A_k$ and the claim follows from (3.18). □

Claim 3.4.5. For any $2 \leqslant k \leqslant n$ we have

$$A_k^2 \leqslant \frac{2\left(g^2(k-1)+g^2(k)\right)}{n-k+1}\sum_{y\in V}\mu\left(\left[f(\omega)-f(\omega^y)\right]^2(1-\omega_y)|N=k-1\right).$$

Proof. Using $f^2(\omega)-f^2(\omega^y) = (f(\omega)-f(\omega^y))(f(\omega)+f(\omega^y)$ and the Cauchy-Schwarz inequality w.r.t. $\mu(\cdot|N=k-1,\omega_y=0)$, we get

$$A_k \leqslant \mathrm{Av}\Big(\mu([f(\omega)-f(\omega^y)]^2|N=k-1,\omega_y=0)^{1/2}\times \\ \mu([f(\omega)+f(\omega^y)]^2|N=k-1,\omega_y=0)^{1/2}\Big),$$

where for any $h : V \to \mathbb{R}$

$$\mathrm{Av}(h) := \frac{1}{n-k+1}\sum_{y\in V}\mu\left((1-\omega_y)|N=k-1\right)h(y).$$

Another application of the Cauchy-Schwarz inequality, this time w.r.t. $\mathrm{Av}(\cdot)$, gives

$$A_k^2 \leqslant \frac{1}{n-k+1}\sum_{y\in V}\mu\left(\left[f(\omega)-f(\omega^y)\right]^2(1-\omega_y)|N=k-1\right)\times \\ \frac{2}{n-k+1}\sum_{z\in V}\mu\left(\left[f^2(\omega)+f^2(\omega^z)\right](1-\omega_z)|N=k-1\right).$$

Inside the second factor in the above r.h.s. the term containing $f^2(\omega)$ is equal to $2\mu(f^2|N=k-1) = 2g^2(k-1)$. Similarly, the term containing $f^2(\omega^y)$, after the change of variable $\eta = \omega^y$ and recalling (3.17), equals

$$\frac{2}{n-k+1}\frac{\gamma(k)k(1-p)}{p\gamma(k-1)}\mu\left(f^2(\eta)|N=k\right) = 2g^2(k).$$

□

Combining Claims 3.4.4 and 3.4.5, we get that

$$(g(k)-g(k-1))^2 \leqslant \frac{A_k^2}{g^2(k-1)+g^2(k)} \leqslant \sum_{y\in V} \frac{2\mu\left([f(\omega)-f(\omega^y)]^2(1-\omega_y)|N=k-1\right)}{n-k+1}. \quad (3.19)$$

Using (3.19) together with (3.17), we get

$$\begin{aligned}\sum_{k=2}^{n} \gamma(k)k[g(k)-g(k-1)]^2 &\leqslant \sum_{k=2}^{n} \frac{2k\gamma(k)}{n-k+1} \sum_{y\in V} \mu\left([f(\omega)-f(\omega^y)]^2(1-\omega_y)|N=k-1\right) \\ &= \frac{2}{1-p} p \sum_{y\in V} \mu\left([f(\omega)-f(\omega^y)]^2(1-\omega_y)\right).\end{aligned}$$

Using the above bound together with Lemma 3.4.3 we get the statement of Proposition 3.4.2. □

The final step in the proof of (3.8) is the following comparison between the quantity $p\sum_{y\in V}\mu\left([f(\omega)-f(\omega^y)]^2(1-\omega_y)\right)$ and the Dirichlet form $\mathcal{D}(f)$ using electrical networks. Recall the definition of the resistance distance and of $\max_y \bar{\mathcal{R}}_y$ given in Section 3.2.2.

Proposition 3.4.6.

$$p\sum_{x\in V} \mu((f(\omega^x)-f(\omega))^2(1-\omega_x)) \leqslant 4n \max_{y\in V} \bar{\mathcal{R}}_y \times \mathcal{D}(f).$$

Proof. We will identify $\omega \in \{0,1\}^V$ with its set of particles $\{x \in V : \omega_x = 1\}$ and we set $F_\omega(u) := f(\omega \cup \{u\}), u \in V$. For each $\vec{e} = (u,v) \in \vec{E}$ we also write $\nabla_{\vec{e}} F_\omega := F_\omega(v) - F_\omega(u)$. Given $x \in V$ and $\omega \in \Omega_+$, let $y_\omega \in V$ be an arbitrarily chosen vertex such that $\omega_{y_\omega} = 1$, and let θ_* be the optimal (i.e. with the smallest energy) unit flow from x to y_ω. By applying [261, Lemma 2.9] to the function F_ω and using the Cauchy-Schwarz inequality, we get that for any $\omega \in \Omega_+$ and $x \in V$ such that $\omega_x = 0$

$$\begin{aligned}(f(\omega^x)-f(\omega))^2 &= (F_\omega(x)-F_\omega(y_\omega))^2 \\ &= \left(\frac{1}{2}\sum_{\vec{e}\in\vec{E}} \theta_*(\vec{e})\nabla_{\vec{e}}F_\omega\right)^2 \leqslant \mathcal{E}_{x,y_\omega} \times \frac{1}{2}\sum_{\vec{e}\in\vec{E}} (\nabla_{\vec{e}}F_\omega)^2.\end{aligned}$$

Hence,

$$\sum_{x\in V}(f(\omega^x)-f(\omega))^2(1-\omega_x)\leqslant n\left(\max_{y\in V}\bar{\mathcal{R}}_y\right)\times\frac{1}{2}\sum_{\vec{e}}(\nabla_{\vec{e}}F_\omega)^2.$$

We next transform the generic term in the sum above into a Dirichlet form term for CBSEP. For any $\vec{e}=(u,v)\in\vec{E}$ we have

$$\begin{aligned}&p\mu(\omega)(\nabla_{\vec{e}}F_\omega)^2\\&=\mu(\omega\cup\{u\})\times\begin{cases}0 & \{u,v\}\subset\omega\\ p(f(\omega\cup\{v\})-f(\omega))^2 & u\in\omega\not\ni v\\ (1-p)(f(\omega\cup\{v\})-f(\omega\cup\{u\}))^2 & \{u,v\}\cap\omega=\varnothing.\end{cases}\end{aligned}$$

Comparing with the expression of $\mathcal{D}(f)$, (1.5), we get immediately that

$$\frac{1}{2}\sum_{\omega\in\Omega_+}p\mu(\omega)\sum_{\vec{e}\in\vec{E}}(\nabla_{\vec{e}}F_\omega)^2\leqslant 4\mathcal{D}(f).$$ □

We are now ready to prove (3.8). Using Proposition 3.4.1 the first term in the r.h.s. of (3.16) is bounded from above by

$$c\log(n)\frac{d_{\mathrm{avg}}d_{\mathrm{max}}^2}{d_{\mathrm{min}}^2}T_{\mathrm{mix}}^{\mathrm{rw}}\mathcal{D}(f).$$

In turn, Proposition 3.4.2 combined with Proposition 3.4.6 gives that the second term in the r.h.s. of (3.16) is bounded from above by

$$c\log(1/p)\times 4n\max_{y\in V}\bar{\mathcal{R}}_y\times\mathcal{D}(f).$$

In conclusion,

$$\mathrm{Ent}(f^2)\leqslant c\max\left(\log(n)\frac{d_{\mathrm{avg}}d_{\mathrm{max}}^2}{d_{\mathrm{min}}^2}T_{\mathrm{mix}}^{\mathrm{rw}},\log(1/p)\times 4n\max_{y\in V}\bar{\mathcal{R}}_y\right)\times\mathcal{D}(f),$$

so that the best constant in the logarithmic Sobolev inequality (3.3) satisfies (3.8).

Turning to (3.9), Proposition 3.4.6 alone is enough to conclude. Indeed, using the two-block argument of [59, Lemma 6.6] (see also Lemma 6.5 and Proposition 6.2 therein) and the well-known fact that the variance w.r.t. a product measure is at most the average of the sum of variances over single spins (see e.g. [20, Chapter 1]), we get

$$\mathrm{Var}(f)\leqslant cp\sum_{x\in V}\mu((f(\omega^x)-f(\omega))^2(1-\omega_x)).$$

The desired bound (3.9) then follows from (1.3) and Proposition 3.4.6.

3.4.2 Lower bounds—Proof of Theorem 3.3.1(c) and (d)

Inject $f = \mathbb{1}_{\{N=1\}}$, the indicator of having exactly one particle, in the logarithmic Sobolev inequality (3.3). For $c > 0$ small enough we have

$$\frac{\mathrm{Ent}(f^2)}{\mathcal{D}(f)} = \frac{\mu(N=1)|\log(\mu(N=1))|}{\frac{2|E|\mu(N=1)}{n}\cdot\frac{p}{2-p}} \geqslant \frac{|\log\mu(N=1)|}{pd_{\mathrm{avg}}} \geqslant c\frac{n}{d_{\mathrm{avg}}},$$

since $\mu(N=1) = np(1-p)^{n-1}/(1-(1-p)^n)$. To check the last inequality, one may distinguish the cases np sufficiently large/of order 1/sufficiently small. This proves (3.10). Using the same function, so that $\mathrm{Var}(f) = \mu(N=1)(1-\mu(N=1))$, we obtain (3.11) in the same way. This concludes the proof of Theorem 3.3.1(c).

The rest of this subsection is dedicated to the proof of the main lower bound—Theorem 3.3.1(d), so we assume that $p_n = O(1/n)$. Let $\lambda_0 > 0$ be the smallest eigenvalue, restricted to the event $\{N \geqslant 2\}$ that there are at least two particles, of $-\mathcal{L}$, where $\mathcal{L}$ is the generator of CBSEP. By [175, Lemma 4.2, Equation (1.4)] we have that

$$\begin{aligned}\alpha^{-1} &\geqslant \lambda_0^{-1}|\log(\mu(N\geqslant 2))|,\\ T_{\mathrm{rel}} &\geqslant \lambda_0^{-1}(1-\mu(N\geqslant 2)),\end{aligned}$$

the second inequality being easy to check from the definition. It is well known (see e.g. [223, Section 3.4]) that

$$\lambda_0^{-1} \geqslant \mathbb{E}_{\mu(\cdot|N\geqslant 2)}(\tau),$$

where τ is the first time when $N = 1$. Putting these together and recalling that $p_n = O(1/n)$, we obtain

$$\begin{aligned}\alpha^{-1} &\geqslant \mathbb{E}_{\mu(\cdot|N\geqslant 2)}(\tau)|\log(\mu(N\geqslant 2))| \geqslant \mathbb{E}_{\mu(\cdot|N\geqslant 2)}(\tau)\Omega(1+|\log(np_n)|),\\ T_{\mathrm{rel}} &\geqslant \mathbb{E}_{\mu(\cdot|N\geqslant 2)}(\tau)\mu(N=1) \qquad\quad \geqslant \mathbb{E}_{\mu(\cdot|N\geqslant 2)}(\tau)\Omega(1).\end{aligned}$$

In turn, again using that $p_n = O(1/n)$, we get

$$\mathbb{E}_{\mu(\cdot|N\geqslant 2)}(\tau) \geqslant \mu(N=2|N\geqslant 2)\mathbb{E}_{\mu(\cdot|N=2)}(\tau) \geqslant \Omega(1)\mathbb{E}_{\mu(\cdot|N=2)}(\tau).$$

It is not hard to see (e.g. via a graphical construction—see Section 3.5.1) that CBSEP stochastically dominates a process of coalescing random walks with birth rate 0, which we will call CSEP. Therefore, $\mathbb{E}_\omega(\tau) \geqslant \mathbb{E}_\omega^{\mathrm{CSEP}}(\tau)$ for any $\omega \in \Omega_+$. Furthermore, CSEP started with two particles has the law of two independent continuous time random walks which jump along each edge with rate $(1-p)/(2-p)$ and coalesce when they meet. Hence, we obtain (3.12) and (3.13), concluding the proof of Theorem 3.3.1(d).

3.5 g-CBSEP—Proof of Theorem 3.3.5

3.5.1 Graphical construction

We start by introducing a graphical construction of g-CBSEP for all initial conditions. The graphical construction of CBSEP can then be immediately deduced by considering the special case $S_1 := \{1\}$ and $S_0 := \{0\}$.

To each edge $e \in E$ we associate a Poisson process of rate $p/(2-p)$ of arrival times $(t_n^e)_{n=1}^{\infty}$. Similarly, to each oriented edge $\vec{e} \in \vec{E}$ we associate a Poisson process of rate $(1-p)/(2-p)$ of arrival times $(t_n^{\vec{e}})_{n=1}^{\infty}$. All the above processes are independent as $e, \vec{e}$ vary in $E, \vec{E}$ respectively. Furthermore, for $e \in E$, $\vec{e} \in \vec{E}$ and $n \geqslant 1$, we define X_n^e and $X_n^{\vec{e}}$ to be mutually independent random variables taking values in S^2. We assume that for all n and $(u, v) \in \vec{E}$, the law of $X_n^{(u,v)}$ is $\rho_u(\cdot|S_1) \otimes \rho_v(\cdot|S_0)$. Similarly, for $\{u, v\} \in E$, the law of $X_n^{\{u,v\}}$ is $\rho_u(\cdot|S_1) \otimes \rho_v(\cdot|S_1)$. Given an initial configuration $\omega(0) \in \Omega^{(g)}$ and a realization of the above variables, we define the realization of g-CBSEP $\omega(t)$ as follows.

Fix $t \geqslant 0$, let t^* be the first arrival time after t, and let $\{x, y\}$ be the endpoints of the edge where it occurs. We set $\omega_z(t^*) = \omega_z(t)$ for all $z \in V \backslash \{x, y\}$. If $E_{\{x,y\}}^{(g)}$ does not occur, that is $\omega_x(t) \in S_0$ and $\omega_y(t) \in S_0$, we set $\omega(t^*) = \omega(t)$. Otherwise, we set

$$(\omega_x(t^*), \omega_y(t^*)) = \begin{cases} X_n^{\{x,y\}} & \text{if } t^* = t_n^{\{x,y\}}, \\ X_n^{(x,y)} & \text{if } t^* = t_n^{(x,y)}. \end{cases}$$

Observation 3.5.1. Let $\omega(t)$ and $\omega'(t)$ be two g-CBSEP processes constructed using the same Poisson processes $(t_n^e)_{n=1}^{\infty}$, $(t_n^{\vec{e}})_{n=1}^{\infty}$ and variables X_n^e, $X_n^{\vec{e}}$ above, but with different initial conditions $\omega, \omega' \in \Omega^{(g)}$ satisfying $\varphi(\omega) = \varphi(\omega') = \eta \in \Omega$. Fix $t \geqslant 0$ and let $\mathcal{F}_t$ be the sigma-algebra generated by the arrival times smaller than or equal to t (but not the X_n^e and $X_n^{\vec{e}}$ variables). Then $\varphi(\omega(t)) = \varphi(\omega'(t)) =: \eta(t)$ is $\mathcal{F}_t$-measurable and only depends on ω through its projection η.

We say that a vertex $v \in V$ is *updated* if $v \in e \in E$ so that there exists $0 \leqslant t^* \leqslant t$ and n such that $t^* \in \{t_n^e, t_n^{\vec{e}}, t_n^{-\vec{e}}\}$ and the event $E_e^{(g)}$ occurs for $\omega(t^*)$, i.e. a successful update occurs at v. Denoting the set of updated vertices by ξ_t, we have

- ξ_t is $\mathcal{F}_t$-measurable and only depends on ω through its projection η,
- if $x \in \xi_t$, then $\omega_x(t) = \omega'_x(t)$ and, conditionally on $\mathcal{F}_t$, the law of $\omega_x(t)$ is $\rho(\cdot|S_{\eta_x(t)})$,
- if $x \in V \backslash \xi_t$, then $\omega_x(t) = \omega_x(0)$ and, in particular, $\eta_x(t) = \eta_x$.

In particular, for all $x \in V$ such that there exists $t_x \leqslant t$ with $(\varphi(\omega(t_x)))_x \neq (\varphi(\omega(0)))_x$, we have $\omega_x(t) = \omega'_x(t)$ (since $x \in \xi_t$).

3.5.2 Proof of Theorem 3.3.5

We are now ready to prove Theorem 3.3.5. The lower bound is an immediate consequence of the fact that the projection chain on the variables $\varphi(\omega)$ coincides with CBSEP.

For the upper bound, let μ_t^η be the law of the CBSEP η_t at time t with parameter $p = \rho(S_1)$ and starting point $\eta \in \Omega_+$. Further denote $\nu^\eta = \rho(\cdot|\varphi(\omega) = \eta)$, the measure ρ conditioned on whether or not a particle is present at each site. Since ρ is itself product, we have

$$\nu^\eta = \bigotimes_{x\in V} \rho_x(\cdot|S_{\eta_x}) = \bigotimes_{x:\eta_x=1} \rho_x(\cdot|S_1) \otimes \bigotimes_{x:\eta_x=0} \rho_x(\cdot|S_0). \tag{3.20}$$

Claim 3.5.2. The law $\rho_t^{\nu^\eta}$ of g-CBSEP with initial law ν^η at time t takes the form

$$\rho_t^{\nu^\eta}(\cdot) = \mu_t^\eta\big(\nu^{\eta_t}(\cdot)\big), \tag{3.21}$$

i.e. it is the average of $\nu^{\eta'}$ over η' distributed as the CBSEP configuration η_t started from η at time t.

Proof. Fix ω satisfying $\varphi(\omega) = \eta$. Denote by $\mathbb{P}^\omega$ the probability w.r.t. the graphical construction of g-CBSEP of Section 3.5.1 with initial condition ω, by $\mathcal{F}_t$ the sigma-algebra generated by the arrival times up to time t, as in Observation 3.5.1, and by $\mathbb{E}_{\mathcal{F}_t}$ the corresponding expectation. Then for any $\omega' \in \Omega^{(g)}$

$$\begin{aligned}\rho_t^\omega(\omega') &= \mathbb{E}_{\mathcal{F}_t}\left[\mathbb{P}^\omega(\omega(t) = \omega'|\mathcal{F}_t)\right] \\ &= \mathbb{E}_{\mathcal{F}_t}\left[\prod_{x\notin\xi_t} \mathbb{1}_{\omega'_x=\omega_x} \prod_{x\in\xi_t} \rho_x(\omega'_x|S_{(\varphi(\omega(t)))_x})\right],\end{aligned} \tag{3.22}$$

the last equality reflecting that by Observation 3.5.1 ξ_t and $\varphi(\omega(t))$ are $\mathcal{F}_t$-measurable. Again by Observation 3.5.1, ξ_t and $\varphi(\omega(t))$ are the same for all ω in the support of ν^η, so we denote the latter by $\eta(t)$. If we now average (3.22) over the initial condition ω w.r.t. ν^η and use (3.20), we obtain $\rho^{\nu^\eta}(\omega') = \mathbb{E}_{\mathcal{F}_t}[\nu^{\eta(t)}(\omega')]$, which is exactly (3.21), since $\eta(t)$ has the law μ_t^η of CBSEP with initial state η, as it is the projection of g-CBSEP with initial condition ω such that $\varphi(\omega) = \eta$. □

Next we write

$$\max_{\omega\in\Omega_+^{(g)}} \|\rho_t^\omega - \rho_+\|_{\mathrm{TV}} \leqslant \max_{\eta\in\Omega_+}\left(\max_{\omega:\varphi(\omega)=\eta} \|\rho_t^\omega - \rho_t^{\nu^\eta}\|_{\mathrm{TV}} + \|\rho_t^{\nu^\eta} - \rho_+\|_{\mathrm{TV}}\right), \tag{3.23}$$

where ρ_+ was defined in Section 3.1.1. Using Claim 3.5.2, it follows that

$$\max_{\eta\in\Omega_+} \|\rho_t^{\nu^\eta} - \rho_+\|_{\mathrm{TV}} = \max_{\eta\in\Omega_+} \|\mu_t^\eta - \mu\|_{\mathrm{TV}}, \tag{3.24}$$

where μ is the reversible measure of CBSEP with parameter p.

To bound the first term in the r.h.s. of (3.23) the key ingredient is to use the graphical construction to embed into g-CBSEP of a suitable continuous time simple random walk $(W_t)_{t\geqslant 0}$ on G with the property that g-CBSEP at time t has a "particle" at the location of W_t.

Given $\omega \in \Omega_+^{(g)}$, let $v \in V$ be such that $\varphi(\omega_v) = 1$, and let $t^* = \min\{t_n^{(u,v)} > 0\}$ be the first time an edge of the form (u, v) is resampled to produce a configuration ω' with $\omega'_u \in S_1$ and $\omega'_v \in S_0$. We then set $W_s = v$ for $s < t^*$ and $W_{t^*} = u$. By iterating the construction we construct $(W_t)_{t\geqslant 0}$ with $W_0 = v$. It is clear that $\varphi(\omega_{W_t(\omega)}(t)) = 1$ for all t and that the law $\mathbb{P}_v(\cdot)$ of $(W_t)_{t\geqslant 0}$ is that of a continuous-time random walk started at v and jumping to a uniformly chosen neighbour at rate $d_{W_t}(1-p)/(2-p)$. We denote by σ_{cov} the cover time of $(W_t)_{t\geqslant 0}$.[3]

Observation 3.5.1 then implies

$$\max_{\eta\in\Omega_+} \max_{\omega:\varphi(\omega)=\eta} \left\|\rho_t^{\omega} - \rho_t^{\nu^\eta}\right\|_{\text{TV}} \leqslant \max_{\substack{\omega,\omega'\in\Omega_+^{(g)} \\ \varphi(\omega)=\varphi(\omega')}} \|\rho_t^{\omega} - \rho_t^{\omega'}\|_{\text{TV}} \leqslant \max_{v\in V} \mathbb{P}_v(\sigma_{\text{cov}} > t). \tag{3.25}$$

The upper bound given in the theorem now follows immediately from (3.23), (3.24), and (3.25) together with a standard comparison between σ_{cov} and the cover time of the *discrete time* simple random walk on G.

Appendix

3.A Proof of Lemma 3.4.3

Recall that

$$\gamma(k) = \binom{n}{k} \frac{p^k}{(1-p)^k} \frac{(1-p)^n}{1-(1-p)^n}$$

and consider the birth and death process on $\{1, \ldots, n\}$ reversible w.r.t. the measure γ with Dirichlet form

$$\mathcal{D}_\gamma(g) = \sum_{k=2}^{n} \gamma(k)k[g(k) - g(k-1)]^2,$$

corresponding to the jump rates $c(1, 0) = 0$ and

$$c(k, k-1) = k \qquad k = 2, \ldots, n$$
$$c(k, k+1) = (n-k)\frac{p}{1-p} \qquad k = 1, \ldots, n-1.$$

[3] We use σ_{cov} to distinguish it from the cover time τ_{cov} of the *discrete time* simple random walk on G.

Let $m = \lceil pn \rceil$ and $i = \max(2, m)$. Using [277, Proposition 4] (see also [367]) the logarithmic Sobolev constant of the above chain is bounded from above, up to an absolute multiplicative constant, by the number $C_* = C_- \vee C_+$, where

$$C_+ = \max_{j \geqslant i+1} \left(\sum_{k=i+1}^{j} \frac{1}{\gamma(k)c(k,k-1)} \right) \gamma(N \geqslant j) |\log(\gamma(N \geqslant j))|,$$
$$C_- = \max_{j \leqslant i-1} \left(\sum_{k=j}^{i-1} \frac{1}{\gamma(k)c(k,k+1)} \right) \gamma(N \leqslant j) |\log(\gamma(N \leqslant j))|. \quad (3.26)$$

Assume first that $i = m$ and let us start with C_+. For $\ell \geqslant 1$ write $a_\ell = \frac{1}{(m+\ell)\gamma(m+\ell)}$ and $S_k = \sum_{\ell=1}^{k} a_\ell$. We have

$$\frac{a_{\ell+1}}{a_\ell} = \frac{1-p}{p} \frac{m+\ell}{n-m-\ell} \geqslant 1,$$

from which it follows that for $0 < \delta < 1$ we have

$$\frac{a_{\ell+1}}{a_\ell} = 1 + \Theta(\ell/m) = e^{\Theta(\ell/m)} \qquad \ell \leqslant m,$$
$$\frac{a_{\ell+1}}{a_\ell} \geqslant \frac{(1-p)(m+\delta m)}{p(n-m)} \geqslant 1+\delta \qquad \ell \geqslant \delta m. \quad (3.27)$$

In particular, for any two integers $s \leqslant t \leqslant m$ such that $t-s \geqslant \min(\sqrt{m}, m/s)$, it holds that for some absolute constant $\beta > 1$

$$\frac{a_t}{a_s} = \prod_{\ell=s}^{t-1} \frac{a_{\ell+1}}{a_\ell} = e^{\Theta((t-s)t/m)} \geqslant \beta. \quad (3.28)$$

We first analyse the behaviour of $S_k \gamma(N \geqslant m+k) |\log(\gamma(N \geqslant m+k))|$ for $k \leqslant \delta m$ where $\delta > 0$ is a sufficiently small constant depending on β.

Lemma 3.A.1. *There exists a constant $c > 0$ such that for $\delta > 0$ small enough and $k \leqslant \delta m$ we have*

$$S_k \gamma(N \geqslant m+k) |\log(\gamma(N \geqslant m+k))| \leqslant c$$

Proof. Let $0 < \delta < 1$. Define recursively

$$k_0 = 1, \quad k_1 = \lceil \sqrt{m} \rceil, \quad k_{t+1} = k_t + \lceil m/k_t \rceil,$$

and let T be the first index such that $k_T \geqslant \delta m$. Using (3.28) together with $a_{\ell+1} \geqslant a_\ell$, $k_{t+1} - k_t \leqslant k_t - k_{t-1}$, and $k_t/m \leqslant \delta$, we claim that for any $2 \leqslant t \leqslant T-1$

$$\frac{(S_{k_{t+1}} - S_{k_t})}{(S_{k_t} - S_{k_{t-1}})} = \frac{\sum_{\ell=k_t+1}^{k_{t+1}} a_\ell}{\sum_{\ell=k_{t-1}+1}^{k_t} a_\ell} \geqslant \beta \frac{k_{t+1} - k_t}{k_t - k_{t-1}}$$
$$\geqslant \beta \left(\frac{k_t}{k_{t-1}} + \frac{k_t}{m} \right)^{-1} \geqslant \beta \left(\frac{k_t}{k_{t-1}} + \delta \right)^{-1}. \quad (3.29)$$

To prove the first inequality in (3.29), observe that for any positive non-decreasing sequence $(a_j)_{j=1}^{\infty}$ and positive integers $m \leqslant n$,

$$\begin{aligned}\frac{a_{n+1}+\cdots+a_{n+m}}{a_1+\cdots+a_n} &\geqslant \min_j\left(\frac{a_{j+m}}{a_j}\right)\left(\frac{a_{n-m+1}+\cdots+a_n}{a_1+\cdots+a_n}\right)\\ &\geqslant \min_j\left(\frac{a_{j+m}}{a_j}\right)\left(\frac{\sum_{j=n-m+1}^n a_j}{(n-m)a_{n-m}+\sum_{j=n-m+1}^n a_j}\right)\\ &\geqslant \min_j\left(\frac{a_{j+m}}{a_j}\right)\frac{m}{n},\end{aligned}$$

because $\sum_{j=n-m+1}^n a_j \geqslant a_{n-m}m$.

If now δ, t are chosen small enough and large enough, respectively, depending on the constant β above, the r.h.s. of (3.29) is greater than e.g. $\beta^{1/2} > 1$. In other words, fixing δ small enough and t_0 large enough, the sequence $\left(\left(S_{k_{t+1}} - S_{k_t}\right)\right)_{t=t_0}^T, t_0 \gg 1$, is exponentially increasing.

Now fix $k \leqslant \delta m$ and t such that $k_t \leqslant k < k_{t+1}$. Assume first that $t_0 \leqslant t < T$. Then, for some positive constant c allowed to depend on β and t_0 and to change from line to line, we have

$$\begin{aligned}S_k &\leqslant \sum_{s=t_0}^{t+1}\left(S_{k_s}-S_{k_{s-1}}\right)+S_{t_0} \leqslant c\left(S_{k_{t+1}}-S_{k_t}\right)+S_{t_0}\\ &\leqslant c\frac{k_{t+1}-k_t}{m\gamma(m+k_{t+1})} \leqslant c\frac{k_{t+1}-k_t}{m\gamma(m+k)}.\end{aligned}$$

If instead $0 \leqslant t < t_0$, we directly have that

$$S_k \leqslant ka_k \leqslant c\frac{k_{t+1}-k_t}{m\gamma(m+k)}.$$

Using the bounds

$$\gamma(N \geqslant m+k) \leqslant c\frac{m+k}{k}\gamma(m+k), \qquad |\log\left(\gamma(N \geqslant m+k)\right)| \leqslant ck^2/m,$$

we finally get that $k_t \leqslant k < k_{t+1}, t < T$,

$$S_k\gamma(N \geqslant m+k)|\log(\gamma(N \geqslant m+k))| \leqslant c\frac{(k_{t+1}-k_t)k_{t+1}}{m} \leqslant c. \qquad \square$$

Let δ be as in Lemma 3.A.1. We next consider the easier case, $k \geqslant \delta m$. By (3.27), for c large enough depending on δ and allowed to change from line to line, we have that $S_k \leqslant S_{k_T} + ca_k \leqslant ca_k$ and $\gamma(N \geqslant m+k) \leqslant c\gamma(m+k)$. Thus, for $k \geqslant \delta m$, we have that

$$\begin{aligned}S_k\gamma(N \geqslant m+k)|\log(\gamma(N \geqslant m+k))| &\leqslant \frac{c}{m+k}|\log(\gamma(N \geqslant m+k))|\\ &\leqslant c\log(1/p),\end{aligned}$$

since for all k we trivially have $\gamma(m+k) \geqslant p^{m+k}$. In conclusion, we have proved that $C_+ \leqslant O(\log(1/p))$ if $m \geqslant 2$. If instead $m = 1$, then the very same computations still give $C_+ \leqslant O(\log(1/p))$, Lemma 3.A.1 being void.

The bound of C_- follows the same pattern. If $m = O(1)$, the reader may readily check that $C_- = O(1)$ because all terms in (3.26) are $O(1)$. If instead $m \gg 1$, we still obtain $C_- = O(1)$, concluding the proof of Lemma 3.4.3.

Chapter 4

Universality for critical KCM: finite number of stable directions

This chapter is based on joint work with Fabio Martinelli and Cristina Toninelli [216], establishing the following result, proving the upper bound of Theorem 1.6.4 for class (c) and Corollary 1.6.5 for families with finite number of stable directions (recall Section 1.6).

Theorem 4.0.1. *Let $\mathcal{U}$ be a critical update family with finite set of stable directions $\mathcal{S}$ and difficulty α. Then*

$$\mathbb{E}_\mu(\tau_0) = e^{O(\log(1/q)^3/q^\alpha)}. \tag{4.1}$$

We start by providing a heuristic explanation of the relaxation mechanism underlying our main result in Section 4.1. In Section 4.2 we fix some notation and gather some preliminary tools from bootstrap percolation that are by now well established in the literature. We will not dwell on the technical aspects of the definitions and invite the reader to refer to Section 4.3 of [269], which we follow closely, for more details. For reader's convenience we have collected in Section 4.2.2 three useful technical lemmas on certain one-dimensional kinetically constrained Markov processes. Although the proof of these lemmas can be found or derived from the existing literature on KCM, we have added the most advanced one in Section 4.A for completeness. Section 4.3 contains the main new technical Poincaré inequality, while Theorem 4.0.1 is proved in Section 4.4.

4.1 Some heuristics behind Theorem 4.0.1

For a high-level and accessible introduction to the main general ideas and techniques involved in bounding from above $\mathbb{E}_\mu(\tau_0)$ we refer to [269, Section

2.4]. There, in particular, it was stressed that while the necessary intuition is developed using dynamical considerations (e.g. by guessing some efficient mechanism to create/heal infection inside the system), the actual mathematical tools are mostly analytic and based on suitable (and, unfortunately, sometimes very technical) Poincaré inequalities. This chapter makes no exception.

In order to go beyond the results of [269] and get the sharp scaling of Theorem 4.0.1 in the case of a *finite* set of stable directions, the following new key input is needed.

For simplicity imagine that $\mathcal{U}$ has only four stable directions coinciding with the four natural directions of $\mathbb{Z}^2$. For a generic model with $|\mathcal{S}| < \infty$ the mechanism is the same, the only difference being that in general 'droplets' have a more complex geometry. Assume further that $\alpha(\vec{e}_1) = 1$ and $\alpha(-\vec{e}_1) = \alpha(\pm\vec{e}_2) = 2$ (see Figure 1.2e). Consider now a *critical droplet*, i.e. a square frame D, centered at the origin, of side length $\approx C\log(1/q)/q$, $C \gg 1$, and $O(1)$-thickness, and suppose that D is infected. Then, w.h.p. (w.r.t. μ) there will be extra infected sites next to D in the $\vec{e}_1$-direction allowing D to infect $D + \vec{e}_1$. However, it will be extremely unlikely to find a pair of infected sites near each other and next to the other three sides of D because of the choice of the side length of D. We conclude that w.h.p. it is easy for D to advance forward in the $\vec{e}_1$-direction but not in the other directions. Moreover, as explained in detail in [269, Section 2.4], an efficient way to effectively realize the motion in the $\vec{e}_1$-direction is via a generalised East path. In its essence the latter can be described by the following game. At every integer time a token is added or removed (if already present) at some integer point according to the following rules:

- each integer can accommodate at most one token;
- a token can be freely added or removed at 1;
- for any $j \geqslant 2$ the operation of adding/removing a token at j is allowed iff there is already a token at $j-1$.

Given $n \in \mathbb{N}$, by an efficient path reaching distance n we mean a way of adding tokens to the original empty configuration to finally place one at n which uses a minimal number of tokens. A combinatorial result (see [112]) says that the optimal number grows like $\log_2(n)$.

The main new idea now is that, while w.h.p. the droplet D will not find a pair of infected sites (which are necessary to grow an extra layer of infection in the $\vec{e}_2$-direction) *next* to e.g. its top side, w.h.p. it will find it at the right height within distance $C\log(1/q)/q^2$ in the $\vec{e}_1$-direction (see Figure 4.1). Hence, a possible efficient way for D to move one step in the $\vec{e}_2$-direction is to:

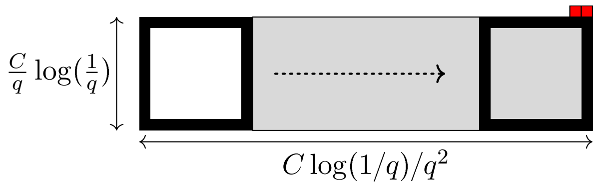

(a) The infected droplet (black frame with width $O(1)$) progressively moves to the right in an East-like way using the extra infected sites present w.h.p. in each column of the gray rectangle. This progression stops when reaching the first infected horizontal pair of sites at the correct height (red pair).

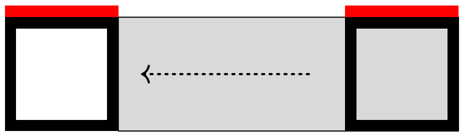

(b) The infected droplet on the right grows into an e_2-extended droplet thanks to the infected pair of sites. The movement is then reverted, progressively retracting the extended droplet in an East way until reaching the original position.

Figure 4.1 – The mechanism for the droplet to grow in the $\vec{e}_2$-direction.

(a) travel in the $\vec{e}_1$-direction in a East-like way until finding the necessary pair of infected sites within distance $C\log(1/q)/q^2$ from the origin;

(b) grow there an extra layer in the $\vec{e}_2$-direction and retrace back to its original position while keeping the acquired extra layer of infection.

A similar mechanism applies to the $-\vec{e}_2$-direction. Slightly more involved is the way in which D can advance in the $-\vec{e}_1$-direction. In this case the extra infected pair needs to be found within distance $C\log(1/q)/q^2$ from the origin in the *vertical* direction (see Figure 4.2). In order to reach it, D performs an East-like movement upwards, each of whose steps is itself realised by the back-and-forth East motion in the e_1 direction described above.
Using the result for the typical time scales of the generalised East process (see [269, formula (3.5)]) it is easy to see that the typical excursion of D for a distance $\ell \equiv C\log(1/q)/q^2$ in the $\vec{e}_1$-direction requires a time lag

$$\Delta t = q^{-|D|O(\log(\ell))} = e^{O(\log^3(1/q))/q}.$$

This time scale also bounds from above the time scale necessary to advance by one step in the "hard" directions $-\vec{e}_1, \pm\vec{e}_2$.

In conclusion, by making a "quasi-local" (*i.e* on a length scale ℓ) East-like motion in the easy direction $\vec{e}_1$, the infected critical droplet D can actually perform a sort of random walk in which each step requires a time Δt. The result of Theorem 4.0.1 becomes now plausible provided that one proves that anomalous regions of missing helping infected sites do not really constitute a serious obstacle.

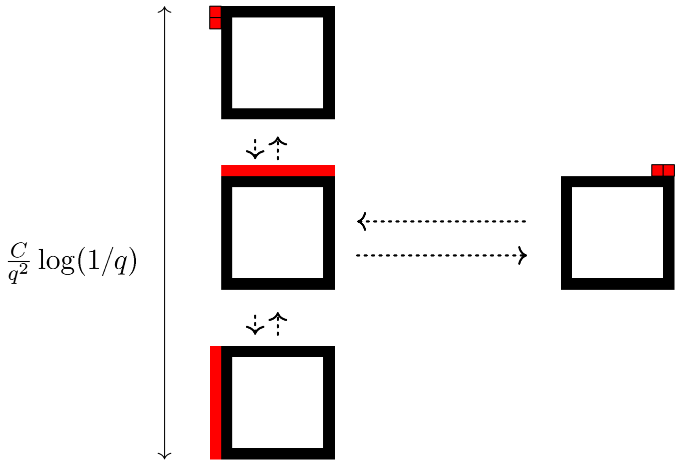

Figure 4.2 – The mechanism for the droplet growth in the $-\vec{e}_1$-direction. The droplet moves in an East way in the $\vec{e}_2$-direction by making long excursions in the $\vec{e}_1$-direction as in Figure 4.1.

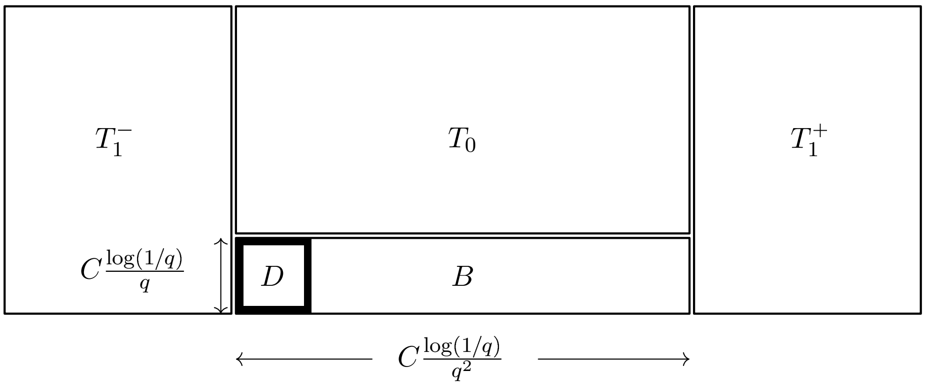

Figure 4.3 – The geometric setting for the toy model of Figure 1.2e.

The above dynamic heuristics can be turned into a rigorous argument using canonical paths. However, a much neater approach is to prove a Poincaré inequality for the $\mathcal{U}$-KCM restricted to a suitable *finite* domain of $\mathbb{Z}^2$ (see Theorem 4.3.6). More precisely, in the toy example discussed above the inequality that we establish is as follows.

Let $V = B \cup T_0 \cup T_1^- \cup T_1^+$ where $B, T_0, T_1^\pm$ are as in Figure 4.3. The set V is an example of a more general geometric construction developed in Section 4.3 and denoted *snail with base B and trapezoids $T_0, T_1^\pm$*. The ratio of the sides of the rectangle B is $\Theta(q)$ while for the other rectangles it is $\Theta(1)$.

Let Ω_0 consist of all configurations of $\{0,1\}^V$ such that:

- each column of B contains an infected site;

- each row of T_0 contains a pair of adjacent infected sites;
- each column of $T_1^{\pm}$ contains a pair of adjacent infected sites.

Notice that by choosing C large enough $\mu(\Omega_0) = 1 - o(1)$ as $q \to 0$. Then, in the key Theorem 4.3.6, we prove that for any $f : \{0,1\}^V \to \mathbb{R}$

$$\mathbb{1}_{\{D \text{ is infected}\}} \operatorname{Var}_V(f \mid \Omega_0) \leqslant e^{O(\log(1/q)^3)/q} \times \mathcal{D}(f),$$

where $\mathcal{D}(f)$ is the Dirichlet form of f (see (4.4)). One can interpret the above inequality as saying that the KCM in V restricted to the good set Ω_0 has a relaxation time at most $e^{O(\log(1/q)^3)/q}$. We prove this by an inductive procedure over $T_0, T_1^{\pm}$ which, in some sense, makes rigorous the dynamic heuristics described above.

4.2 Notation and preliminaries

In this section we gather the relevant notation and basic inputs from bootstrap percolation and KCM theories. We shall always denote spatial regions (either in $\mathbb{Z}^2$ or in $\mathbb{R}^2$) with capital letters and events in the various probability spaces with calligraphic capital letters.

4.2.1 Bootstrap percolation

Stable and quasi-stable directions

For every integer n, we write $[n] := \{0, 1, \dots, n-1\}$. We fix a critical update family $\mathcal{U}$ with difficulty $\alpha = \alpha(\mathcal{U})$ and with a *finite* set $\mathcal{S}$ of stable directions.[1] Using Definition 1.6.1 of $\alpha(\mathcal{U})$ one can fix an open semicircle C with midpoint u_0, one of whose endpoints is in $\mathcal{S}$ and such that $\max_{u \in C} \alpha(u) = \alpha$. Using [74, Lemma 5.3] (see also [70, Lemma 3.5] and [269, Lemma 4.6]) one can choose a set of rational directions[2] $\mathcal{S}' \supset \mathcal{S}$, so that for every two consecutive elements u and v of $\mathcal{S}'$ there exists an update rule $X \in \mathcal{U}$ such that $X \subset \overline{\mathbb{H}}_u \cap \overline{\mathbb{H}}_v$. The elements of $\mathcal{S}'$ are usually referred to as *quasi-stable directions*. Then our fundamental set of directions will be

$$\widehat{\mathcal{S}} = \bigcup_{u \in \mathcal{S}'} (\{u, u_0 - (u - u_0)\} + \{0, \pi/2, \pi, 3\pi/2\}). \tag{4.2}$$

In other words, we start with the stable directions, add to them the quasi-stable ones, reflect them at u_0 and finally make the set obtained invariant by rotation by $\pi/2$. By construction the cardinality of $\widehat{\mathcal{S}}$ is a multiple of 4.

[1] By Lemmas 2.6 and 2.8 of [70] this is equivalent to the fact that all (stable) directions have finite difficulty.

[2] A direction $u \in S^1$ is rational if $\tan(u) \in \mathbb{Q}$ or, equivalently, if $su \in \mathbb{Z}^2\}$ for some $s > 0$.

Remark 4.2.1. Let us note that invariance by rotation and reflection is cosmetic and one could in fact deal directly with the set of quasi-stable directions from [74], though notation would be more laborious and drawings less aesthetic.

We write $u_0, u_1, \dots u_{4k-1}$ for the elements of $\widehat{\mathcal{S}}$ ordered clockwise starting with u_0 and $\widehat{\mathcal{S}}_0$ for those elements of $\widehat{\mathcal{S}}$ belonging to the semicircle C. For all figures we shall take $\widehat{\mathcal{S}} = \{i\pi/4, i \in \{0,1,\dots,7\}\}$ and $u_0 = \pi$. When referring to u_i, the index i will be considered modulo $4k$. With this convention $\widehat{\mathcal{S}}_0 = \{u_{-k+1},\dots,u_{k-1}\}$.

$\mathcal{U}$-bootstrap percolation restricted to $\Lambda \subset \mathbb{Z}^2$

In the sequel, we will sometimes need the following slight variation of the $\mathcal{U}$-bootstrap percolation. Given $\Lambda \subset \mathbb{Z}^2$ and a set $A \subset \Lambda$ of initial infection, we will write $[A]_{\mathcal{U}}^{\Lambda}$ for the closure $\bigcup_{t \geqslant 0} A_t^{\Lambda}$ of the $\mathcal{U}$-bootstrap percolation restricted to Λ, $(A_t^{\Lambda})_{t\geqslant 0}$, defined by

$$A_{t+1}^{\Lambda} = A_t^{\Lambda} \cup \{x \in \Lambda, \exists U \in \mathcal{U}, x + U \subset A_t^{\Lambda}\}.$$

Geometric setup

We next turn to defining the various geometric domains we will need to consider. As the notation is a bit cumbersome, the reader is invited to systematically consult the relevant figures. We fix a large integer w and a small positive number δ depending on $\mathcal{U}$ (e.g. w much larger than the diameter of $\mathcal{U}$ and of the largest difficulty of stable directions), but not depending on q. When using asymptotic notation (as $q \to 0$) we will assume that the implicit constants do not depend on w, δ and q. Throughout the entire chapter we shall consider that q is small, as we are interested in the $q \to 0$ limit. In particular, we shall assume that q is so small that any length scale diverging to $+\infty$ as $q \to 0$ will be (much) larger than the constant w.

Definition 4.2.2. Consider a closed convex polygon P in $\mathbb{R}^2$. Assume that the outward normal vectors to the sides of P belong to $\widehat{\mathcal{S}}$ and that u is one of them. Then we write $\partial_u P$ for the side whose outward normal is u_i.

We can now define the notion of droplet that will be relevant for our setting (see Figure 4.4). In the sequel for $u \in S^1$ we set $\ell_u = \overline{\mathbb{H}}_u \backslash \mathbb{H}_u$ for the boundary of $\mathbb{H}_u$. Moreover, given $x \in \mathbb{R}^2$ and $s \in \mathbb{R}$, we set $\mathbb{H}_u(x) = \mathbb{H}_u + x$, $\mathbb{H}_u(s) = \mathbb{H}_u(su)$ and similarly for $\overline{\mathbb{H}}_u$ and ℓ_u. Finally, for any $u_i \in \hat{\mathcal{S}}$ we set $\rho_i = \inf\{\rho > 0, \exists x \in \mathbb{Z}^2, \langle x, u_i\rangle = \rho\}$ for the smallest positive s such that $\ell_{u_i}(s) \neq \ell_{u_i}$ and $\ell_{u_i}(s) \cap \mathbb{Z}^2 \neq \varnothing$.

Definition 4.2.3 (Quasi-stable annulus and half-annulus)**.** Fix a radius $R = R(q)$ such that $\lim_{q\to 0} R(q) = +\infty$ and let $R_i = \rho_i \left\lfloor \frac{R}{\rho_i} \right\rfloor$ for $i \in [4k]$. We call

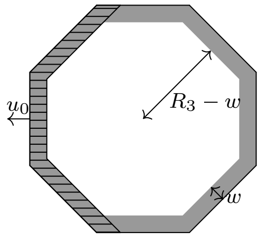

Figure 4.4 – The shaded region is the *quasi-stable annulus* A, while the hatched one is the *quasi-stable half-annulus* HA. As anticipated all the radii R_i are much larger than the width w.

the subset of $\mathbb{R}^2$

$$A = \bigcap_{i\in[4k]} \overline{\mathbb{H}}_{u_i}(R_i) \setminus \bigcap_{i\in[4k]} \mathbb{H}_{u_i}(R_i - w)$$

the *quasi-stable annulus* (or simply the *annulus*) with radius R and width w centered at the origin. We write A^{int} for the region $\bigcap_{i\in[4k]} \mathbb{H}_{u_i}(R_i - w)$ enclosed by A. Clearly, the outer boundary of A is a closed convex polygon P satisfying the assumption of Definition 4.2.2 and we write $\partial_{u_i}A$ for $\partial_{u_i}P$. We also let

$$\partial_{\hat{\mathcal{S}}_0} A = \bigcup_{u\in\hat{\mathcal{S}}_0} \partial_u A, \tag{4.3}$$

and we call

$$HA = \bigcap_{i=-k}^{k} \overline{\mathbb{H}}_{u_i}(R_i) \setminus \bigcap_{i=-k+1}^{k-1} \mathbb{H}_{u_i}(R_i - w)$$

the *quasi-stable half-annulus* of radius R and width w.

Our approach will consist in building progressively larger domains for which we can bound the Poincaré constant of the finite volume KCM process conditionally on the simultaneous occurrence of a certain likely event *and* the presence of an infected annulus. We next define these domains (see Figure 4.5). Recall that δ is a small constant depending on the update family $\mathcal{U}$.

Definition 4.2.4 (Snails)**.** Recall R and R_i from Definition 4.2.3. Let $L = L(q) > 0$ be such that $\lim_{q\to 0} L(q) = +\infty$ and assume that $\frac{L\langle u_0, u_{k-1}\rangle}{\rho_{k-1}} \in \mathbb{N}$ (i.e. $\ell_{u_{k-1}}(Lu_0)$ contains lattice sites). We call a sequence of non-negative numbers $\underline{r} = (r_0, r_1, \dots, r_{2k})$ *admissible* if

$$0 \leqslant r_0 \leqslant \delta L, \qquad r_i \leqslant \delta r_{i-1}, \qquad r_{2k} = 0.$$

Given an admissible $\underline{r}$ we call the set

$$V_L^{R,+}(\underline{r}) = \bigcap_{i=-k+1}^{k-1} \overline{\mathbb{H}}_{u_i}(R_i + L\langle u_0, u_i\rangle) \cap \bigcap_{i=k}^{3k} \overline{\mathbb{H}}_{u_i}(R_i + r_{i-k})$$

the *right-snail* with parameters $(R, L, \underline{r})$. Using the symmetric construction of $\hat{\mathcal{S}}$, the left-snail $V_L^{R,-}(\underline{r})$ with parameters $(R, L, \underline{r})$ is simply defined as

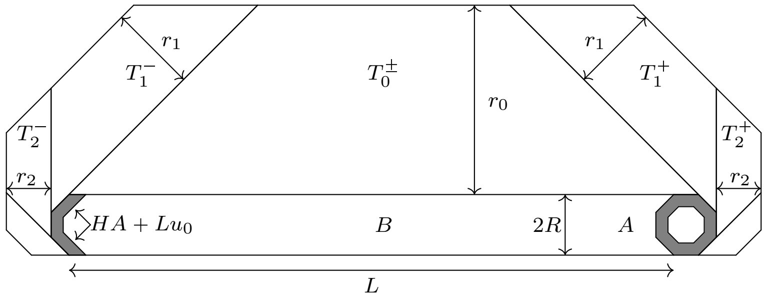

Figure 4.5 – A snail $V = V^+ \cup V^-$ with its base B and its trapezoids $T_0^\pm, T_1^\pm, \dots$. The right-snail of V is $V^+ = B \cup \bigcup_i T_i^+$ while the left-snail is $V^- = B \cup \bigcup_i T_i^-$. In Section 4.3 the shaded quasi-stable annulus A and the half-annulus $HA + Lu_0$ will act as an infected boundary condition.

the reflection of the right-snail w.r.t. the line orthogonal to u_0 and passing through the point $\frac{1}{2}Lu_0$. Finally the *snail* with parameters $(R, L, \underline{r})$ is the set

$$V_L^R(\underline{r}) = V_L^{R,+}(\underline{r}) \cup V_L^{R,-}(\underline{r}).$$

We systematically drop the parameters R, L, and $\underline{r}$ from our notation when no ambiguity arises.

Definition 4.2.5. We observe that any right-snail $V_L^{R,+}(\underline{r})$ can be thought of as the set obtained by stacking together as in Figure 4.5 its *base* defined as

$$B = V_L^R((0, \dots, 0))$$

and its trapezoids defined as

$$\begin{aligned} T_i^+ &= V_L^{R,+}(r_0, \dots, r_i, 0, \dots, 0) \backslash V_L^{R,+}(r_0, \dots, r_{i-1}, 0, \dots, 0) \\ &= (\overline{\mathbb{H}}_{u_{k+i}}(R_{k+i} + r_i) \backslash \overline{\mathbb{H}}_{u_{k+i}}(R_{k+i})) \\ &\qquad \cap \overline{\mathbb{H}}_{u_{k+i-1}}(R_{k+i-1} + r_{i-1}) \cap \overline{\mathbb{H}}_{u_{k+i+1}}(R_{k+i+1}) \end{aligned}$$

with the convention $r_{-1} = L\langle u_0, u_{k-1}\rangle$. Notice that the base B is characterized by two parameters R, L called *radius* and *length* respectively.

With this picture in mind the positive values of $\underline{r}$ coincide with the heights of the corresponding non-empty trapezoids. A similar decomposition holds for the left-snail. In the sequel, it will be convenient to partition the lattice sites in each trapezoid T_i^+ into disjoint *slices* $ST_{i,j}^+, j = 1, 2, \dots$, with each slice consisting of all the lattice sites of the trapezoid lying on a *common line* of $\mathbb{R}^2$ orthogonal to the direction u_{k+i}. Similarly for the lattice sites contained in the *truncated base* $B^\circ := B \backslash (A \cup A^{\text{int}} \cup (HA + Lu_0))$. In this case each slice, denoted $SB_j, j = 1, 2, \dots$, will consist of all the sites belonging to a common suitable translate in the u_0-direction of $\partial_{\hat{\mathcal{S}}_0} A$ defined in (4.3). Recall from Definition 4.2.3 ρ_i and $R_i = \rho_i \lfloor R/\rho_i \rfloor$, where R is the radius of the annulus A.

Definition 4.2.6. Fix a snail and suppose that its trapezoid T_i^+, $i \in [2k]$ is non-empty. The j^{th} slice of $T_i^+ \cap \mathbb{Z}^2$ is the set

$$ST_{i,j}^+ = T_i^+ \cap \ell_{u_{k+i}}(R_{k+i} + j\rho_{k+i}) \cap \mathbb{Z}^2,$$

so that $T_i^+ \cap \mathbb{Z}^2 = \bigcup_{j>0} ST_{i,j}^+$. Similarly, for the left trapezoid T_i^- if non-empty.

Turning to the truncated base B°, we first set $\lambda_{j+1} = \inf\{\lambda > \lambda_j, (\lambda u_0 + \partial_{\hat{\mathcal{S}}_0} A) \cap \mathbb{Z}^2 \neq \varnothing\}$ with $\lambda_0 = 0$. Then we define the j^{th} slice of the truncated base B° of the snail, SB_j, and its i^{th}-side, $SB_{i,j}$, as

$$SB_j = (\lambda_j + \partial_{\hat{\mathcal{S}}_0} A) \cap B^\circ \cap \mathbb{Z}^2,$$
$$SB_{i,j} = (\lambda_j + \partial_{u_i} A) \cap B^\circ \cap \mathbb{Z}^2.$$

Note that for any admissible sequence $\underline{r}$ a non-empty slice of the trapezoid T_i^+ consists of all lattice points of a segment $I \subset \mathbb{R}^2$ orthogonal to u_{k+i} with length $\Omega(r_{i-1})$ and such that $I \cap \mathbb{Z}^2 \neq \varnothing$. Similarly, the number of lattice sites in each slice of B° is $\Theta(R)$. In the sequel we will only consider non-empty slices without explicitly specifying the range of the index $j > 0$.

Helping sets

Recall Definition 1.6.1. If u is a stable direction, then the infected half-plane $\mathbb{H}_u$ needs *finitely* many (exactly $\alpha(u)$) extra infected sites in $\mathbb{R}^2 \backslash \mathbb{H}_u$ in order to infect *infinitely* many sites on the line ℓ_u. If only a finite portion of $\mathbb{H}_u$ is infected, e.g. the dashed region in Figure 4.6, then the propagation of infection to some portion of the line ℓ_u is a delicate problem. A special case which suffices for our purposes is covered in the next lemma (see [70, Lemma 3.4] and [74, Lemma 5.2]).

Lemma 4.2.7. *Fix $u = u_i$, $i \in [4k]$ and recall that w is a large enough integer (depending on $\mathcal{U}$) and let $r \geqslant w^2$. Let*

$$\Lambda := \Lambda(u, w, r) = \overline{\mathbb{H}}_{u_{i-1}}(r) \cap \overline{\mathbb{H}}_{u_i} \cap \overline{\mathbb{H}}_{u_{i+1}}(r) \cap \overline{\mathbb{H}}_{u_{i+2k}}(w)$$

be the (closed) trapezoid in Figure 4.6 of height w and bases orthogonal to u. Note that $\partial_u \Lambda \subset \ell_u$.

(a) Let $Z \subset \mathbb{Z}^2 \backslash \mathbb{H}_u$ be a set of $\alpha(u)$ sites at distance at most $\sqrt{w}$ from the origin such that $[\mathbb{H}_u \cup Z]_{\mathcal{U}} \cap \ell_u$ is infinite. Then there exist finitely many lattice points $a_1, \dots, a_m, b$, on the line ℓ_u such that the following holds. If $\Lambda \backslash \partial_u \Lambda$ and $\bigcup_{j=1}^m (Z + a_j + k_j b)$ are infected, where $k_1, \dots, k_m \in \mathbb{Z}$

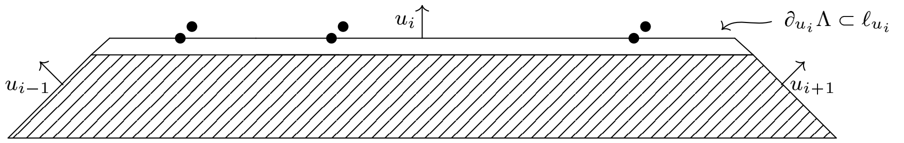

(a) An example helping set (the black dots) consisting of three disjoint copies of Z shifted along ℓ_{u_i}. In the figure $\alpha(u_i) = 2$ with $Z = \{(0,0),(1,1)\}$ and $m = 3$.

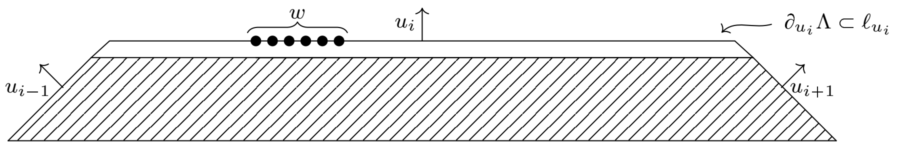

(b) Illustration of w consecutive sites of $\partial_{u_i}\Lambda$.

Figure 4.6 – The setting of Lemma 4.2.7. In the figure u_i is the upwards direction and the hatched trapezoid represents the lattice sites in $\Lambda\backslash\partial_{u_i}\Lambda$. The lemma states that if the hatched region and the black sites are infected, $\partial_{u_i}\Lambda$ also becomes infected (in the $\mathcal{U}$ bootstrap percolation process restricted to a suitable region).

are such that $\{a_1 + k_1 b, a_2 + k_2 b, \dots, a_m + k_m b\}$ form m distinct lattice sites of $\partial_u\Lambda$ at distance at least w from the endpoints of $\partial_u\Lambda$, then the $\mathcal{U}$-bootstrap percolation restricted to the larger trapezoid

$$\overline{\Lambda} = \overline{\mathbb{H}}_{u_{i-1}}(r) \cap \overline{\mathbb{H}}_{u_i}(w/2) \cap \overline{\mathbb{H}}_{u_{i+1}}(r) \cap \overline{\mathbb{H}}_{u_{i+2k}}(w)$$

is able to infect $\partial_u\Lambda$.

(b) If $\Lambda\backslash\partial_u\Lambda$ and w consecutive lattice sites in $\partial_u\Lambda$ are infected, then the $\mathcal{U}$-bootstrap percolation restricted to Λ is able to infect $\partial_u\Lambda$.

Definition 4.2.8 (u-helping sets)**.** Let $i \in [-k+1, k-1]$. Any collection of lattice sites of the form $\{a_1 + k_1 b, a_2 + k_2 b, \dots, a_m + k_m b\}$ satisfying the assumption in (a) above will be referred to as *u_i-helping set* for $\partial_{u_i}\Lambda$ or simply u_i-helping set.

4.2.2 Some KCM tools

For reader's convenience we next collect some general tools from KCM theory that will be applied several times throughout the proof of the main result.

Notation

For every statement $\mathcal{P}$ define $\mathbb{1}_{\{\mathcal{P}\}} = 1$ if $\mathcal{P}$ holds and $\mathbb{1}_{\{\mathcal{P}\}} = 0$ otherwise. For any subset Λ of $\mathbb{R}^2$ we write $(\Omega_\Lambda, \mu_\Lambda)$ for the product probability space $(\{0,1\}^{\Lambda\cap\mathbb{Z}^2}, \bigotimes_{x\in\Lambda\cap\mathbb{Z}^2}\mu_q)$. If $\Lambda = \mathbb{R}^2$, we simply write (Ω, μ). Given $f : \Omega_\Lambda \to \mathbb{R}$ we shall write $\mu_\Lambda(f)$ and $\mathrm{Var}_\Lambda(f)$ for the mean and variance of

f w.r.t. μ_Λ respectively whenever they exist. For any $\omega \in \Omega$ and $\Lambda \subset \mathbb{R}^2$ we write ω_Λ for the collection $(\omega_x)_{x\in\Lambda\cap\mathbb{Z}^2}$. Given a function $f : \Omega \to \mathbb{R}$ depending on finitely many variables we write

$$\mathcal{D}(f) = \sum_{x\in\mathbb{Z}^2} \mu(c_x \operatorname{Var}_x(f)), \tag{4.4}$$

for the KCM *Dirichlet form* of f, where $c_x(\omega)$ is the indicator of the event $\{\exists X \in \mathcal{U} : \forall y \in X, \omega_{x+y} = 0\}$ and $\operatorname{Var}_x(f) := \operatorname{Var}_{\{x\}}(f)$ denotes the conditional variance $\operatorname{Var}(f \mid (\omega_z)_{z\neq x})$. Finally, we shall write $\mathbb{P}_\mu(\cdot)$ for the law of the $\mathcal{U}$-KCM process on $\mathbb{Z}^2$ with initial law μ and $\mathbb{E}_\mu(\cdot)$ for the expectation w.r.t. $\mathbb{P}_\mu(\cdot)$.

Poincaré inequalities

We begin with a well-known general fact on product measures which we state here in ready-to-use form.

Lemma 4.2.9. *Let Λ_i, $i \in \{1,2,3\}$ be three disjoint finite subsets of $\mathbb{Z}^2$ and ν_i be a probability measures on Ω_{Λ_i}. Let ν be the product measure $\bigotimes_{i=1}^3 \nu_i$ on $\bigotimes_{i=1}^3 \Omega_{\Lambda_i}$. Then for any function f we have*

$$\nu_1(\operatorname{Var}_{\nu_2\otimes\nu_3}(f)) \leqslant \operatorname{Var}_\nu(f) \leqslant \nu_1(\operatorname{Var}_{\nu_2\otimes\nu_3}(f)) + \nu_2(\operatorname{Var}_{\nu_1\otimes\nu_3}(f)).$$

Proof of Lemma 4.2.9. The first inequality follows from the total variance formula

$$\operatorname{Var}_\nu(f) = \nu_1(\operatorname{Var}_{\nu_2\otimes\nu_3}(f)) + \operatorname{Var}_{\nu_1}(\nu_2 \otimes \nu_3(f)).$$

For the second inequality we observe that

$$\begin{aligned}\operatorname{Var}_{\nu_1}(\nu_2 \otimes \nu_3(f)) &= \nu_1((\nu_2 \otimes \nu_3(f - \nu(f))^2) \\ &\leqslant \nu((f - \nu_1 \otimes \nu_3(f))^2) = \nu_2(\operatorname{Var}_{\nu_1\otimes\nu_3}(f))\end{aligned}$$

by Jensen's inequality. □

In order to understand the general framework for the last two results, we begin by recalling a standard Poincaré inequality for n independent random variables $X_1, \ldots, X_n$ (for simplicity each one taking finitely many values). For any $f = f(X_1, \ldots, X_n)$

$$\operatorname{Var}(f) \leqslant \sum_i \mathbb{E}(\operatorname{Var}_i(f)),$$

where $\operatorname{Var}_i(f)$ is the conditional variance computed w.r.t. the variable X_i given all the other variables. The sum in the r.h.s. above can be interpreted as the Dirichlet form of the continuous time *Gibbs sampler*, reversible w.r.t. the product law of $(X_i)_i$, which with rate n chooses a random index $i \in [n]$ and resamples X_i w.r.t. its marginal. From this perspective, the above inequality tells us that the relaxation time (see e.g. [256]) of the Gibbs sampler is bounded from above by 1.

Now consider n events $(\mathcal{H}_i)_{i=1}^n$, in the sequel *facilitating events*, and suppose that each $\mathcal{H}_i$ depends only on the variables $(X_j)_{j\neq i}$. An example of a *constrained* Poincaré inequality with facilitating events $(\mathcal{H}_i)_{i=1}^n$ is the inequality

$$\operatorname{Var}(f \mid \Omega_{\mathcal{H}}) \leqslant C \sum_i \frac{\mathbb{E}(\mathbb{1}_{\mathcal{H}_i} \operatorname{Var}_i(f))}{\mathbb{P}(\Omega_{\mathcal{H}})}, \tag{4.5}$$

where $\Omega_{\mathcal{H}} = \bigcup_{i=1}^n \mathcal{H}_i$ and $C \in [1, +\infty]$. Notice that the sum in the r.h.s. above can be interpreted as the Dirichlet form of the continuous time *constrained* Gibbs sampler on $\Omega_{\mathcal{H}}$, which with rate n chooses a random index $i \in [n]$ and resamples X_i w.r.t. its marginal iff $\mathcal{H}_i$ holds. If the facilitating events are such that the constrained Gibbs sampler on $\Omega_{\mathcal{H}}$ is ergodic then $C < +\infty$.

Each one of the two results we are about to discuss next is just a special instance of the above general problem.

Lemma 4.2.10. *Let X_1, X_2 be two independent random variable taking values in two finite sets $\mathbb{X}_1, \mathbb{X}_2$. Let also $\mathcal{H} \subset \mathbb{X}_1$ with $\mathbb{P}(X_1 \in \mathcal{H}) > 0$. Then for any function $f(X_1, X_2)$*

$$\operatorname{Var}(f) \leqslant 2\mathbb{P}(X_1 \in \mathcal{H})^{-1}\mathbb{E}\left(\operatorname{Var}_1(f) + \mathbb{1}_{\{X_1 \in \mathcal{H}\}} \operatorname{Var}_2(f)\right).$$

Remark 4.2.11. The above inequality coincides with (4.5) with $\mathcal{H}_1 = \mathbb{X}_2$, $\mathcal{H}_2 = \{X_1 \in \mathcal{H}\}$ and $C = 2/\mathbb{P}(X_1 \in \mathcal{H})$. Clearly the constrained Gibbs sampler is irreducible because $\mathbb{P}_1(X_1 \in \mathcal{H}) > 0$.

Proof of Lemma 4.2.10. It follows from [88, Proof of Proposition 4.4] that

$$\begin{aligned}\operatorname{Var}(f) &\leqslant \frac{1}{1 - \sqrt{1 - \mathbb{P}(X_1 \in \mathcal{H})}} \mathbb{E}\left(\operatorname{Var}_1(f) + \mathbb{1}_{\{X_1 \in \mathcal{H}\}} \operatorname{Var}_2(f)\right) \\ &\leqslant 2\mathbb{P}(X_1 \in \mathcal{H})^{-1}\mathbb{E}\left(\operatorname{Var}_1(f) + \mathbb{1}_{\{X_1 \in \mathcal{H}\}} \operatorname{Var}_2(f)\right). \qquad \square\end{aligned}$$

The second result concerns a generalisation of the standard (finite volume) constrained Poincaré inequality for the 1-neighbour KCM process, or FA1f KCM, [88].

Let $(\widehat{S}, \widehat{\nu})$ be a finite probability space with $\widehat{\nu}$ a positive probability measure, let $\Omega_n = \widehat{S}^{[n]}$ and $\nu = \bigotimes_{i\in[n]} \nu_i$, where $\nu_i = \widehat{\nu}$ for all $i \in [n]$. Elements of Ω_n are denoted $\omega = (\omega_0, \ldots, \omega_{n-1})$ with $\omega_i \in \widehat{S}$. Fix a single site event $\mathcal{H} \subset \widehat{S}$ and a positive integer $\kappa < n$. Then, according to whether we view the set $[n]$ as the n-cycle or not, we define the *facilitating event* $\mathcal{H}_i$ as follows. If $[n]$ is the n-cycle

$$\mathcal{H}_i = \bigcap_{j=i+1}^{i+\kappa} \{\omega_j \in \mathcal{H}\} \cup \bigcap_{j=i-1}^{i-\kappa} \{\omega_j \in \mathcal{H}\}.$$

If instead $[n]$ is linear

$$\mathcal{H}_i = \begin{cases} \bigcap_{j=i-1}^{i-\kappa}\{\omega_j \in \mathcal{H}\} & \text{if } i+\kappa \geqslant n \\ \bigcap_{j=i+1}^{i+\kappa}\{\omega_j \in \mathcal{H}\} & \text{if } i-\kappa < 0 \\ \bigcap_{j=i+1}^{i+\kappa}\{\omega_j \in \mathcal{H}\} \cup \bigcap_{j=i-1}^{i-\kappa}\{\omega_j \in \mathcal{H}\} & \text{otherwise.} \end{cases}$$

In words, in the periodic case $\mathcal{H}_i$ requires the κ variables immediately after or before i (in the clockwise order) to be in a state belonging to $\mathcal{H}$, while in the linear case the same requirement holds when i is farther than κ from the boundary points of $[n]$. When i is closer than κ to e.g. the the right boundary of $[n]$, then $\mathcal{H}_i$ requires the κ variables immediately before i to be in states belonging to $\mathcal{H}$. The case when $\widehat{S} = \{0, 1\}$, $\widehat{\nu}$ is the Bernoulli$(1-q)$-measure, $\mathcal{H} = \{0\}$ and $\kappa = 1$ is the usual 1-neighbour KCM setting.

Lemma 4.2.12. *Assume that* $(1-\widehat{\nu}(\mathcal{H})^k)^{n/(3\kappa)} < \frac{1}{16}$ *and set* $\Omega_{\mathcal{H}} = \bigcup_{i=0}^{n-1} \mathcal{H}_i$*. Then, for all* $f : \Omega_n \to \mathbb{R}$

$$\mathrm{Var}_\nu(f \mid \Omega_{\mathcal{H}}) \leqslant \left(\frac{2}{\widehat{\nu}(\mathcal{H})}\right)^{O(\kappa)} \sum_{i=1}^{n} \nu\left(\mathbb{1}_{\mathcal{H}_i} \mathrm{Var}_{\nu_i}(f)\right). \tag{4.6}$$

The proof is left to Section 4.A.

Remark 4.2.13. We will apply the lemma with $(\widehat{S}, \widehat{\nu})$ equal to the probability space given by $\{0, 1\}^m$ equipped with the Bernoulli$(1-q)$ product measure conditioned on some event whose probability tends to one as $q \to 0$. The integers $1 \ll m \ll n$ may diverge to infinity as $q \to 0$ while the integer κ will be large but independent of q.

4.3 The core of the proof

In this section we prove a Poincaré inequality which will represent the key step in the proof of Theorem 4.0.1.

4.3.1 Roadmap

Before we dive into the technical details, let us give a hands-on roadmap of the argument. Although it is underlain by the dynamical intuition explained in Section 4.1, the latter is not very transparent in the Poincaré language of the formal proof.

The goal of this section is to prove Theorem 4.3.6. It says that the $\mathcal{U}$-KCM ($\mathcal{U}$ being a fixed critical update family with a finite number of stable directions) on a snail $V = V_L^R(\underline{r})$ (recall Definition 4.2.4 and Figure 4.5), conditioned on a well-chosen *super good* event $\mathcal{SG}(V)$ is able to relax in a time $\exp(\log^3(1/q)/q^\alpha)$, which is the dominating contribution leading to (4.1). For the purposes of the roadmap the reader should think of the snail as having dimensions $R = w^2 \log(1/q)/q^\alpha$, $L = q^{-3w}$ and $r_i = \delta^{i-(2k-1)} q^{-2w}, i \in [2k]$ for some small positive δ. Let us explain the Definition 4.3.2 of $\mathcal{SG}(V)$ before outlining the proof of Theorem 4.3.6.

Good and super good events

The super good event $\mathcal{SG}(V)$ will decompose as a product w.r.t. the partition of V into its annulus A, half-annulus $HA + Lu_0$, annulus interior A^{int}, truncated base B° and trapezoids $T_i^\pm$ from Definitions 4.2.3 to 4.2.5. On A (HA) we require the event $\mathcal{A}$ ($\mathcal{HA}$) that A (HA) is fully infected. These are the only *unlikely* events involved in $\mathcal{SG}(V)$ and we will denote by $\mathcal{SG}$ only events requiring the occurrence of (spatial translates of) $\mathcal{A}$ and $\mathcal{HA}$. Events of type $\mathcal{SG}$ will all have very small probability $\mu(\mathcal{SG})$ of the order of $\exp(-\log^2(1/q)/q^\alpha)$.

In turn, we will write $\mathcal{G}$ for good events, which are *likely* and only involve the presence of appropriate helping sets as in Definition 4.2.8 or sets of w consecutive infections as in Lemma 4.2.7(b). Recall the decomposition of B° into slices SB_j from Definition 4.2.6. We say that the event $\mathcal{SB}_j$ occurs if each side of SB_j (which consists of at most one segment in each direction) has a helping set for the corresponding direction. We then define $\mathcal{G}(B^\circ) = \bigcap_j \mathcal{SB}_j$ and it is not hard to see that this way the occurrence of $\mathcal{SG}(B) = \mathcal{A} \cap \mathcal{HA} \cap \mathcal{G}(B^\circ)$ implies that the infections in B are sufficient to fully infect B.

Notice that in general the event $\mathcal{SB}_j$ depends on the values of ω in the set $\bigcup_{i=0}^k SB_{j+i}$ for some $k \geqslant 0$ depending only on $\mathcal{U}$. In order to avoid this (annoying) technical detail we will use the following simplifying Assumption 4.3.1 implying that $k = 0$.

Assumption 4.3.1. For every stable direction $u \in \hat{\mathcal{S}}_0$ there exists a subset Z_u of the line ℓ_u of cardinality α such that $[\mathbb{H}_u \cup Z_u]_{\mathcal{U}} \cap \ell_u$ has infinite cardinality.

This is by no means restrictive, as the proof applies directly without this assumption up to changing $\mathcal{SB}_{i,j}$ in Definition 4.3.2, following [269, Sec. 7]. We will spare the reader the tedious details, as they already appeared previously in the above-mentioned paper. This assumption is only relevant for treating the base B, for which we will import the result from [269], where the assumption was introduced.

Having defined the good event for the base B, we now define the good event for the trapezoids of the snail V. Let $\mathcal{ST}_{i,j}^\pm$ be the event that the slice $ST_{i,j}^\pm$ in the decomposition of $T_i^\pm$ from Definition 4.2.6 contains a set of w consecutive infected sites. We then define $\mathcal{G}(T_i^\pm) = \bigcap_j \mathcal{ST}_{i,j}^\pm$. Again, by Lemma 4.2.7 it is not hard to see that if B and $T_{i'}^+$ for $i' < i$ are fully infected and $\mathcal{G}(T_i^+)$ occurs, then the $\mathcal{U}$-bootstrap percolation can also infect T_i^+ (and similarly for T_i^-).

Finally, the super good event $\mathcal{SG}(V)$ is defined as $\mathcal{SG}(B) \cap \bigcap_i (\mathcal{G}(T_i^+) \cap \mathcal{G}(T_i^-))$ and it clearly implies that the entire snail V can be infected from within.

Structure of the proof

The fact that V can be fully infected on $\mathcal{SG}(V)$ is reassuring and implies that the relaxation time we are after in Theorem 4.3.6 is finite, but we need an efficient relaxation mechanism to prove the theorem. It is not hard to see that it suffices to treat the right-snail V^+, so we concentrate on it and drop all + superscripts. In the sequel, whenever we refer to the *relaxation in a given region* Λ mathematically this will translate into proving a Poincaré inequality like the one in (4.8) with V replaced by Λ.

The proof proceeds by proving an efficient relaxation in progressively larger and larger volumes always conditioned on a corresponding $\mathcal{SG}$ event. In the process we will often rely on auxiliary constrained block dynamics of several types like those in Section 4.2.2. These auxiliary dynamics allow us to relate the relaxation in a given region to the relaxation in smaller subregions, each subregion having an additional convenient constraint on the configuration outside it. The auxiliary dynamics we will use are of FA1f type (like the one in Lemma 4.2.12) or two-blocks type (like the one in Lemma 4.2.10). By performing such reductions, we reduce the problem of proving an efficient relaxation on a large region to a similar problem on suitable smaller regions. The base case of the above inductive procedure is then treated directly. We now describe the various steps of the above iterative reduction.

The base case: the annulus interior A^{int} First, in Lemma 4.3.10 we treat A^{int} on the event $\mathcal{A}$ that the annulus is fully infected, which serves as a boundary condition. This is fairly easy and can be done in various ways. To give a formal argument, we split A^{int} into strips of bounded width (see Figure 4.7). Fully infected strips perform an FA1f auxiliary dynamics. The boundary condition provides all the sets of w consecutive infections needed for an infected strip to infect its neighbour using Lemma 4.2.7 (see Figure 4.6a).

From A^{int} to the base B Up to now we have a Poincaré inequality on the annulus and its interior. In Proposition 4.3.9 we extend that to a base B. We will not insist on this step, as it was essentially done already in [269]. Indeed, using an East-like dynamics in direction u_0 the relaxation time of B° (on $\mathcal{G}(B^\circ)$) with infected boundary condition in A was shown to be roughly $\exp(\log^3(1/q)/q^\alpha)$. Combining this with the result for A^{int}, we obtain a Poincaré inequality for B.

Adding the first trapezoid to B Our next goal is to consider the relaxation in $B \cup T_0$. In turn, this step is split into two distinct parts.

Adding the first slice $ST_{0,1}$ to B This is achieved in Lemma 4.3.13 (see Figure 4.8). Relaxation in B has already been established in the previous step, so we focus on the relaxation in $ST_{0,1}$. In doing this we are allowed our knowledge of the relaxation in B. We use the FA1f-like dynamics of Lemma 4.2.12, asking for w consecutive infections in $ST_{0,1}$ next to the site to be updated. In other words we have to understand how to efficiently resample a site $x \in ST_{0,1}$ using the $\mathcal{U}$-KCM dynamics when its neighbouring w sites are infected. Using a two-block dynamics (Lemma 4.2.10), resampling B roughly $q^{-O(1)}$ times, we may further impose the condition that the site we wish to resample has a fully infected neighbourhood in B in addition to the next w sites in $ST_{0,1}$, which are already infected. This is exactly the situation in Figure 4.8 and makes the flip of the site we want to update legal for the original $\mathcal{U}$-KCM. Thus, this step produces terms of the Dirichlet form of the $\mathcal{U}$-KCM in (4.4), as well as a term $\mathrm{Var}_B(f \mid \mathcal{SG}(B))$, which we already know how to control.

Adding more slices to B In a sense this part embodies the East-like motion of droplets in direction u_1 hinted in Section 4.1. This connection is rather indirect in the sense that the *bisection method* used to analyse the relaxation in the union of B with several slices of the first trapezoid coincides with the bisection method used to efficiently bound from above the relaxation time of the standard East model in [88].

Consider the problem of the relaxation in a snail consisting of B and $2n$ slices of the first trapezoid. Our aim is to reduce it to the same problem on two similar snails, each one with essentially the same base B but with only n slices. This is achieved in Lemma 4.3.14. We start by introducing an auxiliary constrained two-block dynamics in which B and the first n slices form the first block $\widetilde{V}$, while the second group of n slices form the second block $\overline{T}_0$ (see Figure 4.9). The constraint of the two-blocks dynamics is that a translated base $\overline{B}$ (corresponding to $\overline{V}_{i-1}$ in Figure 4.9) is super good. The base $\overline{B}$ is constructed so that together with $\overline{T}_0$ it forms a snail $\overline{V}$ with size similar to that of $\widetilde{V}$. The relaxation to equilibrium on $\widetilde{V}$ is dealt with by induction on the number of slices, so it remains to analyse the relaxation to equilibrium on $\overline{T}_0$ under the above constraint. The relaxation time of the auxiliary model is $1/\mu(\mathcal{SG})$ (the number of times one needs to update the first block until the constraint becomes satisfied). Then in order to relax on $\overline{T}_0$ it suffices to do so on the larger region $\overline{V}$. We are done since $\overline{V}$ and $\widetilde{V}$ are already treated by the induction.

In Corollary 4.3.16, repeating the above bisection several times, we manage to reproduce the relaxation on a snail with base B and arbitrary number

r_0 of slices of the first trapezoid. Indeed, starting from the snail with a single slice in T_0 provided above, we double its height $\log(1/q)$ times to reach the desired $r_0 \simeq \delta^{-2k+1}q^{-2w}$. Thus, the Poincaré constant of B is multiplied by $1/\mu(\mathcal{SG})^{\log(1/q)}$ in this process.

Adding all trapezoids of the original snail V Finally, repeating the above steps for each trapezoid, we obtain the desired Poincaré constant for the entire snail, concluding Theorem 4.3.6.

4.3.2 Setup

Given a snail $V = V_L^R(\underline{r})$, we shall work in the associated probability space $\Omega_V = \{0,1\}^{V\cap\mathbb{Z}^2}$ endowed with the probability measure $\mu_V(\cdot\,|\,\mathcal{SG}(V))$ conditioned to the simultaneous occurrence of the following events on Ω_V.

Definition 4.3.2 (Good and super good events)**.**

- Recalling Definition 4.2.3, we define $\mathcal{A}$ as the event that A is infected and $\mathcal{HA}$ as the event that $HA + Lu_0$ is infected.
- Recalling Definitions 4.2.6 and 4.2.8, for each SB_j and $u_i \in \widehat{\mathcal{S}}_0$, let $\mathcal{SB}_{i,j}$ denote the event that $SB_{i,j} = \varnothing$ or $SB_{i,j}$ contains an infected u_i-helping set. Then set $\mathcal{SB}_j = \bigcap_{u_i\in\widehat{\mathcal{S}}_0}\mathcal{SB}_{i,j}$.
- Recalling Definition 4.2.5, for each non-empty $ST_{i,j}^{\pm}$ let $\mathcal{ST}_{i,j}^{\pm}$ be the event that there exist w consecutive infected sites in $ST_{i,j}^{\pm}$.

Using the above events, we then define

$$\mathcal{G}(B^\circ) = \bigcap_{j>0}\mathcal{SB}_j, \quad \mathcal{SG}(B) = \mathcal{A}\cap\mathcal{HA}\cap\mathcal{G}(B^\circ), \quad \mathcal{G}(T_i^{\pm}) = \bigcap_{j>0}\mathcal{ST}_{i,j}^{\pm}.$$

Finally, we set $\mathcal{SG}(V^+) = \mathcal{SG}(B)\cap\bigcap_{i\in[2k]}\mathcal{G}(T_i^+)$ and $\mathcal{SG}(V) = \mathcal{SG}(V^+)\cap\mathcal{SG}(V^-)$, with $\mathcal{SG}(V^-)$ the analogue of $\mathcal{SG}(V^+)$ for the left-snail.

We note that the event $\mathcal{HA}$ is there only to ensure the easy removal of the simplifying Assumption 4.3.1.

Remark 4.3.3. The events above are defined so as to preserve as much as possible the original product structure of μ in the conditional measure $\mu_V(\cdot\,|\,\mathcal{SG}(V))$. In fact,

$$\begin{aligned}
\mu_{V^{\pm}}(\cdot\,|\,\mathcal{SG}(V^{\pm})) &= \mu_B(\cdot\,|\,\mathcal{SG}(B))\otimes\left(\bigotimes_{i\in[2k]}\mu_{T_i^{\pm}}(\cdot\,|\,\mathcal{G}(T_i^{\pm}))\right),\\
\mu_{T_i^{\pm}}(\cdot\,|\,\mathcal{G}(T_i^{\pm})) &= \bigotimes_{j>0}\mu_{ST_{i,j}^{\pm}}(\cdot\,|\,\mathcal{ST}_{i,j}^{\pm}),\\
\mu_B(\cdot\,|\,\mathcal{SG}(B)) &= \mu_{A^{\mathrm{int}}}\otimes\delta_{\omega_{A\cup(HA+Lu_0)}=0}\otimes\mu_{B^\circ}(\cdot\,|\,\mathcal{G}(B^\circ)),\\
\mu_{B^\circ}(\cdot\,|\,\mathcal{G}(B^\circ)) &= \bigotimes_{j>0}\mu_{SB_j}(\cdot\,|\,\mathcal{SB}_j),
\end{aligned}$$

since trapezoids and the base are pairwise disjoint by construction and likewise for the slices of the trapezoids, the slices of the base, the annulus, its interior and the translated half-annulus.

Taking into account this product structure, in the next observations we establish that, as claimed in Section 4.3.1, all $\mathcal{G}$ events we will use are likely and all $\mathcal{SG}$ events have roughly the same probability, $q^{\Theta(Rw)}$.

Observation 4.3.4. Let $R \geqslant w^2 \log(1/q)/q^\alpha$ and $L \leqslant q^{4w}$, let $\underline{r}$ be admissible (see Definition 4.2.4) and $r_{i-1} \geqslant q^{-2w}$ for some $i \in [2k]$. Then $\mu(\mathcal{G}(T_i^\pm)) \geqslant 1 - o(1)$ and $\mu(\mathcal{G}(B^\circ)) \geqslant 1 - o(1)$.

Proof. For the first assertion notice that the condition implies that for all $j > 0$, $ST_{i,j}^\pm$ is either empty or has cardinality at least $\Omega(q^{-2w})$. Then by Remark 4.3.3

$$\mu(\mathcal{G}(T_i^\pm)) = \prod_j \mu(\mathcal{ST}_{i,j}^\pm) \geqslant \left(1 - (1 - q^w)^{\Omega(q^{-2w}/w)}\right)^{O(r_i)} \geqslant 1 - o(1),$$

since $r_i \leqslant L \leqslant q^{4w}$ by Definition 4.2.4 and by assumption.

The second assertion is proved similarly (see e.g. [269, Lemma 6.5]). □

Observation 4.3.5. Let $R \geqslant w^2 \log(1/q)/q^\alpha$, $L \leqslant q^{-4w}$, let $\underline{r}$be admissible such that for some $i \in [2k]$, $r_{i+1} = 0$ and $r_{i-1} \geqslant q^{-2w}$ with the convention $r_{-1} = L$. Then

$$\mu(\mathcal{SG}(V_L^R(\underline{r}))) = q^{\Theta(Rw)}. \tag{4.7}$$

Proof. Using Remark 4.3.3 and Observation 4.3.4, it suffices to note that

$$\mu(\mathcal{A} \cap \mathcal{HA}) \geqslant q^{|A|+|HA|} = q^{\Theta(Rw)}.$$ □

4.3.3 Key step

We are ready to state the main result of this section. In the sequel, for any $\Lambda \subset \mathbb{Z}^2$, any $x \in \Lambda$ and any $\omega_\Lambda \in \Omega_\Lambda$ we shall write $c_x^\Lambda(\omega_\Lambda)$ for the constraint $c_x(\omega)$ computed for the configuration ω equal to ω_Λ in Λ and equal to 1 elsewhere. By construction, $c_x^\Lambda(\omega_\Lambda) \leqslant c_x(\omega')$ for any $\omega' \in \Omega$ such that $\omega'_\Lambda = \omega_\Lambda$ and $c_x^\Lambda \geqslant c_x^{\Lambda'}$ for any $\Lambda' \subset \Lambda$. Then for any snail V (or base) we write γ_V for the smallest constant $\gamma \in [1, \infty]$ such that the Poincaré inequality

$$\operatorname{Var}_V(f \mid \mathcal{SG}(V)) \leqslant \gamma \sum_{x \in V} \mu_V\left(c_x^V \operatorname{Var}_x(f)\right) \tag{4.8}$$

holds for every function $f : \Omega \to \mathbb{R}$.

Theorem 4.3.6. *There exist $w_0, \delta_0 > 0$ not depending on q such that for any $0 < \delta \leqslant \delta_0$ and $w \geqslant w_0$ the following holds for any $R = \Theta(w^2 \log(1/q)/q^\alpha)$. Consider the snail $V = V_L^R(\underline{r})$ for admissible $L, \underline{r}$ such that $r_{2k-1} \geqslant q^{-2w}$ and $L \leqslant q^{-4w}$. Then*

$$\gamma_V \leqslant e^{-O(w^4 \log^3(1/q)/q^\alpha)}. \tag{4.9}$$

Claim 4.3.7. We have $\gamma_V \leqslant 3\max(\gamma_{V^+}, \gamma_{V^-})$.

Proof. Set $\Lambda_1 = V^+ \setminus (B \cup T_0^+)$, $\Lambda_2 = V^- \setminus (B \cup T_0^+)$, $\Lambda_3 = B \cup T_0^+ = B \cup T_0^-$, $\nu_1 = \mu_{\Lambda_1}(\cdot \mid \bigcap_{i=1}^{2k-1} \mathcal{G}(T_i^+))$, similarly for ν_2 and $\nu_3 = \mu_{\Lambda_3}(\cdot \mid \mathcal{SG}(B) \cap \mathcal{G}(T_0^+))$. By Remark 4.3.3 we can apply Lemma 4.2.9 to obtain

$$\begin{aligned}\mathrm{Var}_V(f \mid \mathcal{SG}(V)) \leqslant \gamma_{V^-} &\sum_{x\in V^-} \nu_1\left(\mu_{V^-}\left(c_x^{V^-}\,\mathrm{Var}_x(f)\right)\right)\\ &+ \gamma_{V^+}\sum_{x\in V^+} \nu_2\left(\mu_{V^+}\left(c_x^{V^+}\,\mathrm{Var}_x(f)\right)\right)\\ &\leqslant (1+o(1))(\gamma_{V^-}+\gamma_{V^+})\sum_{x\in V}\mu\left(c_x^V\,\mathrm{Var}_x(f)\right),\end{aligned}$$

where in the last inequality we used Observation 4.3.4 to remove the conditioning of ν_1 and ν_2. □

Therefore, in order to prove (4.9) it suffices to prove the analogous statement with V replaced by $V^\pm$. In the sequel we will concentrate on proving (4.9) for the best constant γ_{V^+} in the Poincaré inequality (4.8) with V replaced by its right-snail V^+. The proof is based on comparison methods between Markov processes and induction over right-snails with different L and $\underline{r}$ as outlined in Section 4.3.1. If we exchange right-snails with left-snails the same proof will then apply to the left-snail V^- as well. Since our arguments no longer require a left-snail, for lightness of notation, we drop the superscript "+" from our notation whenever possible.

The proof of the theorem is decomposed into two quite different steps (see Propositions 4.3.9 and 4.3.12 below). In the first one, labelled the *base case*, we consider a right-snail V with no trapezoids ($\underline{r}$=0). In the second step, labelled *reduction step*, roughly speaking we compare the Poincaré constant γ_V of a generic right-snail V with the same constant computed for its base B.

The conclusion of Theorem 4.3.6 follows at once from (4.7), Proposition 4.3.9 and Proposition 4.3.12. In the sequel fix δ, w, R as in the statement of the theorem and recall that $B = V_L^R(0)$.

Remark 4.3.8. For future purposes (see Chapter 6) it is very important to emphasise that it is only in the first step that we use directly the definition of the event $\mathcal{SG}(B)$ entering in the event $\mathcal{SG}(V)$ (cf. Definition 4.3.2). In the second step the only property of the event $\mathcal{SG}(B)$ that is needed is that it is a decreasing event in Ω_B w.r.t. the partial order $\omega \prec \omega'$ iff $\omega_x \leqslant \omega'_x$ for all $x \in B$.

4.3.4 Base case

Proposition 4.3.9. *For any* $f : \Omega_B \to \mathbb{R}$

$$\mathrm{Var}_B(f \mid \mathcal{SG}(B)) \leqslant q^{-O(Rw\log L)} \sum_{x\in B} \mu_B\left(c_x^B\,\mathrm{Var}_x(f)\right).$$

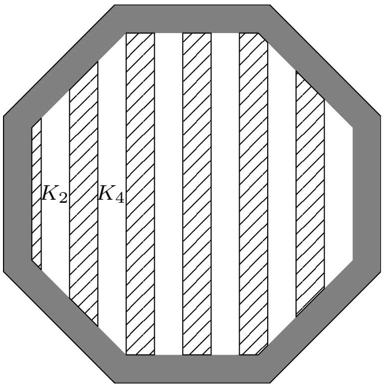

Figure 4.7 – Setting of the proof of Lemma 4.3.10. Every second strip K_i of A^{int} is hatched. The annulus A is shaded.

Proof of Proposition 4.3.9. We first observe that, up to minor modifications, in [269, Proposition 6.6] it was proved that for all $f : \Omega_{B^\circ} \to \mathbb{R}$

$$\mathbb{1}_{\mathcal{A}\cap\mathcal{HA}} \operatorname{Var}_{B^\circ}(f \,|\, \mathcal{G}(B^\circ))) \leqslant q^{-O(Rw\log L)} \mathbb{1}_{\mathcal{A}} \sum_{x\in B^\circ} \mu_{B^\circ}\left(c_x^B \operatorname{Var}_x(f)\right). \quad (4.10)$$

The next step in the proof is an analogous result for A^{int}.

Lemma 4.3.10. *For any $f : \Omega_{A^{int}} \to \mathbb{R}$*

$$\mathbb{1}_{\mathcal{A}} \operatorname{Var}_{A^{int}}(f) \leqslant q^{-O(Rw)} \mathbb{1}_{\mathcal{A}} \sum_{x\in A^{int}} \mu_{A^{int}}\left(c_x^{A\cup A^{int}} \operatorname{Var}_x(f)\right).$$

Proof of Lemma 4.3.10. Let us partition A^{int} into disjoint strips K_i of width w perpendicular to u_0 and number them from left to right (see Figure 4.7).

We can then apply [269, Proposition 3.4] on the generalised FA1f KCM to obtain

$$\operatorname{Var}_{A^{\text{int}}}(f) \leqslant q^{-O(Rw)} \sum_i \mu((\mathbb{1}_{\mathcal{H}_i^+} + \mathbb{1}_{\mathcal{H}_i^-}) \operatorname{Var}_{K_i}(f)),$$

where $\mathcal{H}_i^\pm$ are the events that $K_{i\pm1}$ is fully infected and we use the convention that $\mathcal{H}_i^+$ occurs for the last strip and $\mathcal{H}_i^-$ does for the first one, which corresponds to the boundary condition provided by A. W.l.o.g. it then suffices to bound the generic term $\mu(\mathbb{1}_{\mathcal{H}_i^+} \operatorname{Var}_{K_i}(f))$. But this can be done using Lemma 5.2 of [269] and Lemma 4.2.7(b), which guarantees that if A and K_{i+1} are infected, then K_i can also be infected by the $\mathcal{U}$-bootstrap percolation restricted to $K_i \cup K_{i+1} \cup A$. □

Using Lemma 4.2.9 with parameters $\Lambda_1 = B^\circ$, $\Lambda_2 = A^{\text{int}}$, $\Lambda_3 = \varnothing$, $\nu_1 = \mu_{B^\circ}(\cdot \,|\, \mathcal{G}(B^\circ))$, and $\nu_2 = \mu_{A^{\text{int}}}$, we obtain

$$\begin{aligned}\mathbb{1}_{\mathcal{A}\cap\mathcal{HA}} &\operatorname{Var}_{B^\circ\cup A^{\text{int}}}(f \,|\, \mathcal{G}(B^\circ)) \\ &\leqslant \mathbb{1}_{\mathcal{A}\cap\mathcal{HA}} \left(\mu_{A^{\text{int}}}\left(\operatorname{Var}_{B^\circ}(f \,|\, \mathcal{G}(B^\circ))\right) + \mu_{B^\circ}\left(\operatorname{Var}_{A^{\text{int}}}(f) \,|\, \mathcal{G}(B^\circ)\right)\right).\end{aligned}$$

The first term in the r.h.s. above is bounded by

$$q^{-O(Rw\log L)}\mathbb{1}_{\mathcal{A}\cap\mathcal{H}\mathcal{A}}\sum_{x\in B^\circ}\mu_{B^\circ\cup A^{\mathrm{int}}}(c_x^B\operatorname{Var}_x(f)),$$

using (4.10), while the second one is bounded by

$$q^{-O(Rw\log L)}\mathbb{1}_{\mathcal{A}\cap\mathcal{H}\mathcal{A}}\sum_{x\in A^{\mathrm{int}}}\mu_{B^\circ\cup A^{\mathrm{int}}}\left(c_x^{A\cup A^{\mathrm{int}}}\operatorname{Var}_x(f)\,|\,\mathcal{G}(B^\circ)\right).$$

by Lemma 4.3.10. By Remark 4.3.3 we immediately get that

$$\begin{aligned}\operatorname{Var}_B(f\,|\,\mathcal{SG}(B)) &= \mu_B(\operatorname{Var}_{B^\circ\cup A^{\mathrm{int}}}(f\,|\,\mathcal{G}(B^\circ))\,|\,\mathcal{A}\cap\mathcal{H}\mathcal{A})\\ &\leqslant \mu(\mathcal{SG}(B))^{-1}q^{-O(Rw\log L)}\sum_{x\in B}\mu_B(c_x^B\operatorname{Var}_x(f))\end{aligned}$$

and the proposition follows from Observation 4.3.5. □

4.3.5 Reduction step

Before we can state a relationship between γ_V and γ_B, we need the following notion, which will cover all snail shapes that may arise during the reduction. Recall that δ, w, $V = V_L^{R,+}(\underline{r})$ are fixed as in the statement of Theorem 4.3.6 and that we do not write the $+$ index, though all snails we refer to are right-snails. Also recall that all snails are defined by admissible sequences (see Definition 4.2.4).

Definition 4.3.11. Let C be a constant chosen sufficiently large depending on $\widehat{\mathcal{S}}$, but much smaller than $1/\delta_0$ in Theorem 4.3.6. We say that a snail $\widehat{V} = V_{\widehat{L}}^R(\widehat{\underline{r}})$ is *of type* $i \in [2k]$ if

(a) $\widehat{r}_{i+1} = 0$,

(b) $\widehat{r}_i \leqslant r_i$,

(c) for all $j < i$ it holds that $0 \leqslant r_j - \widehat{r}_j \leqslant C\left(r_i - \widehat{r}_i + \sum_{l=i+1}^{2k-1} r_l\right)$,

(d) $0 \leqslant L - \widehat{L} \leqslant C\left(r_i - \widehat{r}_i + \sum_{l=i+1}^{2k-1} r_l\right)$.

We say that $\widehat{V}$ is *relevant* if there exists $i \in [2k]$ such that $\widehat{V}$ is of type i. In particular, a base $\widehat{B} = V_{\widehat{L}}^R(0)$ is relevant iff $0 \leqslant L - \widehat{L} = O(r_0)$.

In words, $\widehat{V}$ is relevant if all trapezoids except the last one are only slightly shorter than the corresponding ones for V and similarly for the base, while the last trapezoid may be as much shorter as needed. Indeed, observe that by admissibility $\sum_{l=i+1}^{2k-1} r_l < 2r_{i+1}$ for any $i \in [2k]$.

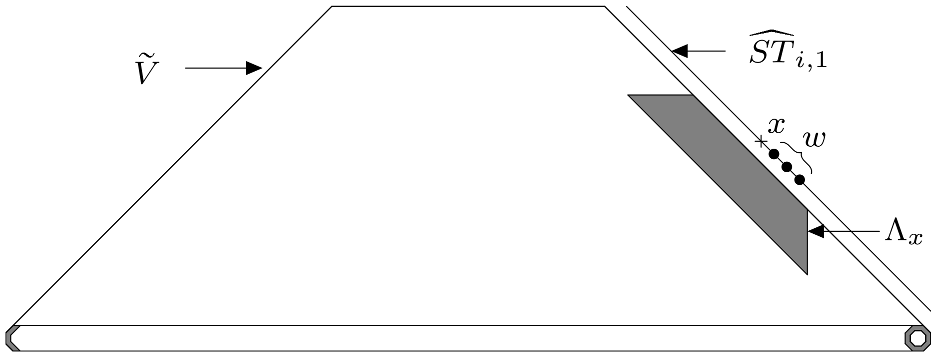

Figure 4.8 – The geometric setting of Lemma 4.3.13. The snail is $\widetilde{V}$, while $\widehat{ST}_{i,1} = \widehat{V}\backslash\widetilde{V}$ is the remaining slice on the top-right. The site x to be updated in (4.13) is marked by a cross. The event $\widetilde{\mathcal{H}}_x$ corresponds to the shaded trapezoid Λ_x being infected and the event $\mathcal{H}_x$ corresponds to the w consecutive sites next to x on one of its sides being infected.

Let us mention that the technical second inequalities in conditions (c) and (d) in the definition above are only needed for the inductive procedure below to always yield relevant snails. We invite the reader to ignore those conditions and admit that all smaller snails arising in our argument have sizes which can be treated by induction.

Proposition 4.3.12. *Let* $\sigma = 1/\min_{\widehat{V}} \mu_{\widehat{V}}(\mathcal{SG}(\widehat{V}))$ *and* $\Gamma = \max_{\widehat{B}} \gamma_{\widehat{B}}$, *where the min and* max *run over relevant snails and relevant bases respectively. Then*

$$\gamma_V \leqslant \left(q^{-w^4}\sigma\right)^{O(\log L)}\Gamma.$$

In the rest of the section we slowly build the proof of this proposition. The first step of reduction consists in removing a trapezoid consisting of a single slice. This is done using Lemma 4.2.12 and may be intuitively understood as an FA1f dynamics of w consecutive infected sites in the slice. Recall ρ_i defined above Definition 4.2.3.

Lemma 4.3.13 (Removing a single slice)**.** *Let* $\widehat{V} = V_{\underline{L}}^{R}(\underline{\widehat{r}})$ *be a snail of type* i *such that* $\widehat{r}_i = \lambda\rho_{k+i}$ *for* $\lambda \in \mathbb{N}$. *In other words, the last non-empty trapezoid,* $\widehat{T}_i$ *of* $\widehat{V}$ *consists of* λ *segments orthogonal to* u_{i+k}. *Then, setting* $\underline{\widetilde{r}} = (\widehat{r}_0, \dots, \widehat{r}_{i-1}, 0, \dots, 0)$, $\widetilde{V} = V_{\underline{L}}^{R}(\underline{\widetilde{r}})$, *we have*

$$\gamma_{\widehat{V}} \leqslant \left(q^{-w^4}\right)^{O(\lambda)} \max\left(\gamma_{\widetilde{V}}, 1/\mu_{\widehat{V}}(\mathcal{SG}(\widehat{V}))\right).$$

Proof of Lemma 4.3.13. By induction on λ it suffices to prove the lemma for $\lambda = 1$, in which case the last trapezoid is simply $\widehat{T}_i = \widehat{ST}_{i,1}$.

We will proceed in two steps. First, we will divide $\widehat{V}$ into $\widetilde{V}$ and $\widehat{ST}_{i,1}$. The $\widetilde{V}$ part is harmless, as it directly relates to $\gamma_{\widetilde{V}}$ appearing in the r.h.s. of the statement of the lemma. In order to reproduce a 'resampling' of $\widehat{ST}_{i,1}$ we will proceed in two steps. First, using the FA1f-like dynamics, Lemma 4.2.12, we will reduce the problem to resampling a single site in $\widehat{ST}_{i,1}$ given that next to it there are w consecutive infections. Then we will use $\widetilde{V}$ to provide additional infections to ensure that $c_x^{\widehat{V}}$ is satisfied and this will yield the x term of the Dirichlet form from (4.4). The lemma is illustrated in Figure 4.8.

Recalling Remark 4.3.3, for any $f : \Omega_{\widehat{V}} \to \mathbb{R}$ Lemma 4.2.9 gives

$$\begin{aligned}\operatorname{Var}_{\widehat{V}}&(f \,|\, \mathcal{SG}(\widehat{V}))\\ &\leqslant \mu_{\widehat{V}}\left(\operatorname{Var}_{\widetilde{V}}(f \,|\, \mathcal{SG}(\widetilde{V})) + \operatorname{Var}_{\widehat{ST}_{i,1}}(f \,|\, \widehat{\mathcal{ST}}_{i,1}) \,|\, \mathcal{SG}(\widehat{V})\right). \quad (4.11)\end{aligned}$$

Since $\mathcal{SG}(\widehat{V}) = \mathcal{SG}(\widetilde{V}) \times \widehat{ST}_{i,1}$, the first term in the r.h.s. above is

$$\begin{aligned}\mu_{\widehat{ST}_{i,1}}\left(\operatorname{Var}_{\widetilde{V}}(f \,|\, \mathcal{SG}(\widetilde{V})) \,|\, \widehat{\mathcal{ST}}_{i,1}\right) &\leqslant \frac{\gamma_{\widetilde{V}}}{\mu_{\widehat{ST}_{i,1}}(\widehat{\mathcal{ST}}_{i,1})} \mu_{\widehat{V}}\left(\sum_{x \in \widetilde{V}} c_x^{\widetilde{V}} \operatorname{Var}_x(f)\right)\\ &= (1 + o(1))\gamma_{\widetilde{V}} \sum_{x \in \widetilde{V}} \mu_{\widehat{V}}(c_x^{\widehat{V}} \operatorname{Var}_x(f))\end{aligned}$$

by the definition (4.8) of $\gamma_{\widetilde{V}}$, Observation 4.3.4, and the fact that $c_x^{\widetilde{V}} \leqslant c_x^{\widehat{V}}$. To bound the second term in (4.11), we use Lemma 4.2.12 for $\mu_{\widehat{ST}_{i,1}}(\cdot \,|\, \widehat{\mathcal{ST}}_{i,1})$ with $\kappa = w$ and constraining event $\mathcal{H} = \{0\} \subset \{0, 1\} = \widehat{S}$, the hypothesis of the lemma following from Observation 4.3.4. This gives

$$\operatorname{Var}_{\widehat{ST}_{i,1}}(f \,|\, \widehat{\mathcal{ST}}_{i,1}) \leqslant q^{-O(w)} \sum_{x \in \widehat{ST}_{i,1}} \mu_{\widehat{ST}_{i,1}}(\mathbb{1}_{\mathcal{H}_x} \operatorname{Var}_x(f)), \quad (4.12)$$

where $\mathcal{H}_x$ is the event that w consecutive sites immediately to the left or to the right of x in $\widehat{ST}_{i,1}$ are infected. Plugging this back in (4.11), we see that we need to bound from above a generic term

$$\mu_{\widehat{V}}\left(\mathbb{1}_{\mathcal{H}_x} \operatorname{Var}_x(f) \,|\, \mathcal{SG}(\widetilde{V})\right), \quad x \in \widehat{ST}_{i,1}. \quad (4.13)$$

At this point we have succeeded in bringing w consecutive infected sites next to the site x, which we want to update. In order to be sure that the constraint $c_x^{\widehat{V}}$ is satisfied, we would like to also bring some infections next to x in $\widetilde{V}$. To do that we first use Lemma 4.2.9 to include $\widetilde{V}$ in the variance, so that we are allowed to 'resample' it and then use the two-block dynamics, Lemma 4.2.10, to indeed obtain the desired infections by resampling $\widetilde{V}$ enough times.

Applying Lemma 4.2.9 with parameters $\Lambda_1 = \widetilde{V}$, $\Lambda_2 = \{x\}$, $\Lambda_3 = \varnothing$, $\nu_1 = \mu_{\widetilde{V}}(\cdot \,|\, \mathcal{SG}(\widetilde{V}))$, and $\nu_2 = \mu_x$, we bound the generic term (4.13) from above by

$$\mu_{\widehat{ST}_{i,1} \setminus \{x\}}\left(\mathbb{1}_{\mathcal{H}_x} \operatorname{Var}_{\widetilde{V} \cup \{x\}}(f \,|\, \mathcal{SG}(\widetilde{V}))\right).$$

We next apply Lemma 4.2.10 to the product space

$$(\Omega_{\widetilde{V}}, \mu_{\widetilde{V}}(\cdot \mid G_{\widetilde{V}})) \otimes (\Omega_{\{x\}}, \mu_x)$$

with constraining event $\widetilde{\mathcal{H}}_x \subset \Omega_{\widetilde{V}}$ that the trapezoid

$$\Lambda_x = x + \left(\overline{\mathbb{H}}_{u_{i+k-1}}(w^2) \cap \mathbb{H}_{u_{i+k}} \cap \overline{\mathbb{H}}_{u_{i+k+1}}(w^2) \cap \mathbb{H}_{u_{i+3k}}(w)\right) \cap \widetilde{V}$$

being infected. It is not hard to check that $\partial_{u_{i+k}}\Lambda_x$ contains x and the w infected sites guaranteed by $\mathcal{H}_x$. In other words, we are in the setting of Figure 4.8. Using $|\Lambda_x \cap \mathbb{Z}^2| = O(w^4)$ and noticing that by the Harris inequality [202] $\mu_{\widetilde{V}}(\widetilde{\mathcal{H}}_x \mid \mathcal{SG}(\widetilde{V})) \geqslant \mu_{\widetilde{V}}(\widetilde{\mathcal{H}}_x)$, Lemma 4.2.10 gives

$$\begin{aligned}\mathbb{1}_{\mathcal{H}_x} \operatorname{Var}_{\widetilde{V}\cup\{x\}}(f \mid \mathcal{SG}(\widetilde{V})) \leqslant q^{-O(w^4)} \mu_{\widetilde{V}\cup\{x\}}\big(&\operatorname{Var}_{\widetilde{V}}(f \mid \mathcal{SG}(\widetilde{V})) \\ &+ \mathbb{1}_{\mathcal{H}_x\cap\widetilde{\mathcal{H}}_x} \operatorname{Var}_x(f) \mid \mathcal{SG}(\widetilde{V})\big).\end{aligned}$$

Finally, since u_{i+k} is an isolated (quasi-)stable direction, it is easily seen (see Figure 4.8 and Lemma 4.2.7) that $\mathbb{1}_{\mathcal{H}_x\cap\widetilde{\mathcal{H}}_x} \leqslant c_x^{\widehat{V}}$. Recalling the definition of the Poincaré constant $\gamma_{\widehat{V}}$ (see (4.8)), we conclude that

$$\begin{aligned}&\mu_{\widehat{ST}_{i,1}\setminus\{x\}}\left(\mathbb{1}_{\mathcal{H}_x} \operatorname{Var}_{\widetilde{V}\cup\{x\}}(f \mid \mathcal{SG}(\widetilde{V}))\right) \\ &\qquad\leqslant q^{-O(w^4)}\left(\gamma_{\widetilde{V}} + 1/\mu(\mathcal{SG}(\widetilde{V}))\right) \sum_{y\in\widetilde{V}\cup\{x\}} \mu_{\widehat{V}}(c_y^{\widehat{V}} \operatorname{Var}_y(f)).\end{aligned}$$

Putting all together, we finally get

$$\begin{aligned}\operatorname{Var}_{\widehat{V}}(f \mid \mathcal{SG}(\widehat{V})) \leqslant |\widehat{ST}_{i,1}| q^{-O(w^4)} \max\left(\gamma_{\widetilde{V}}, 1/\mu_{\widetilde{V}}(\mathcal{SG}(\widetilde{V}))\right)& \\ \times \sum_{x\in\widehat{V}} \mu_{\widehat{V}}\left(c_x^{\widehat{V}} \operatorname{Var}_x(f)\right)&,\end{aligned}$$

where the factor $|\widehat{ST}_{i,1}| = O(\widehat{L}) = O(q^{-4w})$ comes from the fact that each vertex $x \in \widehat{ST}_{i,1}$ produces a term of the form $\sum_{y\in\widetilde{V}} \mu_{\widehat{V}}(c_y^{\widehat{V}} \operatorname{Var}_y(f))$. □

The remaining induction step allows us to reduce the size of the last non-empty trapezoid $\widehat{T}_i$ twice. The proof is illustrated in Figure 4.9.

Lemma 4.3.14 (Bisection of a trapezoid)**.** *Let $\widehat{V} = V_{\widehat{L}}^{R}(\underline{\widehat{r}})$ be a snail of type i such that $\widehat{r}_i$ is larger than some sufficiently large constant. Let $\lambda = \min\{\ell > 0, \ell u_{i+1} \in \mathbb{Z}^2\} = O(1)$ and let $x = u_{i+1}\lambda\lfloor \widehat{r}_i/(2\lambda\langle u_{i+k}, u_{i+1}\rangle)\rfloor$. With this choice $\langle u_{i+k}, x\rangle \simeq \widehat{r}_i/2$. In other words, x is the vector by which the*

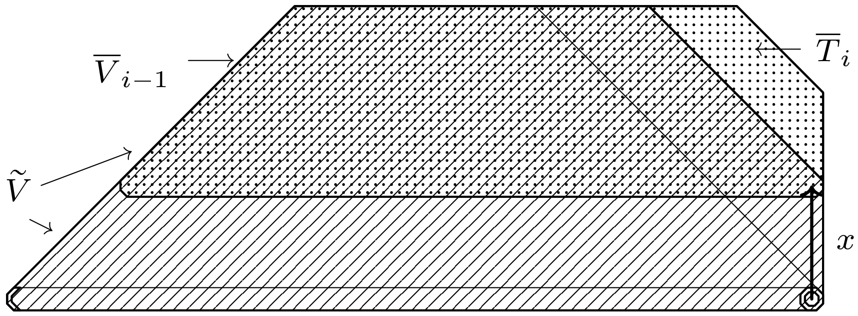

Figure 4.9 – The geometric setting of Lemma 4.3.14. The snail $\widetilde{V}$ is hatched, $\overline{V}$ is dotted and their union is the original snail $\widetilde{V}$. The dotted-hatched region is $\overline{V}_{i-1}$, while the dotted trapezoid is $\overline{T}_i$.

ring should be translated so that half of the last trapezoid, $\widehat{T}_i$, remains above it (see Figure 4.9). Then we set:

$$\begin{aligned}
\widetilde{\underline{r}} &= (\widehat{r}_0, \dots, \widehat{r}_{i-1}, \langle u_{i+k}, x\rangle, 0, \dots, 0),\\
\overline{\underline{r}} &= (\widehat{r}_0 - \langle u_k, x\rangle, \dots, \widehat{r}_i - \langle u_{k+i}, x\rangle, 0, \dots, 0),\\
\overline{L} &= \min\left(\widehat{L}, \widehat{L} - \frac{\langle u_{k-1}, x\rangle}{\langle u_{k-1}, u_0\rangle}\right),\\
\widetilde{V} &= V^R_{\widehat{L}}(\widetilde{\underline{r}}),\\
\overline{V} &= x + V^R_{\overline{L}}(\overline{\underline{r}}).
\end{aligned}$$

In words, $\widetilde{V}$ is $\widehat{V}$ with half of $\widehat{T}_i$ removed, while $\overline{V}$ is the snail such that its last trapezoid $\overline{T}_i$ is exactly that missing half, but with length eventually shortened, so that $\overline{V}$ fits inside $\widehat{V}$ (see Figure 4.9). With these notations,

$$\gamma_{\widehat{V}} \leqslant \gamma_{\widetilde{V}}/\mu(G_{\overline{V}}) + \gamma_{\overline{V}}/\mu(G_{\widetilde{V}})$$

and $\widetilde{V}$ and $\overline{V}$ are snails of type i.

Proof of Lemma 4.3.14. The proof goes as follows. In Claim 4.3.15 we show that the two polygons $\widetilde{V}$ and $\overline{V}$ are indeed snails (defined by admissible sequences) of type i and that they do correspond to their informal definitions in the statement of the lemma. Though technical, this claim hides no subtlety and we invite the reader to skip it. Then we apply Lemma 4.2.10 to reduce the problem of relaxation on $\widehat{V}$ to the one on $\widetilde{V}$ and on $\overline{V}$ which yields the desired result. The event $\mathcal{SG}(\widetilde{V})$ is implied by $\mathcal{SG}(\widehat{V})$ by construction, but the second block, $\overline{V}$, of the dynamics corresponding to Lemma 4.2.10, is updated only when the part of $\mathcal{SG}(\overline{V})$ witnessed in $\widetilde{V}$ occurs.

We begin with some geometric observations following directly from Definitions 4.2.4 and 4.2.5.

Claim 4.3.15. $\widetilde{V}$ and $\overline{V}$ are snails of type i. Furthermore, we have $\widetilde{V} \cup \overline{V} = \widehat{V}$ and $\widehat{V} \backslash \widetilde{V} = \overline{T}_i$ (the last trapezoid of $\overline{V}$).

Proof of the claim. The statement that $\widetilde{V}$ is a snail of type i follows from the definition of $\underline{\widetilde{r}}$ and the same fact for $\widehat{V}$, since $\widetilde{r}_i \leqslant \widehat{r}_i$.

Turning to $\overline{V}$, notice that $\langle u_j, x\rangle \geqslant 0$ for all $i \in [2k]$ and $j \in [k, k+i]$ with equality iff $i = 2k-1$ and $j = k$. Thus, for all $j \in [2k]$ we have $\widehat{r}_j \geqslant \overline{r}_j$ and clearly $\widehat{L} \geqslant \overline{L}$. Thus, recalling the definition of $\overline{V}$ and that $\widehat{V}$ is of type i, conditions (a) and (b) and the left inequalities in (c) and (d) of Definition 4.3.11 are satisfied. Moreover, for $0 \leqslant j < i$ we have

$$\widehat{r}_j - \overline{r}_j = \langle u_{k+j}, x\rangle = \frac{\langle u_{i+1}, u_{k+j}\rangle}{\langle u_{i+1}, u_{k+i}\rangle}(\widehat{r}_i - \overline{r}_i) \leqslant C(\widehat{r}_i - \overline{r}_i), \tag{4.14}$$

so the right inequality of (c) for $\overline{V}$ follows from the one for $\widehat{V}$. Similarly,

$$\widehat{L} - \overline{L} \leqslant \frac{|\langle u_{k-1}, x\rangle|}{\langle u_{k-1}, u_0\rangle} \leqslant C(\widehat{r}_i - \overline{r}_i), \tag{4.15}$$

gives that (d) of Definition 4.3.11 holds for $\overline{V}$.

We next prove that $\overline{V}$ is a snail (with admissible $\overline{L}$ and $\underline{\overline{r}}$). Recalling from Definition 4.2.4 that we need to prove that

(i) $\overline{r}_j \geqslant 0$ for all $j \in [2k]$,

(ii) $\overline{r}_j \leqslant \delta \overline{r}_{j-1}$ for all $1 \leqslant j < 2k$,

(iii) $\overline{r}_0 \leqslant \delta \overline{L}$, and

(iv) $(\widehat{L} - \overline{L})\langle u_0, u_{k-1}\rangle / \rho_{k-1} \in \mathbb{N}$.

To check (i), observe that $\overline{r}_j = \widehat{r}_j - O(\widehat{r}_i) > 0$ for $j < i$ by admissibility of $\underline{\widehat{r}}$ and $\overline{r}_i \simeq \widetilde{r}_i/2 > 0$. By admissibility of $\widehat{V}$, $\widehat{r}_j - \overline{r}_j = \langle u_{k+j}, x\rangle = \Theta(\widehat{r}_i)$ and $C < 1/\delta$ we get (ii). For the last two properties we consider two cases.

First assume that $i \in \{2k-2, 2k-1\}$ (i.e. x corresponds to a horizontal translation to the right—in direction u_{2k}). It is easy to check from the definition of $\overline{L}$ that in this case $\overline{L} = \widehat{L}$, so that (iv) is trivial and (iii) follows from $\overline{r}_0 \leqslant \widehat{r}_0 \leqslant \delta \widehat{L}$. This concludes the proof that $\overline{V}$ is a snail of type i in this case.

Assume that, on the contrary, $i < 2k-2$, so that the $\widehat{L} - \overline{L} = \frac{\langle u_{k-1}, x\rangle}{\langle u_{k-1}, u_0\rangle} = \Theta(\widehat{r}_i)$. Then (iii) follows from the fact that $\widehat{r}_0 - \overline{r}_0 = \Theta(\widehat{r}_i)$ as above. For (iv) simply observe that

$$(\widehat{L} - \overline{L})\langle u_0, u_{k-1}\rangle = \langle u_{k-1}, x\rangle \in \rho_{k-1}\mathbb{N},$$

since $x \in \mathbb{Z}^2$ by the definition of x and λ. This concludes the proof that $\overline{V}$ is a snail of type i.

By Definition 4.2.5 it is clear that

$$\widehat{V}\backslash\overline{V} = \big(\mathbb{H}_{u_{k+i}}(R_{k+i} + \widehat{r}_i)\backslash\mathbb{H}_{u_{k+i}}(R_{k+i} + \langle u_{k+i}, x\rangle)\big) \\ \cap \overline{\mathbb{H}}_{u_{k+i-1}}(R_{k+i-1} + \widehat{r}_{i-1}) \cap \overline{\mathbb{H}}_{u_{k+i+1}}(R_{k+i+1}).$$

It also follows from Definition 4.2.5 that the above trapezoid $\widehat{V}\backslash\overline{V}$ is also equal to $\overline{T}_i$ as claimed. Finally, we have that $\overline{V} \subset \widehat{V}$ using Definition 4.2.4, which completes the proof of the claim. □

Let now

$$\overline{V}_{i-1} = V^R_L(\overline{r}_0, \ldots, \overline{r}_{i-1}, 0, \ldots, 0) = \overline{V}\backslash\overline{T}_i = \overline{V} \cap \widetilde{V}.$$

By Claim 4.3.15 and Remark 4.3.3 we have

$$(\Omega_V, \mu_{\widehat{V}}(\cdot \,|\, G_{\widehat{V}})) = (\Omega_{\widetilde{V}}, \mu_{\widetilde{V}}(\cdot \,|\, \mathcal{SG}(\widetilde{V}))) \otimes (\Omega_{\overline{T}_i}, \mu_{\overline{T}_i}(\cdot \,|\, \mathcal{G}(\overline{T}_i))) \tag{4.16}$$

and we can apply Lemma 4.2.10 with the facilitating event

$$\mathcal{SG}(\overline{V}_{i-1}) = \mathcal{SG}(\overline{B}) \cap \bigcap_{j<i} \mathcal{G}(\overline{T}_j) \subset \Omega_{\widetilde{V}},$$

where $\overline{B}$ and $\overline{T}_j$ are the base and trapezoids of $\overline{V}$. We get

$$\begin{aligned} \operatorname{Var}_{\widehat{V}}(f \,|\, \mathcal{SG}(\widehat{V})) \leqslant \mu(\mathcal{SG}(\overline{V}_{i-1}))^{-1}\mu_{\widehat{V}}\big(&\operatorname{Var}_{\widetilde{V}}(f \,|\, \mathcal{SG}(\widetilde{V})) \\ &+ \mathbb{1}_{\mathcal{SG}(\overline{V}_{i-1})} \operatorname{Var}_{\overline{T}_i}(f \,|\, \mathcal{G}(\overline{T}_i)) \,|\, \mathcal{SG}(\widehat{V})\big), \end{aligned} \tag{4.17}$$

where we used that $\mu_{\widetilde{V}}(\mathcal{SG}(\overline{V}_{i-1}) \,|\, \mathcal{SG}(\widetilde{V})) \geqslant \mu(\mathcal{SG}(\overline{V}_{i-1}))$ by the Harris inequality. Using the definition of the Poincaré constant $\gamma_{\widetilde{V}}$, the fact that $c^{\widetilde{V}}_x \leqslant c^V_x$ together with $\mu(\mathcal{SG}(\overline{V}_{i-1})) \geqslant \mu(\mathcal{SG}(\overline{V}))$ the first term is bounded from above by

$$\frac{\gamma_{\widetilde{V}}}{\mu(\mathcal{SG}(\overline{V}))}\mu_{\widehat{V}}\left(\sum_{x\in\widetilde{V}} c^{\widehat{V}}_x \operatorname{Var}_x(f)\right). \tag{4.18}$$

The term $\mu(\mathcal{SG}(\overline{V}_{i-1}))^{-1}\mu_{\widehat{V}}\left(\mathbb{1}_{\mathcal{SG}(\overline{V}_{i-1})} \operatorname{Var}_{\overline{T}_i}(f \,|\, \mathcal{G}(\overline{T}_i)) \,|\, \mathcal{SG}(\widehat{V})\right)$ from the r.h.s. of (4.17) can be bounded from above by

$$\begin{aligned} \mu(\mathcal{SG}(\widetilde{V}))^{-1}\mu_{\widetilde{V}}&\left(\mu_{\overline{V}_{i-1}}(\operatorname{Var}_{\overline{T}_i}(f \,|\, \mathcal{G}(\overline{T}_i)) \,|\, \mathcal{SG}(\overline{V}_{i-1}))\right) \\ &\leqslant \mu(\mathcal{SG}(\widetilde{V}))^{-1}\mu_{\widetilde{V}}\left(\operatorname{Var}_{\overline{V}}(f \,|\, \mathcal{SG}(\overline{V}))\right) \\ &\leqslant \frac{\gamma_{\overline{V}}}{\mu(\mathcal{G}(\widehat{V}))}\mu_{\widetilde{V}}\left(\sum_{x\in\overline{V}} c^{\widehat{V}}_x \operatorname{Var}_x(f)\right), \end{aligned} \tag{4.19}$$

using (4.16), Lemma 4.2.9, the definition of $\gamma_{\overline{V}}$ and the fact that $c^{\overline{V}}_x \leqslant c^{\widehat{V}}_x$.

If we now combine (4.17), (4.18) and (4.19) we get the statement of the lemma. □

We can now assemble our main induction step from Lemmas 4.3.13 and 4.3.14. Namely, we repeatedly use Lemma 4.3.14 until the last trapezoid is reduced to a bounded number of lines and then apply Lemma 4.3.13 to remove them as well.

Corollary 4.3.16 (Removing a trapezoid). *Let* $\sigma_i = 1/\min_{V_i}\mu(\mathcal{SG}(V^i))$ *with* min *running over all snails of type* i*. Let* $\Gamma_i = \max_{V_i'}\gamma_{V_i'}$*, where the* max *runs over all snails of type* i *with* $r_i = 0$*. Let* $\widehat{V} = V^R_{\widehat{L}}(\widehat{\underline{r}})$ *be a snail of type* i*. Then*

$$\gamma_{\widehat{V}} \leqslant q^{-O(w^4)}\sigma_i^{O(\max(1,\log\widehat{r}_i))}\Gamma_i.$$

Proof of Corollary 4.3.16. Let c be a sufficiently large constant. We prove by induction on $\widehat{r}_i$ that

$$\gamma_{\widehat{V}} \leqslant q^{-cw^4}\sigma_i^{c\max(1,\log\widehat{r}_i)}\Gamma_i.$$

The base of the induction, $\widehat{r}_i \leqslant \sqrt{c}$, follows from Lemma 4.3.13, since $\gamma_{\widetilde{V}} \geqslant 1$ by definition. Assume that $\widehat{r}_i > \sqrt{c}$. Then Lemma 4.3.14 and the induction hypothesis applied to both $\widetilde{V}$ and $\overline{V}$ from that lemma give

$$\gamma_{\widehat{V}} \leqslant \sigma_i q^{-cw^4}\sigma_i^{c\log(2\widehat{r}_i/3)}\Gamma_i \leqslant q^{-cw^4}\sigma_i^{c\log\widehat{r}_i}\Gamma_i,$$

since both $\widetilde{r}_i$ and $\overline{r}_i$ in Lemma 4.3.14 are smaller than $2\widehat{r}_i/3$. This completes the proof of the induction step and the corollary. □

We are now ready to conclude the proof of Proposition 4.3.12 and of Theorem 4.3.6.

Proof of Proposition 4.3.12. Applying Corollary 4.3.16 successively to each non-zero coordinate of $\underline{r}$, we obtain

$$\gamma_V \leqslant \left(q^{-O(w^4)}\sigma\right)^{O(\log L)}\Gamma$$

with the notation of the statement of Proposition 4.3.12. □

Proof of Theorem 4.3.6. Combining Propositions 4.3.9 and 4.3.12 we get

$$\gamma_V \leqslant \left(q^{-w^4}\sigma q^{-Rw}\right)^{O(\log L)} \leqslant e^{-O(w^4\log^3(1/q)/q^\alpha)},$$

where the last equality follows from Observation 4.3.5. □

4.4 Proof of Theorem 4.0.1

Recall that w is a large constant much bigger than the constants in any $O(\cdot)$ notation. Let $t_* = \frac{1}{w}e^{w^5 \log^3(1/q)/q^\alpha}$ and $T = e^{1/q^{3\alpha}}$. Then we have

$$\begin{aligned}
\mathbb{E}_\mu(\tau_0) &= \int_0^{+\infty} \mathbb{P}_\mu(\tau_0 > s) \\
&= \int_0^{t_*} \mathbb{P}_\mu(\tau_0 > s) + \int_{t_*}^{T} \mathbb{P}_\mu(\tau_0 > s) + \int_T^{+\infty} \mathbb{P}_\mu(\tau_0 > s) \\
&\leqslant t_* + T\mathbb{P}_\mu(\tau_0 > t_*) + \int_T^{+\infty} \mathbb{P}_\mu(\tau_0 > s).
\end{aligned}$$

The term t_* has exactly the form required in Theorem 4.0.1. The last term in the r.h.s. above tends to zero as $q \to 0$. Indeed, using [269, Theorem 2] we have that $\forall s > 0, \mathbb{P}_\mu(\tau_0 > s) \leqslant e^{-s\lambda_0}$ with $\lambda_0 \geqslant e^{-\Omega((\log q)^4/q^{2\alpha})}$ and therefore

$$\int_T^{+\infty} \mathbb{P}_\mu(\tau_0 > s) \leqslant \lambda_0^{-1} e^{-T\lambda_0} \to 0 \quad \text{as } q \to 0.$$

In conclusion, the proof of the upper bound in Theorem 4.0.1 boils down to proving

$$\lim_{q\to 0} T\mathbb{P}_\mu(\tau_0 > t_*) = 0. \tag{4.20}$$

That requires a sequence of simple steps ((a)-(d) below) and a more involved part ((e) below). Before turning to the details of the proof of Theorem 4.0.1, let us sketch our approach.

4.4.1 Roadmap

(a) In order to prove that w.h.p. $\tau_0 \leqslant t_*$, it suffices to prove the result for the (stationary) $\mathcal{U}$-KCM process on to the torus Λ and with side $K = 2e^{w^5 \log^3(1/q)/q^\alpha}$ (see (4.21)).

(b) Let $L = \Theta(\lambda)/q^{3w}$ for a large positive constant $\lambda = \lambda(\mathcal{U}, \delta)$, let $R = w^2 \log(1/q)/q^\alpha$, and recall the good and super good events described in Section 4.3.1 and Definition 4.3.2. Given a snail $V = V_L^R(\underline{r}) = B \cup \bigcup_{i\in[2k]} T_i^\pm \subset \Lambda$ (recall Definitions 4.2.4 and 4.2.5) with base B and trapezoids $T_i^\pm$, we will construct a new event $\mathcal{E} \subset \Omega_{\Lambda\cap\mathbb{Z}^2}$ which will guarantee that (in particular) the following occurs.

   (i) For any (translate of) $V \subset \Lambda$ as above, the good events $\mathcal{G}(T_i^\pm)$ occur for all $i \in [2k]$.

   (ii) In every strip of Λ parallel to u_0 and of width $2R$ there exists a translate of the base B for which the super good event holds.

(c) We will prove that $\mu(\mathcal{E}) \geqslant 1 - e^{-1/q^w}$, which will allow us to conclude that it is sufficient to analyse the infection time of the origin of the stationary $\mathcal{U}$-KCM in Λ restricted to $\mathcal{E}$ (see (4.25)).

(d) For the latter process we will follow the standard "variational" approach (see [25, Theorem 2] and also [269, Section 2.2]) and get that

$$T\mathbb{P}_\mu(\tau_0 \geqslant t_*) \leqslant Te^{-t_*\lambda_\mathcal{F}}(1 + o(1)).$$

Here $\lambda_\mathcal{F}$ is related to the Dirichlet problem for the $\mathcal{U}$-KCM on the torus and restricted to $\mathcal{E}$ with boundary condition $f\big|_{\{\omega\in\mathcal{E}:\omega_0=0\}} = 0$. In particular (see (4.27))

$$\lambda_\mathcal{F} \geqslant \inf_f q\frac{\mathcal{D}^{\mathrm{per}}_\Lambda(f)}{\mathrm{Var}_\Lambda(f\,|\,\mathcal{E})},$$

where $\mathcal{D}^{\mathrm{per}}_\Lambda(f)$ is the Dirichlet form of the $\mathcal{U}$-KCM on the torus Λ and the supremum is taken over all $f : \mathcal{E} \to \mathbb{R}$.

(e) The last and most important step will be to prove that

$$\mathrm{Var}_\Lambda(f\,|\,\mathcal{E}) \leqslant e^{O(w^4\log^3(1/q))/q^\alpha}\mathcal{D}^{\mathrm{per}}_\Lambda(f),$$

implying that $t_*\lambda_\mathcal{F}$ diverges as $q \to 0$ rapidly enough. The high-level intuition behind the above Poincaré inequality is as follows. A super good base (i.e. a base B for which the super good $\mathcal{SG}(B)$ event holds), whose presence is guaranteed by (b.ii), will be able to move in Λ using an FA1f-like dynamics as in Lemma 4.2.12 with $\widehat{\nu}(\mathcal{H})$ given by

$$e^{\Theta(w^3\log^2(1/q))/q^\alpha}.$$

Indeed, we will reproduce each step of that dynamics with a resampling of an appropriate super good translate of the snail V, since (b.i) guarantees that the super good base does extend to a super good translate of the snail V. Indeed, the snail (see Figure 4.5) does extend on both sides of the base for a distance $\Theta(r_{2k-1})$, so taking r_{2k-1} of order L, one can induce a change on both sides of the base by resampling the configuration inside the snail. Thanks to Theorem 4.3.6, the relaxation time of the super good snail is $e^{O(w^4\log^3(1/q))/q^\alpha}$. The conclusion of Theorem 4.0.1 then follows rather naturally.

4.4.2 Proof

Let $K = 2\exp(w^5\log^3(1/q)/q^\alpha)$ and let $\Lambda = \mathbb{R}^2/(Ku_0\mathbb{Z} + Ku_k\mathbb{Z})$ be the torus in $\mathbb{R}^2$ of side K directed by u_0, which we think of as centred at 0. Further set

$$R = w^2\log(1/q)/q^\alpha, \qquad W = 1/q^{3w}, \qquad M = K/(2R_0 + W),$$

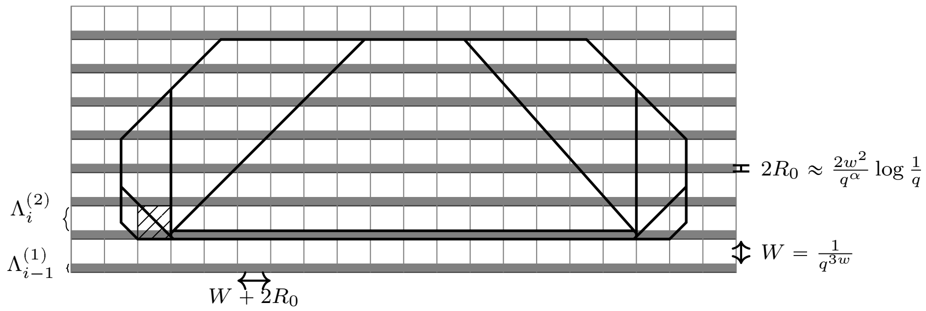

Figure 4.10 – The partition of Λ into strips $\Lambda_i = \Lambda_i^{(1)} \cup \Lambda_i^{(2)}, i \in [M]$. The hatched region represents a square $Q_{i,j}$, which we would like to resample. The thick polygon is the snail V with its trapezoids. Note that its base B does not intersect $Q_{i,j}$ and (almost) spans the squares $Q_{i,j+1}, \dots, Q_{i,j+\lambda}$.

recalling the notation $R_0 = \rho_0 \lfloor R/\rho_0 \rfloor$ from Definition 4.2.3. For simplicity we assume that $u_0(2R_0 + W) \in \mathbb{Z}^2$ and that M is an even integer (W and K can be modified by $O(1)$ and $O(1/q^{3w})$ respectively, so that these both hold).

We partition Λ into alternating strips $\Lambda_i^{(1)}, \Lambda_i^{(2)}, i \in [M]$, of length K and parallel to u_0 (see Figure 4.10). The strips $\Lambda_i^{(1)}$ have width $2R_0$ while the strips $\Lambda_i^{(2)}$ have width W. We write $\Lambda_i = \Lambda_i^{(1)} \cup \Lambda_i^{(2)}$ and we think of the thin strip $\Lambda_i^{(1)}$ as being just below the thick one $\Lambda_i^{(2)}$, when u_0 points left. In turn, we partition Λ_i into consecutive squares $Q_{i,j}, j \in [M]$, of side length equal to $2R_0 + W$ and sides parallel to u_0 and u_k and we write $Q_{i,j}^{(a)} = Q_{i,j} \cap \Lambda_i^{(a)}, a \in \{1, 2\}$.

Remark 4.4.1. Recalling Definition 4.2.3, the width of the thin strips is chosen so that an annulus A of radius R would fit tightly inside.

We are now ready to detail the steps (a)-(e) sketched in the roadmap above.

Step (a)

Notice that $t_* = K/(2w)$ and let τ_0, τ_0^Λ denote the infection times of the origin for the $\mathcal{U}$-KCM process on $\mathbb{Z}^2$ and for the $\mathcal{U}$-KCM process on the discrete torus $\Lambda \cap \mathbb{Z}^2$ respectively. Using the fact that the jump rates of the KCM are bounded, a standard argument of finite speed of information propagation (see e.g. [258]) implies that

$$\mathbb{P}_\mu(\tau_0 \geqslant t_*) \leqslant \mathbb{P}_{\mu_\Lambda}\left(\tau_0^\Lambda \geqslant t_*\right) + e^{-\Omega(K)} \quad \text{as } q \to 0. \tag{4.21}$$

Step (b)

Given a small positive constant $\varepsilon = \varepsilon(\mathcal{U}) = \Omega(1)$ and a large one $\lambda = \lambda(\mathcal{U}, \delta)$ to be specified later (recall the constant δ from Definition 4.2.4) let

(i) $\mathcal{SG}\left(Q^{(1)}_{i,j}\right)$ be the event that the rightmost and leftmost annuli A in $Q^{(1)}_{i,j}$ are infected and any segment $I \subset Q^{(1)}_{i,j}$ intersecting $\mathbb{Z}^2$, of length εR and orthogonal to some $u_i \in \widehat{\mathcal{S}}_0$ contains an infected u_i-helping set in $Q^{(1)}_{i,j}$;

(ii) $\mathcal{G}(Q_{i,j})$ be the event that any segment $I \subset Q_{i,j}$ intersecting $\mathbb{Z}^2$, of length εW and orthogonal to some $u \in \widehat{\mathcal{S}}$ contains w infected consecutive sites;

(iii) $\mathcal{E}_i$ be the event that for all the squares $Q_{i,j} \subset \Lambda_i$ the event $\mathcal{G}(Q_{i,j})$ holds and moreover there exists $j \in [M]$ such that $\bigcap_{j'=j+1}^{j+\lambda} \mathcal{SG}\left(Q^{(1)}_{i,j'}\right)$ also holds;

(iv) $\mathcal{E} = \bigcap_{i \in [M]} \mathcal{E}_i$.

Remark 4.4.2. Similarly to Definition 4.3.2, (i) needs to be modified slightly if Assumption 4.3.1 is not satisfied, but we keep working under that assumption.

Step (c)

With our choice of K, R, W, as in Observation 4.3.5, we get $\mu\left(\mathcal{SG}\left(Q^{(1)}_{i,j}\right)\right) = q^{O(Rw)}$. Moreover, using the Harris inequality

$$\mu\left(\bigcap_{j \in [\lambda]} \mathcal{SG}\left(Q^{(1)}_{i,j}\right)\right) \geqslant q^{O(\lambda R w)}. \tag{4.22}$$

Also,

$$\mu\left(\bigcap_{j \in [M]} \mathcal{G}(Q_{i,j})\right) \geqslant 1 - O\left(MW^2\right) e^{-q^w \varepsilon W / w^2} \geqslant 1 - e^{-q^{-2w+o(1)}}.$$

In conclusion,

$$1 - \mu(\mathcal{E}_i) \leqslant e^{-q^{-2w+o(1)}} + \left(1 - q^{O(\lambda R w)}\right)^{\lfloor M/\lambda \rfloor} \leqslant e^{-q^{-2w+o(1)}} \tag{4.23}$$

and $\mu(\mathcal{E}) \geqslant 1 - M(1 - \mu(\mathcal{E}_i)) \geqslant 1 - e^{-q^{-2w+o(1)}}$. Therefore, writing $\tau^{\Lambda}_{\mathcal{E}^c}$ for the hitting time of $\mathcal{E}^c$ for the $\mathcal{U}$-KCM process in Λ and recalling that $t_* = K/(2w)$, we obtain

$$\mathbb{P}_{\mu_\Lambda}(\tau^{\Lambda}_{\mathcal{E}^c} \leqslant t_*) \leqslant O(K^2 t_*) \mu(\mathcal{E}^c) + e^{-\Omega(K^2 t_*)} \leqslant e^{-q^{-2w+o(1)}}. \tag{4.24}$$

In the second inequality above we used a simple union bound over the updates for the $\mathcal{U}$-KCM in Λ together with the fact that the law of the $\mathcal{U}$-KCM process in Λ started from μ_Λ is equal to μ_Λ at any given time and a simple large deviations result on the number of updates.

Thus, if $\mathcal{F} = \{\omega : \omega_0 = 0\} \cup \mathcal{E}^c$ then (4.21) together with (4.24) imply that

$$\begin{aligned}\mathbb{P}_\mu(\tau_0 \geqslant t_*) &\leqslant \mathbb{P}_\mu(\tau_0^\Lambda \geqslant t_*) + e^{-\Omega(K)}\\ &\leqslant \mathbb{P}_{\mu_\Lambda}(\tau_{\mathcal{F}}^\Lambda \geqslant t_*) + \mathbb{P}_{\mu_\Lambda}(\tau_{\mathcal{E}^c}^\Lambda \leqslant t_*) + e^{-\Omega(K)}\\ &\leqslant \mathbb{P}_{\mu_\Lambda}(\tau_{\mathcal{F}}^\Lambda \geqslant t_*) + e^{-q^{-2w+o(1)}}.\end{aligned} \tag{4.25}$$

Step (d)

As in [25, Theorem 2],

$$\mathbb{P}_{\mu_\Lambda}(\tau_{\mathcal{F}}^\Lambda \geqslant t_*) \leqslant e^{-\lambda_{\mathcal{F}} t_*}, \tag{4.26}$$

with

$$\lambda_{\mathcal{F}} = \inf\left\{\frac{\mathcal{D}_\Lambda^{\mathrm{per}}(f)}{\mu_\Lambda(f^2)}, f|_{\mathcal{F}} = 0\right\},$$

where $\mathcal{D}_\Lambda^{\mathrm{per}}(f)$ denotes the Dirichlet form of the $\mathcal{U}$-KCM process on the torus Λ (see (4.4)). Observe now that for any $f : \Omega_\Lambda \to \mathbb{R}$ such that $f|_{\mathcal{F}} = 0$

$$\begin{aligned}\mathrm{Var}_\Lambda(f \,|\, \mathcal{E}) &= \frac{1}{2}\sum_{\omega}\sum_{\omega'} \mu_\Lambda(\omega \,|\, \mathcal{E})\mu_\Lambda(\omega' \,|\, \mathcal{E})(f(\omega) - f(\omega'))^2\\ &\geqslant \mu_\Lambda(\omega_0 = 0 \,|\, \mathcal{E})\mu_\Lambda(f^2 \,|\, \mathcal{E}) \geqslant q\mu_\Lambda(f^2),\end{aligned}$$

where for the last inequality we used the Harris inequality ($\{\omega : \omega_0 = 0\}$ and $\mathcal{E}$ are both decreasing events) and the fact that $f^2 \mathbb{1}_{\mathcal{E}} = f^2$. Hence,

$$\lambda_{\mathcal{F}} \geqslant q \inf_f \frac{\mathcal{D}_\Lambda^{\mathrm{per}}(f)}{\mathrm{Var}_\Lambda(f \,|\, \mathcal{E})}. \tag{4.27}$$

Notice the absence of the conditioning event $\mathcal{E}$ in the Dirichlet form $\mathcal{D}_\Lambda^{\mathrm{per}}(f)$.

Step (e)

Our main result on the above variational problem is as follows.

Theorem 4.4.3. *For all $w > 0$ large enough, all $\varepsilon > 0$ small enough and all $f : \Omega_\Lambda \to \mathbb{R}$*

$$\mathrm{Var}_\Lambda(f \,|\, \mathcal{E}) \leqslant e^{O(w^4(\log(1/q))^3)/q^\alpha} \mathcal{D}_\Lambda^{\mathrm{per}}(f), \tag{4.28}$$

i.e.

$$\lambda_{\mathcal{F}} \geqslant e^{-O(w^4(\log(1/q))^3)/q^\alpha}. \tag{4.29}$$

Before proving this theorem, let us first complete the proof of (4.20). Using $t_* = w^{-1}\exp(w^5(\log(1/q))^3/q^\alpha)$ and (4.29) we get that for any w large enough

$$t_*\lambda_{\mathcal{F}} \geqslant 1/q^w$$

which, together with (4.25) and (4.26) and the choice of $T = e^{1/q^{3\alpha}}$, gives

$$T\mathbb{P}_\mu(\tau_0 \geqslant t_*) \leqslant T\left(e^{-\lambda_{\mathcal{F}}t_*} + e^{-q^{-2w+o(1)}}\right) \to 0. \tag{4.30}$$

This proves (4.20) and therefore Theorem 4.0.1 modulo Theorem 4.4.3.

Proof of Theorem 4.4.3. The two main ingredients of the proof are Lemma 4.2.12 and Theorem 4.3.6. The definition of the event $\mathcal{E} = \bigcap_i \mathcal{E}_i$ and the fact that the strips Λ_i are disjoint imply that $\mu_\Lambda(\cdot \,|\, \mathcal{E}) = \bigotimes_i \mu_{\Lambda_i}(\cdot \,|\, \mathcal{E}_i)$. In turn, Lemma 4.2.9 gives

$$\mathrm{Var}_\Lambda(f \,|\, \mathcal{E}) \leqslant \sum_i \mu_\Lambda(\mathrm{Var}_{\Lambda_i}(f \,|\, \mathcal{E}_i) \,|\, \mathcal{E}). \tag{4.31}$$

Hence, it is enough to analyse a generic term $\mu_\Lambda(\mathrm{Var}_{\Lambda_i}(f \,|\, \mathcal{E}_i) \,|\, \mathcal{E})$ and for this purpose we plan to apply Lemma 4.2.12 to bound from above $\mathrm{Var}_{\Lambda_i}(f \,|\, \mathcal{E}_i)$.

Recall that the strip Λ_i is the disjoint union of M squares $(Q_{i,j})_{j=1}^M$ and recall the definition of the "single square" events $\mathcal{SG}\left(Q_{i,j}^{(1)}\right)$ and $\mathcal{G}(Q_{i,j})$ given in (i) and (ii) above. Those definitions allow us to write (in what follows the index i of the strip is fixed)

$$\mu_{\Lambda_i}(\cdot \,|\, \mathcal{E}_i) = \nu_i\left(\cdot \,\Big|\, \bigcup_j \bigcap_{j'=j+1}^{j+\lambda} \mathcal{SG}\left(Q_{i,j'}^{(1)}\right)\right)$$

where $\nu_{i,j} = \mu_{Q_{i,j}}(\cdot \,|\, \mathcal{G}(Q_{i,j}))$ and $\nu_i = \bigotimes_j \nu_{i,j}$. We can now apply Lemma 4.2.12 to the product measure ν_i with $\mathcal{SG}\left(Q_{i,j}^{(1)}\right)$ as the event $\mathcal{H}$, M as the parameter n, and λ as the parameter κ. The choice of the key parameter κ entering the definition of the associated facilitating events $\mathcal{H}_{i,j}$ in the periodic case,

$$\mathcal{H}_{i,j} = \bigcap_{j'=j+1}^{j+k} \mathcal{SG}\left(Q_{i,j'}^{(1)}\right) \cup \bigcap_{j'=j-1}^{j-k} \mathcal{SG}\left(Q_{i,j'}^{(1)}\right),$$

will be postponed to Lemma 4.4.4 below. There κ will be chosen to be large enough but independent of q. The requirement of Lemma 4.2.12 that $(1-\hat{\nu}(\mathcal{H})^\kappa)^{n/(3\kappa)} < 1/16$ is implied by (4.23).

In the above setting, Lemma 4.2.12 gives

$$\mathrm{Var}_{\Lambda_i}(f \,|\, \mathcal{E}_i) \leqslant q^{-O(Rw\lambda)} \sum_j \nu_i\left(\mathbb{1}_{\mathcal{H}_{i,j}} \mathrm{Var}_{Q_{i,j}}(f \,|\, \mathcal{G}(Q_{i,j}))\right).$$

By combining the above with (4.31) we finally get

$$\mathrm{Var}_\Lambda(f \mid \mathcal{E}) \leqslant q^{-O(Rw\lambda)} \sum_{i,j} \mu_\Lambda \left(\mathbb{1}_{\mathcal{H}_{i,j}} \mathrm{Var}_{Q_{i,j}}(f \mid \mathcal{G}(Q_{i,j})) \mid \bigcap_{i',j' \in [M]} \mathcal{G}(Q_{i',j'}) \right).$$

We shall now analyse a generic term in the sum above with the help of Theorem 4.3.6.

Lemma 4.4.4. *There exists an constant $\lambda = \lambda(\mathcal{U}, \delta)$ such that the following holds. If the parameter κ of the facilitating events $\mathcal{H}_{i,j}$ is taken equal to λ then, for any function $f : \Omega_\Lambda \to \mathbb{R}$ and any i, j,*

$$\mu_\Lambda \left(\mathbb{1}_{\mathcal{H}_{i,j}} \mathrm{Var}_{Q_{i,j}}(f \mid \mathcal{G}(Q_{i,j})) \mid \bigcap_{i',j'} \mathcal{G}(Q_{i',j'}) \right) \leqslant q^{-O(w^4 \log^3(1/q)/q^\alpha)} \sum_{\substack{x \in \Lambda \\ d(x, Q_{i,j}) \leqslant O(\lambda W)}} \mu_\Lambda \left(c_x \mathrm{Var}_x(f) \right).$$

If we assume the lemma, we immediately recover (4.28), concluding the proof of Theorem 4.4.3. □

Proof of Lemma 4.4.4. We assume that $\mathbb{1}_{\mathcal{H}_{i,j}} = 1$ and that w.l.o.g. the event $\mathcal{H}^+_{i,j} = \bigcap_{j'=j+1}^{j+\lambda} \mathcal{SG}\left(Q^{(1)}_{i,j'}\right)$ occurs. Next, we recall Definition 4.2.4 of the snail $V^R_L(\underline{r})$ and we choose $r_l = \rho_{k+l} \lfloor \delta r_{l-1}/\rho_{k+l} \rfloor$ for all $l \in [2k]$, setting $r_{-1} = L = \lambda(W + 2R_0) - 2R_0$. We choose λ sufficiently large, depending on δ and $\mathcal{U}$ but not on w and q, in such a way that $Q_{i,j} \subset x + V^R_L(\underline{r})$, where x is the center of the rightmost annulus in $Q^{(1)}_{i,j+\lambda}$. We write $V = x + V^R_L(\underline{r})$ and observe that, by construction, $Q_{i,j} \cap B = \varnothing$, where B is the base V. Finally, we recall Definition 4.3.2 of the events $\mathcal{SG}(B), \mathcal{G}(T^\pm_l)$ and $\mathcal{SG}(V) = \mathcal{SG}(B) \cap \bigcap_{l \in [2k]} (\mathcal{G}(T^+_l) \cap \mathcal{G}(T^-_l))$ for the snail V. It is easy to verify the following implications (see Figure 4.10):

$$\mathcal{H}^+_{i,j} \subset \mathcal{SG}(B) \qquad \bigcap_{i',j'} \mathcal{G}(Q_{i',j'}) \subset \bigcap_{l \in [2k]} (\mathcal{G}(T^+_l) \cap \mathcal{G}(T^-_l)). \tag{4.32}$$

Indeed, for the first inclusion, recalling (i) it is clear that $\mathcal{A}$ and $\mathcal{HA}$ occur (since the leftmost annulus in $Q_{i,j+1}$ contains HA and the rightmost one in $Q_{i,j+\lambda}$ contains A) and that all $\mathcal{SB}_{m,p}$ occur (for $SB_{m,p}$ contained in two consecutive squares $Q_{i,j'}, Q_{i,j'+1}$ at least in one of them we are guaranteed to have the helping sets; for $SB_{0,p}$ close to the left boundary of $Q_{i,j'}$ the infected rightmost annulus provides the desired helping sets). To see the second one, observe that for all l, m, $ST^\pm_{l,m}$ intersects at least one of the squares $Q_{i',j'}$ in a segment of length at least εW.

Using (4.32) and $\mu_\Lambda(\mathcal{E}) = 1 - o(1)$, we have that

$$\begin{aligned}
\mu_\Lambda &\left(\mathbb{1}_{\mathcal{H}^+_{i,j}} \operatorname{Var}_{Q_{i,j}}(f \mid \mathcal{G}(Q_{i,j})) \mid \bigcap_{i',j'} \mathcal{G}(Q_{i',j'}) \right) \\
&\leqslant (1+o(1))\mu_\Lambda \left(\mathbb{1}_{\mathcal{SG}(B)} \mathbb{1}_{\bigcap_{i',j'} \mathcal{G}(Q_{i',j'})} \inf_a \mu_{Q_{i,j}} \left(\mathbb{1}_{\mathcal{G}(Q_{i,j})}(f-a)^2 \right) \right) \\
&\leqslant (1+o(1))\mu_\Lambda \left(\inf_a \mu_{Q_{i,j}} \left(\mathbb{1}_{\mathcal{SG}(V)}(f-a)^2 \right) \right) \\
&\leqslant \mu_\Lambda \left(\mathbb{1}_{\mathcal{SG}(V)} \left(f - \mu_V(f \mid \mathcal{SG}(V)) \right)^2 \right) / \mu(\mathcal{SG}(V)) \\
&= \mu_\Lambda \left(\operatorname{Var}_V(f \mid \mathcal{SG}(V)) \right).
\end{aligned}$$

If we now apply the bound (4.9) of Theorem 4.3.6 and use the fact that V is contained in a deterministic $O(\lambda W)$-neighborhood of the square $Q_{i,j}$ we get the conclusion of the lemma, once we observe that $c_x^V \leqslant c_x$, where c_x are the constraints on the torus Λ. □

Appendix

4.A Proof of Lemma 4.2.12

We will consider the linear case—the periodic one is treated identically. For simplicity we assume that $2k$ divides n. Partition $[n]$ into blocks $I_0, \ldots, I_{N-1}$ where $I_i := \{i\kappa, \ldots, (i+1)\kappa - 1\}$ and $N = n/\kappa$. Let $\mathcal{H}^{(\ell)}$ be the event that there exists i in the *left half* $[N]^{(\ell)} := [N/2]$ of $[N]$ such that $\omega_j \in \mathcal{H}$ for all $j \in I_i$. Let $\mathcal{H}^{(r)}$ be defined similarly but for the blocks with index in the right half $[N]^{(r)} := [N] \backslash [N]^{(\ell)}$. Using the assumption of the lemma $\nu(\mathcal{H}^{(\ell)}) = \nu(\mathcal{H}^{(r)}) > 15/16$ and [59, Lemma 6.5], we get

$$\operatorname{Var}_\nu(f \mid \Omega_\mathcal{H}) \leqslant 24\nu \left(\mathbb{1}_{\mathcal{H}^{(r)}} \operatorname{Var}^{(\ell)}(f) + \mathbb{1}_{\mathcal{H}^{(\ell)}} \operatorname{Var}^{(r)}(f) \mid \Omega_\mathcal{H} \right),$$

where $\operatorname{Var}^{(\ell)}$ denotes the variance computed w.r.t. the variables corresponding to the blocks in the left half and similarly for $\operatorname{Var}^{(r)}$.

Given $\mathcal{H}^{(\ell)}$, let ξ be the smallest label in $[N]^{(\ell)}$ such that $\omega_j \in \mathcal{H}$ for all $j \in I_\xi$. Using Lemma 4.2.9 and the fact that the event $\{\xi = i\}$ is independent of the variables $(\omega_j)_{j \geqslant (i+1)\kappa}$, we get that

$$\begin{aligned}
\nu \left(\mathbb{1}_{\mathcal{H}^{(\ell)}} \operatorname{Var}^{(r)}(f) \mid \Omega_\mathcal{H} \right) &\leqslant \sum_{i \in [N]^{(\ell)}} \nu \left(\mathbb{1}_{\{\xi = i\}} \operatorname{Var}_{\geqslant (i+1)\kappa}(f) \mid \Omega_\mathcal{H} \right) \\
&\leqslant \frac{1}{\hat{\nu}(\mathcal{H})^\kappa} \sum_{i \in [N]^{(\ell)}} \nu \left(\mathbb{1}_{\{\xi = i\}} \operatorname{Var}_{\geqslant (i+1)\kappa}(f) \right),
\end{aligned} \tag{4.33}$$

where $\operatorname{Var}_{\geqslant (i+1)\kappa}(f)$ is the variance w.r.t. the variables $(\omega_j)_{j \geqslant (i+1)\kappa}$. The r.h.s. above can now be bounded above using [269, Proposition 3.4]. If $\mathcal{H}_{I_j}$

is the event that $\omega_l \in \mathcal{H}$ for all $l \in I_j$, with the convention that $\mathcal{H}_{I_N}$ and $\mathcal{H}_{I_{-1}}$ do not occur, we get that

$$\mathbb{1}_{\{\xi=i\}} \operatorname{Var}_{j\geqslant(i+1)\kappa}(f) \leqslant \frac{1}{\widehat{\nu}(\mathcal{H})^{O(\kappa)}} \sum_{j=i+1}^{N-1} \nu_{\geqslant(i+1)\kappa}\left(\mathbb{1}_{\{\xi=i\}}\mathbb{1}_{\mathcal{H}_j^{\pm}} \operatorname{Var}_{I_j}(f)\right),$$

where $\mathcal{H}_j^{\pm} = \mathcal{H}_{I_{j-1}} \cup \mathcal{H}_{I_{j+1}}$ and Var_{I_j} is the variance w.r.t. the variables in I_j. By inserting the r.h.s. above into the r.h.s. of (4.33), we obtain that $\nu\left(\mathbb{1}_{\mathcal{H}^{(\ell)}} \operatorname{Var}^{(r)}(f) \,|\, \Omega_{\mathcal{H}}\right)$ is smaller than

$$\frac{1}{\widehat{\nu}(\mathcal{H})^{O(\kappa)}}\nu\left(\sum_{j=1}^{N-1}\sum_{i=0}^{j-2} \mathbb{1}_{\{\xi=i\}}\mathbb{1}_{\mathcal{H}_j^{\pm}} \operatorname{Var}_{I_j}(f) + \sum_{i\in[N]^{\ell}} \mathbb{1}_{\{\xi=i\}}\mathbb{1}_{\mathcal{H}_{i+1}^{\pm}} \operatorname{Var}_{I_{i+1}}(f)\right)$$
$$\leqslant \frac{2}{\widehat{\nu}(\mathcal{H})^{O(\kappa)}} \sum_{j=1}^{N-1} \nu\left(\mathbb{1}_{\mathcal{H}_j^{\pm}} \operatorname{Var}_{I_j}(f)\right),$$

where we have isolated the term $j = i+1$ and used $\sum_i \mathbb{1}_{\{\xi=i\}} \leqslant 1$ and $\mathbb{1}_{\{\xi=i\}} \leqslant 1$ for the two terms respectively. Exactly the same argument can be applied to the term $\nu\left(\mathbb{1}_{\mathcal{H}^{(r)}} \operatorname{Var}^{(\ell)}(f) \,|\, \Omega_{\mathcal{H}}\right)$ to conclude that

$$\operatorname{Var}_{\nu}(f \,|\, \Omega_{\mathcal{H}}) \leqslant \frac{96}{\widehat{\nu}(\mathcal{H})^{O(\kappa)}} \sum_{j=0}^{N-1} \nu\left(\left(\mathbb{1}_{\mathcal{H}_{I_{j+1}}} + \mathbb{1}_{\mathcal{H}_{I_{j-1}}}\right) \operatorname{Var}_{I_j}(f)\right). \tag{4.34}$$

We finally bound from above a generic term, considering $\nu\left(\mathbb{1}_{\mathcal{H}_{I_1}} \operatorname{Var}_{I_0}(f)\right)$ for concreteness.

We apply Lemma 4.2.10 with $X_1 = \omega_{\kappa-1}, X_2 = (\omega_0, \omega_1, \ldots, \omega_{\kappa-2})$ and facilitating event $\{\omega_{\kappa-1} \in \mathcal{H}\}$ to $\operatorname{Var}_{I_0}(f)$ in order to get

$$\operatorname{Var}_{I_0}(f) \leqslant \frac{2}{\widehat{\nu}(\mathcal{H})}\nu_{I_0}\left(\operatorname{Var}_{\kappa-1}(f) + \mathbb{1}_{\{\omega_{\kappa-1}\in\mathcal{H}\}} \operatorname{Var}_{I_0\setminus\{\kappa-1\}}(f)\right). \tag{4.35}$$

Thus, we obtain

$$\begin{aligned}
&\nu\left(\mathbb{1}_{\mathcal{H}_{I_1}} \operatorname{Var}_{I_0}(f)\right)\\
&\quad\leqslant \frac{2}{\widehat{\nu}(\mathcal{H})}\left(\nu\left(\mathbb{1}_{\mathcal{H}_{I_1}} \operatorname{Var}_{\kappa-1}(f)\right) + \nu\left(\mathbb{1}_{\mathcal{H}_{I_1}}\mathbb{1}_{\{\omega_{\kappa-1}\in\mathcal{H}\}} \operatorname{Var}_{I_0\setminus\{\kappa-1\}}(f)\right)\right)\\
&\quad\leqslant \frac{2}{\widehat{\nu}(\mathcal{H})}\left(\nu\left(\mathbb{1}_{\mathcal{H}_{\kappa-1}} \operatorname{Var}_{\kappa-1}(f)\right) + \nu\left(\mathbb{1}_{\mathcal{H}_{\kappa-2}} \operatorname{Var}_{I_0\setminus\{\kappa-1\}}(f)\right)\right)
\end{aligned}$$

We can repeat the step leading to (4.35) for $X_1 = \omega_{\kappa-2}, X_2 = (\omega_0, \ldots, \omega_{\kappa-3})$ and facilitating event $\{\omega_{\kappa-2} \in \mathcal{H}\}$ and so on. At the end of the iteration we finally get that

$$\nu\left(\mathbb{1}_{\mathcal{H}_{I_1}} \operatorname{Var}_{I_0}(f)\right) \leqslant \left(\frac{2}{\widehat{\nu}(\mathcal{H})}\right)^{\kappa} \sum_{i\in I_0} \nu\left(\mathbb{1}_{\mathcal{H}_i} \operatorname{Var}_i(f)\right).$$

Putting all together, we have finally proved that

$$\begin{aligned}\operatorname{Var}_\nu(f) &\leqslant \frac{96}{\widehat{\nu}(\mathcal{H})^{O(\kappa)}}\sum_{j\in[N]}\nu\left(\left(\mathbb{1}_{\mathcal{H}_{I_{j+1}}}+\mathbb{1}_{\mathcal{H}_{I_{j-1}}}\right)\operatorname{Var}_{I_j}(f)\right)\\ &\leqslant\left(\frac{2}{\widehat{\nu}(\mathcal{H})}\right)^{O(\kappa)}\sum_j\sum_{i\in I_j}\nu\left(\mathbb{1}_{\mathcal{H}_i}\operatorname{Var}_i(f)\right).\end{aligned}$$

Chapter 5

Sharp threshold for the FA-2f kinetically constrained model

This chapter is based on joint work with Fabio Martinelli and Cristina Toninelli [215], proving Theorem 1.4.6 (recall Section 1.4.4).

5.1 Proof of Theorem 1.4.6: lower bound

In this section we establish the lower bounds (1.14) of Theorem 1.4.6. Our proof is actually a procedure to establish a general lower bound for $\mathbb{E}_\mu(\tau_0)$ based on bootstrap percolation which improves upon a previous general result [272, Lemma 4.3] which lower bounds $\mathbb{E}_\mu(\tau_0)$ with the mean infection time for the corresponding bootstrap percolation model.

We begin with an auxiliary statement. For a rectangle $R \subset \mathbb{Z}^2$ and $\eta \in \Omega$ we denote by $[\eta]_R$ the set of sites $x \in R$ which can become infected by legal moves only using infections in R. Note that $[\eta]_R$ is a union of disjoint cuboids with sides parallel to the lattice directions. For $x, y \in R$ we write $\{x \overset{R}{\longleftrightarrow} y\}$ for the event that $[\eta]_R$ contains a rectangle containing x and y.

Proposition 5.1.1. *Let $V = [-\ell, \ell]^2$ with $\ell = \ell(q)$ be such that*

$$\mu(0 \in [\eta]_V) = o(1) \tag{5.1}$$

and let

$$\rho := \sup_{x \in V : d(x, V^c) = 1} \mu\big(\{x \overset{V}{\longleftrightarrow} 0\}\big). \tag{5.2}$$

Then

$$\mathbb{E}_\mu(\tau_0) \geqslant \frac{\Omega(1)}{\rho|V|}$$

and $\tau_0 \geqslant q/(|V|\rho)$ w.h.p.

Proof. Let $\{\eta_t\}_{t\geqslant 0}$ denote the stationary KCM on $\mathbb{Z}^2$ and let $\mathcal{I} = \{\omega : 0 \in [\omega]_V\}$. By assumption $\mu(\mathcal{I}) = o(1)$. Let $\tau = \inf\{t \geqslant 0, \eta_t \in \mathcal{I}\}$ and observe that $\tau \leqslant \tau_0$ and that $\mathbb{P}_\mu(\tau > 0) = 1 - o(1)$. Notice also that at $t = \tau$ a flip at a site $x \in V$ with $d(x, V^c) = 1$ which is pivotal for $\mathcal{I}$ must occur. In particular, $\eta_\tau \in \{x \overset{V}{\longleftrightarrow} 0\}$. For $s > 0$ let $N_V(s)$ be the number of clock rings in V up to time s as defined in Section 1.2.2. By stationarity, at each of those updates the KCM configuration is distributed according to μ. Thus,

$$\mathbb{P}_\mu(\tau \leqslant s) = \mathbb{E}\big(\mathbb{P}_\mu(\tau \leqslant s \,|\, N_V(s))\big) \leqslant \mathbb{P}_\mu(\tau = 0) + \mathbb{E}(N_V(s))\rho \leqslant o(1) + s|V|\rho.$$

Above we used a union bound to write

$$\mathbb{P}_\mu(0 < \tau \leqslant s \,|\, N_V(s)) \leqslant N_V(s)\rho,$$

together with $\mathbb{E}(N_V(s)) = s|V|$. In conclusion,

$$\lim_{\varepsilon\to 0} \limsup_{q\to 0} \mathbb{P}_\mu\big(\tau_0|V|\rho < \varepsilon\big) = 0,$$

which concludes the proof by Markov's inequality. □

We can now easily deduce the lower bound of Theorem 1.4.6 from Proposition 5.1.1 and bootstrap percolation results.

Proof of the lower bound (1.14) *in Theorem 1.4.6.* Let $\ell = \frac{\log(1/q)}{4q}$. Condition (5.1) of Proposition 5.1.1 holds by Eq. (1.7) due to [7]. Then Theorem 10.5.1 and Lemma 10.2.9 give $\rho \leqslant \exp(-\frac{\pi^2}{9q} + \frac{O(1)}{\sqrt{q}})$ and (1.14) follows. □

5.2 Mobile droplets

This section, which represents the core of the chapter, is split into two distinct parts:

- the definition of mobile droplets together with the choice of the mesoscopic critical length scale L_{D} characterising their linear size;
- the analysis of two key properties of mobile droplets namely:
  - their equilibrium probability ρ_{D};
  - the relaxation time of FA-2f in a box of linear size $\Theta(L_D)$ *conditionally* on the presence of a mobile droplet.

Mobile droplets are defined as boxes of suitable linear size in which the configuration of infection is *super-good* (see Definition 5.2.6). In turn, the super-good event (see Section 5.2.2) is constructed recursively via a multi-scale procedure on a sequence of exponentially increasing length scales $(\ell_n)_{n=1}^N$

(see Definition 5.2.3). While clearly inspired by the classical procedure used in bootstrap percolation [225], an important novelty in our construction is the freedom that we allow for the position of the super-good core of scale ℓ_n inside the super-good region of scale ℓ_{n+1}. The final scale ℓ_N corresponds to the critical scale L_{D} mentioned above and a convenient choice is $L_{\mathrm{D}} \sim q^{-17/2}$ (see (5.6)). There is nothing special in the exponent 17/2: as long as we choose a sufficiently large exponent our results would not change. The choice of L_D is in fact only dictated by the requirement that w.h.p. there exist no L_{D} consecutive lattice sites at distance $e^{O(\log(1/q))^{\Theta(1)}/q}$ from the origin which are healthy and $L_{\mathrm{D}} \ll e^{o(1/q)}$. Finally, similarly to their bootstrap percolation counterparts, the probability ρ_{D} of mobile droplets crucially satisfies $\rho_{\mathrm{D}} \simeq (\tau_0^{\mathrm{BP}})^{-2}$ (see Proposition 5.2.7).

The extra degree of freedom in the construction of the super-good event provides a much more flexible structure that can be moved around using the FA-2f moves *without* going through the bottleneck corresponding to the creation of a brand new additional droplet nearby. The main consequence of this feature (see Proposition 5.2.9) is that the relaxation time of the FA-2f dynamics in a box of side L_{D} conditioned on being super-good is sub-leading w.r.t. ρ_{D}^{-1} as $q \to 0$ and it contributes only to the second order term in Theorem 1.4.6.

5.2.1 Notation

For any integer n, we write $[n]$ for the set $\{1, \dots, n\}$. We denote by $\vec{e}_1, \vec{e}_2$ the standard basis of $\mathbb{Z}^2$, and write $d(x, y)$ for the Euclidean distance between $x, y \in \mathbb{Z}^2$. Given a set $\Lambda \subset \mathbb{Z}^2$, we set $\partial\Lambda := \{y \in \mathbb{Z}^d \backslash \Lambda, d(y, \Lambda) = 1\}$. Given two positive integers a_1, a_2, we write $R(a_1, a_2) \subset \mathbb{Z}^2$ for the rectangle $[a_1] \times [a_2]$ and we refer to a_1, a_2 as the *width* and *height* of R respectively. We also write $\partial_r R$ ($\partial_l R$) for the column $\{a_1 + 1\} \times [a_2]$ (the column $\{0\} \times [a_2]$), and $\partial_u R$ ($\partial_d R$) for the the row $[a_1] \times \{a_2 + 1\}$ (the row $[a_1] \times \{0\}$). Similarly for any rectangle of the form $R + x, x \in \mathbb{Z}^2$.

Given $\Lambda \subset \mathbb{Z}^2$ and $\omega \in \Omega$, we write $\omega_\Lambda \in \Omega_\Lambda := \{0, 1\}^\Lambda$ for the restriction of ω to Λ. The configuration (in Ω or Ω_Λ) identically equal to one is denoted by $\mathbf{1}$. Given disjoint $\Lambda_1, \Lambda_2 \subset \mathbb{Z}^2$, $\omega^{(1)} \in \Omega_{\Lambda_1}$ and $\omega^{(2)} \in \Omega_{\Lambda_2}$, we write $\omega^{(1)} \cdot \omega^{(2)} \in \Omega_{\Lambda_1 \cup \Lambda_2}$ for the configuration equal to $\omega^{(1)}$ in Λ_1 and to $\omega^{(2)}$ in Λ_2. We write μ_Λ for the marginal of μ on Ω_Λ and $\mathrm{Var}_\Lambda(f)$ for the variance of f w.r.t. μ_Λ, given the variables $(\omega_x)_{x \notin \Lambda}$.

5.2.2 Super-good event and mobile droplets

As anticipated, mobile droplets will be defined as square regions of a certain side length in which the infection configuration satisfies a specific condition dubbed *super-good*. The latter requires in turn the definition of a key event for rectangles—*ω-traversability* (see also [225])—together with a sequence of

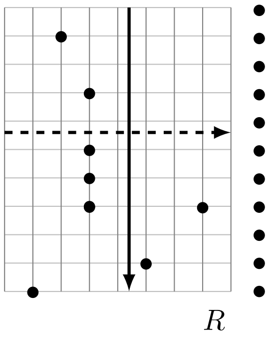

Figure 5.1 – Black circles denote infected sites. The boundary condition ω in the figure is fully infected on $\partial_r R$ and fully healthy elsewhere. The rectangle R is quasi-right-traversable under ω but neither quasi-up-, nor quasi-left-traversable. It is also down-traversable but not traversable in any other direction.

exponentially increasing length scales.

Definition 5.2.1 (ω-Traversability)**.** Fix a rectangle $R = R(a_1, a_2) + x$ together with $\eta \in \Omega_R$ and a boundary configuration $\omega \in \Omega_{\partial R}$. We say that R is *ω-right-traversable (ω-left-traversable)* for η if each couple of adjacent columns of $R \cup \partial_r R$ (of $R \cup \partial_l R$) contains at least one infection. We denote the corresponding event by $\mathcal{T}^{\omega}_{\rightarrow}(R)$ ($\mathcal{T}^{\omega}_{\leftarrow}(R)$) and we depict it in our drawings with a dashed horizontal right (left) arrow (see Figure 5.1). The ω-up(down)-traversability is defined similarly by looking at couples of adjacent rows of $R \cup \partial_u R$ ($R \cup \partial_d R$). The corresponding events will be denoted by $\mathcal{T}^{\omega}_{\uparrow}(R)$ ($\mathcal{T}^{\omega}_{\downarrow}(R)$) and they will be depicted in our drawings with a dashed up (down) arrow.

We say that R is *right-traversable* for η if it is $\mathbf{1}$-right-traversable or, equivalently, if it is ω-right-traversable for all ω. We denote the corresponding event by $\mathcal{T}_{\rightarrow}(R) \equiv \mathcal{T}^{\mathbf{1}}_{\rightarrow}(R)$ and we depict it in our drawings with a solid horizontal right arrow. We define analogously the *left-traversable*, *up-traversable* and *down-traversable* events, $\mathcal{T}_{\leftarrow}(R)$, $\mathcal{T}_{\uparrow}(R)$ and $\mathcal{T}_{\downarrow}(R)$ respectively (see Figure 5.1).

Remark 5.2.2. Notice that right-traversability requires that the rightmost column contains an infection. Similarly for the other directions.

Definition 5.2.3 (Length scales and nested rectangles)**.** For all integer n we set[1]

$$\ell_n = \begin{cases} 1 & \text{if } n = 0, \\ \left\lfloor \frac{\exp(n\sqrt{q})}{\sqrt{q}} \right\rfloor & \text{if } n \geqslant 1 \end{cases} \tag{5.3}$$

and

$$\Lambda^{(n)} = \begin{cases} R(\ell_{n/2}, \ell_{n/2}) & \text{if } n \text{ is even}, \\ R(\ell_{(n+1)/2}, \ell_{(n-1)/2}) & \text{if } n \text{ is odd}. \end{cases} \tag{5.4}$$

Notice that $(\Lambda^{(2m)})_{m \geqslant 0}$ is a sequence of squares, while $(\Lambda^{(2m+1)})_{m \geqslant 0}$ is a sequence of rectangles elongated in the horizontal direction and $\Lambda^{(n_1)} \subset \Lambda^{(n_2)}$

[1]This choice of geometrically increasing length scales is inspired by [190].

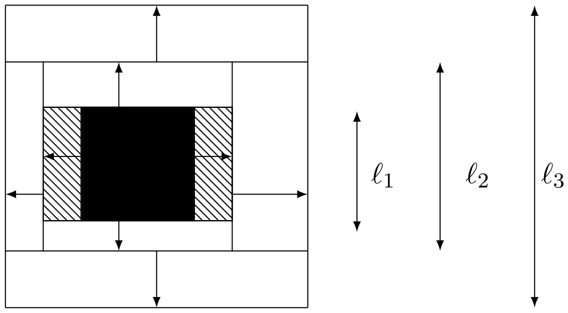

Figure 5.2 – An example of super-good configuration in the square $\Lambda^{(6)}$. The black square, of type $\Lambda^{(2)} + x$, is completely infected and it is a super-good core for the rectangle of type $\Lambda^{(3)}$ formed by it together with the two hatched rectangles. This $\Lambda^{(3)}$-type rectangle is also super-good because of the right/left-traversability of the hatched parts (black arrows) and it is a super-good core for the square containing it and so on.

if $n_1 < n_2$. We also say that a rectangle R is of *class* n if there exist $w, z \in \mathbb{Z}^2$ such that $\Lambda^{(n-1)} + w \subsetneq R \subset \Lambda^{(n)} + z$. Thus, for $n = 0$ R is a single site, for $n = 2m > 0$ it is a rectangle of width ℓ_m and height $a_2 \in (\ell_{m-1}, \ell_m]$ and for $n = 2m + 1$, it is a rectangle of height ℓ_m and width $a_1 \in (\ell_m, \ell_{m+1}]$.

We are now ready to introduce the key notion of the ω-super-good event on different scales. This event is defined recursively on n and it has a hierarchical structure. Roughly speaking, a rectangle R of the form $R = \Lambda^{(n)} + x, x \in \mathbb{Z}^2$, is ω-super-good if it contains a $\mathbf{1}$-super-good rectangle R' of the form $R' = \Lambda^{(n-1)} + x'$ called the *core* and outside the core it satisfies certain ω-traversability conditions (see Figure 5.2).

Definition 5.2.4 (ω-Super-good rectangles)**.** Let us fix an integer $n \geqslant 0$, a rectangle $R = R(a_1, a_2) + x$ of class n and $\omega \in \Omega_{\partial R}$. We say that R is ω-super-good and denote the corresponding event by $\mathcal{SG}^\omega(R)$ if the following occurs.

- $n = 0$. In this case R consists of a single site and $\mathcal{SG}^\omega(R)$ is the event that this site is infected.
- $n = 2m$. For any $s \in [0, \ell_m - \ell_{m-1}]$ write $R = C_s \cup (\Lambda^{(n-1)} + x + s\vec{e}_2) \cup D_s$, where C_s (D_s) is the part of R below (above) $\Lambda^{(n-1)} + x + s\vec{e}_2$. With this notation we set
$$\mathcal{SG}^\omega_s(R) := \mathcal{T}^\omega_\downarrow(C_s) \cap \mathcal{SG}^{\mathbf{1}}(\Lambda^{(n-1)} + x + s\vec{e}_2) \cap \mathcal{T}^\omega_\uparrow(D_s)$$
and let $\mathcal{SG}^\omega(R) = \bigcup_{s\in[0,\ell_m-\ell_{m-1}]} \mathcal{SG}^\omega_s(R)$.
- $n = 2m + 1$. In this case $\mathcal{SG}^\omega(R)$ requires that there is a core in R of the form $\Lambda^{(n-1)} + x + s\vec{e}_1, s \in [0, \ell_{m+1} - \ell_m]$, which is $\mathbf{1}$-super-good and the

two remaining rectangles forming R to the left and to the right of the core are ω-left-traversable and ω-right-traversable respectively.

We will say that R is super-good if it is $\mathbf{1}$-super-good and denote the corresponding event by $\mathcal{SG}(R)$. We have monotonicity in the boundary condition in the sense that if R is super-good then R is ω-super-good for all $\omega \in \Omega_{\partial R}$.

Remark 5.2.5 (Irreducibility of the FA-2f chain in $\mathcal{SG}^\omega(R)$)**.** It is not difficult to verify that for all $\eta \in \mathcal{SG}^\omega(R)$, there exists a path of legal moves that connects η to the fully infected configuration. The above property implies in particular that if we consider the FA-2f chain in R restricted to $\mathcal{SG}^\omega(R)$, then the chain is irreducible.

Now let

$$N := \left\lceil \frac{8\log(1/q)}{\sqrt{q}} \right\rceil \tag{5.5}$$

and observe that

$$\ell_N = q^{-17/2+o(1)}. \tag{5.6}$$

In the sequel we will refer to ℓ_N as the *droplet scale.*

Definition 5.2.6 (Mobile droplets)**.** Given $\omega \in \Omega$, a *mobile droplet* for ω is any square R of the form $R = \Lambda^{(2N)} + x$ for which $\omega_R \in \mathcal{SG}(R)$.

The two key properties of mobile droplets are the following.

Proposition 5.2.7 (Probability of mobile droplets)**.** *For all $n \leqslant 2N$,*

$$\mu_{\Lambda^{(n)}}(\mathcal{SG}) \geqslant \exp\Big(-\frac{\pi^2}{9q}\big(1 + O\big(\sqrt{q}\log^2(1/q)\big)\big)\Big).$$

In particular, this holds for $\rho_{\mathrm{D}} := \mu(\Lambda^{(2N)}$ is a mobile droplet).

Remark 5.2.8. The lower bound of Proposition 5.2.7 is saturated on the droplet scale. Indeed, it is essentially sharp for $n = 2N$.

The proof of Proposition 5.2.7 follows from standard 2-neighbour bootstrap percolation techniques and it is deferred to Appendix 5.A. The second property of mobile droplets requires a bit of preparation.

Given R of class n and $\omega \in \Omega_{\mathbb{Z}^2\setminus R}$, let $\gamma^\omega(R)$ be the best constant C in the Poincaré inequality

$$\mathrm{Var}_R(f|\mathcal{SG}^\omega(R)) \leqslant C \sum_{x\in R} \mu_R(c_x^{R,\omega}\,\mathrm{Var}_x(f)\,|\,\mathcal{SG}^\omega(R)), \tag{5.7}$$

where, for all $\Lambda \subset \mathbb{Z}^2$, all $\omega \in \Omega_{\partial\Lambda}$ and $x \in \Lambda$

$$c_x^{\Lambda,\omega}(\eta) = c_x(\eta\cdot\omega),$$

and

$$c_x(\omega) = \begin{cases} 1 & \text{if } \sum_{y\sim x}(1-\omega_y) \geqslant 2 \\ 0 & \text{otherwise} \end{cases} \tag{5.8}$$

with $y \sim x$ iff x, y are nearest neighbours. The fact that FA-2f restricted to $\mathcal{SG}^\omega(R)$ is irreducible (see Remark 5.2.5) implies that $\gamma^\omega(R)$ is finite. However, proving a good upper bound on $\gamma^\omega(R)$ is quite hard. In the sequel we will sometimes refer to $\gamma^\omega(R)$ as the *relaxation time of* $\mathcal{SG}^\omega(R)$.

Proposition 5.2.9 (Relaxation time of mobile droplets)**.** *For all* $n \leqslant 2N$

$$\max_\omega \gamma^\omega(\Lambda^{(n)}) \leqslant \exp\big(O(\log^2(1/q)n)\big).$$

In particular, on the droplet scale we get

$$\max_\omega \gamma^\omega(\Lambda^{(2N)}) \leqslant e^{O(\log^3(1/q))/\sqrt{q}}.$$

Remark 5.2.10. We stress an important difference in the definition of $\gamma^\omega(\Lambda^{(n)})$ w.r.t. a similar definition in (4.8). Indeed, in (5.7) the conditioning w.r.t. the super-good event $\mathcal{SG}^\omega(R)$ appears in the l.h.s. *and* in the r.h.s. of the inequality, while in (4.8) the conditioning was absent in the r.h.s. Keeping the conditioning also in the r.h.s. is a delicate and important point if one wants to get a Poincaré constant which is *sub-leading* w.r.t. $\rho_{\rm D}^{-1}$. Theorem 4.3.6 in the context of FA-2f would give a Poincaré constant bounded from above by $\exp(\log(1/q)^3/q)$, much bigger than $\rho_{\rm D}^{-1}$.

5.2.3 Proof of Proposition 5.2.9

The proof is unfortunately rather long and technical but the main idea and technical ingredients can be explained as follows.

Given the recursive definition of the super-good event $\mathcal{SG}^\omega(\Lambda^{(n)})$ it is quite natural to try bounding from above its relaxation time in progressively larger and larger volumes. A high-level "dynamical intuition" here goes as follows. After every time interval $\Theta(\gamma^{\mathbf{1}}(\Lambda^{(n-1)}))$ the core of $\Lambda^{(n)}$, namely a super-good rectangle of type $\Lambda^{(n-1)}$ inside $\Lambda^{(n)}$, will equilibrate under the FA-2f dynamics. Therefore, the relaxation time of $\mathcal{SG}(\Lambda^{(n)})$ should be at most $T_{\rm eff}^{(n)} \times \gamma^{\mathbf{1}}(\Lambda^{(n-1)})$, where $T_{\rm eff}^{(n)}$ is the time that it takes for the core to equilibrate its position inside $\Lambda^{(n)}$, assuming that at each time the infections inside it are at equilibrium. The main step necessary to transform this rather vague idea into a proof is as follows.

In order to analyse the characteristic time scale of the effective dynamics of a core, we need to improve and expand a well established mathematical technique for KCM to relate the relaxation times of two ω-super-good regions on different scales. Such a technique introduces various types of *auxiliary constrained block chains* and a large part of our argument is devoted to proving good bounds on their relaxation times (see Section 5.3). The main application of this technique to our concrete problem is summarised in Lemmas 5.2.11 and 5.2.12 below which easily imply Proposition 5.2.9. Let

$$\Lambda^{(n,+)} = \begin{cases} R(\ell_m + 1, \ell_m) & \text{if } n = 2m, \\ R(\ell_{m+1}, \ell_m + 1) & \text{if } n = 2m + 1. \end{cases}$$

The two key steps connecting the relaxation times of super-good rectangles of increasing length scale are as follows.

Lemma 5.2.11 (From $\ell_{\lfloor n/2 \rfloor} + 1$ to $\ell_{\lfloor n/2 \rfloor + 1}$)**.** *For all* $0 \leqslant n \leqslant 2N - 1$

$$\max_\omega \gamma^\omega(\Lambda^{(n+1)}) \leqslant \max_\omega \gamma^\omega(\Lambda^{(n,+)}) \exp(O(\log^2(q))).$$

Lemma 5.2.12 (From $\ell_{\lfloor n/2 \rfloor}$ to $\ell_{\lfloor n/2 \rfloor} + 1$)**.** *For all* $0 \leqslant n \leqslant 2N - 1$

$$\max_\omega \gamma^\omega(\Lambda^{(n,+)}) \leqslant q^{-O(1)} \max_\omega \gamma^\omega(\Lambda^{(n)}).$$

The lemmas imply that

$$\max_\omega \gamma^\omega(\Lambda^{(n+1)}) \leqslant \exp(O(\log(q)^2)) \max_\omega \gamma^\omega(\Lambda^{(n)})$$

and Proposition 5.2.9 follows by induction over n.

Proof of Lemma 5.2.11. Given $n \leqslant 2N - 1$ let K_n be the smallest integer K such that $\lceil (2/3)^K (\ell_{\lfloor n/2 \rfloor + 1} - \ell_{\lfloor n/2 \rfloor}) \rceil = 1$. Definition (5.3) and (5.5) imply that $\max_{n \leqslant 2N-1} K_n \leqslant O(\log(1/q))$. Consider the (exponentially increasing) sequence

$$d_k = \lceil (2/3)^{K_n - k} (\ell_{\lfloor n/2 \rfloor + 1} - \ell_{\lfloor n/2 \rfloor}) \rceil, \quad k \leqslant K_n, \tag{5.9}$$

and let $s_k = d_{k+1} - d_k$ for $k \leqslant K_n - 1$. Next consider the collection $(R^{(k)})_{k=0}^{K_n}$ of rectangles of class $n + 1$ interpolating between $\Lambda^{(n,+)}$ and $\Lambda^{(n+1)}$ defined by

$$R^{(k)} = \begin{cases} R(\ell_m + d_k, \ell_m) & \text{if } n = 2m, \\ R(\ell_{m+1}, \ell_m + d_k) & \text{if } n = 2m + 1. \end{cases}$$

By construction, $R^{(k)} \subset R^{(k+1)}$, $R^{(0)} = \Lambda^{(n,+)}$ and $R^{(K_n)} = \Lambda^{(n+1)}$. Finally, recall the events $\mathcal{SG}^\omega(R)$ and $\mathcal{SG}^\omega_s(R)$ constructed in Definition 5.2.4) for any rectangle R of class $n + 1 \leqslant 2N$ and let

$$a_k = \max_\omega \left(\mu_{R^{(k)}}(\mathcal{SG}^{\mathbf{1}}_{d_k} \,|\, \mathcal{SG}^\omega)\right)^{-2} \max_\omega \left(\mu_{R^{(k)}}(\mathcal{SG}^\omega_0 \,|\, \mathcal{SG}^\omega)\right)^{-1}, \tag{5.10}$$

where $\max_\omega$ is over all $\omega \in \Omega_{\partial R^{(k)}}$. In Corollary 5.A.3 we prove that

$$\mu_R\big(\mathcal{SG}^\omega_s(R) \,|\, \mathcal{SG}^{\omega'}(R)\big) \geqslant q^{O(1)}$$

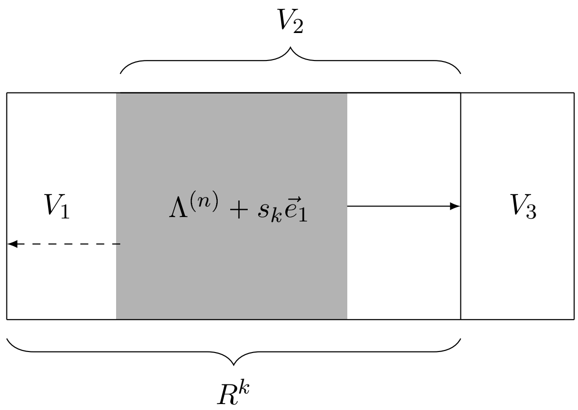

Figure 5.3 – The partition of $R^{(k+1)}$ into the rectangles V_1, V_2, V_3. We illustrate here the event $\mathcal{F}_{1,2}$. The grey region $\Lambda^{(n)}+s_k\vec{e}_1$ to the left boundary of V_2 is $\mathcal{SG}$ and the *dashed* arrow in V_1 indicates its ω-traversability. The *solid* arrow in $V_2\backslash(\Lambda^{(n)} + s_k\vec{e}_1)$ indicates instead the **1**-traversability of $V_2\backslash(\Lambda^{(n)} + s_k\vec{e}_1)$. Clearly the entire configuration belongs to the events $\mathcal{H}$ and $\mathcal{K}$ defined in (5.14), (5.15).

uniformly over all rectangles R of class $n + 1 \leqslant 2N$, all possible values of the offset s and all choices of the boundary configurations $\omega, \omega' \in \Omega_{\partial R}$. As a consequence

$$\max_{n\leqslant 2N-1}\max_{k\leqslant K_n} a_k \leqslant (1/q)^{O(1)}. \tag{5.11}$$

With the above notation the key inequality for proving Lemma 5.2.11 is

$$\max_{\omega}\gamma^{\omega}(R^{(k+1)}) \leqslant Ca_k \times \max_{\omega}\gamma^{\omega}(R^{(k)}), \quad k \in [0, K_n - 1], \tag{5.12}$$

for some universal constant $C > 0$. Recalling that $R^{(0)} = \Lambda^{(n,+)}$ and $R^{(K_n)} = \Lambda^{(n+1)}$, from (5.12) it follows that

$$\max_{\omega}\gamma^{\omega}(\Lambda^{(n+1)}) \leqslant \Big(C^{K_n}\prod_{k=0}^{K_n-1} a_k\Big) \times \max_{\omega}\gamma^{\omega}(\Lambda^{(n,+)}) \tag{5.13}$$

which in turn implies Lemma 5.2.11 by (5.11) and $K_n \leqslant O(\log(1/q))$.

The proof of (5.12), which is detailed for simplicity only in the even case $n = 2m$, relies on the Poincaré inequality for a properly chosen *auxiliary block chain* proved in Proposition 5.3.5. In order to exploit that proposition we partition $R^{(k+1)}$ into three disjoint rectangles V_1, V_2, V_3 as follows (see Figure 5.3):

$$V_1 = R(s_k, \ell_m), \qquad V_2 = R^{(k)}\backslash V_1, \qquad V_3 = R^{(k+1)}\backslash R^{(k)}.$$

Then, given a boundary configuration $\omega \in \Omega_{\partial R^{(k+1)}}$, let

$$\mathcal{H} = \{\eta : \eta_3 \in \mathcal{T}^{\omega}_{\rightarrow}(V_3) \text{ and } \eta_1 \cdot \eta_2 \in \mathcal{SG}^{\eta\cdot\omega}(V_1 \cup V_2)\}, \tag{5.14}$$
$$\mathcal{K} = \{\eta : \eta_1 \in \mathcal{T}^{\omega}_{\leftarrow}(V_1) \text{ and } \eta_2 \cdot \eta_3 \in \mathcal{SG}^{\eta\cdot\omega}(V_2 \cup V_3)\}, \tag{5.15}$$

$\eta_i := \eta_{V_i}$ and $\eta \cdot \omega$ denotes the configuration equal to ω on $\partial R^{(k+1)}$ and equal to η on $R^{(k+1)}$. Notice that $\mathcal{H} \cup \mathcal{K} = \mathcal{SG}^{\omega}(R^{(k+1)})$. The width of V_2 is in fact $\ell_m + 2d_k - d_{k+1} \geqslant \ell_m$ and therefore any configuration in $\mathcal{SG}^{\omega}(R^{k+1})$ necessarily contains a super-good core in either $V_1 \cup V_2$ or $V_2 \cup V_3$.

We next introduce two additional events

$$\mathcal{F}_{1,2} = \mathcal{SG}^{\mathbf{1}\cdot\omega}_{s_k}(V_1 \cup V_2) \qquad\qquad \mathcal{F}_{2,3} = \mathcal{SG}^{\mathbf{1}\cdot\omega}_{0}(V_2 \cup V_3),$$

where, with a slight abuse of notation, $\mathbf{1} \cdot \omega$ equals $\mathbf{1}_{R^{(k+1)}} \cdot \omega$. In words, $\mathcal{F}_{1,2}$ ($\mathcal{F}_{2,3}$) consists of super-good configurations in $V_1 \cup V_2$ ($V_2 \cup V_3$) with a super-good core of type $\Lambda^{(n)}$ inside V_2 in the *leftmost* possible position. Monotonicity in the boundary condition easily implies that

$$\{\eta : \eta_3 \in \mathcal{T}^{\omega}_{\rightarrow}(V_3) \text{ and } \eta_1 \cdot \eta_2 \in \mathcal{F}_{1,2}\} \subset \mathcal{H} \cap \mathcal{K},$$

and similarly for $\mathcal{F}_{2,3}$.

At this stage we can apply Proposition 5.3.5 with parameters $\Omega_i = \Omega_{V_i}$ for $i \in \{1,2,3\}$, $\mathcal{A}_1 = \mathcal{T}^{\omega}_{\leftarrow}(V_1)$, $\mathcal{A}_3 = \mathcal{T}^{\omega}_{\rightarrow}(V_3)$, $\mathcal{B}^{\eta_3}_{1,2} = \mathcal{SG}^{\eta\cdot\omega}(V_1 \cup V_2)$ and $\mathcal{B}^{\eta_1}_{2,3} = \mathcal{SG}^{\eta\cdot\omega}(V_2 \cup V_3)$ to get

$$\begin{aligned}&\operatorname{Var}_{R^{(k+1)}}\big(f \,|\, \mathcal{SG}^{\omega}(R^{(k+1)})\big) = \operatorname{Var}_{R^{(k+1)}}\big(f \,|\, \mathcal{H} \cup \mathcal{K}\big)\\ &\leqslant cT^{(1)}_{\text{aux}} \times \mu_{R^{(k+1)}}\big(\mathbb{1}_{\mathcal{H}} \operatorname{Var}(f \,|\, \mathcal{H}, \eta_3) + \mathbb{1}_{\mathcal{K}} \operatorname{Var}(f \,|\, \mathcal{K}, \eta_1) \,|\, \mathcal{H} \cup \mathcal{K}\big),\end{aligned} \tag{5.16}$$

for some universal constant $c > 0$, where

$$T^{(1)}_{\text{aux}} = \max_{\substack{\eta_{V_1} \in \mathcal{T}^{\omega}_{\leftarrow}(V_1)\\ \eta_{V_3} \in \mathcal{T}^{\omega}_{\rightarrow}(V_3)}} \left(\frac{\mu_{R^{(k+1)}}\big(\mathcal{SG}^{\eta\cdot\omega}(V_1 \cup V_2)\big)}{\mu_{R^{(k+1)}}\big(\mathcal{F}_{1,2}\big)}\right)^2 \times \frac{\mu_{R^{(k+1)}}\big(\mathcal{SG}^{\eta\cdot\omega}(V_2 \cup V_3)\big)}{\mu_{R^{(k+1)}}\big(\mathcal{F}_{2,3}\big)}.$$

Using (5.10) and the fact that $V_1 \cup V_2 = R^{(k)}$, $V_2 \cup V_3 = R^{(k)} + s_k$, one easily sees that $T^{(1)}_{\text{aux}} \leqslant a_k$.

In order to conclude the proof of (5.12) we are left with the analysis of the average w.r.t. $\mu_{R^{(k+1)}}(\cdot \,|\, \mathcal{H} \cup \mathcal{K})$ in the r.h.s. of (5.16). Recalling (5.7), for any $\eta \in \Omega_{R^{(k+1)}}$ such that $\eta_3 \in \mathcal{T}^{\omega}_{\rightarrow}(V_3)$ we get

$$\begin{aligned}\operatorname{Var}(f \,|\, \mathcal{H}, \eta_3) \leqslant &\max_{\omega' \in \Omega_{\partial R^{(k)}}} \gamma^{\omega'}(R^{(k)})\\ &\times \sum_{y \in R^{(k)}} \mu_{R^{(k)}}\big(c_y^{R^{(k)},\eta_3\cdot\omega} \operatorname{Var}_y(f) | \mathcal{SG}^{\eta\cdot\omega}(R^{(k)})\big).\end{aligned} \tag{5.17}$$

An analogous inequality holds for $\operatorname{Var}(f\,|\,\mathcal{K},\eta_1)$ when $\eta_1 \in \mathcal{T}^\omega_{\leftarrow}$. Finally,

$$\mu_{R^{(k+1)}}\big(\mathbb{1}_{\mathcal{H}}\mu_{R^{(k)}}\big(c_y^{R^{(k)},\eta_3\cdot\omega}\operatorname{Var}_y(f)\,|\,\mathcal{SG}^{\eta\cdot\omega}(R^{(k)})\big)\,|\,\mathcal{SG}^\omega(R^{(k+1)})\big)$$
$$= \mu_{R^{(k+1)}}\big(\mathbb{1}_{\mathcal{H}}c_y^{R^{(k+1)},\omega}\operatorname{Var}_y(f)\,|\,\mathcal{SG}^\omega(R^{(k+1)})\big)$$

and the similarly for $\mathcal{K}$. Inserting the above into (5.16), we get

$$\operatorname{Var}_{R^{(k+1)}}\big(f\,|\,SG^\omega(R^{(k+1)})\big) \leqslant O(a_k)\times\max_{\omega'}\gamma^{\omega'}(R^{(k)})$$
$$\times\sum_{x\in R^{(k+1)}}\mu_{R^{(k+1}}\big(c_x^{R^{(k+1)},\omega}\operatorname{Var}_x(f)\,|\,SG^\omega(R^{(k+1)})\big),$$

which proves (5.12). □

Proof of Lemma 5.2.12. The proof is similar to the proof of Lemma 5.2.11, but in this case we plan to use Proposition 5.3.7. Again we provide the details only in the case $n = 2m$.

The result for the case $m = 0$ follows immediately since $\Lambda^{(0,+)}$ contains only two sites. If $m \geqslant 1$ we begin by writing $\Lambda^{(n,+)} = R(\ell_m + 1, \ell_m) = V_1 \cup V_2 \cup V_3$, where V_1 denotes the leftmost column, V_3 the rightmost column and V_2 all the remaining columns (see Fig 5.4). By construction $V_1 \cup V_2$ and $V_2 \cup V_3$ are translates of $\Lambda^{(n)}$. Then, for any given $\omega \in \Omega_{\partial\Lambda^{(n,+)}}$, we introduce the events

$$\mathcal{M} = \mathcal{T}^\omega_{\rightarrow}(V_3) \cap \mathcal{SG}(V_1 \cup V_2) \qquad \mathcal{N} = \mathcal{T}^\omega_{\leftarrow}(V_1) \cap \mathcal{SG}(V_2 \cup V_3)$$

and observe that $\mathcal{SG}^\omega(\Lambda^{(n,+)}) = \mathcal{M} \cup \mathcal{N}$. In order to be able to use Proposition 5.3.7 we need some further events. The first one is the event $\overline{\mathcal{SG}}(V_2)$ which is best explained by Figure 5.4. It corresponds to requiring that inside the rectangle $V_2 = R(\ell_m - 1, \ell_m) + \vec{e}_1$ there exists a **1**-super-good square $R(\ell_{m-1}, \ell_{m-1}) + x$ and the remaining rectangles in $V_2 \backslash R(\ell_{m-1}, \ell_{m-1}) + x$ which sandwich $R(\ell_{m-1}, \ell_{m-1}) + x$ are **1**-traversable. The formal Definition 5.A.4 is left to Appendix 5.A.

It is immediate to verify that for any $\eta_2 \in \overline{\mathcal{SG}}(V_2)$ there exist two vertical intervals $I_1 = I_1(\eta_2) \subset V_1$ and $I_3 = I_3(\eta_2) \subset V_3$ such that $\eta_{I_1} \neq \mathbf{1}$ implies that $\eta_1 \cdot \eta_2 \in \mathcal{SG}(V_1 \cup V_2)$ and similarly if $\eta_{I_3} \neq \mathbf{1}$. Here, as before, $\eta_i := \eta_{V_i}$. We then set

$$\hat{\mathcal{C}}_{1,2} := \big\{\eta : \eta_2 \in \overline{\mathcal{SG}}(V_2), \eta_{I_1(\eta_2)} \neq \mathbf{1}\big\} \tag{5.18}$$

and for $\eta \in \hat{\mathcal{C}}_{1,2}$ we let

$$\mathcal{A}_3^{\eta_1\cdot\eta_2} = \{\eta_{I_3(\eta_2)} \neq \mathbf{1}\}.$$

By construction

$$\big\{\eta :\ \eta_1 \cdot \eta_2 \in \hat{\mathcal{C}}_{1,2} \text{ and } \eta_3 \in \mathcal{A}_3^{\eta_1\cdot\eta_2}\big\} \subset \mathcal{M} \cap \mathcal{N}.$$

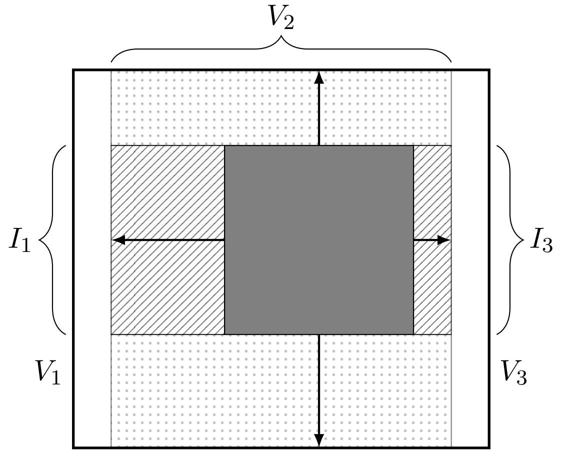

Figure 5.4 – The partition of $\Lambda^{(n,+)}$ into the rectangle V_2 and the two columns V_1 and V_3. Here we illustrate the event $\overline{\mathcal{SG}}(V_2)$: the grey region is a super-good rectangle of the type $\Lambda^{(n-2)}$, while the patterned rectangles are **1**-traversable in the arrow directions. If there is at least one infection in I_3 then the rectangle $V_2 \cup V_3$ is super-good. Analogously for I_1.

We can finally apply Proposition 5.3.7 with parameters $\mathcal{C}_{1,2} = \mathcal{SG}(V_1 \cup V_2)$, $\mathcal{C}_{2,3} = \mathcal{SG}(V_2 \cup V_3)$, $\mathcal{A}_1 = \mathcal{T}^{\omega}_{\leftarrow}(V_1)$ and $\mathcal{A}_3 = \mathcal{T}^{\omega}_{\rightarrow}(V_3)$ to get that

$$\begin{aligned}
&\mathrm{Var}_{\Lambda^{(n,+)}}(f \,|\, \mathcal{M} \cup \mathcal{N}) \\
&\quad \leqslant cT^{(2)}_{\mathrm{aux}} \times \mu_{\Lambda^{(n,+)}}\Big(\mathbb{1}_{\mathcal{M}} \mathrm{Var}(f \,|\, \mathcal{C}_{1,2}, \eta_3) + \mathbb{1}_{\mathcal{N}} \mathrm{Var}(f \,|\, \mathcal{C}_{2,3}, \eta_1) \\
&\quad + \mathbb{1}_{\mathcal{M}} \mathrm{Var}(f \,|\, \mathcal{A}_3, \eta_1, \eta_2) + \mathbb{1}_{\mathcal{N}} \mathrm{Var}(f \,|\, \mathcal{A}_1, \eta_2, \eta_3) \,|\, \mathcal{M} \cup \mathcal{N}\Big),
\end{aligned} \tag{5.19}$$

for some universal constant $c > 0$, with

$$T^{(2)}_{\mathrm{aux}} = \max_{\eta \in \hat{\mathcal{C}}_{1,2}} \frac{\mu_{\Lambda^{(n,+)}}(\mathcal{A}_3)}{\mu_{\Lambda^{(n,+)}}(\mathcal{A}_3^{\eta_1,\eta_2})} \times \frac{\mu_{\Lambda^{(n,+)}}(\mathcal{C}_{1,2})}{\mu_{\Lambda^{(n,+)}}(\hat{\mathcal{C}}_{1,2})}.$$

Clearly, $\min_{\eta \in \hat{\mathcal{C}}_{1,2}} \mu_{\Lambda^{(n,+)}}(\mathcal{A}_3^{\eta_1,\eta_2}) \geqslant q$. Furthermore, in Lemma 5.A.5 we will establish that $\mu_{V_1 \cup V_2}(\hat{\mathcal{C}}_{1,2} \,|\, \mathcal{C}_{1,2}) \geqslant q^{O(1)}$. All together

$$T^{(2)}_{\mathrm{aux}} \leqslant q^{-O(1)}. \tag{5.20}$$

We now turn to examine the four averages w.r.t. $\mu_{\Lambda^{(n,+)}}(\cdot \,|\, \mathcal{M} \cup \mathcal{N})$ appearing in the r.h.s. of (5.19). Recall that $\mathcal{M} \cup \mathcal{N} = \mathcal{SG}^{\omega}(\Lambda^{(n,+)})$. For the first and second average we can mimic what we did for (5.17) and get that they are both bounded from above by

$$\max_{\omega'} \gamma^{\omega'}(\Lambda^{(n)}) \sum_{x \in V_1 \cup V_2} \mu_{\Lambda^{(n,+)}}\big(c_x^{\Lambda^{(n,+)},\omega} \mathrm{Var}_x(f) \,|\, \mathcal{SG}^{\omega}(\Lambda^{(n,+)})\big). \tag{5.21}$$

We will now explain how to upper bound the third average,

$$\mu_{\Lambda^{(n,+)}}\Big(\mathbb{1}_{\mathcal{M}}\operatorname{Var}(f\,|\,\mathcal{A}_3,\eta_1,\eta_2)\,|\,\mathcal{M}\cup\mathcal{N}\Big),$$

the forth one being similar. We need to distinguish two cases, according to whether the boundary condition ω has an infection on the column $V_3+\vec{e}_1$ or not.

Assume $\omega_{V_3+\vec{e}_1}=\mathbf{1}$ In this case $\mathcal{A}_3=\mathcal{T}_{\rightarrow}(V_3)=\Omega_{V_3}\backslash\{\mathbf{1}\}$ and Proposition 5.3.2 (1), gives that

$$\begin{aligned}\operatorname{Var}(f\,|\,\mathcal{A}_3,\eta_1,\eta_2)&=\operatorname{Var}_{V_3}(f\,|\,\mathcal{T}_{\rightarrow}(V_3))\\&\leqslant q^{-O(1)}\sum_{x\in V_3}\mu_{V_3}\big(\tilde{c}_x\operatorname{Var}_x(f)\,|\,\mathcal{T}_{\rightarrow}(V_3)\big),\end{aligned}\tag{5.22}$$

with $\tilde{c}_x(\eta)=1$ if x has at least one infected neighbour inside V_3 and $\tilde{c}_x(\eta)=0$ otherwise. For $x\in V_3$ let

$$A_x=\mu_{\Lambda^{(n,+)}}\Big(\mathbb{1}_{\mathcal{M}}\ \mu_{V_3}\big(\tilde{c}_x\operatorname{Var}_x(f)\,|\,\mathcal{T}_{\rightarrow}(V_3)\big)\,|\,\mathcal{M}\cup\mathcal{N}\Big).$$

Using $\mathcal{M}\subset\mathcal{SG}^{\omega}(\Lambda^{(n,+)})$, $\mu_{\Lambda^{(n,+)}}(\mathcal{T}_{\rightarrow}(V_3)\,|\,\mathcal{SG}^{\omega}(\Lambda^{(n,+)}))=1$, $V_1\cup V_2=\Lambda^{(n)}$ and the fact that the average of a conditional variance is not more than the total variance, we get

$$\begin{aligned}A_x&=\mu_{\Lambda^{(n,+)}}\big(\mathbb{1}_{\mathcal{M}}\tilde{c}_x\operatorname{Var}_x(f)\,|\,\mathcal{SG}^{\omega}(\Lambda^{(n,+)})\big)\\&\leqslant\mu_{\Lambda^{(n,+)}}\big(\mathbb{1}_{\mathcal{M}}\tilde{c}_x\operatorname{Var}_{\{x\}\cup\Lambda^{(n)}}(f\,|\,\mathcal{SG}(\Lambda^{(n)}))\,|\,\mathcal{SG}^{\omega}(\Lambda^{(n,+)})\big).\end{aligned}$$

Next, we use Proposition 5.3.3 to write

$$\operatorname{Var}_{\{x\}\cup\Lambda^{(n)}}(f|\mathcal{SG}(\Lambda^{(n)}))\leqslant\frac{2}{q}\big(\operatorname{Var}_{\Lambda^{(n)}}(f|\mathcal{SG}(\Lambda^{(n)}))+\mathbb{1}_{\{\eta_{x-\vec{e}_1}=0\}}\operatorname{Var}_x(f)\big).$$

Recalling (5.7), we get

$$\begin{aligned}\operatorname{Var}_{\Lambda^{(n)}}(f|\mathcal{SG}(\Lambda^{(n)}))&\leqslant\gamma^{\mathbf{1}}(\Lambda^{(n)})\sum_{y\in\Lambda^{(n)}}\mu_{\Lambda^{(n)}}\big(c_y^{\Lambda^{(n)}\mathbf{1}}\operatorname{Var}_y(f)\,|\,\mathcal{SG}(\Lambda^{(n)})\big)\\&\leqslant\gamma^{\mathbf{1}}(\Lambda^{(n)})\sum_{y\in\Lambda^{(n)}}\mu_{\Lambda^{(n)}}\big(c_y^{\Lambda^{(n,+)},\omega}\operatorname{Var}_y(f)\,|\,\mathcal{SG}(\Lambda^{(n)})\big),\end{aligned}$$

because $c_y^{\Lambda^{(n)},\mathbf{1}}\leqslant c_y^{\Lambda^{(n,+)},\omega}$. Finally, observe that $\mathbb{1}_{\{\eta_{x-\vec{e}_1}=0\}}\tilde{c}_x\leqslant c_x^{\Lambda^{(n,+)},\omega}$, because if $x\in V_3$ has an infected neighbour in V_3 (the constraint $\tilde{c}_x$) and $x-\vec{e}_1\in V_2$ is also infected then x has two infected neighbours in $\Lambda^{(n,+)}$. Putting all together, we conclude that

$$\begin{aligned}A_x&\leqslant\frac{2}{q}\gamma^{\mathbf{1}}(\Lambda^{(n)})\\&\times\mu_{\Lambda^{(n,+)}}\Big(c_x^{\Lambda^{(n,+)},\omega}\operatorname{Var}_x(f)+\sum_{y\in\Lambda^{(n)}}c_y^{\Lambda^{(n,+)},\omega}\operatorname{Var}_y(f)\,|\,\mathcal{SG}^{\omega}(\Lambda^{(n,+)})\Big).\end{aligned}$$

In conclusion, using $|V_3| = \ell_m = q^{-O(1)}$ we get that when $\omega_{V_3+\vec{e}_1} = \mathbf{1}$ the third average in the r.h.s. of (5.19) satisfies

$$\begin{aligned}\mu_{\Lambda^{(n,+)}}\Big(\mathbb{1}_{\mathcal{M}} \operatorname{Var}(f \,|\, \mathcal{A}_3, \eta_1, \eta_2) \,|\, \mathcal{M} \cup \mathcal{N}\Big) &\leqslant \sum_{x \in V_3} A_x \\ &\leqslant \frac{\gamma^{\mathbf{1}}(\Lambda^{(n)})}{q^{-O(1)}} \sum_{y \in \Lambda^{(n,+)}} \mu_{\Lambda^{(n,+)}}\Big(c_y^{\Lambda^{(n,+)},\omega} \operatorname{Var}_y(f) \,|\, \mathcal{SG}^{\omega}(\Lambda^{(n,+)})\Big). \quad (5.23)\end{aligned}$$

Combining (5.19), (5.20), (5.21) and (5.23) (and its analogue for the forth average in the r.h.s. of (5.19)), we conclude the proof of Lemma 5.2.12 in this case.

Assume $\omega_{V_3+\vec{e}_1} \neq \mathbf{1}$ In this case $\mathcal{T}_{\rightarrow}^{\omega}(V_3) = \Omega_{V_3}$, so $\operatorname{Var}_{V_3}(f \,|\, \mathcal{T}_{\rightarrow}^{\omega}(V_3)) = \operatorname{Var}_{V_3}(f)$. The proof is then identical to the previous one except for inequality (5.22) which now follows from Proposition 5.3.2 (2) with any site in V_3 neighbouring an infection of $\omega_{V_3+\vec{e}_1}$ being unconstrained. □

5.3 Constrained Poincaré inequalities

In this section we state and prove various Poincaré inequalities for the auxiliary chains that were instrumental for the proofs of Lemmas 5.2.11 and 5.2.12.

5.3.1 FA-1f-type Poincaré inequalities

Fix $\Lambda \subset \mathbb{Z}^2$ a connected set and let $\widetilde{\Omega}_\Lambda = \Omega_\Lambda \backslash \mathbf{1}$. Given $x \in \Lambda$ let N_x^Λ be the set of neighbours of x in Λ and let $\mathcal{N}_x^\Lambda$ be the event that N_x^Λ contains at least one infection. For any $z \in \Lambda$ consider the two Dirichlet forms

$$\begin{aligned}\mathcal{D}_\Lambda^{\text{FA-1f}}(f) &= \mu_\Lambda\Big(\sum_{x \in \Lambda} \mathbb{1}_{\mathcal{N}_x^\Lambda} \operatorname{Var}_x(f) \,|\, \widetilde{\Omega}_\Lambda\Big), \qquad & f : \widetilde{\Omega}_\Lambda \to \mathbb{R}, \\ \mathcal{D}_\Lambda^{\text{FA-1f},z}(f) &= \mu_\Lambda\Big(\sum_{\substack{x \in \Lambda \\ x \neq z}} \mathbb{1}_{\mathcal{N}_x^\Lambda} \operatorname{Var}_x(f) + \operatorname{Var}_z(f)\Big), \qquad & f : \Omega_\Lambda \to \mathbb{R}.\end{aligned}$$

Remark 5.3.1. The alert reader will recognise the above expressions as the Dirichlet forms of the FA-1f process on $\widetilde{\Omega}_\Lambda$ or on Ω_Λ with the site z unconstrained.

Our first tool is a Poincaré inequality for these Dirichlet forms.

Proposition 5.3.2. *Let Λ be a connected subset of $\mathbb{Z}^2$ and let $z \in \Lambda$ be an arbitrary site. Then:*

(1) for any $f : \widetilde{\Omega}_\Lambda \to \mathbb{R}$,

$$\mathrm{Var}_\Lambda(f \,|\, \widetilde{\Omega}_\Lambda) \leqslant \frac{1}{q^{O(1)}} \mathcal{D}_\Lambda^{\mathrm{FA-1f}}(f); \tag{5.24}$$

(2) for any $f : \Omega_\Lambda \to \mathbb{R}$,

$$\mathrm{Var}_\Lambda(f) \leqslant \frac{1}{q^{O(1)}} \mathcal{D}_\Lambda^{\mathrm{FA-1f},z}(f), \tag{5.25}$$

where the constants in the $O(1)$ *do not depend on* z *or* Λ.

Proof. Inequality (5.24) is proved in [59, Theorem 6.1]. In order to prove (5.25), consider the auxiliary Dirichlet form

$$\mu_\Lambda(\mathrm{Var}_z(f)) + \mu_\Lambda\big(\mathbb{1}_{\widetilde{\Omega}_\Lambda} \mathrm{Var}_\Lambda(f \,|\, \widetilde{\Omega}_\Lambda)\big).$$

The corresponding ergodic, continuous time Markov chain on Ω_Λ, reversible w.r.t. μ_Λ, updates the state of z at rate 1 and, if $\omega \in \widetilde{\Omega}_\Lambda$, it updates the entire configuration w.r.t. $\pi(\cdot \,|\, \widetilde{\Omega}_\Lambda)$. Since the chain enters $\widetilde{\Omega}_\Lambda$ at rate q (by flipping ω_z), a simple coupling argument shows that its relaxation time is $O(1/q)$. Hence,

$$\begin{aligned} \mathrm{Var}_\Lambda(f) &\leqslant O(1)/q\Big(\mu_\Lambda(\mathrm{Var}_z(f)) + \mu_\Lambda\big(\mathbb{1}_{\widetilde{\Omega}_\Lambda} \mathrm{Var}_\Lambda(f \,|\, \widetilde{\Omega}_\Lambda)\big)\Big) \\ &\leqslant \frac{1}{q^{O(1)}}\Big(\mu_\Lambda(\mathrm{Var}_z(f)) + \mu_\Lambda(\widetilde{\Omega}_\Lambda)\mathcal{D}_\Lambda^{\mathrm{FA-1f}}(f)\Big), \end{aligned}$$

where the second inequality follows from (5.24). We may then conclude by observing that $\mu_\Lambda(\mathrm{Var}_z(f)) + \mu_\Lambda(\widetilde{\Omega}_\Lambda)\mathcal{D}_\Lambda^{\mathrm{FA-1f}}(f) \leqslant 2\mathcal{D}_\Lambda^{\mathrm{FA-1f},z}(f)$. □

Our second tool is a general constrained Poincaré inequality for two independent random variables.

Proposition 5.3.3 (See Lemma 4.2.10)**.** *Let* X_1, X_2 *be two independent random variable taking values in two finite sets* $\mathbb{X}_1, \mathbb{X}_2$ *respectively. Let also* $\mathcal{H} \subset \mathbb{X}_1$ *with* $\mathbb{P}(X_1 \in \mathcal{H}) > 0$. *Then for any* $f : \mathbb{X}_1 \times \mathbb{X}_2 \to \mathbb{R}$ *it holds*

$$\mathrm{Var}(f) \leqslant 2\mathbb{P}(X_1 \in \mathcal{H})^{-1}\mathbb{E}\big(\mathrm{Var}_1(f) + \mathbb{1}_{\{X_1 \in \mathcal{H}\}} \mathrm{Var}_2(f)\big).$$

with $\mathrm{Var}_i(f) = \mathrm{Var}(f(X_1, X_2) \,|\, X_i)$.

5.3.2 Constrained block chains

In this section we define two auxiliary constrained reversible Markov chains and give an upper bound for the corresponding Poincaré constants (Lemmas 5.3.5 and 5.3.7).

Let $(\Omega_i, \pi_i)_{i=1}^3$ be finite probability spaces and let (Ω, π) denote the associated product space. For $\omega \in \Omega$ we write $\omega_i \in \Omega_i$ for its i^{th} coordinate and we assume for simplicity that $\pi_i(\omega_i) > 0$ for each ω_i. Fix $\mathcal{A}_3 \subset \Omega_3$ and for each $\omega_3 \in \mathcal{A}_3$ consider an event $\mathcal{B}_{1,2}^{\omega_3} \subset \Omega_1 \times \Omega_2$. Analogously, fix $\mathcal{A}_1 \subset \Omega_1$ and for each $\omega_1 \in \mathcal{A}_1$ consider an event $\mathcal{B}_{2,3}^{\omega_1} \subset \Omega_2 \times \Omega_3$. We then set

$$\begin{aligned}\mathcal{H} &= \{\omega : \omega_3 \in \mathcal{A}_3 \text{ and } (\omega_1, \omega_2) \in \mathcal{B}_{1,2}^{\omega_3}\},\\ \mathcal{K} &= \{\omega : \omega_1 \in \mathcal{A}_1 \text{ and } (\omega_2, \omega_3) \in \mathcal{B}_{2,3}^{\omega_1}\}\end{aligned}$$

and let for any $f : \mathcal{H} \cup \mathcal{K} \to \mathbb{R}$

$$\mathcal{D}_{\text{aux}}^{(1)}(f) = \pi\big(\mathbb{1}_{\mathcal{H}} \operatorname{Var}_\pi(f \,|\, \mathcal{H}, \omega_3) + \mathbb{1}_{\mathcal{K}} \operatorname{Var}_\pi(f \,|\, \mathcal{K}, \omega_1) \,|\, \mathcal{H} \cup \mathcal{K}\big).$$

Remark 5.3.4. It is easy to check that $\mathcal{D}_{\text{aux}}^{(1)}(f)$ is the Dirichlet form of the continuous time Markov chain on $\mathcal{H} \cup \mathcal{K}$ in which if $\omega \in \mathcal{H}$ the pair (ω_1, ω_2) is resampled with rate one from $\pi_1 \otimes \pi_2(\cdot \,|\, \mathcal{B}_{1,2}^{\omega_3})$ and if $\omega \in \mathcal{K}$ the pair (ω_2, ω_3) is resampled with rate one from $\pi_2 \otimes \pi_3(\cdot \,|\, \mathcal{B}_{2,3}^{\omega_1})$. This chain is reversible w.r.t. $\pi(\cdot \,|\, \mathcal{H} \cup \mathcal{K})$ and its constraints, contrary to what happens for general KCM, depend on the to-be-updated variables.

Proposition 5.3.5. *There exists a universal constant c such that the following holds. Suppose that there exist two events $\mathcal{F}_{1,2}$, $\mathcal{F}_{2,3}$ such that*

$$\{\omega : \omega_3 \in \mathcal{A}_3 \text{ and } (\omega_1, \omega_2) \in \mathcal{F}_{1,2}\} \subset \mathcal{H} \cap \mathcal{K}, \tag{5.26}$$

$$\{\omega : \omega_1 \in \mathcal{A}_1 \text{ and } (\omega_2, \omega_3) \in \mathcal{F}_{2,3}\} \subset \mathcal{H} \cap \mathcal{K} \tag{5.27}$$

and let

$$T_{\text{aux}}^{(1)} = \max_{\omega_3 \in \mathcal{A}_3,} \left(\frac{\pi(\mathcal{B}_{1,2}^{\omega_3})}{\pi(\mathcal{F}_{1,2})}\right)^2 \max_{\omega_1 \in \mathcal{A}_1} \frac{\pi(\mathcal{B}_{2,3}^{\omega_1})}{\pi(\mathcal{F}_{2,3})}.$$

Then, for all $f : \mathcal{H} \cup \mathcal{K} \to \mathbb{R}$,

$$\operatorname{Var}_\pi(f \,|\, \mathcal{H} \cup \mathcal{K}) \leqslant c T_{\text{aux}}^{(1)} \mathcal{D}_{\text{aux}}^{(1)}(f).$$

Proof. Consider the Markov chain $(\omega(t))_{t \geqslant 0}$ determined by the Dirichlet form $\mathcal{D}_{\text{aux}}^{(1)}$ as described in Remark 5.3.4. Given two arbitrary initial conditions $\omega(0)$ an $\omega'(0)$ we will construct a coupling of the two chains such that with probability $\Omega(1)$ we have $\omega(t) = \omega'(t)$ for any $t > T_{\text{aux}}^{(1)}$. Standard arguments [256] then prove that the mixing time of the chain is $O(T_{\text{aux}}^{(1)})$ and the conclusion of the proposition follows. To construct our coupling we use the following representation of the Markov chain.

We are given two independent Poisson clocks with rate one and the chain transitions occur only at the clock rings. Suppose that the first clock rings. If the current configuration ω does not belong to $\mathcal{H}$ the ring is ignored. Otherwise, a Bernoulli variable ξ with probability of success $\pi(\mathcal{F}_{1,2} \,|\, \mathcal{B}_{1,2}^{\omega_3})$ is sampled. If $\xi = 1$, then the pair (ω_1, ω_2) is resampled w.r.t. the measure

$\pi(\cdot \mid \mathcal{F}_{1,2}, \mathcal{B}_{1,2}^{\omega_3})$, while if $\xi = 0$, then (ω_1, ω_2) is resampled w.r.t. the measure $\pi(\cdot \mid \mathcal{F}_{1,2}^c, \mathcal{B}_{1,2}^{\omega_3})$. Clearly, in doing so the couple (ω_1, ω_2) is resampled w.r.t. $\pi(\cdot \mid \mathcal{B}_{1,2}^{\omega_3})$. Similarly if the second clock rings but with $\mathcal{H}$, (ω_1, ω_2), $\mathcal{F}_{1,2}$ and $\mathcal{B}_{1,2}^{\omega_3}$ replaced by $\mathcal{K}$, (ω_2, ω_3), $\mathcal{F}_{2,3}$ and $\mathcal{B}_{2,3}^{\omega_1}$ respectively. It is important to notice that $\pi(\cdot \mid \mathcal{F}_{1,2}, \mathcal{B}_{1,2}^{\omega_3}) = \pi(\cdot \mid \mathcal{F}_{1,2})$ for all $\omega_3 \in \mathcal{A}_3$, as, by assumption, $\mathcal{F}_{1,2} \subset \bigcap_{\omega_3 \in \mathcal{A}_3} \mathcal{B}_{1,2}^{\omega_3}$. Similarly, $\pi(\cdot \mid \mathcal{F}_{2,3}, \mathcal{B}_{2,3}^{\omega_1}) = \pi(\cdot \mid \mathcal{F}_{2,3})$ for all $\omega_1 \in \mathcal{A}_1$.

In our coupling both chains use the same clocks. Suppose that the first clock rings and that the current pair of configurations is (ω, ω'). Assume also that at least one of them, say ω, is in $\mathcal{H}$ (otherwise, both remain unchanged). In order to construct the coupling update we proceed as follows.

- If $\omega' \notin \mathcal{H}$ then ω is updated as described above, while ω' stays still.
- If $\omega' \in \mathcal{H}$ we first maximally couple the two Bernoulli variables ξ, ξ' corresponding to ω, ω' respectively. Then:
  - if $\xi = \xi' = 1$, we update both (ω_1, ω_2) and (ω_1', ω_2') to the *same* couple $(\eta_1, \eta_2) \in \mathcal{F}_{1,2}$ with probability $\pi((\eta_1, \eta_2) \mid \mathcal{F}_{1,2})$;
  - otherwise we resample (ω_1, ω_2) and (ω_1', ω_2') independently from their respective law given ξ, ξ'.

Similarly if the ring comes from the second clock. The final coupling is then equal to the Markov chain on $\Omega \times \Omega$ with the transition rates described above. Suppose now that there are three consecutive rings occurring at times $t_1 < t_2 < t_3$ such that:

- the first and last ring come from the first clock while the second ring comes from the second clock, and
- the sampling of the Bernoulli variables (if any) at times t_1, t_2 and t_3 all produce the value one.

Then we claim that at time t_3 the two copies are coupled.

To prove the claim, we begin by observing that after the first update at t_1 both copies of the coupled chain belong to $\mathcal{K}$. Here we use (5.26). Indeed, if the first update is successful for ω (*i.e.* $\omega \in \mathcal{H}$) then the updated configuration belongs to $\mathcal{F}_{1,2} \times \{\omega_3\} \subset \mathcal{K}$, because of our assumption $\xi = 1$. If, on the contrary, the first update fails (*i.e.* $\omega \notin \mathcal{H}$) then $\omega \in \mathcal{K} \backslash \mathcal{H}$ before and after the update. The same applies to ω'.

Next, using again the assumption on the Bernoulli variables together with the previous observation, we get that after the second ring the new pair of current configurations agree on the second and third coordinate. Moreover both copies belong to $\mathcal{H}$ thanks to (5.27). Finally, after the third ring the two copies couple on the first and second coordinates using again the assumption on the outcome for the Bernoulli variables.

In order to conclude the proof of the proposition it is enough to observe that for any given time interval Δ of length one the probability that there exist $t_1 < t_2 < t_3$ in Δ satisfying the requirements of the claim is bounded from below by

$$c \min_{\omega_3 \in \mathcal{A}_3} \pi\big(\mathcal{F}_{1,2} | \mathcal{B}_{1,2}^{\omega_3}\big)^2 \min_{\omega_1 \in \mathcal{A}_1} \pi\big(\mathcal{F}_{2,3} | \mathcal{B}_{2,3}^{\omega_1}\big),$$

for some constant $c > 0$. □

In the same setting consider two other events $\mathcal{C}_{1,2} \subset \Omega_1 \otimes \Omega_2$, $\mathcal{C}_{2,3} \subset \Omega_2 \otimes \Omega_3$ and let

$$\mathcal{M} = \mathcal{A}_3 \cap \mathcal{C}_{1,2}, \qquad \mathcal{N} = \mathcal{A}_1 \cap \mathcal{C}_{2,3}.$$

The Dirichlet form of our second Markov chain on $\mathcal{M} \cup \mathcal{N}$ is then

$$\mathcal{D}_{\text{aux}}^{(2)}(f) = \pi\Big(\mathbb{1}_{\mathcal{M}} \operatorname{Var}(f \,|\, \mathcal{C}_{1,2}, \omega_3) + \mathbb{1}_{\mathcal{M}} \operatorname{Var}(f \,|\, \mathcal{A}_3, \omega_1, \omega_2) + \mathbb{1}_{\mathcal{N}} \operatorname{Var}(f \,|\, \mathcal{C}_{2,3}, \omega_1) + \mathbb{1}_{\mathcal{N}} \operatorname{Var}(f \,|\, \mathcal{A}_1, \omega_2, \omega_3) \,|\, \mathcal{M} \cup \mathcal{N}\Big). \quad (5.28)$$

Remark 5.3.6. Similarly to the first case, the continuous time chain defined by (5.28) is reversible w.r.t. $\pi(\cdot \,|\, \mathcal{M} \cup \mathcal{N})$ and it can be described as follows. If $\omega \in \mathcal{M}$ then with rate one (ω_1, ω_2) is resampled w.r.t. $\pi_1 \otimes \pi_2(\cdot \,|\, \mathcal{C}_{1,2})$ and, independently at unit rate, ω_3 is resampled w.r.t. $\pi_3(\cdot \,|\, \mathcal{A}_3)$. Similarly, independently from the previous updates at rate one, if $\omega \in \mathcal{N}$ then (ω_2, ω_3) is resampled w.r.t. $\pi_2 \otimes \pi_3(\cdot \,|\, \mathcal{C}_{2,3})$ and, independently, ω_1 is resampled from $\pi_1(\cdot \,|\, \mathcal{A}_1)$.

Proposition 5.3.7. *There exists a universal constant c such that the following holds. Suppose that there exist an event $\hat{\mathcal{C}}_{1,2} \subset \mathcal{C}_{1,2}$ and a collection $(\mathcal{A}_3^{\omega_1,\omega_2})_{(\omega_1,\omega_2) \in \hat{\mathcal{C}}_{1,2}}$ of subsets of $\mathcal{A}_3$ such that*

$$\big\{\omega : (\omega_1, \omega_2) \in \hat{\mathcal{C}}_{1,2} \text{ and } \omega_3 \in \mathcal{A}_3^{\omega_1,\omega_2}\big\} \subset \mathcal{M} \cap \mathcal{N}, \quad (5.29)$$

and let

$$T_{\text{aux}}^{(2)} = \max_{(\omega_1,\omega_2) \in \hat{\mathcal{C}}_{1,2}} \frac{\pi(\mathcal{A}_3)}{\pi(\mathcal{A}_3^{\omega_1,\omega_2})} \times \frac{\pi(\mathcal{C}_{1,2})}{\pi(\hat{\mathcal{C}}_{1,2})}.$$

Then there exists $c > 0$ such that for all $f : \mathcal{M} \cup \mathcal{N} \to \mathbb{R}$,

$$\operatorname{Var}(f \,|\, \mathcal{M} \cup \mathcal{N}) \leqslant c T_{\text{aux}}^{(2)} \mathcal{D}_{\text{aux}}^{(2)}(f).$$

Proof. We proceed as in the proof of Proposition 5.3.5 with the following representation for the Markov chain. We are given four independent Poisson clocks of rate one and each clock comes equipped with a collection of i.i.d. random variables. The four independent collections, the first being for the first clock, etc., are

$$\big((\omega_1^{(i)}, \omega_2^{(i)})\big)_{i=1}^{\infty}, \quad \big(\eta_3^{(i)}\big)_{i=1}^{\infty}, \quad \big((\omega_2^{(i)}, \omega_3^{(i)})\big)_{i=1}^{\infty}, \quad \big(\eta_1^{(i)}\big)_{i=1}^{\infty},$$

where the laws of the collections are $\pi_1 \otimes \pi_2(\cdot \mid \mathcal{C}_{1,2})$, $\pi_3(\cdot \mid \mathcal{A}_3)$, $\pi_2 \otimes \pi_3(\cdot \mid \mathcal{C}_{2,3})$ and $\pi_1(\cdot \mid \mathcal{A}_1)$ respectively.

At each ring of the first and second clocks the configuration is updated with the variables from the corresponding collection iff $\omega \in \mathcal{M}$. Similarly for the third and fourth clocks with $\mathcal{N}$. In order to couple different initial conditions, we use the same collections of clock rings and update configurations.

Suppose now that there are four consecutive rings $t_1 < t_2 < t_3 < t_4$, coming from the first, second, third and forth clocks in that order, such that:

- at t_1 the proposed update (η_1, η_2) of the first two coordinates belongs to $\hat{\mathcal{C}}_{1,2}$, and
- at t_2 the proposed update η_3 of the third coordinate belongs to $\mathcal{A}_3^{(\eta_1,\eta_2)}$.

We then claim that after t_4 all initial conditions ω are coupled. To prove this, we first observe that after the second ring each chain belongs to $\mathcal{N}$. Indeed, if $\omega \notin \mathcal{M}$, then the first two proposed updates are ignored and the configuration $\omega \in \mathcal{N} \setminus \mathcal{M}$. If, on the contrary, $\omega \in \mathcal{M}$, then both updates are successful and the configuration is updated to $(\eta_1, \eta_2, \eta_3) \in \hat{\mathcal{C}}_{1,2} \times \mathcal{A}_3^{\eta_1,\eta_2} \subset \mathcal{M} \cap \mathcal{N}$ by (5.29).

Since after t_2 the state of the chain is necessarily in $\mathcal{N}$, the third and fourth updates to states (η'_2, η'_3) and η'_1 respectively are both successful and thus any initial condition leads to the state $(\eta'_1, \eta'_2, \eta'_3)$ after t_4, which proves the claim. The proof is then completed as in Proposition 5.3.5. □

5.4 Proof of Theorem 1.4.6: upper bound

The starting point is as in Section 4.4. Let κ be a large enough constant, let

$$t_* = \exp\left(\frac{\pi^2}{9q}\left(1 + \kappa\sqrt{q}\log^3(1/q)\right)\right) \tag{5.30}$$

and let $T = \lfloor \exp(\log^4(1/q)/q) \rfloor$. Then

$$\begin{aligned}
\mathbb{E}_\mu(\tau_0) &= \int_0^{+\infty} ds\, \mathbb{P}_\mu(\tau_0 > s) \\
&= \int_0^{t_*} ds\, \mathbb{P}_\mu(\tau_0 > s) + \int_{t_*}^{T} ds\, \mathbb{P}_\mu(\tau_0 > s) + \int_T^{+\infty} ds\, \mathbb{P}_\mu(\tau_0 > s) \\
&\leqslant t_* + T\mathbb{P}_\mu(\tau_0 > t_*) + \int_T^{+\infty} ds\, \mathbb{P}_\mu(\tau_0 > s).
\end{aligned}$$

The term t_* has exactly the form required in (1.15). The last term in the r.h.s. above tends to zero as $q \to 0$ if c is large enough. Indeed, using [269, Theorem 2] (see also [269] and Chapter 4) we have that for any $s \geqslant 0$

$\mathbb{P}_\mu(\tau_0 > s) \leqslant e^{-s\lambda_0}$ with $\lambda_0 \geqslant e^{-\Omega((\log q)^3/q)}$. In conclusion, the proof of (1.15) boils down to proving

$$\lim_{q\to 0} T\mathbb{P}_\mu(\tau_0 > t_*) = 0. \tag{5.31}$$

The key ingredients to prove (5.31) are Propositions 5.2.7 and 5.2.9 and Proposition 5.4.1 below. The latter is a Poincaré inequality for an auxiliary process, the *generalised coalescencing and branching symmetric exclusion process* (g-CBSEP), preliminarily studied in Chapter 3. Once we have these key ingredients, the strategy to prove (5.31) is similar to the one in Section 4.4. In particular, for the first part of the proof (Section 5.4.2) we will omit most of the details and refer to Section 4.4 for a more detailed explanation.

5.4.1 The g-CBSEP process

Given a finite connected graph $G = (V, E)$ and a finite probability space $(\mathcal{S}, \pi)$, assign a variable $\sigma_x \in \mathcal{S}$ to each vertex $x \in V$ and write $\sigma = (\sigma_x)_{x\in V}$ and $\pi_G(\sigma) = \prod_x \pi(\sigma_x)$. Fix also a bipartition $\mathcal{S}_1 \sqcup \mathcal{S}_0 = \mathcal{S}$ such that $\pi(\mathcal{S}_1) > 0$ and define the projection $\varphi : \mathcal{S}^V \to \{0,1\}^V$ by $\varphi(\sigma) = (\mathbb{1}_{\{\sigma_x\in\mathcal{S}_1\}})_{x\in V}$. We will say that a vertex x is occupied by a *particle* if $\sigma_x \in \mathcal{S}_1$ and we will write $\Omega_G^+ \subset \Omega_G = \mathcal{S}^V$ for the set of configurations σ with at least one particle. Finally, for any edge $e = \{x, y\} \in E$ let $\mathcal{E}_e$ be the event that there exists a particle at x or at y.

The g-CBSEP continuous time Markov chain on Ω_G^+ with parameters $(\mathcal{S}, \mathcal{S}_1, \pi)$ runs as follows. The state $\{\sigma_x, \sigma_y\}$ of every edge $e = \{x, y\}$ for which $\mathcal{E}_e$ holds is resampled with rate one (independently of all the other edges) w.r.t. $\pi_x \otimes \pi_y(\cdot \,|\, \mathcal{E}_e)$. Thus, an edge containing exactly one particle can swap the position of the particle between its endpoints or can create a new particle at the empty endpoint (a branching transition). An edge with two particles can kill one of them (a coalescing transition) with equal probability or keep them untouched. It is immediate to check that g-CBSEP is ergodic on Ω_G^+ with reversible stationary measure $\pi_G^+ := \pi_G(\cdot \,|\, \Omega_G^+)$ and that its Dirichlet form $\mathcal{D}^{g\text{-CBSEP}}(f)$ for $f : \Omega_G^+ \to \mathbb{R}$, takes the form

$$\mathcal{D}^{g\text{-CBSEP}}(f) = \sum_{e\in E} \pi_G^+\big(\mathbb{1}_{\mathcal{E}_e} \operatorname{Var}_e(f \,|\, \mathcal{E}_e)\big),$$

where $\operatorname{Var}_e(f \,|\, \mathcal{E}_e)$ is the variance w.r.t. σ_x, σ_y conditioned on $\mathcal{E}_e$ if $e = \{x, y\}$. Let now $T_{\text{rel}}^{g\text{-CBSEP}}$ be the *relaxation time* of g-CBSEP on Ω_G^+ defined as the best constant C in the Poincaré inequality

$$\operatorname{Var}_{\pi_G^+}(f) \leqslant C\mathcal{D}^{g\text{-CBSEP}}(f).$$

In the above setting the main result needed to prove (5.31) is as follows. For any positive integers d and L set $n = L^d$ and let $\mathbb{Z}_L = \{0, 1, \ldots, L-1\}$ be the set of remainders modulo L. The d- dimensional discrete torus with n vertices, $\mathbb{T}_n^d$ in the sequel, is the set $\mathbb{Z}_L^d$ endowed with the graph structure inherited from $\mathbb{Z}^d$.

Proposition 5.4.1. *Let $G = \mathbb{T}_n^d$ and assume that $\mathcal{S}$, $\mathcal{S}_1$ and π depend on n in such a way that* $\lim_{n\to\infty} \pi(\mathcal{S}_1) = 0$ *and* $\lim_{n\to\infty} n\pi(\mathcal{S}_1) = +\infty$. *Then, as $n \to \infty$,*

$$T_{\mathrm{rel}}^{g\text{-}CBSEP} \leqslant O\big(\pi(\mathcal{S}_1)^{-1} \log\big(\pi(\mathcal{S}_1)^{-1}\big)\big).$$

5.4.2 Transforming (5.31) into a Poincaré inequality

Using standard "finite speed of propagation" bounds (see Step (a) in Section 4.4.2) it is enough to prove (5.31) for FA-2f on the discrete torus $\mathbb{T}_n^2$ with linear size $\sqrt{n} = 2T$. Next we fix a small positive constant $\delta < 1/2$ and choose $N_\delta = N - \lfloor \log(1/\delta)/\sqrt{q} \rfloor$ where $N = \lceil \frac{8\log(1/q)}{\sqrt{q}} \rceil$ is the final scale in the droplet construction (see (5.5)). With this choice $\ell_{N_\delta} \simeq \delta \ell_N = \delta/q^{17/2+o(1)}$ (cf. (5.3)) and w.l.o.g. we assume that ℓ_{N_δ} divides $2T$.

We partition the torus $\mathbb{T}_n^2$ into $M = n/\ell_{N_\delta}^2$ equal mesoscopic disjoint boxes $(Q_j)_{j=1}^M$, where each Q_j is a suitable lattice translation by a vector in $\mathbb{T}_n^2$ of the box $Q = [\ell_{N_\delta}]^2 = \Lambda^{(2N_\delta)}$ (see (5.4)). The labels of the boxes can be thought of as belonging to the new torus $\mathbb{T}_M^2$ and we assume that Q_i, Q_j are neighbouring boxes in $\mathbb{T}_n^2$ iff i, j are neighbouring sites in $\mathbb{T}_M^2$. In $\Omega_{\mathbb{T}_n^2}$ we consider the event

$$\mathcal{E} = \bigcup_{j\in\mathbb{T}_M^2} \mathcal{SG}_j \cap \bigcap_{i\in\mathbb{T}_M^2} \mathcal{G}_i$$

where $\mathcal{SG}_i$ is the event that Q_i is super-good (see Definition 5.2.6) and $\mathcal{G}_i$ is the event that any row and any column (of lattice sites) of Q_i contains an infected site.

In order to apply the same strategy as Section 4.4 it is crucial to have that the "environment" characterised by $\mathcal{E}$ is so likely that (cf. (4.24))

$$\lim_{q\to 0} \mu(\mathcal{E}^c) T^3 t_* = 0. \tag{5.32}$$

It is such a requirement that guided us in the choice of the side length ℓ_{N_δ} of the mesoscopic boxes. Using Proposition 5.2.7 together with trivial bounds on the probability that a row/column of a box Q_i does not contain an infection it is immediate to verify (5.32). An easy consequence of (5.32) (see Step (c) in Section 4.4.2) is that it is sufficient to prove (5.31) for the stationary FA-2f in $\mathbb{T}_n^2$ restricted to $\mathcal{E}$. For the latter process we follow the standard "variational" approach (see Step (d) in Section 4.4.2) and the overall conclusion is that

$$T\mathbb{P}_\mu(\tau_0 \geqslant t_*) \leqslant T e^{-t_* \lambda_{\mathcal{F}}} + o(1),$$

with

$$\lambda_{\mathcal{F}} \geqslant \inf_f q \frac{\mathcal{D}_{\mathbb{T}_n^2}(f)}{\mathrm{Var}_{\mathbb{T}_n^2}(f \mid \mathcal{E})},$$

where $\mathcal{D}_{\mathbb{T}_n^2}(f)$ is the Dirichlet form of FA-2f on the torus $\mathbb{T}_n^2$ and the supremum is taken over all $f : \mathcal{E} \to \mathbb{R}$.

5.4.3 Bounding $\lambda_{\mathcal{F}}$ from below

The last and most important step is to prove that

$$\lambda_{\mathcal{F}} \geqslant e^{-O(\log^3(1/q)/\sqrt{q})} \rho_{\mathrm{D}}, \tag{5.33}$$

where $\rho_{\mathrm{D}} \geqslant \exp(-\frac{\pi^2}{9q}(1 + O(\sqrt{q}\log^2(1/q))))$ is the probability that a box $[\ell_N]^2$ is super-good (cf. Proposition 5.2.7). Once (5.33) is established, the proof of (5.31) is complete because $t_* \lambda_{\mathcal{F}}$ diverges rapidly enough as $q \to 0$ if the constant κ in the definition (5.30) of t_* is chosen large enough.

The proof of (5.33) is crucially based on Propositions 5.2.9 and 5.4.1 and is divided into two parts.

Application of Proposition 5.4.1 Write G for the graph $\mathbb{T}_M^2$, $\mathcal{S}$ for the state space $\mathcal{G}_i \subset \Omega_{Q_i}$ with $i \in \mathbb{T}_M^2$, π for $\mu_{Q_i}(\cdot \mid \mathcal{G}_i)$ and $\mathcal{S}_1 \subset \mathcal{S}$ for the event $\mathcal{SG}_i \cap \mathcal{G}_i$. Then, using $\lim_{q\to 0} \mu(\mathcal{G}_i) = 1$ and Proposition 5.2.7, it is easy to check that

$$\pi(\mathcal{S}_1) = \exp\Big(-\frac{\pi^2}{9q}\big(1 + O\big(\sqrt{q}\log^2(1/q)\big)\big)\Big).$$

Recalling $M = n/\ell_{N_\delta}^2$ with $n = 4T^2$, $\lim_{q\to 0} M\rho(S_1) = +\infty$ and the requirement of Proposition 5.4.1 holds.

With this notation we consider the g-CBSEP on Ω_G^+ with parameters $(\mathcal{S}, \mathcal{S}_1, \pi)$ and we identify any function $f : \mathcal{E} \to \mathbb{R}$ with a function $f_G : \Omega_G^+ \to \mathbb{R}$ via the obvious bijection between $\mathcal{E}$ and Ω_G^+: $f(\omega) = f_G(\omega_{Q_1}, \dots, \omega_{Q_M})$. Under this bijection

$$\begin{aligned}
\mathrm{Var}_{\pi_G^+}(f_G) &= \mathrm{Var}_{\mathbb{T}_n^2}(f \mid \mathcal{E}), \\
\mathcal{D}^{g\text{-CBSEP}}(f_G) &= \sum_{i\sim j} \mu_{\mathbb{T}_n^2}\big(\mathbb{1}_{\mathcal{SG}_{i,j}} \mathrm{Var}_{Q_i\cup Q_j}(f \mid \mathcal{SG}_{i,j}) \mid \mathcal{E}\big),
\end{aligned}$$

where $\mathcal{SG}_{i,j}$ is a shorthand notation for the event $(\mathcal{SG}_i \cup \mathcal{SG}_j) \cap \mathcal{G}_i \cap \mathcal{G}_j$ and $\sum_{i\sim j}$ denotes the sum over pairs, each counted once, of adjacent boxes. Using Proposition 5.4.1 we conclude that

$$\begin{aligned}
\mathrm{Var}_{\mathbb{T}_n^2}(f \mid \mathcal{E}) &= \mathrm{Var}_{\pi_G^+}(f_G) \leqslant O(\pi(\mathcal{S}_1)^{-1}\log(1/\pi(\mathcal{S}_1)))\mathcal{D}^{g\text{-CBSEP}}(f_G) \\
&= \exp\Big(\frac{\pi^2}{9q}(1 + O(\sqrt{q}\log^2(1/q)))\Big) \qquad (5.34) \\
&\quad \times \sum_{i\sim j} \mu_{\mathbb{T}_n^2}\big(\mathbb{1}_{\mathcal{SG}_{i,j}} \mathrm{Var}_{Q_i\cup Q_j}(f \mid \mathcal{SG}_{i,j}) \mid \mathcal{E}\big).
\end{aligned}$$

Application of Proposition 5.2.9 We next compare the sum appearing in the r.h.s. of (5.34) to the FA-2f Dirichlet form $\mathcal{D}_{\mathbb{T}_n^2}(f \mid \mathcal{E})$ and prove that the "comparison cost" is at most $\exp\big(O\big(\log^3(1/q)/\sqrt{q}\big)\big)$, so sub-leading w.r.t. the main term $\exp(\frac{\pi^2}{9q})$ above.

Lemma 5.4.2.

$$\sum_{i\sim j} \mu_{\mathbb{T}_n^2}\big(\mathbb{1}_{\mathcal{SG}_{i,j}} \operatorname{Var}_{Q_i\cup Q_j}(f \mid \mathcal{SG}_{i,j}) \mid \mathcal{E}\big) \leqslant e^{O(\log^3(1/q)/\sqrt{q})} \sum_{x\in\mathbb{T}_n^2} \mu_{\mathbb{T}_n^2}\big(c_x^{\mathbb{T}_n^2} \operatorname{Var}_x(f)\big),$$

where $c_x^{\mathbb{T}_n^2}$ is the FA-2f constraint at x for the torus $\mathbb{T}_n^2$ (see (5.8)*).*

Remark 5.4.3. As it will be clear from the proof, we actually prove a stronger statement, namely the constraint $c_x^{\mathbb{T}_n^2}$ above will appear multiplied by the indicator that x belongs to a droplet. While for many choices of f the presence of this additional constraint may completely change the average $\mu_{\mathbb{T}_n^2}(c_x^{\mathbb{T}_n^2} \operatorname{Var}_x(f))$, it is possible to exhibit choices of f, for which $\mathbb{1}_{\{x \text{ belongs to a "droplet"}\}} c_x^{\mathbb{T}_n^2} \operatorname{Var}_x(f) \simeq c_x^{\mathbb{T}_n^2} \operatorname{Var}_x(f)$.

Proof of Lemma 5.4.2. The lemma follows by summing Claim 5.4.4. □

Claim 5.4.4. Fix two adjacent boxes Q_i, Q_j and let $\Lambda_{i,j} \supset Q_i \cup Q_j$ be a translate of the box $\Lambda^{(2N)}$. Then

$$\mu_{\mathbb{T}_n^2}\big(\mathbb{1}_{\mathcal{SG}_{i,j}} \operatorname{Var}_{Q_i\cup Q_j}(f \mid \mathcal{SG}_{i,j}) \mid \mathcal{E}\big) \leqslant e^{O(\log^3(1/q)/\sqrt{q})} \sum_{x\in\Lambda_{i,j}} \mu_{\mathbb{T}_n^2}\big(\mathbb{1}_{\mathcal{SG}(\Lambda_{i,j})} c_x^{\mathbb{T}_n^2} \operatorname{Var}_x(f)\big).$$

Proof Claim 5.4.4. Let $\mathcal{G} = \bigcap_{k\in\mathbb{T}_M^2} \mathcal{G}_k \supset \mathcal{E}$ and recall that $\mu(\mathcal{E}) = 1 - o(1)$. Next, let $\mathcal{G}(\Lambda_{i,j})$ be the event that any ℓ_{N_δ} lattice sites contained in $\Lambda_{i,j}$ forming either a row or a column of some Q_k contain an infection.

Observation 5.4.5. The event $\mathcal{SG}_{i,j} \cap \mathcal{G}(\Lambda_{i,j})$ implies the event $\mathcal{SG}(\Lambda_{i,j})$.

The formal proof of this observation, as illustrated in Figure 5.5 is left to the reader.

Write $\rho_{i,j}$ for $\mu(\mathcal{SG}_{i,j} \mid \mathcal{G})$ and observe that the term $\operatorname{Var}_{Q_i\cup Q_j}(\cdot)$ does not depend on the variables $\omega_{Q_i}, \omega_{Q_j}$. This fact together with the Cauchy-

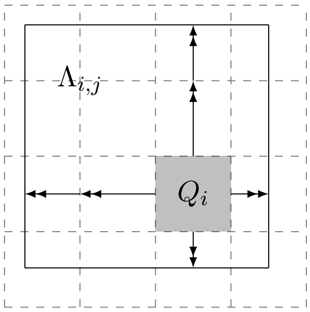

Figure 5.5 – Illustration of Observation 5.4.5. The shaded square of shape $\Lambda^{(2N_\delta)}$ is $\mathcal{SG}$ and the arrows indicate the presence of an infection in each row/column, as guaranteed by $\mathcal{G}(\Lambda_{i,j})$ with $\Lambda_{i,j}$ being the larger square of shape $\Lambda^{(2N)}$. Observation 5.4.5 asserts that these events combined imply $\mathcal{SG}(\Lambda_{i,j})$ (see Figure 5.2).

Schwartz inequality allows us to write

$$\begin{aligned}
&\mu_{\mathbb{T}_n^2}\big(\mathbb{1}_{\mathcal{SG}_{i,j}}\operatorname{Var}_{Q_i\cup Q_j}(f\,|\,\mathcal{SG}_{i,j})\,|\,\mathcal{E}\big)\\
&\quad= (1+o(1))\mu_{\mathbb{T}_n^2}\big(\mathbb{1}_{\mathcal{SG}_{i,j}}\operatorname{Var}_{Q_i\cup Q_j}(f\,|\,\mathcal{SG}_{i,j})\,|\,\mathcal{G}\big)\\
&\quad= (1+o(1))\rho_{i,j}\mu_{\mathbb{T}_n^2}\big(\mu_{\Lambda_{i,j}}\big(\operatorname{Var}_{Q_i\cup Q_j}(f\,|\,\mathcal{SG}_{i,j})\,|\,\mathcal{G}(\Lambda_{i,j})\big)\,|\,\mathcal{G}\big)\\
&\quad\leqslant (1+o(1))\rho_{i,j}\mu_{\mathbb{T}_n^2}\big(\operatorname{Var}_{\Lambda_{i,j}}(f\,|\,\mathcal{SG}_{i,j}\cap\mathcal{G}(\Lambda_{i,j}))\,|\,\mathcal{G}\big)\\
&\quad\leqslant (1+o(1))\rho_{i,j}\frac{\mu_{\mathbb{T}_n^2}(\mathcal{SG}(\Lambda_{i,j})\,|\,\mathcal{G})}{\mu_{\mathbb{T}_n^2}(\mathcal{SG}_{i,j}\cap\mathcal{G}(\Lambda_{i,j})\,|\,\mathcal{G})}\mu_{\mathbb{T}_n^2}\big(\operatorname{Var}_{\Lambda_{i,j}}(f\,|\,\mathcal{SG}(\Lambda_{i,j}))\,|\,\mathcal{G}\big)\\
&\quad= (1+o(1))\mu_{\mathbb{T}_n^2}(\mathcal{SG}(\Lambda_{i,j})\,|\,\mathcal{G})\mu_{\mathbb{T}_n^2}\big(\operatorname{Var}_{\Lambda_{i,j}}(f\,|\,\mathcal{SG}(\Lambda_{i,j}))\,|\,\mathcal{G}\big).
\end{aligned}$$

In the last inequality we used Observation 5.4.5 together with the inequality $\operatorname{Var}(X\,|\,A)\leqslant\mathbb{P}(B)/\mathbb{P}(A)\operatorname{Var}(X\,|\,B)$ valid for any (finite) random variable X and any two events $A\subset B$ with $\mathbb{P}(A)>0$. By applying now Proposition 5.2.9 to the term $\operatorname{Var}_{\Lambda_{i,j}}(f\,|\,\mathcal{SG}(\Lambda_{i,j}))$ and, using $c_x^{\Lambda_{i,j}}\leqslant c_x^{\mathbb{T}_n^2}$, we conclude that

$$\begin{aligned}
&\mu_{\mathbb{T}_n^2}\big(\mathbb{1}_{\mathcal{SG}_{i,j}}\operatorname{Var}_{Q_i\cup Q_j}(f\,|\,\mathcal{SG}_{i,j})\,|\,\mathcal{E}\big)\\
&\qquad\leqslant e^{O(\log^3(1/q)/\sqrt{q})}\mu_{\mathbb{T}_n^2}(\mathcal{SG}(\Lambda_{i,j})\,|\,\mathcal{G})\\
&\qquad\quad\times\sum_{x\in\Lambda_{i,j}}\mu_{\mathbb{T}_n^2}\big(\mu_{\Lambda_{i,j}}\big(c_x^{\Lambda_{i,j}}\operatorname{Var}_x(f)\,|\,\mathcal{SG}(\Lambda_{i,j})\big)\,|\,\mathcal{G}\big)\\
&\qquad= e^{O(\log^3(1/q)/\sqrt{q})}\sum_{x\in\Lambda_{i,j}}\mu_{\mathbb{T}_n^2}\big(\mathbb{1}_{\mathcal{SG}(\Lambda_{i,j})}c_x^{\mathbb{T}_n^2}\operatorname{Var}_x(f)\,|\,\mathcal{G}\big).
\end{aligned}$$

The proof is complete because $\mu_{\mathbb{T}_n^2}(\mathcal{G})=1-o(1)$. □

Appendix

5.A Probability of super-good events

In this appendix we prove Proposition 5.2.7 and we gather several more technical and relatively standard bootstrap percolation estimates on the probability of super-good events used in Section 5.2.

For $z > 0$ we define

$$g(z) = -\log\big(\beta(1-e^{-z})\big),$$

where $\beta(u) = (u + \sqrt{u(4-3u)})/2$. It is known [225, Proposition 5(ii)] that $\int_0^\infty g(z)\,dz = \pi^2/18$. We next recall some straightforward properties of g.

Fact 5.A.1. *The function g is positive, decreasing, differentiable and convex on $(0,\infty)$. Moreover, the following asymptotic behaviour holds:*

$$\begin{aligned} g(z) &\sim \frac{1}{2}\log(1/z), & g'(z) &\sim \frac{-1}{2z}, & &\text{as } z \to 0,\\ g(z) &\sim e^{-2z}, & g'(z) &\sim -2e^{-2z}, & &\text{as } z \to \infty, \end{aligned}$$

where $x \sim y$ stands for $x = (1+o(1))y$.

The relevance of this function comes from its link to the probability of traversability. Recalling Definition 5.2.1, for any positive integers a and b we set

$$T^{\mathbf{1}}(a,b) = \mu(\mathcal{T}^{\mathbf{1}}_{\rightarrow}(R(a,b))), \qquad T^{\mathbf{0}}(a,b) = \mu(\mathcal{T}^{\mathbf{0}}_{\rightarrow}(R(a,b))),$$

where $\mathbf{0}$ stands for the fully infected configuration. Note that these probabilities are the same for left-traversability, while for up or down-traversability a and b are inverted in the r.h.s. The next lemma follows easily from [225, Lemma 8]. Let $q' = -\log(1-q) = q + O(q^2)$.

Lemma 5.A.2. *For any positive integers a and b and $\omega \in \{\mathbf{0},\mathbf{1}\}$ we have*

$$T^\omega(a,b) = q^{O(1)}e^{-ag(bq')}.$$

Corollary 5.A.3. *For any positive integers a and b we have*

$$\max_{0\leqslant s,s'\leqslant b}\frac{T^{\mathbf{0}}(a,s)T^{\mathbf{0}}(a,b-s)}{T^{\mathbf{1}}(a,s')T^{\mathbf{1}}(a,b-s')} \leqslant q^{-O(1)}. \tag{5.35}$$

In particular, for any boundary conditions ω,ω' and rectangle R of class $1 \leqslant n \leqslant 2N$ with n odd, we have

$$\mu_R(\mathcal{SG}^\omega_s(R)\,|\,\mathcal{SG}^{\omega'}(R)) \geqslant q^{O(1)} \tag{5.36}$$

uniformly over all possible values of s and boundary conditions ω,ω' and similarly for even n.

Proof. (5.35) follows immediately from Lemma 5.A.2. To obtain (5.36), recall that

$$\mathcal{SG}^{\omega'}(R) = \bigcup_{s'}\mathcal{SG}^{\omega'}_{s'}(R);$$

there are $q^{-O(1)}$ possible values of s'; by (5.35), for all s, s', ω and ω',

$$\mu_R(\mathcal{SG}^\omega_s)/\mu_R(\mathcal{SG}^{\omega'}_{s'}) \geqslant q^{O(1)}.$$

□

We are now ready for the main result of this appendix.

Proof of Proposition 5.2.7. We will prove the same bound for the super-good event occurring with all $s = 0$ in Definition 5.2.4 on all scales, *i.e.* the initial infection $\Lambda^{(0)}$ being in the bottom-left corner of $\Lambda^{(n)}$. Once the offsets are fixed, it suffices to prove the bound on this probability for $n = 2N$, in which case it reads

$$q \prod_{m=1}^{N} T^{\mathbf{1}}(\ell_m - \ell_{m-1}, \ell_m) T^{\mathbf{1}}(\ell_m - \ell_{m-1}, \ell_{m-1}) \\ = q^{O(N)} \exp\Big(- \sum_{m=1}^{N} (\ell_m - \ell_{m-1})(g(q'\ell_m) + g(q'\ell_{m-1}))\Big),$$

by Lemma 5.A.2 and symmetry. Since g is decreasing, the last sum is at most

$$2 \sum_{m=1}^{\infty} (\ell_m - \ell_{m-1}) g(q'\ell_{m-1}).$$

The term for $m = 1$ is $O(\log(1/q)/\sqrt{q})$ by Fact 5.A.1. For the other terms we use that by convexity for any $0 < a < b$

$$(b - a)g(a) \leqslant \int_a^b g(z)\, dz - O((b - a)^2 g'(a)).$$

Using Fact 5.A.1, we get

$$-(b - a)^2 g'(a) \leqslant O((b - a))^2 \times \begin{cases} 1/a & \text{if } a = O(1) \\ e^{-a} & \text{if } a = \Omega(1). \end{cases}$$

Finally, we have $\ell_m - \ell_{m-1} \leqslant 2\sqrt{q}\ell_{m-1}$ by (5.3), so

$$q' \sum_{m=2}^{m_0} \frac{(\ell_m - \ell_{m-1})^2}{\ell_{m-1}} \leqslant O(q'\sqrt{q}\ell_{m_0}) = O(\sqrt{q}) \\ (q')^2 \sum_{m=m_0}^{\infty} (\ell_m - \ell_{m-1})^2 e^{-q'\ell_{m-1}} = O(\sqrt{q}),$$

setting $m_0 = \max\{m, \ell_m \leqslant 1/q\} = O(\log(1/q))/\sqrt{q})$. Putting these bounds together and recalling (5.5), we conclude the proof of Proposition 5.2.7. $\square$

We next turn to defining the event $\overline{\mathcal{SG}}(V_2)$ required in the proof of Lemma 5.2.12, so we fix $n = 2m \in [2, 2N)$ and $R = R(\ell_m - 1, \ell_m) = V_2 - \vec{e}_1$.

Definition 5.A.4. We say that $\overline{\mathcal{SG}}(R)$ occurs if there exist integers $0 \leqslant s_1 \leqslant \ell_m - \ell_{m-1} - 1$ and $0 \leqslant s_2 \leqslant \ell_m - \ell_{m-1}$ such that the intersection of the following events, in the sequel $\overline{\mathcal{SG}}_{s_1,s_2}(R)$, occurs (see Figure 5.4)

$$\begin{aligned}
&\mathcal{SG}(\Lambda^{(n-2)} + s_1\vec{e}_1 + s_2\vec{e}_2);\\
&\mathcal{T}_{\leftarrow}(R(s_1, \ell_{m-1}) + s_2\vec{e}_2);\\
&\mathcal{T}_{\rightarrow}(R(\ell_m - \ell_{m-1} - 1 - s_1, \ell_{m-1}) + (s_1 + \ell_{m-1})\vec{e}_1 + s_2\vec{e}_2);\\
&\mathcal{T}_{\downarrow}(R(\ell_m - 1, s_2));\\
&\mathcal{T}_{\uparrow}(R(\ell_m - 1, \ell_m - \ell_{m-1} - s_2) + (s_2 + \ell_{m-1})\vec{e}_2).
\end{aligned}$$

The event $\overline{\mathcal{SG}}(V_2)$ is defined by translation of $\overline{\mathcal{SG}}(R)$. Then for any $\omega_2 \in \overline{\mathcal{SG}}_{s_1,s_2}(V_2)$, the segments I_1 and I_3 are given by

$$\begin{aligned}
I_1(\eta_2) &= R(1, \ell_{m-1}) + s_2(\omega_2)\vec{e}_2 \subset V_1 = R(1, \ell_m),\\
I_3(\eta_2) &= R(1, \ell_{m-1}) + s_2(\omega_2)\vec{e}_2 + \ell_m\vec{e}_1 \subset V_3 = V_1 + \ell_m\vec{e}_1.
\end{aligned}$$

Lemma 5.A.5. *Recalling* (5.18)*, we have*

$$\mu_{V_1\cup V_2}(\hat{\mathcal{C}}_{1,2} \,|\, \mathcal{SG}) \geqslant q^{-O(1)}.$$

Proof. Recall that $V_1 \cup V_2 = \Lambda^{(n)}$ and assume $\mathcal{SG}(\Lambda^{(n)})$ occurs. For any $0 \leqslant s_1, s_2 \leqslant \ell_m - \ell_{m-1}$ we write

$$\mathcal{SG}_{s_1,s_2}(\Lambda^{(n)}) = \mathcal{SG}_{s_2}(\Lambda^{(n)}) \cap \mathcal{SG}_{s_1}(\Lambda^{(n-1)} + s_2\vec{e}_2).$$

Then by Corollary 5.A.3 for any such s_1, s_2 we have

$$\mu_{\Lambda^{(n)}}(\mathcal{SG}_{s_1,s_2}(\Lambda^{(n)})) = \mu_{\Lambda^{(n)}}(\mathcal{SG}(\Lambda^{(n)}))q^{O(1)},$$

so it suffices to show that

$$\mu_{V_2}(\overline{\mathcal{SG}}_{0,0}(V_2)) \geqslant \mu_{\Lambda^{(n)}}(\mathcal{SG}_{1,0}(\Lambda^{(n)}))q^{O(1)},$$

since $\mu(\mathcal{T}_{\leftarrow}(I_1(\eta_{V_2}))) \geqslant q$ for any $\omega_{V_2} \in \overline{\mathcal{SG}}(V_2)$.

However, by Definitions 5.2.4 and 5.A.4 and symmetry we have

$$\begin{aligned}
&\frac{\mu_{V_2}(\overline{\mathcal{SG}}_{0,0}(V_2))}{\mu_{\Lambda^{(n)}}(\mathcal{SG}_{1,0}(\Lambda^{(n)}))}\\
&\qquad= \frac{T^{\mathbf{1}}(\ell_m - \ell_{m-1} - 1, \ell_{m-1})T^{\mathbf{1}}(\ell_m - \ell_{m-1}, \ell_m - 1)}{T^{\mathbf{1}}(\ell_m - \ell_{m-1} - 1, \ell_{m-1})T^{\mathbf{1}}(\ell_m - \ell_{m-1}, \ell_m)T^{\mathbf{1}}(1, \ell_{m-1})}\\
&\qquad\geqslant \frac{T^{\mathbf{1}}(\ell_m - \ell_{m-1}, \ell_m - 1)}{T^{\mathbf{1}}(\ell_m - \ell_{m-1}, \ell_m)} = q^{O(1)}e^{-(\ell_m - \ell_{m-1})(g((\ell_m-1)q') - g(\ell_m q'))},
\end{aligned}$$

the last equality following from Lemma 5.A.2.

By convexity of g we get

$$g((\ell_m - 1)q') - g(\ell_m q') \leqslant -q' g'((\ell_m - 1)q'). \tag{5.37}$$

By Fact 5.A.1 we have that the r.h.s. of (5.37) is $O(1/\ell_m)$. Putting this together we obtain

$$\frac{\mu_{V_2}(\overline{\mathcal{SG}}_{0,0}(V_2))}{\mu_{\Lambda^{(n)}}(\mathcal{SG}_{1,0}(\Lambda^{(n)}))} \geqslant q^{O(1)} e^{-O(\ell_m - \ell_{m-1})/\ell_m} = q^{O(1)}, \tag{5.38}$$

as desired, the last equality coming from (5.3). □

5.B Proof of Proposition 5.4.1

Let $(\mathcal{S}, \mathcal{S}_1, \pi)$ be the parameters of the g-CBSEP on $\mathbb{T}_n^d$ and $\ell = \lceil \pi(\mathcal{S}_1)^{-1/d} \rceil$. For simplicity we assume that $n^{1/d}/\ell \in \mathbb{N}$ and we partition the torus $\mathbb{T}_n^d$ into $M = (n/\ell)^d$ equal boxes $(B_j)_{j=1}^M$, where each B_j is a suitable lattice translation by a vector in $\mathbb{T}_n^d$ of the box $B = [\ell]^d$. The labels of the boxes can be thought of as belonging to $\mathbb{T}_M^d$ and we assume that B_i, B_j are neighbouring boxes in $\mathbb{T}_n^d$ iff $i \sim j$ in $\mathbb{T}_M^d$.

We then set $\hat{\mathcal{S}} = \mathcal{S}^B, \hat{\pi}((\sigma_x)_{x \in B}) = \bigotimes_{x \in B} \pi(\sigma_x), \hat{\mathcal{S}}_1 = \bigcup_{x \in B} \{\sigma_x \in \mathcal{S}_1\}$ and we consider the auxiliary renormalized g-CBSEP (denoted $\hat{g}$-CBSEP in the sequel) on the graph $\hat{G} = \mathbb{T}_M^d$ with parameters $(\hat{\mathcal{S}}, \hat{\mathcal{S}}_1, \hat{\pi})$. Using the assumption $\lim_{n \to \infty} \pi(\mathcal{S}_1) = 0$, we have that $\hat{\pi}(\hat{\mathcal{S}}_1) = (1 - \pi(\mathcal{S}_1))^{\ell^d} \to 1/e$ as $n \to +\infty$.

Lemma 5.B.1. *Let $T_{\text{rel}}^{\hat{g}\text{-CBSEP}}$ be the relaxation time of $\hat{g}$-CBSEP on $\widehat{G}$. Then there exists a constant $C > 0$ depending on d such that $T_{\text{rel}}^{\hat{g}\text{-CBSEP}} \leqslant C$.*

Proof. This follows by comparison between g-CBSEP and generalised FA-1f analogous to (1.6) combined with [269, Proposition 3.5]. □

Proof of Proposition 5.4.1. For any pair of neighbouring boxes B_i and B_j write $\hat{\mathcal{E}}_{i,j}$ for the event $\bigcup_{x \in B_i \cup B_j} \{\sigma_x \in \mathcal{S}_1\}$. Using Lemma 5.B.1 and the definition of $T_{\text{rel}}^{\hat{g}\text{-CBSEP}}$ we get that that

$$\operatorname{Var}_{\pi^+_{\mathbb{T}_n^d}}(f) \leqslant C \sum_{i \sim j} \pi^+_{\mathbb{T}_n^d}\big(\mathbb{1}_{\hat{\mathcal{E}}_{i,j}} \operatorname{Var}_{B_i \cup B_j}(f \mid \hat{\mathcal{E}}_{i,j})\big),$$

where the sum in the r.h.s. is an equivalent way to express the Dirichlet form of $\hat{g}$-CBSEP. Now fix a pair of adjacent boxes B_i, B_j and let $T_{\text{rel}}^{g\text{-CBSEP}}(i,j)$ be the relaxation time of our original g-CBSEP with parameters $(\mathcal{S}, \mathcal{S}_1, \pi)$ on $B_i \cup B_j$. By symmetry $T_{\text{rel}}^{g\text{-CBSEP}}(i,j)$ does not depend on i, j and the common value will be denoted by $\widetilde{T}_{\text{rel}}$. If we plug the Poincaré inequality for g-CBSEP on $B_i \cup B_j$

$$\operatorname{Var}_{B_i \cup B_j}(f \mid \hat{\mathcal{E}}_{i,j}) \leqslant \widetilde{T}_{\text{rel}} \sum_{x \sim y \in B_i \cup B_j} \pi^+_{B_i \cup B_j}\big(\mathbb{1}_{\mathcal{E}_{x,y}} \operatorname{Var}_{x,y}(f \mid \mathcal{E}_{x,y})\big).$$

into the r.h.s. above of we get

$$\begin{aligned}\operatorname{Var}_{\pi^+_{\mathbb{T}^d_n}}(f) &\leqslant C\widetilde{T}_{\mathrm{rel}}\sum_{i\sim j}\sum_{x\sim y\in B_i\cup B_j}\pi^+_{\mathbb{T}^d_n}\left(\mathbb{1}_{\hat{\mathcal{E}}_{i,j}}\pi^+_{B_i\cup B_j}\left(\mathbb{1}_{\mathcal{E}_{x,y}}\operatorname{Var}_{x,y}(f\,|\,\mathcal{E}_{x,y})\right)\right)\\ &\leqslant Cd\widetilde{T}_{\mathrm{rel}}\sum_{x\sim y\in\mathbb{T}^d_n}\pi^+_{\mathbb{T}^d_n}\left(\mathbb{1}_{\mathcal{E}_{x,y}}\operatorname{Var}_{x,y}(f\,|\,\mathcal{E}_{x,y})\right)\\ &= Cd\widetilde{T}_{\mathrm{rel}}\mathcal{D}^{g\text{-CBSEP}}(f),\end{aligned}$$

i.e. $T^{g\text{-CBSEP}}_{\mathrm{rel}} \leqslant O\big(\widetilde{T}_{\mathrm{rel}}\big)$. It remains to bound $\widetilde{T}_{\mathrm{rel}}$ from above.

Let $T^{\mathrm{CBSEP}}_{\mathrm{mix}}$ be the mixing time of g-CBSEP on $G_{i,j}$ with parameters $\mathcal{S}' = \{0,1\}, \mathcal{S}'_1 = \{1\}$ and $\pi'(1) = \pi(\mathcal{S}_1) = 1-\pi'(0)$. Let $T^{\mathrm{rw}}_{\mathrm{cov}}$ be the cover time of the continuous-time random walk on $G_{i,j}$. Theorem 3.3.5 implies that $\widetilde{T}_{\mathrm{rel}} \leqslant O(T^{\mathrm{CBSEP}}_{\mathrm{mix}} + T^{\mathrm{rw}}_{\mathrm{cov}})$. It is well known (see *e.g.* [256]) that $T^{\mathrm{rw}}_{\mathrm{cov}} \leqslant O\big(\ell^d\log(\ell)\big) = O\big(\pi(\mathcal{S}_1)^{-1}\log(1/\pi(\mathcal{S}_1))\big)$ and Corollary 3.3.2 proves[2] the same bound for $T^{\mathrm{CBSEP}}_{\mathrm{mix}}$. In conclusion,

$$\widetilde{T}_{\mathrm{rel}} \leqslant O\big(\pi(\mathcal{S}_1)^{-1}\log(1/\pi(\mathcal{S}_1))\big). \qquad \square$$

[2]Strictly speaking Corollary 3.3.2 deals with the torus of cardinality $\pi(\mathcal{S}_1)^{-1}$ but the same proof extends to our case of the graph $G_{i,j}$.

Chapter 6

Refined universality for critical KCM: upper bounds

This chapter is based on [210], proving the upper bound of Theorem 1.6.4 for classes (b) and (d)-(g) (recall Section 1.6 and particularly Section 1.6.5). Before embarking on the proof, let us make a few remarks.

Remark 6.0.1. Firstly, for reasons of extremely technical nature, we do not provide a full proof of (the upper bound of) Theorem 1.6.4(e). More precisely, we prove it as stated for models with rules contained in the axes of the lattice. We also prove a fully general upper bound of the form

$$\exp\left(\frac{O(\log(1/q))\log\log\log(1/q)}{q^{\alpha}}\right).$$

Furthermore, with very minor modifications (see Remark 6.3.1), the error factor can be reduced from log log log to $\log_*$, where $\log_*$ denotes the number of iterations of the logarithm before the result becomes negative (the inverse of the tower function). Unfortunately, removing this minuscule error term requires further work, which we omit for the sake of concision. Instead, we provide a sketch of how to achieve this at the end of Section 6.3.1.

Remark 6.0.2. Secondly, although we will not do this, it is possible to circumvent the use of the core of Chapter 4 and establish the upper bounds of Theorem 1.6.4(c) and (e) independently. This approach has the merit of making all but Lemma 4.3.10 and Section 4.4 of Chapter 4 redundant in the proof of Theorem 1.6.4, a considerable gain. However, since the present chapter is already sufficiently involved, we have chosen to directly import all of Chapter 4 in Section 6.5.3.

Remark 6.0.3. Finally, for the sake of verifiablity, in Section 6.A.1 we have left the rather cumbersome approach to microscopic dynamics used to prove Lemma 6.A.1 originally in [210]. However, Lemma 6.A.1 can be seen as a direct corollary of the more elegant Theorem 2.1.1 developed for this purpose.

The chapter is organised as follows. In Section 6.2 we formally state the two fundamental techniques we use to move from one scale to the next, namely East-extensions and CBSEP-extensions, which import and generalise ideas of Chapter 5. They will be used in various combinations throughout the rest of the chapter. The proofs of the results about those extensions, including the microscopic dynamics are deferred to Section 6.A, since they are quite technical and do not require new ideas. Sections 6.3 and 6.4 are the core of this chapter and establish estimates on the relaxation times of "super good droplets" from microscopic scales up to and past the "critical" scale by means of East-extensions and CBSEP-extensions respectively. In Section 6.5 we recall, adapt and apply mechansims from Chapters 4 and 5, allowing us to go from "post-critical" to infinite scales. It is only at that point that we assemble Theorem 1.6.4 class by class from the general tools gathered in the previous sections. Finally, in Section 6.B we establish bounds on conditional probabilities, which, although technical and not particularly conceptual, serve a key role in both Sections 6.3 and 6.4. They establish general analogues of Section 5.A in an entirely new way and may be of independent interest for bootstrap percolation.

6.1 Preliminaries

6.1.1 Correlation inequalities

Let us recall two well-known correlation inequalities due to Harris [202] and van den Berg–Kesten [357]. The Harris inequality will be used throughout and we state some particular formulations that will be useful for us. The BK inequality is not natural to use for an upper bound in our setting and has not been employed to this purpose until now. Nevertheless, it will prove crucial in Section 6.B to estimate certain conditional probabilities.

For Section 6.1.1 we fix a finite $\Lambda \subset \mathbb{Z}^2$. We say that an event $\mathcal{A} \subset \Omega_\Lambda$ is *decreasing* if adding infections does not destroy its occurrence.

Proposition 6.1.1 (Harris inequality)**.** *Let $\mathcal{A}, \mathcal{B} \subset \Omega_\Lambda$ be decreasing. Then*

$$\mu(\mathcal{A} \cap \mathcal{B}) \geqslant \mu(\mathcal{A})\mu(\mathcal{B}).$$

We collectively refer to this proposition and the following corollaries as *Harris inequality.*

Corollary 6.1.2. *Let $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D} \subset \Omega_\Lambda$ be nonempty and decreasing events*[1] *such that $\mathcal{B}$ and $\mathcal{D}$ are independent, then*

$$\mu(\mathcal{A}|\mathcal{B} \cap \mathcal{D}) \geqslant \mu(\mathcal{A}|\mathcal{B}) \geqslant \mu(\mathcal{A}),$$
$$\mu(\mathcal{A} \cap \mathcal{C}|\mathcal{B} \cap \mathcal{D}) \geqslant \mu(\mathcal{A}|\mathcal{B})\mu(\mathcal{C}|\mathcal{D}).$$

[1] Though we will not mention this, events we consider are assumed nonempty.

Given two decreasing events $\mathcal{A}, \mathcal{B} \subset \Omega_\Lambda$, we say that $\mathcal{A}$ and $\mathcal{B}$ *occur disjointly* in $\omega \in \Omega_\Lambda$ if there exist disjoint sets $X, Y \subset \Lambda$, such that $\omega_{X \cup Y} = \mathbf{0}$; $\omega'_X = \mathbf{0}$ implies $\omega' \in \mathcal{A}$; and $\omega'_Y = \mathbf{0}$ implies $\omega' \in \mathcal{B}$ for $\omega' \in \Omega_\Lambda$.

Proposition 6.1.3 (BK inequality)**.** *Let $\mathcal{A}, \mathcal{B} \subset \Omega_\Lambda$ be decreasing events. Then*

$$\mu(\mathcal{A} \text{ and } \mathcal{B} \text{ occur disjointly}) \leqslant \mu(\mathcal{A})\mu(\mathcal{B}).$$

6.1.2 Directions

Throughout this chapter we fix a critical update family $\mathcal{U}$ with difficulty α. We call a direction $u \in S^1$ *rational* if $u\mathbb{R} \cap \mathbb{Z}^2 \neq \varnothing$. By the definition of α there exists a semicircle with rational midpoint u_0 such that all directions in the semicircle have difficulty at most α. We may assume without loss of generality that the direction $u_0 + \pi/2$ is hard unless $\mathcal{U}$ is isotropic. It is not difficult to show (see e.g. [74, Lemma 5.3]) that one can find a set $\mathcal{S}'$ of rational directions such that:

- all isolated and semi-isolated stable directions are in $\mathcal{S}'$;
- $u_0 \in \mathcal{S}'$;
- for every two consecutive directions u, v in $\mathcal{S}'$ there exists a rule $X \in \mathcal{U}$ such that $X \subset \overline{\mathbb{H}}_u \cap \overline{\mathbb{H}}_v$.

We further consider the set $\widehat{\mathcal{S}} = \mathcal{S}' + \{0, \pi/2, \pi, 3\pi/2\}$ obtained by making $\mathcal{S}'$ invariant by rotation by $\pi/2$. We will refer to the elements of $\widehat{\mathcal{S}}$ as *quasi-stable directions* or simply directions, as they are the only ones of interest to us. We label the elements of $\widehat{\mathcal{S}} = (u_i)_{i \in [4k]}$ clockwise and consider their indices modulo $4k$ (we write $[n]$ for $\{0, \dots, n-1\}$), so that $u_{i+2k} = -u_i$ is perpendicular to u_{i+k}. In figures we take $\widehat{\mathcal{S}} = \frac{\pi}{4}\mathbb{Z}$ and $u_0 = (-1, 0)$. Further observe that if all $U \in \mathcal{U}$ are contained in the axes of $\mathbb{Z}^2$, then $\widehat{\mathcal{S}} = \frac{\pi}{2}\mathbb{Z}$.

For $i \in [4k]$ we introduce $\rho_i = \min\{\rho > 0 : \exists x \in \mathbb{Z}^2, \langle x, u_i \rangle = \rho\}$ and $\lambda_i = \min\{\lambda > 0 : \lambda u_i \in \mathbb{Z}^2\}$, which are both well-defined, as the directions are rational (in fact $\rho_i \lambda_i = 1$, but we will use both notations for transparency).

6.1.3 Droplets

We next define the geometry of the droplets we will use.

Definition 6.1.4 (Droplet)**.** A *droplet* is a set of the form

$$\Lambda(\underline{r}) = \bigcap_{i \in [4k]} \overline{\mathbb{H}}_{u_i}(r_i)$$

for $\underline{r}$ with positive coordinates (see the black regions in Fig. 6.2). We say that a droplet is *symmetric* if it is of the form $x + \Lambda(\underline{r})$ with $2x \in \mathbb{Z}^2$ and $r_i = r_{i+2k}$ for all $i \in [2k]$. For a set of radii $\underline{r}$ we define the *side lengths* $\underline{s} = (s_i)_{i \in [4k]}$ with s_i the length of the side of $\Lambda(\underline{r})$ with outer normal u_i.

Note that if all $U \in \mathcal{U}$ are contained in the axes of $\mathbb{Z}^2$, then droplets are simply rectangles with sides parallel to the axes.

We write $(\underline{e}_i)_{i\in[4k]}$ for the canonical basis of $\mathbb{R}^{4k}$ and we write $\underline{1} = \sum_{i\in[4k]} \underline{e}_i$, so that $\Lambda(r\underline{1})$ is a polygon with inscribed circle of radius r and sides perpendicular to $\widehat{\mathcal{S}}$. It will often be more convenient to parametrise dimensions of droplets differently. For $i \in [4k]$ we set

$$\underline{v}_i = \sum_{j=i-k+1}^{i+k-1} \langle u_i, u_j\rangle \underline{e}_j.$$

This way $\Lambda(\underline{r} + \underline{v}_i)$ is obtained from $\Lambda(\underline{r})$ by extending the two sides parallel to u_i by 1 in direction u_i and leaving all other sides unchanged. Note that if $\Lambda(\underline{r})$ is symmetric, then so is $\Lambda(\underline{r} + \lambda_i\underline{v}_i)$ for $i \in [4k]$.

Definition 6.1.5 (Tube)**.** Given $i \in [4k]$, $\underline{r}$ and a multiple l of λ_i, we define the *tube of length l, direction i and radii $\underline{r}$* (see the thickened regions in Fig. 6.2)

$$T(\underline{r}, l, i) = \Lambda(\underline{r} + l\underline{v}_i)\backslash\Lambda(\underline{r}).$$

We will often need to consider boundary conditions for our events on droplets and tubes. Given two disjoint finite regions $A, B \subset \mathbb{Z}^2$ and two configurations $\eta \in \Omega_A$ and $\omega \in \Omega_B$, we define $\eta \cdot \omega \in \Omega_{A\cup B}$ as

$$(\eta \cdot \omega)_x = \begin{cases} \eta_x & x \in A, \\ \omega_x & x \in B. \end{cases}$$

6.1.4 Scales

Throughout the chapter we consider the positive integer constants

$$1/\varepsilon \gg 1/\delta \gg C \gg W.$$

Each one is assumed to be large enough depending on $\mathcal{U}$ and, therefore, $\widehat{\mathcal{S}}$ and α (e.g. $W > \alpha$), and much larger than the next. These constants are not allowed to depend on q. Whenever asymptotic notation is used, its implicit constants are not allowed to depend on the above ones, but only on $\mathcal{U}$.

The following are our main scales corresponding to the mesoscopic and internal dynamics.

$$\ell^{\mathrm{m}+} = q^{-C}/\sqrt{\delta}, \qquad \ell^{\mathrm{m}} = q^{-C},$$
$$\ell^{\mathrm{m}-} = q^{-C}\cdot\sqrt{\delta}, \qquad \ell^{\mathrm{i}} = C^2\log(1/q)/q^{\alpha}.$$

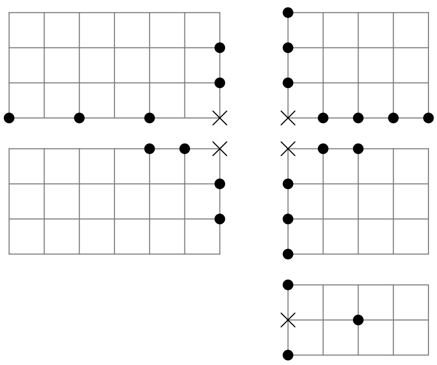

(a) The five update rules $U \in \mathcal{U}$ given as dots. The cross marks the origin.

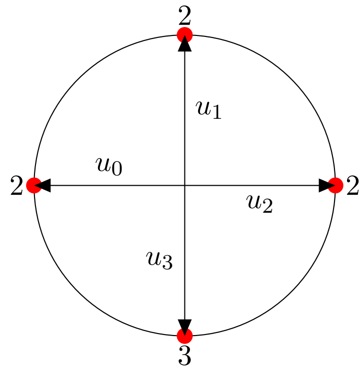

(b) The four stable directions, which coincide with $\hat{\mathcal{S}}$, and their difficulties.

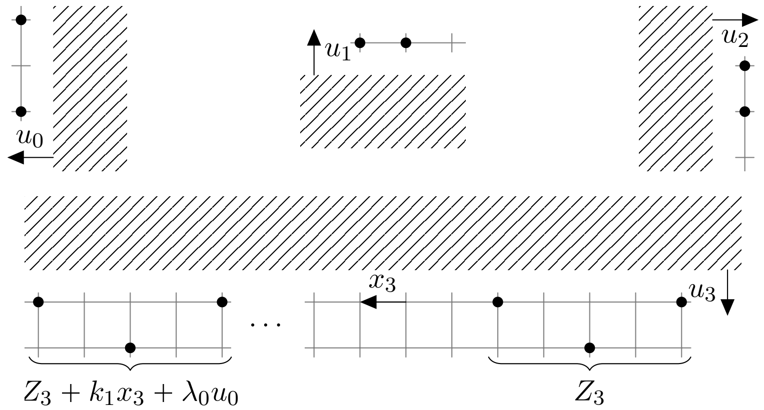

(c) Possible choice of u_i-helping sets. The hatched region represents $\mathbb{H}_{u_i} \cap \mathbb{Z}$.

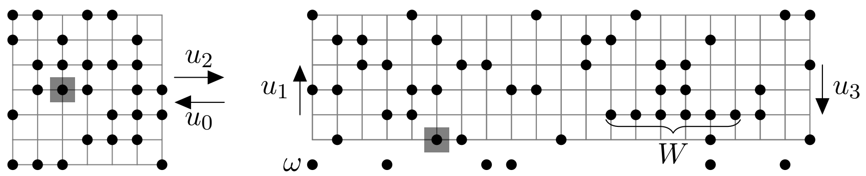

(d) Tubes traversable in the horizontal and vertical directions respectively. Note that for non-rectangular geometry, i.e. $k > 1$, a tube's shape uniquely determines the direction (see Fig. 6.2). Without the highlighted infection traversability is lost in both directions. Only the left tube is $\mathcal{ST}$ ($\mathcal{ST} = \mathcal{T}$ for u_0 and u_2). The right one is $\mathcal{T}^\omega$ in direction u_3, but not $\mathcal{T} = \mathcal{T}^{\mathbf{1}}$. Boundary conditions are irrelevant in other directions. The $C^2 + d$ offset (see Fig. 6.2) of Definition 6.2.1 is omitted.

Figure 6.1 – An intricate semi-directed example.

6.1.5 Helping sets

We next introduce various constant-sized sets of infections sufficient to induce growth. As the definitions are quite technical in general, in Fig. 6.1 we introduce a deliberately complicated example, on which to illustrate them.

Let us fix a direction $u_i \in \widehat{\mathcal{S}}$ with $\alpha(u_i) < \infty$. Let S be a nonempty discrete segment perpendicular to u_i. We will assume that S is of the form

$$\left\{x \in \mathbb{Z}^2 : \langle x, u_i\rangle = 0, \|x\| \leqslant r\right\}$$

for some $r \geqslant W$, but all definitions are extended by translation.

Definition 6.1.6 (W-helping set)**.** A *W-helping set* for S is any set of W consecutive infected sites in S, that is a set of the form $x + [W]\lambda_{i+k}u_{i+k}$ for some $x \in S$. We denote by $\mathcal{H}_d^W(S)$ the event that there is a W-helping set for S at distance at least d from its endpoints.

The relevance of W-helping sets is that, since W is large enough, together with a suitable neighbourhood of S in $\mathbb{H}_{u_i}$ they fully infect S by expanding the infected interval one site at a time. We next define some smaller sets which are sufficient to induce such growth but have the annoying feature that they are not necessarily contained in S and do not necessarily induce growth in a simple sequential way like W-helping sets. Let us note that except in Section 6.A.1 the reader will not lose anything conceptual by thinking that u_i and α-helping sets defined below are simply single infected sites in S.

Definition 6.1.7 (u_i-helping set)**.** Assume that $\alpha(u_i) \leqslant \alpha$. By Definition 1.6.1 (see [70, Lemma 3.3]) we may fix a set $Z_i \subset \mathbb{Z}^2\backslash\mathbb{H}_{u_i}$ and $x_i \in \mathbb{Z}^2\backslash\{0\}$ such that $\langle x_i, u_i\rangle = 0$ such that $|Z_i| = \alpha$, $[Z_i \cup \mathbb{H}_{u_i}]_{\mathcal{U}} \supset x_i\mathbb{N}$, where $\mathbb{N} = \{0, 1, \ldots\}$, and $|[Z_i \cup \mathbb{H}_{u_i}]_{\mathcal{U}}\backslash\overline{\mathbb{H}}_{u_i}| < \infty$. A *$u_i$-helping set* is a set of the form

$$\bigcup_{j \in \left[\frac{\|x_i\|}{\lambda_{i+k}}\right]} (Z_i + j\lambda_{i+k}u_{i+k} + k_j x_i), \tag{6.1}$$

for some integers k_j. Moreover, we choose x_i so that the *period*

$$Q := \frac{\|x_i\|}{\lambda_{i+k}}$$

is independent of i and sufficiently large so that for all $i \in [4k]$ with $\alpha(u_i) \leqslant \alpha$, the diameter of $\{0\} \cup Z_i$ is much smaller than Q. We may choose $Q = O(1)$ and all Z_i within distance $O(1)$ from the origin.

In the example of Fig. 6.1 only the u_3 direction has Z_3 such that $[Z_3 \cup \mathbb{H}_{u_3}]_{\mathcal{U}}$ only contains every second site of the line $\overline{\mathbb{H}}_{u_i}\backslash\mathbb{H}_{u_i}$. This is indeed necessary, since at least 4 sites are needed to infect the full line. For this model we might take $Q = 2$, corresponding to the fact that we need two translates of Z_3 with suitable residues modulo 2, in order to infect the entire line.

Definition 6.1.8 (α-helping set)**.**

- If $\alpha(u_i) \leqslant \alpha$ and $\alpha(u_{i+2k}) > \alpha$, then a *$\alpha$-helping set* is a u_i-helping set.
- If $\alpha(u_i) \leqslant \alpha$ and $\alpha(u_{i+2k}) \leqslant \alpha$, then a *$\alpha$-helping set* for S is a set of the form $H \cup H'$ with H a u_i-helping set and $-H' = \{-h : h \in H'\}$ a u_{i+2k}-helping set.
- If $\alpha < \alpha(u_i) \leqslant \infty$, there are no *$\alpha$-helping sets.*

If $\alpha(u_i) < \infty$, any set which is either a W-helping set or a α-helping set is called *helping set.* If $\alpha(u_i) = \infty$, there are no *helping sets.*

In the example of Fig. 6.1 u_0 and u_2 are both of difficulty $\alpha = 2$, so α-helping sets correspond to a pair of consecutive infections *and* a pair of infections at distance 2. These two pairs may be distant from each other within S. The consecutive infections are not u_0-helping sets, but we include them in α-helping sets in order for α-helping sets in direction u_0 to be the symmetric ones of those in direction u_2.

Definition 6.1.9. Let $\omega \in \Omega_{\mathbb{Z}^2 \setminus S}$ be a boundary condition. The event $\mathcal{H}_d^\omega(S)$ occurs if S has a helping set such that the vectors by which the sets Z_i (and, possibly, $-Z_{i+2k}$) are translated in Eq. (6.1) are contained in S and are at distance at least d from the endpoints of S. It will be convenient, given a domain $\Lambda \supset S$ and a boundary condition $\omega \in \Omega_{\mathbb{Z}^2 \setminus \Lambda}$ to define $\mathcal{H}_d^\omega(S) = \{\eta \in \Omega_\Lambda : \eta_S \in \mathcal{H}_d^{\omega \cdot \eta_{\Lambda \setminus S}}(S)\}$ by abuse. We write simply $\mathcal{H}(S)$ if $d = 0$ and the domain Λ is such that the boundary condition ω is irrelevant.

The following observation will be used systematically in probability estimates. It follows easily from the definitions above (see e.g. [70, Lemma 4.2]).

Observation 6.1.10. For any ω we have: if $\alpha(u_i) < \infty$, then

$$\mu\left(\mathcal{H}^\omega(S)\right) \geqslant \mu\left(\mathcal{H}^W(S)\right) \geqslant 1 - \left(1 - q^W\right)^{\lfloor |S|/W \rfloor} \geqslant q^W;$$

if $\alpha(u_i) \leqslant \alpha$, then

$$\mu(\mathcal{H}(S)) \geqslant \left(1 - (1 - q^\alpha)^{\Omega(|S|)}\right)^{O(1)}.$$

6.1.6 Super good events

Throughout the chapter we will refer to SG events for various droplets but we will usually not need to know exactly how they are constructed. However, we will systematically assume that for any sequence of radii $\underline{r}$ and boundary condition ω it holds that

- for $x \in \mathbb{Z}^2$ we define $\mathcal{SG}^\omega(x + \Lambda(\underline{r}))$ by translating $\mathcal{SG}^{\omega(\cdot - x)}(\Lambda(\underline{r}))$ by x;

- any SG event is nonempty and decreasing both in the configuration and in the boundary condition;
- $\mu(\mathcal{SG}^\omega(\Lambda(\underline{r}))) \leqslant q^{-O(W)}\mu(\mathcal{SG}(\Lambda(\underline{r})))$, systematically writing $\mathcal{SG}$ for $\mathcal{SG}^{\mathbf{1}}$.

6.1.7 Constrained Poincaré inequalities

Finally, we define (constrained) the Poincaré constants of various regions. Henceforth we will use the shorthand notation $\mu_\Lambda(\cdot|\mathcal{SG}^\omega) = \mu_\Lambda(\cdot|\mathcal{SG}^\omega(\Lambda))$ and similarly for traversable tubes (see Definition 6.2.1 below), as well as for conditional variances. Given a finite $\Lambda \subset \mathbb{Z}^2$ such that $\mathcal{SG}(\Lambda)$ is defined, let $\gamma(\Lambda)$ be the smallest constant $\gamma \geqslant 1$ such that the inequality

$$\mathrm{Var}_\Lambda(f|\mathcal{SG}) \leqslant \gamma \sum_{x\in\Lambda} \mu_\Lambda\left(c_x^{\mathbf{1}}\,\mathrm{Var}_x(f)\right), \tag{6.2}$$

holds for all $f : \Omega \to \mathbb{R}$, where $c_x^{\Lambda,\omega}(\eta) = c_x(\eta_\Lambda \cdot \omega_{\mathbb{Z}^2\setminus\Lambda})$ (recall Eq. (1.2)) and we omit Λ, when $c_x^{\Lambda,\omega}$ is inside μ_Λ like we do for $\mathcal{SG}$, etc.

Remark 6.1.11. It is important to take note of the absence of conditioning on $\mathcal{SG}$ in the r.h.s. of Eq. (6.2). This definition follows Chapter 4 and differs from the one in Chapter 5. Although this nuance is not important most of the time, this choice is crucial for the proof of Theorem 6.3.8 below. Unfortunately, this enforces some minor adaptations when importing intermediate results from Chapter 5, as dealt with by the following observation.

Observation 6.1.12. Let $A \subset B \subset \mathbb{Z}^2$ be finite sets, $\mathcal{A} \subset \Omega_A$ and $\mathcal{B} \subset \Omega_B$ be events and $f : \Omega_B \to [0,\infty)$ be a function. Then

$$\mu_B\left(\mathbb{1}_{\mathcal{A}}\mu_A(f)|\,\mathcal{B}\right) \leqslant \frac{\mu(\mathcal{A})}{\mu(\mathcal{B})}\mu_B(f).$$

Typically, in Chapter 5 we transform terms of the form $\mu_B(\mathbb{1}_{\mathcal{A}}\mu_A(f|\mathcal{A})|\mathcal{B})$ into $\mu_B(f|\mathcal{B})$, relying on additional information about the relative structure of $\mathcal{A}$ and $\mathcal{B}$, which will not always be available to us in the present chapter. Instead, in Chapter 4 we would simply disregard the numerator in Observation 6.1.12, which is too rough for our purposes. Therefore, corresponding amendments are needed for adapting arguments from the latter chapter as well, although the definition Eq. (6.2) is unchanged.

6.2 One-directional extensions

We first need the following traversability $\mathcal{T}$ and symmetric traversability $\mathcal{ST}$ events, for which we make the same conventions as for SG events (see Section 6.1.6). The definition is illustrated in Figs. 6.1 and 6.2.

Definition 6.2.1 (Traversability)**.** Fix $i \in [4k]$, $\underline{r}$ and l multiple of ρ_i. Assume that $\alpha(u_j) < \infty$ for all $j \in (i-k, i+k)$. For $m \geqslant 0$ and $j \in (i-k, i+k)$ write $S_{j,m} = \Lambda(\underline{r} + m\underline{v}_i + \rho_j \underline{e}_j) \backslash \Lambda(\underline{r} + m\underline{v}_i)$ and implicitly always assume that the indices are such that $S_{j,m} \subset T(\underline{r}, l, i) =: T$. For $\omega \in \Omega_{\mathbb{Z}^2 \backslash \Lambda(\underline{r}+l\underline{v}_i)}$ we say that T is (ω, d)-*traversable* if for all m and j the event $\mathcal{H}^{\omega}_{C^2+d}(S_{j,m})$ occurs. We denote by $\mathcal{T}^{\omega}_d(T)$ the event that the tube is (ω, d)-traversable and omit ω if it is $\mathbf{1}$ and d if it is 0.

We define $\mathcal{ST}^{\omega}_d(T)$ (the tube is (ω, d)-*symmetrically traversable*) identically to $\mathcal{T}^{\omega}_d(T)$, except that we replace $\mathcal{H}^{\omega}_{C^2+d}(S_{j,m})$ by $\mathcal{H}^{W}_d(S_{j,m})$ for all j such that $\max(\alpha(u_j), \alpha(u_{j+2k})) > \alpha$. In particular, if no such j exist, $\mathcal{ST} = \mathcal{T}$.

We will use two different ways to enlarge droplets to larger scales based on the East and CBSEP-extensions from Definitions 6.2.2 and 6.2.4. Both share the following setting.

Let $\underline{r} = q^{-O(C)}$, $i \in [4k]$ and $l \in (0, \ell^{\mathrm{m}+}]$ be a multiple of λ_i. Following Chapter 5, define

$$d_m = \lambda_i \lfloor (3/2)^m \rfloor$$

for $m \in [1, M)$ and $M = \min\{m : \lambda_i (3/2)^m \geqslant l\}$. Let $d_M = l$, $\Lambda^m = \Lambda(\underline{r} + d_m \underline{v}_i)$ and $s_{m-1} = d_m - d_{m-1}$ for $m \in [2, M]$.

6.2.1 East-extension

Definition 6.2.2 (East-extension)**.** Fix $i \in [4k]$, $\underline{r}$ and l multiple of λ_i. Assume that $\mathcal{SG}(\Lambda(\underline{r}))$ is defined[2] and that $\alpha(u_j) < \infty$ for all $j \in (i-k, i+k)$. We say that we *East-extend* $\Lambda(\underline{r})$ *by* l *in direction* u_i (see Fig. 6.2a) if for all $s \in (0, l]$ divisible by λ_i and $\omega \in \Omega_{\mathbb{Z}^2 \backslash \Lambda(\underline{r}+s\underline{v}_i)}$ we have $\eta \in \mathcal{SG}^{\omega}(\Lambda(\underline{r} + s\underline{v}_i))$ iff

$$\eta_{\Lambda(\underline{r})} \in \mathcal{SG}(\Lambda(\underline{r})), \qquad\qquad \eta_{T(\underline{r},s,i)} \in \mathcal{T}^{\omega}(T(\underline{r}, s, i)).$$

Recall γ from Section 6.1.7. The following is proved in Section 6.A.3.

Proposition 6.2.3. *Assume that we East-extend* $\Lambda(\underline{r})$ *by* l *in direction* u_i. *Then*

$$\gamma(\Lambda(\underline{r} + l\underline{v}_i)) \leqslant \max\left(\gamma(\Lambda(\underline{r})), \mu^{-1}_{\Lambda(\underline{r})}(\mathcal{SG})\right) e^{O(C^2)\log^2(1/q)} \prod_{m=1}^{M-1} a_m,$$

with

$$a_m = \mu^{-1}\left(\mathcal{SG}\left(\Lambda^m + s_m u_i\right) \middle| \mathcal{SG}(\Lambda^m)\right). \tag{6.3}$$

[2]Note that we do not require $\mathcal{SG}^{\omega'}(\underline{r})$ to be defined for $\omega' \in \Omega_{\mathbb{Z}^2 \backslash \Lambda(\underline{r})}$, but, if it is, as noted in Section 6.1.6, we always write $\mathcal{SG}(\Lambda(\underline{r}))$ for $\mathcal{SG}^{\mathbf{1}}(\Lambda(\underline{r}))$.

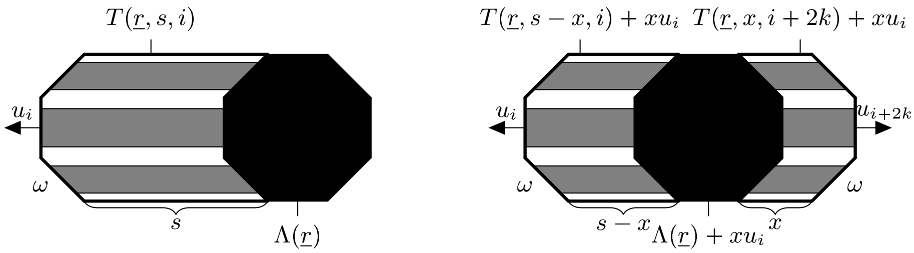

(a) East-extension. The thickened tube is traversable ($\mathcal{T}$).

(b) CBSEP-extension. Thickened tubes are symmetrically traversable ($\mathcal{ST}$).

Figure 6.2 – One-directional extensions. The black droplet is SG. Helping sets appear in the shaded parallelograms. White strips have width $\Theta(C^2)$.

6.2.2 CBSEP-extension

Definition 6.2.4 (CBSEP-extension)**.** Fix $i \in [4k]$, $\underline{r}$ and l divisible by λ_i. Assume that $\mathcal{SG}(\Lambda(\underline{r}))$ is defined[2] and that $\mathcal{U}$ has a finite number of stable directions. We say that we *CBSEP-extend* $\Lambda(\underline{r})$ *by* l *in direction* u_i (see Fig. 6.2b) if for all $s \in (0, l]$ divisible by λ_i and $\omega \in \Omega_{\mathbb{Z}^2 \setminus \Lambda(\underline{r}+s\underline{v}_i)}$ we have $\mathcal{SG}^\omega(\Lambda(\underline{r}+s\underline{v}_i)) = \bigcup_x \mathcal{SG}^\omega_x(\Lambda(\underline{r}+s\underline{v}_i))$ and for *offsets* $x \in [0, s]$ divisible by λ_i we define $\eta \in \mathcal{SG}^\omega_x(\Lambda(\underline{r}+s\underline{v}_i))$ iff the following all hold:

$$\begin{aligned}\eta_{T(\underline{r},s-x,i)+xu_i} &\in \mathcal{ST}^\omega(T(\underline{r}, s-x, i) + xu_i);\\ \eta_{\Lambda(\underline{r})+xu_i} &\in \mathcal{SG}(\Lambda(\underline{r}) + xu_i);\\ \eta_{T(\underline{r},x,i+2k)+xu_i} &\in \mathcal{ST}^\omega(T(\underline{r}, x, i+2k) + xu_i).\end{aligned}$$

The following is proved in Section 6.A.3 based on Chapter 5.

Proposition 6.2.5. *Assume that we CBSEP-extend* $\Lambda(\underline{r})$ *by* l *in direction* u_i*. Denote* $\Lambda_1 = T(\underline{r}, \lambda_i, i+2k)$, $\Lambda_2 = \Lambda(\underline{r}-\lambda_i\underline{v}_i)$ *and* $\Lambda_3 = T(\underline{r}-\lambda_i\underline{v}_i, \lambda_i, i)$, *so that* $\Lambda(\underline{r}+\lambda_i\underline{v}_i)-\lambda_i u_i = \Lambda_1 \sqcup \Lambda_2 \sqcup \Lambda_3$ *and* $\Lambda_2 \cup \Lambda_3 = \Lambda(\underline{r}) = (\Lambda_1 \cup \Lambda_2) + \lambda_i u_i$. *Suppose that events* $\overline{\mathcal{SG}}(\Lambda_2)$, $\overline{\mathcal{ST}}_{\eta_2}(\Lambda_1)$ *and* $\overline{\mathcal{ST}}_{\eta_2}(\Lambda_3)$ *are defined for all* $\eta_2 \in \overline{\mathcal{SG}}(\Lambda_2)$ *so that*

$$\left\{\eta : \eta_{\Lambda_1} \in \overline{\mathcal{ST}}_{\eta_{\Lambda_2}}(\Lambda_1), \eta_{\Lambda_2} \in \overline{\mathcal{SG}}(\Lambda_2), \eta_{\Lambda_3} \in \overline{\mathcal{ST}}_{\eta_{\Lambda_2}}(\Lambda_3)\right\} \subset \mathcal{SG}(\Lambda_1 \cup \Lambda_2) \cap \mathcal{SG}(\Lambda_2 \cup \Lambda_3). \quad (6.4)$$

Then

$$\gamma(\Lambda(\underline{r}+l\underline{v}_i)) \leqslant \frac{\mu_{\Lambda(\underline{r})}(\mathcal{SG}) \max(\mu^{-1}_{\Lambda(\underline{r})}(\mathcal{SG}), \gamma(\Lambda(\underline{r})))e^{O(C^2)\log^2(1/q)}}{\mu_{\Lambda(\underline{r}+l\underline{v}_i)}(\mathcal{SG})\mu_{\Lambda_1\cup\Lambda_2}(\overline{\mathcal{SG}}|\mathcal{SG}) \min_{\eta_2 \in \overline{\mathcal{SG}}(\Lambda_2)} \mu_{\Lambda_3}(\overline{\mathcal{ST}}_{\eta_2}|\mathcal{ST}^{\mathbf{0}})},$$

where $\overline{\mathcal{SG}}(\Lambda_1 \cap \Lambda_2) = \{\eta : \eta_{\Lambda_2} \in \overline{\mathcal{SG}}(\Lambda_2), \eta_{\Lambda_1} \in \overline{\mathcal{ST}}_{\eta_{\Lambda_2}}\}$.

6.3 East-type dynamics

In this section we treat various East-type dynamics on all scales. This is the most novel and central part of this chapter, albeit the most technical.

6.3.1 Internal East dynamics

For this section we assume that $\mathcal{U}$ is balanced and without loss of generality that we may fix $i \in (0, 2k)$ such that $\alpha(u_j) \leqslant \alpha$ for all $j \in (-k, i+k)$. Let $\underline{r}^{(0)} = (r_j^{(0)})_{j\in[4k]}$ be a symmetric sequence of radii such that $r_j = \Theta(1/\varepsilon)$ is a multiple of λ_j for all j, the vertices of $\Lambda(\underline{r}^{(0)})$ are in $\lambda_i u_i \mathbb{Z} + \lambda_0 u_0 \mathbb{Z}$ and the corresponding side lengths $\underline{s}^{(0)}$ are also $\Theta(1/\varepsilon)$. We define

$$s_j^{(n)} = \begin{cases} s_j^{(0)}\ell^{(n)} & -k < j < i+k, \\ s_j^{(0)} & i+k < j < 3k \end{cases} \tag{6.5}$$

and $s_{-k}^{(n)}$ and $s_{i+k}^{(n)}$ as required to be the sides of a droplet, where

$$\ell^{(n)} = \begin{cases} W^n & n \leqslant N^{\mathrm{c}}, \\ \lceil W^{\exp(n-N^{\mathrm{c}})}/q^\alpha \rceil & N^{\mathrm{c}} < n \leqslant N^{\mathrm{i}}, \end{cases} \tag{6.6}$$

$$\begin{aligned} N^{\mathrm{c}} &= \min\{n : W^n \geqslant q^{-\alpha}\} \\ N^{\mathrm{i}} &= \min\left\{n : \ell^{(n)} \geqslant \ell^{\mathrm{i}}\varepsilon\right\} = N^{\mathrm{c}} + \log\log\log(1/q) + O(\log\log W). \end{aligned}$$

We denote $\Lambda^{(n)} = \Lambda(\underline{r}^{(n)})$, where $\underline{r}^{(n)}$ is the sequence of radii corresponding to $\underline{s}^{(n)}$ such that $r_{-k}^{(n)} = r_{-k}^{(0)}$ and $r_{i+k}^{(n)} = r_{i+k}^{(0)}$ (see Fig. 6.3a).

Remark 6.3.1. Note that despite the extremely fast divergence of $\ell^{(n)}q^\alpha$, for $n \in (N^{\mathrm{c}}, N^{\mathrm{i}}]$ it holds that $W \leqslant \ell^{(n+1)}/\ell^{(n)} < (\ell^{(n)}q^\alpha)^2 < \log^4(1/q)$. The sharp divergence will ensure that some error terms below sum to the largest one, so as to avoid additional factors of the order of $N^{\mathrm{i}} - N^{\mathrm{c}}$ in the final answer, particularly for the semi-directed class (f). This technique was introduced in Eq. (10.14), while the geometrically increasing scale choice relevant for small n originates from [190]. It should be noted that this divergence can be further amplified up to a tower of exponentials of height linear in $n - N^{\mathrm{c}}$. In that case the $\log\log\log(1/q)$ error term in Theorem 6.3.8 becomes $\log_*(1/q)$, but is, alas, still divergent.

Semi-directed models

We now assume that $\mathcal{U}$ is semi-directed. In that case we may set $i = 2k-1$. This is the only feature of semi-directed models used in this subsection. Hence, the reasoning applies equally well to all balanced models with rules

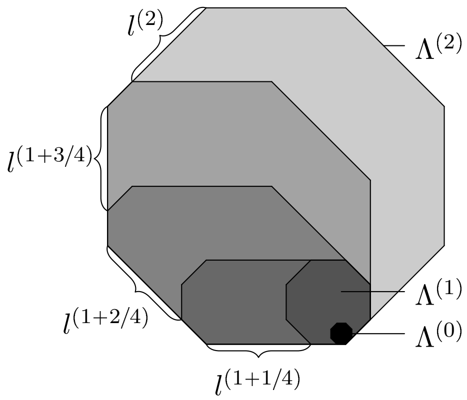

(a) Case case $i = 2k - 1$ of semi-directed models in Section 6.3.1. For $n \in \mathbb{N}$ droplets are symmetric and homothetic to the black $\Lambda^{(0)}$. Intermediate ones $\Lambda^{(1+1/4)}$, $\Lambda^{(1+2/4)}$ and $\Lambda^{(1+3/4)}$ obtained by East-extensions (see Fig. 6.2a) in directions u_0, u_1 and u_2 respectively are drawn in progressive shades of grey.

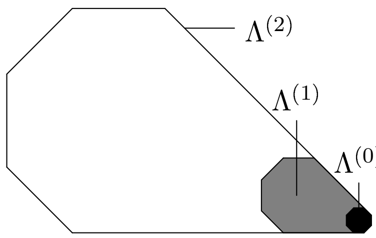

(b) Case $i = 1$ of balanced rooted models in Section 6.3.1. The black, grey and white droplets are $\Lambda^{(0)}$, $\Lambda^{(1)}$ and $\Lambda^{(2)}$ respectively. In this case no fractional scales are introduced.

Figure 6.3 – Geometry of the nested droplets $\Lambda^{(n)}$ used in Section 6.3.1 for $k = 2$.

contained in the axes of the lattice, since then $k = 1$ and we can always set $i = 1$ for balanced models.

Observe that, since $i = 2k - 1$, we may obtain $\Lambda^{(n+1)}$ from $\Lambda^{(n)}$ by $2k$ successive extensions in directions u_0 through u_i (see Fig. 6.3a). We denote the droplets obtained this way by $\Lambda^{(n+j/(2k))}$ for $j \in (0, 2k)$ and denote their radii and side lengths by $\underline{r}^{(n+j/(2k))}$ and $\underline{s}^{(n+j/(2k))}$ respectively. We write $l^{(n+j/(2k))} = s_{j+k}^{(n+1)} - s_{j+k}^{(n)} = \Theta(\ell^{(n+1)}/\varepsilon)$ for the length l such that $\underline{r}^{(n+(j+1)/(2k))} = \underline{r}^{(n+j/(2k))} + l\underline{v}_j$.

Definition 6.3.2 (Semi-directed internal SG)**.** Let $\mathcal{U}$ be semi-directed. Let $\mathcal{SG}(\Lambda^{(0)})$ be the event that $\Lambda^{(0)}$ is fully infected. Recursively, for $n \in [N^{\mathrm{i}}]$ and $j \in [2k]$, we define $\mathcal{SG}(\Lambda^{(n+(j+1)/(2k))})$ by East-extending $\Lambda^{(n+j/(2k))}$ in direction u_j by $l^{(n+j/(2k))}$ (recall Definition 6.2.2).

Theorem 6.3.3. *Let $\mathcal{U}$ be semi-directed. Then*

$$\gamma\left(\Lambda^{(N^{\mathrm{i}})}\right) \leqslant \exp\left(\frac{\log\log(1/q)}{\varepsilon^6 q^\alpha}\right), \qquad \mu_{\Lambda^{(N^{\mathrm{i}})}}(\mathcal{SG}) \geqslant \exp\left(\frac{-1}{\varepsilon^2 q^\alpha}\right).$$

Proof of Theorem 6.3.3. For $n \in 1/(2k)\mathbb{N}$, $j \in [2k]$ and $m \geqslant 2$, such that $n < N^{\mathrm{i}}$ and $(n - j/(2k)) \in \mathbb{N}$ set

$$a_m^{(n)} = \mu^{-1}\left(\mathcal{SG}\left(\Lambda^{(n)} + \left(\lfloor (3/2)^{m+1} \rfloor - \lfloor (3/2)^m \rfloor\right)\lambda_j u_j\right) \middle| \mathcal{SG}\left(\Lambda^{(n)}\right)\right). \quad (6.7)$$

For the sake of simplifying expressions we will abusively assume that for all $l^{(n)}$ are of the form $\lambda_j\lfloor(3/2)^m\rfloor$ with integer m. Without this assumption, one would need to treat the term corresponding to $m = M-1$ in Proposition 6.2.3 separately, but identically.

From Proposition 6.2.3 we get

$$\gamma\left(\Lambda^{(N^{\mathrm{i}})}\right) \leqslant \frac{\exp(\log^{O(1)}(1/q))}{\mu^2_{\Lambda^{(N^{\mathrm{i}})}}(\mathcal{SG})} \prod_{n=0}^{N^{\mathrm{i}}-1/(2k)} \prod_{m=1}^{M^{(n)}} a_m^{(n)}, \tag{6.8}$$

where the product is over $n \in 1/(2k)\mathbb{N}$ and $M^{(n)} = \log l^{(n)}/\log(3/2) + O(1)$. Indeed, by the Harris inequality a_m in Eq. (6.3) for $\underline{r} = \underline{r}^{(n)}$ is at most

$$a_m^{(n)}/\mu\left(\mathcal{T}\left(T(\underline{r}^{(n)}, \left(\lfloor(3/2)^{m+1}\rfloor - \lfloor(3/2)^m\rfloor)\lambda_j, j\right)\right)\right),$$

while by Definition 6.2.2

$$\prod_{m=1}^{M^{(n)}} \mu\left(\mathcal{T}\left(T\left(\underline{r}^{(n)}, (\lfloor(3/2)^{m+1}\rfloor - \lfloor(3/2)^m\rfloor)\lambda_j, j\right)\right)\right)$$
$$\geqslant q^{O(WM^{(n)})}\frac{\mu(\mathcal{SG}(\Lambda^{(n+1/(2k))}))}{\mu(\mathcal{SG}(\Lambda^{(n)}))}.$$

To evaluate the r.h.s. of Eq. (6.8) we will need the following lemma.

Lemma 6.3.4. *Let $n \in 1/(2k)\mathbb{N}$ be such that $n \leqslant N^{\mathrm{i}}$ and $m \geqslant 1$. Then*

$$a_m^{(n)} \leqslant \mu^{-1}_{\Lambda^{(n)}}(\mathcal{SG}) \leqslant \min\left((\delta q^\alpha W^n)^{-W^n/\varepsilon^2}, e^{1/(\varepsilon^2 q^\alpha)}\right). \tag{6.9}$$

Moreover, if

$$\ell^{(\lfloor n\rfloor)} \geqslant 1/\left(q^\alpha \log^W(1/q)\right), \qquad M^{(n)} \geqslant m+W, \qquad (3/2)^m \leqslant 1/q^\alpha,$$

setting

$$n_m = \min\left\{n' \in \mathbb{N} : \ell^{(n')} \geqslant 1/\left(q^\alpha \log^W(1/q)\right), M^{(n')} \geqslant m+W\right\} \leqslant n, \tag{6.10}$$

the following improvements holds

$$a_m^{(n)} \leqslant \exp\left(\frac{(3/2)^m}{\varepsilon^4}\left((N^{\mathrm{c}} - n_m)^2 + \mathbb{1}_{n\geqslant N^{\mathrm{c}}}\log^{2/3}\log(1/q)\right)\right) \tag{6.11}$$
$$\times\begin{cases}\exp\left(1/\left(q^\alpha \log^{W-O(1)}(1/q)\right)\right) & \text{if } m \leqslant \frac{\log(1/(q^\alpha\log^W(1/q)))}{\log(3/2)}\\ \exp\left(1/\left(q^\alpha\log^{W-O(1)}\log(1/q)\right)\right) & \text{otherwise.}\end{cases}$$

Let us finish the proof of Theorem 6.3.3 before proving Lemma 6.3.4. Using the trivial bound $a_m^{(n)} \leqslant \exp(1/(\varepsilon^2 q^\alpha))$ from Eq. (6.9) we get

$$\prod_{n=N^{\mathrm{c}}-\lceil 1/\varepsilon\rceil}^{N^{\mathrm{i}}-1/(2k)} \prod_{m=\left\lceil \frac{\log(1/q^\alpha)}{\log(3/2)}\right\rceil}^{M^{(n)}} a_m^{(n)} \leqslant \exp\left(\frac{1+\sum_{n=N^{\mathrm{c}}}^{N^{\mathrm{i}}} e^{n+1-N^{\mathrm{c}}}}{\varepsilon^4 q^\alpha}\right) = \exp\left(\frac{\log\log(1/q)}{\varepsilon^5 q^\alpha}\right), \tag{6.12}$$

which is the main contribution. Note that $n < N^{\mathrm{c}} - 1/\varepsilon$ implies $M^{(n)} < \log(1/q^\alpha)/\log(3/2)$, so the above product exhausts the terms in Eq. (6.8) with large m.

Next, using the first bound on $a_m^{(n)}$ from Eq. (6.9), we obtain

$$\prod_{n=0}^{\frac{-\log(q^\alpha \log^W(1/q))}{\log W}} \prod_{m=1}^{M^{(n)}} a_m^{(n)} \leqslant \exp\left(1/\left(q^\alpha \log^{W-O(1)}(1/q)\right)\right); \tag{6.13}$$

$$\prod_{m=1}^{\frac{\log(1/q^\alpha)}{\log(3/2)}} \prod_{n=\left\lceil \frac{\log(1/(q^\alpha \log^W(1/q)))}{\log W}\right\rceil}^{n_m-1/(2k)} a_m^{(n)} \leqslant \exp\left(-\sum_{m=1}^{\frac{\log(1/q^\alpha)}{\log(3/2)}} \frac{(3/2)^m \log(\delta q^\alpha (3/2)^m)}{\varepsilon^3}\right) \leqslant e^{1/(q^\alpha\varepsilon^4)}. \tag{6.14}$$

Finally, we use Eq. (6.11) to show

$$\begin{aligned}\sum_{m=1}^{\frac{\log(1/q^\alpha)}{\log(3/2)}} \sum_{n=n_m}^{N^{\mathrm{i}}-1/(2k)} \log a_m^{(n)} \leqslant{}& \log\log\log(1/q) \sum_{m=1}^{\frac{\log(1/q^\alpha)}{\log(3/2)}} \frac{(3/2)^m (N^{\mathrm{c}}-n_m)^3}{\varepsilon^4} + \\ & \log\log\log(1/q)\log^{2/3}\log(1/q) \sum_{m=1}^{\frac{\log(1/q^\alpha)}{\log(3/2)}} \frac{(3/2)^m}{\varepsilon^4} + \\ & \frac{N^{\mathrm{i}}}{q^\alpha \log^{W-O(1)}(1/q)} \frac{\log(1/(q^\alpha \log^W(1/q)))}{\log(3/2)} + \\ & \frac{\log^{O(1)}\log(1/q))}{q^\alpha \log^{W-O(1)}\log(1/q)} \\ \leqslant{}& \frac{\log\log\log(1/q)}{q^\alpha \varepsilon^{O(1)}} + \frac{\log\log\log(1/q)\log^{2/3}\log(1/q)}{q^\alpha\varepsilon^{O(1)}} + \\ & \frac{1}{q^\alpha \log^{W-O(1)}(1/q)} + \frac{1}{q^\alpha \log^{W-O(1)}\log(1/q)} \\ \leqslant{}& \frac{\log^{3/4}\log(1/q)}{q^\alpha}.\end{aligned}$$

Plugging the last result and Eqs. (6.12) to (6.14) in Eq. (6.8) and recalling Eq. (6.9), we conclude the proof of Theorem 6.3.3. □

Proof of Lemma 6.3.4. Let us fix m and n as in the statement for Eq. (6.9). The bound $a_m^{(n)} \leqslant \mu_{\Lambda^{(n)}}^{-1}(\mathcal{SG})$ follows from the Harris inequality. To upper bound the latter term we note that by Definition 6.2.2,

$$\mu_{\Lambda^{(n)}}(\mathcal{SG}) = \mu_{\Lambda^{(0)}}(\mathcal{SG}) \prod_{p=0}^{n-1/(2k)} \mu_{T(\underline{r}^{(p)},l^{(p)},j(p))}(\mathcal{T}), \tag{6.15}$$

setting $j(p) \in [2k]$ such that $p - j(p)/(2k) \in \mathbb{N}$. Clearly,

$$\mu_{\Lambda^{(0)}}(\mathcal{SG}) = q^{|\Lambda^{(0)}|} = q^{\Theta(1/\varepsilon^2)}. \tag{6.16}$$

Let us fix $p \in 1/(2k)\mathbb{N}$, $p < N^{\mathrm{i}}$. Then, using Definition 6.2.1, Observation 6.1.10 and the Harris inequality, we get

$$\begin{aligned} \mu_{T(\underline{r}^{(p)},l^{(p)},j(p))}(\mathcal{T}) &\geqslant q^{O(W)} \left(1 - e^{-q^{\alpha}\ell^{(p)}/O(\varepsilon)}\right)^{O(l^{(p)})} \qquad (6.17)\\ &\geqslant q^{O(W)} \cdot \begin{cases} (\delta q^{\alpha} W^{p})^{W^{p}/(\delta\varepsilon)} & p \leqslant N^{\mathrm{c}}, \\ \exp\left(-1/\left(q^{\alpha}\exp\left(W^{\exp(p-N^{\mathrm{c}})}/\delta\right)\right)\right) & p > N^{\mathrm{c}}, \end{cases} \end{aligned}$$

the last inequality taking into account $1/\varepsilon \gg 1/\delta \gg W \gg 1$, $\ell^{(N^{\mathrm{c}})} = W^{O(1)}/q^{\alpha}$ and the explicit expressions Eq. (6.6). From Eqs. (6.15) to (6.17) it is elementary to check Eq. (6.9).

We next turn to proving Eq. (6.11), so we fix n and m satisfying the corresponding hypotheses of Lemma 6.3.4. Denote $s_m = (\lfloor(3/2)^{m+1}\rfloor - \lfloor(3/2)^{m}\rfloor)\lambda_j u_j$ for $j = j(n)$ as in Eq. (6.7), so that

$$a_m^{(n)} = \mu^{-1}\left(\mathcal{SG}\left(\Lambda^{(n)} + s_m\right)\middle|\mathcal{SG}\left(\Lambda^{(n)}\right)\right).$$

By the Harris inequality, Definitions 6.2.2 and 6.3.2 we have

$$\begin{aligned} a_m^{(n)} &\leqslant \mu_{\Lambda^{(n_m)}}^{-1}(\mathcal{SG}) \qquad (6.18)\\ &\times \prod_{p=n_m}^{n-1/(2k)} \mu^{-1}\left(\mathcal{T}\left(T\left(\underline{r}^{(p)},l^{(p)},j(p)\right)+s_m\right)\middle|\mathcal{T}\left(T\left(\underline{r}^{(p)},l^{(p)},j(p)\right)\right)\right). \end{aligned}$$

Our goal is then to bound the last factor, using Corollary 6.B.4, which quantifies the fact that "small perturbations s_m do not modify traversability much."

Let us fix p as above, denote $T = T(\underline{r}^{(p)},l^{(p)},j(p))$ and $T' = T + s_m$. From Eq. (6.10) it is not hard to check that the hypotheses of Corollary 6.B.4

are satisfied, so that

$$\begin{aligned}\mu(\mathcal{T}(T')|\mathcal{T}(T)) &\geqslant q^{-O(W)} \times \left(1-(1-q^{\alpha})^{\ell^{(\lfloor p\rfloor)}/O(\varepsilon)}\right)^{O((3/2)^m)} \\ &\quad\times\left(1-O(W\varepsilon)(3/2)^m/\ell^{(\lfloor p\rfloor)}-q^{1-o(1)}\right)^{O(\ell^{(\lfloor p\rfloor+1)}/\varepsilon)} \\ &\geqslant q^{-O(W)}\times\begin{cases}(\delta q^{\alpha}W^{p})^{O((3/2)^m)} & p\leqslant N^{\mathrm{c}}\\ \exp(-(3/2)^m\exp\left(-W^{\exp(p-N^{\mathrm{c}})}/\delta\right)) & p>N^{\mathrm{c}}\end{cases}\\ &\quad\times\begin{cases}\exp\left(-q^{-\alpha+1/2-o(1)}\right) & (3/2)^m\leqslant q^{-\alpha+1/2-o(1)}\\ \exp\left(-W^2(3/2)^m\frac{\ell^{(\lfloor p\rfloor+1)}}{\ell^{(\lfloor p\rfloor)}}\right) & \text{otherwise.}\end{cases}\end{aligned} \tag{6.19}$$

Further notice that if $(3/2)^m \leqslant q^{-\alpha+1/2-o(1)}$ or $p > N^{\mathrm{c}}$, the third term dominates, while otherwise the second one does. Moreover, if $p > N^{\mathrm{c}}+\Delta$ with

$$\Delta := \log\frac{\log\log\log(1/q)}{3\log W},$$

then the Harris inequality and Eq. (6.17) directly give the bound

$$\mu(\mathcal{T}(T')|\mathcal{T}(T)) \geqslant \exp\left(-1/\left(q^{\alpha}\log^{W}\log(1/q)\right)\right). \tag{6.20}$$

Finally, we can plug Eqs. (6.9), (6.19) and (6.20) in Eq. (6.18) to obtain the following bounds. If $(3/2)^m \leqslant q^{-\alpha+1/2-o(1)}$, then

$$a_m^{(n)} \leqslant \exp\left(1/\left(q^{\alpha}\log^{W}(1/q)\right)\right).$$

Otherwise,

$$\begin{aligned}a_m^{(n)} \leqslant{} &\begin{cases}\exp\left(1/\left(q^{\alpha}\log^{W-O(1)}(1/q)\right)\right) & (3/2)^m\leqslant 1/\left(q^{\alpha}\log^{W}(1/q)\right)\\ (\delta q^{\alpha}W^{n_m})^{-(3/2)^m/\varepsilon^3} & \text{otherwise}\end{cases}\\ &\times\prod_{p=n_m}^{\min(n,N^{\mathrm{c}})}(\delta q^{\alpha}W^{p})^{-O((3/2)^m)}\\ &\times\begin{cases}1 & n\leqslant N^{\mathrm{c}}\\ \exp\left((3/2)^m W^{2\exp(\Delta)}/\delta\right) & n>N^{\mathrm{c}}\end{cases}\\ &\times\begin{cases}\exp\left(1/\left(q^{\alpha}\log^{W-O(1)}(1/q)\right)\right) & (3/2)^m\leqslant 1/\left(q^{\alpha}\log^{W}(1/q)\right)\\ \exp\left(1/\left(q^{\alpha}\log^{W-O(1)}\log(1/q)\right)\right) & \text{otherwise,}\end{cases}\end{aligned}$$

the terms corresponding to $\mu^{-1}_{\Lambda^{(n_m)}}(\mathcal{SG})$ and to values of p in the intervals $[n_m, N^{\mathrm{c}}]$, $(N^{\mathrm{c}}, N^{\mathrm{c}}+\Delta]$ and $(N^{\mathrm{c}}+\Delta, N^{\mathrm{i}})$ respectively. Indeed, in the last term for small m we used Eq. (6.19), while for large m, we directly applied

Eq. (6.20). Observing that the product of the second case for the first term, the second term and the third term can be bounded by

$$\exp\left(\frac{(3/2)^m}{\varepsilon^4}\left((N^{\mathrm{c}}-n_m)^2+\mathbb{1}_{n\geqslant N^{\mathrm{c}}}\log^{2/3}\log(1/q)\right)\right),$$

we obtain the desired Eq. (6.11). □

Balanced rooted models

We now assume that $\mathcal{U}$ is balanced rooted (the rooted character is assumed only in order not to override definitions for semi-directed families above and isotropic ones in Section 6.4.1 below, but is not needed otherwise). We set $i=1$ in this case, which is always valid, since the model is balanced (recall i from the beginning of Section 6.3.1).

We need to define a two-directional East-extension which is morally the concatenation of one in direction u_0 and one in direction u_1, but whose actual definition is much more technical, so as to respect the homothetic relation between the $\Lambda^{(n)}$ and yet maintain a product structure.

We begin with some geometric preparations. Fix $n\in[N^{\mathrm{i}}]$ (since the definitions for semi-directed models no longer apply, but only the ones from the beginning of Section 6.3.1 do, n is an integer here). Observe that we can cover $\Lambda^{(n+1)}$ with droplets $(D_\kappa)_{\kappa\in[K]}$ so that the following conditions hold (see Fig. 6.4).

- For all $\kappa\in[K]$, $D_\kappa\subset\Lambda^{(n+1)}$;
- $\bigcup_{\kappa=2}^{K-1}D_\kappa=\Lambda^{(n+1)}$;
- $K=O(\ell^{(n+1)}/\ell^{(n)})$;
- any segment of length $\ell^{(n)}/(C\varepsilon)$ perpendicular to u_j for some $j\in(-k,k]$ intersects at most $O(1)$ of the D_κ;
- droplets are assigned a *generation* $g\in\{0,1,2\}$, so that only $D_0:=\Lambda^{(n)}$ is of generation $g=0$, only $D_1:=\Lambda(\underline{r}^{(n)}+l_1\underline{v}_1)$ is of generation $g=1$, where
$$l_1=\frac{r_k^{(n+1)}-r_k^{(n)}}{\langle u_1,u_k\rangle},$$
so that D_1 spans the u_{k+1}-side of $\Lambda^{(n+1)}$;
- if $\kappa\geqslant 2$, then D_κ is of generation $g=2$, and is of the form
$$D_\kappa=y_\kappa u_1+\Lambda(\underline{r}^{(n)}+l_\kappa\underline{v}_0)$$
for certain $l_\kappa\geqslant 0$ and $y_\kappa\in[0,l_1]$ multiple of λ_1.

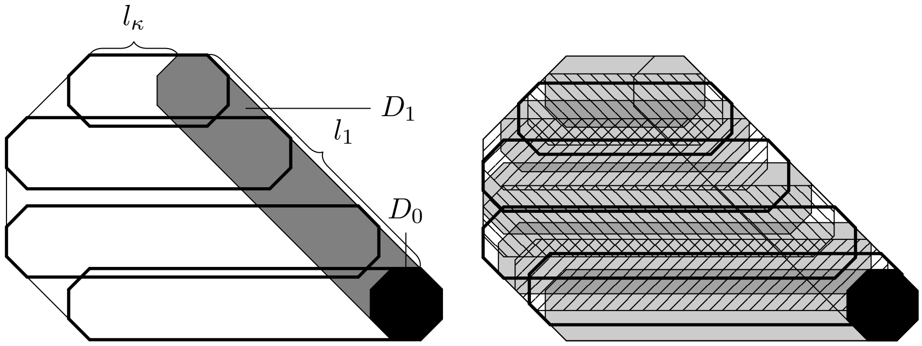

(a) The droplets D_κ corresponding to corners of $\Lambda^{(n+1)}$. The generation 0 droplet is given in black, while the one of generation 1 is shaded.

(b) All droplets D_κ. In the second generation, for visibility, droplets alternate between shaded, thickened and hatched.

Figure 6.4 – Geometry of the droplets $(D_\kappa)_{\kappa\in[K]}$ used in the two-dimensional East-extension in Definition 6.3.5. Also recall Fig. 6.3b.

To construct the D_κ of generation 2, it essentially suffices to increment y_κ by $\Theta(\ell^{(n)}/\varepsilon)$ and define l_κ to be the largest possible, so that $D_\kappa \subset \Lambda^{(n+1)}$. Finally, we add to our collection of droplets the ones with y_κ corresponding to a corner of $\Lambda^{(n+1)}$ and again take l_κ maximal (see Fig. 6.4).

Definition 6.3.5 (Two-dimensional East-extension)**.** Fix $n \in [N^{\mathrm{i}}]$ and let $R \subset \Lambda^{(n+1)}$ be a region of the form

$$\bigcup_{I\in\mathcal{I}}\left(\bigcap_{\kappa\in I} D_\kappa \setminus \bigcup_{\kappa\in[K]\setminus I} D_\kappa\right) \tag{6.21}$$

for some family $\mathcal{I}$ of subsets of $[K]$. We say that R is *n-traversable* ($\mathcal{T}_n(R)$ occurs) if for all $j \in (-k,k)$ and every segment $S \subset R$ perpendicular to u_j of length at least $\delta\ell^{(n)}/\varepsilon$ at distance

- at least W from the boundary of all D_κ, the event $\mathcal{H}(S)$ occurs;
- at most W from a side of a D_κ parallel to S for some $\kappa \in [K]$, but S does not intersect any non-parallel side of any $D_{\kappa'}$, the event $\mathcal{H}^W(S)$ occurs.

We say that we *East-extend* $\Lambda^{(n)}$ to $\Lambda^{(n+1)}$ if the event $\mathcal{SG}(D_1)$ is defined by East-extending $\Lambda^{(n)}$ by l_1 in direction u_1 and $\mathcal{SG}(\Lambda^{(n+1)})$ is defined as $\mathcal{SG}(D_1) \cap \mathcal{T}_n(\Lambda^{(n+1)}\setminus D_1)$. Indeed,

$$\Lambda^{(n+1)}\setminus D_1 = \bigcup_{\kappa\in[K]} D_\kappa\setminus D_1 = \bigcup_{I\subset[K]\setminus\{0,1\}}\left(\bigcap_{\kappa\in I} D_\kappa\setminus\bigcup_{\kappa\notin I} D_\kappa\right).$$

Remark 6.3.6. Note that these n-traversability events are product over the disjoint regions into which all the boundaries of $(D_\kappa)_{\kappa\in[K]}$ partition $\Lambda^{(n+1)}$.

Armed with this notion, we are ready to define our SG events up to the internal scale for our models of interest.

Definition 6.3.7 (Balanced rooted internal SG)**.** Let $\mathcal{U}$ be balanced rooted. Let $\mathcal{SG}(\Lambda^{(0)})$ be the event that $\Lambda^{(0)}$ is fully infected. We define $\mathcal{SG}(\Lambda^{(n)})$ for $n \in [N^{\mathrm{i}}]$ by successively East-extending $\Lambda^{(p)}$ to $\Lambda^{(p+1)}$.

Theorem 6.3.8. *Let $\mathcal{U}$ be balanced rooted. Then*

$$\gamma\left(\Lambda^{(N^{\mathrm{i}})}\right) \leqslant \exp\left(\frac{\log(1/q)\log\log\log(1/q)}{\varepsilon^3 q^\alpha}\right), \quad \mu_{\Lambda^{(N^{\mathrm{i}})}}(\mathcal{SG}) \geqslant \exp\left(\frac{-1}{\varepsilon^2 q^\alpha}\right).$$

Proof of Theorem 6.3.8. For $m \geqslant 1$ and $n \in [N^{\mathrm{i}}]$ denote

$$a_m^{(n)} = \max_{j\in\{0,1\}} \mu^{-1}\left(\mathcal{SG}\left(\Lambda^{(n)} + \left(\lfloor(3/2)^{m+1}\rfloor - \lfloor(3/2)^m\rfloor\right)\lambda_j u_j\right)\middle|\mathcal{SG}\left(\Lambda^{(n)}\right)\right). \tag{6.22}$$

For the sake of simplifying expressions we will abusively assume that for all $\kappa \in [K]$ the length l_κ is of the form $\lambda_0\lfloor(3/2)^m\rfloor$ with integer m. Without this assumption, one would need to treat the term corresponding to $m = M-1$ in Proposition 6.2.3 separately, but identically. We next deduce Theorem 6.3.8 from the following two lemmas.

Lemma 6.3.9. *For $n < N^{\mathrm{i}}$ we have*

$$\gamma\left(\Lambda^{(n+1)}\right) \leqslant \frac{\gamma(\Lambda^{(n)})e^{O(C^2)\log^2(1/q)}}{(\mu_{\Lambda^{(n+1)}}(\mathcal{SG})\mu_{\Lambda^{(n+1)}}(\mathcal{T}_n))^{O(1)}}\prod_{m=1}^{M^{(n)}} a_m^{(n)},$$

where $M^{(n)} = \lceil 1/\varepsilon\rceil + \lceil\log \ell^{(n+1)}/\log(3/2)\rceil$.

Lemma 6.3.10. *For any $n \leqslant N^{\mathrm{i}}$ and $m \geqslant 1$ we have*

$$\begin{aligned} a_m^{(n)} &\leqslant \mu^{-1}_{\Lambda^{(n)}}(\mathcal{SG}) \leqslant \mu^{-1}_{\Lambda^{(n)}}(\mathcal{SG})\mu^{-1}_{\Lambda^{(n)}}(\mathcal{T}_{n-1}) \\ &\leqslant \min\left((\delta q^\alpha W^n)^{-W^n/\varepsilon^2}, e^{1/(\varepsilon^2 q^\alpha)}\right). \end{aligned} \tag{6.23}$$

From Lemmas 6.3.9 and 6.3.10 we get

$$\begin{aligned} \gamma\left(\Lambda^{(N^{\mathrm{i}})}\right) &\leqslant e^{\log^{O(1)}(1/q)}\prod_{n=0}^{N^{\mathrm{i}}-1}(\mu_{\Lambda^{(n+1)}}(\mathcal{SG})\mu_{\Lambda^{(n+1)}}(\mathcal{T}_n))^{-O(1)}\prod_{m=1}^{M^{(n)}} a_m^{(n)} \\ &\leqslant e^{\log^{O(1)}(1/q)}\prod_{n=0}^{N^{\mathrm{i}}-1}(\mu_{\Lambda^{(n+1)}}(\mathcal{SG})\mu_{\Lambda^{(n+1)}}(\mathcal{T}_n))^{-O(\log(1/q))} \\ &\leqslant \exp\left(\frac{\log(1/q)\log\log\log(1/q)}{\varepsilon^3 q^\alpha}\right). \end{aligned}$$

□

Proof of Lemma 6.3.9. Let us start with a general observation. Consider two regions $A, B \subset \mathbb{Z}^2$ and a measure ν on $\Omega_{A\cup B}$. The law of total variance reads

$$\mathrm{Var}_{\nu_{A\cup B}}(f) = \nu_{B\setminus A}\left(\mathrm{Var}_{\nu_A}(f)\right) + \mathrm{Var}_{\nu_{B\setminus A}}\left(\nu_A(f)\right).$$

The latter term can be bounded from above by

$$\begin{aligned}\mathrm{Var}_{\nu_{B\setminus A}}\left(\nu_A(f)\right) &\leqslant \nu_A\left(\mathrm{Var}_{\nu_{B\setminus A}}(f)\right)\\ &= \nu_{A\setminus B}\left(\nu_{A\cap B}(\mathrm{Var}_{\nu_{B\setminus A}}(f))\right) \leqslant \nu_{A\setminus B}\left(\mathrm{Var}_{\nu_B}(f)\right),\end{aligned}$$

using the convexity of the variance and the law of total variance. Hence,

$$\mathrm{Var}_{\nu_{A\cup B}} \leqslant \nu_{A\cup B}\left(\mathrm{Var}_{\nu_A} + \mathrm{Var}_{\nu_B}\right). \tag{6.24}$$

Fix $n \in [N^{\mathrm{i}}]$. Applying the above inequality several times, we obtain that

$$\mathrm{Var}_{\Lambda^{(n+1)}}(f|\mathcal{SG}) \leqslant \sum_{\kappa=2}^{K-1} \mu_{\Lambda^{(n+1)}}\left(\mathrm{Var}_{D_\kappa}\left(f|\mathcal{SG}\left(\Lambda^{(n+1)}\right)\right)|\mathcal{SG}\right) \tag{6.25}$$

and turn to bounding a generic summand.

We East-extend $\Lambda^{(n)}$ in direction u_0 by an arbitrarily large amount, which defines $\mathcal{SG}(D_\kappa)$ for all $\kappa \geqslant 2$ (it was already defined in Definition 6.3.5 for D_1 by East-extending in direction u_1). Observe that by the Harris inequality and the product structure of Definition 6.2.2, as in Eq. (6.8), for any $\kappa \in [2, K)$ Proposition 6.2.3 gives

$$\gamma(D_\kappa) \leqslant \max\left(\gamma\left(\Lambda^{(n)}\right), \mu^{-1}_{\Lambda^{(n)}}(\mathcal{SG})\right) e^{O(C^2)\log^2(1/q)} \frac{\mu_{\Lambda^{(n)}}(\mathcal{SG})}{\mu_{D_\kappa}(\mathcal{SG})} \prod_{m=1}^{M(\kappa)} a_m^{(n)} \tag{6.26}$$

with $M(\kappa) = \min\{m : \lambda_0(3/2)^m \geqslant l_\kappa\}$ and the same holds for D_1 with $M(1) = \min\{m : \lambda_1(3/2)^m \geqslant l_1\}$. Without loss of generality fix $\kappa = 2$, since all droplets of generation 2 are treated identically. Our goal is to show

$$\begin{aligned}\mu_{\Lambda^{(n+1)}}\left(\mathrm{Var}_{D_2}\left(f|\mathcal{SG}\left(\Lambda^{(n+1)}\right)\right)\Big|\mathcal{SG}\right) &\leqslant \frac{\gamma(\Lambda^{(n)})e^{O(C^2)\log^2(1/q)}}{(\mu_{\Lambda^{(n+1)}}(\mathcal{SG})\mu_{\Lambda^{(n+1)}}(\mathcal{T}_n))^{O(1)}}\\ &\times \prod_{m=1}^{M^{(n)}} a_m^{(n)} \sum_{y\in\Lambda^{(n+1)}} \mu_{\Lambda^{(n+1)}}\left(c_y^{\mathbf{1}}\,\mathrm{Var}_y(f)\right),\end{aligned} \tag{6.27}$$

from which Lemma 6.3.9 clearly follows in view of Eq. (6.25).

Let $V = D_1 \cup D_2$ (that is a $\dashv$-shaped region in Fig. 6.4) and

$$\mathcal{SG}(V) := \mathcal{SG}(D_1) \cap \mathcal{T}_n(D_2\setminus D_1).$$

By the product structure of traversability (see Definition 6.3.5), it is clear that

$$\mu_V\left(\cdot|\mathcal{SG}\left(\Lambda^{(n+1)}\right)\right)=\mu_V(\cdot|\mathcal{SG}(V))=\mu_{D_2\setminus D_1}(\cdot|\mathcal{T}_n)\otimes\mu_{D_1}(\cdot|\mathcal{SG}),$$

so that by convexity of the variance

$$\mu_{\Lambda^{(n+1)}}\left(\operatorname{Var}_{D_2}\left(f|\mathcal{SG}\left(\Lambda^{(n+1)}\right)\right)\middle|\,\mathcal{SG}\right)\leqslant\mu_{\Lambda^{(n+1)}}\left(\operatorname{Var}_V(f|\mathcal{SG})|\,\mathcal{SG}\right). \tag{6.28}$$

Further using a two-block dynamics (see e.g. Lemma 6.A.3), we have

$$\operatorname{Var}_V(f|\mathcal{SG})\leqslant\frac{\mu_V(\operatorname{Var}_{D_1}(f|\mathcal{SG})+\mathbb{1}_{\mathcal{E}}\operatorname{Var}_{D_2\setminus D_1}(f|\mathcal{T}_n)|\mathcal{SG})}{\Omega(\mu_V(\mathcal{E}|\mathcal{SG}))}, \tag{6.29}$$

where

$$\mathcal{E}=\mathcal{SG}(\Lambda^{(n)}+y_2u_1)\cap\mathcal{T}_n(D_1\cap D_2).$$

By convexity of the variance and the fact that

$$\mathcal{E}\cap\mathcal{T}_n(D_2\setminus D_1)\subset\mathcal{SG}(\Lambda^{(n)}+y_2u_1)\cap\mathcal{T}(D_2\setminus(\Lambda^{(n)}+y_2u_1))=\mathcal{SG}(D_2) \tag{6.30}$$

(recall Definitions 6.2.2 and 6.3.5 and the fact that each segment of length $\ell^{(n)}/(\varepsilon C)\gg\delta\ell^{(n)}/\varepsilon$ intersects at most $O(1)$ droplets), we have

$$\begin{aligned}\mu_V\left(\mathbb{1}_{\mathcal{E}}\operatorname{Var}_{D_2\setminus D_1}(f|\mathcal{T}_n)\middle|\,\mathcal{SG}\right)&\leqslant\frac{\mu(\mathcal{E})}{\mu_V(\mathcal{SG})}\mu_V(\operatorname{Var}_{D_2}(f|\mathcal{E}\cap\mathcal{T}_n(D_2\setminus D_1)))\\&\leqslant\frac{\mu(\mathcal{E})\mu_{D_2}(\mathcal{SG})\mu_V(\operatorname{Var}_{D_2}(f|\mathcal{SG}))}{\mu_V(\mathcal{SG})\mu(\mathcal{E}\cap\mathcal{T}_n(D_2\setminus D_1))}\\&\leqslant\frac{\mu_V(\operatorname{Var}_{D_2}(f|\mathcal{SG}))}{\mu^2_{\Lambda^{(n+1)}}(\mathcal{T}_n)}.\end{aligned} \tag{6.31}$$

Indeed, in the last line we recalled the definitions of $\mathcal{SG}(D_2)$ and $\mathcal{SG}(V)$, while in the second one we took into account that for any events $\mathcal{A}\subset\mathcal{B}$ with $\mu(\mathcal{A})>0$ it holds that

$$\operatorname{Var}(f|\mathcal{A})=\min_{c\in\mathbb{R}}\mu\left((f-c)^2\middle|\,\mathcal{A}\right)\leqslant\frac{\mu((f-\mu(f|\mathcal{B}))^2\mathbb{1}_{\mathcal{A}})}{\mu(\mathcal{A})}\leqslant\frac{\mu(\mathcal{B})}{\mu(\mathcal{A})}\operatorname{Var}(f|\mathcal{B}) \tag{6.32}$$

and Eq. (6.30). Plugging Eq. (6.31) in Eq. (6.29) and noting again that by the Harris inequality $\mu_V(\mathcal{E}|\mathcal{SG})\geqslant\mu(\mathcal{E})\geqslant\mu_{\Lambda^{(n)}}(\mathcal{SG})\mu_{\Lambda^{(n+1)}}(\mathcal{T}_n)$, to get

$$\operatorname{Var}_V(f|\mathcal{SG})\leqslant\frac{O(1)\mu_V(\operatorname{Var}_{D_1}(f|\mathcal{SG})+\operatorname{Var}_{D_2}(f|\mathcal{SG}))}{\mu_{\Lambda^{(n)}}(\mathcal{SG})\mu^3_{\Lambda^{(n+1)}}(\mathcal{T}_n)}. \tag{6.33}$$

Applying Eq. (6.26), we obtain

$$\gamma(V)\leqslant\frac{\gamma(\Lambda^{(n)})e^{O(C^2)\log^2(1/q)}}{\mu^2_{\Lambda^{(n)}}(\mathcal{SG})\mu^4_{\Lambda^{(n+1)}}(\mathcal{T}_n)}\prod_{m=1}^{M^{(n)}}a_m^{(n)},$$

since $M^{(n)} \geqslant \max_{\kappa\in[K]} M(\kappa)$. Inserting this in Eq. (6.28), we complete the proof of Eq. (6.27) and of Lemma 6.3.9. □

Proof of Lemma 6.3.10. The first inequality in Eq. (6.23) follows from the Harris inequality, while the second one is trivial, so we turn to the last one and fix $n \in [N^{\mathrm{i}}]$. Note that by Definition 6.3.5

$$\mu_{\Lambda^{(n+1)}}(\mathcal{SG}) \geqslant \mu_{\Lambda^{(n)}}(\mathcal{SG})\mu_{\Lambda^{(n+1)}}(\mathcal{T}_n)\mu_{D_1\setminus D_0}(\mathcal{T}). \tag{6.34}$$

We will therefore proceed by induction starting with

$$\mu_{\Lambda^{(0)}}(\mathcal{SG}) = q^{|\Lambda^{(0)}|} = q^{\Theta(1/\varepsilon^2)}. \tag{6.35}$$

Moreover, from Definition 6.3.5, to ensure the occurrence of $\mathcal{T}_n(\Lambda^{(n+1)})$, it suffices to have $O(WK\ell^{(n+1)})/(\ell^{(n)}\delta)$ well-placed W-helping sets, as well as $O((\ell^{(n+1)})^2)/(\ell^{(n)}\delta\varepsilon)$ helping sets for segments of length $\delta\ell^{(n)}/(3\varepsilon)$. Indeed, we may split lines perpendicular to each u_j for $j \in (-k,k)$ into successive disjoint segments of length $\delta\ell^{(n)}/(3\varepsilon)$ with a possible smaller leftover and place W-helping sets or helping sets depending on whether the segment under consideration is close to a parallel boundary of one of the D_κ. Recalling that $1/\varepsilon \gg 1/\delta \gg W \gg 1$, $\ell^{(N^{\mathrm{c}})} = W^{O(1)}q^\alpha$, $K = O(\ell^{(n+1)}/\ell^{(n)})$, the explicit expressions Eq. (6.6), the Harris inequality and Observation 6.1.10, we obtain

$$\begin{aligned}\mu_{\Lambda^{(n+1)}}(\mathcal{T}_n) &\geqslant q^{O(W^2K\ell^{(n+1)})/(\ell^{(n)}\delta)}\left(1-(1-q^\alpha)^{\delta\ell^{(n)}/O(\varepsilon)}\right)^{O((\ell^{(n+1)})^2/(\ell^{(n)}\delta\varepsilon))}\\ &\geqslant e^{-\log^{O(1)}(1/q)}\left(1-e^{-q^\alpha\delta\ell^{(n)}/O(\varepsilon)}\right)^{O((\ell^{(n+1)})^2/(\ell^{(n)}\delta\varepsilon))}\\ &\geqslant e^{-\log^{O(1)}(1/q)}\times\begin{cases}(\delta q^\alpha W^n)^{W^n/(\delta^2\varepsilon)} & n\leqslant N^{\mathrm{c}}\\ \exp\left(-1/\left(q^\alpha\exp\left(W^{\exp(n-N^{\mathrm{c}})}\right)\right)\right) & n>N^{\mathrm{c}}.\end{cases}\end{aligned} \tag{6.36}$$

Essentially the same computation leads to the same bound for $\mu_{D_1\setminus D_0}(\mathcal{T})$, the only difference being that only $O(1)$ W-helping sets and $O(\ell^{(n+1)}/\varepsilon)$ helping sets are needed. Further recalling Eqs. (6.34) and (6.35), it is elementary to check Eq. (6.23). □

Removing the surplus factor To conclude, let us briefly sketch how to remove the $\log\log\log(1/q)$ factor appearing in Theorem 6.3.8, which would also propagate to pollute Theorem 1.6.4(e).

Theorem 6.3.11. *Let $\mathcal{U}$ be balanced rooted. Instead of Definition 6.3.7, it is possible to define $\mathcal{SG}(\Lambda^{(N^{\mathrm{i}})})$ in such a way that*

$$\gamma\left(\Lambda^{(N^{\mathrm{i}})}\right) \leqslant \exp\left(\frac{\log^{O(1)}\log(1/q)}{q^\alpha}\right), \qquad \mu_{\Lambda^{(N^i)}}(\mathcal{SG}) \geqslant \exp\left(\frac{-1}{\varepsilon^2 q^\alpha}\right).$$

Sketch proof of Theorem 6.3.11. For a proof one should combine the techniques of both parts of Section 6.3.1. More precisely, a less crude bound on $a_m^{(n)}$ than Eq. (6.23) should be established along the lines of Eq. (6.11). As in Eq. (6.18), we may further decompose $a_m^{(n)}$ into a product over scales $p \leqslant n$. The relevant values of the parameters correspond to $(3/2)^m \leqslant 1/(\log^W(1/q)q^\alpha)$, say, and $p \in [N^c, n]$, as other cases can be dealt with using the crude bound Eq. (6.23). Further, as in Eq. (6.20), we can also discard $p \geqslant N^c + \Delta$. Hence, we need to focus for the remaining values of m and p on lower bounding

$$\mu\left(\mathcal{T}_p\left(\left(\Lambda^{(p+1)}\backslash D_1\right) + s_m\right)\middle| \mathcal{T}_p\left(\Lambda^{(p+1)}\backslash D_1\right)\right)$$

and $\mu(\mathcal{T}((D_1\backslash D_0) + s_m)|\mathcal{T}(D_1\backslash D_0))$, the latter being treated exactly like $\mu(\mathcal{T}(T')|\mathcal{T}(T))$ in Eq. (6.19). Turning to the former conditional probability, it can be further decomposed as a product over elementary regions delimited by the boundaries of the $(D_\kappa)_{\kappa\in[K]}$.

Unfortunately, for such (non-convex) polygonal regions R, bounding

$$\mu\left(\mathcal{T}_p\left(R + s_m\right)\middle| \mathcal{T}_p(R)\right)$$

is no easy feat. Indeed, Corollary 6.B.4 only treats tubes and, more importantly deals with helping sets for one direction only in each part of the tube (recall Fig. 6.2a), while $\mathcal{T}_p(R)$ requires helping sets in various directions, which are all dependent. To make matters worse, for certain families $\mathcal{U}$ it may happen that a single set of α infections is simultaneously a helping set for different directions and this would create complex and heavy dependency among different directions, which could, *a priori*, make boundary regions attract such sets.

To deal with this issue, one could further elaborate Definition 6.3.5. Indeed, we may split $\Lambda^{(p+1)}\backslash D_1$ into disjoint horizontal strips (recall Fig. 6.4b) of width $\ell^{(p)}/(W\varepsilon)$. Each strip is assigned a direction u_j, $j \in (-k, k)$ and we will only ask for helping sets for this direction to be present. These requirements are again cut essentially along the boundaries of all D_κ into parallelograms as in Lemma 6.B.3, placing W-helping sets on segments close to the boundaries. Naturally, some leftover regions remain without helping sets as in Definition 6.3.5, but they are unimportant as in for balanced rooted models.

By doing this, we make the event $\mathcal{T}_n(R)$ the intersection of traversability events of parallelograms in the sense of Lemma 6.B.3, so that its result can be applied as in the proof of Corollary 6.B.4, leading to a calculation similar to the one in Theorem 6.3.3. The only significant change is that now there are $O(W\ell^{(p+1)}/\ell^{(p)})$ parallelograms instead of a constant number. This is not really a problem, but, if one wishes to avoid careful computations, given that we are interested in the range $p \in (N^c, N^c + \Delta)$, we can brutally bound this by its maximum, which is $\log^{O(1)} \log(1/q)$ by the definition of Δ. □

6.3.2 Mesoscopic East dynamics

We next treat the East mesoscopic dynamics, which is essentially an extension of the internal one. Although they actually apply to all balanced models, the results of this section will only be used for balanced models with infinite stable directions, so the definitions will only apply to that class. As for that class there is a lot of margin, our reasoning will be far from tight for the sake of simplicity.

Extending the notation from Section 6.3.1 for balanced rooted models, for $n > N^{\mathrm{i}}$, we set $\ell^{(n)} = W^{n-N^{\mathrm{i}}}\ell^{(N^{\mathrm{i}})}$ and define $\underline{s}^{(n)}, \underline{r}^{(n)}, \Lambda^{(n)}$ as in that section. Further let $N^{\mathrm{m}} = \inf\{n : \ell^{(n)}/\varepsilon \geqslant \ell^{\mathrm{m}} = q^{-C}\} = \Theta(C\log(1/q)/\log W)$ and assume for simplicity that $\ell^{(N^{\mathrm{m}})} = q^{-C}\varepsilon$. We will only be interested in $n \leqslant N^{\mathrm{m}}$ and Definitions 6.3.5 and 6.3.7 remain unchanged for such n.

Theorem 6.3.12. *Let $\mathcal{U}$ be a balanced model with infinitely many stable directions (class (b)). Then*

$$\gamma\left(\Lambda^{(N^{\mathrm{m}})}\right) \leqslant \exp\left(\frac{\log^2(1/q)}{\varepsilon^3 q^\alpha}\right), \qquad \mu_{\Lambda^{(N^{\mathrm{m}})}}(\mathcal{SG}) \geqslant \exp\left(\frac{-2}{\varepsilon^2 q^\alpha}\right).$$

Proof. The proof is essentially identical to the one of Theorem 6.3.8, so we will only indicate the necessary changes. To start with, Lemma 6.3.9 applies without change for $n \in [N^{\mathrm{i}}, N^{\mathrm{m}})$. Also, the Harris inequality still implies that $a_m^{(n)} \leqslant \mu^{-1}_{\Lambda^{(n)}}(\mathcal{SG}) \leqslant \mu^{-1}_{\Lambda^{(N^{\mathrm{m}})}}(\mathcal{SG})$. Therefore,

$$\gamma\left(\Lambda^{(N^m)}\right) \leqslant \frac{\gamma(\Lambda^{(N^{\mathrm{i}})})e^{\log^{O(1)}(1/q)}}{(\mu_{\Lambda^{(N^{\mathrm{m}})}}(\mathcal{SG})\min_{n\in[N^{\mathrm{m}}]}\mu_{\Lambda^{n+1}}(\mathcal{T}_n))^{O(N^{\mathrm{m}}M^{(N^{\mathrm{m}}-1)})}}.$$

Recalling the bound of $\gamma(\Lambda^{(N^{\mathrm{i}})})$ established in Theorem 6.3.8, together with the fact that $N^{\mathrm{m}} \leqslant C\log(1/q)$ and $M^{(N^{\mathrm{m}}-1)} \leqslant O(C\log(1/q))$, it suffices to prove that

$$\mu_{\Lambda^{(N^{\mathrm{m}})}}(\mathcal{SG}) \min_{n\in[N^{\mathrm{m}}]} \mu_{\Lambda^{n+1}}(\mathcal{T}_n) \geqslant \exp(-2/(\varepsilon^2 q^\alpha)), \tag{6.37}$$

in order to conclude the proof of Theorem 6.3.12.

Once again, the proof of Eq. (6.37) proceeds similarly to the one of Eq. (6.23) in Lemma 6.3.10. Indeed, the same computation as Eq. (6.36) in the present setting gives that for $n \in [N^{\mathrm{i}}, N^{\mathrm{m}})$ we have

$$\mu_{\Lambda^{(n+1)}}(\mathcal{T}_n) \geqslant q^{O(W^3/\delta)} \exp\left(-e^{-q^\alpha\delta\ell^{(n)}/O(\varepsilon)} O\left(W^2\ell^{(n)}/(\delta\varepsilon)\right)\right). \tag{6.38}$$

From Eq. (6.34) it follows that

$$\mu_{\Lambda^{(N^{\mathrm{m}})}}(\mathcal{SG}) \geqslant \mu_{\Lambda^{(N^{\mathrm{i}})}}(\mathcal{SG}) \prod_{n=N^{\mathrm{i}}}^{N^{\mathrm{m}}-1} \mu_{\Lambda^{(n+1)}\setminus\Lambda^{(n)}}(\mathcal{T}_n).$$

Plugging Eqs. (6.23) and (6.38) in the r.h.s., this yields Eq. (6.37) as desired. □

6.4 CBSEP-type dynamics

In this section we establish results involving CBSEP-type dynamics. It is relevant only for isotropic models up to the internal scale and for unrooted models on mesoscopic level.

6.4.1 Isotropic internal and mesoscopic dynamics

Let $\mathcal{U}$ be isotropic. We follow and generalise Chapter 5.

Let $\underline{r}^{(0)}$ be a symmetric sequence of radii with $r_i^{(0)} = r_{i+2k}^{(0)}$ for all $i \in [2k]$, such that $r_i^{(0)} = \Theta(1/\varepsilon)$ is a multiple of λ_i for all $i \in [4k]$ and the corresponding side lengths $s_i^{(0)}$ are also $\Theta(1/\varepsilon)$. For any $i \in [2k]$ and $n = 2km + r$ with $r \in [2k]$ we define

$$s_i^{(n)} = s_{i+2k}^{(n)} = s_i^{(0)} 2^m \times \begin{cases} 2 & k \leqslant i < k + r \\ 1 & \text{otherwise} \end{cases}$$

and $\Lambda^{(n)} = \Lambda(\underline{r}^{(n)})$ with $\underline{r}^{(n)}$ the sequence of radii associated to $\underline{s}^{(n)}$ satisfying $r_i^{(n)} = r_{i+2k}^{(n)}$ for all $i \in [2k]$. Further set $N^{\mathrm{m}+} = 2k\lceil \log(\varepsilon\ell^{\mathrm{m}+})/\log 2\rceil$ (recall $\ell^{\mathrm{m}+}$ from Section 6.1.4, where m is not a variable and stands for 'mesoscopic').

Note that $\Lambda^{(n)}$ are nested symmetric droplets extended in one direction at each step satisfying $\Lambda^{(2km)} = 2^m \Lambda^{(0)}$. Recall Definition 6.2.4 and Fig. 6.2b.

Definition 6.4.1 (Isotropic SG)**.** Let $\mathcal{U}$ be isotropic. We say that $\Lambda^{(0)}$ is *SG* if it is fully infected. We then recursively define $\mathcal{SG}(\Lambda^{(n+1)})$ for $n \geqslant 0$ by CBSEP-extending $\Lambda^{(n)}$ in direction u_n by $l^{(n)} := s_{n+k}^{(n)} = \Theta(2^{n/2k}/\varepsilon)$ (recall that indices of directions and sequences are considered modulo $4k$ as needed and see Fig. 6.5a).

Theorem 6.4.2. *Let $\mathcal{U}$ be isotropic. Then for all $n \leqslant N^{\mathrm{m}+}$*

$$\gamma\left(\Lambda^{(N^{\mathrm{m}+})}\right) \leqslant \frac{\exp(1/(\log^{C/2}(1/q)q^\alpha))}{\mu(\mathcal{SG}(\Lambda^{(N^{\mathrm{m}+})}))}, \quad \mu\left(\mathcal{SG}\left(\Lambda^{(n)}\right)\right) \geqslant \exp\left(1/(q^\alpha\varepsilon^2)\right).$$

Proof of Theorem 6.4.2. We seek to apply Proposition 6.2.5, in order to recursively upper bound $\gamma(\Lambda^{(n)})$ for all $n \leqslant N^{\mathrm{m}+}$. To that end, we need the following definition.

Definition 6.4.3. Fix $2km + r = n \leqslant N^{\mathrm{m}+}$ with $r \in [2k]$ and let $\Lambda_1^{(n)} = T(\underline{r}^{(n)}, \lambda_r, r + 2k)$, $\Lambda_2^{(n)} = \Lambda(\underline{r}^{(n)} - \lambda_r \underline{v}_r)$ and $\Lambda_3^{(n)} = T(\underline{r}^{(n)}, \lambda_r, r) - \lambda_r u_r$ (as in Proposition 6.2.5 with $\underline{r} = \underline{r}^{(n)}$, $l = l^{(n)}$ and $i = r$). If $n < 2k$, we define $\overline{\mathcal{SG}}(\Lambda_2^{(n)})$ to occur if $\Lambda_2^{(n)}$ is fully infected and similarly for $\overline{\mathcal{ST}}(\Lambda_1^{(n)})$ and $\overline{\mathcal{ST}}(\Lambda_3^{(n)})$.

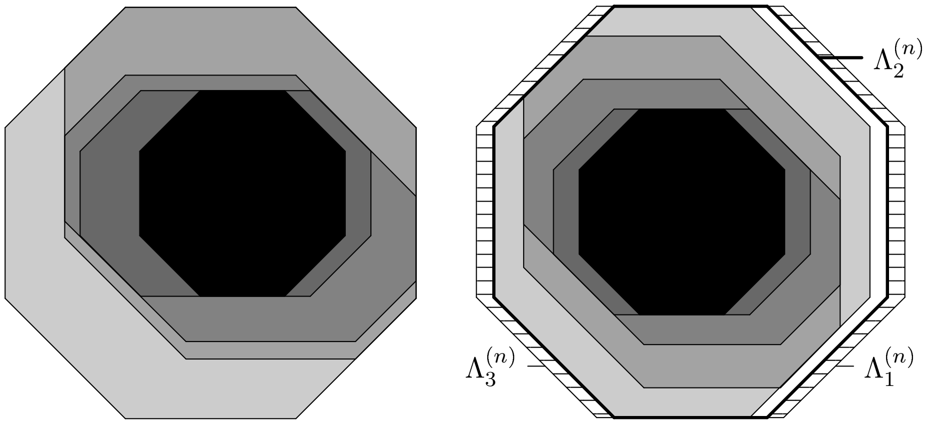

(a) A generic realisation of $\mathcal{SG}(\Lambda^{(n)})$ depicting the SG translates of $\Lambda^{(n)}, \dots, \Lambda^{(n-2k)}$ involved in progressive shades of grey. Each extension is as in Fig. 6.2b.

(b) The setting of Definition 6.4.3. The tubes $\Lambda_1^{(n)}$ and $\Lambda_3^{(n)}$ of width λ_r are hatched, $\Lambda_2^{(n)} = \Lambda^{(n)} \setminus \Lambda_3^{(n)}$ is thickened, while the droplets $\Lambda(\underline{r}^i)$, $0 \leqslant i \leqslant 2k$ are in progressive shades of grey, starting from the black $\Lambda(\underline{r}^0) = \Lambda^{(n-2k)}$.

Figure 6.5 – Geometry of isotropic $\mathcal{SG}$ and $\overline{\mathcal{SG}}$ events.

For $n \geqslant 2k$, we define $\overline{\mathcal{ST}}(\Lambda_1^{(n)})$ to be the event that for every segment $S \subset \Lambda_1^{(n)}$ perpendicular to some u_j with $j \neq r \pm k$ of length $2^m/(W\varepsilon)$ the event $\mathcal{H}^W(S)$ occurs. $\overline{\mathcal{ST}}(\Lambda_3^{(n)})$ is defined analogously and also does not depends on the configuration outside of $\Lambda_3^{(n)}$. Finally, for $n \geqslant 2k$, we define $\overline{\mathcal{SG}}(\Lambda_2^{(n)})$ as the intersection of the following events (see Fig. 6.5b).

- $\mathcal{SG}(\Lambda(\underline{r}^0))$, where $\underline{r}^0 = \underline{r}^{(n-2k)}$;
- $\mathcal{ST}_W(T(\underline{r}^0, l^{(n-2k)}/2 - \lambda_r, r)) \cap \mathcal{ST}_W(T(\underline{r}^0, l^{(n-2k)}/2 - \lambda_r, r + 2k))$;
- for all $i \in (0, 2k)$

$$\begin{aligned}&\mathcal{ST}_W\left(T\left(\underline{r}^{(n-2k+i)} - \lambda_r(\underline{v}_r + \underline{v}_{r+2k}), l^{(n-2k+i)}/2, r+i\right)\right)\\&\quad\cap \mathcal{ST}_W\left(T\left(\underline{r}^{(n-2k+i)} - \lambda_r(\underline{v}_r + \underline{v}_{r+2k}), l^{(n-2k+i)}/2, r+i+2k\right)\right).\end{aligned}$$

- for every $i \in [n-2k, n]$, $j \in [4k]$ and segment $S \subset \Lambda_2^{(n)}$ perpendicular to u_j of length $2^m/(W\varepsilon)$ at distance at most W from the u_j-side (parallel to S) of $\Lambda(\underline{r}^{(i)})$, the event $\mathcal{H}^W(S)$ holds.

In words, $\overline{\mathcal{SG}}(\Lambda_2^{(n)})$ is close to being the event that the central copy of $\Lambda^{(n-2k)}$ in $\Lambda_2^{(n)}$ is SG, the two tubes of equal length around it corresponding to a CBSEP-extension by $l^{(n-2k)}$ in direction u_r and so on until we reach

$\Lambda^{(n)}$ after $2k$ extensions. However, we have modified this event in the following ways. Firstly, the first extension is shortened by $2\lambda_r$, so that the final result after the $2k$ extensions fits inside $\Lambda_2^{(n)}$ and actually only its u_{r+k} and u_{r-k}-sides are shorter than those of $\Lambda_2^{(n)}$ by λ_r. Secondly, the traversability events for tubes are required to occur with segments shortened by W (recall Definition 6.2.1) on each side. Finally, we require W helping sets for the last $O(W)$ lines of each tube, as well as the first $O(W)$ outside the tube.

Taking this into account, we claim that Eq. (6.4) is verified. Indeed, if $\overline{\mathcal{SG}}(\Lambda_2^{(n)})$ and $\overline{\mathcal{ST}}(\Lambda_3^{(n)})$ occur, then the droplets $\Lambda^{(i)}$ for $i \in [n-2k, n]$ are all SG. To see this, proceed by induction and observe that each traversability appearing in Definition 6.4.3 together with the W-helping sets implies the corresponding traversability for the droplets $\Lambda^{(i)}$, since the droplets are perturbed (shifted position and modified side lengths) by $O(1)$, which is less than the amount, W, by which we the segments required in $\mathcal{ST}_W$ are contracted compared to those in $\mathcal{ST}$. In total, for each segment appearing implicitly in the $\mathcal{ST}$ events defining $\mathcal{SG}(\Lambda^{(i)})$ (via Definition 6.2.4), we have asked either for a W-helping set or a helping set in a slightly shorter segment.

Hence, we may apply Proposition 6.2.5 to get

$$\gamma\left(\Lambda^{(n+1)}\right) \leqslant \frac{\max(1, \mu_{\Lambda^{(n)}}(\mathcal{SG})\gamma(\Lambda^{(n)}))e^{O(C^2)\log^2(1/q)}}{\mu_{\Lambda^{(n+1)}}(\mathcal{SG})\mu_{\Lambda_1^{(n)}\cup\Lambda_2^{(n)}}(\overline{\mathcal{ST}}(\Lambda_1^{(n)}) \cap \overline{\mathcal{SG}}(\Lambda_2^{(n)})|\mathcal{SG})\mu_{\Lambda_3^{(n)}}(\overline{\mathcal{ST}}|\mathcal{ST}^{\mathbf{0}})} \tag{6.39}$$

for $n \geqslant 2k$ and $\gamma(\Lambda^{(n)}) \leqslant e^{O(C^2)\log^2(1/q)}$ for $n < 2k$. By the Harris inequality and the fact that both $\overline{\mathcal{ST}}(\Lambda_1^{(n)})$ and $\overline{\mathcal{ST}}(\Lambda_3^{(n)})$ can be guaranteed by $O(W^2)$ W-helping sets, we have

$$\mu_{\Lambda_3^{(n)}}\left(\overline{\mathcal{ST}}\middle|\,\mathcal{ST}^{\mathbf{0}}\right) \geqslant q^{W^{O(1)}}, \tag{6.40}$$

$$\mu_{\Lambda_1^{(n)}\cup\Lambda_2^{(n)}}\left(\overline{\mathcal{ST}}\left(\Lambda_1^{(n)}\right) \cap \overline{\mathcal{SG}}\left(\Lambda_2^{(n)}\right)\middle|\,\mathcal{SG}\right) \geqslant q^{W^{O(1)}}\mu\left(\overline{\mathcal{SG}}\left(\Lambda_2^{(n)}\right)\middle|\,\mathcal{SG}\left(\Lambda_1^{(n)}\cup\Lambda_2^{(n)}\right)\right). \tag{6.41}$$

Furthermore, by the Harris inequality, Lemma 6.B.1 and Corollary 6.B.4,

$$\mu\left(\overline{\mathcal{SG}}\left(\Lambda_2^{(n)}\right)\middle|\,\mathcal{SG}\left(\Lambda_1^{(n)}\cup\Lambda_2^{(n)}\right)\right) \geqslant \begin{cases}\mu\left(\overline{\mathcal{SG}}\left(\Lambda_2^{(n)}\right)\right) & 2^m \leqslant 1/\left(\log^C(1/q)q^\alpha\right), \\ q^{O(C)}\frac{\mu(\overline{\mathcal{SG}}(\Lambda_2^{(n)}))}{\mu\left(\mathcal{SG}\left(\Lambda^{(n-2k)}\right)\right)} & 2^m \geqslant \log^C(1/q)/q^\alpha, \\ \exp(-2^m q^{1-o(1)}) & \text{otherwise.}\end{cases} \tag{6.42}$$

Indeed, applying Lemma 6.B.1 $2k$ times gives that, conditionally on the event $\mathcal{SG}(\Lambda_1^{(n)}\cup\Lambda_2^{(n)})$, with probability $q^{O(C)}$ we have $\mathcal{SG}(\Lambda^{(n-2k)})$ and the traversability events corresponding to symmetrically extending it in $2k$ steps to $\Lambda_1^{(n)}\cup\Lambda_2^{(n)}$ all occur; in the range of m for which it applies, Corollary 6.B.4 deals with the conditional probability of the more restrictive traversability events required in Definition 6.4.3 given the original ones; W-helping sets in Definition 6.4.3 are paid for directly via the Harris inequality, since only $O(W^2)$ of them are needed. We direct the reader to Section 5.A for the details of an analogous argument in a simpler setting.

Iterating Eq. (6.39) and plugging Eqs. (6.40) to (6.42), we obtain

$$\gamma\left(\Lambda^{(N^{\mathrm{m}+})}\right) \leqslant \frac{e^{q^{-\alpha+1-o(1)}}}{\mu_{\Lambda^{(N^{\mathrm{m}+})}}(\mathcal{SG})} \prod_{n=2k}^{2^m\leqslant 1/(\log^C(1/q)q^\alpha)} \mu^{-1}\left(\overline{\mathcal{SG}}\left(\Lambda_2^{(n)}\right)\right) \times \prod_{n:2^m\geqslant \log^C(1/q)/q^\alpha}^{N^{\mathrm{m}+}} \frac{\mu(\mathcal{SG}(\Lambda^{(n-2k)}))}{\mu(\overline{\mathcal{SG}}(\Lambda_2^{(n)}))}.$$

Hence, Theorem 6.4.2 follows immediately from Lemma 6.4.4 below. □

Lemma 6.4.4. *The following bounds hold for $n\in[2k,N^{\mathrm{m}+}]$ and $m=\lfloor n/(2k)\rfloor$.*

$$\mu\left(\overline{\mathcal{SG}}\left(\Lambda_2^{(n)}\right)\right)\geqslant \exp\left(\frac{-1}{\log^{C-3}(1/q)q^\alpha}\right) \qquad \textit{if } 2\leqslant 2^m\leqslant \frac{1}{\log^C(1/q)q^\alpha}, \tag{6.43}$$

$$\frac{\mu(\overline{\mathcal{SG}}(\Lambda_2^{(n)}))}{\mu(\mathcal{SG}(\Lambda^{(n-2k)}))}\geqslant q^{W^{O(1)}} \qquad \textit{if } 2^m\geqslant \frac{\log^C(1/q)}{q^\alpha}, \tag{6.44}$$

$$\mu\left(\mathcal{SG}\left(\Lambda^{(n)}\right)\right)\geqslant \exp\left(\frac{-1}{q^\alpha\varepsilon^2}\right). \tag{6.45}$$

Proof. Let us first bound $\mu(\mathcal{SG}(\Lambda^{(n)}))$ for $n\leqslant N^{\mathrm{m}+}$ before easily deducing a bound on $\mu(\overline{\mathcal{SG}}(\Lambda_2^{(n)}))$. From Lemma 6.B.1 and Definition 6.2.4 we have

$$\mu_{\Lambda^{(n+1)}}(\mathcal{SG}) = q^{O(C)}\mu_{\Lambda^{(n)}}(\mathcal{SG})\mu_{T(\underline{r}^{(n)},l^{(n)},n)}(\mathcal{ST}), \tag{6.46}$$

so we need to bound the last term. Applying Observation 6.1.10 and Definition 6.2.1 as in Eq. (6.17), we get

$$\begin{aligned}\mu_{T(\underline{r}^{(n)},l^{(n)},n)}(\mathcal{ST}) &\geqslant q^{O(W)}\left(1-e^{-q^\alpha 2^m/O(\varepsilon)}\right)^{O(2^m/\varepsilon)} \\ &\geqslant q^{O(W)}\begin{cases}(\delta q^\alpha 2^m)^{2^m/(\delta\varepsilon)} & 2^m\leqslant 1/q^\alpha,\\ \exp\left(-2^m\exp\left(-q^\alpha 2^m\right)\right) & \text{otherwise.}\end{cases}\end{aligned} \tag{6.47}$$

Plugging this in Eq. (6.46) and iterating, we get

$$\mu\left(\Lambda^{(n)}\right) \leqslant \begin{cases} \exp\left(-1/\left(\log^{C-2}(1/q)q^{\alpha}\right)\right) & 2^m \leqslant 1/\left(\log^{C}(1/q)q^{\alpha}\right), \\ \exp\left(-1/\left(q^{\alpha}\varepsilon^2\right)\right) & \text{otherwise,} \end{cases}$$

proving Eq. (6.45).

Recalling Definition 6.4.3, we have that for any $n \in [2k, N^{\mathrm{m}+}]$

$$\mu\left(\overline{\mathcal{SG}}\left(\Lambda_2^{(n)}\right)\right) = q^{W^{O(1)}}\mu\left(\mathcal{SG}\left(\Lambda^{(n-2k)}\right)\right) \times \prod_{i\in[2k]} \mu^2\left(\mathcal{ST}_W\left(T\left(\underline{r}^i, l^{(n-2k+i)}/2, n+i\right)\right)\right).$$

However, the terms in the product can be bounded exactly as in Eq. (6.47), entailing Eqs. (6.43) and (6.44). □

6.4.2 Unbalanced unrooted models

Unbalanced internal dynamics

For unbalanced $\mathcal{U}$ with finite number of stable directions the internal dynamics is essentially trivial and so is the SG event up to the internal scale.

Definition 6.4.5 (Unbalanced internal SG)**.** If $\mathcal{U}$ is unbalanced with finite number of stable directions, we say that $\Lambda^{(0)} := \Lambda(\underline{r}^{(0)})$, defined by the side lengths $s_j^{(0)} = \lambda_j\lceil \ell^{\mathrm{i}}/\lambda_j\rceil$, is super good if all sites in $\Lambda^{(0)}\setminus\Lambda(\underline{r}^{(0)} - W\underline{1})$ are infected.

The following straightforward result was proved in Lemma 4.3.10.

Proposition 6.4.6. *For unbalanced $\mathcal{U}$ with finite number of stable directions we have*

$$\max\left(\gamma\left(\Lambda^{(0)}\right), \mu^{-1}\left(\mathcal{SG}\left(\Lambda^{(0)}\right)\right)\right) \leqslant q^{-O(W\ell^{\mathrm{i}})} \leqslant \exp\left(-C^3\log^2(1/q)/q^{\alpha}\right).$$

CBSEP mesoscopic dynamics

Let $\mathcal{U}$ be unbalanced unrooted with finite number of stable directions (the unbalanced character is only assumed, so as not to override definitions for semi-directed and isotropic models). W.l.o.g. let $\alpha(u_j) \leqslant \alpha$ for all $j \neq \pm k$ and $\min(\alpha(u_k), \alpha(u_{-k})) > \alpha$. We will only use $4k$ scales for the mesoscopic dynamics. Recall Section 6.1.4. For $i \in [0, 2k]$ let $\Lambda^{(i)} = \Lambda(\underline{r}^{(i)})$ be centered at 0 with $\underline{r}^{(i)}$ defined by

$$s_j^{(i)} = s_{j+2k}^{(i)} = \begin{cases} \lambda_j\lceil \ell^{\mathrm{i}}/\lambda_j\rceil & i-k \leqslant j < k \\ \lambda_j\lceil \ell^{\mathrm{m}-}/\lambda_j\rceil & -k \leqslant j < i-k. \end{cases}$$

For $i \in (2k, 4k]$, we define $\Lambda^{(i)}$ similarly by

$$s_j^{(i)} = s_{j+2k}^{(i)} = \begin{cases} \lambda_j \lceil \ell^{\mathrm{m}-}/\lambda_j \rceil & i-3k \leqslant j < k \\ \lambda_j \lceil \ell^{\mathrm{m}+}/\lambda_j \rceil & -k \leqslant j < i-3k. \end{cases} \tag{6.48}$$

These droplets are exactly as in Fig. 6.5a, except that the extensions are much longer. More precisely, we have $\Lambda^{(i+1)} = \Lambda(\underline{r}^{(i)} + l^{(i)}(\underline{v}_i + \underline{v}_{i+2k})/2)$ with $l^{(i)} = (1-o(1))\ell^{\mathrm{m}-}$ if $i \in [2k]$ and $l^{(i)} = (1-O(\delta))\ell^{\mathrm{m}+}$ if $i \in [2k, 4k)$.

Theorem 6.4.7. *Let $\mathcal{U}$ be unbalanced unrooted with finite number of stable directions. Then*

$$\max\left(\gamma\left(\Lambda^{(4k)}\right), \mu^{-1}\left(\mathcal{SG}\left(\Lambda^{(2k)}\right)\right)\right) \leqslant \exp\left(\frac{\log^2(1/q)}{\delta q^{\alpha}}\right).$$

Proof. We will proceed similarly to Theorem 6.4.2, but the first two steps are more special (see Fig. 6.6). For $i \in [4k]$, as in Proposition 6.2.5, let

$$\begin{aligned} \Lambda_1^{(i)} &= T\left(\underline{r}^{(i)}, \lambda_i, i+2k\right) \\ \Lambda_2^{(i)} &= \Lambda\left(\underline{r}^{(i)} - \lambda_i \underline{v}_i\right) \\ \Lambda_3^{(i)} &= T\left(\underline{r}^{(i)}, \lambda_i, i\right) - \lambda_r u_r. \end{aligned} \tag{6.49}$$

Definition 6.4.8. Let $\overline{\mathcal{ST}}(\Lambda_1^{(0)})$ (resp. $\overline{\mathcal{ST}}(\Lambda_3^{(0)})$) be the events that $\Lambda_1^{(0)}$ (resp. $\Lambda_3^{(0)}$) is fully infected and $\overline{\mathcal{SG}}(\Lambda_2^{(0)})$ be the event that $\Lambda_2^{(0)} \backslash \Lambda(\underline{r}^{(0)} - 2W\underline{1})$ is fully infected.

Let $\overline{\mathcal{SG}}(\Lambda_2^{(1)})$ occur if:

- $\mathcal{ST}_W(T(\underline{r}^{(0)} - \lambda_1 \underline{v}_1, l^{(0)}/2, 0))$ occurs,
- $(\Lambda((\ell^{\mathrm{i}} + W)\underline{1}) \backslash \Lambda((\ell^{\mathrm{i}} - 2W)\underline{1})) \cap \Lambda_2^{(1)}$ is fully infected,
- $\mathcal{ST}_W(T(\underline{r}^{(0)} - \lambda_1 \underline{v}_1, l^{(0)}/2, 2k))$ occurs,
- for all $j \neq \pm k$ and segment $S \subset \Lambda_2^{(1)}$ perpendicular to u_j at distance at most W from the u_j-side of $\Lambda_2^{(1)}$ and of length ℓ^{i}/W the event $\mathcal{H}^W(S)$ occurs.

Further let $\overline{\mathcal{ST}}(\Lambda_1^{(1)})$ occur if $\Lambda((\ell^{\mathrm{i}} + W)\underline{1}) \cap \Lambda_1^{(1)}$ is fully infected and for all $j \neq \pm k$ and segment $S \subset \Lambda_1^{(1)}$ perpendicular to u_j of length ℓ^{i}/W the event $\mathcal{H}^W(S)$ occurs. We define $\overline{\mathcal{ST}}(\Lambda_3^{(1)})$ analogously.

Let $i \in [2, 4k)$. We say that $\overline{\mathcal{ST}}(\Lambda_1^{(i)})$ occurs if for all $j \in [4k]$ and $m \in \{i, i-1\}$ every segment $S \subset \Lambda_1^{(i)}$ perpendicular to u_j at distance at most W from the u_j-side of $\Lambda^{(m)}$ of length $s_j^{(m)}/W$ the event $\mathcal{H}^W(S)$ occurs. We define $\overline{\mathcal{ST}}(\Lambda_3^{(i)})$ similarly and let $\overline{\mathcal{SG}}(\Lambda_2^{(i)})$ occur if:

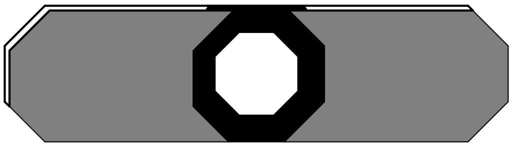

(a) Case $i = 1$.

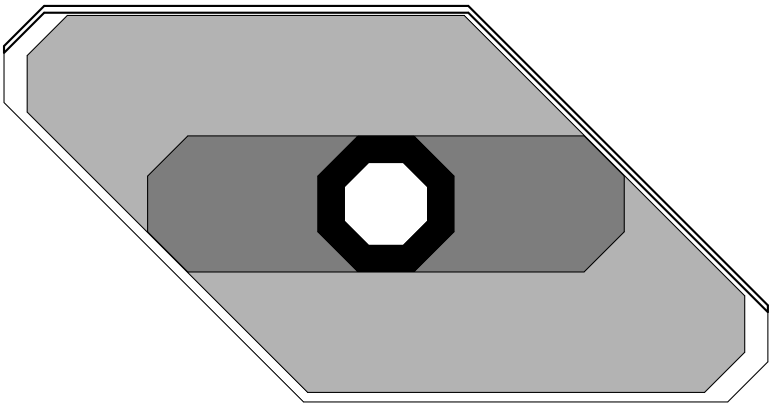

(b) Case $i > 1$. Regions around all boundaries contain W-helping sets.

Figure 6.6 – The events $\overline{\mathcal{SG}}(\Lambda_2^{(i)})$ and $\overline{\mathcal{ST}}(\Lambda_3^{(i)})$ of Definition 6.4.8. $\Lambda_3^{(i)}$ is thickened. Black regions are entirely infected. Shaded tubes are W-symmetrically traversable.

- $\mathcal{SG}(\Lambda^{(i-2)})$ occurs;
- for each $m \in \{0, 2k\}$ the following occurs

$$\mathcal{ST}_W\left(T\left(\underline{r}^{(i-2)}, l^{(i-2)}/2 - \sqrt{W}, i-2+m\right)\right)$$
$$\cap\, \mathcal{ST}_W\left(T\left(\underline{r}^{(i-1)} - \sqrt{W}\left(\underline{v}_i + \underline{v}_{i+2k}\right), l^{(i-1)}/2 - \sqrt{W}, i-1+m\right)\right);$$

- for all $j \in [4k]$, $m \in \{i-1, i\}$ and segment $S \subset \Lambda_2^{(i)}$ perpendicular to u_j of length $s_j^{(m)}/W$ at distance at most W from the u_j-side of $\Lambda^{(m)}$ the event $\mathcal{H}^W(S)$ holds.

With these definitions it is again not hard to verify Eq. (6.4) (see Fig. 6.6), so that Proposition 6.2.5, Proposition 6.4.6 and the Harris inequality give

$$\gamma\left(\Lambda^{(4k)}\right) \leqslant \frac{\exp(C^{O(1)}\log^2(1/q)/q^{\alpha})}{\prod_{i\in[4k]}\mu_{\Lambda^{(i+1)}}(\mathcal{SG})\mu_{\Lambda_1^{(i)}}(\overline{\mathcal{ST}})\mu_{\Lambda_2^{(i)}}(\overline{\mathcal{SG}})\mu_{\Lambda_3^{(i)}}(\overline{\mathcal{ST}})}. \qquad (6.50)$$

It is not hard to check from Definition 6.4.8 that each $\mathcal{SG}$, $\overline{\mathcal{SG}}$ and $\overline{\mathcal{ST}}$ event in Eq. (6.50) requires at most $C\ell^{\mathrm{i}}$ fixed infections, $W^{O(1)}$ W-helping sets and $O(1)$ W-symmetrically traversable tubes. We claim that the probability of each tube being W-symmetrically traversable is $q^{O(W)}$, which allows us to conclude, given Eq. (6.50). Indeed, as in Eq. (6.46), we have e.g.

$$\mu_{T(\underline{r}^{(0)}-\lambda_1\underline{v}_1,l^{(0)},0)}\left(\mathcal{ST}_W\right) \geqslant q^{O(W)}\exp\left(-O(\ell^{\mathrm{m}-})\exp\left(-q^{\alpha}\ell^{\mathrm{i}}/W\right)\right) \geqslant q^{O(W)}$$

by the choice of scales in Section 6.1.4. Traversability for $i > 1$ is slightly more subtle, since some of the parallelograms (recall Fig. 6.2b) require W-helping sets, since $\alpha(u_k) > \alpha$ and $\alpha(u_{-k}) > \alpha$. However, the u_k-side of $\Lambda^{(i)}$ for $i > 0$ has length $\Omega(\ell^{\mathrm{m}-})$, which is much larger than q^{-W}, so we can still conclude the proof of our claim in the same way, using Observation 6.1.10. □

6.4.3 Semi-directed mesoscopic dynamics

Let $\mathcal{U}$ be semi-directed and w.l.o.g. $\alpha(u_i) \leqslant \alpha$ for all $i \neq -k$. Recall from Section 6.3.1 that we defined $\Lambda^{(N^{\mathrm{i}})}$, a symmetric droplet with side lengths $\underline{s}^{(N^{\mathrm{i}})}$ equal to $\Theta(\ell^{(N^{\mathrm{i}})}/\varepsilon)$, as well as $\mathcal{SG}(\Lambda^{(N^{\mathrm{i}})})$ in Definition 6.3.2. As in Section 6.4.2 for unbalanced unrooted models with finite number of stable directions, for $i \in [N^{\mathrm{i}}+1, N^{\mathrm{i}}+2k]$ we define

$$s_j^{(i)} = s_{j+2k}^{(i)} = \begin{cases} s_j^{(N^{\mathrm{i}})} & i-N^{\mathrm{i}}-k \leqslant j < k, \\ \lambda_j\lceil \ell^{\mathrm{m}-}/\lambda_j\rceil & -k \leqslant j < i-N^{\mathrm{i}}-k, \end{cases}$$

while for $i \in (N^{\mathrm{i}}+2k, N^{\mathrm{i}}+4k]$, $\underline{s}^{(i)}$ is given by Eq. (6.48). We then define $\Lambda^{(N^{\mathrm{i}}+i)} = \Lambda(\underline{r}^{(N^{\mathrm{i}}+i)})$ with $\underline{r}^{(N^{\mathrm{i}}+i)}$ the sequence of radii associated to $\underline{s}^{(N^{\mathrm{i}}+i)}$ satisfying

$$\Lambda\left(\underline{r}^{(N_i+i)}\right) = \Lambda\left(\underline{r}^{(N^{\mathrm{i}}+i-1)} + l^{(N^{\mathrm{i}}+i-1)}\left(\underline{v}_{i-1}+\underline{v}_{i+2k-1}\right)/2\right),$$

with $l^{(N^{\mathrm{i}}+i-1)} = s_{i+k-1}^{(N^{\mathrm{i}}+i)} - s_{i+k-1}^{(N^{\mathrm{i}}+i-1)}$, which is $(1-o(1))\ell^{\mathrm{m}-}$ for $i \in [1, 2k]$ and $(1-O(\delta))\ell^{\mathrm{m}+}$ for $i \in (2k, 4k]$.

Theorem 6.4.9. *Let $\mathcal{U}$ be semi-directed. Then*

$$\gamma\left(\Lambda^{(N^{\mathrm{i}}+4k)}\right) \leqslant \exp\left(\frac{\log\log(1/q)}{\varepsilon^{O(1)}q^{\alpha}}\right), \quad \mu_{\Lambda^{(N^{\mathrm{i}}+2k)}}(\mathcal{SG}) \geqslant \exp\left(\frac{-1}{\varepsilon^{O(1)}q^{\alpha}}\right).$$

Proof. The proof proceeds exactly like Theorem 6.4.7, except that the first two steps are much more delicate and require taking into account the internal structure of $\mathcal{SG}(\Lambda^{(N^{\mathrm{i}})})$ on all scales down to 0, which is, alas, rather complex (recall Fig. 6.3a) and also not symmetric w.r.t. Λ_1 and Λ_3. This is not unexpected and is to some extent the crux of semi-directed models. As previously, for $i \in [N^{\mathrm{i}}, N^{\mathrm{i}}+4k)$ we define $\Lambda_1^{(i)}, \Lambda_2^{(i)}, \Lambda_3^{(i)}$ by Eq. (6.49). The next definitions are illustrated in Fig. 6.7.

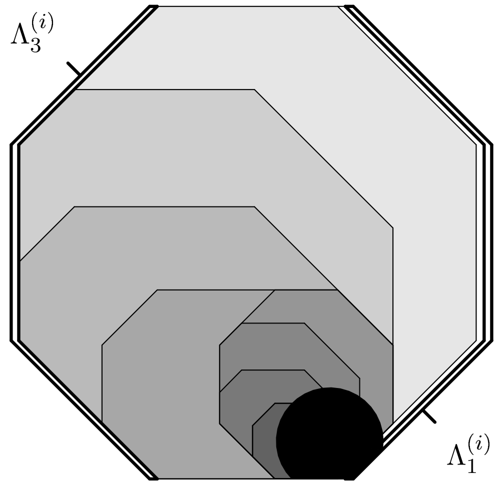

(a) Case $i = N^{\mathrm{i}}$ of Definition 6.4.10.

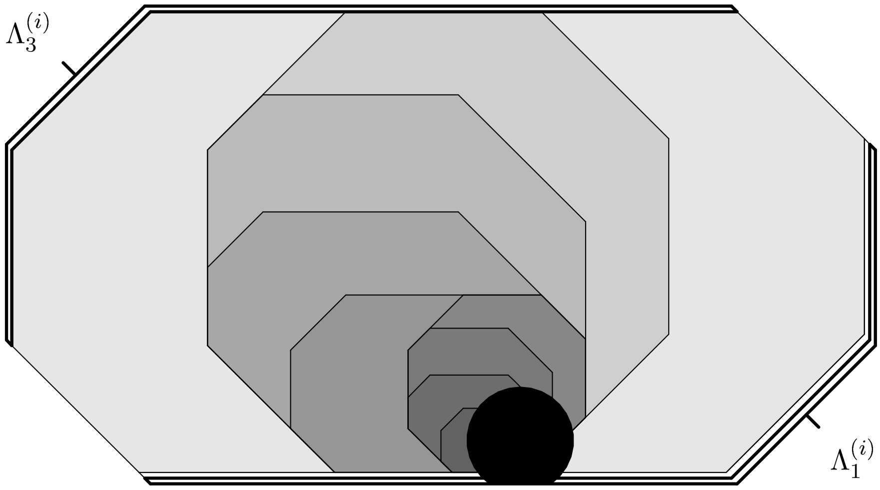

(b) Case $i = N^{\mathrm{i}} + 1$ of Definition 6.4.11.

Figure 6.7 – The events $\overline{\mathcal{ST}}(\Lambda_1^{(i)})$, $\overline{\mathcal{SG}}(\Lambda_2^{(i)})$ and $\overline{\mathcal{ST}}(\Lambda_3^{(i)})$. The microscopic black regions are entirely infected. Shaded tubes are W-traversable. W-helping sets are required close to all all boundaries.

Definition 6.4.10. Let $\overline{\mathcal{ST}}(\Lambda_3^{(N^{\mathrm{i}})})$ be the event that for all $j \neq \pm k$ every segment $S \subset \Lambda_3^{(N^{\mathrm{i}})}$ perpendicular to u_j of length $s_j^{(N^{\mathrm{i}})}/W$ the event $\mathcal{H}^W(S)$ occurs.

Let $\overline{\mathcal{ST}}(\Lambda_1^{(N^{\mathrm{i}})})$ be the event that for all $j \in (-k, k-1)$ every segment $S \subset \Lambda_1^{(N^{\mathrm{i}})}$ of length $s_j^{(N^{\mathrm{i}})}/W$ perpendicular to u_j the event $\mathcal{H}^W(S)$ occurs and all sites in $\Lambda_1^{(N^{\mathrm{i}})}$ at distance at most $\sqrt{W}/\varepsilon$ from the origin are infected.

For $n \in [0, N^{\mathrm{i}}]$ such that $2kn \in \mathbb{N}$ let $\Lambda'^{(n)} = \Lambda(\underline{r}^{(n)} - \lambda_0(\underline{v}_0 + \underline{v}_{2k}))$ and define $\mathcal{SG}'(\Lambda'^{(n)})$ recursively exactly like $\mathcal{SG}(\Lambda^{(n)})$ in Definition 6.3.2 with all droplets replaced by their contracted versions Λ' and all traversability events required in East-extensions replaced by the corresponding W-traversability events ($\mathcal{T}_W$, see Definition 6.2.1). Let $\mathcal{W}'$ be the event that for every $n \in [0, N^{\mathrm{i}}]$, $j \in [4k]$ and segment $S \subset \Lambda_2^{(N^{\mathrm{i}})}$ perpendicular to u_j of length $s_j^{(n)}/W$ at distance at most W from the u_j-side of $\Lambda^{(n)}$ the event $\mathcal{H}^W(S)$ holds. Let $\mathcal{I}'$ be the event that all sites in $\Lambda_2^{(N^{\mathrm{i}})}$ at distance at most $\sqrt{W}/\varepsilon$ from the origin are infected. Finally, set

$$\overline{\mathcal{SG}}\left(\Lambda_2^{(N^{\mathrm{i}})}\right) = \mathcal{SG}'\left(\Lambda'^{(N^{\mathrm{i}})}\right) \cap \mathcal{W}' \cap \mathcal{I}'.$$

Definition 6.4.11. Let $\overline{\mathcal{ST}}(\Lambda_1^{(N^{\mathrm{i}}+1)})$ be the event that for all $j \in (-k+1, k)$ and every segment $S \subset \Lambda_1^{(N^{\mathrm{i}}+1)}$ perpendicular to u_j of length $s_j^{(N^{\mathrm{i}})}/W$ the event $\mathcal{H}^W(S)$ occurs and all sites in $\Lambda_1^{(N^{\mathrm{i}}+1)}$ at distance at most $\sqrt{W}/\varepsilon$ from the origin are infected. Further let $\overline{\mathcal{ST}}(\Lambda_3^{(N^{\mathrm{i}}+1)})$ be the event that for all $j \in [4k]$, $m \in \{N^{\mathrm{i}}, N^{\mathrm{i}}+1\}$ and every segment $S \subset \Lambda_3^{(N^{\mathrm{i}}+1)}$ perpendicular to u_j of length $s_j^{(m)}/W$ at distance at most W from the u_j-side of $\Lambda^{(m)}$ the event $\mathcal{H}^W(S)$ occurs.

For $n \in [0, N^{\mathrm{i}}]$ such that $2kn \in \mathbb{N}$ let

$$\Lambda''^{(n)} = \Lambda\left(\underline{r}''^{(n)}\right) = \Lambda\left(\underline{r}^{(n)} - \lambda_1\left(\underline{v}_1 + \underline{v}_{2k+1}\right)\right)$$

and define $\mathcal{SG}''(\Lambda''^{(n)})$ like $\mathcal{SG}'(\Lambda'^{(n)})$ in Definition 6.4.10. Further let

$$\mathcal{SG}''\left(\Lambda''^{(N^{\mathrm{i}}+1)}\right) = \mathcal{SG}''\left(\Lambda''^{(N^{\mathrm{i}})}\right) \cap \bigcap_{j \in \{0, 2k\}} \mathcal{ST}_W\left(T\left(\underline{r}''^{(N^{\mathrm{i}})}, l^{(N^{\mathrm{i}})}/2, j\right)\right).$$

Let $\mathcal{W}''$ (resp. $\mathcal{I}''$) be defined like $\mathcal{W}'$ (resp. $\mathcal{I}'$) in Definition 6.4.10 with Λ' replaced by Λ'' and N^{i} replaced by $N^{\mathrm{i}}+1$. Finally, we set

$$\overline{\mathcal{SG}}\left(\Lambda_2^{(N^{\mathrm{i}}+1)}\right) = \mathcal{SG}''\left(\Lambda''^{(N^{\mathrm{i}}+1)}\right) \cap \mathcal{W}'' \cap \mathcal{I}''.$$

Furthermore, since Definition 6.4.8 for $i \in [2, 4k)$ does not inspect the internal structure of $\mathcal{SG}(\Lambda^{(0)})$ (see Fig. 6.6b), we may use the exact same definition for $\overline{\mathcal{ST}}(\Lambda_1^{(N^{\mathrm{i}}+i)})$, $\overline{\mathcal{SG}}(\Lambda_2^{(N^{\mathrm{i}}+i)})$ and $\overline{\mathcal{ST}}(\Lambda_3^{(N^{\mathrm{i}}+i)})$ with $i \in (1, 4k)$. Once again we can check from Definitions 6.4.10 and 6.4.11 and Fig. 6.7 that Eq. (6.4) holds. Hence, Proposition 6.2.5, Theorem 6.3.3 and the Harris inequality give

$$\gamma\left(\Lambda^{(N^{\mathrm{i}}+4k)}\right) \leqslant \frac{\exp(\log\log(1/q)/(\varepsilon^{O(1)}q^{\alpha}))}{\prod_{i=N^{\mathrm{i}}}^{N^{\mathrm{i}}+4k-1} \mu_{\Lambda^{(i+1)}}(\mathcal{SG})\mu_{\Lambda_1^{(i)}}(\overline{\mathcal{ST}})\mu_{\Lambda_2^{(i)}}(\overline{\mathcal{SG}})\mu_{\Lambda_3^{(i)}}(\overline{\mathcal{ST}})}. \tag{6.51}$$

It therefore remains to bound each of the terms in the denominator by $\exp(-1/(\varepsilon^{O(1)}q^{\alpha}))$ in order to conclude the proof of Theorem 6.4.9.

Notice that a total of $\varepsilon^{-O(1)}$ fixed infections and $W^{O(1)}N^{\mathrm{i}} = q^{o(1)}$ W-helping sets are required in all the events in Eq. (6.51), which amounts to a negligible factor. The probability of $\mathcal{SG}'(\Lambda'^{(N^{\mathrm{i}})})$ and $\mathcal{SG}''(\Lambda''^{(N^{\mathrm{i}})})$ can be bounded exactly like $\mathcal{SG}(\Lambda^{(N^{\mathrm{i}})})$ in Lemma 6.3.4, yielding a contribution of $\exp(-1/(\varepsilon^{O(1)}q^{\alpha}))$. Finally, the remaining bounded number of $\mathcal{ST}_W$ events are treated as in Theorem 6.4.7 to give a negligible $q^{-O(W)}$ factor. □

6.5 Global dynamics: assembly of Theorem 1.6.4

In this section we recall and adapt global dynamics mechanisms from Chapters 4 and 5 and assemble the pieces to prove our main result Theorem 1.6.4 for each refined universality class. As already noted, all lower bounds are proved in Chapter 8 and the upper ones for classes (a) and (c) were proved in [269] and Chapter 4 respectively, so we only need to establish the upper bounds for classes (b) and (d)-(g).

6.5.1 Global CBSEP dynamics

Let us recall the global CBSEP mechanism introduced in Chapter 5.

Let $\Lambda^{\mathrm{m}-}$ and $\Lambda^{\mathrm{m}+}$ be droplets with sides $\Theta(\ell^{\mathrm{m}-})$ and $\Theta(\ell^{\mathrm{m}+})$ respectively. Consider a tiling of $\mathbb{R}^2$ with square boxes $Q_{i,j} = [0, \ell^{\mathrm{m}}) \times [0, \ell^{\mathrm{m}}) + \ell^{\mathrm{m}}(i, j)$. We say that the box $Q_{i,j}$ is *good* if for every segment $S \subset Q_{i,j}$ perpendicular to some $u \in \widehat{\mathcal{S}}$ of length at least $\varepsilon\ell^{\mathrm{m}-}$, $\mathcal{H}^W(S)$ occurs and denote the corresponding event by $\mathcal{G}_{i,j}$. We say that it is *super good* if $Q_{i,j}$ contains a super good translate of $\Lambda^{\mathrm{m}-}$ and denote the corresponding event by $\mathcal{SG}_{i,j}$.

Proposition 6.5.1. *Let $T = \exp(-\log^4(1/q)/q^{\alpha})$. Assume that $\mathcal{SG}(\Lambda^{\mathrm{m}+})$ and $\mathcal{SG}(\Lambda^{\mathrm{m}-})$ are defined so that $(1 - \mu(\mathcal{SG}(\Lambda^{\mathrm{m}-})))^T T^4 = o(1)$ and for all $x \in \mathbb{Z}^2$ such that $x + \Lambda^{\mathrm{m}-} \subset \Lambda^{\mathrm{m}+}$ we have $\mathcal{SG}(x + \Lambda^{\mathrm{m}-}) \cap \mathcal{G}(\Lambda^{\mathrm{m}+}) \subset \mathcal{SG}(\Lambda^{\mathrm{m}+})$, where $\mathcal{G}(\Lambda^{\mathrm{m}+})$ stands for the event that for every segment $S \subset \Lambda^{\mathrm{m}+}$ perpendicular to some $u \in \widehat{\mathcal{S}}$ of length at least $3\varepsilon\ell^{\mathrm{m}-}$ the event $\mathcal{H}^W(S)$ occurs. Then*

$$\mathbb{E}_\mu(\tau_0) \leqslant \gamma\left(\Lambda^{\mathrm{m}+}\right) \frac{\log(1/\mu(\mathcal{SG}(\Lambda^{\mathrm{m}-})))}{q^{O(C)}}.$$

We omit the proof, which is identical to Section 5.4,[3] and turn to the proof of Theorem 1.6.4 for classes (d), (f) and (g).

Proof of Theorem 1.6.4(d). Let $\mathcal{U}$ be an unbalanced unrooted update family with finite number of stable directions. Recalling Section 6.4.2 for such families, let $\Lambda^{\mathrm{m}+} = \Lambda^{(4k)}$ and $\Lambda^{\mathrm{m}-} = \Lambda^{(2k)}$. By Theorem 6.4.7 (and Definitions 6.2.1 and 6.2.4) the hypotheses of Proposition 6.5.1 are satisfied and it yields

$$\mathbb{E}_\mu(\tau_0) \leqslant \exp\left(\frac{\log^2(1/q)}{\varepsilon q^\alpha}\right),$$

concluding the proof. □

Proof of Theorem 1.6.4(f). Let $\mathcal{U}$ be semi-directed. Recalling Section 6.4.3, let $\Lambda^{\mathrm{m}+} = \Lambda^{(N^{\mathrm{i}}+4k)}$ and $\Lambda^{\mathrm{m}-} = \Lambda^{(N^{\mathrm{i}}+2k)}$. By Theorem 6.4.9 (and Definitions 6.2.1 and 6.2.4) the hypotheses of Proposition 6.5.1 are satisfied and it yields

$$\mathbb{E}_\mu(\tau_0) \leqslant \exp\left(\frac{\log\log(1/q)}{\varepsilon^{O(1)} q^\alpha}\right),$$

concluding the proof. □

Proof of Theorem 1.6.4(g). Let $\mathcal{U}$ be isotropic. Recalling Section 6.4.1, let $\Lambda^{\mathrm{m}+} = \Lambda^{(N^{\mathrm{m}+})}$, $n^{\mathrm{m}-} = 2k\lceil\log(\varepsilon\ell^{\mathrm{m}-})/\log 2\rceil$ and $\Lambda^{\mathrm{m}-} = \Lambda^{(n^{\mathrm{m}-})}$. By Theorem 6.4.2 (and Definitions 6.2.1 and 6.2.4) the hypotheses of Proposition 6.5.1 are satisfied and it yields

$$\mathbb{E}_\mu(\tau_0) \leqslant \frac{\exp(1/(\log^{C/3}(1/q)q^\alpha))}{\mu(\mathcal{SG}(\Lambda^{(N^{\mathrm{m}+})}))} = \exp\left(\frac{1+o(1)}{\varepsilon^2 q^\alpha}\right),$$

concluding the proof. □

6.5.2 Global FA-1f dynamics

We next import the global FA-1f dynamics together with much of the mesoscopic multi-directional East one simultaneously from Chapter 4.

Proposition 6.5.2. *Let $\mathcal{U}$ have a finite number of stable directions, $T = \exp(-\log^4(1/q)/q^\alpha)$ and $\underline{r}^{\mathrm{i}}$ be such that the associated side lengths satisfy $C \leqslant s_j^{\mathrm{i}} \leqslant O(\ell^{\mathrm{i}})$ for all $j \in [4k]$. Assume that for all $l = \Theta(\ell^{\mathrm{m}})$ multiple of λ_0 the event $\mathcal{SG}(\Lambda(\underline{r}^{\mathrm{i}} + l\underline{v}_0))$ is defined so that $(1-\mu(\mathcal{SG}(\Lambda(\underline{r}^{\mathrm{i}}+l\underline{v}_0))))^T T^W = o(1)$. Then,*

$$\mathbb{E}_\mu(\tau_0) \leqslant \frac{\max_{l=\Theta(\ell^{\mathrm{m}})}\gamma(\Lambda(\underline{r}^{\mathrm{i}}+l\underline{v}_0))}{(q^{1/\delta}\min_{l=\Theta(\ell^{\mathrm{m}})}\mu(\mathcal{SG}(\Lambda(\underline{r}^{\mathrm{i}}+l\underline{v}_0))))^{\log(1/q)/\delta}}.$$

[3] Due to the difference between Eq. (6.2) and Eq. (5.7), the factor $\mu(\mathcal{SG}(\Lambda_{i,j}|\mathcal{G}))$ in the last display of Section 5.4 cancels out with $\pi(\mathcal{S}_1)^{-1}$ in Eq. (5.34) up to a $q^{O(C)}$ factor, rather than compensating the conditioning in the last display of Section 5.4, which is absent in our setting.

The proof is as in Chapter 4, the only difference being that one needs to replace the base of the snail there by $\Lambda^{\mathrm{m}} := \Lambda(\underline{r}^{\mathrm{i}} + \lambda_0\lceil \ell^{\mathrm{m}}/\lambda_0\rceil \underline{v}_0)$, which has a similar shape by hypothesis; the corresponding event that the base is super good by $\mathcal{SG}(\Lambda^{\mathrm{m}})$; and Proposition 4.3.9 by the definition Eq. (6.2) of $\gamma(\Lambda^{\mathrm{m}})$. As this proposition is essentially the entire content of Chapter 4 (see particularly Proposition 4.3.12 and Remark 4.3.8), we refer the reader to that chapter for the details.

Proof of Theorem 1.6.4(e). Let $\mathcal{U}$ be balanced rooted with finite number of stable directions. Recall $\Lambda^{(N^{\mathrm{i}})} = \Lambda(\underline{r}^{(N^{\mathrm{i}})})$ and $\underline{r}^{(N^{\mathrm{i}})} =: \underline{r}^{\mathrm{i}}$ from Section 6.3.1. Fix $l = \Theta(\ell^{\mathrm{m}})$ multiple of λ_0 and East-extend $\Lambda^{(N^{\mathrm{i}})}$ by l in direction u_0. It is easy to check from Definition 6.2.2 and Observation 6.1.10 that

$$\frac{\mu(\mathcal{SG}(\Lambda(\underline{r}^{\mathrm{i}} + l\underline{v}_0)))}{\mu(\mathcal{SG}(\Lambda(\underline{r}^{\mathrm{i}})))} = \mu\left(\mathcal{T}\left(T\left(\underline{r}^{\mathrm{i}}, l, 0\right)\right)\right) = q^{O(W)}.$$

Then, by Proposition 6.2.3, Theorem 6.3.11 and the Harris inequality, we obtain

$$\gamma\left(\Lambda(\underline{r}^{\mathrm{i}} + l\underline{v}_0)\right) \leqslant \exp\left(\frac{\log(1/q)}{\varepsilon^3 q^\alpha}\right), \qquad \mu_{\Lambda(\underline{r}^{\mathrm{i}}+l\underline{v}_0)}(\mathcal{SG}) \geqslant \exp\left(\frac{-2}{\varepsilon^2 q^\alpha}\right).$$

Plugging this in Proposition 6.5.2, we obtain

$$\mathbb{E}_\mu(\tau_0) \leqslant \exp\left(\frac{2\log(1/q)}{\varepsilon^3 q^\alpha}\right),$$

which concludes the proof. □

6.5.3 Global East dynamics

Finally, for class (b) we will need a simpler version of the procedure of Section 4.4 with East dynamics instead of FA-1f.

Proof of Theorem 1.6.4(b). Let $\mathcal{U}$ be balanced with infinite number of stable directions, $T = \exp(1/q^{3\alpha})$ and $\underline{s}^{\mathrm{m}} = \underline{s}^{(N^{\mathrm{m}})}$, $\underline{r}^{\mathrm{m}} = \underline{r}^{(N^{\mathrm{m}})}$ and $\Lambda^{\mathrm{m}} = \Lambda^{(N^{\mathrm{m}})}$ with the notation of Section 6.3.2. In particular, $s_j^{\mathrm{m}} = \Theta(\ell^{\mathrm{m}})$ for $j \in [-k, k+1]$ and $s_j^{\mathrm{m}} = O(\ell^{\mathrm{m}})$ for $j \in [k+2, 3k-1]$. We East-extend Λ^{m} by $2l = 2(\lambda_0 + r_0^{\mathrm{m}} + r_{2k}^{\mathrm{m}})$ in direction u_0 to obtain $\Lambda = \Lambda(\underline{r}^{\mathrm{m}} + 2l\underline{v}_0)$. Proposition 6.2.3, Theorem 6.3.12, the Harris inequality and the simple fact that $\mu(\mathcal{T}(T(\underline{r}^{\mathrm{m}}, 2l, 0))) = q^{O(W)}$ (recall Observation 6.1.10) give

$$\gamma(\Lambda) \leqslant \exp\left(\frac{\log^2(1/q)}{\varepsilon^{O(1)} q^\alpha}\right), \qquad \mu_\Lambda(\mathcal{SG}) \geqslant \exp\left(\frac{-3}{\varepsilon^2 q^\alpha}\right). \tag{6.52}$$

A similar argument to the rest of the proof was already discussed thoroughly in Section 4.4 and then in Section 5.4, so we will only provide a sketch. The adapted approach of Section 4.4 proceeds as follows.

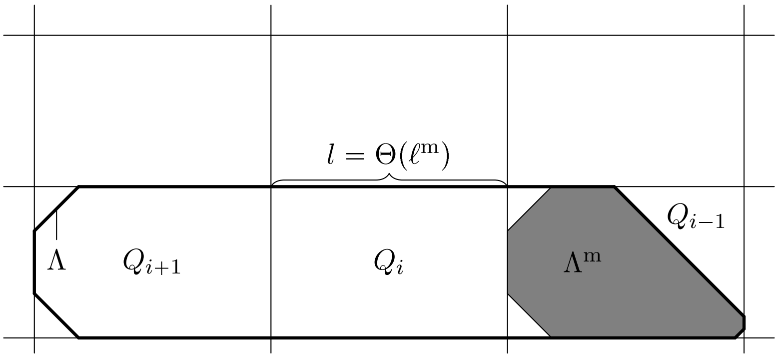

Figure 6.8 – Illustration of the global East dynamics (Section 6.5.3). The shaded droplet Λ^{m} inscribed in the box Q is extended by $2l$ to the thickened one Λ.

(1) Denoting $t_* = \exp(-1/(\varepsilon^W q^{2\alpha}))$, by the main result of [269] it suffices to show that $T\mathbb{P}_\mu(\tau_0 > t_*) = o(1)$, in order to deduce $\mathbb{E}_\mu(\tau_0) \leqslant t_* + o(1)$.

(2) By finite speed of propagation we may work with the $\mathcal{U}$-KCM on a large the discrete torus of size $T \gg t_*$.

(3) We partition the torus into strips and the strips into translates of the box $Q = \mathbb{H}_{u_0}(\lambda_0 + r_0^{\mathrm{m}}) \cap \mathbb{H}_{u_k}(\rho_k + r_k^{\mathrm{m}}) \cap \overline{\mathbb{H}}_{u_{-k}}(r_{-k}^{\mathrm{m}}) \cap \overline{\mathbb{H}}_{u_{2k}}(r_{2k}^{\mathrm{m}})$ as shown in Fig. 6.8. We say Q is *good* ($\mathcal{G}(Q)$ occurs) if for each segment $S \subset Q$ perpendicular to some $u \in \widehat{\mathcal{S}}$ of length $\varepsilon\ell^{\mathrm{m}}$ the event $\mathcal{H}^W(S)$ occurs. Further define $\mathcal{SG}(Q)$ to occur if the (only integer) translate of Λ^{m} contained in Q is SG. We say that *the environment is good* ($\mathcal{E}$ occurs) if all boxes are and in each strip at least one box is super good. The sizes are chosen so that it is sufficiently likely for this event to always occur up to t_*. Indeed, we have $(1-\mu(\mathcal{SG}(\Lambda^{\mathrm{m}})))^T T^W = o(1)$ by Theorem 6.3.12 and $(1-\mu_Q(\mathcal{G}))T^W = o(1)$ by Observation 6.1.10.

(4) By a standard variational technique it then suffices to prove a Poincaré inequality, bounding the variance of a function conditionally on $\mathcal{E}$ by the Dirichlet form on the torus. Moreover, since μ and $\mathcal{E}$ are product w.r.t. the partition of Fig. 6.8, it suffices to prove this inequality on a single strip.

(5) Finally, we prove such a bound, using an auxiliary East dynamics for the boxes and the definition of γ to reproduce the resampling of the state of a box by moves of the original $\mathcal{U}$-KCM.

Let us explain the last step above in more detail, as it is the only one that genuinely differs from Chapter 4.

Let $Q_i = Q + ilu_0$ and $\mathbb{T} = \bigcup_{i\in[T]} Q_i$ be our strip of interest (indices are considered modulo T, since the strip is on the torus). As explained above, our goal is to prove that for all $f : \Omega_{\mathbb{T}} \to \mathbb{R}$ it holds that

$$\operatorname{Var}_{\mathbb{T}}(f|\mathcal{E}) \leqslant \exp\left(1/\left(\varepsilon^{O(1)}q^{2\alpha}\right)\right)\sum_{x\in\mathbb{T}}\mu_{\mathbb{T}}\left(c_x^{\mathbf{1}}\operatorname{Var}_x(f)\right), \tag{6.53}$$

where $c_x^{\mathbb{T},\mathbf{1}}$ takes into account the circular geometry of $\mathbb{T}$.

By [269, Proposition 3.4] on the generalised East chain we have

$$\operatorname{Var}_{\mathbb{T}}(f|\mathcal{E}) \leqslant \exp\left(1/\left(\varepsilon^5 q^{2\alpha}\right)\right)\sum_{i\in[T]}\mu_{\mathbb{T}}\left(\mathbb{1}_{\mathcal{SG}(Q_{i-1})}\operatorname{Var}_{Q_i}(f|\mathcal{G})\middle|\mathcal{E}\right), \tag{6.54}$$

since Theorem 6.3.12 gives $\mu_Q(\mathcal{SG}) \geqslant \exp(-2/(\varepsilon^2 q^\alpha))$.[4]

Next observe that $\Lambda_i \supset Q_i$, where $\Lambda_i := \Lambda + (i-1)lu_0$ (see Fig. 6.8). Hence, by convexity of the variance and the fact that $\mu(\mathcal{E}) = 1 - o(1)$ and we have

$$\begin{aligned}\mu_{\mathbb{T}}&\left(\mathbb{1}_{\mathcal{SG}(Q_{i-1})}\operatorname{Var}_{Q_i}(f|\mathcal{G})\middle|\mathcal{E}\right)\\ &\leqslant (1+o(1))\mu_{\mathbb{T}}\left(\operatorname{Var}_{\Lambda_i}(f|\mathcal{SG}(Q_{i-1})\cap\mathcal{G}(Q_i)\cap\mathcal{G}(\Lambda_i\backslash Q_i))\right),\\ &\leqslant (1+o(1))\mu_{\mathbb{T}}\left(\operatorname{Var}_{\Lambda_i}(f|\mathcal{SG})\right),\end{aligned}$$

writing $\mathcal{G}(\Lambda_i\backslash Q_i) \subset \mathcal{G}(Q_{i+1}) \cap \mathcal{G}(Q_{i-1})$ for the event that $\mathcal{H}^W(S)$ holds for all segments $S \subset \Lambda_i\backslash Q_i$ of length $2\varepsilon\ell^{\mathrm{m}}$ perpendicular to some $u \in \widehat{\mathcal{S}}$ and using Eq. (6.32) and $\mathcal{SG}(Q_{i-1}) \cap \mathcal{G}(Q_i) \cap \mathcal{G}(\Lambda_i\backslash Q_i) \subset \mathcal{SG}(\Lambda_i)$ (recall Definition 6.2.2) for the second inequality. Finally, recalling Eqs. (6.2), (6.52) and (6.54), we obtain Eq. (6.53) as desired. □

Appendix

6.A Extensions

Recall Definition 6.1.7. Let $\underline{r}$ be a sequence of radii and $\Lambda = \Lambda(\underline{r})$. Given $\omega \in \Omega_{\mathbb{Z}^2\backslash\Lambda}$ and $i \in [4k]$, we define $\Lambda \subset \Lambda_i^\omega \subset \Lambda(\underline{r} + O(1)\underline{v}_i)$ as

$$\Lambda_i^\omega = \begin{cases}\Lambda & \alpha(u_i) > \alpha,\\ \Lambda \cup \bigcup_x \left(\left([Z_i \cup \mathbb{H}_{u_i}]_{\mathcal{U}}\backslash\overline{\mathbb{H}}_{u_i}\right) + x\right)\backslash\left\{y \in \mathbb{Z}^2\backslash\Lambda : \omega_y = 0\right\} & \text{otherwise,}\end{cases}$$

the union being over $x \in \Lambda$ such that: $\omega_{x+Z_i\backslash\Lambda} = \mathbf{0}$ and x is at distance at least C from all sides of Λ except the u_i-side. In words, we essentially look at pieces of u_i-helping sets for the last few lines of the droplet and add to Λ the sites which each piece can infect. The reason for introducing this is that helping sets may need to infect a few sites outside Λ before creating their periodic infections on the corresponding line and it is those sites that we wish to include in Λ_i^ω. We set $\Lambda_I^\omega = \bigcup_{i\in I}\Lambda_i^\omega$ for $I \subset [4k]$.

[4] Strictly speaking [269] does not deal with the torus conditioned on having an infection, but this issue is easily dealt with by the method of [59].

6.A.1 Microscopic dynamics

Let $i \in [4k]$ be such that $\alpha(u_j) < \infty$ for all $j \in I = \{i-k+1, \ldots, i+k-1\}$. Fix $\Lambda = \Lambda(\underline{r})$ with sides at least $\Omega(1/\delta)$ and at most $q^{-O(C)}$. Let $l \in [0, O(1)]$, $\omega \in \Omega_{\mathbb{Z}^2\setminus\Lambda(\underline{r}+l\underline{v}_i)}$, $\Lambda^+ = (\Lambda(\underline{r}+l\underline{v}_i))_I^\omega$ and $T = T(\underline{r}, l, i)$. Our goal is to provide a relaxation mechanism for an East-extension of bounded length.

Lemma 6.A.1. *In the above setting we have*

$$\mu_{\Lambda^+\setminus\Lambda}(\mathrm{Var}_T(f|\mathcal{T}^\omega)) \leqslant \exp\left(O\left(C^2\right)\log^2(1/q)\right) \times \sum_{x\in\Lambda^+\setminus\Lambda} \mu_{\Lambda^+\setminus\Lambda}\left(c_x^{\mathbf{0}_\Lambda\cdot\omega_{\mathbb{Z}^2\setminus\Lambda^+}} \mathrm{Var}_x(f)\right) \tag{6.55}$$

and the same holds for $\mathcal{ST}$ instead of $\mathcal{T}$.

The proof is both standard and messy, so we only provide a sketch.

Sketch proof. By convexity of the variance, it is enough to upper bound $\mathrm{Var}_{\Lambda^+\setminus\Lambda}(f|\mathcal{T}^\omega(T))$. We will use the canonical path technique (see e.g. [310, Theorem 4.2.1]), so we need to define for any two configurations $\eta, \eta' \in \Omega_{\Lambda^+\setminus\Lambda} \cap \mathcal{T}^\omega(T) =: \mathcal{A}$ a sequence $\Gamma(\eta, \eta')$ of configurations in $\mathcal{A}$ differing by single legal updates of the $\mathcal{U}$-KCM with boundary condition $\mathbf{0}_\Lambda \cdot \omega_{\mathbb{Z}^2\setminus\Lambda^+}$, leading from η to η'. We call such sequences *canonical paths.*

Recalling the notation of Definition 6.1.7, for $j \in I$ such that $\alpha(u_j) \leqslant \alpha$ let X_j denote the intersection of $[Z_j \cup \mathbb{H}_{u_j}]_\mathcal{U}$ with a sufficiently long segment of $\overline{\mathbb{H}}_{u_j}\setminus\mathbb{H}_{u_j}$ such that for all $x \in X_j$ we have that $\sqrt{W} \leqslant \langle x, x_j\rangle \leqslant 2\sqrt{W}$. It is easy to see that if $\mathbb{H}_{u_j}$ is fully infected, then X_j can infect $X_j + Qx_j$ only modifying states in $\overline{\mathbb{H}}_{u_j}\setminus\mathbb{H}_{u_j}$. Moreover, if $Z_j \cup \mathbb{H}_{u_j}$ is infected, then X_j can be infected in at most $O(\sqrt{W})$ steps. Finally, observe that a W-helping set in $\overline{\mathbb{H}}_{u_j}\setminus\mathbb{H}_{u_j}$ can move freely in both directions along the line.

Let us first describe the path in the case when T consists of a single line perpendicular to, say, u_j. Our paths will proceed in four stages. First, starting from η, we infect $O(\sqrt{W})$ sites until we infect translates of X_j with all possible residues modulo Q along the line. We next infect the last W sites of the line. Then we change η with that W-helping set to η' with the same W-helping set. Finally, we reach η', which can be done as in the reverse of the first two stages, so we will only describe the first three stages. Note that if $\alpha(u_j) > \alpha$, the first two stages are not needed, as W-helping sets are guaranteed by $\mathcal{T}^\omega$ (which does not depend on ω in that case).

In the first stage we simply add the infections of $[Z_j \cup \mathbb{H}_{u_j}]_\mathcal{U}\setminus\mathbb{H}_{u_j}$ translated appropriately one by one until we infect the translate of X_j in $O(\sqrt{W})$ steps. Naturally, we do this for Q different translates, so as to obtain each residue.

In the second stage we perform an East motion of translates of X_j, starting from the ones we infected in the first stage. As noted above, thinking of $\mathbb{H}_{u_j}$ as infected, using X_j we can infect $X_j + Qx_j$ (which may intersect X_j or other infected sites of η, in which case we only infect the additional sites), then use $X_j + Qx_j$ to infect $X_j + 2Qx_j$, then use $X_j + Qx_j$ to remove any infections in $X_j + Qx_j$ which are not in X_j, $X_j + 2Q_j$ or η. We continue similarly, as described for the East model in Section 1.6.5 (also see [8, Fig. 2]). Doing this, we can eventually infect the last W/Q translates of X_j of the form $X_j + kQx_j$ in T. Repeating this for all Q translates with different residues modulo Q, we obtain the desired last W infections in the line. We finally remove all other auxiliary infections by reversing the same path. Note that by our choice of directions $\widehat{\mathcal{S}}$, it is indifferent whether $\mathbb{H}_{u_j}$ is infected or only Λ.

In the third stage we move the W-helping set to the other extremity of the segment, leaving behind η'. More precisely, if the W infections are currently at position x, we infect $x - \lambda_{j+k}u_{j+k}$ and then remove the infection at $x + (W-1)\lambda_{j+k}u_{j+k}$ (the last site of the W-helping set) if and only if it is not present in η'. Moreover, with a finite number of infections we can also infect any site $\Lambda_j^\omega \setminus \Lambda$ at distance $O(1)$ from the W-helping set, but at large enough distance from its extremities. Thus, as we move the W-helping set, we can ensure that all sites in $\Lambda_j^\omega \setminus \Lambda$ on one side of its midpoint have the state η' and the others have state η. Finally, when η is completely replaced by η', we move the W infections back to the end of the segment, still leaving η' behind.

To prove the lemma in the case of a single line, it suffices to bound the length of these paths, as well as

$$\max_{\eta'' \in \mathcal{A}} \sum_{\eta'' \in \Gamma(\eta, \eta')} \frac{\mu(\eta)\mu(\eta')}{\mu(\eta'')}, \tag{6.56}$$

where we neglected $|T|$ (which is only polynomial in q) and used that by the Harris inequality and Observation 6.1.10 $\mu(\eta) = q^{O(W)}\mu(\eta|\mathcal{A})$. The length of the paths is polynomial in $|T|$, as for the East process (see [8]). Hence, we turn to upper bounding (6.56). Observe that in the first and second stages the configurations η'' only differ from η in at most $|X_j|O(\log|T|)$ sites (this is the fundamental property of East paths), so $\mu(\eta)/\mu(\eta'') \leqslant q^{-O(\log|T|)}$ and there are at most $(O(|T|))^{O(\log|T|)}$ possible choices for the discrepancies, hence the contribution to Eq. (6.56) is at most $\exp(O(C^2)\log^2(1/q))$. In the third stage at all times there are at most $O(W)$ sites on which the configuration differs from both η and η' and $O(|T|)$ possible positions of the W-helping set (we know that on one of its sides we have η and η' on the other). Recalling that μ is product, we obtain a contribution of only $q^{O(C)}$ to Eq. (6.56), which concludes the proof.

In order to treat an arbitrary number of lines and take into account that the tube T is composed of segments in different directions, we proceed iteratively. Firstly, the structure of $\mathcal{T}^\omega$ (recall Fig. 6.2) makes different directions independent, and they do not overlap, so we may proceed one direction at a time. To treat several consecutive lines in the same direction, we first produce the W-helping set for the first line as above, then use it to act as an infected boundary condition for the second line, which we may place next to the site we want to update. This way we can also create a W-helping set for the second line and so on. Eventually, we have W-helping sets for all lines and we may perform the third stage, moving all of them simultaneously to change η into η'. We finally remove these W-helping sets by reversing the path from η' that would create them. The computation of the congestion of the path identical to the simpler one-line case. $\square$

Corollary 6.A.2. *In the same setting as above, we have*

$$\mu_{\Lambda^+}\left(\mathrm{Var}_T(f|\mathcal{T}^\omega)|\,\mathcal{SG}(\Lambda)\right) \leqslant \exp\left(O\left(C^2\right)\log^2(1/q)\right) \times \max\left(\gamma(\Lambda),\mu_\Lambda^{-1}(\mathcal{SG})\right)\sum_{x\in\Lambda^+}\mu_{\Lambda^+}\left(c_x^\omega\,\mathrm{Var}_x(f)\right)$$

and the same holds with $\mathcal{ST}$ instead of $\mathcal{T}$.

Proof. By a standard two-block result (see e.g. Lemma 4.2.10) and convexity of the variance, for $x\in\Lambda^+\backslash\Lambda$ we get

$$c_x^{\Lambda^+\backslash\Lambda,\mathbf{0}_\Lambda\cdot\omega_{\mathbb{Z}^2\backslash\Lambda^+}}\mu_{\Lambda\cup\{x\}}\left(\mathrm{Var}_x(f)|\,\mathcal{SG}(\Lambda)\right) \leqslant q^{O(1)}\mu_{\Lambda\cup\{x\}}\left(\mathrm{Var}_\Lambda(f|\mathcal{SG})+c_x^{\Lambda^+,\omega}\,\mathrm{Var}_x(f)\middle|\,\mathcal{SG}(\Lambda)\right),$$

since it suffices to infect a $O(1)$ neighbourhood of x in Λ (see e.g. Lemma 4.3.13 for an analogous reasoning). Plugging this in Eq. (6.55) and recalling Eq. (6.2), this gives exactly the desired result. $\square$

6.A.2 Auxiliary three-block chain

We next prove a non-product variant of the standard two-block technique for the purposes of the proof of the East-extension Proposition 6.2.3. Let $(\Omega_i,\pi_i)_{i=1}^3$ be finite positive probability spaces, (Ω,π) denote the associated product space and $\nu=\pi(\cdot|\mathcal{H})$ for some event $\mathcal{H}\subset\Omega$. For $\omega\in\Omega$ we write $\omega_i\in\Omega_i$ for its i^{th} coordinate. Consider an event $\mathcal{F}\subset\Omega_1\times\Omega_2$ and set

$$\mathcal{D}(f)=\nu\left(\mathrm{Var}_\nu(f|\omega_3)+\mathbb{1}_\mathcal{F}\,\mathrm{Var}_\nu(f|\omega_1,\omega_2)\right)$$

for any $f:\mathcal{H}\to\mathbb{R}$. Observe that $\mathcal{D}$ is the Dirichlet form of the continuous time Markov chain on $\mathcal{H}$ in which the couple (ω_1,ω_2) is resampled at rate one from $\nu(\cdot|\omega_3)$ and, if $(\omega_1,\omega_2)\in\mathcal{F}$, then ω_3 is resampled with rate one from $\nu(\cdot|\omega_1,\omega_2)$. This chain is reversible w.r.t. ν.

Lemma 6.A.3. *Assume that $\mathcal{F} \times \Omega_3 \subset \mathcal{H}$. Then, for all $f : \mathcal{H} \to \mathbb{R}$ we have*

$$\mathrm{Var}_\nu(f) \leqslant O(1) \max_{\omega_3 \in \Omega_3} \nu^{-1}(\mathcal{F}|\omega_3)\mathcal{D}(f).$$

Proof. We follow Proposition 5.3.5. Consider the Markov chain $(\omega(t))_{t\geqslant 0}$ described above. Given two arbitrary initial conditions $\omega(0)$ an $\omega'(0)$ we will construct a coupling of the two chains such that with probability $\Omega(1)$ we have $\omega(t) = \omega'(t)$ for $t > T := \max_{\omega_3\in\Omega_3} \nu^{-1}(\mathcal{F}|\omega_3)$. Standard arguments [256] then prove that the mixing time of the chain is $O(T)$ and the lemma follows.

To construct our coupling, we use the following representation of the Markov chain. We are given two independent Poisson clocks with rate one and the chain transitions occur only at the clock rings. When the first clock rings, a Bernoulli variable ξ with probability of success $\nu(\mathcal{F}|\omega_3)$ is sampled. If $\xi = 1$, then the couple (ω_1, ω_2) is resampled w.r.t. the measure $\pi(\cdot|\mathcal{F}) = \nu(\cdot|\mathcal{F}, \omega_3)$, while if $\xi = 0$, then (ω_1, ω_2) is resampled w.r.t. the measure $\nu(\cdot|\mathcal{F}^c, \omega_3)$. Clearly, in doing so the couple (ω_1, ω_2) is resampled w.r.t. $\nu(\cdot|\omega_3)$. If the second clock rings, we resample ω_3 from π_3 if $\omega \in \mathcal{F}$ and ignore the ring otherwise.

Both chains will use the same clocks. When the first clock rings and the current couple of configurations is (ω, ω'), we first maximally couple the two Bernoulli variables ξ, ξ' corresponding to ω, ω' respectively. Then:

- if $\xi = \xi' = 1$, we update both (ω_1, ω_2) and (ω_1', ω_2') to the *same* couple $(\eta_1, \eta_2) \in \mathcal{F}$ with probability $\pi((\eta_1, \eta_2)|\mathcal{F})$;
- otherwise, we resample (ω_1, ω_2) and (ω_1', ω_2') independently from their respective law given ξ, ξ'.

When the second clock rings, the two chins attempt to update to two maximally coupled couples of configurations with the corresponding distributions.

Suppose now that two consecutive rings occur at times $t_1 < t_2$ at the first and second clocks respectively and the Bernoulli variables at time t_1 are both 1. Then the two configurations are clearly identical at t_2. To conclude the proof, observe that for any time interval Δ of length one the probability that there exist $t_1 < t_2$ in Δ as above is at least $1/(4T)$. □

6.A.3 Proofs of the one-directional extensions

We will require a more technical version of Eq. (6.2) accounting for a boundary condition. Let $\gamma_I^\omega(\Lambda)$ be the smallest constant $\gamma \geqslant 1$ such that for all $f : \Omega \to \mathbb{R}$

$$\mu_{\Lambda_I^\omega}\left(\mathrm{Var}_\Lambda\left(f|\mathcal{SG}^\omega\right)\right) \leqslant \gamma \sum_{x\in\Lambda_I^\omega} \mu_{\Lambda_I^\omega}\left(c_x^\omega \mathrm{Var}_x(f)\right). \tag{6.57}$$

For the rest of the section we assume the setting of Section 6.2 and set $I = \{i - k + 1, \ldots, i + k - 1\}$.

Lemma 6.A.4. *Assume that we East-extend $\Lambda(\underline{r})$ by l in direction u_i. Then*

$$\gamma\left(\Lambda^M\right) \leqslant \max_{\omega} \gamma_I^{\omega}\left(\Lambda^1\right) \prod_{m=1}^{M-1} \frac{a_m}{q^{O(W)}}$$

where a_m is defined in Eq. (6.3).

Proof. We will loosely follow Eq. (5.13). Proceeding by induction it suffices to prove that for any $m \in [1, M)$ and $\omega \in \Omega_{\mathbb{Z}^2 \setminus \Lambda^{m+1}}$

$$\gamma_I^{\omega}\left(\Lambda^{m+1}\right) \leqslant \max_{\omega' \in \Omega_{\mathbb{Z}^2 \setminus \Lambda^m}} \gamma_I^{\omega'}\left(\Lambda^m\right) \frac{a_m}{q^{O(W)}}. \tag{6.58}$$

Let us fix m and ω as above. We partition $\Lambda^{m+1} = V_1 \sqcup V_2 \sqcup V_3$, so that

$$V_1 \cup V_2 = \Lambda^m, \qquad V_2 \cup V_3 = \Lambda^m + s_m u_i.$$

In order to apply Lemma 6.A.3, we define $\Omega_1 = \Omega_{V_1}$, $\Omega_2 = \Omega_{V_2}$, $\Omega_3 = \mathcal{T}^{\omega}(V_3)$ (note that V_3 is a translate of $T(\underline{r}, s_m, i)$) and equip them with $\pi_1 = \mu_{V_1}$, $\pi_2 = \mu_{V_2}$ and $\pi_3 = \mu_{V_3}(\cdot|\mathcal{T}^{\omega})$ respectively. We set $\mathcal{H} = \mathcal{SG}^{\omega}(\Lambda^{m+1})$ and $\mathcal{F} = \mathcal{SG}(\Lambda^m) \cap \mathcal{SG}(V_2)$ (which were defined by East-extending $\Lambda(\underline{r})$). Indeed, it holds that $\mathcal{F} \times \Omega_3 \subset \mathcal{H}$, so Lemma 6.A.3 gives

$$\begin{aligned}&\operatorname{Var}_{\Lambda^{m+1}}(f|\mathcal{SG}^{\omega}) \leqslant \max_{\omega'} O\left(\mu^{-1}\left(\mathcal{SG}(V_2)|\,\mathcal{SG}^{\omega'}(\Lambda^m)\right)\right)\\ &\times \mu_{\Lambda^{m+1}}\left(\operatorname{Var}_{\Lambda^m}(f|\mathcal{SG}^{\eta_{V_3}\cdot\omega}) + \mathbb{1}_{\mathcal{SG}(V_2)}\operatorname{Var}_{V_3}(f|\mathcal{T}^{\omega})\middle|\,\mathcal{SG}^{\omega}\left(\Lambda^{m+1}\right)\right).\end{aligned} \tag{6.59}$$

From Eq. (6.57) we have

$$\begin{aligned}&\mu_{(\Lambda^{(m+1)})_I^{\omega}}\left(\operatorname{Var}_{\Lambda^m}(f|\mathcal{SG}^{\eta_{V_3}\cdot\omega})\right)\\ &\qquad\leqslant \max_{\omega'} \gamma_I^{\omega'}\left(\Lambda^m\right) \sum_{x \in (\Lambda^{m+1})_I^{\omega}} \mu_{(\Lambda^{m+1})_I^{\omega}}\left(c_x^{\omega} \operatorname{Var}_x(f)\right).\end{aligned}$$

On the other hand, recalling Observation 6.1.12,

$$\begin{aligned}&\mu_{(\Lambda^{(m+1)})_I^{\omega}}\left(\mathbb{1}_{\mathcal{SG}(V_2)}\operatorname{Var}_{V_3}(f|\mathcal{T}^{\omega})\middle|\,\mathcal{SG}^{\omega}\left(\Lambda^{m+1}\right)\right)\\ &\quad\leqslant \frac{\mu(\mathcal{SG}(V_2) \cap \mathcal{T}^{\omega}(V_3))}{\mathcal{SG}^{\omega}(\Lambda^{m+1})} \mu_{(\Lambda^{m+1})_I^{\omega}}\left(\operatorname{Var}_{V_3}(f|\mathcal{T}^{\omega})\middle|\,\mathcal{SG}(V_2) \cap \mathcal{T}^{\omega}(V_3)\right)\\ &\quad\leqslant \frac{\mu_{(\Lambda^{m+1})_I^{\omega}}(\operatorname{Var}_{\Lambda^m + s_m u_i}(f|\mathcal{SG}(V_2) \cap \mathcal{T}^{\omega}(V_3)))}{\mu(\mathcal{T}(V_3))q^{O(W)}}\\ &\quad\leqslant \frac{\gamma_I^{\omega}(\Lambda^m)}{\mu(\mathcal{T}(V_3))q^{O(W)}} \sum_{x \in (\Lambda^m + s_m u_i)_I^{\omega}} \mu_{(\Lambda^{m+1})_I^{\omega}}\left(c_x^{(\Lambda^m + s_m u_i)_I^{\omega},\omega} \operatorname{Var}_x(f)\right),\end{aligned}$$

where we used Definition 6.2.2 in the second inequality and Eqs. (6.32) and (6.57) in the third one. Plugging these bounds into Eq. (6.59), we obtain

$$\gamma_I^\omega\left(\Lambda^{m+1}\right) \leqslant \frac{\max_{\omega'} \gamma_I^{\omega'}(\Lambda^m)}{q^{O(W)}\mu(\mathcal{T}(V_3)) \min_{\omega'} \mu(\mathcal{SG}(V_2)|\mathcal{SG}^{\omega'}(\Lambda^m))}.$$

Since by Definition 6.2.2 (and Section 6.1.6) we have

$$\mu\left(\mathcal{T}(V_3)\right)\mu\left(\mathcal{SG}(V_2)\middle|\,\mathcal{SG}^{\omega'}\left(\Lambda^m\right)\right) = q^{O(W)}/a_m$$

for all ω', the proof of Eq. (6.58) and the lemma are complete. □

Proof of Proposition 6.2.3. By Lemma 6.A.4 it suffices to relate $\gamma(\Lambda(\underline{r}))$ and $\max_\omega \gamma_I^\omega(\Lambda^1)$, using Corollary 6.A.2. Notice that by Definition 6.2.2 we have

$$\mathcal{SG}^\omega(\Lambda^1) = \mathcal{SG}(\Lambda(\underline{r})) \times \mathcal{T}^\omega(T(\underline{r}, \lambda_i, i)). \tag{6.60}$$

Therefore,

$$\begin{aligned}\operatorname{Var}_{\Lambda^1}\left(f|\mathcal{SG}^\omega\right) \leqslant \mu_{\Lambda(\underline{r})}\left(\operatorname{Var}_{T(\underline{r},\lambda_i,i)}\left(f|\mathcal{T}^\omega\right)\middle|\,\mathcal{SG}\right)&\\ + \mu_{T(\underline{r},\lambda_i,i)}\left(\operatorname{Var}_{\Lambda(\underline{r})}(f|\mathcal{SG})\middle|\,\mathcal{T}^\omega\right)&\end{aligned}$$

(this is a straightforward property of conditional variances w.r.t. a product measure, see e.g. Lemma 4.2.9 or Eq. (6.24)). The former term above is treated by Corollary 6.A.2, while the latter is dealt with by Eq. (6.2). □

We next turn to CBSEP-extensions, setting $J = [4k]\backslash\{i+k, i-k\}$.

Lemma 6.A.5. *Assume we CBSEP-extend $\Lambda(\underline{r})$ by l in direction u_i. Then*

$$\gamma\left(\Lambda(\underline{r} + l\underline{v}_i)\right) \leqslant \max_\omega \gamma_J^\omega\left(\Lambda(\underline{r} + \lambda_i\underline{v}_i)\right) \frac{\mu_{\Lambda(\underline{r}+\lambda_i\underline{v}_i)}(\mathcal{SG})}{\mu_{\Lambda(\underline{r}+l\underline{v}_i)}(\mathcal{SG})} e^{O(C^2)\log^2(1/q)}.$$

Proof. As in Eq. (5.13) (with minor amendments as in Lemma 6.A.4), we have

$$\gamma\left(\Lambda^M\right) \leqslant \max_\omega \gamma_J^\omega\left(\Lambda^1\right) \frac{\mu_{\Lambda^1}(\mathcal{SG})}{\mu_{\Lambda^M}(\mathcal{SG})} \prod_{m=1}^{M-1} \frac{b_m}{q^{O(W)}}$$

with

$$b_m = \max_\omega \mu_{\Lambda^{m+1}}^{-2}\left(\mathcal{SG}_{s_m}\middle|\,\mathcal{SG}^\omega\right) \max_\omega \mu_{\Lambda^{m+1}}^{-1}\left(\mathcal{SG}_0^\omega\middle|\,\mathcal{SG}^\omega\right),$$

so we are left with proving $b_m \leqslant q^{-O(C)}$ for all m. The last statement is simply Lemma 6.B.1—the analogue of Corollary 5.A.3. □

Proof of Proposition 6.2.5. By Lemma 6.A.5 it suffices to relate $\gamma_J^\omega(\Lambda^1)$ and $\gamma\left(\Lambda(\underline{r})\right)$. This is done exactly as in Lemma 5.2.12 (see particularly Eq. (5.19) there), replacing Eq. (5.23) by Corollary 6.A.2. □

6.B Conditional probabilities

Recall Definition 6.2.4. The next result generalises Corollary 5.A.3, which relied on explicit computations unavailable in our setting. This result is the reason for the somewhat artificial Definition 6.1.8 of helping sets and Definition 6.2.1 of $\mathcal{ST}$.

Lemma 6.B.1. *Let $\mathcal{U}$ have a finite number of stable directions. Fix $i \in [4k]$ and a symmetric droplet $\Lambda = \Lambda(\underline{r} + l\underline{v}_i)$ obtained by CBSEP-extension by l in direction u_i. Assume that $l \leqslant \ell^{\mathrm{m}+}$ is divisible by λ_i. Then for all $s \in [0, l]$ divisible by λ_i and $\omega, \omega' \in \Omega_{\mathbb{Z}^2\setminus\Lambda}$*

$$\mu\left(\mathcal{SG}_s^{\omega}(\Lambda) \middle| \mathcal{SG}^{\omega'}(\Lambda)\right) \geqslant q^{O(C)}.$$

Proof. We will prove that for all $s, s' \in [0, l]$ divisible by λ_i and $\omega, \omega' \in \Omega_{\mathbb{Z}^2\setminus\Lambda}$ we have

$$\frac{\mu_\Lambda(\mathcal{SG}_s^{\omega})}{\mu_\Lambda(\mathcal{SG}_{s'}^{\omega'})} = q^{O(W)}. \tag{6.61}$$

Once this is established, given that

$$\max_{s'} \mu_\Lambda\left(\mathcal{SG}_{s'}^{\omega'}\right) \leqslant \mu_\Lambda\left(\mathcal{SG}^{\omega'}\right) = \mu_\Lambda\left(\bigcup_{s'} \mathcal{SG}_{s'}^{\omega'}\right) \leqslant O(l) \max_{s'} \mu_\Lambda\left(\mathcal{SG}_{s'}^{\omega'}\right),$$

we immediately deduce the desired result. Moreover, it clearly suffices to establish Eq. (6.61) for $s' = 0$.

To prove Eq. (6.61), let us first observe that by the symmetry of Definitions 6.1.8 and 6.2.4,

$$\frac{\mu_\Lambda(\mathcal{SG}_s^{\omega})}{\mu_\Lambda(\mathcal{SG}_0^{\omega'})} = \frac{\mu_{T_s}(\mathcal{ST}^{\omega_s})\mu_{T_{l-s}}(\mathcal{ST}^{\omega_{l-s}})}{\mu_{T_l}(\mathcal{ST}^{\omega_l})}, \tag{6.62}$$

where $T_x = T(\underline{r}, x, i)$ and the ω_x are certain boundary conditions that can be expressed in terms of ω, ω'. Further note that for any x and ω''

$$\mu_{T_x}\left(\mathcal{ST}^{\omega''}\right) = q^{O(W)}\mu_{T_x}(\mathcal{ST}), \tag{6.63}$$

by the Harris inequality, since it suffices to add W-helping sets on the last $O(1)$ lines of the tube. Finally, observing that $T_s \cup (T_{l-s} + su_i) = T_l$, we get

$$\mu_{T_s}(\mathcal{ST})\mu_{T_{l-s}}(\mathcal{ST}) \leqslant \mu_{T_l}(\mathcal{ST}) \leqslant \mu_{T_s}(\mathcal{ST}^{\mathbf{0}})\mu_{T_{l-s}}(\mathcal{ST}). \tag{6.64}$$

Putting Eqs. (6.62) to (6.64) together, we obtain Eq. (6.61) as desired. □

We next treat certain perturbations of traversability events, building them progressively from segments and parallelograms in the next lemmas.

Lemma 6.B.2. *Fix $i \in [4k]$ such that $\alpha(u_i) \leqslant \alpha$. Let S be a segment perpendicular to u_i and $S', S'' \subset S$ be segments partitioning S. Assume that $|S| \geqslant W|S''|$ and $|S| = q^{-\alpha+o(1)}$. Then*

$$\mu\left(\mathcal{H}(S')|\mathcal{H}(S)\right) \geqslant 1 - \frac{W^{1/3}|S''|}{|S|} - q^{1-o(1)}.$$

Proof. Let us note that a stronger version of this result can be proved more easily by counting circular shifts of the configuration in a $O(1)$ neighbourhood of S such that a given helping set remains at distance at least some constant from S''. We prefer to give the proof below as a preparation for Lemma 6.B.3.

For concreteness, let us assume that $\alpha(u_{i+2k}) > \alpha$, other cases being treated similarly. Thus, helping sets are just u_i-helping sets or W-helping sets. Recall from Definition 6.1.7 that a u_i-helping set is composed of Q translates of the set Z_i. Further let $S \subset \overline{\mathbb{H}}_{u_i} \backslash \mathbb{H}_{u_i}$. For $r \in [Q]$ we denote by $\mathcal{H}_r(S)$ the event that S has a translate of Z_i by a vector of the form $(r + k_r Q)\lambda_{i+k} u_{i+k}$ with $k_r \in \mathbb{Z}$ and similarly define $\mathcal{H}_r(S')$. In words, we look for the part of the helping set with a specified reminder r modulo Q.

Since $|S| = q^{-\alpha+o(1)}$, the probability that there are $\alpha + 1$ infected sites at distance $O(1)$ from each other and from S is $q^{1-o(1)}$. Furthermore, if this does not happen, but $\mathcal{H}(S)$ occurs, then all $\mathcal{H}_r(S)$ for $r \in [Q]$ occur disjointly. Therefore, by the BK inequality Proposition 6.1.3,

$$\mu(\mathcal{H}(S)) \leqslant q^{1-o(1)} + \prod_{r\in[Q]} \mu(\mathcal{H}_r(S)) \leqslant \left(1 + q^{1-o(1)}\right) \prod_{r\in[Q]} \mu(\mathcal{H}_r(S)), \quad (6.65)$$

since, as in Observation 6.1.10, we have

$$\mu(\mathcal{H}_r(S)) \geqslant 1 - (1 - q^{\alpha})^{\Omega(|S|)} \geqslant q^{o(1)}. \quad (6.66)$$

Using Eq. (6.65) and applying the Harris inequality, we get

$$\begin{aligned}\frac{\mu(\mathcal{H}(S'))}{\mu(\mathcal{H}(S))} &\geqslant \left(1 - q^{1-o(1)}\right) \prod_{r\in[Q]} \frac{\mu(\mathcal{H}_r(S'))}{\mu(\mathcal{H}_r(S))} \\ &\geqslant \left(1 - q^{1-o(1)}\right) \left(\frac{|S'| - O(1)}{|S|}\right)^{Q},\end{aligned}$$

where in the last inequality we used that $\mathcal{H}_r(S)$ and $\mathcal{H}_r(S')$ can be expressed in terms of the i.i.d. (and therefore exchangeable) Bernoulli variables corresponding to each translate of the helping set being infected. Recalling that $|S| \geqslant W|S''|$, this concludes the proof. □

Lemma 6.B.3. *Let $i, j \in [4k]$ be such that $\alpha(u_i) \leqslant \alpha$ and $j \notin \{i, i + 2k\}$. Consider the parallelogram*

$$R = R(l, h) = \overline{\mathbb{H}}_{u_i}(l) \cap \overline{\mathbb{H}}_{u_j}(h) \cap \overline{\mathbb{H}}_{u_{j+2k}}(0) \cap \overline{\mathbb{H}}_{u_{i+2k}}(0)$$

for $l \in [\rho_i, e^{q^{-o(1)}}]$ *and* $h = q^{-\alpha+o(1)}$. *We say that* R *is* traversable *in direction* u_i *(*$\mathcal{T}(R)$ *occurs), if for each nonempty segment of the form*

$$S = R \cap \overline{\mathbb{H}}_{u_i}(h') \backslash \mathbb{H}_{u_i}(h')$$

the event $\mathcal{H}_{C^2}^{\mathbf{1}_{\mathbb{Z}^2 \backslash R(l+W,h)}}(S)$ *occurs. Let* $R' = R(l, h')$ *with* $1 > h'/h \geqslant 1 - 1/W$. *Then*

$$\mu\left(\mathcal{T}(R')|\mathcal{T}(R)\right) \geqslant \left(1 - W^{1/2}\left(1 - \frac{h'}{h}\right) - q^{1-o(1)}\right)^{O(l)}$$

Proof. Let us write simply $\mathcal{H}_m$ for $\mathcal{H}_{C^2}^{\mathbf{1}_{\mathbb{Z}^2 \backslash R}}(R \cap \overline{\mathbb{H}}(m\rho_i) \backslash \mathbb{H}_{u_i}(m\rho_i))$ and similarly define $\mathcal{H}'_m$ for R'. Let m take its values in $[M]$ for some integer M. Separate R into its lower and upper halves R_1 and R_2, consisting of $\lfloor M/2 \rfloor$ and $\lceil M/2 \rceil$ segments perpendicular to u_i respectively. If $\mathcal{T}(R)$ occurs, then one of the following occurs.

- There is a set of $\alpha + 1$ infections at distance $O(1)$ from each other and from both R_1 and R_2, and the rectangles, formed by removing in each of R_1 and R_2 the $O(1)$ lines closest to their common boundary, are both traversable.
- The rectangles R_1 and R_2 are disjointly traversable.

Using the BK inequality Proposition 6.1.3, this gives

$$\begin{aligned} \mu(\mathcal{T}(R)) &\leqslant q^{1-o(1)}\mu^2\left(\mathcal{T}(R(l/2 - O(1), h))\right) + \mu\left(\mathcal{T}(R_1)\right)\mu(\mathcal{T}(R_2)) \\ &= \left(1 + q^{1-o(1)}\right)\mu(\mathcal{T}(R_1))\mu(\mathcal{T}(R_2)), \end{aligned}$$

the last estimate following as in Eq. (6.66) from the fact that traversing the $O(1)$ lines at the boundary of R_1 and R_2 happens with probability $q^{o(1)}$. Iterating the same reasoning, we obtain

$$\mathcal{T}(R) \leqslant \left(1 + q^{1-o(1)}\right) \prod_{m \in [M]} \mu(\mathcal{H}_m),$$

since $l = e^{q^{-o(1)}}$. Hence, by the Harris inequality

$$\frac{\mu(\mathcal{T}(R'))}{\mu(\mathcal{T}(R))} \geqslant \left(1 - q^{1-o(1)}\right) \prod_{m \in [M]} \frac{\mu(\mathcal{H}'_m)}{\mu(\mathcal{H}_m)}.$$

The last fraction can be bounded, using Lemma 6.B.2, to obtain

$$\mu\left(\mathcal{T}(R')|\mathcal{T}(R)\right) \geqslant \left(1 - W^{1/3}\left(1 - \frac{h'}{h}\right) - q^{1-o(1)}\right)^{M}. \qquad \square$$

As a result, we are able to prove the following result, which vastly generalises Lemma 5.A.5.

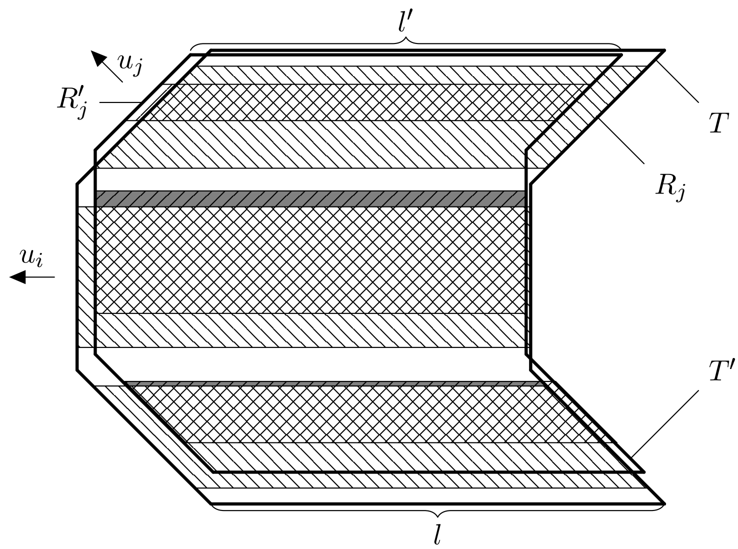

Figure 6.9 – Illustration of the perturbation of Corollary 6.B.4. The two thickened tubes are T and T'. The parallelograms R_j are North-West hatched, while R'_j are North-East hatched. Thus, R''_j are double hatched. The shaded rectangles are R_j^2, while R_j^1 are the remainder of the area which is North-East but not double hatched.

Corollary 6.B.4 (Perturbation cost). *Let $T = T(\underline{r}, l, i)$ be a tube with $i \in [4k]$ such that $\alpha(u_j) \leqslant \alpha$ for all $j \in (i-k, i+k)$. Denote the side lengths of $\Lambda(\underline{r})$ by $\underline{s}$ as usual. Assume that $l \in [\Omega(1), e^{q^{-o(1)}}]$, $s := \min_{i-k<j<i+k} s_j = q^{-\alpha+o(1)}$ and $\max_{i-k<j<i+k} s_j = q^{-\alpha+o(1)}$. For some $\Delta \in (0, s/\sqrt{W}]$, let $\underline{r}'$ and l' be such that $0 \leqslant s_j - s'_j \leqslant O(\Delta)$ for all $j \in [i-k, i+k]$ and $0 \leqslant l - l' \leqslant O(\Delta)$. Further let $x \in \mathbb{R}^2$ be such that $\|x\| = O(\Delta)$ and $d, d' \in [0, O(\Delta)]$ with $d \leqslant d'$. Denoting $T' = T(\underline{r}', l', i) + x$, for any boundary conditions $\omega, \in \Omega_{\mathbb{Z}^2 \setminus T}$ and $\omega' \in \Omega_{\mathbb{Z}^2 \setminus T'}$, we have*

$$\mu\left(\mathcal{T}_{d'}^{\omega'}(T')\middle|\, \mathcal{T}_d^{\omega}(T)\right) \geqslant q^{-O(W)}\left(1-(1-q^{\alpha})^{\Omega(s_j)}\right)^{O(\Delta)} \times \left(1 - W\Delta/s - q^{1-o(1)}\right)^{O(l)}$$

and the same holds for $\mathcal{ST}$ instead of $\mathcal{T}$.

Proof. Recalling Definition 6.2.1, it is clear that $\mathcal{T}_d^{\omega}(T)$ is the intersection of $2k-1$ independent traversability events for parallelograms of length l in the sense of Lemma 6.B.3. Let us denote them by $(R_j)_{j=i-k+1}^{i+k-1}$ and, similarly, $(R'_j)_{j=i-k+1}^{i+k-1}$ for T' with R_j and R'_j having sides perpendicular to u_j (see Fig. 6.9). Finally, set $R''_j = R_j \cap R'_j = R(l - O(\Delta), s_j - O(\Delta + C^2))$ for

$j \in (i-k, i+k)$ and observe that $R'_j \backslash R''_j$ consists of two disjoint parallelograms $R^1_j = R(O(\Delta), s_j - O(\Delta + C^2))$ and $R^2_j = R(l - O(\Delta), O(\Delta))$ with the notation of Lemma 6.B.3 (up to translation).

Observe that that $\mathcal{T}^{\omega'}_{d'}(T')$ is implied by the presence of W-helping sets on the last $O(1)$ lines of each R'_j and R^1_j and the traversability of all R^1_j and R''_j. Then by the Harris inequality and the independence of $\mathcal{T}(R_j)$ we have that

$$\mu\left(\mathcal{T}^{\omega'}_{d'}(T')\middle|\mathcal{T}^{\omega}_{d}(T)\right) \geqslant q^{O(W)} \prod_j \mu\left(\mathcal{T}\left(R^1_j\right)\right) \mu\left(\mathcal{T}\left(R''_j\right)\middle|\mathcal{T}\left(R_j\right)\right).$$

We may then conclude, using Lemma 6.B.3 and that by Observation 6.1.10

$$\mu\left(\mathcal{T}\left(R^1_j\right)\right) \geqslant \left(1 - (1-q^{\alpha})^{\Omega(s_j)}\right)^{O(\Delta)}. \qquad \square$$

Part II

Combinatorics

Chapter 7

Universality for critical KCM: infinite number of stable directions

This chapter is based on joint work with Laure Marêché and Cristina Toninelli [214], establishing the following result, proving the lower bound of Theorem 1.6.4 for class (b) and Corollary 1.6.5 for families with infinite number of stable directions (recall Section 1.6).

Theorem 7.0.1. *Let $\mathcal{U}$ be a critical update family with an infinite number of stable directions and difficulty α. Then*

$$\mathbb{E}_\mu(\tau_0) \geqslant e^{\Omega(1)/q^{2\alpha}}$$

and the same holds for T_{rel}.

7.1 Sketch of the proof

In this section we outline roughly the strategy to derive our main result, Theorem 7.0.1. The hypothesis of infinite number of stable directions provides us with an interval of stable directions. We can then construct stable 'droplets' of shape as in Figure 7.3 (see Definitions 7.3.5 and 7.3.6), where we recall from Section 1.2.1 that a set is stable if it coincides with its closure. Thus, if all infections are initially inside a droplet, this will be true at any time under the KCM dynamics. The relevance and advantage of such shapes come from the fact that only infections situated to the left of a droplet can induce growth left. This is manifestly not feasible without the hypothesis of having an interval of stable directions. It is worth noting that these shapes, which may seem strange at first sight, are actually very natural and intrinsically present in the dynamics. Indeed, such is the shape of the stable sets for a representative model of this class—the modified 2-neighbour model with one (any) rule removed, that is the three-rule update family with rules $\{(-1,0),(0,1)\}$,$\{(-1,0),(0,-1)\}$,$\{(0,-1),(1,0)\}$ (it can also be seen as the modified Duarte model with an additional rule). The stable sets in this case are actually Young diagrams.

We construct a collection of such droplets covering the initial configuration of infections, so that it gives an upper bound on the closure. To do this, we devise an improvement of the α-covering algorithm of Bollobás, Duminil-Copin, Morris and Smith [70]. It is important for us not to overestimate the closure as brutally. Indeed, a key step and the main difficulty of the present chapter is the Closure Proposition 7.3.20, which roughly states that the collections of droplets associated to the closure of the initial infections is equal to the collection for the initial infections. This is highly non-trivial, as in order not to overshoot in defining the droplets, one is forced to ignore small patches of infections (larger than the ones in [70]), which can possibly grow significantly when we take the closure for the bootstrap percolation process and especially so if they are close to a large infected droplet. In order to remedy this problem, we introduce a relatively intrinsic notion of 'crumb' (see Definition 7.3.1) such that its closure remains one and does not differ too much from it. A further advantage of our algorithm for creating the droplets over the one of [70] is that it is somewhat canonical, with a well-defined unique output, which has particularly nice 'algebraic' description and properties (see Remark 7.3.10). Another notable difficulty we face is systematically working in roughly a half-plane (see Remark 7.3.21 for generalisations) with a fully infected boundary condition, but we manage to extend our reasoning to this setting very coherently.

Finally, having established the Closure Proposition 7.3.20 alongside standard and straightforward results like an Aizenman–Lebowitz Lemma 7.3.13 and an exponential decay of the probability of occurrence of large droplets (Lemma 7.3.15), we finish the proof via the following approach, inspired by the one developed by Marêché, Martinelli and Toninelli [267] for the Duarte model. The key step here (see Section 7.4) is mapping the KCM legal paths to those of an East dynamics via a suitable renormalisation. Roughly speaking, we say that a renormalised site is infected if it contains a large droplet of infections. However, for the renormalised configuration to be mostly invariant under the original KCM dynamics, we rather look for the droplets in the closure of the original set of infections instead. This is where the Closure Proposition 7.3.20 is used to compensate the fact that the closure of equilibrium is not equilibrium. In turn, this mapping together with the combinatorial result for the East model recalled in Section 1.3.2 (Proposition 1.3.7), yield a bottleneck for our dynamics corresponding to the creation of $\log(1/q_{\mathrm{eff}})$ droplets, where $1/q_{\mathrm{eff}}$ is the equilibrium distance between two infected sites in the renormalized lattice, and $q_{\mathrm{eff}} \sim e^{-1/q^{\alpha}}$. This provides for the time scales the desired lower bound $q_{\mathrm{eff}}^{\log(q_{\mathrm{eff}})} \sim e^{1/q^{2\alpha}}$ of Theorem 7.0.1.

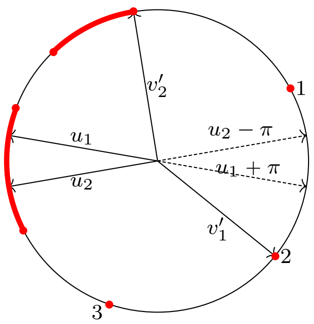

Figure 7.1 – Illustration of Lemma 7.2.1 and its proof. Thickened arcs represent intervals of strongly stable directions. Solid dots represent isolated and semi-isolated stable directions. The difficulties of the isolated stable directions are indicated next to them and yield that the difficulty of the model is $\alpha = 2$. The directions chosen in Lemma 7.2.1 are the solid vectors u_1, u_2, $v_1 = v_1'$ and a direction v_2 in the strongly stable interval ending at v_2' sufficiently close to v_2'. Note that the definition of v_2' (and v_1') disregards stable directions with difficulty smaller than α as present on the figure.

The last part of the proof follows very closely the ideas put forward in [267] for the Duarte model. However, in [267], there was no need to develop a subtle droplet algorithm since, owing to the oriented character of the Duarte constraint, droplets could simply be identified with some large infected vertical segments. It is also worth noting that, thanks to the less rigid notion of droplets that we develop in the general setting, some of the difficulties faced in [267] for Duarte are no longer present here.

7.2 Preliminaries and notation

Let us fix a critical update family $\mathcal{U}$ with an infinite number of stable directions for the rest of the chapter. We will omit $\mathcal{U}$ from all notation, such as $\alpha(\mathcal{U})$.

Directions The next lemma establishes that one can make a suitable choice of 4 stable directions, which we will use for all our droplets. At this point the statement should look very odd and technical, but it simply reflects the fact that we have a lot of freedom for the choice and we make one which will simplify a few of the more technical points in later stages. Nevertheless, this is to a large extent not needed besides for concision and clarity.

A direction $u \in S^1$ is called *rational* if $\tan u \in \mathbb{Q} \cup \{\infty\}$.

Lemma 7.2.1. *There exist rational stable directions $\mathcal{S} = \{u_1, u_2, v_1, v_2\}$ (see Figure 7.1) with difficulty at least α such that*

- *The directions appear in counter-clockwise order u_1, u_2, v_1, v_2.*
- *No $u \in \mathcal{S}$ is a semi-isolated stable direction.*

- *u_{3-i} belongs to the cone spanned by v_i and u_i for $i \in \{1,2\}$ i.e. the strictly smaller interval among $[v_i, u_i]$ and $[u_i, v_i]$ contains u_{3-i}.*
- *0 is contained in the interior of the convex envelope of $\mathcal{S}$.*
- *Either $u_2 < v_1 - \pi/2$ or $u_1 > v_2 + \pi/2$.*
- *$(\mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}) \cap \mathbb{Z}^2$ is stable or, equivalently, $\nexists U \in \mathcal{U}, U \subset \mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}$.*
- *the directions*

$$\begin{aligned} u' &= (u_1 + u_2)/2, \\ u_1' &= (3u_1 + u_2)/4, \\ u_2' &= (u_1 + 3u_2)/4 \end{aligned}$$

  are rational.

Proof. Since $\mathcal{U}$ has an infinite number of stable directions and they form a finite union of closed intervals with rational endpoints [74, Theorem 1.10], there exists a non-empty open interval I''' of stable directions. Further note that the set J of directions u such that there exists a rule $U \in \mathcal{U}$ and $x \in U$ with $\langle x, u \rangle = 0$ is finite, so one can find a non-trivial closed subinterval $I'' \subset I'''$ which does not intersect J. The directions u_1 and u_2 will be chosen in I'', which clearly implies that they are strongly stable and thus with infinite difficulty. Moreover, if there exists $U \in \mathcal{U}$ with $U \subset \mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}$, by stability of u_2, we have $U \cap (\mathbb{H}_{u_1} \backslash \mathbb{H}_{u_2}) \neq \varnothing$, which contradicts $I'' \cap J = \varnothing$.

Since $\mathcal{U}$ is critical it does not have two opposite strongly stable directions, so there is no strongly stable direction in $I'' + \pi$. If there are any (isolated or semi-isolated) stable directions in $I'' + \pi$, we can further choose a non-trivial open subinterval $I' \subset I''$, for which this is not the case (there is a finite number of isolated and semi-isolated stable directions). Let $\pi > \delta > 0$ be such that the angle between any two consecutive directions of difficulty at least α is at most $\pi - \delta$ (it is well defined by Definition 1.6.1). We then choose a non-trivial closed subinterval $I' \supset I = [u_1, u_2]$ with u_1 rational and $u_1' = (3u_1 + u_2)/4$ rational and with $0 < u_2 - u_1 < \delta < \pi$. It easily follows from the sum and difference formulas for the tangent function that u', u_2' and u_2 are also rational.

Let

$$\begin{aligned} v_1' &= \max\{v \in (u_2, u_1 + \pi) \colon \alpha(v) \geqslant \alpha\}, \\ v_2' &= \min\{v \in (u_2 - \pi, u_1) \colon \alpha(v) \geqslant \alpha\}. \end{aligned}$$

These both exist, since $I + \pi$ does not contain stable directions, both $(u_2, u_2 + \pi)$ and $(u_1 - \pi, u_1)$ contain directions with difficulty at least α by Definition 1.6.1 and the set of such directions is closed. If v_1' is not semi-isolated, we set $v_1 = v_1'$ and similarly for v_2. Otherwise, we choose a rational strongly stable direction sufficiently close to v_1' as v_1 and similarly for v_2. We claim that this choice satisfies all the desired conditions. Indeed, all directions in $\mathcal{S}$ are stable non-semi-isolated rational with difficulty at least α and the last but one condition was already verified.

One does have that u_1 is in the cone spanned by v_2 and u_2, which is implied by $v_2 \in (u_2-\pi, u_1)$ and similarly for u_2, so the third condition is also verified. If $v_2'-v_1' \geqslant \pi$, then there is an open half circle contained in (v_1', v_2') with no direction of difficulty at least α, which contradicts Definition 1.6.1, so $v_2 - v_1 < \pi$ and the same holds for $u_1 - v_2$, $u_2 - u_1$ and $v_1 - u_2$ by the definition of v_1' and v_2', the fact that v_1 and v_2 are sufficiently close to them and the fact that I was chosen smaller than π. Thus 0 is in the convex envelope of $\mathcal{S}$.

Finally, if one has both $v_1-u_2 \leqslant \pi/2$ and $u_1-v_2 \leqslant \pi/2$, then one obtains $v_2'-v_1' > \pi-\delta$, since I is smaller than δ. However, v_1' and v_2' are consecutive directions of difficulty at least α, which contradicts the definition of δ. □

Notation For the rest of the chapter we fix directions $\mathcal{S} = \{u_1, u_2, v_1, v_2\}$ as in Lemma 7.2.1 and assume without loss of generality that $u_2 < v_1 - \pi/2$.

Let us fix large constants

$$1 \ll C_1 \ll C_2' \ll C_2 \ll C_3 \ll C_4' \ll C_4 \ll C_5,$$

each of which can depend on previous ones as well as on $\mathcal{U}$ and $\mathcal{S}$. We will also use asymptotic notation whose constants can depend on $\mathcal{U}$ and $\mathcal{S}$, but not on C_1 or the other constants above. All asymptotic notation is with respect to $q \to 0$, so we assume throughout that $q > 0$ is sufficiently small.

For any two sets $K, \partial \subset \mathbb{R}^2$ we define $[K]_\partial = [(K \cup \partial) \cap \mathbb{Z}^2]\backslash\partial$.

Finally, we make the convention that throughout the chapter all distances, balls and diameters are Euclidean unless otherwise stated. We say that a set $X \subset \mathbb{R}^2$ is *within distance* δ of a set $Y \subset \mathbb{R}^2$ if $d(x, Y) \leqslant \delta$ for all $x \in X$ where d is the Euclidean distance.

7.3 Droplet algorithm

In this section we define our main tool—the droplet algorithm. It can be seen as a significant improvement on the α-covering and u-iceberg algorithms [70, Definitions 6.6 and 6.22], many of whose techniques we adapt to our setting.

We will work in an infinite domain Λ defined as follows (see Figure 7.2). Fix some vector $a_0 \in \mathbb{R}^2$ and let

$$\begin{aligned} \partial =& \mathbb{H}_{u'} \cup \mathbb{H}_{u_1'}(a_0) \cup \mathbb{H}_{u_2'}(a_0), \\ \Lambda =& \mathbb{R}^2\backslash\partial, \end{aligned} \tag{7.1}$$

where the directions u', u_1' and u_2' are those defined in Lemma 7.2.1. In other words, Λ is a cone with sides perpendicular to u_1' and u_2' cut along a

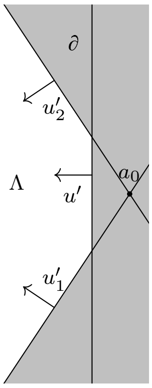

Figure 7.2 – The open domain ∂ defined in (7.1) is shaded, while its complement Λ is not. The lines are the boundaries of the three half-planes defining ∂. Note that if $a_0 \notin \mathbb{H}_{u'}$, then Λ becomes simply a cone.

line perpendicular to u'. The reader is invited to simply think that ∂ is a half-plane directed by u', which will not change the reasoning.

7.3.1 Clusters and crumbs

Let Γ be the graph with vertex set $\mathbb{Z}^2$ but with $x \sim y$ if and only if $\|x-y\| \leqslant C_2$. Let Γ' be defined similarly with C_2 replaced by C_2'. Given a finite $K \subset \Lambda \cap \mathbb{Z}^2$, we say that $\kappa \subset K$ is a *connected component* of K in Γ if the subgraph of Γ induced by the vertex set κ is connected and there do not exist vertices $x \in K\backslash\kappa$ and $y \in \kappa$ such that $x \sim y$ in Γ.

Crumbs For a given finite set $K \subset \Lambda \cap \mathbb{Z}^2$ of infections we would like to have a notion of a connected component being 'big' or 'small.' 'Small' components will be dubbed 'crumbs' and will play a negligible perturbative role in the bootstrap percolation process, by inducing only 'very localised' growth and being 'well isolated' from the rest of the infections. A sufficient condition for this, as identified in [70], is that $|\kappa| < \alpha$. However, contrary to what was the case in [70], we need the notion of 'crumb' to be stable under the closure (with respect to the bootstrap percolation process), i.e. the closure of a 'crumb' to still be a 'crumb.' We thus identify as 'crumb' any component, which is the closure of a set of size less than α. Also taking into account the boundary, this leads us to the following notion.

Definition 7.3.1 (Crumb)**.** Fix a finite set $K \subset \Lambda \cap \mathbb{Z}^2$ and let κ be a connected component of K in Γ. We say that κ is a *crumb* for K if the following conditions hold.

- For all $x \in \kappa$ we have $d(x, \partial) > C_2$.
- There exists a set $P_\kappa \subset \mathbb{Z}^2$ such that $[P_\kappa] \supset \kappa$ and $|P_\kappa| = \alpha - 1$.

First properties of crumbs It follows from the definition that a crumb κ for K is at distance more than C_2 from $\partial \cup (K\backslash\kappa)$. Moreover, the closure of a crumb is within bounded distance from the crumb, as we shall see in Corollary 7.3.17 (see Figure 7.5a). Also, crumbs have diameters much smaller than C_3, as we shall see in Corollary 7.3.17. The proofs of this corollary and Observation 7.3.16, which it follows from, are both independent of the rest of the argument and are only postponed for convenience. Nevertheless, we allow ourselves to use these (easy) results ahead of their proofs.

These properties justify and quantify the idea that crumbs are 'small,' that they only grow 'locally,' and it is clear that (if we disregard the boundary) the closure of a crumb is a crumb.

Modified crumbs Unfortunately, if K is the union of two crumbs at distance slightly larger than C_2, it is not necessarily true that $[K]$ is still composed of crumbs (recall that, albeit locally, crumbs can grow under the bootstrap percolation process), which can be disastrous. This is the reason for introducing 'modified crumbs' with $C_2' \ll C_2$, so that in the scenario above all connected components of $[K]$ in Γ' are 'modified crumbs' (there may now be more than two of them).

Definition 7.3.2 (Modified crumb)**.** We define a *modified crumb* by replacing in Definition 7.3.1 Γ by Γ' and C_2 by C_2'.

In the sequel we will encounter more 'modified' notions and constants (like C_2'). These will be applied to K equal to the closure $[K']_\partial$ of some K', which is our initial set of infections. Our ultimate goal is to ensure that simply using these modified notions based on (much smaller) modified constants will compensate the closure operation.

Clusters We next consider connected components which are not crumbs. Since they can be very large (particularly so if we are working with the closure of a set), we cut them up into (possibly overlapping) pieces termed 'clusters,' which have bounded size. Roughly speaking, a 'cluster' is any 'big, but not too big' connected set of infections.

Definition 7.3.3 (Cluster)**.** Fix a finite set $K \subset \Lambda \cap \mathbb{Z}^2$. Let κ be a connected component of K in Γ which is not a crumb. We say that a subset C of κ is a *cluster* for K if the following conditions hold.

- $\operatorname{diam}(C) \leqslant C_3$.
- C is connected in Γ (i.e. C is a connected component of C in Γ).
- Either $C = \kappa$ or for all $x \in \kappa \backslash C$ and $y \in C$ such that $x \sim y$ in Γ we have $\operatorname{diam}(C \cup \{x\}) > C_3$.

A cluster is called *boundary cluster* if it is at distance at most C_2 from ∂. For a cluster C we denote by $Q(C)$ the smallest open quadrilateral with sides perpendicular to $\mathcal{S}$ containing the set $\{x \in \mathbb{R}^2 : d(x, C) < C_4\}$.

We similarly define *modified cluster* and *modified boundary cluster* by replacing Γ by Γ' and C_2 by C_2'. For a cluster or modified cluster C we denote by $Q'(C)$ the smallest open quadrilateral with sides perpendicular to $\mathcal{S}$ containing the set $\{x \in \mathbb{R}^2 : d(x, C) < C_4'\}$.

Identifying clusters and crumbs In order to identify the clusters and crumbs of K, one may proceed as follows. Determine the connected components of K in Γ and consider each of them separately. For a given component κ first check if it is at distance at most C_2 from ∂. If so, then it is not a crumb and will give rise to clusters. If not, then check if κ is the closure of at most $\alpha - 1$ sites. If this second verification succeeds, then κ is determined to be a crumb and, as mentioned above, it must have diameter much smaller than C_3.

If κ is thus determined not to be a crumb, we proceed to identify its clusters. If $\mathrm{diam}(\kappa) \leqslant C_3$, then there is a single cluster—κ—and we are done. If not, we construct the clusters of κ by the following algorithm. Initialise the set $C = \varnothing$. If there exists $y \in \kappa \backslash C$ such that $C \cup \{y\}$ is connected in Γ and has diameter at most C_3, then replace C by $C \cup \{y\}$ and repeat. If several such y exist, then we do this for each possible y in parallel. The clusters containing x are all possible sets C obtained via this algorithm to which no y can be added.

In particular, this provides us with a partition of K into well separated crumbs, single clusters equal to their corresponding connected component and sets of overlapping clusters whose union is a connected component of diameter larger than C_3.

First properties of clusters Following the algorithm above, we obtain some basic properties of clusters.

Observation 7.3.4. Let C be a non-boundary cluster or non-boundary modified cluster for a finite $K \subset \Lambda \cap \mathbb{Z}^2$. Then $|C| \geqslant \alpha$.

Proof. Let κ be the connected component of K in Γ containing C. If $\mathrm{diam}(\kappa) \leqslant C_3$, then $C = \kappa$ and κ would be a crumb if we had $|\kappa| \leqslant \alpha - 1$, by taking $P_\kappa \supset \kappa$. If, on the contrary, $\mathrm{diam}(\kappa) > C_3$, then $\mathrm{diam}(C) \geqslant C_3 - C_2$ (by the third condition of Definition 7.3.3) and we can choose C_3 large enough to have $\frac{C_3 - C_2}{C_2} \geqslant \alpha$. □

Finally, for every cluster C we have $\mathrm{diam}(C) \leqslant C_3$, so C intersects at most $2^{5C_3^2}$ other clusters. Also, $Q(C) \supset [C]$, since $Q(C) \cap \mathbb{Z}^2 \supset C$ is stable. Furthermore, $\mathrm{diam}(Q(C)) = \Theta(C_4)$, as $\mathrm{diam}(C) \leqslant C_3$. Analogous statements hold for modified clusters.

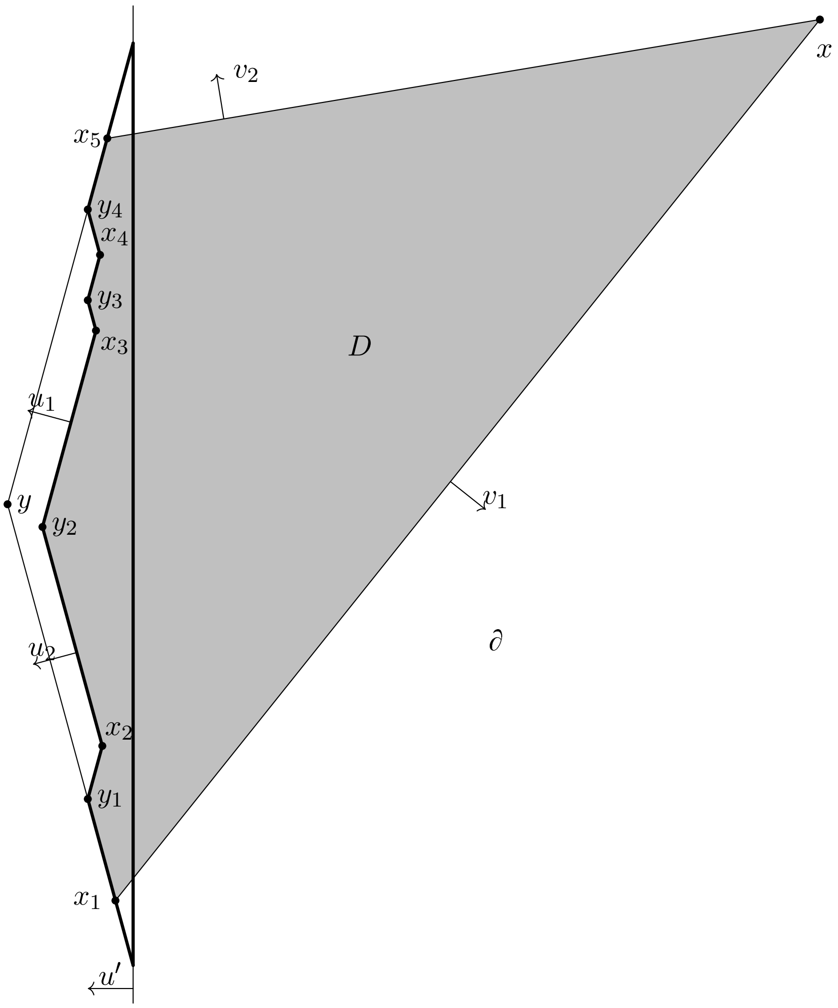

Figure 7.3 – The shaded region D is a distorted Young diagram (DYD) as in Definition 7.3.5. The larger quadrilateral with vertices x, x_1, y and x_5 is $Q(D)$. Note that $Q(D)$ can degenerate into a triangle, but we call it a quadrilateral nevertheless. On the figure $|D|$ is the length of the v_1 side, but this is not always the case. The thickened region is the cut distorted Young diagram (CDYD) $C(D)$ of D. The vertical line is the boundary between Λ on its left and ∂ on its right.

7.3.2 Distorted Young diagrams

We now define the shape that our 'droplets' will have, which resembles Young diagrams[1]. The following definitions are illustrated in Figure 7.3.

Definition 7.3.5 (DYD)**.** A *distorted Young diagram* (DYD) is a subset of $\mathbb{R}^2$ of the form

$$(\mathbb{H}_{v_1}(x) \cap \mathbb{H}_{v_2}(x)) \cap \bigcap_{i\in I}(\mathbb{H}_{u_1}(x_i) \cup \mathbb{H}_{u_2}(x_i)) \tag{7.2}$$

for a finite set I, some set $X = \{x_i : i \in I\}$ of vectors $x_i \in \mathbb{R}^2$ and $x \in \mathbb{R}^2$. The vectors x_i and x are uniquely defined up to redundancy (and up to the convention that all x_i are on the topological boundary of the DYD). Alternatively, a DYD can also be defined by

$$(\mathbb{H}_{v_1}(x) \cap \mathbb{H}_{v_2}(x)) \cap \bigcup_{i\in I}(\mathbb{H}_{u_1}(y_i) \cap \mathbb{H}_{u_2}(y_i)), \tag{7.3}$$

where y_i are the convex corners of the diagram rather than the concave ones.

For any DYD D we denote by y the vector such that

$$\langle y, u_j\rangle = \sup_{a\in D}\langle a, u_j\rangle = \max_{i\in I}\langle y_i, u_j\rangle$$

for $j \in \{1, 2\}$. We further denote

$$Q(D) = \mathbb{H}_{u_1}(y) \cap \mathbb{H}_{u_2}(y) \cap \mathbb{H}_{v_1}(x) \cap \mathbb{H}_{v_2}(x),$$

i.e. the minimal open quadrilateral containing D with sides directed by $\mathcal{S}$. In these terms, for any cluster or modified cluster C we have that $Q(C)$ and $Q'(C)$ are DYD, $Q(Q(C)) = Q(C)$ and $Q(Q'(C)) = Q'(C)$.

Definition 7.3.6 (CDYD)**.** A *cut distorted Young diagram* (CDYD) is a subset of $\mathbb{R}^2$ of the form

$$\Lambda \cap (\mathbb{H}_{u_1}(y) \cap \mathbb{H}_{u_2}(y)) \cap \bigcap_{i\in I}(\mathbb{H}_{u_1}(x_i) \cup \mathbb{H}_{u_2}(x_i))$$

for a finite set I and some vectors $x_i \in \mathbb{R}^2$ and $y \in \Lambda$. Alternatively, one can write

$$\Lambda \cap \bigcup_{i\in I}(\mathbb{H}_{u_1}(y_i) \cap \mathbb{H}_{u_2}(y_i)),$$

where $y_i \in \Lambda$ are the convex corners.

[1] For the 3-rule model alluded to in Section 7.1 stable sets consist precisely of Young diagrams and the directions $\mathcal{S}$ provided by Lemma 7.2.1 can be arbitrarily close to the four axis directions, yielding Young diagrams.

For a DYD, D, we denote by $C(D)$ the CDYD defined by the same x_i and y or the same y_i. We extend the notation $C(D)$ to CDYD by setting $C(D) = D$ if D is a CDYD. Note that by Lemma 7.2.1 all DYD and CDYD are stable for the bootstrap percolation dynamics (restricted to Λ). Also pay attention to the fact that CDYD are not necessarily connected, contrary to DYD.

Definition 7.3.7 (Size)**.** For a DYD D we set $\pi(D) = \{x \in \mathbb{R} : \exists\, y \in D, \langle y, v_1 + \pi/2\rangle = x\}$ to be its *projection* (parallel to v_1) and $|D| = \sup \pi(D) - \inf \pi(D)$ to be its *size*—the length of the projection. For a CDYD D we denote its *size* $\mathrm{diam}(D)/C_1$ by $|D|$.

Note that if D is a DYD, then $|D| = |Q(D)|$ by Lemma 7.2.1 and the assumption we made that $u_2 < v_1 - \pi/2$. Furthermore, for all DYD $\mathrm{diam}(D) = \Theta(|D|)$ again by Lemma 7.2.1 with constants depending only on $\mathcal{S}$. One should be careful with the meaning of size for disconnected CDYD, but it will not cause problems, as all CDYD arising in our forthcoming algorithm are connected.

Observation 7.3.8. Note that for any $d \geqslant 1$ the number of discretised DYD and CDYD (i.e. intersections of a DYD or CDYD with $\mathbb{Z}^2$) containing a fixed point $a \in \mathbb{R}^2$ of diameter at most d is less than c^d for some constant c depending only on $\mathcal{S}$.

Proof. Note that a DYD or CDYD is uniquely determined by its rugged edge formed by its u_1 and u_2-sides. However, this edge injectively defines an oriented percolation path with directions perpendicular to u_1 and u_2 on the lattice

$$\{x \in \mathbb{R}^2 : \exists x_1, x_2 \in \mathbb{Z}^2, \langle x, u_1\rangle = \langle x_1, u_1\rangle, \langle x, u_2\rangle = \langle x_2, u_2\rangle\}$$

(except its endpoints, which lie on similar lattices). Since the graph-length of this path is bounded by $O(d)$ and its endpoints are within distance d from a, the result follows. ☐

7.3.3 Span

We next introduce a procedure of merging DYD and CDYD. This will be used only for couples of intersecting ones, but can be defined regardless of whether they intersect. The operation is illustrated in Figure 7.4.

Lemma 7.3.9. *For any two DYD, D_1 and D_2, the minimal DYD containing $D_1 \cup D_2$ is well defined. We denote it by $D_1 \vee D_2$ and call it their* span. *The operation $\vee$ is associative*[2] *and commutative.*

[2] Associativity was referred to as commutativity by previous authors [74].

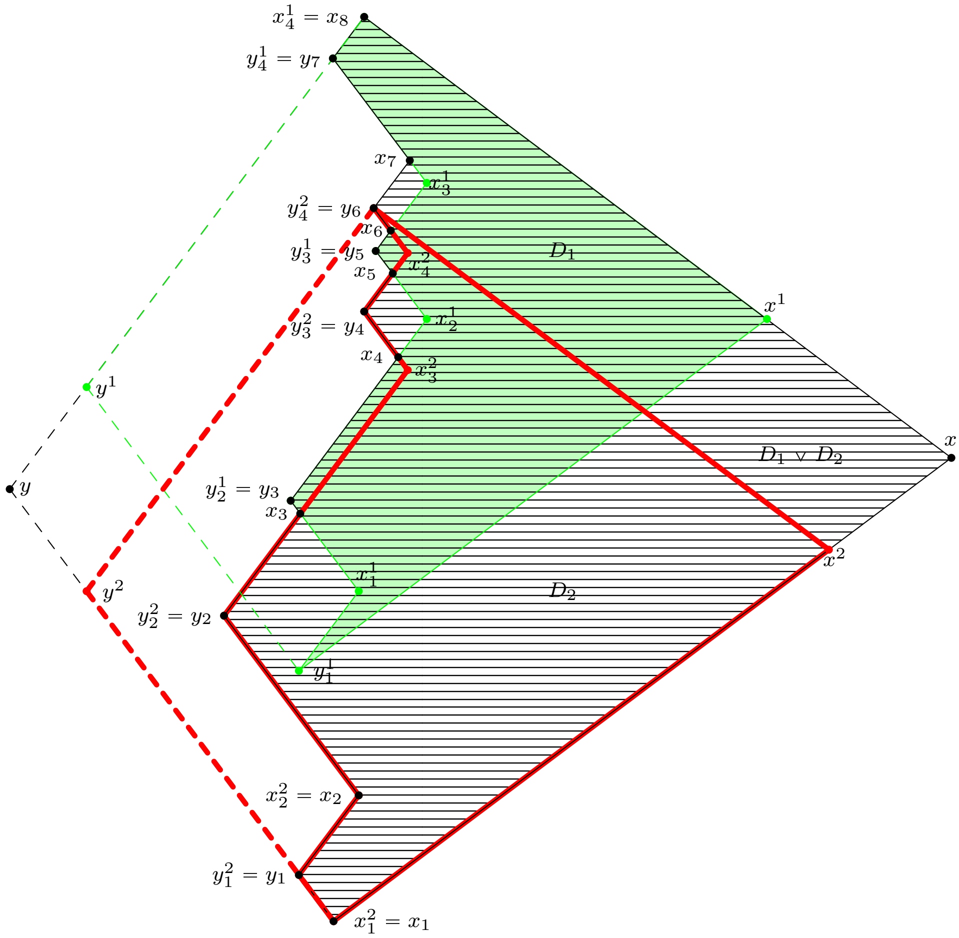

Figure 7.4 – The shaded region D_1 and thickened region D_2 are DYD. Their respective quadrilaterals $Q(D_i)$ are completed by dashed lines. Their span $D_1 \vee D_2$ is hatched and its quadrilateral $Q(D_1 \vee D_2)$ is also completed by dashed lines.

Proof. Let D_1 be defined by $Y^1 = \{y_i^1 : i \in I\}, x^1$ (see (7.3)) and similarly for D_2. Let $x \in \mathbb{R}^2$ be the vector such that $\mathbb{H}_{v_i}(x^1) \cup \mathbb{H}_{v_i}(x^2) = \mathbb{H}_{v_i}(x)$ for $i \in \{1, 2\}$. Let Y be the set of $y_i \in Y^1 \cup Y^2$ such that for all $y_j \in Y^1 \cup Y^2$ with $y_i \neq y_j$ we have $\mathbb{H}_{u_1}(y_j) \cap \mathbb{H}_{u_2}(y_j) \not\supset \mathbb{H}_{u_1}(y_i) \cap \mathbb{H}_{u_2}(y_i)$. We denote by D the DYD defined by Y, x and claim that for any DYD $D' \supset D_1 \cup D_2$ we have $D' \supset D$, which is enough to conclude that $D = D_1 \vee D_2$ is well defined. Let D' be defined by Y', x'.

Note that for each $y_i \in Y$ (and in fact in $Y_1 \cup Y_2$) there is a sequence of points in D_1 or D_2 converging to y_i, so that (by extraction of a subsequence) there exists y'_j with $\mathbb{H}_{u_1}(y'_j) \cap \mathbb{H}_{u_2}(y'_j) \supset \mathbb{H}_{u_1}(y_i) \cap \mathbb{H}_{u_2}(y_i)$. Similarly, there is a sequence of points in D_1 or D_2 converging to the boundary of $\mathbb{H}_{v_1}(x)$, so that $\mathbb{H}_{v_1}(x') \supset \mathbb{H}_{v_1}(x)$ and similarly for v_2. Thus, we do have $D' \supset D$.

Finally, the commutativity is obvious and the associativity follows from the characterisation of $D_1 \vee D_2$ as the minimal DYD containing both D_1 and D_2. □

We analogously define the *span* $D_1 \vee D_2$ of two CDYD D_1 and D_2—the minimal CDYD containing both—and note that it coincides with their union (which is also commutative and associative). We also define the *span* $C \vee D$ of a DYD D and a CDYD C as the minimal CDYD containing $(C \cup D)\backslash\partial$, which coincides with $C \vee C(D)$. The proof that it is well defined is analogous to Lemma 7.3.9.

We have thus defined an associative and commutative binary operation $\vee$ on all DYD and CDYD. Moreover, the idempotent unary operation $C(\cdot)$ is distributive with respect to $\vee$ and $C(D_1) \vee D_2 = C(D_1 \vee D_2)$. Furthermore, the span of several DYD is the minimal DYD containing all of them, while the span of several DYD and at least one CDYD is the minimal CDYD containing all the corresponding CDYD.

7.3.4 Droplet algorithm and spanned droplets

A *droplet* is any DYD contained in Λ or CDYD. We are now ready to define our *droplet algorithm*, which takes as input a finite set $K \subset \Lambda \cap \mathbb{Z}^2$ of infections and outputs a set $\mathcal{D}$ of disjoint connected droplets. It proceeds as follows.

- Form an initial collection of DYD $\mathcal{D}$ consisting of $Q(C)$ for all clusters C of K. If a DYD $D \in \mathcal{D}$ intersects ∂, replace it by its CDYD, $C(D)$, to obtain a droplet.

- As long as it is possible, replace two intersecting droplets of $\mathcal{D}$ by their span. If the span intersects ∂, replace it by its CDYD to obtain a droplet.

- Output the collection $\mathcal{D}$ obtained when all droplets are disjoint.

We similarly define the *modified droplet algorithm* by replacing $Q(C)$ by $Q'(C)$ and clusters by modified clusters above.

The output $\mathcal{D}$ is clearly a collection of disjoint connected droplets. Indeed, by induction all x_i corners of droplets remain in Λ (see Figure 7.4), so that DYD remain connected when replaced by CDYD.

Remark 7.3.10. From the results of Section 7.3.3 it is clear that the order of merging does not impact the output of the algorithm, which is thus well defined. It can also be expressed as the minimal collection of disjoint droplets containing the intersection with Λ of the original collection of quadrilaterals. This minimal collection is well defined. Consequently, the union of the output is increasing in the input.

Definition 7.3.11 (Spanned droplets)**.** Let D be a droplet and $K \subset \mathbb{Z}^2$. We say that D is *spanned* for K with boundary ∂ if the output of the droplet algorithm for $K \cap D$ has a droplet containing D. We omit K and ∂ if they are clear from the context. Similarly, D is *modified spanned* if the output of the modified droplet algorithm for $K \cap D$ has a droplet containing D.

Note that, when seen as an event, a droplet being spanned is monotone. It is also clear that each droplet appearing in (the intermediate or final stages of) the droplet algorithm is spanned and similarly for the modified droplet algorithm. Indeed, the clusters responsible for creating a droplet in the course of the algorithm are contained in the droplet, so each of them is still a cluster of $K \cap D$ (recall that crumbs have diameter much smaller than C_3).

7.3.5 Properties of the algorithm

We next establish several properties of the algorithm. The approach is similar to the one of [70] with the notable exception of the key Closure Proposition 7.3.20. We start with the following purely geometric statement.

Lemma 7.3.12 (Subadditivity)**.** *Let D_1 and D_2 be two DYD or CDYD with non-empty intersection. Then*

$$|D_1 \vee D_2| \leqslant |D_1| + |D_2|.$$

Furthermore, if D is a DYD intersecting ∂, then $|C(D)| \leqslant |D|$.

Proof. First assume that D_1 and D_2 are DYD. Since $|D| = |Q(D)|$ for any DYD D and $D_1 \vee D_2 \subset Q(Q(D_1) \vee Q(D_2))$, it suffices to prove the assertion for merging quadrilaterals instead of DYD. But in that case it is not hard to check directly and is a particular case of Lemma 15 of the first arXiv version of [74] (or Lemma 23 of the second version). Since similar (but actually slightly more involved) details were omitted in the proof of the corresponding

Lemma 4.6 of [74] and differed to earlier versions, we will not go into useless detail here either. To give a sketch of a possible argument, one can check that for fixed shapes of $Q(D_1)$ and $Q(D_2)$ the maximal $Q(Q(D_1) \vee Q(D_2))$ is achieved when their intersection is reduced to a vertex. Yet, in those configurations one can obtain the v_1 and v_2 sides of $Q(Q(D_1) \vee Q(D_2))$ as the union of those of $Q(D_1)$ and translates of those of $Q(D_2)$ (see Figure 7.4). This concludes the proof, as only v_1 and (possibly) v_2 sides contribute to $|\cdot|$ by Lemma 7.2.1.

Next assume that D_1 is a DYD and D_2 is a CDYD. Let $Y = \{y_i : i \in I\}$ be the set of vectors defining $C(D_1)$ and let $a \in D_1 \cap D_2$. Since $Y \subset \overline{D_1}$, we have that $d(y_i, a) \leqslant \mathrm{diam}(D_1)$. It then easily follows that the CDYD defined by only one corner, y_i, which we denote $C(y_i)$, is within distance $O(\mathrm{diam}(D_1))$ from $C(a)$. But then $C(D_1) = \bigcup_{i \in I} C(y_i)$ is within distance $O(\mathrm{diam}(D_1))$ from $C(a)$. Thus, $|D_1 \vee D_2| \leqslant (\mathrm{diam}(D_2) + O(\mathrm{diam}(D_1)))/C_1 \leqslant |D_2| + |D_1|$, since $\mathrm{diam}(D_1) = O(|D_1|)$ and all implicit constants depend only on $\mathcal{S}$ and are thus much smaller than C_1.

Next assume that D_1 and D_2 are CDYD. Then the statement is trivial, because $D_1 \vee D_2 = D_1 \cup D_2$, so $\mathrm{diam}(D_1) + \mathrm{diam}(D_2) \geqslant \mathrm{diam}(D_1 \vee D_2)$ by the triangle inequality.

Finally, let D be a DYD intersecting ∂. Then, $|C(Q(D))| \geqslant |C(D)|$ and $|Q(D)| = |D|$, so we may assume that $D = Q(D)$ and prove $|C(D)| \leqslant |D|$. But in this case it is easy to see that $\mathrm{diam}(C(D)) = O(\mathrm{diam}(D)) = O(|D|)$ with constants depending only on $\mathcal{S}$, which concludes the proof. □

The subadditivity lemma will be used to prove the next two adaptations of classical results.

Lemma 7.3.13 (Aizenman–Lebowitz)**.** *Let K be a finite set and let D be a spanned droplet with $|D| \geqslant C_4^2$. Then for all $C_4^2/C_1 \leqslant k \leqslant |D|/C_1$ there exists a connected spanned droplet D' with $k \leqslant |D'| \leqslant 2k$. The same statement holds for modified spanned droplets.*

Proof. By Lemma 7.3.12 at each step of the droplet algorithm the largest size of a droplet appearing in the collection at most doubles. Initially the largest size is at most C_1C_4 and in the end there is a (unique) droplet $D'' \supset D$, so that $|D''| \geqslant |D|/C_1 \geqslant C_4^2/C_1 > C_1C_4$. Then there is a stage of the algorithm at which the maximal size of a droplet in $\mathcal{D}$ is between k and $2k$, which is enough since all droplets appearing in the droplet algorithm are connected and spanned. The proof for modified spanned droplets is identical, using the modified droplet algorithm. □

Lemma 7.3.14 (Extremal)**.** *Let $K \subset \mathbb{Z}^2$ and let D be a droplet spanned for K. Then the total number of disjoint clusters for $K \cap D$ in D is at least* $\mathrm{diam}(D)/C_4^2$.

Proof. In this proof all clusters will be clusters for $K \cap D$. Assume that at the initial stage of the algorithm there are k clusters (not disjoint). One can then find k/C_4' disjoint ones, since their diameter is at most C_3. Furthermore, by Lemma 7.3.12 the total size of droplets in the collection $\mathcal{D}$ is decreasing, so that $|D|/C_1 \leqslant |D'| \leqslant kC_1C_4$, where $D' \supset D$ is some droplet in the output of the algorithm. Indeed, $|Q(C)| \leqslant C_1C_4$ for all clusters C. This concludes the proof, since $|D| \geqslant \mathrm{diam}(D)/C_1$ for all DYD and CDYD. □

We next transform this extremal bound into an exponential decay of the probability that a droplet is spanned until saturation at the critical size. In the following lemma, we identify the configuration ω having law μ and the set of its zeroes.

Lemma 7.3.15 (Exponential decay)**.** *Let D be a droplet such that $|D| \leqslant 2/(C_5q^\alpha)$. Then*

$$\mu(D \textit{ is spanned for } \omega) < \exp(-C_4|D|).$$

Proof. Let D be a droplet with $|D| \leqslant 2/(C_5q^\alpha)$, so that $\mathrm{diam}(D) = d \leqslant 2C_1/(C_5q^\alpha)$. By Lemma 7.3.14 if D is spanned for ω, it contains at least d/C_4^2 disjoint clusters for $\omega \cap D$, each one having diameter at most C_3. Each non-boundary cluster has at least α sites by Observation 7.3.4, while boundary clusters are non-empty and located at distance at most C_2 from ∂. Thus, we have the union bound

$$\begin{aligned}
&\mu(D \text{ is spanned for } \omega)\\
&\leqslant \sum_{l=0}^{d/C_4^2} \binom{C_3^{2\alpha}d^2}{l}\binom{C_3d}{d/C_4^2-l} q^{l\alpha+(d/C_4^2-l)}\\
&\leqslant \sum_{l=d/(2C_4^2)}^{d/C_4^2} (C_4'q^\alpha d^2/l)^l.e^d + \sum_{l'=d/(2C_4^2)}^{d/C_4^2} (C_4'qd/l')^{l'}.e^d\\
&\leqslant \sum_{l=d/(2C_4^2)}^{d/C_4^2} \left(\frac{C_4'e^{2C_4^2}q^\alpha}{1/(2C_4^2)}\cdot\frac{2C_1}{C_5q^\alpha}\right)^l + \sum_{l'=d/(2C_4^2)}^{d/C_4^2}\left(2C_4^2C_4'e^{2C_4^2}q\right)^{l'}\\
&\leqslant \exp(-C_4d),
\end{aligned}$$

recalling that C_5 is sufficiently large depending on C_4, C_4' and C_1. □

Our next aim is to prove that the closure of a set is contained in its droplet collection up to very local infections next to initial ones. To that end we will need some preliminary results, similar to those used by Bollobás, Duminil-Copin, Morris and Smith [70].

Observation 7.3.16 (Lemma 6.5 of [70])**.** Let u be a rational non-semi-isolated stable direction. Let $K \subset \mathbb{Z}^2$ with $|K| < \alpha(u)$ (if $\alpha(u) = \infty$ the condition is that K is finite, but there is no a priori bound on its size). Then there exists a constant $C(\mathcal{U}, u, |K|)$ not depending on K such that $[K]_{\mathbb{H}_u}$ is within distance $C(\mathcal{U}, u, |K|)$ from K.

Since we will require some improvements later, we spell out a proof of the above result for completeness (actually our proof is slightly different from the one in [70]).

Proof of Observation 7.3.16. We prove the statement by induction on $|K|$. For a $K = \{x\}$ this is easy, since if $\langle x, u\rangle$ is sufficiently large $[K]_{\mathbb{H}_u} = K$ and otherwise there is a single possible configuration for each value of $\langle x, u\rangle$ up to translation. Assume the result holds for $|K| < n$. If one can write $K = K_1 \sqcup K_2$ with $K_1, K_2 \neq \varnothing$ and $d(K_1, K_2) > 2C(\mathcal{U}, u, n-1) + O(1)$, then $[K]_{\mathbb{H}_u} = [K_1]_{\mathbb{H}_u} \sqcup [K_2]_{\mathbb{H}_u}$, since $[K_1]_{\mathbb{H}_u}$ and $[K_2]_{\mathbb{H}_u}$ are at sufficiently large distance, hence no site can use both to become infected. Assume that, on the contrary, there are no large gaps between parts of K. There is a finite number of such K up to translation and for each of these $[K]$ is finite (e.g. since K is contained in a quadrilateral with sides perpendicular to $\mathcal{S}$), so within uniformly bounded distance from K. Therefore, if $\mathbb{H}_u$ is sufficiently far from K, $[K]_{\mathbb{H}_u} = [K]$. Otherwise, there is a finite number of possible K up to translation perpendicular to u and for each of them $[K]_{\mathbb{H}_u}$ is finite, so that one can indeed find a finite uniform constant $C(\mathcal{U}, u, n)$ as claimed. □

A quantitative version of this result was proved in Chapter 9. An easy corollary of Observation 7.3.16 is the fact that crumbs can only grow very locally (see Figure 7.5a).

Corollary 7.3.17. *Let C_1 be sufficiently large depending on $\mathcal{U}$. Let $K \subset \mathbb{Z}^2$ with $|K| < \alpha$. Then $[K]$ is within distance $C_1/(6\alpha)$ from K. Also, for a (modified) crumb κ we have that* $\operatorname{diam}([\kappa]) \leqslant \alpha C_2$ *and $[\kappa]$ is within distance C_1 from κ.*

Proof. The first assertion follows from Observation 7.3.16, since if it were wrong, one could simply translate a set K sufficiently far from a half-plane yielding a contradiction with the observation.

Next consider a (modified) crumb κ and P_κ minimal with $|P_\kappa| < \alpha$ and $[P_\kappa] \supset \kappa$. Then $[\kappa] \subset [P_\kappa]$ is within distance $C_1/(6\alpha)$ from P_κ. If the sites of P_κ are not connected in the graph Γ'' on $\mathbb{Z}^2$ with connections at distance at most $C_1 + C_2$, then either κ is not connected in Γ or P_κ is not minimal, which are both contradictions. Similarly, if there is no site of κ at distance smaller than $C_1/(2\alpha)$ from a $C_1/(2\alpha)$-connected component of P_κ, that component can be removed from P_κ, contradicting minimality. Hence, P_κ is within distance $C_1/2$ from κ. The result is then immediate, as

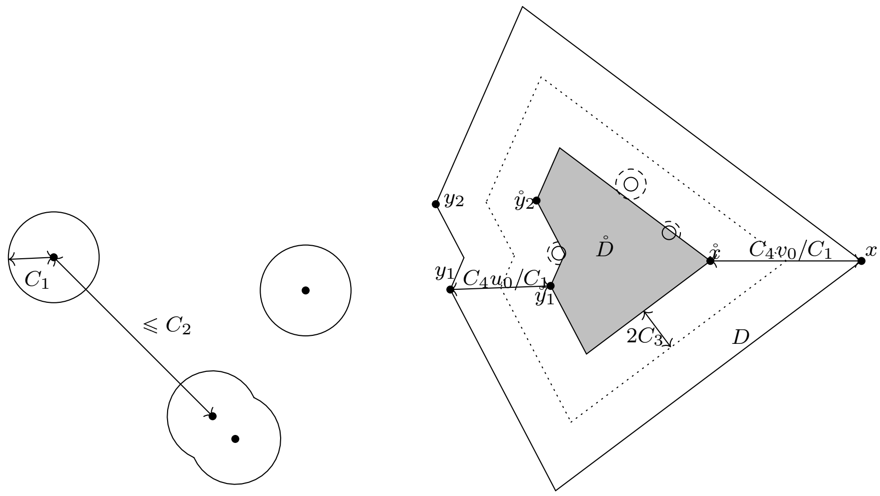

(a) The dots represent the sites of a crumb. The (disconnected) circled shape bounds its closure. Note that crumbs may have gaps of size C_2 while the growth allowed is only $C_1 \ll C_2$.

(b) The shaded region is the shrunken DYD $\mathring{D}$ of the largest DYD D. The solid circles represent crumbs and the dashed arcs are the bound for their growth provided by Lemma 7.3.19. The modified clusters of the closure are included in the dotted DYD.

Figure 7.5 – Illustrations of Corollary 7.3.17, Lemma 7.3.19 and Proposition 7.3.20.

$[\kappa]$ is within distance $C_1/2 + C_1/(6\alpha)$ from κ and its diameter is at most $C_1/(3\alpha) + \mathrm{diam}(P_\kappa)$, while $\mathrm{diam}(P_\kappa) \leqslant (\alpha - 1)(C_1 + C_2)$. □

In order to treat infection at the concave corners of droplets we will need the following modification of Observation 7.3.16.

Corollary 7.3.18. *Let u_1 and u_2 be rational strongly stable directions such that $\mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}$ is stable for the bootstrap percolation dynamics i.e. $\nexists U \in \mathcal{U}, U \subset \mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}$. Let $K \subset \mathbb{Z}^2$ with $|K| \leqslant \alpha - 1$. Then $[K]_{\mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}}$ is within distance $C(\mathcal{U}, u_1, u_2)$ from K.*

Proof. We apply a similar induction to the one in the proof of Observation 7.3.16. The only difference is that we can no longer use translation invariance. If $d(K, \mathbb{H}_{u_2}) > C(\mathcal{U}, u_1, |K|) + O(1)$, by Observation 7.3.16, we have $[K]_{\mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}} = [K]_{\mathbb{H}_{u_1}}$ and similarly for u_1 and u_2 interchanged. We can thus assume that K is within distance $C'(\mathcal{U}, u_1, u_2)$ from the origin. But then $[K \cup \mathbb{H}_{u_1} \cup \mathbb{H}_{u_2}] \subset \mathbb{H}_{u_1} \cup \mathbb{H}_{u_2} \cup \mathbb{H}_{u'}(C''(\mathcal{U}, u_1, u_2)u')$, where $u' = (u_1 + u_2)/2$, since the latter region is stable by the hypothesis on u_1, u_2. □

We next transform these results for infinite regions into a result for droplets. It states that a crumb next to a droplet cannot grow significantly (see

Figure 7.5b).

Lemma 7.3.19. *Let C_1 be sufficiently large depending on $\mathcal{U}$ and $\mathcal{S}$. Let D be a DYD at distance at least C_3 from ∂ or be a CDYD and let κ be a crumb. Then $[\kappa]_{D\cup\partial} = [\kappa]_D$ is within distance C_1 of κ.*

Proof. Assume that D is a DYD at distance at least C_3 from ∂. The proof of [70, Lemma 6.10] applies using (7.2), Observation 7.3.16, Corollary 7.3.18 and the arguments in the proof of Corollary 7.3.17 to give the result for $[\kappa]_D$, which is therefore at distance at least $C_2 - C_1$ from ∂ since $d(\kappa, \partial) \geqslant C_2$, so that in fact $[\kappa]_D = [\kappa]_{D\cup\partial}$.

Assume next that D is a CDYD. Then actually $D \cup \partial$ can be viewed as a DYD on the entire plane without boundary specified by an infinite number of vectors x_i, so that we are in the previous case. In order to avoid introducing the corresponding notion of infinite DYD, one can consider an increasing exhaustive sequence of DYD D_i converging to $D \cup \partial$ in the product topology and apply the previous result for $[\kappa]_{D_i}$, which will thereby apply to $D \cup \partial$. Finally, $[\kappa]_D = [\kappa]_{D\cup\partial}$ follows, since $d([\kappa]_{D\cup\partial}, \partial) \geqslant C_2 - C_1$. □

The next proposition is key to making the output of the algorithm essentially invariant under the KCM dynamics without having to pay for the fact that the closure for the bootstrap percolation dynamics of infections at equilibrium is not at all at equilibrium itself. The proof is illustrated in Figure 7.5b.

Proposition 7.3.20 (Closure)**.** *Let K be a finite set and $\mathcal{D}'$ be the collection of droplets given by the modified droplet algorithm with input $[K]_\partial$. Let $\mathcal{D}$ be the output of the droplet algorithm for K. Then*

$$\forall D' \in \mathcal{D}'\, \exists D \in \mathcal{D}, D' \subset D.$$

Proof. Let $\mathcal{K}$ be the set of crumbs for K. Set $\kappa_0 = \bigcup_{\kappa\in\mathcal{K}} \kappa$.

Claim 1. For each crumb $\kappa \in \mathcal{K}$ its closure $[\kappa] = [\kappa]_\partial$ consists of at most $\alpha - 1$ modified crumbs of $[\kappa]$ all contained within distance C_1 from κ.

Proof of Claim 1. There exists a set P_κ as in Definition 7.3.1, such that $[P_\kappa] \supset \kappa$ and thus $[P_\kappa] \supset [\kappa]$, which proves that all connected components of $[\kappa]$ for Γ' are modified crumbs. The fact that $[\kappa]$ is within distance C_1 of κ (and thus at distance at least C_2' from ∂) was proved in Corollary 7.3.17, which also shows that $[\kappa] = [\kappa]_\partial$, since κ is at distance more than C_2 from ∂. □

We can thus define $\mathcal{K}'(\kappa)$ to be the set of modified crumbs of $[\kappa]_\partial$, so that their union is disjoint and equal to $[\kappa]_\partial$. Moreover, crumbs in $\mathcal{K}$ are at distance at least C_2 from each other, so for any two of them $\kappa_1 \neq \kappa_2$ we have that any $\kappa_1' \in \mathcal{K}'(\kappa_1)$ and $\kappa_2' \in \mathcal{K}'(\kappa_2)$ are at distance at least $C_2 - 2C_1 \gg C_2'$ and also at such distance from ∂, so that $[\kappa_0]_\partial = \bigcup_{\kappa\in\mathcal{K}}[\kappa]_\partial$ has no modified cluster and consists of modified crumbs at distance at most C_1 from κ_0.

For a droplet $D \in \mathcal{D}$ consider the set of vectors Y and x (x is absent for CDYD) defining it. Then define $\mathring{Y} = Y + C_4u_0/C_1$ and $\mathring{x} = x + C_4v_0/C_1$, where $u_0 \in \mathbb{R}^2$ is the vector such that $\langle u_0, u_1\rangle = \langle u_0, u_2\rangle = -1$ and v_0 is defined identically in terms of v_1 and v_2. We denote by $\mathring{D}$ the droplet defined by $\mathring{Y}$ and $\mathring{x}$ and call it a *shrunken droplet*. Let $D_0 = \bigcup_{D\in\mathcal{D}} D$ and $\mathring{D}_0 = \bigcup_{D\in\mathcal{D}} \mathring{D}$. It is clear that $\mathring{D}$ is at distance at least C_4/C_1 from $\Lambda\backslash D$ for all droplets D. In particular, all shrunken droplets are at distance at least C_4/C_1 from each other and shrunken DYD are at distance at least C_4/C_1 from ∂, so that Lemma 7.3.19 applies to them and $[\mathring{D}_0]_\partial = \mathring{D}_0$.

Claim 2. $\mathring{D}_0 \cup \kappa_0 \supset K$.

Proof of Claim 2. Note that it is enough to prove that the clusters of K are contained in $\mathring{D}_0$. Assume that there exists $a \in K\backslash\mathring{D}_0$ and $a \in C$ for some cluster. Then, $Q(C) \cap \Lambda$ is contained in some $D \in \mathcal{D}$, which is defined by Y and x (x is absent for CDYD). Then since $a \notin \mathring{D}$, either for all $\mathring{y}_i \in \mathring{Y}$ we have $a \notin \mathbb{H}_{u_1}(\mathring{y}_i) \cap \mathbb{H}_{u_2}(\mathring{y}_i)$ or $a \notin \mathbb{H}_{v_1}(\mathring{x}) \cap \mathbb{H}_{v_2}(\mathring{x})$. In the former case, $a - C_4u_0/C_1 \notin \mathbb{H}_{u_1}(y_i) \cap \mathbb{H}_{u_2}(y_i)$ for all $y_i \in Y$. However, $Q(C)$ contains the ball of radius C_4 centered at a and $\|u_0\| = O(1)$, so we get a contradiction. If $a \notin \mathbb{H}_{v_1}(\mathring{x}) \cap \mathbb{H}_{v_2}(\mathring{x})$, the first point on the segment from a to $a - C_4v_0/C_1$ that is not in D is in Λ and in $Q(C)$, hence a contradiction. □

Claim 3. The set $[K]_\partial\backslash[\kappa_0]_\partial$ is within distance C_3 of $\mathring{D}_0$.

Proof of Claim 3. By Claim 2 we have $K_0 = \mathring{D}_0 \cup \kappa_0 \supset K$. It then clearly suffices to prove that $[K_0]_\partial\backslash[\kappa_0]_\partial$ is within distance C_3 of $\mathring{D}_0$.

Consider a crumb $\kappa \in \mathcal{K}$ at distance at most C_2 from $\mathring{D}_0$, so at distance at most C_2 from a shrunken droplet $\mathring{D}$ and necessarily at distance at least $C_4/C_1 - C_2 - C_3$ from any other shrunken droplet and from ∂ if D is a DYD. By Lemma 7.3.19 $[\kappa]_{\mathring{D}} = [\kappa]_{\mathring{D}\cup\partial}$ is within distance C_1 of κ. Hence,

$$[K_0 \cup \partial] = \mathring{D}_0 \cup \partial \cup [\kappa_0] \cup \bigcup_{\kappa,D}[\kappa]_{\mathring{D}}, \tag{7.4}$$

where the last union is on couples (κ, D) as above. Indeed, all $[\kappa]_{\mathring{D}}$ and $[\kappa]$ (for different κ) are at distance at least $C_2 - 2C_1$ from each other and from $\mathring{D}_0\backslash\mathring{D}$ (by the reasoning above), so for each site of Λ the intersection of the ball of radius $O(1)$ centered at it with the set on the right-hand side of (7.4) coincides with the intersection with one of the sets $[\kappa \cup \mathring{D}]$, $[\kappa]$ or $\mathring{D}_0 \cup \partial$, which are all stable, so no infections occur, which proves (7.4).

The claim follows easily from (7.4), since for every couple κ, D the set $[\kappa]_{\mathring{D}}$ is within distance C_1 of κ, which is itself at distance at most C_2 from $\mathring{D}_0$, and κ has diameter much smaller than C_3 by Corollary 7.3.17. □

Let C' be a modified cluster of $[K]_\partial$ and assume for a contradiction that $C' \subset [\kappa_0]_\partial$. From Definition 7.3.3 we get that C' is also a modified cluster of $[\kappa_0]_\partial$, but this is a contradiction, since $[\kappa_0]_\partial$ only consists of modified crumbs.

Since any modified cluster C' of $[K]_\partial$ has diameter at most C_3 (by Definition 7.3.3) and intersects $[K]_\partial \backslash [\kappa_0]_\partial$, which is within distance C_3 of $\mathring{D}_0$ by Claim 3, we get that C' is within distance $2C_3$ of $\mathring{D}_0$. Therefore, $\bigcup_{C' \in \mathcal{C}'([K]_\partial)} Q'(C') \subset D_0 \cup \partial$, where the union is over all modified clusters of $[K]_\partial$, since $\operatorname{diam}(Q'(C')) \ll C_4/C_1 \leqslant d(\mathring{D}_0, \Lambda \backslash D_0)$. As $\mathcal{D}$ is the output of the droplet algorithm, D_0 is the union of disjoint DYD non-intersecting ∂ and CDYD, so it necessarily contains $\bigcup_{D' \in \mathcal{D}'} D'$ (see Remark 7.3.10), which concludes the proof. □

Remark 7.3.21. It should be noted that the algorithm is more easily and naturally defined with no boundary, but that will not be sufficient for our purposes. However, this 'free' algorithm is trivially obtained as a specialisation of ours. It is also possible to deal with more general boundaries, with infinite input sets, as well as with droplets defined by more directions and possibly with several rugged sides.

7.4 Renormalised East dynamics

In this section we map the original dynamics into an East one and conclude the proof of our main result. In Section 7.4.1 we introduce the necessary notation for the relevant geometry. In Section 7.4.2 we consider a renormalised dynamics on the slices of Figure 7.6 by algorithmically selecting certain modified spanned droplets of size $\Omega(1/q^\alpha)$. In Section 7.4.3 we further renormalise to recover an exact East dynamics where q is replaced by q_{eff} corresponding to the probability of spanning such a droplet. Finally, in Section 7.4.4 we prove Theorem 7.0.1 roughly as in [267].

7.4.1 Geometric setup

Let us start by defining the domain V we will work in, recalling the notation from Lemma 7.2.1. Roughly speaking, V is an isosceles triangle with height $e^{1/(C_5 q^\alpha)}$ directed by u' (see Figure 7.6). It is divided into 'columns' $\mathcal{C}_i$ perpendicular to u' of width roughly $1/q^\alpha$, so that the origin of $\mathbb{Z}^2$ is in the middle of the last column, close to the tip of V.

More formally, set $L = 1/(C_5 q^\alpha)$ and let ι be the smallest $x \geqslant 1$ such that the site $\frac{x}{2q^\alpha} u'$ is in $\mathbb{Z}^2$, so that $\iota = 1 + O(q^\alpha)$. This way our columns will have width ι/q^α and be separated along rational lines. We define the domain

$$V = \mathbb{H}_{u'}(e^L u') \backslash \left(\mathbb{H}_{u'_2}(-\iota/(2q^\alpha)u') \cup \mathbb{H}_{u'_1}(-\iota/(2q^\alpha)u') \right).$$

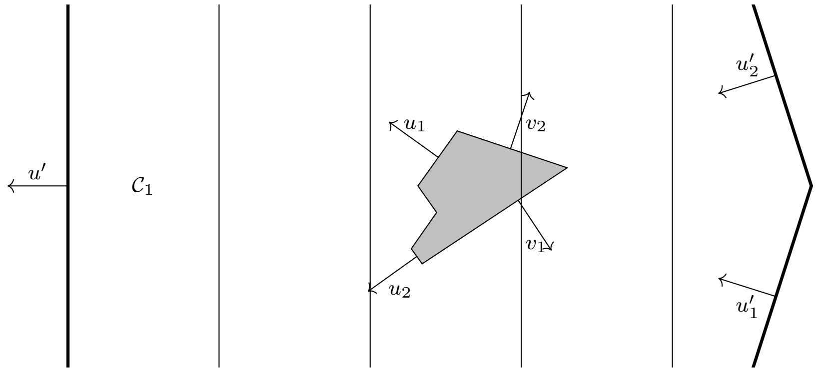

Figure 7.6 – The domain V is the thickened triangle, a portion of which is displayed. Solid lines separate columns $\mathcal{C}_i$. Inside the domain is drawn a DYD, which witnesses $\Phi(\omega)_3 = \uparrow$.

Let us choose C_5 so that half the number of columns

$$N = e^L q^\alpha/(2\iota) + 1/4 = e^L q^\alpha(1/2 + O(q^\alpha))$$

is an integer. We then partition the domain $V = \bigcup_{i=1}^{2N} \mathcal{C}_i$ into columns with

$$\mathcal{C}_i = \{x \in V : e^L - \iota(i-1)/q^\alpha > \langle x, u' \rangle \geqslant e^L - \iota i/q^\alpha\},$$

so that 0 is in the middle of $\mathcal{C}_{2N}$ and $e^L u' \in \mathbb{Z}^2$. We shall refer to $\mathcal{C}_i$ as the *i-th column*. Finally, define the half-plane containing $\mathcal{C}_{i+1}$, but not intersecting $\mathcal{C}_i$

$$\mathbb{H}_i = \mathbb{H}_{u'}((e^L - \iota i/q^\alpha)u')$$

and the natural boundary for $\mathcal{C}_i$

$$\partial_i = \mathbb{H}_i \cup \bar{\partial},$$

obtained by considering $\mathcal{C}_j, j \geqslant i+1$ as fully infected, where

$$\bar{\partial} = \mathbb{H}_{u'_2}(-\iota/(2q^\alpha)u') \cup \mathbb{H}_{u'_1}(-\iota/(2q^\alpha)u').$$

Note that these boundaries are of the form considered in Section 7.3.

7.4.2 Arrow variables

Let $\omega \in \Omega$. We will now define a collection of arrow variables which depend only on the restriction of ω to V. We naturally identify the restriction of ω to V with the subset of V where ω is 0 and we use the notation $\omega = \varnothing$ to indicate that all sites are filled (healthy) in V, namely $\omega_x = 1$ for all $x \in V$. Let $\omega^{(0)} = \omega \cap V$. We define the position of the first *up-arrow* as the smallest index $i_1(\omega) \in \{1, 2, \ldots, 2N\}$ such that there is a modified spanned droplet of size at least L for $[\omega^{(0)}]_{\partial_{i_1(\omega)}}$ with boundary $\partial_{i_1(\omega)}$. If no such i_1 exists, we say that there are no up-arrows and set $i_1(\omega) = \infty$. We further denote $\omega^{(1)} = \omega^{(0)} \cap \mathbb{H}_{i_1(\omega)}$ as soon as $i_1(\omega) < \infty$, while otherwise $\omega^{(1)} = \varnothing$.

We define the set $I(\omega) = \{i_1(\omega), i_2(\omega), \dots\} \subset \{1, \dots, 2N\}$ containing the positions of up-arrows recursively as follows. If there are no up-arrows, then $I = \varnothing$. Otherwise, we set $I(\omega) = \{i_1(\omega)\} \cup I(\omega^{(1)})$ and $\omega^{(k)} = (\omega^{(k-1)})^{(1)}$, which defines $\omega^{(k)}$ for all k. Let us note that if $i_1(\omega) \neq \infty$, then $i_1(\omega) < i_1(\omega^{(1)})$, since by definition $[\omega^{(1)}]_{\partial_{i_1(\omega)}} = \varnothing$. Finally, we may define $\Phi(\omega) \in \{\uparrow, \downarrow\}^{\{1,\dots,2N\}}$ as

$$\Phi(\omega)_k = \begin{cases} \uparrow & \text{if } k \in I(\omega), \\ \downarrow & \text{otherwise.} \end{cases}$$

The next Lemma states that the probability to find at least one up-arrow decays as

$$q_{\text{eff}} = e^{-L}.$$

Lemma 7.4.1.

$$\mu(i_1 < \infty) \leqslant q_{\text{eff}}.$$

Proof. Fix $1 \leqslant i \leqslant 2N$ and consider the event $i_1 = i$. It is clearly included in the event E_i that there is a modified spanned droplet of size at least L for $[\omega^{(0)}]_{\partial_i}$ with boundary ∂_i. By Proposition 7.3.20 there is also a spanned droplet of size at least L/C_1 for $\omega^{(0)} \backslash \partial_i$ with boundary ∂_i. By Lemma 7.3.13 this implies that there is also a spanned connected droplet of size between L/C_1^2 and $2L/C_1^2$. Then one can rewrite E_i as the union over all such droplets D of the event that D is spanned. Note that for each discretised DYD $D \cap \mathbb{Z}^2$ the event that there exists a spanned DYD D' with $D' \cap \mathbb{Z}^2 = D \cap \mathbb{Z}^2$ coincides with the event that a suitably chosen such D'_0 is spanned. Indeed, the intersection of two DYD is a DYD by (7.2) and the spanning of all D' depend only on the finite number of sites in $D \cap \mathbb{Z}^2$, so there is a finite number of possible events associated to different D' and one can consider the intersection of a D' defining each of these events. The same reasoning holds for CDYD and so for each discretised droplet $D \cap \mathbb{Z}^2$ one can bound the probability that there exists a spanned droplet with such discretisation using Lemma 7.3.15. Thus, by the union bound on discretised droplets counted in Observation 7.3.8, one obtains

$$\mu(E_i) \leqslant |V|.e^L 2e^{-C_4 L/C_1^2} \leqslant q_{\text{eff}}/(2N). \qquad \square$$

We next consider the event of having at least n up-arrows

$$\mathcal{B}(n) = \{\omega \in \Omega : |I(\omega)| \geqslant n\}.$$

Corollary 7.4.2. *For any* $1 \leqslant n \leqslant 2N$ *we have*

$$\mu(\mathcal{B}(n)) \leqslant q_{\text{eff}}^n.$$

Proof. We prove the statement by induction on n. The base, $n = 1$, is given by Lemma 7.4.1. For $n > 1$ we have

$$\begin{aligned}\mu(|I| \geqslant n) &= \sum_{i=1}^{2N} \mu(i_1(\omega) = i; |I(\omega \cap \mathbb{H}_i)| \geqslant n-1) \\ &\leqslant \sum_{i=1}^{2N} \mu(i_1 = i)\mu(|I| \geqslant n-1) \\ &\leqslant q_{\text{eff}}^n,\end{aligned}$$

where we used that the event $i_1 = i$ only depends on $\omega \backslash \mathbb{H}_i$ (i_1 is a stopping time for the filtration induced by the columns) and that the event $|I| \geqslant n-1$ is increasing for the order defined by $\omega \preceq \omega'$ when $\omega \subset \omega'$. □

We will now state a key deterministic property of the arrows under legal moves of the KCM dynamics.

Lemma 7.4.3. *Let $\omega \in \Omega$. Let $x \in \mathcal{C}_i$ be such that $\omega_x = 1$ and the constraint at x is satisfied by $\omega \cup \bar{\partial}$. Assume that $\Phi(\omega) \neq \Phi(\omega^x)$. Let $j = \max\{k\colon \Phi(\omega)_k \neq \Phi(\omega^x)_k\}$. Then*

$$\begin{aligned}\Phi(\omega)_{[0,i-1]} &= \Phi(\omega^x)_{[0,i-1]}, & \Phi(\omega)_{[i-1,j]} &= (\uparrow,\downarrow,\uparrow,\downarrow,\uparrow,\dots),\\ \Phi(\omega)_{[j+1,2N]} &= \Phi(\omega^x)_{[j+1,2N]}, & \Phi(\omega^x)_{[i-1,j]} &= (\uparrow,\uparrow,\downarrow,\uparrow,\downarrow,\dots)\end{aligned}$$

with the convention that $\Phi(\omega)_0 =\uparrow$ for all ω.

Proof. We denote $\Phi := \Phi(\omega)$ and $\Phi' := \Phi(\omega^x)$. Clearly, $\Phi_{[0,i-1]} = \Phi'_{[0,i-1]}$, since those values do not depend on $\omega \cap \mathbb{H}_{i-1}$.

Claim 1. Let $k \geqslant i$. If $\Phi_k =\uparrow$, then $\Phi_{[k+1,2N]} \geqslant \Phi'_{[k+1,2N]}$ for the lexicographic order associated to $\uparrow<\downarrow$. If $\Phi'_k =\uparrow$, then $\Phi_{[k+1,2N]} \leqslant \Phi'_{[k+1,2N]}$.

Proof of Claim 1. The two assertions being analogous, we only prove the first one, so assume that $\Phi_k =\uparrow$. Let $j' = \min\{l > k\colon \Phi_l =\uparrow\}$. Then there is a modified spanned droplet of size at least L for $[\omega^{(0)} \cap \mathbb{H}_k]_{\partial_{j'}}$ with boundary $\partial_{j'}$. But this is also true for ω^x instead of ω, as they coincide in $\mathbb{H}_k$, and in particular the position of the first up-arrow of Φ' after k is at most j'. □

Claim 2. Let $k \geqslant i-1$ be such that $\Phi_k = \Phi'_k =\downarrow$. Then $k > j$ i.e. $\Phi_{[k,2N]} = \Phi'_{[k,2N]}$.

Proof of Claim 2. We can clearly assume that $k < 2N$. Further assume for a contradiction that $\Phi_{k+1} =\uparrow$ and $\Phi'_{k+1} =\downarrow$. Let $i' = \max\{l < k\colon \Phi_l =\uparrow\}$. Then there exists a modified spanned droplet D of size at least L for $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_{k+1}}$ with boundary ∂_{k+1}. By Lemma 7.3.13 we can assume that $L \leqslant |D| \leqslant C_1 L$. However, if $d(D, \mathcal{C}_{k+1}) > C_5$, then D is also modified spanned for $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_k}$ with boundary ∂_k, contradicting the definition of i'. Indeed, from the output of the modified droplet algorithm for $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_k} \cap D$ with boundary ∂_k we can create a collection $\hat{\mathcal{D}}$ of droplets for ∂_{k+1} by extending CDYD appropriately, thus $\hat{\mathcal{D}}$ contains $Q'(C')\backslash\partial_k = Q'(C')\backslash\partial_{k+1}$ for every modified cluster C' of $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_k} \cap D$ with boundary ∂_k. Moreover, the modified clusters of $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_{k+1}} \cap D$ with boundary ∂_{k+1} are contained in the modified clusters of $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_k} \cap D$ with boundary ∂_k, so $\hat{\mathcal{D}}$ contains the output of the modified droplet algorithm for $[\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_{k+1}} \cap D$ with boundary ∂_{k+1} by Remark 7.3.10, itself containing D.

Therefore, $d(D, \mathcal{C}_{k+1}) \leqslant C_5$. Moreover, D is not modified spanned for $[(\omega^x)^{(0)} \cap \mathbb{H}_{k-1}]_{\partial_{k+1}}$ with boundary ∂_{k+1} (otherwise $\Phi'_{[k,k+1]} \neq (\downarrow, \downarrow)$). Therefore, there exists a site $y \in D$ such that

$$y \in [\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_{k+1}} \backslash [(\omega^x)^{(0)} \cap \mathbb{H}_{k-1}]_{\partial_{k+1}}.$$

We consider two subcases. First assume that $d(x, \mathbb{R}^2\backslash\mathbb{H}_{i-1}) \geqslant C_1$. Then, the constraint at x is satisfied by $(\omega \cap \mathbb{H}_{i-1}) \cup \bar{\partial}$, so $[\omega^{(0)} \cap \mathbb{H}_{k-1}]_{\partial_{k+1}} = [(\omega^x)^{(0)} \cap \mathbb{H}_{k-1}]_{\partial_{k+1}}$, and there is a path

$$P \subset [\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_{k+1}} \backslash [(\omega^x)^{(0)} \cap \mathbb{H}_{k-1}]_{\partial_{k+1}}$$

from $\mathbb{R}^2\backslash\mathbb{H}_{k-1}$ to y such that each two consecutive sites are at distance at most $O(1)$. But $d(y, \mathbb{R}^2\backslash\mathbb{H}_{k-1}) \geqslant \iota/q^\alpha - \operatorname{diam}(D) - C_5 \geqslant C_2(L+1)$, so one can find a subpath $P' \subset \mathcal{C}_k \cap P$ of diameter at least $C_2 L$. Yet, it is clear that $P' \subset [\omega^{(0)} \cap \mathbb{H}_{i'}]_{\partial_k}$ implies the existence of a modified spanned droplet of size larger than L with boundary ∂_k, so one would have an up-arrow of Φ in $[i'+1, k]$—a contradiction. If, on the contrary, $d(x, \mathbb{R}^2\backslash\mathbb{H}_{i-1}) \leqslant C_1$, we can redo the same reasoning, but P needs to extend to either $\mathbb{R}^2\backslash\mathbb{H}_{k-1}$ or x, both of which are sufficiently far from y.

Thus, $\Phi_{k+1} = \Phi'_{k+1}$, as the case $\Phi_{k+1} = \downarrow, \Phi'_{k+1} = \uparrow$ is treated identically. But then either both are $\uparrow$, in which case we are done by Claim 1 or both are $\downarrow$ and we are done by induction. □

It is easy to see that the only non-identical arrow sequences $\Phi_{[i-1,j]}$ and $\Phi'_{[i-1,j]}$ satisfying the two claims are $(\uparrow, \downarrow, \uparrow, \downarrow, \dots)$ and $(\uparrow, \uparrow, \downarrow, \uparrow, \dots)$ (in this order using that $\omega_x = 1$). Indeed, by Claims 1 and 2 $\Phi_k \neq \Phi'_k$ for all $i \leqslant k \leqslant j$, by Claim 1 one cannot have two consecutive up arrows neither in Φ nor in Φ' in the interval $[i, j]$ and by Claim 2 $\Phi_{i-1} = \Phi'_{i-1} = \uparrow$. □

7.4.3 Renormalised East dynamics

We partition $\{1, \dots, 2N\}$ into blocks $B_i = \{2i-1, 2i\}$ for $1 \leqslant i \leqslant N$. Given $\omega \in \Omega$, we define $\eta(\omega) \in \{0,1\}^{\{1,\dots,N\}}$ by

$$\eta(\omega)_i = \mathbb{1}_{\{\forall j \in B_i : \Phi(\omega)_j = \downarrow\}}$$

for all $i \in \{1, \dots N\}$. Let

$$n = \lfloor L \rfloor = \left\lfloor \frac{1}{C_5 q^\alpha} \right\rfloor < \lfloor \log_2 N \rfloor.$$

Recall the definition of legal paths, Definition 1.3.6. Given an event $\mathcal{E} \subset \Omega$ and a legal path $\gamma = (\omega_{(0)}, \dots, \omega_{(k)})$ we will say that $\gamma \cap \mathcal{E} = \varnothing$ if $\omega_{(i)} \notin \mathcal{E}$ for all $i \in \{0, \dots, k\}$. Also, given $\omega \in \Omega$ and $\mathcal{A} \subset \Omega$, we say that γ connects ω to $\mathcal{A}$ if $\omega_{(0)} = \omega$ and $\omega_{(k)} \in \mathcal{A}$. Recall that $\mathcal{B}(n) \subset \Omega$ is the set of configurations with at least n up-arrows. The following is a straightforward but important corollary of Lemma 7.4.3.

Corollary 7.4.4. *For any legal path $(\omega_{(0)}, \dots, \omega_{(k)})$, the path formed by $(\eta(\omega_{(0)}), \dots, \eta(\omega_{(k)}))$ is legal for the East model on $\{1, \dots, N\}$ with fixing $\eta_0 = 0$.*

Proof. Lemma 7.4.3 gives that $\eta(\omega_{(j)}) \neq \eta(\omega_{(j+1)})$ implies that $\Phi(\omega_{(j)})$ and $\Phi(\omega_{(j+1)})$ only differ on an alternating chain of arrows ending in some B_i, preceded by $\uparrow$. Then clearly $\eta(\omega_{(j)})_l = \eta(\omega_{(j+1)})_l$ for all $l \neq i$ and $\eta(\omega_{(j)})_{i-1} = 0$. □

Let $\Omega_\downarrow$ and $\Omega_\uparrow^{2N}$ be respectively the set of configurations which do not have up-arrows, and the set of configurations with an up-arrow in the $2N$-th column, namely

$$\begin{aligned} \Omega_\downarrow &= \{\omega \in \Omega : \Phi(\omega) = (\downarrow, \dots, \downarrow)\}, \\ \Omega_\uparrow^{2N} &= \{\omega \in \Omega : \Phi(\omega)_{2N} = \uparrow\}. \end{aligned}$$

Combining the last corollary with Proposition 1.3.7, we obtain the most important input for the proof of the main result.

Corollary 7.4.5. *For any $\omega \in \Omega_\downarrow$ there does not exist a legal path γ with $\gamma \cap \mathcal{B}(n+1) = \varnothing$ connecting ω to $\Omega_\uparrow^{2N}$.*

7.4.4 Proof of Theorem 7.0.1

To prove Theorem 7.0.1 it is sufficient to prove the lower bound for the mean infection time and use the following inequality (see [89, Theorem 4.4] and also [269, Section 2.2])

$$T_{\mathrm{rel}} \geqslant q \mathbb{E}_\mu(\tau_0). \tag{7.5}$$

However, it is instructive to construct at this stage a test function that directly gives the desired lower bound on T_{rel} without going through the comparison with the mean infection time. Indeed, the mechanism will appear more clearly this way.

Proof of Theorem 7.0.1 for T_{rel} We define the event $\tilde{\mathcal{A}}$ as

$= \{\omega \in \Omega \colon \exists$ legal path γ with $\gamma \cap \mathcal{B}(n) = \varnothing$ connecting $\omega \cup (\mathbb{Z}^2 \backslash V)$ to $\Omega_\downarrow\}$

and the test function $f : \Omega \to \{0,1\}$

$$f = \mathbb{1}_{\tilde{\mathcal{A}}}.$$

Then, by Definition 1.3 we get

$$T_{\mathrm{rel}} \geqslant \frac{\mu(\tilde{\mathcal{A}})(1-\mu(\tilde{\mathcal{A}}))}{\mathcal{D}(f)}, \tag{7.6}$$

where the Dirichlet form $\mathcal{D}(f)$ is defined in (1.4).

Lemma 7.4.6 (Bounds on $\mu(\tilde{\mathcal{A}})$)**.**

$$\mu(\tilde{\mathcal{A}})\left(1-\mu(\tilde{\mathcal{A}})\right) \geqslant \exp\left(\frac{\log q}{C_4 q^\alpha}\right).$$

Proof. By Lemma 7.4.1 we have

$$\mu(\tilde{\mathcal{A}}) \geqslant \mu(\Omega_\downarrow) \geqslant 1 - q_{\mathrm{eff}} \geqslant 1/2.$$

On the other hand,

$$1 - \mu(\tilde{\mathcal{A}}) \geqslant \mu(\Omega_\uparrow^{2N}) \geqslant q^{C_1 L} \geqslant \exp(C_1 \log q/(C_5 q^\alpha)),$$

where we used Corollary 7.4.5 for the first inequality as well as the fact that if $(\omega_{(0)}, \ldots, \omega_{(k)})$ is a legal path, then $(\omega_{(k)}, \ldots, \omega_{(0)})$ is one as well, and for the second inequality we notice that for the $2N$-th arrow to be up it is sufficient to have an infected segment of length $C_1 L$ in $\mathcal{C}_{2N}$. □

Lemma 7.4.7 (Estimate of the Dirichlet form)**.** $\mathcal{D}(f) \leqslant \exp\left(-1/(C_5^3 q^{2\alpha})\right)$.

Proof. Using the fact that $f(\omega)$ depends only on the values of ω in V, we get

$$\begin{aligned}\mathcal{D}(f) &= \sum_{x \in V} \mu(c_x \mathrm{Var}_x(f)) \\ &= q(1-q) \sum_{x\in V} \mu\left(c_x \mathbb{1}_{\{\omega \in \tilde{\mathcal{A}},\, \omega^x \notin \tilde{\mathcal{A}}\}} + c_x \mathbb{1}_{\{\omega \notin \tilde{\mathcal{A}}, \omega^x \in \tilde{\mathcal{A}}\}}\right) \leqslant |V| \mu(\mathcal{B}(n-1)),\end{aligned} \tag{7.7}$$

since, by Lemma 7.4.3 $||I(\omega)| - |I(\omega^x)|| \leqslant 1$ when $c_x = 1$, so the indicators both imply $\omega \in \mathcal{B}(n-1)$. Indeed, $\omega \in \tilde{\mathcal{A}}$ implies the existence of a legal path γ from $\Omega_\downarrow$ to $\omega \cup (\mathbb{Z}^2 \backslash V)$ with each configuration not in $\mathcal{B}(n)$. Since $c_x = 1$, the path $\bar{\gamma}$ obtained by adding the transition from $\omega \cup (\mathbb{Z}^2\backslash V)$ to $\omega^x \cup (\mathbb{Z}^2 \backslash V)$ is also legal, thus the hypothesis $\omega^x \notin \tilde{\mathcal{A}}$ is not satisfied unless $\omega^x \in \mathcal{B}(n)$ (and similarly for $\omega \notin \tilde{\mathcal{A}}, \omega^x \in \tilde{\mathcal{A}}$). Thus, the result follows by using Corollary 7.4.2. □

Then the lower bound for T_{rel} of Theorem 7.0.1 follows from (7.6), Lemma 7.4.6 and Lemma 7.4.7.

The above proof, together with the matching upper bound of Theorem 2(a) of [269] indicate that the bottleneck dominating the time scales is the creation of $\Theta(\log(1/q_{\mathrm{eff}}))$ simultaneous droplets of probability q_{eff}.

Proof of Theorem 7.0.1 for $\mathbb{E}_\mu(\tau_0)$ The proof of the lower bound for the infection time follows a similar route, with some complications due to the fact that we have to identify a (sufficiently likely) initial set starting from which we have to go through the bottleneck configurations before infecting the origin.

By [267, Corollary 3.4], to prove the desired lower bound on $\mathbb{E}_\mu(\tau_0)$ it suffices to construct a local function $\phi = \phi_q$ such that

(i) $\mu(\phi^2) = 1$,

(ii) $\frac{\mu(\phi)^4}{\mathcal{D}(\phi)} \geqslant \exp(1/(C_5^4 q^{2\alpha}))$,

(iii) $\phi(\omega) = 0$ if $\omega_0 = 0$.

Inspired by [267] we let

$$\Omega_g = \Omega_\downarrow \cap \{\omega \in \Omega : \omega_{\Lambda_0} = 1\}$$

where $\Lambda_0 = \{x \in \mathbb{Z}^2 : d(x,0) \leqslant 1/(4q^\alpha)\} \subset \mathcal{C}_{2N}$ and $\mathcal{A}$ be the event

$$\{\omega \in \Omega : \exists \text{ a legal path } \gamma \text{ with } \gamma \cap \mathcal{B}(n) = \varnothing \text{ connecting } \omega \cup (\mathbb{Z}^2 \backslash V) \text{ to } \Omega_g\}.$$

Then we set

$$\phi(\cdot) = \mathbb{1}_{\mathcal{A}}(\cdot)/\mu(\mathcal{A})^{1/2}. \tag{7.8}$$

We are now left with proving that this function satisfies (i)-(iii) above.

Property (i) follows immediately from (7.8). In order to verify (ii) we start by establishing a lower bound on $\mu(\mathcal{A})$. By definition it holds that

$$\mu(\mathcal{A}) \geqslant \mu(\Omega_g) \geqslant \mu(\omega_{\Lambda_0} = 1)\mu(\Omega_\downarrow) \geqslant e^{-O(1)/q^{2\alpha-1}}(1 - q_{\mathrm{eff}}) = e^{-O(1)/q^{2\alpha-1}}, \tag{7.9}$$

where we used Harris' inequality [202] ($\{\omega_{\Lambda_0} = 1\}$ and $\Omega_\downarrow$ are increasing events if we consider that $\omega \leqslant \omega'$ when $\omega_x \leqslant \omega'_x$ for all $x \in \mathbb{Z}^2$), Lemma 7.4.1 and $|\Lambda_0| = O(1/q^{2\alpha})$.

Furthermore, one can repeat the proof of Lemma 7.4.7 to obtain

$$\mathcal{D}(\phi) \leqslant e^{-1/(C_5^3 q^{2\alpha})}. \tag{7.10}$$

Thus, recalling (7.9), Property (ii) holds.

We are therefore only left with proving the next lemma establishing Property (iii), completing the proof of Theorem 7.0.1.

Lemma 7.4.8. *Let ω be such that $\omega_0 = 0$. Then any legal path connecting Ω_g to ω intersects $\mathcal{B}(n)$.*

As in the lower bound on $1 - \mu(\tilde{\mathcal{A}})$ for T_{rel}, the proof relies on Corollary 7.4.5, but an additional complication arises due to the fact that infecting the origin does not a priori require creating a critical droplet nearby.

Proof of Lemma 7.4.8. Suppose for a contradiction that there exists a configuration ω with $\omega_0 = 0$, a configuration $\omega_{(0)} \in \Omega_g$ and a legal path $\gamma = (\omega_{(0)}, \dots, \omega_{(k)})$ with $\omega_{(k)} = \omega$ and $\omega_{(j)} \notin \mathcal{B}(n)$ for all $j \in \{0, \dots, k\}$. Assuming without loss of generality that $\omega_{(j)} \neq \omega_{(j-1)}$ for all j, let x_j be such that $\omega_{(j)} = (\omega_{(j-1)})^{x_j}$. Consider the path $\tilde{\gamma} = (\tilde{\omega}_{(0)}, \dots, \tilde{\omega}_{(k)})$ obtained by performing the same updates as for γ except for flips in the column $\mathcal{C}_{2N}$, which are performed only if they correspond to infecting sites. More precisely, we let $\tilde{\omega}_{(0)} = \omega_{(0)}$ and

$$\tilde{\omega}_{(j)} = \begin{cases} (\tilde{\omega}_{(j-1)})^{x_j} & \text{if } x_j \notin \mathcal{C}_{2N} \text{ or } (\tilde{\omega}_{(j-1)})_{x_j} = 1, \\ \tilde{\omega}_{(j-1)} & \text{otherwise.} \end{cases}$$

It is not difficult to verify by induction that $\tilde{\gamma}$ is also a legal path with $\tilde{\omega}_{(j)} \leqslant \omega_{(j)}$ for all j (where $\omega \leqslant \omega'$ when $\omega_x \leqslant \omega'_x$ for all $x \in \mathbb{Z}^2$) and that $\tilde{\omega}_{(j)}$ and $\omega_{(j)}$ coincide outside of $\mathcal{C}_{2N}$. Then $(\tilde{\omega}_{(k)})_0 \leqslant (\omega_{(k)})_0 = 0$ and by definition $(\tilde{\omega}_{(0)})_{\Lambda_0} = 1$. Therefore, since inside $\mathcal{C}_{2N}$ each site that has been infected in γ is also infected in $\tilde{\omega}_{(k)}$, we conclude that necessarily $\tilde{\omega}_{(k)} \cap \mathcal{C}_{2N}$ contains a (modified) spanned droplet of size $1/(4C_1 q^{\alpha}) > L$ with boundary $\partial_{2N} = \bar{\partial}$. Indeed, there is a path of sites x with steps of size $O(1)$ from $\mathbb{Z}^2 \backslash \Lambda_0$ to 0 such that $(\tilde{\omega}_{(k)})_x = 0$. This means that $\tilde{\omega}_{(k)} \in \Omega_{\uparrow}^{2N}$. Furthermore, for all j we have $\Phi(\tilde{\omega}_{(j)})_{[1,2N-1]} = \Phi(\omega_{(j)})_{[1,2N-1]}$, as those do not depend on the sites in $\mathcal{C}_{2N}$. Thus, using Corollary 7.4.5, together with the facts that $\tilde{\omega}_{(0)} \in \Omega_g \subset \Omega_{\downarrow}$, $\tilde{\omega}_{(k)} \in \Omega_{\uparrow}^{2N}$ and $\tilde{\gamma} \cap \mathcal{B}(n+1) = \varnothing$, we reach a contradiction. □

Chapter 8

Refined universality for critical KCM: lower bounds

This chapter is based on joint work with Laure Marêché [213], proving the lower bounds of Theorem 1.6.4, leaving out the classes (d) and (g), which follow directly from Theorem 1.6.3 (recall Section 1.6 and particularly Section 1.6.4).

8.1 Supercritical rooted dynamics of droplets

8.1.1 Setting and preliminaries

Let $\mathcal{U}$ be an update family. Assuming they exist, we further fix two non-collinear rational stable directions u_1 and u_2. We set $u_3 = u_1+\pi$, $u_4 = u_2+\pi$ and $\mathcal{T} = \{u_1, u_2, u_3, u_4\}$. We will simply call *parallelogram* a set of the form

$$\begin{aligned} R(a, b; c, d) &= \left\{x \in \mathbb{R}^2 \mid \langle x, u_3\rangle \in [a, c], \langle x, u_4\rangle \in [b, d]\right\} \\ &= \overline{\mathbb{H}}_{u_1}(-a) \cap \overline{\mathbb{H}}_{u_2}(-b) \cap \overline{\mathbb{H}}_{u_3}(c) \cap \overline{\mathbb{H}}_{u_4}(d) \end{aligned}$$

for real numbers $a \leqslant b, c \leqslant d$ and denote by $\mathring{R}(a, b; c, d)$ its topological interior. For parallelograms we will systematically extend definitions by translation and interchange of u_1 and u_2 (resp. u_3 and u_4).

Finally,

$$C_6 \gg C_5 \gg C_2' \gg C_1 \gg r = \max\{\|s - s'\| \mid s, s' \in U \cup \{0\}, U \in \mathcal{U}\}$$

are constants not depending on q, but only on $\mathcal{U}$ and $\mathcal{T}$, each one sufficiently large with respect to functions of the next.[1] Furthermore, we systematically assume that q is small enough, as we are interested in $q \to 0$.

[1] We use C_2' instead of C_2 and avoid C_3 and C_4 for coherence with the appendices.

Definition 8.1.1 ([70, Definitions 2.3 and 2.4])**.** A set $Z \subset \mathbb{Z}^2$ is *strongly connected* if it is connected in the graph with vertex set $\mathbb{Z}^2$ defined by $x \sim y$ if $\|x - y\| \leqslant C_2'$.

Given $K \geqslant C_1^2 C_2'$ to be specified in Section 8.2, we say a parallelogram is *critical* when its diameter is contained between K/C_1 and K.

A parallelogram D is *spanned* in η if there exists a strongly connected set $X \subset [D \cap \eta]$ such that the smallest parallelogram containing X is D.

If η, η' are two configurations, we say that $\eta \leqslant \eta'$ when $\eta_s \leqslant \eta'_s$ for all s. For instance, if a parallelogram D is spanned for η, then is is also spanned for any $\eta' \leqslant \eta$. This order should not be confused with the (inverted) one induced by inclusion when viewing η as its set of infections.

Notice that the event that a given parallelogram D is spanned depends only on η_D and does not occur when η_D contains no infections. We further state two immediate consequences of Definition 8.1.1 for future reference.

Observation 8.1.2. Let $R \subset \mathbb{R}^2$. Then every parallelogram D spanned in η_R intersects R.

Observation 8.1.3. Let η be a configuration and X be a strongly connected component of $[\eta]$. Then $X = [\eta \cap X]$.

Proof. By maximality of a strongly connected component, $[\eta \cap X] \subset X$ is at distance at least $C_2' > 2r$ from other strongly connected components X of $[\eta]$. Thus,

$$[\eta] = \bigsqcup_Y [\eta \cap Y],$$

where the union is on all strongly connected components of $[\eta]$. □

Another standard fact is the following Aizenman-Lebowitz lemma originating from [7], whose proof can be found in the appendix (Lemma 8.A.9).

Lemma 8.1.4. *Let D be a spanned parallelogram with diameter $d \geqslant C_1 C_2'$ and let $C_1 d \geqslant k \geqslant C_1 C_2'$. Then there exists a spanned parallelogram with diameter d' such that $k/C_1 \leqslant d' \leqslant k$. In particular, if $d \geqslant K/C_1$, then there exists a spanned critical parallelogram.*

We next import and adapt the notion of crossing from [70, Definition 8.17].

Definition 8.1.5 (Crossing)**.** We say that a parallelogram $R = R(a, b; 0, d)$ is *u_1-crossed* if there exists a strongly connected set in $[\mathbb{H}_{u_1} \cup (R \cap \eta)]$ intersecting both $\mathbb{H}_{u_1}$ and $\overline{\mathbb{H}}_{u_3}(a)$.

Let $C_R^{u_1}$ denote the event that there exists $\eta' \geqslant \eta$ such that R is u_1-crossed for η', but there is no spanned critical parallelogram for η'_R.

We say that a parallelogram $\Lambda = R(0, 0; L, H)$ *has no (ℓ, h)-crossing* (or simply crossing) if the event $C_R^{u_1}$ does not occur for any $R \subset \Lambda$ of the form

$R(a,0;a+\ell,H)$ and the event $C_R^{u_2}$ does not occur for any $R\subset\Lambda$ of the form $R(0,b;L,b+h)$.

Finally, we say that a site $s\in\mathbb{Z}^2$ is *locally infectable* in a configuration η if $s\in[\eta\cap(s+R(-2K,-2K;2K,2K))]$. We also denote η^s the configuration that is equal to η everywhere except at s, i.e. $\eta^s_s=1-\eta_s$ and $\eta^s_{s'}=\eta_{s'}$ for any $s'\neq s$. We then have the following property, originating from Section 5.1.

Lemma 8.1.6. *Let $\eta\in\{0,1\}^{\mathbb{Z}^2}$, $s\in\mathbb{Z}^2$, $U\in\mathcal{U}$ be such that $s+U\subset\eta$ and let $R=R(-2K,-2K;2K,2K)$. Assume that the origin is not locally infectable in η, but is locally infectable in η^s. Then there exists a critical parallelogram D spanned in η^s_R such that $D\subset R(-3K,-3K;3K,3K)$.*

Proof. By definition, $0\in[\eta^s\cap R]\backslash[\eta\cap R]$. Therefore, $s\in R$ and $s+U\not\subset R$. In particular, $d(s,0)>K$. Let X be the strongly connected component of 0 in $[\eta^s\cap R]$. By Observation 8.1.3 we have that $s\in X$, since otherwise we would have $0\in[X]\subset[\eta\cap R]$. Therefore, the smallest parallelogram containing X is spanned and has diameter at least K. In particular, by Lemma 8.1.4 there exists a critical parallelogram D spanned in η^s_R and we are done by Observation 8.1.2, since $\mathrm{diam}(D)\leqslant K$ by Definition 8.1.1. □

8.1.2 The combinatorial bottleneck

With the notation above we are ready to prove a very general deterministic bottleneck (Lemma 8.1.10 below), constituting the core this chapter, which relatively straightforwardly translates into the following bound on $\mathbb{E}_\mu(\tau_0)$.

Proposition 8.1.7. *Let T,L,H,K,ℓ,h be positive real numbers sufficiently large with respect to C_2'. Denote*

$$\rho=\max_D\mu(D\text{ is spanned}),$$

where the max is over all critical parallelograms. Also set

$$p_\leftarrow=\max_{a,b\in\mathbb{R}}\mu\left(C^{u_1}_{R(a,b;a+\ell,b+H)}\right)$$
$$p_\downarrow=\max_{a,b\in\mathbb{R}}\mu\left(C^{u_2}_{R(a,b;a+L,b+h)}\right).$$

Assume that for some integer $n\geqslant 0$ we have the following inequalities on geometry:

$$L\geqslant 3^n(11K+\ell)\qquad\qquad H\geqslant 3^n(11K+h),$$

and probability:

$$1/8\geqslant\mu(0\text{ is locally infectable})$$
$$1\geqslant T(LH)^2\max(p_\downarrow,p_\leftarrow)$$
$$1\geqslant TLH(LHK^3\rho)^{n+1}.$$

Then the $\mathcal{U}$-KCM on $\mathbb{Z}^2$ satisfies $\mathbb{E}_\mu(\tau_0)\geqslant T$.

Remark 8.1.8. Although the bootstrap percolation estimates needed to make use of this statement in higher dimensions are not yet available, let us mention that our argument is not dimension sensitive.

The proof of Proposition 8.1.7 will occupy the rest of the present section. We start by fixing ℓ, h, K as in the statement and introducing the following definitions.

Recall that for any $R \subset \mathbb{Z}^2$, we identify configurations $\eta \in \{0,1\}^R$ with their zero set $\{x \in R, \eta_x = 0\}$. Unless otherwise specified, configurations $\eta \in \{0,1\}^R$ are extended to $\{0,1\}^{\mathbb{Z}^2}$ by keeping the same zero set.

Definition 8.1.9 (Good paths and configurations)**.** For any parallelogram $R \subset \mathbb{R}^2$, configuration $\eta \in \{0,1\}^{R\cap\mathbb{Z}^2}$ and integer $n \geqslant 0$, we say that η is *n-good* when the maximum number of critical parallelograms that are disjointly[2] spanned in η is at most n and R has no crossing for η.

A *legal path* in R is a sequence $(\eta^{(j)})_{0\leqslant j\leqslant m}$ of configurations in $\{0,1\}^{R\cap\mathbb{Z}^2}$ such that for every $j \in \{0,\dots,m-1\}$, there exists $s \in R\cap\mathbb{Z}^2$ such that $\eta^{(j+1)} = (\eta^{(j)})^s$ and $(s+U)\cap R \subset \eta^{(j)}$ for some $U \in \mathcal{U}$. For any integer $n \geqslant 0$, the path is *n-good* if for every $j \in \{0,\dots,m\}$, $\eta^{(j)}$ is n-good. For any $A, B \subset \{0,1\}^{R\cap\mathbb{Z}^2}$, we say $(\eta^{(j)})_{0\leqslant j\leqslant m}$ is a path from A to B when $\eta^{(0)} \in A$ and $\eta^{(m)} \in B$ (if A or $B = \{\eta\}$, we will write η to simplify).

We denote by $G(R)$ the set of configurations in $\{0,1\}^{R\cap\mathbb{Z}^2}$ that contain no spanned critical parallelogram and such that R contains no crossing, i.e. the 0-good configurations. For any $n \in \mathbb{N}$, we define

$$V(n,R) = \{\eta \in \{0,1\}^{R\cap\mathbb{Z}^2} \mid \text{there is an } n\text{-good legal path from } G(R) \text{ to } \eta\}.$$

Finally, we define our domain sizes for the induction to come:

$$L_n = \frac{3^n-1}{2}(9K+\ell) + 3^nK \qquad H_n = \frac{3^n-1}{2}(9K+h) + 3^nK,$$

so that $L_n - L_{n-1} = 2L_{n-1} + 9K + \ell$ and $H_n - H_{n-1} = 2H_{n-1} + 9K + h$.

Lemma 8.1.10. *For any non-negative integer n, for any parallelogram $R = R(a,b;c,d)$ such that $c-a \geqslant 2L_n$ and $d-b \geqslant 2H_n$, we have that for all $\eta \in V(n,R)$, there is no spanned critical parallelogram in η intersecting $R(a+L_n, b+H_n; c-L_n, d-H_n)$.*

We first deduce Proposition 8.1.7 from Lemma 8.1.10.

Proof of Proposition 8.1.7, assuming Lemma 8.1.10. It suffices to prove that $\mathbb{P}_\mu(\tau_0 \geqslant 2T) \geqslant 1/2$. Let $\tau' = \inf\{t \geqslant 0, 0 \text{ is locally infectable in } \eta(t)\}$. Clearly, $\tau' \leqslant \tau_0$. We denote $R = R(-L/2,-H/2;L/2,H/2)$. By Lemmas 8.1.6 and 8.1.10 and Definition 8.1.9 it follows that the event $\tau' \geqslant 2T$ contains the event G defined as the intersection of the following.

[2] As is standard [357], we say that the parallelograms $R_1,\dots,R_k$ are disjointly spanned in η if one can find disjoint sets $X_1,\dots,X_k \subset \eta$ such that $\eta'_{X_i} = 0$ implies that R_i is spanned in η' for all $1 \leqslant i \leqslant k$.

G_1: $\tau' > 0$, i.e. 0 is not locally infectable in $\eta(0)$.

G_2: There is no critical parallelogram spanned in $(\eta(0))_R$.

G_3: For all $0 \leqslant t \leqslant 2T$, no $n+1$ critical parallelograms are disjointly spanned in $(\eta(t))_R$.

G_4: For all $0 \leqslant t \leqslant 2T$, R has no crossing for $(\eta(t))_R$.

We next brutally bound the probability of these four events.

By assumption $1 - \mathbb{P}_\mu(G_1) \leqslant 1/8$. By the union bound on all (discrete) critical parallelograms intersecting R (recall Observation 8.1.2) we have

$$1 - \mathbb{P}_\mu(G_2) \leqslant O(LHK^2)\rho \leqslant \frac{LHK^3\rho}{8} \leqslant \frac{1}{8(TLH)^{1/(n+1)}} \leqslant \frac{1}{8}.$$

Let N denote the number of clock rings in R between 0 and $2T$ (the clock rings are the times at which updates are attempted by the KCM, see Section 1.2.2, at each site these times form a Poisson point process of parameter 1 independent from those of other sites). Let $\eta^{(j)}$ denote the restriction of the configuration to $R \cap \mathbb{Z}^2$ after the j-th such clock ring. By the weak law of large numbers we have $\mathbb{P}_\mu(N \geqslant O(TLH)) \leqslant 1/16$. If for any $\eta \in \{0,1\}^{R\cap\mathbb{Z}^2}$ we write $\mathcal{D}_{n+1}(\eta) = \{\text{there are } n+1 \text{ parallelograms disjointly spanned in } \eta\}$, then by stationarity, the BK inequality [357] and the union bound, we get

$$\begin{aligned} 1 - \mathbb{P}_\mu(G_3) &\leqslant \mathbb{P}_\mu(N \geqslant O(TLH)) + \sum_{j=0}^{O(TLH)} \mathbb{P}_\mu(\mathcal{D}_{n+1}(\eta^{(j)})) \\ &\leqslant \frac{1}{16} + O(TLH)\mu(\mathcal{D}_{n+1}((\eta)_R)) \\ &\leqslant \frac{1}{16} + TLH(O(LHK^2)\rho)^{n+1} \leqslant \frac{1}{8}. \end{aligned}$$

Similarly, we have

$$1 - \mathbb{P}_\mu(G_4) \leqslant \frac{1}{16} + TLH.O(L+H)\max(p_\downarrow, p_\leftarrow) \leqslant \frac{1}{8}.$$

We then conclude that

$$1 - \mathbb{P}_\mu(\tau_0 \geqslant 2T) \leqslant \sum_{i=1}^{4}(1 - \mathbb{P}_\mu(G_i)) \leqslant \frac{1}{2}$$

by the assumptions of the proposition. □

Proof of Lemma 8.1.10. We will prove the lemma by induction on n. For any n let us call $\mathcal{H}_n$ the statement of the lemma for n. $\mathcal{H}_0$ holds by definition. Let $n \geqslant 1$ and assume that $\mathcal{H}_{n-1}$ holds. Let $R = R(a,b;c,d)$ be such that $c - a \geqslant 2L_n$ and $d - b \geqslant 2H_n$. We will prove Lemma 8.1.10 by showing $\mathcal{H}_n$ using the following result, whose proof we postpone for the moment.

Lemma 8.1.11. *For all $\eta \in V(n,R)\backslash G(R)$, there exists a critical parallelogram not intersecting*

$$R' = R(a + L_n - L_{n-1}, b + H_n - H_{n-1}; c - (L_n - L_{n-1}), d - (H_n - H_{n-1}))$$

that is spanned in η (see Figure 8.1).

From Lemma 8.1.11 we have that a spanned critical parallelogram remains outside R', so there are at most $n-1$ critical parallelograms disjointly spanned in $\eta_{R'}$ for any $\eta \in V(n,R)$. Hence, for any $\eta \in V(n,R)$ we have $\eta_{R'} \in V(n-1,R')$, so we can apply $\mathcal{H}_{n-1}$ to R', which directly yields $\mathcal{H}_n$ and concludes the proof of Lemma 8.1.10. □

Consequently, it is enough to prove Lemma 8.1.11.

Proof of Lemma 8.1.11. Let us start by introducing the following geometric regions, represented in Figure 8.1.

$$\begin{aligned} R_\ell &= \mathring{R}(a + 2L_{n-1} + 7K, b + 2H_{n-1} + 7K; \\ &\qquad c - L_{n-1} - 7K, d - 2H_{n-1} - 7K) \\ R_{\ell,\downarrow} &= R(a, b; a + 2L_{n-1} + 7K, d) \\ R_{\ell,\uparrow} &= R(c - 2L_{n-1} - 7K, b; c, d) \\ R'_{\ell,\downarrow} &= R(a + L_{n-1}, b + H_{n-1}, a + L_{n-1} + 7K, d - H_{n-1}) \\ R'_{\ell,\uparrow} &= R(c - L_{n-1} - 7K, b + H_{n-1}, c - L_{n-1}, d - H_{n-1}) \end{aligned}$$

We also define similar regions with index h instead of ℓ. Moreover, we recall R' from the statement of Lemma 8.1.11. We further define the two frames (see Figure 8.1)

$$\begin{aligned} B &= R'_{\ell,\downarrow} \cup R'_{\ell,\uparrow} \cup R'_{h,\downarrow} \cup R'_{h,\uparrow} \\ B' &= R(a + L_{n-1} + 3K, b + H_{n-1} + 3K; c - L_{n-1} - 3K, d - H_{n-1} - 3K)\backslash \\ &\qquad \mathring{R}(a + L_{n-1} + 4K, b + H_{n-1} + 4K; c - L_{n-1} - 4K, d - H_{n-1} - 4K). \end{aligned}$$

We next define some more subtle regions depending on the configuration.

Claim 8.1.12. Let $\eta \in \{0,1\}^{R\cap\mathbb{Z}^2}$ be such that every critical parallelogram spanned in η intersects R'. Then there exists a closed contour $\gamma \subset \mathbb{R}^2$ (that is, a self-avoiding and closed path obtained by connecting sites of $\mathbb{Z}^2$ by straight lines linking a site to its left, top, right or bottom neighbour) satisfying the following properties:

- $\gamma \subset B'$.
- $d(\gamma, R\backslash B') \geqslant C_1^2$.
- Every $s \in \bar{\gamma}$ is not locally infectable in η, where

$$\bar{\gamma} = \{s \in \mathbb{Z}^2 | d(s,\gamma) \leqslant C_1\}.$$

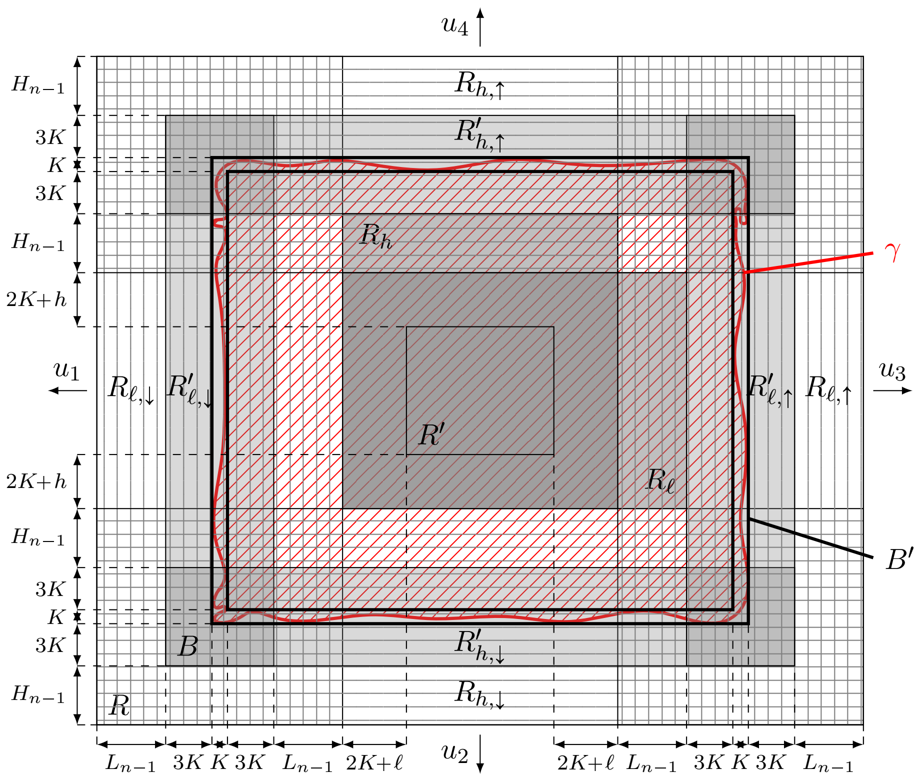

Figure 8.1 – The setting of the proof of Lemma 8.1.11. For the figure we assume that $u_3 = 0$ and $u_4 = \pi/2$. B' is the frame with thickened boundary, R_ℓ and R_h are the overlapping regions in dark gray. The regions $R'_{\ell,\downarrow}$, $R'_{h,\downarrow}$, $R'_{\ell,\uparrow}$ and $R'_{h,\uparrow}$ are in lighter gray and the frame formed by their union is B. The horizontally (resp. vertically) hatched regions are $R_{h,\downarrow}$ and $R_{h,\uparrow}$ (resp. $R_{\ell,\downarrow}$ and $R_{\ell,\uparrow}$). The contour inside B' is γ and its diagonally hatched interior is $\mathring{\gamma}$. All the regions drawn are closed subsets of $\mathbb{R}^2$ with the exception of R_ℓ, R_h and $\mathring{\gamma}$, which are open. The thicker version, $\bar{\gamma}$, of γ and the set $F \subset \mathbb{Z}^2$ defined in (8.1) are not drawn.

- The (topologically open) interior, $\mathring{\gamma} \subset \mathbb{R}^2$, defined by γ contains R'.

Proof. We proceed by renormalization. Let H denote the regular hexagon centred at the origin with diameter C_1^3 and having two horizontal sides. Consider the tiling of the plane with translates of H and denote by $\mathbb{T}$ the triangular lattice formed by their centres. Let $T = \{t \in \mathbb{T} | H + t \subset B'\}$ be the sites of $\mathbb{T}$ corresponding to B'. We say that a site $t \in T$ is *open* if no site in $(t + H) \cap \mathbb{Z}^2$ is locally infectable in η.

If there exists a contour of open sites in T surrounding R' (where a contour in $\mathbb{T}$ is a self-avoiding and closed path in the graph $(\mathbb{T}, \{(t, t') \in \mathbb{T}^2 | t + H \text{ and } t' + H \text{ share a side}\})$, we may choose γ approximating this contour, which clearly satisfies the conditions of the claim. Assume that such a contour does not exist. Then there is a path of closed sites in T from the inner to the outer boundary of T. In particular, this path yields a strongly connected set X of sites of $\mathbb{Z}^2$ that are locally infectable in η, with diameter at least $K - 4C_1^3$, contained in either the "left part" of the frame B', defined as

$$R''_{\ell,\downarrow} = R(a + L_{n-1} + 3K, b + H_{n-1} + 3K; a + L_{n-1} + 4K, d - H_{n-1} - 3K),$$

(see Figure 8.1) or in the top, right or bottom part of B', defined similarly. Without loss of generality, assume that X is contained in $R''_{\ell,\downarrow}$. Since the sites of X are locally infectable in η, they are infectable in $\eta_{R'''_{\ell,\downarrow}}$, where

$$R'''_{\ell,\downarrow} = R(a + L_{n-1} + K, b + H_{n-1} + K; a + L_{n-1} + 6K, d - H_{n-1} - K).$$

We denote X' the strongly connected component of $[\eta_{R'''_{\ell,\downarrow}}]$ containing X, and we consider the smallest parallelogram containing X'. By Observation 8.1.3, it is spanned in $\eta_{R'''_{\ell,\downarrow}}$, thus by Lemma 8.1.4 there exists a critical parallelogram spanned in $\eta_{R'''_{\ell,\downarrow}}$, hence spanned in η and contained in $R'_{\ell,\downarrow}$. We deduce the existence of a critical parallelogram spanned in η not intersecting R', hence a contradiction. □

We fix $\eta \in \{0,1\}^{R \cap \mathbb{Z}^2} \backslash G(R)$ such that R has no crossing and every critical parallelogram spanned in η intersects R'. Let

$$F = \{s \in B' \cap \mathbb{Z}^2 | s \text{ is not locally infectable in } \eta\} \tag{8.1}$$

and fix a contour γ, its thickened version $\bar{\gamma} \subset F$ and its interior $\mathring{\gamma}$ as provided by Claim 8.1.12 (see Figure 8.1).

We will prove that there is no n-good legal path from η to $G(R)$. Since legal paths can be reversed, this implies that there $\eta \notin V(n, R)$, which proves Lemma 8.1.11.

Let $(\eta^{(j)})_{0 \leqslant j \leqslant m}$ be an n-good legal path with $\eta^{(0)} = \eta$. We will use an induction on $j \in \{0, \ldots, m\}$ to prove that $\eta^{(m)} \notin G(R)$. More precisely, we will prove by induction on j that the following properties hold for $j \in \{0, \ldots, m\}$.

$\mathcal{P}_j^1$ For every $\zeta \in \{\ell, h\}$, there exists a critical parallelogram contained in R_ζ spanned in $\eta^{(j)}$.

$\mathcal{P}_j^2$ The sites of F are not locally infectable in $\eta^{(j)}$.

$\mathcal{P}_j^3$ For every $(\zeta, \xi) \in \{\ell, h\} \times \{\downarrow, \uparrow\}$, $\eta^{(j)}_{R_{\zeta,\xi}} \in V(n-1, R_{\zeta,\xi})$.

$\mathcal{P}_j^4$ $[\eta^{(j)}_{\mathring{\gamma}}] = [\eta^{(0)}_{\mathring{\gamma}}]$.

If $\mathcal{P}_m^1$ is satisfied, then there exists a critical parallelogram spanned in $\eta^{(m)}$, so $\eta^{(m)} \notin G(R)$, which proves Lemma 8.1.11, so it suffices to establish the induction.

Base: j=0. Firstly, by assumption, every critical parallelogram spanned in η intersects R', so it is contained in $R_\ell \cap R_h$ and therefore $\mathcal{P}_0^1$ holds, since $\eta \notin G(R)$. Secondly, the definition of F, (8.1), implies $\mathcal{P}_0^2$. In addition, $\mathcal{P}_0^4$ is trivial, so we only need to prove $\mathcal{P}_0^3$.

Let $(\zeta, \xi) \in \{\ell, h\} \times \{\downarrow, \uparrow\}$. By the assumption on η and Observation 8.1.2 there is no critical parallelogram spanned in $\eta_{R_{\zeta,\xi}}$. Moreover, $\eta_{R_{\zeta,\xi}}$ has no crossing in $R_{\zeta,\xi}$, since η has no crossing in R. Thus, $\eta_{R_{\zeta,\xi}} \in G(R_{\zeta,\xi})$, hence $\eta_{R_{\zeta,\xi}} \in V(n-1, R_{\zeta,\xi})$, so $\mathcal{P}_0^3$ holds.

Induction step. Let $j \in \{0, \dots, m-1\}$, and suppose that $\mathcal{P}_j^1$, $\mathcal{P}_j^2$, $\mathcal{P}_j^3$ and $\mathcal{P}_j^4$ hold. Since $(\eta^{(k)})_{0 \leqslant k \leqslant m}$ is a legal path, we have $\eta^{(j+1)} = (\eta^{(j)})^s$ and $(s+U) \cap R \subset \eta^{(j)}$ for some $s \in R \cap \mathbb{Z}^2$ and $U \in \mathcal{U}$.

Our first claim means that "F shields $\mathring{\gamma}$ from the influence of the exterior."

Claim 8.1.13. $\mathcal{P}_{j+1}^4$ holds.

Proof. If $s \notin \mathring{\gamma}$, then $\eta^{(j+1)}_{\mathring{\gamma}} = \eta^{(j)}_{\mathring{\gamma}}$, so $[\eta^{(j+1)}_{\mathring{\gamma}}] = [\eta^{(j)}_{\mathring{\gamma}}] = [\eta^{(0)}_{\mathring{\gamma}}]$ by $\mathcal{P}_j^4$. Moreover, if $s \in \mathring{\gamma}$, then $s + U \subset \mathring{\gamma} \cup \bar{\gamma}$. Furthermore, the sites of $\bar{\gamma}$ are in F, so by $\mathcal{P}_j^2$ they are not locally infectable in $\eta^{(j)}$ and, in particular, $\eta^{(j)}_{\bar{\gamma}} = 1$. Thus, $s + U \subset \mathring{\gamma}$ and so $[\eta^{(j+1)}_{\mathring{\gamma}}] = [\eta^{(j)}_{\mathring{\gamma}}] = [\eta^{(0)}_{\mathring{\gamma}}]$ by $\mathcal{P}_j^4$. □

We next establish that at least one spanned critical parallelogram of $\eta^{(j+1)}_{\mathring{\gamma}}$ remains "to the left" of $R_{\ell,\uparrow}$ (see Figure 8.1), as well as one "below" $R_{h,\uparrow}$ (these two parallelograms may be the same).

Claim 8.1.14. There exists a critical parallelogram contained in $\mathbb{H}_{u_3}(c - 2L_{n-1} - 7K)$ that is spanned in $\eta^{(j+1)}_{\mathring{\gamma}}$ and similarly for $\mathbb{H}_{u_4}(d - 2H_{n-1} - 7K)$.

Proof. We will only treat $H = \mathbb{H}_{u_3}(c - 2L_{n-1} - 7K)$ as the argument for the other half-plane is the same. Assume for a contradiction that there is no critical parallelogram contained in H that is spanned in $\eta^{(j+1)}_{\mathring{\gamma}}$.

By Claim 8.1.13 we have that $[\eta^{(j+1)}_{\mathring{\gamma}}] = [\eta^{(0)}_{\mathring{\gamma}}]$ and by assumption there exists a critical parallelogram D spanned in $\eta^{(0)}$ intersecting R'. D is then contained in $\mathring{\gamma}$, hence spanned in $\eta^{(0)}_{\mathring{\gamma}}$. We consider X a strongly connected set of $[D\cap\eta^{(0)}_{\mathring{\gamma}}]$ such that the smallest parallelogram containing X is D, then X intersects $\overline{\mathbb{H}}_{u_3}(c-2L_{n-1}-9K-\ell)$. Therefore, if we call X' the strongly connected component of $[\eta^{(j+1)}_{\mathring{\gamma}}]$ containing X, then X' intersects $\overline{\mathbb{H}}_{u_3}(c-2L_{n-1}-9K-\ell)$, and the smallest parallelogram D' containing X' contains X, thus it contains D hence it has diameter at least K/C_1, since D is critical. By Observation 8.1.3, D' is then spanned in $\eta^{(j+1)}_{\mathring{\gamma}\cap X'}$, hence by Lemma 8.1.4 there exists a critical parallelogram spanned for $\eta^{(j+1)}_{\mathring{\gamma}\cap X'}$. By the assumption it cannot be contained in H, so X' intersects $\overline{\mathbb{H}}_{u_1}(-(c-2L_{n-1}-8K))$. Now, if we denote $a_0=\max\{a'\mid X'\subset\overline{\mathbb{H}}_{u_1}(-a')\}$, then $a_0\geqslant a$, as $X'\subset[\eta^{(j+1)}_{\mathring{\gamma}}]$ and u_1 is a stable direction, and $a_0\leqslant c-2L_{n-1}-9K-\ell$. However, by Observation 8.1.3, $X'=[\eta^{(j+1)}_{\mathring{\gamma}\cap X'}]$, so $R_{X'}=R(a_0,b,a_0+\ell,d)$ is u_1-crossed for $\eta^{(j+1)}_{\mathring{\gamma}}$. Furthermore, any critical parallelogram intersecting $R_{X'}$ is contained in H, so by Observation 8.1.2 and the assumption, there is no critical parallelogram spanned for $\eta^{(j+1)}_{\mathring{\gamma}}\cap R_{X'}$. This yields that $C^{u_1}_{R_{X'}}$ occurs for $\eta^{(j+1)}$, which is impossible since $\eta^{(j+1)}$ has no crossing as a configuration in a n-good legal path. □

Since at least one spanned critical parallelogram is "to the left" of $R_{\ell,\uparrow}$, we will see that there are at most $n-1$ inside it (and similarly for $R_{h,\uparrow}$), which will yield the following.

Claim 8.1.15. $\eta^{(j+1)}_{R_{\ell,\uparrow}}\in V(n-1,R_{\ell,\uparrow})$ and $\eta^{(j+1)}_{R_{h,\uparrow}}\in V(n-1,R_{h,\uparrow})$.

Proof. We will only prove $\eta^{(j+1)}_{R_{\ell,\uparrow}}\in V(n-1,R_{\ell,\uparrow})$, as the other proof is similar. By $\mathcal{P}^3_j$ it suffices to prove that $\eta^{(j+1)}_{R_{\ell,\uparrow}}$ is $(n-1)$-good.

$R_{\ell,\uparrow}$ has no crossing for $\eta^{(j+1)}_{R_{\ell,\uparrow}}$ because R has no crossing for $\eta^{(j+1)}$. Let $D_1,\dots,D_k$ be critical parallelograms that are disjointly spanned in $\eta^{(j+1)}_{R_{\ell,\uparrow}}$ with maximal k. By Claim 8.1.14, there exists a critical parallelogram $D\subset\mathbb{H}_{u_3}(c-2L_{n-1}-7K)$ that is spanned in $\eta^{(j+1)}_{\mathring{\gamma}}$ and, therefore, also in $\eta^{(j+1)}$. Since D is disjoint from $R_{\ell,\uparrow}$, we deduce that $D_1,\dots,D_k,D$ are disjointly spanned in $\eta^{(j+1)}$, so $\eta^{(j+1)}$ contains $k+1$ disjointly spanned critical parallelograms. Since $\eta^{(j+1)}$ is n-good, we get $k\leqslant n-1$. □

We are now ready to prove that a spanned critical parallelogram in $\mathbb{H}_{u_4}(d-2H_{n-1}-7K)$ provided by Claim 8.1.14 is in fact in R_ℓ (and similarly for R_h). In a nutshell, Claim 8.1.14 prevents the parallelogram from being "too far up," being "too far left or down" induces a crossing and if the parallelogram is "too far right" we may apply $\mathcal{H}_{n-1}$ by Claim 8.1.15.

Claim 8.1.16. $\mathcal{P}^1_{j+1}$ holds.

Proof. We only treat R_ℓ, as R_h is similar. Assume for a contradiction that there is no critical parallelogram contained in R_ℓ spanned in $\eta^{(j+1)}$.

By Claim 8.1.14, there exists a critical parallelogram D such that $D \cap R_{h,\uparrow} = \varnothing$ that is spanned in $\eta^{(j+1)}_{\mathring{\gamma}}$. By assumption $D \not\subset R_\ell$. There are three possibilities:

a) $D \cap R_{\ell,\downarrow} \neq \varnothing$, i.e. D is "too far left;"

b) $D \cap R_{h,\downarrow} \neq \varnothing$, i.e. D is "too far down;"

c) $D \cap \overline{\mathbb{H}}_{u_1}(-(c - L_{n-1} - 7K)) \neq \varnothing$, i.e. D is "too far right."

We first assume for a contradiction that case c) occurs. Since D is spanned in $\eta^{(j+1)}_{\mathring{\gamma}}$, by Observation 8.1.2 D intersects $\mathring{\gamma}$, hence since D is critical, the intersection of D and $\overline{\mathbb{H}}_{u_1}(-(c - L_{n-1} - 7K))$ is in $R'_{\ell,\uparrow}$. However, this contradicts $\mathcal{H}_{n-1}$ applied to $\eta^{(j+1)}_{R_{\ell,\uparrow}} \in V(n-1, R_{\ell,\uparrow})$ by Claim 8.1.15.

Cases a) and b) being analogous, we assume for a contradiction that case a) occurs. Since D is spanned in $\eta^{(j+1)}_{\mathring{\gamma}}$, there exists a strongly connected set $X \subset [D \cap \eta^{(j+1)}_{\mathring{\gamma}}]$ such that D is the smallest parallelogram containing X. Let X' be the strongly connected component of $[\eta^{(j+1)}_{\mathring{\gamma}}]$ containing X, and D' be the smallest parallelogram containing X'. Then D' contains X, hence it contains D and $\operatorname{diam}(D') \geqslant K/C_1$, since D is critical. Furthermore, X' is a strongly connected component of $[\eta^{(j+1)}_{\mathring{\gamma}}] = [\eta^{(0)}_{\mathring{\gamma}}]$ by Claim 8.1.13, so $X' = [\eta^{(0)}_{\mathring{\gamma}} \cap X']$ by Observation 8.1.3, hence D' is spanned in $\eta^{(0)}_{\mathring{\gamma}\cap X'}$. By Lemma 8.1.4 there exists a critical parallelogram D'' that is spanned in $\eta^{(0)}_{\mathring{\gamma}\cap X'}$. Then D'' is spanned in $\eta^{(0)} = \eta$ and therefore intersects R', so by Observation 8.1.2, X' intersects $\overline{\mathbb{H}}_{u_1}(-a - 2L_{n-1} - 8K - \ell)$. Since D intersects $R_{\ell,\downarrow}$ by assumption, X intersects $\overline{\mathbb{H}}_{u_3}(a + 2L_{n-1} + 7K)$, thus X' intersects this half-plane as well. Now, if we denote $a_0 = \max\{a' \mid X' \subset \overline{\mathbb{H}}_{u_1}(-a')\}$, then $a_0 \geqslant a$, as $X' \subset [\eta^{(0)}_{\mathring{\gamma}}]$ and u_1 is a stable direction, and $a_0 \leqslant a + 2L_{n-1} + 7K$. However, $X' = [\eta^{(0)}_{\mathring{\gamma}\cap X'}]$, so $R_{X'} = R(a_0, b, a_0 + \ell, d)$ is u_1-crossed for $\eta^{(0)}$. However, $\eta^{(0)}$ has no crossing, so $C^{u_1}_{R_{X'}}$ does not occur for $\eta^{(0)}$. Hence, recalling Definition 8.1.5 and Observation 8.1.2, there exists a critical parallelogram spanned in $\eta^{(0)}_{R_{X'}}$, hence in $\eta^{(0)}$, intersecting $R_{X'}$. Since critical parallelograms have diameter at most K, this contradicts the fact that all critical parallelograms spanned in $\eta^{(0)} = \eta$ intersect R'. □

We are now able to use an argument similar to that of Claim 8.1.15: having a spanned critical parallelogram in R_ℓ entails that at most $n-1$ of them are in $R_{\ell,\downarrow}$ (and similarly for $R_{h,\downarrow}$), hence the following.

Claim 8.1.17. $\mathcal{P}^3_{j+1}$ holds.

Proof. Thanks to Claim 8.1.15, it remains only to prove $\eta^{(j+1)}_{R_{\ell,\downarrow}} \in V(n-1, R_{\ell,\downarrow})$ and $\eta^{(j+1)}_{R_{h,\downarrow}} \in V(n-1, R_{h,\downarrow})$. The proof is actually the same as in Claim 8.1.15, replacing $\mathbb{H}_{u_3}(c - 2L_{n-1} - 7K)$ by R_ℓ and Claim 8.1.14 by Claim 8.1.16. □

Finally, we will see that for a site of F to be locally infectable in $\eta^{(j+1)}$, by Lemma 8.1.6 one would need a critical parallelogram spanned in $\eta^{(j+1)}$ near the site, which is impossible by $\mathcal{H}_{n-1}$ since F is contained in the middle parts of the $R_{\zeta,\xi}$ and $\mathcal{P}^3_{j+1}$ holds.

Claim 8.1.18. $\mathcal{P}^2_{j+1}$ holds.

Proof. Assume for a contradiction that there exists $s' \in F \subset B'$ that is locally infectable in $\eta^{(j+1)}$. Then $\mathcal{P}^2_j$ and Lemma 8.1.6 imply that there exist $\zeta \in \{\ell, h\}$, $\xi \in \{\downarrow, \uparrow\}$ and a critical parallelogram $D \subset R'_{\zeta,\xi}$ spanned in $\eta^{(j+1)}_{R_{\zeta,\xi}}$. However, by Claim 8.1.17 $\eta^{(j+1)}_{R_{\zeta,\xi}} \in V(n-1, R_{\zeta,\xi})$, so $\mathcal{H}_{n-1}$ yields the desired contradiction. □

Claims 8.1.13, 8.1.16, 8.1.17 and 8.1.18 together establish the induction step, which completes the proof of Lemma 8.1.11. □

8.2 Application of Proposition 8.1.7

In this section we derive the lower bounds of Theorem 1.6.4 from Proposition 8.1.7. For that purpose we require some estimates on the probabilities appearing in the statement of the proposition, which are mostly proved in the appendices. We restate those results below as needed. Throughout the section $\mathcal{U}$ is a critical update family with difficulty α subject to further assumptions recalled in each subsection. Such a family admits two non-collinear rational stable directions. We set u_1 and u_2 to be two arbitrary such directions, which will be chosen differently for each class of update families. We will use the definitions of Section 8.1 with u_1 and u_2.

Let us start with the easiest estimate.

Lemma 8.2.1. *Let $q^{-1/2} \leqslant K \leqslant q^{-2\alpha}$. Then,*

$$\mu(0 \textit{ is locally infectable}) \leqslant 1/8.$$

Proof. Let $R = R(-2K, -2K; 2K, 2K)$. Since $\mathcal{U}$ is critical, $\mathrm{diam}([\eta_R]) \leqslant C_1 K$, so, starting the bootstrap percolation dynamics with η_R, the origin is either infected in time at most $C_1^3 K^2$ or not at all. We conclude using e.g. [74, Theorem 1.4]. □

Turning to the probability of spanning there are two cases to consider. For unbalanced models the following is essentially a reformulation of the most difficult result of [70].

Lemma 8.2.2. *Assume that $\mathcal{U}$ is unbalanced and $K = q^{-\alpha-1/4}$. Then, for any critical parallelogram D we have*

$$\mu(D \text{ is spanned}) \leqslant \exp\left(-\frac{(\log(1/q))^2}{C_5 q^\alpha}\right).$$

Proof. By definition if D is a spanned critical parallelogram, then there exists a strongly connected set $X \subset [\eta \cap D]$ with diameter at least $\operatorname{diam}(D)/C_1$. If X' is the strongly connected component of X in $[\eta \cap D]$, then by Observation 8.1.3, $X' = [\eta \cap D \cap X']$. Therefore, if D' is the smallest S_U-droplet in the sense of [70, Definition 2.1] with the S_U defined in [70, Lemma 6.2] such that D' contains X', then D' is internally spanned by $\eta \cap D$ (in the sense of [70, Definition 2.4][3]), and $\operatorname{diam}(D') \geqslant \operatorname{diam}(D)/C_1 \geqslant K/C_1^2$. Repeating the proof of [70, Lemma 8.37], we get that there is a critical droplet in the sense of [70, Definition 2.5] internally spanned by $\eta \cap D$. Then the union bound over such droplets and [70, Lemma 8.36] yield the desired result. □

Concerning balanced models, in Appendix 8.A (Corollary 8.A.11) we establish the following.

Lemma 8.2.3. *Assume that $\mathcal{U}$ is balanced and $q^{-\alpha}/C_5 \leqslant K \leqslant q^{-2\alpha}$. Then for any critical parallelogram D we have*

$$\mu(D \text{ is spanned}) \leqslant \exp\left(-\frac{1}{q^\alpha C_5}\right).$$

For the remaining conditions of Proposition 8.1.7 we need to distinguish the different classes of models.

8.2.1 Infinite number of stable directions

In this section we assume that $\mathcal{U}$ has an infinite number of stable directions. We then choose two rational directions $u_1 < u_2 < u_1 + \pi$ sufficiently close to each other, such that all directions in $[2u_1 - u_2, 2u_2 - u_1]$ are stable and u_1, u_2 satisfy a technical condition which the reader is advised to ignore, namely that u_1 and u_2 are constructed like the eponymous directions in the proof of Lemma 7.2.1.

[3] Technically, it is not exactly the case, as [70] uses a different choice of constants. However, their results that we use still hold in our setting.

Proof of Theorem 1.6.4(a)

For this section we further assume that $\mathcal{U}$ is unbalanced. We fix the values of the parameters of Proposition 8.1.7 as follows.

$$K = q^{-\alpha-1/4} \qquad \ell = q^{-4\alpha} \qquad L = \exp\left(\frac{(\log(1/q))^2}{C_6 q^\alpha}\right)$$
$$T = \exp\left(\frac{(\log(1/q))^4}{C_6^2 q^{2\alpha}}\right) \qquad h = q^{-4\alpha} \qquad H = \exp\left(\frac{(\log(1/q))^2}{C_6 q^\alpha}\right).$$

Then Theorem 1.6.4(a) follows directly from Proposition 8.1.7 and the upper bound of [269, Theorem 2(a)]. The hypotheses of the proposition follow from the choice of parameters, Lemmas 8.2.1 and 8.2.2, and the following bound on the crossing probabilities proved in Appendix 8.B (Lemma 8.B.4).

Lemma 8.2.4. *With the notation and assumptions above we have*

$$\max(p_{\leftarrow}, p_{\downarrow}) \leqslant \exp\left(-q^{-3\alpha}\right).$$

Proof of Theorem 1.6.4(b)

For this section we further assume that $\mathcal{U}$ is balanced. We fix the values of the parameters of Proposition 8.1.7 as follows.

$$K = q^{-\alpha} \qquad \ell = q^{-4\alpha} \qquad L = \exp\left(\frac{1}{C_6 q^\alpha}\right)$$
$$T = \exp\left(\frac{1}{C_6^2 q^{2\alpha}}\right) \qquad h = q^{-4\alpha} \qquad H = \exp\left(\frac{1}{C_6 q^\alpha}\right).$$

Then Theorem 1.6.4(b) follows directly from Proposition 8.1.7. The hypotheses of the proposition follow from the choice of parameters, Lemmas 8.2.1 and 8.2.3, and Lemma 8.2.4, which still applies.

8.2.2 Finite number of stable directions

Proof of Theorem 1.6.4(c)

In this section we assume that $\mathcal{U}$ is unbalanced, rooted and has a finite number of stable directions. Therefore, we can find rational directions u_1, u_2, u_3 such that $u_1 + \pi = u_3$, $u_2 \in (u_1, u_3)$ and $\alpha(u_i) \geqslant \alpha + 1$ for all $i \in \{1, 2, 3\}$.

We fix the values of the parameters of Proposition 8.1.7 as follows.

$$K = q^{-\alpha-1/4} \qquad \ell = q^{-\alpha-5/8} \qquad L = q^{-\alpha-3/4}$$
$$T = \exp\left(\frac{(\log(1/q))^3}{C_6 q^\alpha}\right) \qquad h = q^{-\alpha-5/8} \qquad H = q^{-\alpha-3/4}.$$

Then Theorem 1.6.4(c) follows directly from Proposition 8.1.7 and the upper bound of Theorem 4.0.1. The hypotheses of the proposition follow from the choice of parameters, Lemmas 8.2.1 and 8.2.2, and the following bound on the crossing probabilities proved in Appendix 8.B (Lemma 8.B.3).

Lemma 8.2.5. *With the notation and assumptions above we have*

$$\max(p_\leftarrow, p_\downarrow) \leqslant \exp\left(-q^{-\alpha-1/4}\right).$$

Proof of Theorem 1.6.4(e)

In this section we assume that $\mathcal{U}$ is balanced, rooted and with a finite number of stable directions. Therefore, we can find non-opposite rational directions u_1, u_2 such that $\alpha(u_1) \geqslant \alpha + 1$ and $\alpha(u_2) \geqslant \alpha + 1$. We fix the values of the parameters of Proposition 8.1.7 as follows.

$$K = 1/(C_5 q^\alpha) \qquad \ell = q^{-\alpha-1/2} \qquad L = q^{-\alpha-3/4}$$
$$T = \exp\left(\frac{\log(1/q)}{C_6 q^\alpha}\right) \qquad h = q^{-\alpha-1/2} \qquad H = q^{-\alpha-3/4}.$$

Then Theorem 1.6.4(e) follows directly from Proposition 8.1.7. The hypotheses of the proposition follow from the choice of parameters, Lemmas 8.2.1 and 8.2.3, and Lemma 8.2.5, which still holds.

Proof of Theorem 1.6.4(f)

In this section we assume that $\mathcal{U}$ is semi-directed. Therefore, we can find non-opposite rational directions u_1, u_2 such that $\alpha(u_1) = \alpha$ and $\alpha(u_2) \geqslant \alpha + 1$. We fix the values of the parameters of Proposition 8.1.7 as follows.

$$K = 1/(C_5 q^\alpha) \qquad \ell = q^{-\alpha-1/2} \qquad L = q^{-\alpha-3/4}$$
$$T = \exp\left(\frac{\log\log(1/q)}{C_6^3 q^\alpha}\right) \qquad h = \frac{\log\log(1/q)}{q^\alpha} \qquad H = \frac{(\log(1/q))^{1/4}}{q^\alpha}.$$

Then Theorem 1.6.4(f) follows directly from Proposition 8.1.7. The hypotheses of the proposition follow immediately from the choice of parameters, Lemmas 8.2.1 and 8.2.3, and the following bound on the crossing probabilities proved in Appendix 8.B (Lemma 8.B.3).

Lemma 8.2.6. *With the notation and assumptions above we have*

$$p_\leftarrow \leqslant \exp\left(-q^{-\alpha-1/4}\right) \qquad p_\downarrow \leqslant \exp\left(-\frac{\log\log(1/q)}{2C_6^2 q^\alpha}\right).$$

Appendix

8.A Bounds on spanning

Relation to previous works Let us start by explaining why additional arguments are needed, as specialists would probably expect such bounds to be automatic. In [70] two main algorithms were used—the covering and the spanning ones. The former provides bounds of the type we need but for a notion of covered droplet invoking only the initial configuration. Furthermore, its treatment involves a somewhat technical geometric result on subadditivity of the size of droplets. On the other hand, the spanning algorithm works with the closure of the initial configuration inside droplets, which obstructs obtaining results analogous to those for the covering algorithm in the same way. Yet, it is the spanning algorithm which is the most useful and particularly so for unbalanced models. In [70] an inductive multi-scale scheme was used to bootstrap the bounds on the probability of droplets being spanned from a size which is easily controlled by the more rudimentary predecessor of the covering algorithm developed in [74]. This fairly technical procedure can be circumvented using our method. Indeed, if one has bounds analogous to the ones for covered droplets up to size $1/q^{\alpha(\mathcal{U})}$, one can directly prove the result of [70] in one step, which was made there as well.

The reason why in [70] one could not directly transfer the easier bounds on covering, which were established there anyway, to spanning is that the covering algorithm there lacks the key property of being essentially closure-invariant in a sense made precise below. This property was one of the main features gained in Chapter 7 by using a less wasteful notion of cluster. Therefore, we accomplish our goal as follows. We carry through (a simplified version of) the scheme of Chapter 7 to obtain general bounds for droplets covered in the sense of Chapter 7 and we use the key closure lemma (see below) to directly transfer those to spanning. On the more technical level, we should mention that analogous bounds on spanning were established in [70] in the course of their induction, but the proof needlessly uses that the model is unbalanced and constrains the choice of directions used for defining droplets, which we will need to choose freely. Moreover, Chapter 7 made unnecessary use of the existence of strongly stable directions[4], which is only needed for treating the algorithm with boundary condition. We are thus obliged to review the proofs. The reader familiar with the details of Chapter 7 would probably be satisfied by skipping directly to Appendix 8.B and consulting the statements as needed there.

Outline The appendix is structured as follows. In section 8.A.1 we recall several results from Chapter 7, leading up to Lemma 8.A.6 providing good bounds on the probability of being covered (in a sense made precise below) and to the Closure lemma 8.A.7 relating the results of the covering algorithm for a set and for its closure. In section 8.A.2, using the latter lemma we transfer the bounds of the former one to the notion of spanning used throughout the body of the chapter. Section 8.A.3 establishes, yet

[4]Strongly stable directions are those contained in the topological interior of the set of stable directions.

again, the same bounds on the probability of spanned droplets occurring, but in the presence of an infected boundary, following the same reasoning and relying more closely on Chapter 7.

Notation For the remainder of the chapter we fix an arbitrary critical update family $\mathcal{U}$ with difficulty α. Following Chapter 7 we consider constants

$$1 \ll C_1 \ll C_2' \ll C_2 \ll C_3 \ll C_4' \ll C_4 \ll C_5 \ll C_6$$

such that each one is larger than a suitable function of the previous ones, depending on $\mathcal{T}$, $\mathcal{T}_0$, $\mathcal{S}_u$, etc. to be defined below and on $\mathcal{U}$. These constants do not depend on q, which is always assumed small enough, as we are interested in $q \to 0$.

For any finite set of directions $\mathcal{V} \subset S^1$ a $\mathcal{V}$-droplet is a set of the form $\bigcap_{v\in\mathcal{V}} \overline{\mathbb{H}}_v(a_v)$ for some $a_v \in \mathbb{R}$.

8.A.1 Covering and closure

We start by studying the covering algorithm in the spirit of Section 7.3 (but without the boundary and rugged edge present there). The reader is invited to consult that chapter for most proofs and more details, as indicated below. By definition 1.6.1 we can fix a set of non semi-isolated rational stable directions[5] $\mathcal{T}_0$ with difficulty at least α, such that the convex envelope of the elements of $\mathcal{T}_0$ contain 0 in its interior and either

- $|\mathcal{T}_0| = 3$ or
- $|\mathcal{T}_0| = 4$ and one has $\mathcal{T}_0 = \{u, v, u + \pi, v + \pi\}$ for some $u, v \in S^1$.

Let Γ be the graph with vertex set $\mathbb{Z}^2$ but with $x \sim y$ iff $\|x - y\| \leqslant C_2$.

Definition 8.A.1 (Definitions 7.3.1 and 7.3.3)**.** Fix a finite set $Z \subset \mathbb{Z}^2$. Let κ be a connected component of the subgraph of Γ induced by the vertex set Z.

- κ is a *crumb* for Z if there exists a set $P_\kappa \subset \mathbb{Z}^2$ such that $[P_\kappa] \supset \kappa$ and $|P_\kappa| = \alpha - 1$.
- If κ is not a crumb for Z, we say that a $C \subset \kappa$ is a α-*cluster* (or simply cluster) of Z if the following conditions hold
  - $\mathrm{diam}(C) \leqslant C_3$.
  - C is connected in Γ.
  - For all $x \in \kappa \backslash C$ and $y \in C$ such that $x \sim y$ in Γ we have $\mathrm{diam}(C \cup \{x\}) > C_3$.

[5]Semi-isolated stable directions are the endpoints of intervals of stable directions with nonempty interior.

It can be proved Observation 7.3.4 that any cluster contains at least α sites. Moreover, Corollary 7.3.17 yields that a crumb has diameter at most αC_2. For a cluster C we denote by $Q(C)$ the smallest $\mathcal{T}_0$-droplet containing the set $\{x \in \mathbb{R}^2 : d(x, C) \leqslant C_4\}$.

We next define the covering algorithm we will use. It is an adaptation of the droplet algorithm of Chapter 7 and should not be confused with the covering algorithms of [70, 74].

Definition 8.A.2 (Covering algorithm)**.** Given a finite set $Z \subset \mathbb{Z}^2$ of infections the *covering algorithm* outputs a set $\mathcal{D}$ of disjoint $\mathcal{T}_0$-droplets as follows.

- Form an initial collection $\mathcal{D}$ of $\mathcal{T}_0$-droplets consisting of $Q(C)$ for all clusters C of Z.
- Whenever there exist $D_1, D_2 \in \mathcal{D}$ with $D_1 \cap D_2 \neq \varnothing$, replace them with the smallest $\mathcal{T}_0$-droplet containing their union, which we denote by $D_1 \vee D_2$.
- Output the collection $\mathcal{D}$ obtained when all $\mathcal{T}_0$ droplets in $\mathcal{D}$ are disjoint.

Equivalently, $\mathcal{D}$ is the minimal collection (with respect to inclusion of the union of its elements) of disjoint $\mathcal{T}_0$-droplets containing the union of $Q(C)$ for all clusters C of Z. In particular, $\mathcal{D}$ does not depend on the order in which droplets are merged.

We say that a $\mathcal{T}_0$-droplet D is *covered* by a set Z of infections if the above algorithm for $Z \cap D$ outputs a $\mathcal{T}_0$-droplet containing D.

We make the convention that all $\mathcal{T}_0$-droplets have diameter at least C_4' and contain a site of $\mathbb{Z}^2$.

We next state some properties of the covering algorithm.

Lemma 8.A.3 (Lemma 4.6 of [74])**.** *Let D_1 and D_2 be $\mathcal{T}_0$-droplets such that $D_1 \cap D_2 \neq \varnothing$. Then*

$$\operatorname{diam}(D_1 \vee D_2) \leqslant \operatorname{diam}(D_1) + \operatorname{diam}(D_2).$$

This immediately implies the Aizenman-Lebowitz lemma (see e.g. [74, Lemma 4.8]).

Lemma 8.A.4 (Aizenman-Lebowitz)**.** *Let Z be a set of infections and D be a $\mathcal{T}_0$-droplet covered by Z. Then for all $C_1C_4 \leqslant k \leqslant \operatorname{diam}(D)$ there exists a $\mathcal{T}_0$-droplet D' covered by Z with $k \leqslant \operatorname{diam}(D') \leqslant 2k$.*

A further consequence of Lemma 8.A.3 is the following.

Lemma 8.A.5 (Lemma 7.3.14)**.** *Let Z be a set of infections and D be a $\mathcal{T}_0$-droplet covered by Z. Then D contains at least $\lceil \operatorname{diam}(D)/C_4^2 \rceil$ disjoint clusters of $Z \cap D$.*

We are now able to deduce the relevant bounds on covering following Lemma 7.3.15.

Lemma 8.A.6. *Let D be a $\mathcal{T}_0$-droplet with $d = \mathrm{diam}(D)$. Let $1 > \epsilon > 0$. Then we have*

$$\mu(D \text{ is covered}) \leqslant \begin{cases} q^{d\epsilon/(3C_4^2)} & \text{if } d \leqslant \frac{C_1}{q^{\alpha-\epsilon}} \\ e^{-C_1C_4d} & \text{if } \frac{1}{C_1 q^{\alpha-\epsilon}} \leqslant d \leqslant \frac{C_1}{e^{C_4^4}q^{\alpha}} \\ d^2 e^{-C_1/(C_5q^\alpha)} & \text{if } \frac{1}{C_1 e^{C_4^4}q^\alpha} \leqslant d. \end{cases} \tag{8.2}$$

Proof. Let Z be the (random) set of infections in D. By Lemma 8.A.5 we have that if D is covered, it contains at least $\lceil d/C_4^2 \rceil$ disjoint clusters of Z, each one having diameter at most C_3 and at least α sites. Thus, the union bound gives

$$\mu(D \text{ is covered}) \leqslant \binom{C_3^{2\alpha}d^2}{\lceil d/C_4^2\rceil} q^{\alpha\lceil d/C_4^2\rceil}.$$

For $d \leqslant C_4^2$ this gives $C_3^{2\alpha}d^2q^\alpha$, which concludes the proof. For $C_4^2 \leqslant d \leqslant C_1/(e^{C_4^4}q^\alpha)$ we use the inequality $\binom{n}{k} \leqslant (ne/k)^k$ to obtain the desired bounds. For the case $d \geqslant 1/(e^{C_4^4}q^\alpha)$ we use Lemma 8.A.4 to extract a smaller $\mathcal{T}_0$-droplet D' covered by Z (hence intersecting D) with $1/(2e^{C_4^4}q^\alpha) \leqslant \mathrm{diam}(D') \leqslant 1/(e^{C_4^4}q^\alpha)$. We then apply the second bound to D' and use the union bound to conclude. □

We would now like to use analogous bounds on the probability of $\mathcal{T}_0$-droplets being covered with initial condition $[Z]$ instead of Z. Unfortunately, we do not have access to the law of $[Z]$ when Z follows μ. Therefore, we rather bound the output of the covering algorithm for the closure using the original output. For that purpose, we define parallel notions of Γ', *modified clusters* and *modified covering*, by replacing C_2 by C_2' and C_4 by C_4'.

We then have the following key property, whose proof is identical to the one of Proposition 7.3.20, up to the relevant simplifications (we do not have rugged edges and there is no boundary).

Lemma 8.A.7 (Closure)**.** *Let $Z \subset \mathbb{Z}^2$ be a finite set and let $\mathcal{D}'$ be the collection of $\mathcal{T}_0$-droplets given by the modified covering algorithm with input $[Z]$. Let $\mathcal{D}$ be the output of the covering algorithm for Z. Then*

$$\forall D' \in \mathcal{D}' \, \exists D \in \mathcal{D}, D' \subset D.$$

8.A.2 Spanning

Let $\mathcal{T}$ be an arbitrary finite set of rational directions containing the origin in the interior of its convex envelope. We then generalise the notion of spanning from Definition 8.1.1.

Definition 8.A.8 (Spanning)**.** Let D be a $\mathcal{T}$-droplet. We say that D is *spanned* by $Z \subset \mathbb{Z}^2$ if there exists a set $C \subset [Z \cap D]$ connected in Γ' such that the smallest $\mathcal{T}$-droplet containing C is D.

We will need the following Aizenman-Lebowitz type lemma. Though this is a very classical result, some additional arguments are needed to prove it, because $\mathcal{T}$ is not composed of stable directions.

Lemma 8.A.9 (Aizenman-Lebowitz)**.** *Let $Z \subset \mathbb{Z}^2$ and D be a $\mathcal{T}$-droplet spanned by Z with $\operatorname{diam}(D) \geqslant C_1 C_2'$. Then for any $C_2' \leqslant k \leqslant \operatorname{diam}(D)$, there exists a $\mathcal{T}$-droplet D' spanned by $Z \cap D$ with $k \leqslant \operatorname{diam}(D') \leqslant C_1 k$.*

Proof. If $k \geqslant \operatorname{diam}(D)/C_1$ there is nothing to prove, as $D' = D$ is as desired. Assume $k \leqslant \operatorname{diam}(D)/C_1$. Let C be a connected component of $[Z \cap D]$ in Γ' with maximal diameter. By Definition 8.A.8 $\operatorname{diam}(C) \geqslant \operatorname{diam}(D)/\sqrt{C_1}$. By Observation 8.1.3 and [70, Lemma 6.18] (we use it although definitions slightly differ from [70], see Footnote 3) there exists $C' \subset C$ connected in Γ' such that $C' \subset [C' \cap Z \cap D]$ and $k \leqslant \operatorname{diam}(C') \leqslant \sqrt{C_1} k$. Denoting D' the smallest $\mathcal{T}$-droplet containing C', we are done. □

Observation 8.A.10. Let D be a $\mathcal{T}$-droplet spanned by $Z \subset \mathbb{Z}^2$ with $\operatorname{diam}(D) \geqslant C_4$. Then there exists a $\mathcal{T}_0$-droplet $\bar{D}$ covered by Z, intersecting D and such that $\operatorname{diam}(\bar{D}) = \Theta(\operatorname{diam}(D))$.

Proof. Let C be as in Definition 8.A.8. Notice that, since $\operatorname{diam}(D) \geqslant C_3$, we can find modified clusters for $[Z \cap D]$ whose union is a connected set in Γ' containing C. Then there is a $\mathcal{T}_0$-droplet in the output of the modified covering algorithm for $[Z \cap D]$ containing C. By Lemma 8.A.7 there is also a $\mathcal{T}_0$-droplet $\bar{D}$ in the output of the covering algorithm for $Z \cap D$ containing C, so that $\operatorname{diam}(\bar{D}) \geqslant \operatorname{diam}(C) = \Omega(\operatorname{diam}(D))$. But $\bar{D}$ is at most the smallest $\mathcal{T}_0$-droplet containing $\{x \in \mathbb{R}^2 : d(x, D) \leqslant C_4\}$, so $\operatorname{diam}(\bar{D}) = \Theta(\operatorname{diam}(D))$. Moreover, since $\bar{D}$ is in the output of the covering algorithm for $Z \cap D$, it is covered by Z and intersects D. □

We immediately deduce from this observation and Lemma 8.A.6 the desired bounds on spanning.

Corollary 8.A.11. *Let D be a $\mathcal{T}$-droplet with $d = \operatorname{diam}(D)$ and let $1 > \epsilon > 0$. Then*

$$\mu(D \text{ is spanned}) \leqslant \begin{cases} q^{d\epsilon/C_5} & \text{if } d \leqslant q^{-\alpha+\epsilon} \\ e^{-2C_4 d} & \text{if } q^{-\alpha+\epsilon} \leqslant d \leqslant \frac{C_1}{C_5 q^{\alpha}} \\ d^{O(1)} e^{-2/(C_5 q^{\alpha})} & \text{if } \frac{1}{C_5 q^{\alpha}} \leqslant d. \end{cases} \tag{8.3}$$

8.A.3 Boundary and spanning

We next turn to the treatment of an infinite infected boundary condition, following Chapter 7, which is applicable only for models with an infinite number of stable directions. Indeed, for a model with a finite number of stable directions a bounded set of infections next to the boundary can induce a set of supplementary infections and, thereby, a droplet of the size of the boundary, making similar algorithms useless. We therefore fix an update family $\mathcal{U}$ with an infinite number of stable directions and difficulty α, to which the treatment of Chapter 7 applies.

For the rest of this section let $\mathcal{S}_u = \{u^-, u^+, v_1, v_2\}$ be a set of 4 directions chosen as in Lemma 7.2.1[6] (we rename (u_1, u_2) from that chapter into (u^-, u^+) to avoid notational conflict) with $u = (u^- + u^+)/2$. The proof of Chapter 7 allows us to choose u^- and u^+ as close as we want, even depending on v_1 and v_2. We will choose them close enough for our results to hold. Let $\partial = \mathbb{H}_u$. For any set $Z \subset \mathbb{Z}^2$ we write $[Z]_\partial = [Z \cup \partial]\backslash\partial$. We will use the term *cluster* in the sense of Definition 7.3.3, extending Definition 8.A.1 (crumbs close to ∂ are considered as clusters instead and Z is replaced by $Z\backslash\partial$). We replace the notion of DYD from Chapter 7 by that of $\mathcal{S}_u$-droplet and the notion of CDYD becomes that of *cut* $\mathcal{S}_u$*-droplet*—a nonempty set of the form

$$\left(\overline{\mathbb{H}}_{u^-}(x) \cap \overline{\mathbb{H}}_{u^+}(y)\right)\backslash\partial \tag{8.4}$$

for some $x, y \in \mathbb{R}$, which is a geometric triangle. We further replace the use of the diameter by considering the size $|\cdot|$ from Definition 7.3.7. Namely for a cut $\mathcal{S}_u$-droplet D we denote $|D| = \mathrm{diam}(D)/C_1$, while for an $\mathcal{S}_u$-droplet D, $|D|$ denotes the length of its projection parallel to v_1. We then define correspondingly an extension of the covering algorithm as in Section 7.3.4 and a notion of covered (cut) $\mathcal{S}_u$-droplet. For the reader unfamiliar with Chapter 7, let us indicate that the change with respect to the covering algorithm of Definition 8.A.2 corresponds to replacing at each stage of the algorithm any $\mathcal{S}_u$-droplet D intersecting ∂ by the smallest cut $\mathcal{S}_u$-droplet containing $D\backslash\partial$. The properties of Section 7.3.5, analogous to Lemmas 8.A.3-8.A.5 and 8.A.7, remain valid for this setting. Furthermore, combining the proofs of Lemma 8.A.6 and Lemma 7.3.15 shows that the following holds.

Lemma 8.A.12. *Let D be a cut $\mathcal{S}_u$-droplet or an $\mathcal{S}_u$-droplet not intersecting ∂ with $d = |D|$. Let $1 > \epsilon > 0$. Then* (8.2) *holds.*

We similarly extend Definition 8.A.8 to the setting with boundary.

Definition 8.A.13 (Spanning with boundary)**.** We call *whole $\mathcal{S}_u$-droplet* any $\mathcal{S}_u$-droplet at distance at least C_3 from ∂ and, by abuse, we call collectively *$\mathcal{S}_u$-droplet* any cut or whole $\mathcal{S}_u$-droplet. We say that an $\mathcal{S}_u$-droplet D is

[6] It is not hard to see that in Lemma 7.2.1, with a finite number of exceptions, given any rational strongly stable direction $u \in S^1$ we can define $\mathcal{S}_u$ correspondingly.

spanned by $Z \subset \mathbb{Z}^2$ if there exists a set $C \subset [Z \cap D]_\partial$ connected in Γ' such that the smallest $\mathcal{S}_u$-droplet containing C is D.

We next recall several properties of the spanning algorithm following closely [70].

Definition 8.A.14 (Definition 6.15 of [70]). Let $Z = \{z_1, \dots, z_{k_0}\}$ be a finite set of infections. Set $\mathcal{Z}^0 = \{Z_1^0, \dots, Z_{k_0}^0\}$ with $Z_i^0 = \{z_i\}$. For each $t \geqslant 0$ do the following.

- If there exist Z_i^t and Z_j^t such that $[Z_i^t]_\partial \cup [Z_j^t]_\partial$ is connected in Γ', then set $\mathcal{Z}^{t+1} = (\mathcal{Z}^t \backslash \{Z_i^t, Z_j^t\}) \cup \{Z_i^t \cup Z_j^t\}$.
- Otherwise, define the *span of* Z by $\langle Z \rangle = \{D(Z^t), Z^t \in \mathcal{Z}^t\}$, where $D(Z')$ denotes the smallest $\mathcal{S}_u$-droplet containing Z', and terminate the algorithm.

Similarly, for any $A \subset \mathbb{R}^2$ we denote $\langle A \rangle = \langle A \cap \mathbb{Z}^2 \rangle$.

Observation 8.A.15 (Lemma 6.16 of [70]). Writing κ_i for the connected components of $[Z]_\partial$ in Γ', we have $\langle Z \rangle = \{D(\kappa_1), \dots, D(\kappa_k)\}$.

Observation 8.A.16 (Lemma 6.17 of [70]). A nonempty $\mathcal{S}_u$-droplet is spanned iff $D \in \langle D \cap Z \rangle$.

Lemma 8.A.17 (Lemma 6.21 of [70]). *Let Z be a finite set of at least two infections such that $[Z]_\partial$ is connected in Γ'. Then there exists a nontrivial partition $Z = Z_1 \sqcup Z_2$ such that $[Z_1]_\partial$, $[Z_2]_\partial$ and $[Z_1]_\partial \cup [Z_2]_\partial$ are connected in Γ'.*

The next lemma follows from the definition of size and Lemma 7.3.12.

Lemma 8.A.18. *For any $\mathcal{S}_u$-droplets D, D_1, D_2 with $|D_1| \geqslant C_3$ or $|D_2| \geqslant C_3$ such that $\langle D_1 \rangle = \{D_1\}$, $\langle D_2 \rangle = \{D_2\}$ and $\langle D_1 \cup D_2 \rangle = \{D\}$ we have $|D_1|/C_1 \leqslant |D| \leqslant |D_1| + |D_2| + O(C_2')$.*

This standardly implies (see e.g. [70, Lemma 6.18]) the following.

Lemma 8.A.19 (Aizenman-Lebowitz). *Let D be a spanned $\mathcal{S}_u$-droplet and $C_3 \leqslant k \leqslant |D|$. Then there exists a spanned $\mathcal{S}_u$-droplet $D' \subset D$ with $k \leqslant |D'| \leqslant 3k$.*

Similarly to Corollary 8.A.11 we obtain the following.

Corollary 8.A.20. *Let D be an $\mathcal{S}_u$-droplet with $d = |D| \geqslant 1/(C_5 q^\alpha)$. Then*

$$\mu(D \text{ is spanned}) \leqslant d^{O(1)} e^{-2/(C_5 q^\alpha)}.$$

Remark 8.A.21. Let us note that the results of this section remain valid if ∂ is replaced by any sufficiently regular boundary condition. Namely, if $u_\perp = u + \pi/2$ and f is a δ-Lipschitz function for $\delta < \tan((u^+ - u^-)/2)$, then we can use any ∂ with topological interior

$$\{x \in \mathbb{R}^2, \langle x, u\rangle < f(\langle x, u_\perp\rangle)\}$$

such that ∂, $\partial \cup D$ are stable for any cut $\mathcal{S}_u$-droplet.

Finally, one can also remove the boundary by considering infections sufficiently far from it to recover the setting of the previous section for the directions under consideration.

8.B Bound on crossing

For this appendix we place ourselves in the context of Section 8.1 (in particular, $\mathcal{T}$-droplets will be parallelograms). In sections 8.B.1 and 8.B.2 we show that crossings are unlikely in directions with respectively finite and infinite difficulty. Of course, though we treat u_1, the results are also valid for u_2.

8.B.1 Crossing in a direction with finite difficulty

One can use Corollary 8.A.11 to show that if u_1 has finite difficulty, a u_1-crossing without large droplets is extremely unlikely. To do that, we will use a concept of *partition* close to the one from [70, Definition 8.20].

Definition 8.B.1. Assume that $0 < \alpha(u_1) < \infty$. Let $R = R(a, b; c, d)$ be a parallelogram and $Z \subset R \cap \mathbb{Z}^2$. Set $m = \lfloor (c - a)/(C_1C_6)\rfloor \geqslant 1$ and

$$S_i = \mathbb{H}_{u_1}(-(c - iC_1C_6)) \cap \overline{\mathbb{H}}_{u_2}(-b) \cap \overline{\mathbb{H}}_{u_3}(c - (i - 1)C_1C_6) \cap \overline{\mathbb{H}}_{u_4}(d)$$

for $1 \leqslant i \leqslant m - 1$ and $S_m = R(a, b; c - (m - 1)C_1C_6, d)$. A u_1-partition of R for Z is a sequence $a_1, \dots, a_k$ of positive integers with $m = a_1 + \dots + a_k$ such that, setting $t_j = a_1 + \dots + a_j$, we have either

- $a_j = 1$ and S_{t_j} contains an $\alpha(u_1)$-cluster for Z (see Definition 8.A.1) or
- there exists a $\mathcal{T}$-droplet D spanned by $Z \cap \bigcup_{i=t_{j-1}+1}^{t_j} S_i$, with $C_1C_6a_j \geqslant \operatorname{diam}(D) \geqslant a_jC_6$.

The following lemma is close to [70, Lemma 8.21].

Lemma 8.B.2. *Let R be a parallelogram. If $0 < \alpha(u_1) < \infty$ and R is u_1-crossed then there exists a u_1-partition for $\eta \cap R$.*

Proof. For notational convenience we assume that $R = R(-a, 0; 0, d)$. In this proof, all clusters and crumbs are with respect to $\alpha(u_1)$. The proof is by induction on m.

Suppose that the property holds for any $m' \leqslant m-1$. If S_1 contains a cluster of $\eta \cap R$, we set $a_1 = 1$ and we are done, since $R(-a,0;-C_1C_6,d)$ is u_1-crossed. Let us assume S_1 contains no cluster of $\eta \cap R$. Then $S_1' = R(-C_1C_6 + C_1C_3, 0; 0, d)$ intersects no cluster of $\eta \cap R$, so if $\mathcal{K}$ is the set of connected components of $\eta \cap S_1'$ in Γ, each $\kappa \in \mathcal{K}$ is a crumb of $\eta \cap S_1'$. In particular, all elements of $[\kappa]$ and $[\kappa \cup \mathbb{H}_{u_1}]\backslash\mathbb{H}_{u_1}$ are at distance at most C_1 of κ (see Observation 7.3.16 and the proof of Corollary 7.3.17). As elements of $\mathcal{K}$ are at distance at least C_2 from one another, this means that $[\eta \cap S_1'] = \bigcup_{\kappa\in\mathcal{K}}[\kappa]$, and that $\bar{Z} = \bigcup_{\kappa\in\mathcal{K}'}[\kappa \cup \mathbb{H}_{u_1}]$ is closed, where $\mathcal{K}' = \{\kappa \in \mathcal{K} : d(\kappa, \mathbb{H}_{u_1}) \leqslant C_2\}$. Moreover, the diameter of a crumb is at most $\alpha(u_1)C_2$, so all elements of $\bar{Z}$ are at distance at most $(\alpha(u_1)+2)C_2$ of $\mathbb{H}_{u_1}$. Since R is u_1-crossed, this implies that there exists $z \in \bar{Z}$ and $w \in [(\eta \cap R)\backslash\bar{Z}]$ such that $d(z,w) \leqslant C_2'$. Then $d(w, \mathbb{H}_{u_1}) \leqslant (\alpha(u_1)+2)C_2 + C_2'$.

Let X be the connected component in Γ' of $[(\eta \cap R)\backslash\bar{Z}]$ containing w. If $X \subset [\eta \cap S_1']$, then $X \subset \bigcup_{\kappa\in\mathcal{K}}[\kappa]$, so $X \subset [\kappa]$ for some $\kappa \in \mathcal{K}$, since they are at distance more than C_2' from one another. Moreover, by Observation 8.1.3, $X = [((\eta \cap R)\backslash\bar{Z}) \cap X]$, so $X \not\subset \bar{Z}$, so $\kappa \notin \mathcal{K}'$. However, this contradicts the fact that $d(w,z) \leqslant C_2'$, as $d(\bar{Z}, [\kappa]) \geqslant C_2 - 2C_1$.

Therefore, $X \not\subset [\eta \cap S_1']$, so X intersects $R\backslash S_1'$. Let $a_1 = \max\{i \geqslant 1, X \cap S_i \neq \varnothing\}$ and D be the smallest $\mathcal{T}$-droplet containing X. Clearly, $\mathrm{diam}(D) \geqslant \mathrm{diam}(X) \geqslant a_1C_6$, since $d(w, \mathbb{H}_{u_1}) \leqslant C_3$. Furthermore, since $X = [((\eta \cap R)\backslash\bar{Z}) \cap X]$, D is spanned by $\eta \cap \bigcup_{i=1}^{a_1} S_i$. We then conclude by Lemma 8.A.9 and the induction hypothesis for $R(-a,0;-a_1C_1C_6,b)$. □

We next require a more sophisticated version of [70, Lemma 8.23].

Lemma 8.B.3. *Fix K in Definitions 8.1.1 and 8.1.5 by*

$$K = \begin{cases} 1/(C_5 q^{\alpha}) & \text{if } \mathcal{U} \text{ is balanced} \\ q^{-\alpha-1/4} & \text{if } \mathcal{U} \text{ is unbalanced.} \end{cases}$$

Assume that $0 < \alpha(u_1) < \infty$. Let $R = R(a,b;c,d)$ with $d - b \leqslant \frac{\log(1/q)}{C_6^3 q^{\alpha(u_1)}}$ and $1/q^{C_1} \geqslant c - a \geqslant 1/q$. Then

$$\mu(C_R^{u_1}) \leqslant \begin{cases} \exp\left(-(c-a)\exp\left(-2C_6^2(d-b)q^{\alpha(u_1)}\right)/C_6^2\right) & \text{if } \mathcal{U} \text{ is balanced} \\ \exp\left(-(c-a)q^{1/4}/C_6\right) & \text{if } \mathcal{U} \text{ is unbalanced.} \end{cases}$$

Proof. For notational convenience, we assume that $R = R(-a,0;0,d)$. If $C_R^{u_1}$ holds, there exists $\eta' \geqslant \eta$ such that R is u_1-crossed for η' and there is no spanned critical $\mathcal{T}$-droplet for $\eta' \cap R$. By Lemma 8.B.2, there exists a u_1-partition for $\eta' \cap R$ and, by Lemma 8.A.9, all corresponding spanned $\mathcal{T}$-droplets have diameter at most K/C_1. We notice that any infected site or spanned droplet for η' is still an infected site or spanned droplet for η.

We first assume that $\mathcal{U}$ is balanced. Given a partition $\mathcal{P}$ we define its numbers and total sizes of big/small/cluster parts by

$$\mathcal{B} = \{j : 1/\sqrt{q} < a_j \leqslant 1/(C_5C_6q^\alpha)\} \qquad b = |\mathcal{B}| \qquad B = \sum_{j\in\mathcal{B}} a_j$$
$$\mathcal{S} = \{j : 1 < a_j \leqslant 1/\sqrt{q}\} \qquad s = |\mathcal{S}| \qquad S = \sum_{j\in\mathcal{S}} a_j$$
$$\mathcal{C} = \{j : a_j = 1\} \qquad c = |\mathcal{C}|.$$

We denote by $\mathcal{P}(b, s, c, B, S)$ the set of partitions $\mathcal{P}$ with the corresponding numbers and total sizes of parts.

Then, using Corollary 8.A.11, we get that the probability of a given $\mathcal{P}$ occurring is at most

$$\begin{aligned}\Pi(\mathcal{P}) &= \prod_{j\in\mathcal{C}}(1-(1-q^{\alpha(u_1)})^{C_6^2d})\prod_{j\in\mathcal{S}}q^{a_j\sqrt{C_6}}\prod_{j\in\mathcal{B}}e^{-C_3C_6a_j}\\ &= (1-(1-q^{\alpha(u_1)})^{C_6^2d})^c q^{S\sqrt{C_6}}e^{-C_3C_6B}\end{aligned}$$

by the union bound on all possible droplets and their positions, recalling that $d = q^{-O(1)}$. Indeed, the probability that there is no set of $\alpha(u_1)$ zeroes connected in Γ' in a given S_i is the probability that for any possible such set C, $\eta_C \neq 0$, which, by the Harris inequality, is bigger than the product of this probability for each set C.

Assuming for simplicity that $1/\sqrt{q}$ and $1/(C_5C_6q^\alpha)$ are integers, we can count $\mathcal{P}(b, s, c, B, S)$ in the following way (the first binomial coefficient corresponds to the decomposition of $\mathcal{B}$ into ordered parts, the second one to the decomposition of $\mathcal{S}$, and the last two to the ordering of the parts of $\mathcal{B}$, $\mathcal{S}$, $\mathcal{C}$):

$$\begin{aligned}|\mathcal{P}(b,s,c,B,S)| &\leqslant \binom{B-b/\sqrt{q}-1}{b-1}\binom{S-s-1}{s-1}\binom{b+s+c}{b}\binom{s+c}{s}\\ &\leqslant 2^{B+S}(b+s+c)^b(s+c)^s \leqslant e^{B+S}q^{-C_1s},\end{aligned}$$

recalling that $C_6(B + S + c) < a \leqslant 1/q^{C_1}$. Therefore, denoting by $m = \lfloor a/(C_1C_6)\rfloor = B + S + c$ the total number of strips, we have

$$\begin{aligned}&\sum_{B,S,b,s}\ \sum_{\mathcal{P}\in\mathcal{P}(b,s,m-B-S,B,S)}\Pi(\mathcal{P})\\ &\qquad\leqslant m^4\max_{B,S}\left(1-\left(1-q^{\alpha(u_1)}\right)^{C_6^2d}\right)^{m-B-S}q^{S\sqrt{C_6}/2}e^{-C_3C_6B/2}\\ &\qquad\leqslant m^4\max_{0\leqslant c\leqslant m}e^{-c\exp\left(-2C_6^2dq^{\alpha(u_1)}\right)}e^{-C_2C_6(m-c)}\\ &\qquad\leqslant \exp\left(-\frac{m}{2}\exp\left(-2C_6^2dq^{\alpha(u_1)}\right)\right),\end{aligned}$$

which concludes the proof in the balanced case, recalling the hypotheses of the lemma.

We next consider $\mathcal{U}$ to be unbalanced. Notice that, since $K = q^{-\alpha-1/4}$, there may be droplets with diameter larger than $1/(C_5q^\alpha)$. Therefore, we further set

$$\mathcal{H} = \left\{j : 1/(C_5C_6q^\alpha) < a_j \leqslant 1/(C_6q^{\alpha+1/4})\right\} \qquad h = |\mathcal{H}| \qquad H = \sum_{j\in\mathcal{H}} a_j.$$

Then Corollary 8.A.11 gives that the probability of a given $\mathcal{P}$ occurring is at most

$$\Pi(\mathcal{P}) \times \left(q^{O(1)}e^{-2/(C_5q^\alpha)}\right)^h \leqslant \Pi(\mathcal{P}) \times \exp\left(-Hq^{1/4}C_6/C_5\right).$$

We further easily check that

$$\binom{H - h/(C_5C_6q^\alpha) - 1}{h-1} \leqslant e^{H\sqrt{q}} \qquad \binom{h+b+s+c}{h} \leqslant e^{H\sqrt{q}},$$

so, as above the probability of any $\mathcal{P}$ occurring is at most

$$m^6 \exp\left(-m.\min\left(C_6q^{1/4}/(2C_5), \exp\left(-2C_6^2dq^{\alpha(u_1)}\right)\right)\right) \leqslant \exp\left(-\frac{C_6mq^{1/4}}{3C_5}\right),$$

which concludes the proof. □

8.B.2 Crossing in a direction with infinite difficulty

If $\mathcal{U}$ has an infinite number of stable directions, we need to treat an infected boundary condition. This is essential, as we will work in exponentially large regions, for which the bounds from the previous section cannot be applied.

We place ourselves in the setting of Section 8.2.1. We will write (cut, whole or either) droplet for (cut, whole or either) $\mathcal{S}_{u_1}$-droplet in the sense of Definition 8.A.13, with u_1^- and u_1^+ sufficiently close to u_1. These should not be confused with $\mathcal{T}$-droplets, which are called parallelograms to avoid any confusion.

We will seek to apply Corollary 8.A.20 rather than 8.A.11 to prove the following.

Lemma 8.B.4. *Fix*

$$K = \begin{cases} q^{-\alpha} & \text{if } \mathcal{U} \text{ is balanced} \\ q^{-\alpha-1/4} & \text{if } \mathcal{U} \text{ is unbalanced} \end{cases}$$

for Definitions 8.1.1 and 8.1.5. Let $R = R(a,b;c,d)$ with $C_1 \leqslant d-b \leqslant \exp(q^{-3\alpha})$ and $c-a \geqslant q^{-4\alpha}$. Then

$$\mu(C_R^{u_1}) \leqslant \exp\left(-q^{-3\alpha}\right).$$

Our strategy is as follows. Instead of considering u_1-partitions, we directly retrace the spanning algorithm to obtain a hierarchy of droplets reaching a cut droplet of size roughly $c-a$. We reassure the reader familiar with [225] that our hierarchies will be very simple and imprecise, as the a priori hypothesis that there are no critical parallelograms removes the metastability (it is no longer easy for large droplets to grow) together with the need of fine tuning. Namely, their seeds will be of size roughly K which will also be the increment of the size of unary vertices (the reader unfamiliar with hierarchies is invited to consult the definitions below). The lack of critical parallelograms entails that all droplets in the hierarchy are cut (so they are simply very flat triangles). The bound on the probability of seeds being spanned is provided by Corollary 8.A.20 and entropy is easily subdominant, so we can focus on the probability that the infections around a cut droplet are such that if that droplet is infected, the infection can expand to fill a slightly larger cut droplet. However, this would imply that there is a (smaller scale) $u_1^{\pm}$-crossing from the side of the smaller one to side of the larger one (see Figure 8.2b). The probability of this event is again bounded directly by Corollary 8.A.20, taking into account Remark 8.A.21.

Let us begin by introducing our hierarchies following Holroyd [225]. Let $T = q^{-\alpha-1/4}$. Fix a droplet D. A *hierarchy* $\mathcal{H}$ for D is a rooted unary-binary tree with each vertex x labelled by a droplet $D_x \subset D$, so that the label of the root is D. We denote by $N(x)$ the set of children of $x \in V(\mathcal{H})$, so that $|N(x)| \in \{0,1,2\}$ for all x. The leaves are called *seeds* and the binary vertices are called *splitters*. We alert the reader that in reality there will only be cut droplets in our hierarchies, but for technical reasons we define them in general. A hierarchy is defined to satisfy the following conditions.

- If $y \in N(x)$, then $D_y \subset D_x$.
- If D_x is a whole droplet, then x is a seed and $T/3 \leqslant |D_x|$.
- If D_x is a cut droplet, then $T/3 \leqslant |D_x| \leqslant T$ if and only if x is a seed.
- If $N(x) = \{y\}$ and $|N(y)| = 1$, then $T < |D_x| - |D_y| \leqslant 2T$.
- If $N(x) = \{y\}$, then $|D_x| - |D_y| \leqslant 2T$.
- If $N(x) = \{y, z\}$, then $|D_x| - |D_y| > T$ and $\langle D_y \cup D_z \rangle = \{D_x\}$.

We set

$$S(\mathcal{H}) = \{x \in V(\mathcal{H}) : |N(x)| = 0\}$$
$$N(\mathcal{H}) = \{(x,y) \in (V(\mathcal{H}))^2 : N(x) = \{y\}, |N(y)| = 1\}$$

and remark that $|S(\mathcal{H})|-1$ is the number of splitters. We say that a hierarchy $\mathcal{H}$ *occurs* if the following events occur disjointly (are witnessed by disjoint sets of infections, see [357]).

- For every seed $x \in S(\mathcal{H})$ we have that D_x is spanned.
- For every $x \in V(\mathcal{H})$ such that $N(x) = \{y\}$ we have $D_x \in \langle D_y \cup (\eta \cap D_x)\rangle$.

Lemma 8.B.5. *If D is a spanned droplet with $|D| \geqslant T/3$, then some hierarchy for D occurs.*

Proof. The proof is very similar e.g. to [70, Lemma 8.7]. Assuming that D_0 is a spanned droplet with $|D_0| \geqslant T/3$, we construct an occurring hierarchy by induction on D_0 with respect to inclusion. If $|D_0| \leqslant T$ or D_0 is a whole droplet, the hierarchy with only vertex labelled by D_0 is as desired.

Assume that D_0 is a cut droplet and $|D_0| > T$. Let Z be a connected component for Γ' of $[D_0 \cap \eta]_\partial$ such that the smallest droplet containing Z is D_0, and let $Z_0 = Z \cap \eta$. We then have $[Z_0]_\partial = Z$. By Lemma 8.A.17 there exist sequences $Z_1, \dots, Z_m$ and $Z'_1, \dots, Z'_m$ of subsets of Z_0 and $D_1, \dots, D_m$ and $D'_1, \dots, D'_m$ such that the following conditions hold for all $0 < i \leqslant m$.

- $Z_{i-1} = Z_i \sqcup Z'_i$, $[Z_i]_\partial$, $[Z'_i]_\partial$ and $[Z_i]_\partial \cup [Z'_i]_\partial$ are connected in Γ'.
- $D_i = D([Z_i]_\partial)$ and $D'_i = D([Z'_i]_\partial)$.
- $\langle D_i \cup D'_i\rangle = \{D_{i-1}\}$ and $|D_i| \geqslant |D'_i|$
- $m \geqslant 1$ is the minimal index such that one of the following holds:
  (1) $|D_0| - |D_m| > T$;
  (2) D_m is a whole droplet;
  (3) $|D_m| \leqslant T$.

If (1) does not hold, then we attach a seed labelled by D_m to the root and we are done, as $|D_m| \geqslant |D_{m-1}|/3 > T/3$ by Lemma 8.A.18 and minimality of m. Indeed, D_m being spanned is witnessed by Z_m, while $D_0 \in \langle D_m \cup (\eta \cap D_0)\rangle$ is witnessed by $Z_0 \backslash Z_m$.

Assume that (1) holds. Then we consider two cases. If $T < |D_0| - |D_m| \leqslant 2T$, we attach a hierarchy for D_m (occurring for Z_m) to the root D_0 and we are done using Lemma 8.A.18 to get that $|D_m| > T/3$ as above. Otherwise we attach a splitter labelled by D_{m-1} to the root D_0 (if $m = 1$, then D_0 is the splitter) and hierarchies for D_m and D'_m to that splitter. Then we are done, recalling Lemma 8.A.18, to get that $|D_m| \geqslant |D'_m| \geqslant |D_{m-1}| - |D_m| - O(C'_2) \geqslant T - O(C'_2)$. □

In order to bound the probability that a hierarchy occurs, we will need the following result.

Lemma 8.B.6. *Let $D_1 \subset D_2$ be two cut droplets for $\partial = \mathbb{H}_{u_1}$ such that $T < |D_2| - |D_1| \leqslant 2T$ and $|D_2| \leqslant q^{-4\alpha}$. Then*

$$\mu(D_2 \in \langle D_1 \cup (\eta \cap D_2)\rangle) \leqslant e^{-q^{-\alpha}/C_5}.$$

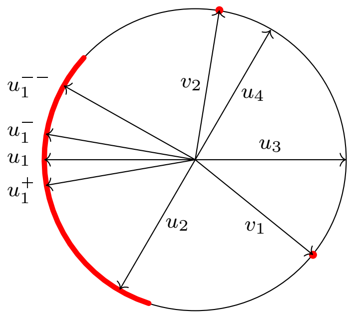

(a) We recall all directions in the context of Lemma 8.B.6, stable ones being thickened. The only relevant ones for a u_1-crossing are u_1, u_1^+, u_1^- and u_1^{--} (or similarly u_1^{++}). Indeed, v_1 and v_2 are only used for whole droplets and to bound the probability of a u_1-crossing we consider cut droplets with directions u_1^+, u_1^-. Note that u_1 and u_2 should be very close and $u_1^\pm$ even closer, but this is avoided here for visibility.

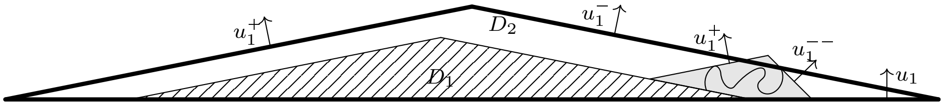

(b) Growth of the infection in the hatched cut droplet D_1 to the thickened one, D_2, requires a path of infections such as the one on the right, inducing the shaded spanned cut $\mathcal{S}_{u_1^-}$-droplet.

Figure 8.2 – Illustration of the proof of Lemma 8.B.6 bounding the probability that the infections around a cut droplet, D_1, allow an infection filling D_1 to grow and fill the slightly larger cut droplet, D_2.

Proof. The proof is illustrated in Figure 8.2. Let us denote $D_i = (\overline{\mathbb{H}}_{u_1^+}(x_i) \cap \overline{\mathbb{H}}_{u_1^-}(y_i))\backslash\mathbb{H}_{u_1}$ for $i \in \{1,2\}$. Define the strips $X = \overline{\mathbb{H}}_{u_1^+}(x_2)\backslash\overline{\mathbb{H}}_{u_1^+}(x_1)$ and $Y = \overline{\mathbb{H}}_{u_1^-}(y_2)\backslash\overline{\mathbb{H}}_{u_1^-}(y_1)$ and assume without loss of generality that $y_2 - y_1 = \Omega(T)$. Assume that $D_2 \in \langle D_1 \cup (\eta \cap D_2)\rangle$ occurs. Setting $\eta' = \eta \cap Y \cap D_2$, this implies $D_2 \in \langle (D_2\backslash Y) \cup \eta'\rangle$. We consider two cases.

Assume that $D_2 \in \langle \eta' \rangle$. By Corollary 8.A.20 the probability of this event is at most $q^{-O(1)}e^{-2/(C_5 q^\alpha)}$.

Assume that, on the contrary, $D_2 \notin \langle \eta' \rangle$ and set $\partial' = \mathbb{H}_{u_1} \cup \overline{\mathbb{H}}_{u_1^-}(y_1)$. Then by Observation 8.1.3 there exists a set $C \subset [D_2 \cap \eta]_{\partial'}$ connected in Γ' such that $d(C, \mathbb{H}_{u_1^-}(y_1)) \leqslant C_2'$ and $C \not\subset \mathbb{H}_{u_1^-}(y_2)$. By definition this implies the existence of a cut $\mathcal{S}_{u_1^-}$-droplet spanned by $D_2 \cap \eta$ with boundary ∂', where the two directions of cut droplets in $\mathcal{S}_{u_1^-}$ are u_1^+ and $u_1^{--} = 2u_1^- - u_1^+$ (recall Remark 8.A.21). Hence, by the union bound over all possible such droplets and Corollary 8.A.20 we obtain the desired result. □

We are now ready to assemble the proof of Lemma 8.B.4 as outlined at the beginning of the section.

Proof of Lemma 8.B.4. Assume that $C_R^{u_1}$ occurs and let $\eta' \geqslant \eta$ be as in Definition 8.1.5. Then there exists a spanned cut droplet for $\eta' \cap R$ with boundary $\mathbb{H}_{u_1}$ of diameter at least $c - a$. By Lemma 8.A.19 this implies the existence of a droplet D spanned for $\eta' \cap R$ with $q^{-4\alpha}/C_1 \leqslant |D| \leqslant 3q^{-4\alpha}/C_1$. We set $Z = \eta' \cap R \cap D$.

Let us assume for a contradiction that there exists a whole droplet of size at least $q^{-\alpha-1/4}/3$ spanned for Z. It is easy to check that there exists a parallelogram of diameter at least $q^{-\alpha-1/4}/C_1$ spanned by Z (consider a connected component satisfying Definition 8.A.13, take the smallest parallelogram containing it and use Observation 8.1.3). By Lemma 8.1.4 this contradicts the absence of spanned critical parallelograms for $\eta' \cap R$.

Therefore, D is a cut droplet and by Lemma 8.B.5 there exists a hierarchy for D occurring for the zero set Z, whose labels are all cut droplets. Let $\mathscr{H}(D)$ denote the set of such hierarchies. Now, by the BK inequality [357], for any hierarchy $\mathcal{H}$ we have the following analogue of [225, Equation (37)]:

$$\mu(\mathcal{H} \text{ occurs}) \leqslant \prod_{x\in S(\mathcal{H})} \mu(D_x \text{ is spanned}) \prod_{(x,y)\in N(\mathcal{H})} \mu(D_x \in \langle D_y \cup (\eta \cap D_x)\rangle).$$

Thanks to Corollary 8.A.20 and Lemma 8.B.6, we deduce

$$\mu(C_R^{u_1}) \leqslant \sum_D \sum_{\mathcal{H}\in\mathscr{H}(D)} \exp\left(-q^{-\alpha}(|S(\mathcal{H})| + |N(\mathcal{H})|)/C_5\right).$$

The number of choices for D is $O((d-b)q^{-4\alpha})$. We separate the sum over hierarchies according to their number of vertices $v(\mathcal{H}) = \Theta(|S(\mathcal{H})|+|N(\mathcal{H})|)$. By Lemma 8.A.18 we have that $v(\mathcal{H}) = \Omega(|D|/T) = \Omega(q^{-3\alpha+1/4}/C_1)$. Finally, the number of hierarchies for a given cut droplet D with v vertices is at most $q^{-O(v)}$. Combining these bounds we have

$$\mu(C_R^{u_1}) \leqslant (d-b) \sum_{v=\Omega(q^{-3\alpha+1/4}/C_1)} \exp\left(-q^{-\alpha}\Omega(v)/C_5\right),$$

which concludes the proof. □

Chapter 9

Complexity of two-dimensional bootstrap percolation difficulty: algorithm and NP-hardness

This chapter is based on joint work with Tamás Mezei [219].

For convenience we introduce the notation

$$\bar{Z} := [Z \cup \mathbb{H}_u] \backslash \mathbb{H}_u, \tag{9.1}$$

whenever the direction u is clear from the context, and $l_u = \overline{\mathbb{H}}_u \backslash \mathbb{H}_u = \{x \in \mathbb{Z}^2 : \langle x, u\rangle = 0\}$.

9.1 Results

So far it has not been investigated how one could determine the difficulty α in practice (recall Section 1.6.1), mainly owing to the simple definition and to the fact that for simple models such as the ones in Figure 1.2 this is straightforward. In this chapter we consider α from a computational perspective.

Throughout the chapter, we assume that $\mathcal{U}$ is described as a family of sets of pairs of integer coordinates represented in binary. Therefore the size of the input is proportional to

$$\|\mathcal{U}\| := \log D \cdot \sum_{U \in \mathcal{U}} |U|, \tag{9.2}$$

where D is the "diameter" of $\mathcal{U}$:

$$D = 2 \cdot \max\left\{ \|x\|_\infty \ : \ x \in \bigcup_{U \in \mathcal{U}} U \right\}. \tag{9.3}$$

A further justification of the need to take D into account in $\|\mathcal{U}\|$ is provided in Section 9.A showing that the difficulty α is not bounded in terms of $\sum_{U\in\mathcal{U}} |U|$

only. Our first result is that α is computable. We prove this by giving an explicit algorithm and bounding its complexity.

Theorem 9.1.1. *There exists an algorithm which, given a critical bootstrap percolation update family $\mathcal{U}$, computes its difficulty α.*[1]

Remark 9.1.2. In fact, it is not hard to check that our algorithm runs in time at most

$$|\mathcal{U}|^2 \cdot 2^{D^2(1+o(1))} = \exp(O(D^2)),$$

which is in the worst case at most doubly exponential in $\|\mathcal{U}\|$. This bound is as sharp as a bound in terms of D only can be. Indeed, $|\mathcal{U}| = e^{O(D^2)}$ and $|\mathcal{U}|$ can be as large as 2^{D^2}.

Explicit bounds analogous to the ones derived in the proof of Theorem 9.1.1 are the only missing ingredient causing the constants appearing in the main results of [70] and Chapter 7 to be implicit (cf [70, Lemma 6.5] and Observation 7.3.16).

Moreover, a corresponding uncomputability result in higher dimensions based on supercritical models in two dimensions has been announced by Balister, Bollobás, Morris and Smith [27] prior to our work. As that could lead one to expect, Theorem 9.1.1 is not at all automatic.

On a high level, the main idea behind our proof is that if a half-plane $\mathbb{H}_u$ is infected, the process restricted to the line l_u is a 1-dimensional bootstrap percolation process. Owing to the bounded range of the rules and translation invariance, the final state of this process is either periodic with bounded period or finite, which two possibilities can be distinguished in a correspondingly bounded time.

On the other hand, we also prove the following negative result.

Theorem 9.1.3. *The problem of computing the difficulty α of a critical bootstrap percolation update family $\mathcal{U}$ is NP-hard.*

This result is proved by a fairly technical reduction to the Set Cover decision problem in Section 9.3. Besides the result of [27], another reason to expect that the problem of determining α is hard in a sense made clear in Theorem 9.1.3 is a recent parallel notion of difficulties adapted to subcritical models termed “critical densities.” Those are introduced by the author in Chapter 12 and they are clearly far too complicated for one to expect to be able to compute them. From this point of view the result of Theorem 9.1.3 is not unexpected.

[1] This result is proved independently by Balister, Bollobás, Morris and Smith [27].

9.2 Decidability: proof of Theorem 9.1.1

In this section we provide an algorithm to compute the difficulty of a critical model. Let us stress that it is not optimized and is only meant to prove Theorem 9.1.1.

Proof of Theorem 9.1.1. Fix an update family $\mathcal{U}$. To start, let us see how to determine the set of stable directions in time polynomial in the size of the input $\|\mathcal{U}\|$. Indeed, for each site x in each rule U we determine its polar coordinates $(r_x, \theta_x) = (\|x\|_2, x/\|x\|_2) \in \mathbb{R}_+ \times S^1$. On the practical side, r_x can be represented as the square root of an integer bounded by D^2 and θ_x can be encoded by its tangent, which is rational with numerator and denominator bounded by D, and one boolean indicating whether $\theta_x \in (-\pi/2, \pi/2)$. Then for each rule U we take an arbitrary $x_0 \in U$ and compute $\theta_x - \theta_{x_0}$ for all $x \in U$ (its tangent is still rational and its numerator and denominator are bounded by D^2). We determine the largest and smallest such values, δ_+, δ_-, considering differences in $(-\pi, \pi]$. Finally, the unstable interval of U is $(\theta_{x_0} + \delta_+ + \pi/2, \theta_{x_0} - \delta_- + 3\pi/2) \subset S^1$ (which is empty if $\delta_+ - \delta_- \geqslant \pi$). The set of unstable directions is then the union of these intervals for all $U \in \mathcal{U}$. In particular, the isolated stable directions and, more generally, the endpoints of the intervals of stable directions for $\mathcal{U}$ are among the endpoints of the intervals for different U, so there are at most $2|\mathcal{U}|$ of them. In order to determine this union in practice it suffices to check for each of these endpoints whether it is stable (not contained in any of the unstable intervals for other $U \in \mathcal{U}$) and keep the information whether it was a left or right endpoint of the associated interval. Hence, the preliminary step of determining the (isolated) stable directions is completed in polynomial time in $\|\mathcal{U}\|$. It is also not hard to verify for each of the $|\mathcal{U}|$ right-endpoints whether there exists a stable direction in the half-circle starting there and whether there are finitely many of them (i.e. all are isolated), which allows one to decide if $\mathcal{U}$ is supercritical, critical or subcritical in polynomial time.

Assuming that $\mathcal{U}$ is determined to be critical, we can use Definition 1.6.1 to compute the difficulty, α, once we know all $\alpha(u) \in \{1, 2, \dots\}$ for isolated stable directions. Indeed, for each of the open semi-circles with one endpoint among those considered above, we only need to calculate the maximum of $\alpha(u)$ for isolated stable directions u (if there are any non-isolated directions, we do not need to consider the semi-circle). As this can also be done in time polynomial in $\|\mathcal{U}\|$, we will now fix an isolated stable direction u and provide an algorithm for determining $\alpha(u)$.

We shall assume that D is sufficiently large throughout the proof. Indeed, given D, all $U \in \mathcal{U}$ are distinct subsets of $\{-D/2, \dots, D/2\}^2$, so there are at most $2^{2^{(D+1)^2}}$ possible $\mathcal{U}$ and $|\mathcal{U}| \leqslant 2^{(D+1)^2}$. Therefore, the algorithm's asymptotic complexity is only determined by families with large values of D, as one can directly list the difficulties for isolated stable directions with "small" values of D in constant time.

Recall the notation $\bar{Z}$ from (9.1), which we shall use without specifying u, as it will be clear from the context. In order to determine $\alpha(u)$ we will use the following lemmas to bound the size of the set Z in Definition 1.6.1. The first of these is a one-dimensional result which we shall reduce the problem to.

Lemma 9.2.1. *Let $\mathcal{U}$ be an update family, let $u \in S^1$ be an isolated stable direction and let A be a finite subset of l_u. Then the set $\bar{A}$ is either infinite or its maximal distance from A is at most $D^3 \cdot 2^D$.*

Proof. Observe that by stability of u we have $\bar{A} \subset l_u$. Then the dynamics started from $\mathbb{H}_u \cup A$ can be viewed as a dynamics on l_u only. Note that l_u consists of integer sites on a line, so it is naturally identified with $\mathbb{Z}$ by the composition of a homothety and a rotation. Furthermore, we know that u is an isolated stable direction and, thereby, $l_{u+\pi/2}$ (which is simply a rotation of l_u) contains a site x in some $U \in \mathcal{U}$ with $\|x\|_\infty \leqslant D/2$ by (9.3). Hence, the homothety ratio is between $1/D$ and 1.

Notice that the dynamics restricted to l_u is simply a 1-dimensional bootstrap percolation process, where each rule $U \in \mathcal{U}$ is replaced by $U \cap l_u$ if $U \subset (\mathbb{H}_u \cup l_u)$ and removed otherwise. It therefore suffices to prove the following claim, which concludes the proof. □

Claim. For a one-dimensional bootstrap percolation family and a finite set $A \subset \mathbb{Z}$, we have that $\bar{A}$ is either infinite or its maximal distance from A is at most $D^2 \cdot 2^D$.

Proof. Denote $A = \{a_1, \dots, a_n\}$ with $a_1 < \cdots < a_n$. Let us denote by P the property that the following three conditions hold:

- $|[A]| < \infty$, $d(s, A) \leqslant D \cdot 2^{D+1}$ for all $s \in [A]$,
- $\max[A] - a_n \leqslant D \cdot 2^{D+1} - D$,
- $a_1 - \min[A] \leqslant D \cdot 2^{D+1} - D$.

Let A be minimal with respect to inclusion violating P. We next prove that $|[A]| = \infty$.

Base. Assume that $|A| = 1$, without loss of generality $A = \{0\}$. If $[A] = A$, we have nothing to prove, as P clearly holds. Otherwise, assume that $x \in \mathbb{Z}$ becomes infected on the first step. Then, since $\{0\}$ is the only infected site initially, $\{x\}$ is a rule in the update family. However, that entails that $k.x$ becomes infected on the k-th iteration at the latest and, in particular, $[A]$ is infinite.

Step. Assume that $|A| > 1$. Assume for a contradiction that there exists $0 < i < n$ and $b \in [A]$ such that $a_{i+1} > b > a_i$ and $\min(b - a_i, a_{i+1} - b) > D \cdot 2^{D+1}$. Then, by minimality of A, both $A' = \{a_1, \dots, a_i\}$ and $A'' = A \backslash A'$ satisfy P. Therefore,

$$\min[A''] - \max[A'] > D \cdot 2^{D+2} - 2(D \cdot 2^{D+1} - D) > D,$$

so that $[A] = [A'] \cup [A'']$, which contradicts the existence of $b \in [A]$. Indeed, there is no site in $\mathbb{Z}$ such that a rule translated by it intersects both $[A']$ and $[A'']$ and by definition of the closure those do not evolve under the dynamics.

Assume next that $\max([A]) > a_n + D \cdot 2^{D+1} - D$ (the corresponding case for $\min([A])$ is treated identically). Then, by the pigeon-hole principle, there exist $b, c \in \mathbb{Z}$ with $a_n + D < b < c - D < \max([A]) - 2D$ such that

$$\varnothing \neq [A] \cap [b, b + D - 1] = ([A] \cap [c, c + D - 1]) - (c - b)$$

(since no infection can cross a region of size D not intersecting $[A]$ to reach $\max([A])$). Therefore, $[A] \cap [b, b + D - 1]$ infects a translate of itself, since the dynamics to the right of $b + D$ is not affected by infections to the left of b, once we fix the state of $b, \dots, b + D - 1$. Similarly to the case $|A| = 1$, this is a contradiction with $|[A]| < \infty$, which concludes the proof. □

The next lemma is an application of the covering algorithm of [74]. For the sake of completeness, we will include a sketch of it in the proof.

Lemma 9.2.2. *Let $\mathcal{U}$ be a critical update family and u be an isolated stable direction. Let $Z \subset \mathbb{H}_{u+\pi}$ be a set of size at most D. Then for every $z \in [Z]$ we have $\langle z, u \rangle \geqslant -O(D^4)$.*

Proof. First, we prove the following claim.

Claim. There exists a set $\mathcal{T} \supset \{u\}$ of three or four stable directions containing the origin in their convex envelope (if viewed as a subset of $\mathbb{R}^2$) such that for each $v \in \mathcal{T}$ there exists $x \in \mathbb{Z}^2 \cap v\mathbb{R}$ such that $\|x\|_\infty \leqslant D/2$ and such that for every $v, w \in \mathcal{T}$ we have $|v - w + \pi| \geqslant 2/D^2$.

Proof. First assume that $u + \pi$ is unstable. Let $\mathcal{T}$ consist of u and the stable directions, $u + \pi + \delta_+$ and $u + \pi - \delta_-$ ($\delta_\pm \in (0, \pi]$), closest to $u + \pi$ in both semi-circles with endpoint $u + \pi$ (these exist as the set of stable directions is closed). Furthermore, recalling that $\mathcal{U}$ is not supercritical, there is no semi-circle of unstable directions, so $\delta_+ + \delta_- < \pi$. This implies that indeed $\mathcal{T}$ contains 0 in its convex envelope.

Assume that, on the contrary, $u + \pi$ is stable. Consider the semi-circle $(u, u + \pi) \subset S^1$. In it there exists a stable direction (since $\mathcal{U}$ is not supercritical). If there are no unstable directions, we pick $u_- = u + \pi/2$, otherwise, we set u_- to be an isolated or semi-isolated stable direction in that semi-circle. We define u_+ similarly in the opposite semi-circle. We set $\mathcal{T} = \{u, u + \pi, u_-, u_+\}$. It is clear that 0 is in the convex envelope of $\mathcal{T}$.

In both cases $\mathcal{T}$ consists of directions which are either isolated, semi-isolated or a rotation by $\pi/2$ of such a direction. Therefore, as in the proof of Lemma 9.2.1, there exists a site x as in the statement of the claim.

Finally, let us bound the difference between two directions $v \neq w$ such that there exist $x \in \mathbb{Z}^2 \cap v\mathbb{R}$ and $y \in \mathbb{Z}^2 \cap w\mathbb{R}$ with $\max(\|x\|_\infty, \|y\|_\infty) \leqslant D/2$. Indeed, $\det(x, y) \in \mathbb{Z}\backslash\{0\}$, so

$$|\sin(v - w)| = \frac{|\det(x, y)|}{\|x\|_2\|y\|_2} \geqslant \frac{2}{D^2}$$

and therefore $|v - w| \geqslant 2/D^2$. □

We fix a set $\mathcal{T}$ as in the claim. We call a $\mathcal{T}$-droplet a polygon with sides perpendicular to the directions in $\mathcal{T}$. Since $\mathcal{T}$ contains 0 in its convex envelope there exist $\mathcal{T}$-droplets. Since the difference between consecutive directions in $\mathcal{T}$ are at most $\pi - 2/D^2$, we can find a $\mathcal{T}$-droplet P with diameter $O(D^3)$ containing $[-D/2, D/2]^2 \supseteq \bigcup_{U \in \mathcal{U}} U$ (e.g. a $\mathcal{T}$-droplet circumscribed around a circle with D).

We can then directly apply the covering algorithm of [74] to conclude the proof. Let us outline that algorithm in our setting. We start with a set of translates of P, namely $\{z + P, z \in Z\}$. At each step if two of the current droplets P_1, P_2 satisfy that there exists $x \in \mathbb{Z}^2$ such that $(P + x) \cap P_1 \neq \varnothing$ and $(P + x) \cap P_2 \neq \varnothing$, then we replace them by the smallest $\mathcal{T}$-droplet containing their union. We repeat this as long as possible.

By Lemma 4.6 of [74] (stating that the diameter of the smallest droplet containing two intersecting ones is at most the sum of their respective diameters) the sum of diameters of droplets increases by at most $\operatorname{diam}(P) = O(D^3)$. Therefore, in the final set of droplets the total diameter is $O(D^4)$, as the number of droplets decreases by 1 at each step. Moreover, by Lemma 4.5 of [74] the union of the final droplets contains $[Z]$, so the proof is complete, as each of the output droplets contains at least one site of $Z \subset \mathbb{H}_{-u}$. □

Algorithm. Let us first describe an algorithm to determine $\alpha(u)$ and postpone its analysis. For each integer k from 1 to D we successively perform the following operations to determine if there exists a set Z of size k as in Definition 1.6.1. We stop as soon as such a set is found and return the corresponding (minimal) value of k. For each fixed k we start by choosing a set Z_0. The first site is 0 and each new one z is picked within distance $D^{11} \cdot 2^D$ from some of the previous ones and such that

$$0 \leqslant \langle z' - z, u \rangle = O(D^4) \tag{9.4}$$

for some z' among the previous ones. There are at most

$$\binom{D^{O(1)} \cdot 2^D}{D} = 2^{D^2+o(D^2)} = \exp(O(D^2))$$

such choices. For each of them we successively inspect different translations $t \in \mathbb{Z}^2$, such that $0 \leqslant \langle t, u\rangle = O(D^5)$ and

$$0 \leqslant \langle t, (-y, x)\rangle < x^2 + y^2, \tag{9.5}$$

where $(-y, x) \in \mathbb{Z}^2$ is such that $(x, y) \in u\mathbb{R}$ and x and y are co-prime, in the (total) order given by $\langle t, u\rangle$ starting from $t = 0$. Finally, fix $Z = Z_0 + t$.

For each Z we run the bootstrap dynamics with initial set of infections $Z \cup \mathbb{H}_u$ until it either stops infecting new sites or infects a site s with $\|s\|_\infty \geqslant D^{13} \cdot 2^D$ and $\langle s, u\rangle = O(D^5)$. This can be done by checking at each step each site at distance $D^{13} \cdot 2^D + D$ from the origin for each rule and repeating this for 5^D time steps. If the dynamics becomes stationary, we continue to the next choice of Z, while otherwise we return $|Z|$ for the value of $\alpha(u)$.

Correctness. We now turn to proving that the algorithm does return an output and it is precisely $\alpha(u)$. The first assertion is easy. Indeed, as u is an isolated stable direction, (by [70, Lemma 2.8]) there exists a rule $U \in \mathcal{U}$ with

$$U \subset \mathbb{H}_u \cup \{x \in l_u, \langle x, u + \pi/2\rangle > 0\},$$

so that adding D consecutive sites on l_u to $\mathbb{H}_u$ is enough to infect a half-line of l_u, only taking U into account. Thus, we know that $\alpha(u) \leqslant D$ and the algorithm will eventually check such a configuration when $k = D$, unless it has returned a smaller value, and infections will propagate to distance $D^{13} \cdot 2^D$ (and in fact to infinity). Let us then prove that the output is $\alpha(u)$.

Denote by t_j the values of t considered by the algorithm, so that $t_0 = 0$. Note that by (9.5) there exists a single $t \in \mathbb{Z}^2$ with a given value of $\langle t, u\rangle$, so that this scalar product indeed defines a total order on the values of t and we can also extend our notation to $j < 0$ for convenience, though those are not examined by the algorithm. Further define $l_j := \{s \in \mathbb{Z}^2, \langle s, u\rangle = \langle t_j, u\rangle\}$ and $Z_j = Z_0 + t_j$ for some Z_0 considered by the algorithm, so that $l_0 = l_u$ by abuse of notation.[2]

Claim. Assume that a set Z_i considered by the algorithm is of size $k \leqslant \alpha(u)$ and such that $\bar{Z}_j$ (recall (9.1)) is finite for all $0 \leqslant j \leqslant i$. Then the maximal distance between a site from $\bar{Z}_i$ and Z_i is at most $D^5 \cdot 2^D \max(0, \langle t_i, u\rangle)$.

[2]Here we view 0 as an element of $\mathbb{Z}$, possible value of j, while u is an element of S^1. As we will not make reference to l_v with $v = 0 \in S^1$, we hope that this will not lead to confusion.

Proof. We prove the statement by induction on $i \in \mathbb{Z}$. For $i < 0$, i.e. $\langle t_i, u\rangle < 0$, then $Z_i \subset \mathbb{H}_u$ by (9.4) and there is nothing to prove, since $\bar{Z}_i = \varnothing$ – no additional infections take place. Assume the property to hold for all t_j with $j \leqslant i$. We aim prove the same for $i+1$.

Observe that for each $0 < j \leqslant i+1$ we have that

$$\bar{Z}_{i+1} \cap l_j \subseteq (\bar{Z}_{i+1-j} \cap l_0) + t_{i+1} - t_{i+1-j}. \tag{9.6}$$

Indeed, $Z_{i+1} \cup \mathbb{H}_u \subseteq (Z_{i+1-j} \cup \mathbb{H}_u) + t_{i+1} - t_{i+1-j}$, since $Z_{i+1} = Z_{i+1-j} + t_{i+1} - t_{i+1-j}$ and $\mathbb{H}_u + t_{i+1} - t_{i+1-j} \supset \mathbb{H}_u$. Furthermore, by stability of u we have that $\bar{Z}_{i+1} \cap l_j = \varnothing$ for $j > i+1$. Also, by (9.6) and the induction hypothesis we have that $\bar{Z}_{i+1} \backslash l_0$ is at distance at most $D^5 \cdot 2^D \langle t_i, u\rangle$ from Z_{i+1}, so we are left with proving that sites in $\bar{Z}_{i+1} \cap l_0$ are at distance at most $D^5 \cdot 2^D \langle t_{i+1}, u\rangle$ from Z_{i+1}.

Consider the set

$$Z' = \{z \in \bar{Z}_{i+1} \cap l_0, d(z, Z_{i+1}) \leqslant D + D^5 \cdot 2^D \langle t_i, u\rangle\}.$$

By the reasoning above we have that $\bar{Z}_{i+1} \cap l_0 = Z' \cup \bar{Z}'$. However, by Lemma 9.2.1, $\bar{Z}'$ cannot be at distance more than $2^D \cdot D^3$ from Z', as $\mathbb{Z}_{i+1} \backslash l_0$ is at distance at least D from all sites in $\bar{Z}_{i+1} \backslash Z'$. Recalling the definition of Z', we get that $\bar{Z}_{i+1}$ is at distance at most $D + D^3 \cdot 2^D + D^5 \cdot 2^D \langle t_i, u\rangle$ and we are done. Indeed, $\langle t_{i+1} - t_i, u\rangle \geqslant 1/D$, since there exists a site $x \in \mathbb{Z}^2 \cap u\mathbb{R}$ with $\|x\|_\infty \leqslant D/2$ and $\langle t_{i+1} - t_i, x\rangle > 0$ is an integer. □

The claim clearly implies that the algorithm cannot return a value smaller than $\alpha(u)$. In order to conclude, we need to show that when $k = \alpha(u)$ among the sets examined by the algorithm there will be a set Z such that there exists $z \in \bar{Z}$ with $\|z\|_\infty \geqslant D^{13} \cdot 2^D$ and therefore the output will be $\alpha(u)$.

Consider a set $Z \subset \mathbb{Z}^2 \backslash \mathbb{H}_u$ as in Definition 1.6.1 of size $\alpha(u)$ (and therefore minimal). Recall that by Lemma 9.2.2 every $z \in Z$ satisfies $\langle z, u\rangle = O(D^4)$ (otherwise $\bar{Z} = [Z]$ is finite, as $\mathcal{U}$ is not supercritical) and, by stability of u, the same holds for $\bar{Z}$. Let $\mathcal{P} = \{x \in \mathbb{R}, \exists z \in Z, \langle z, u\rangle = x\}$ and define $\bar{\mathcal{P}}$ similarly for $\bar{Z}$. These are discrete subsets of $\mathbb{R}$. Note that by minimality of Z and Lemma 9.2.2, $\mathcal{P} \subset \mathbb{R}$ cannot have a gap of length larger than $O(D^4)$. Indeed, there exists $x \in \bar{\mathcal{P}}$ such that infinitely many points of $\bar{Z}$ project to it and those are all in $\bar{Z}'$ where Z' are the sites in Z that project to $x' \in \mathcal{P}$ such that there exist n and $x_0 = x, x_1, \dots, x_n = x'$ in $\mathcal{P}$ with $|x_{j+1} - x_j| = O(D^4)$ and if $Z' \neq Z$, we obtain a contradiction with the minimality of Z.

Analogously, let $\mathcal{P}^\perp = \{x \in \mathbb{R}, \exists z \in Z, \langle z, (u + \pi/2)\rangle = x\}$ and define $\bar{\mathcal{P}}^\perp$ similarly for $\bar{Z}$. We claim that its $\mathcal{P}^\perp$ cannot have a gap of length larger than $O(D^{10} \cdot 2^D)$. This time $\bar{\mathcal{P}}^\perp$ is necessarily infinite, as only a finite number of points $z \in \mathbb{Z}^2$ with $\langle z, u\rangle = O(D^4)$ have the same $(u + \pi/2)$-projection. Considering a set $Z' \subset Z$ inducing the corresponding distance $O(D^{10} \cdot 2^D)$-connected component of $\mathcal{P}^\perp$ and using the claim instead of Lemma 9.2.2 as in the previous paragraph, we reach a contradiction with the minimality of Z.

Hence, all Z of size $\alpha(u)$ as in Definition 1.6.1 are in fact considered by the algorithm. Since such a Z with infinite $\bar{Z}$ exists, the algorithm does indeed output $\alpha(u)$. □

9.3 NP-hardness: proof of Theorem 9.1.3

In this section we prove Theorem 9.1.3 by providing a reduction from SET COVER to 2D CRITICAL BOOTSTRAP DIFFICULTY. For the SET COVER problem we consider a *universe* $\{1,\ldots,N\}$ and a collection $\mathcal{S}$ of subsets of the universe and assume that $|\mathcal{S}| \geqslant 4$ and $N \geqslant 4$. The SET COVER problem asks for determining the minimum cardinality of a subset of $\mathcal{S}$ which covers the universe. It is one of the first NP-complete problems described by Karp [242].

We fix an instance

$$\mathcal{S} = \{S_i \ : \ i \in \mathbb{Z}, 1 \leqslant i \leqslant |\mathcal{S}|\}.$$

Our goal is to define a critical bootstrap percolation update family whose difficulty α is (up to a simple transformation) the solution to SET COVER. Let the set of rules associated to $\mathcal{S}$ be

$$\mathcal{U}_{\mathcal{S}} = \{U_0, U_1\} \cup \{U_{i,j}^k \ : \ 1 \leqslant i \leqslant |\mathcal{S}|,\, 1 \leqslant k \leqslant |\mathcal{S}|^2,\, i,k \in \mathbb{Z},\, j \in S_i\},$$

where

$$\begin{aligned} U_0 &= \left\{(-k,0),(0,-k) \ : \ 1 \leqslant k \leqslant N|\mathcal{S}|^2\right\}, \\ U_1 &= \left\{(+k,0),(0,-k) \ : \ 1 \leqslant k \leqslant N|\mathcal{S}|^2\right\} \end{aligned}$$

and the rules $U_{i,j}^k$, defined as follows, share a large portion of their structure (see Figure 9.1).

$$\begin{aligned} T &= \left\{(0,-y) \ : \ 1 \leqslant y \leqslant N \cdot |\mathcal{S}|^2\right\}, \\ W &= \{(x,0) \ : \ 1 \leqslant x \leqslant |\mathcal{S}|^2\} \cup \{(l \cdot |\mathcal{S}|, 1) \ : \ 1 \leqslant l \leqslant |\mathcal{S}|\}, \\ U_{i,j}^k &= T \cup \big((W \cup \{(i \cdot |\mathcal{S}|, 2)\}) - (k + (N+j) \cdot |\mathcal{S}|^2, 0)\big). \end{aligned}$$

First we claim that the only isolated stable direction is $u = \pi/2$ and $[-\pi, 0]$ contains the rest of the stable directions. The unstable intervals corresponding to the rules U_0 and U_1 are $(0,\pi/2)$ and $(\pi/2,\pi)$, respectively. The unstable interval of $U_{i,j}^k$ is contained in $(0,\pi/2)$ for all i,j,k. Thus, $\mathcal{U}_{\mathcal{S}}$ is indeed critical and $\alpha(\mathcal{U}_{\mathcal{S}}) = \alpha(u)$, so that we can focus on this direction.

Let $M \subseteq \{1,\ldots,|\mathcal{S}|\}$ be an optimal solution to the SET COVER problem given by $\mathcal{S}$ i.e. a set of minimal size such that

$$\bigcup_{i \in M} S_i = \{1,\ldots,N\}.$$

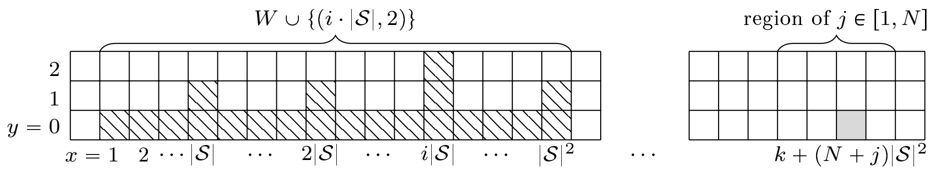

Figure 9.1 – A visualisation of $(U_{i,j}^k \backslash T) + (k + (N + j)|\mathcal{S}|^2, 0)$; the shaded cell indicates where the origin is shifted to.

We first claim that, setting

$$Z_0 = W \cup \{(i \cdot |\mathcal{S}|, 2) \ : \ i \in M\}$$

we have $[Z_0 \cup \mathbb{H}_u] \supset l_u$, so that

$$\alpha(u) \leqslant |Z_0| = |W| + |M| = |\mathcal{S}|^2 + |\mathcal{S}| + |M|. \tag{9.7}$$

Indeed, using once each of the rules $U_{i,j}^k$ for all $i \in M$, $j \in S_i$ and $1 \leqslant k \leqslant |\mathcal{S}|^2$, one infects all sites in

$$\left[1 + (N + 1) \cdot |\mathcal{S}|^2, (2N + 1) \cdot |\mathcal{S}|^2\right] \times \{0\},$$

since M is a cover, and those are enough to infect l_u using U_0 and U_1.

For any $Z \subseteq \mathbb{Z}^2$ recall the notation $\bar{Z} = [Z \cup \mathbb{H}_u] \backslash \mathbb{H}_u$ from (9.1). To prove that (9.7) is actually an equality, we suppose that there exists a set $Z \subset \mathbb{Z}^2 \backslash \mathbb{H}_u$ for which $|\bar{Z}| = \infty$ and $|Z| < |Z_0|$. Fix a minimal such set Z.

First note that $|U_0 \backslash \mathbb{H}_u| = N|\mathcal{S}|^2$ and similarly for U_1. Therefore, if there exists $p \in \mathbb{Z}^2 \backslash \mathbb{H}_u$ such that one of $p + U_0$ and $p + U_1$ is a subset of $Z \cup \mathbb{H}_u$, then $|Z| \geqslant N|\mathcal{S}|^2 > |Z_0|$ – a contradiction. However, in order not to have $\bar{Z} = \varnothing$ some of the rules must be applicable to $Z \cup \mathbb{H}_u$ and therefore there exists $p \in \mathbb{Z}^2 \backslash \mathbb{H}_u$ such that $p + W \subseteq Z$.

Observation 9.3.1. For any $q \in \mathbb{Z}^2 \backslash \{0\}$ we have $|(q + W) \backslash W| > |\mathcal{S}|$.

Although the verification is immediate, calling this fact an observation is deceptive, since W is designed to possess this property. It follows that p is unique, otherwise $|Z| > |W| + |\mathcal{S}| \geqslant |Z_0|$ (since any minimal cover is smaller than the universe), a contradiction.

Lemma 9.3.2. *Every point $q \in \bar{Z} \backslash Z$ has the same y-coordinate as p.*

Proof. Suppose that there exists $q \in \bar{Z} \backslash Z$ contradicting the statement of the lemma and consider such a q with minimal infection time for the process with initial set of infections $Z \cup \mathbb{H}_u$. Then Z contains at least $|W| - |\mathcal{S}|$ sites on the row of q, as all rules contain at least as many and by minimality of q. Therefore, $|Z| \geqslant 2(|W| - |\mathcal{S}|) > |Z_0|$, a contradiction. □

By Lemma 9.3.2 and the fact that $(Z \cup \mathbb{H}_u) - (0,1) \subseteq (Z - (0,1)) \cup \mathbb{H}_u$ and $(Z \cup \mathbb{H}_u) + (1,0) = (Z + (1,0)) \cup \mathbb{H}_u$, we can assume without loss of generality that $p = 0$.

By the minimality of Z and Lemma 9.3.2, the y-coordinate of any site in Z is 0, 1, or 2. Indeed, in order to infect each of the sites $q \in \bar{Z} \subseteq l_u$, we use one of the rules, but those are all contained in $\{x \in \mathbb{Z}^2, \langle x, u\rangle \leqslant 2\}$, so one can remove any other sites from Z without changing $\bar{Z}$.

Lemma 9.3.3. *There does not exist $q \in \mathbb{Z}^2\backslash\{0\}$ such that $q + W \subseteq \bar{Z}$.*

Proof. Let q be as in the statement of the lemma such that no other $q' + W$ becomes fully infected before $q + W$ for the process with initial infections $Z \cup \mathbb{H}_u$. By Lemma 9.3.2 we have that $q \in l_u$.

If $|x| \geqslant |\mathcal{S}|^2$, then by Lemma 9.3.2 the set $Z\backslash W$ contains at least $|W\backslash l_u| = |\mathcal{S}|$ elements (with y-coordinate 1), therefore $|Z| \geqslant |W| + |\mathcal{S}| \geqslant |Z_0|$, a contradiction.

Assume that $|x| < |\mathcal{S}|^2$. If $l_u \cap (q+W)\backslash W \subseteq Z$, then by Observation 9.3.1 we have $|Z| \geqslant |W| + |\mathcal{S}|$ – a contradiction. Therefore, some of the sites in $l_u \cap (q + W)\backslash W \subseteq \bar{Z}$ are infected by the process. However, by minimality of q they can only be infected using U_0 or U_1. Yet, as soon as one can use rule U_0 or U_1 to infect a site in l_u, the entire l_u can be infected using those rules only. Thus, removing from Z every site in $Z\backslash W$ with y-coordinate 1 (and in particular $(q + W)\backslash(l_u \cup W) \neq \varnothing$) does not prevent the infection of infinitely many sites, which contradicts the minimality of Z. □

By Lemma 9.3.3 we have that until a rule U_0 or U_1 is used the only possible infections are of the form "$k + (N + j)|\mathcal{S}|^2$ becomes infected via rule $U^k_{i,j}$". Therefore, all sites $(x, 2) \in Z$ are either redundant (which contradicts the minimality of Z) or satisfy $x = i \cdot |\mathcal{S}|$ with $1 \leqslant i \leqslant |\mathcal{S}|$.

Finally, set $I = \{i \ : \ (i \cdot |\mathcal{S}|, 2) \in Z\}$ and

$$J = \{1, \dots, N\}\backslash \bigcup_{i \in I} S_i.$$

Then, in order to have $|\bar{Z}| = \infty$, it is necessary (and sufficient) to have a sequence of $N|\mathcal{S}|^2$ consecutive sites in

$$(Z \cap l_u) \cup \{(k + (N + j)|\mathcal{S}|^2, 0) \ : \ i \in I, 1 \leqslant k \leqslant |\mathcal{S}|^2, j \in S_i\}.$$

However, such a sequence is either disjoint from the infections of the form $(k + (N + j)|\mathcal{S}|^2, 0)$, in which case $|Z| \geqslant N|\mathcal{S}|^2 > |Z_0|$ – a contradiction, or disjoint from W. In the latter case the sequence contains at most

$$|Z| - |W| - |I| + (N - |J|) \cdot |\mathcal{S}|^2 < (|Z_0| - |W|) + (N - |J|)|\mathcal{S}|^2$$

infected sites. If $|J| \neq N$, i.e. I is not a cover, the number of sites is at most $|\mathcal{S}| + (N - 1)|\mathcal{S}|^2 < N|\mathcal{S}|^2$ – a contradiction. Otherwise, I is a cover and $|Z| \geqslant |W| + |I| \geqslant |Z_0|$, as M is a minimal cover. This contradiction completes the proof that $\alpha(u)$ is indeed equal to $|W|+|M| = |\mathcal{S}|^2+|\mathcal{S}|+|M|$ as claimed.

The set $\mathcal{U}_{\mathcal{S}}$ (to which we reduced the SET COVER problem $\mathcal{S}$) contains $|\mathcal{S}|^3 \sum_{S_i \in \mathcal{S}} |S_i|$ rules, each of which has cardinality at most $O(N|\mathcal{S}|^2)$, thus the reduction is indeed polynomial. This concludes the proof of Theorem 9.1.3, because $\alpha(\mathcal{U}_{\mathcal{S}}) - |\mathcal{S}|^2 - |\mathcal{S}|$ is the size of an optimal set cover from $\mathcal{S}$.

9.4 Open problems

Let us conclude with a few open questions naturally suggested by the present chapter. Of course, many more complexity issues arise systematically for hard problems, but let us mention the foremost ones.

Question 9.4.1. *Can one find a good approximation of α in time polynomial of the input size $\|\mathcal{U}\|$ (defined in (9.2))?*

Question 9.4.2. *Are there interesting subfamilies of critical models for which the difficulty is computable in polynomial time $\|\mathcal{U}\|$?*

Question 9.4.3. *In view of Remark 9.1.2, can one find an algorithm which computes α in $e^{O(\|\mathcal{U}\|)}$ time?*

In Section 9.A we provide an example showing the α itself can be exponentially large in $\|\mathcal{U}\|$, suggesting that one should not hope for a subexponential complexity algorithm to compute it.

Question 9.4.4. *Is the* 2D CRITICAL BOOTSTRAP DIFFICULTY *problem in* NP *(and thus* NP*-complete)?*

Appendix

9.A Relevance of the diameter

In this appendix we provide a sequence $(\mathcal{U}_k)_{k=2}^{\infty}$ of update families such that $\sum_{U\in\mathcal{U}_k} |U|$ is constant and $\alpha(\mathcal{U}_k)$ is exponential in $\|\mathcal{U}_k\|$. This answers a question raised during the preparation of this chapter. The example gives some relevance to the questions in Section 9.4 as well as further justifying the definition of $\|\mathcal{U}\|$ in equation (9.2). For any integer $k \geqslant 2$ let $\mathcal{U}_k = \{U_1, U_2\}$ with

$$U_1 = \{(0,-1),(k,0),(k-1,0)\}$$
$$U_2 = \{(0,-1),(-k,0),(-k+1,0)\}.$$

Proposition 9.A.1. *For any integer $k \geqslant 2$ the update family $\mathcal{U}_k$ is critical and*

$$\alpha(\mathcal{U}_k) = k = \frac{D}{2} = \frac{1}{2} \cdot e^{\|\mathcal{U}_k\|/6}.$$

Proof. It is not hard to check as in the examples in Figure 1.2 that (similarly to the Duarte model) the set of stable directions for $\mathcal{U}_k$ is $[-\pi, 0] \cup \{\pi/2\}$, so the model is critical. Moreover, $\alpha := \alpha(\mathcal{U}_k) = \alpha(u)$ where $u := \pi/2$ is the only isolated stable direction.

It suffices to prove that $\alpha = k$. Consider $Z_0 = \{(i, 0) : i \in \{1, \ldots, k\}\}$ and observe that $[Z_0 \cup \mathbb{H}_u] = \mathbb{H}_u \cup l_u$. Indeed, by stability of u we have $[Z_0 \cup \mathbb{H}_u] \subseteq \mathbb{H}_u \cup l_u$, while using U_1 one can infect successively $(-i, 0)$ for all $i \leqslant 0$. Similarly, using U_2, one can infect $(k + i, 0)$ for $i > 0$.

We are thus left with proving that for any $Z \subset \mathbb{Z}^2$ with $|Z| < k$ we have $|\bar{Z}| < \infty$. Consider a minimal set Z contradicting this statement.

Let $p(i, j) = (i, 0)$ be the projection onto l_u and let $p(Z) = \{p(z) : z \in Z\}$ be the projection of Z. We claim that

$$\overline{p(Z)} \supseteq p(\bar{Z}). \tag{9.8}$$

Let $l_j = \{(i, j) : i \in \mathbb{Z}\}$ and let $m = \max\{j : \bar{Z} \cap l_m \neq \varnothing\}$. By stability of u we have that $l_m \cap Z \neq \varnothing$. As $(0, -1) \in U_1 \cap U_2$, we have that $p((\bar{Z} \backslash Z) \cap l_m) \subseteq p(\bar{Z} \cap l_{m-1})$. Moreover, since $U_1 \cup U_2 \subset \mathbb{H}_u \cup l_u$, we have $\bar{Z} \cap l_{m-1} = \overline{(Z \backslash l_m)} \cap l_{m-1}$. Therefore, if we consider $Z' = (Z \backslash l_m) \cup ((Z \cap l_m) - (0, 1))$, i.e. we decrease the y-coordinates of all sites in $Z \cap l_m$ by 1, we have that

$$p(\bar{Z}') \supseteq p(\bar{Z} \cap l_m). \tag{9.9}$$

Furthermore, as $U_1 \cup U_2 \subset \mathbb{H}_u \cup l_u$ and $Z' \cap (\mathbb{H}_u \cup \bigcup_{j<m} l_j) \supseteq Z \cap (\mathbb{H}_u \cup \bigcup_{j<m} l_j)$, we have

$$\overline{Z' \cap (\mathbb{H}_u \cup \bigcup_{j<m} l_j)} \supseteq \overline{Z \cap (\mathbb{H}_u \cup \bigcup_{j<m} l_j)}.$$

Combining this with (9.9), we get that $p(\bar{Z}') \supseteq p(\bar{Z})$. Repeating this procedure until $m = 0$, we obtain (9.8).

By stability of u we have that $\bar{Z} \subseteq \bigcup_{0 \leqslant j \leqslant m} l_j$, so $\bar{Z}$ is infinite if and only if $p(\bar{Z})$ is. Since $|p(Z)| \leqslant |Z|$, we may replace Z by $p(Z)$ and assume without loss of generality that $Z \subset l_u$. As l_u identifies with $\mathbb{Z}$ by $(i, 0) \mapsto i$, the following lemma concludes the proof. □

Lemma 9.A.2. *Consider the* 1*-dimensional update family consisting of the rules $U_1 = \{k, k-1\}$ and $U_2 = \{-k, -k+1\}$. There does not exist $Z \subset \mathbb{Z}$ with $|Z| < k$ such that $|[Z]| = \infty$.*

Proof. Notice that if $z \in [Z]\backslash Z$ is used to infect another site using rule U_1, then either $z - k$ or $z - (k-1)$ gets infected after z, so z is infected using rule U_1. Therefore, $z + k$ and $z + k - 1$ are infected before z.

Let Z be a counterexample of the statement of the lemma. Without loss of generality, we may assume that $\inf([Z]) = -\infty$. Necessarily, there exists $z \in [Z]$ with $z < \min Z - k^2$, which is infected using rule U_1. By the argument above, $z + k$ and $z + k - 1$ are infected via rule U_1 (before z gets infected). Iterating this argument we obtain that $X_0 = \{z + k^2 - k + 1, \ldots, z + k^2\}$ are all infected by rule U_1.

Let $X_i = X_0 + k \cdot i$ and

$$Y_i = \{x - k \cdot i - z - k^2 + k \ : \ x \in X_i,\ x \text{ is infected using } U_1\},$$

so that $Y_0 = \{1, \ldots, k\} \supseteq Y_i$ for all $i \geqslant 0$. As in the proof of Proposition 9.A.1, one can check that $[X_0] = \mathbb{Z}$, so $[Z] = \mathbb{Z}$. Therefore, by an analogous reasoning for U_2, we have that all sites to the right of Z are infected using rule U_2. Thus, $Y_{i_0} = \varnothing$ for i_0 sufficiently large. For any $y \in Y_{i-1}\backslash Y_i$ the site $y + k \cdot i + z + k^2 - k$ is contained in Z, because, by definition, it does not get infected by U_1, and the first argument of this proof shows that it cannot be infected via U_2. Hence, $k = |Y_0\backslash Y_{i_0}| \leqslant |Z|$, a contradiction. $\square$

Chapter 10

The second term for two-neighbour bootstrap percolation in two dimensions

This chapter is based on joint work with Robert Morris [220], proving the lower bound of Theorem 1.4.5 in a slightly different form (recall Section 1.4.3).

In this chapter we exceptionally denote the parameter q of bootstrap percolation by p, say that the set of initial infections A is *p-random* and denote its law by $\mathbb{P}_p$ instead of μ. Rather than working with $\mu\left(\tau_0^{\mathrm{BP}}\right)$, we consider the critical probability

$$p_c\big([n]^d, r\big) := \inf\Big\{p \in (0,1)\ :\ \mathbb{P}_p\big([A] = [n]^d\big) \geqslant 1/2\Big\},$$

where $[n]^d = \{1, \dots n\}^d$ is endowed with the r-neighbour bootstrap percolation model (all sites outside $[n]^d$ being considered healthy). Our goal is to prove the following lower bound.

Theorem 10.0.1. *There exists a constant $C > 0$ such that*

$$p_c\big([n]^2, 2\big) \ \geqslant\ \frac{\pi^2}{18\log n} - \frac{C}{(\log n)^{3/2}}\,.$$

The rest of the chapter is organised as follows. In Section 10.1 we give an outline of the proof of Theorem 10.0.1, and in Section 10.2 we recall some basic tools and facts that we will need later, and set up some useful notation and conventions used throughout the chapter. In Section 10.3 we state (and give an extended sketch of the proof of) our key bounds on the probability that a rectangle is internally filled by A together with a sub-rectangle (the full details of the proof are postponed to Section 10.A) In Section 10.4 we introduce the hierarchies we will use in the proof, prove some standard facts about the family of hierarchies, and describe a partition of this family which plays an important role in the analysis. Finally, in Section 10.5, we prove

Theorem 10.0.1. We finish the chapter, in Section 10.6, by mentioning a couple of natural open problems.

10.1 An outline of the proof

The proof of Theorem 10.0.1 is very technical, so in this section we will attempt to give the reader an easily-digestible outline of the main ideas behind the proof. The main step will be to bound the probability that a 'critical droplet' R (a rectangle with sides of length between $1/p$ and $(1/p)\log(1/p)$) is 'internally filled' by the p-random set A. The claimed lower bound on $p_c([n]^2, 2)$ will follow easily from this bound via a standard argument (using a lemma due to Aizenman and Lebowitz [7] and the union bound). In order to state this theorem precisely, we will need to introduce a little notation.

A *rectangle* is a non-empty set $R \subset \mathbb{Z}^2$ of the form $[a,b] \times [c,d]$; we write $\dim(R) = (b-a+1, d-c+1)$ for the *dimensions* of R. We say that a rectangle R is *internally filled* by A if $[A \cap R] = R$. We also need the function

$$g(z) := -\log\Big(\beta\big(1-e^{-z}\big)\Big) \tag{10.1}$$

where $\beta(u) := \frac{1}{2}\big(u + \sqrt{u(4-3u)}\big)$, which was defined by Holroyd [225], who also proved that

$$\int_0^\infty g(z)\,\mathrm{d}z \;=\; \lambda := \frac{\pi^2}{18}. \tag{10.2}$$

Finally, set $q := -\log(1-p)$, and note that $q \geqslant p$, and that $q \sim p$ as $p \to 0$. (This notation is convenient, because the probability that a set of size a contains no element of the p-random set A is e^{-aq}. We will assume throughout that $p \to 0$.) We can now state our main bound on the probability that a critical droplet is internally filled.

Theorem 10.1.1. *There exists a constant $C > 0$ such that the following holds. Let R be a rectangle with dimensions $\dim(R) = (a,b)$, and suppose that $a \leqslant b$, and*

$$\frac{C}{q} \;\leqslant\; b \;\leqslant\; \frac{1}{2q}\log\frac{1}{q}. \tag{10.3}$$

Then

$$\mathbb{P}_p\big([A \cap R] = R\big) \leqslant \exp\left(-\min\left\{\frac{2\lambda}{q} + \frac{1}{q^{3/4}},\; (b-a)g(aq) + \frac{2}{q}\int_0^{aq} g(z)\,\mathrm{d}z - \frac{C}{\sqrt{q}}\right\}\right).$$

We remark that the first term in the minimum is easily large enough for our purposes, and is only needed for technical reasons; the reader should therefore focus her attention on the second term. Let us write $\mathrm{long}(R)$ and $\mathrm{short}(R)$ for the maximum and minimum (respectively) of the dimensions of

R. In order to deduce Theorem 10.0.1 from Theorem 10.1.1, we will need the following fundamental lemma of Aizenman and Lebowitz [7].

The Aizenman–Lebowitz lemma. *If* $[A] = [n]^2$*, then for each* $1 \leqslant k \leqslant n$ *there exists a rectangle* R *with*

$$k \leqslant \mathrm{long}(R) \leqslant 2k$$

that is internally filled by A*.*

To deduce a lower bound on $p_c([n]^2, 2)$, we simply apply the Aizenman–Lebowitz lemma with $k = (1/(4q))\log(1/q)$, and take a union bound over choices of R, using Theorem 10.1.1 to bound the probability that R is internally filled, and the (straightforward) fact that

$$(b-a)g(aq) + \frac{2}{q}\int_0^{aq} g(z)\,\mathrm{d}z \geqslant \frac{2\lambda}{q} - \frac{O(1)}{\sqrt{q}}$$

if $a \leqslant b$ and $b \geqslant (1/(4q))\log(1/q)$, see Lemma 10.2.9.

Our main challenge will therefore be to prove Theorem 10.1.1. As has become standard in the area since their introduction by Holroyd [225], we will do so using hierarchies; however, our definition will differ in various important ways from that used in [225], and also from the various notions of hierarchy used in, for example, [69, 70, 129, 190]. These were discussed in detail in Section 1.4.3.

10.2 Basic facts and definitions

In this section we will recall a few basic facts about two-neighbour bootstrap percolation on $[n]^2$, state a few simple properties of the function $g(z)$, and introduce some further notation. For convenience, let us fix (for the rest of the chapter) sufficiently large constants $B > 0$ and $C = C(B) > 0$, and a sufficiently small constant $\delta = \delta(B, C) > 0$.

10.2.1 Preliminaries

To begin, recall the following simple and well-known fact (see, e.g., [68, Problem 34]). We write $\phi(R)$ for the semi-perimeter of a rectangle R, so $\phi(R) = \mathrm{long}(R) + \mathrm{short}(R)$.

Lemma 10.2.1. *If* $[A \cap R] = R$*, then* $|A \cap R| \geqslant \dfrac{\phi(R)}{2}$*.*

Now, recall from (10.1) the definition of the function $g(z)$. The next lemma, which bounds the probability that a sufficiently small rectangle is internally filled, follows easily from Lemma 10.2.1 (see, e.g., [190, Lemma 2]).

Lemma 10.2.2. *There exists $\delta > 0$ such that for any $p > 0$ and any rectangle R with $\dim(R) = (a,b)$, where $a \leqslant b$ and $ap \leqslant \delta$,*

$$\mathbb{P}_p\big([A \cap R] = R\big) \leqslant 3^{\phi(R)} \exp\Big(- \phi(R) g(aq)\Big) .$$

In order to control the growth of a droplet, we will need to bound various probabilities relating to the existence of double gaps. To be precise, let us say that a rectangle $R = [a,b] \times [c,d]$ has a *vertical double gap* if there exists $j \in [a, b-1]$ such that

$$A \cap \big([j, j+1] \times [c,d]\big) = \varnothing ,$$

and similarly for a *horizontal double gap*. (We will say that R has a double gap if it has a horizontal or vertical double gap.) We will say that R is *crossed from left to right*[1] if it has no vertical double gap and the rightmost column $\{b\} \times [c,d]$ is *occupied*, that is, has non-empty intersection with A. Note that if the column to the left of R is already infected, and R is crossed from left to right, then R will also be infected by the process. The following simple estimates were proved in [225, Lemma 8].

Lemma 10.2.3. *If R is a rectangle with $\dim(R) = (a,b)$, then*

$$\mathbb{P}_p\big(R \text{ has no vertical double gap}\big) \leqslant e^{-(a-1)g(bq)}$$

and

$$\mathbb{P}_p\big(R \text{ is crossed from left to right}\big) \leqslant e^{-ag(bq)} .$$

We remark that the function g is positive, decreasing, convex and differentiable on $(0,\infty)$, that $g(z) \sim e^{-2z}$ as $z \to \infty$, that

$$-\frac{1}{2}\log z - \sqrt{z} \;\leqslant\; g(z) \;\leqslant\; -\frac{1}{2}\log z + z \tag{10.4}$$

for all sufficiently small $z > 0$ (see [190, Observation 4]), that

$$e^{2g(z)} \leqslant \frac{C}{z} \tag{10.5}$$

for all $0 < z \leqslant 3e^{2B}$ (see [190, Observation 10]), and that

$$-g'(z) \leqslant \begin{cases} B/z & \text{if } z \leqslant B \\ 3e^{-2z} & \text{if } z \geqslant B/2 \end{cases} \tag{10.6}$$

since B and $C = C(B)$ were chosen sufficiently large.

[1] We define similarly the notions of being crossed from right to left, bottom to top, and top to bottom.

10.2.2 Analytic estimates

We will use the following definition from [225] to control the growth of a droplet.

Definition 10.2.4. For each $\mathbf{a} \leqslant \mathbf{b} \in \mathbb{R}^2_+$, define

$$W(\mathbf{a},\mathbf{b}) = \inf_{\gamma\,:\,\mathbf{a}\to\mathbf{b}} \int_\gamma \big(g(y)\,\mathrm{d}x + g(x)\,\mathrm{d}y\big)\,, \tag{10.7}$$

where the infimum is taken over all piecewise linear increasing paths from $\mathbf{a}$ to $\mathbf{b}$ in $\mathbb{R}^2$.

Now, for any pair $S \subset R$ of rectangles, define

$$U(S,R) := W\big(q\dim(S), q\dim(R)\big)\,. \tag{10.8}$$

One of the key lemmas from [225] states that the integral in (10.7) is minimized when the path γ is chosen as close to the diagonal as possible. We will use the following immediate consequence of this fact.

Lemma 10.2.5 (Lemma 16 of [225])**.** *Let $S \subset R$ be rectangles with* $\mathrm{long}(S) \leqslant \mathrm{short}(R)$. *Then*

$$\frac{U(S,R)}{q} \;=\; (d-c)g(dq) + \frac{2}{q}\int_{dq}^{aq} g(z)\,\mathrm{d}z + (b-a)g(aq)\,,$$

where $a = \mathrm{short}(R)$, $b = \mathrm{long}(R)$, $c = \mathrm{short}(S)$ *and* $d = \mathrm{long}(S)$.

When $\mathrm{long}(S) > \mathrm{short}(R)$ we will use the following easy consequence of the fact that $g(z)$ is decreasing (it also follows immediately from [225, Lemma 16]).

Lemma 10.2.6. *Let $S \subset R$ be rectangles with* $\mathrm{long}(S) \geqslant \mathrm{short}(R)$. *Then*

$$\frac{U(S,R)}{q} \;\geqslant\; (b-d)g(aq)\,,$$

where $a = \mathrm{short}(R)$, $b = \mathrm{long}(R)$ *and* $d = \mathrm{long}(S)$.

We will also need the following straightforward bound from [190].

Lemma 10.2.7 (Lemma 14 of [190])**.** *Let $S \subset R$ be rectangles, such that* $\mathrm{long}(S) \leqslant \mathrm{short}(R)$. *Then*

$$\frac{U(S,R)}{q} \;\geqslant\; \frac{2}{q}\int_0^{aq} g(z)\,\mathrm{d}z + (b-a)g(aq) - \frac{\phi(S)}{2}\log\left(1 + \frac{1}{\phi(S)q}\right) - O\big(\phi(S)\big)\,,$$

where $a = \mathrm{short}(R)$ *and* $b = \mathrm{long}(R)$.

In order to transition between $U(S,R)$ and the bounds of Section 10.3, below, we will also need the following simple upper bound. If S and R are rectangles with dimensions $\dim(R) = (a,b)$ and $\dim S = (a-s, b-t)$, then set

$$Q(S,R) := sg\big((b-t)q\big) + tg\big((a-s)q\big)\,, \tag{10.9}$$

The following lemma follows immediately from the fact that $g(z)$ is decreasing.

Lemma 10.2.8 (Proposition 13 of [225]). *Let $S \subset R$ be rectangles. Then*

$$\frac{U(S,R)}{q} \leqslant Q(S,R)\,.$$

We will also need a couple of additional technical lemmas, each of which follows easily from simple properties of the function g. The first is a variant of [190, Observation 19], with slightly weaker assumptions and conclusion.

Lemma 10.2.9. *If $a \leqslant b$ and $b \geqslant (1/(4q))\log(1/q)$, then*

$$\frac{2}{q}\int_0^{aq} g(z)\,\mathrm{d}z + (b-a)g(aq) \;\geqslant\; \frac{2\lambda}{q} - \frac{4e^4}{\sqrt{q}}\,.$$

Proof. Recall that $B > 0$ is a sufficiently large constant, and note that if $a \leqslant B/q$ then

$$(b-a)g(aq) \;\geqslant\; \frac{g(B)}{5q}\log\frac{1}{q} \;\geqslant\; \frac{2\lambda}{q}\,,$$

since g is decreasing and $q \to 0$. Let us therefore assume that $a \geqslant B/q$, and observe that therefore

$$\int_{aq}^{\infty} g(z)\,\mathrm{d}z \;\leqslant\; g(aq)\,, \tag{10.10}$$

since $g(z) \sim e^{-2z}$ as $z \to \infty$, and hence, recalling the definition (10.2) of λ,

$$\frac{2}{q}\int_0^{aq} g(z)\,\mathrm{d}z + (b-a)g(aq) \;\geqslant\; \frac{2}{q}\int_0^{\infty} g(z)\,\mathrm{d}z \;=\; \frac{2\lambda}{q}$$

if $b-a \geqslant 2/q$. Finally, if $b-a \leqslant 2/q$, then $a \geqslant (1/(4q))\log(1/q) - 2/q$, and so

$$\frac{2}{q}\int_0^{aq} g(z)\,\mathrm{d}z \;\geqslant\; \frac{2\lambda}{q} - \frac{2g(aq)}{q} \;\geqslant\; \frac{2\lambda}{q} - \frac{4e^4}{\sqrt{q}}\,,$$

by (10.10), and since $g(z) \leqslant 2e^{-2z}$ if $z \geqslant B$. □

The next lemma quantifies how much harder it is for a droplet to grow far from the diagonal. To state it, we need to introduce a further large constant $L_1 = L_1(B, C, \delta) > 0$.

Lemma 10.2.10. *If $L_1 a \leqslant b \leqslant B/q$, then*

$$\frac{2}{q}\int_{aq}^{bq} g(z)\,\mathrm{d}z \;\leqslant\; (b-a)\big(g(aq)+g(bq)\big) - 4Cb\,.$$

Proof. We claim first that $-B/z \leqslant g'(z) \leqslant -\delta/z$ for every $0 < z < B$. Indeed, this follows since $g'(z) \sim -1/(2z)$ as $z \to 0$ and $g'(z) \sim -2e^{-2z}$ as $z \to \infty$, and since B was chosen sufficiently large, and $\delta = \delta(B)$ sufficiently small. Now, integrating by parts, we obtain

$$\frac{2}{q}\int_{aq}^{bq} g(z)\,\mathrm{d}z \;\leqslant\; 2\big(bg(bq) - ag(aq)\big) + 2B(b-a)\,.$$

It follows that

$$\frac{2}{q}\int_{aq}^{bq} g(z)\,\mathrm{d}z - (b-a)\big(g(aq)+g(bq)\big) \leqslant (a+b)\big(g(bq)-g(aq)+2B\big)\,.$$

Now, since $g'(z) \leqslant -\delta/z$ for every $z < B$, and $b/a \geqslant L_1$, we have

$$g(aq) - g(bq) \;=\; -\int_{aq}^{bq} g'(z)\,\mathrm{d}z \;\geqslant\; \delta \log L_1 \;\geqslant\; 5C\,,$$

and so the claimed bound follows. □

We will also need some larger constants, denoted by $L_2, L_3, \ldots$, where each L_i is chosen to be sufficiently large depending on B, C, δ, and all of $L_1, \ldots, L_{i-1}$. We will use $O(\cdot)$ to denote the existence of an absolute constant, that is, a constant that does *not* depend on any of the aforementioned ones.

10.2.3 Correlation inequalities

To finish this section, we will state the fundamental inequalities of van den Berg and Kesten [357] and Reimer [305], which we will use in Section 10.4 to bound the probability that a hierarchy is 'satisfied' by A, the p-random set of infected sites, see Definition 10.4.3 and Lemma 10.4.7.

In our setting, an *event* $\mathcal{E}$ is simply a family of subsets of $[n]^2$, and the event $\mathcal{E}$ is said to *occur* if $A \in \mathcal{E}$. Two events $\mathcal{E}$ and $\mathcal{F}$ are said to *occur disjointly* for A if there exist disjoint sets $X, Y \subset [n]^2$ depending on A such that $S \in \mathcal{E}$ for any S such that $S \cap X = A \cap X$, and $T \in \mathcal{F}$ for any T such that $T \cap Y = A \cap Y$. We write $\mathcal{E} \circ \mathcal{F}$ for the event that $\mathcal{E}$ and $\mathcal{F}$ occur disjointly.

Recall that we write $\mathbb{P}_p$ to indicate that A is a p-random subset of $[n]^2$. The following fundamental lemma was proved in 1985 by van den Berg and Kesten [357].

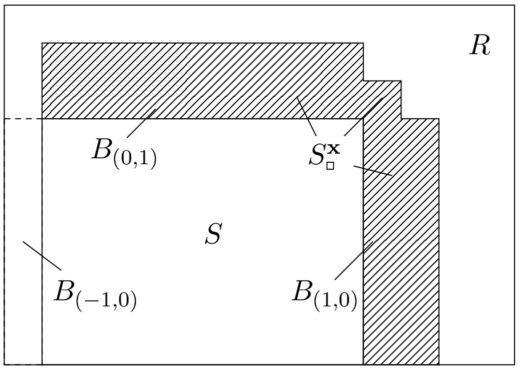

Figure 10.1 – An example of a frame. The non-empty buffers and frame $S^{\mathbf{x}}_{\square}$ (hatched) of S in R with $x_{1,0} = x_{0,1} = 1$, $x_{-1,0} = x_{0,-1} = 0$. Note that the buffers may have width 1.

The van den Berg–Kesten Lemma. *Let $\mathcal{E}$ and $\mathcal{F}$ be any two increasing events and let $p \in (0,1)$. Then*

$$\mathbb{P}_p(\mathcal{E} \circ \mathcal{F}) \leqslant \mathbb{P}_p(\mathcal{E})\,\mathbb{P}_p(\mathcal{F})\,.$$

The authors of [357] also conjectured that their inequality holds in the following more general setting; this was proved 15 years later by Reimer [305].

Reimer's Theorem. *Let $\mathcal{E}$ and $\mathcal{F}$ be any two events and let $p \in (0,1)$. Then*

$$\mathbb{P}_p(\mathcal{E} \circ \mathcal{F}) \leqslant \mathbb{P}_p(\mathcal{E})\,\mathbb{P}_p(\mathcal{F})\,.$$

We remark that the events which we will need to consider will not all be increasing (or decreasing); however, they will all be obtained by intersecting an increasing event with a decreasing event. For such events the conclusion of Reimer's theorem was proved earlier, by van den Berg and Fiebig [356], and the proof is significantly simpler.

10.3 The key lemmas

In this section we will state our key bounds (Lemmas 10.3.3 and 10.3.4, below) on the probability that a rectangle R is internally filled by the union of A (chosen according to $\mathbb{P}_p$) and a rectangle $S \subset R$. In order to simplify the statement somewhat, we will begin by giving some rather technical definitions, which are illustrated in Figure 10.1.

Throughout this section, we will assume that $S \subset R$ are rectangles with $\text{short}(S) \geqslant 2$.

Definition 10.3.1. The *buffers* of S in R are the sets

$$B_{(i,j)}(S,R) := \big\{ v \in R \backslash S \,:\, v - (2i,2j) \in S \big\}\,,$$

where $(i,j) \in \mathcal{I} := \big\{(1,0),(0,1),(-1,0),(0,-1)\big\}$. We call the elements of $\mathcal{I}$ *directions*, define

$$Z(S,R) := \big\{ \mathbf{d} \in \mathcal{I} \,:\, B_{\mathbf{d}}(S,R) \neq \varnothing \big\}$$

to be the collection of non-empty buffers of S in R, and set $z(S,R) = |Z(S,R)|$.

Given $\mathbf{x} = (x_{\mathbf{d}})_{\mathbf{d}\in\mathcal{I}} \in \{0,1\}^{\mathcal{I}}$, define the $\mathbf{x}$-*buffer* of S in R to be

$$B^{\mathbf{x}}(S,R) := \bigcup_{\mathbf{d}\in\mathcal{I}\,:\,x_{\mathbf{d}}=1} B_{\mathbf{d}}(S,R)\,,$$

and the $\mathbf{x}$-*frame* of S in R to be the set

$$S^{\mathbf{x}}_{\square} := B^{\mathbf{x}}(S,R) \cup \big\{ v \in R\backslash S \,:\, |N(v) \cap B^{\mathbf{x}}(S,R)| \geqslant 2 \big\}\,.$$

and set $S^{\mathbf{x}}_{\blacksquare} := S \cup S^{\mathbf{x}}_{\square}$. Thus $\mathbf{x}$ encodes the inclusion in $S^{\mathbf{x}}_{\square}$ of some of the (non-empty) buffers, and also the 'corner' site in between two selected buffers. We will write x and y for the number of non-empty horizontal and vertical buffers included in $B^{\mathbf{x}}(S,R)$, i.e.,

$$x := x'_{(1,0)} + x'_{(-1,0)} \qquad \text{and} \qquad y := x'_{(0,1)} + x'_{(0,-1)}\,, \tag{10.11}$$

where $\mathbf{x}' = \mathbf{x}\cdot\mathbb{1}_{Z(S,R)}$ (i.e., $x'_{\mathbf{d}} := x_{\mathbf{d}}$ if $\mathbf{d} \in Z(S,R)$, and $x'_{\mathbf{d}} := 0$ otherwise).

We are now ready to define our key technical events, which will appear in our hierarchies (see Section 10.4, below), and are designed to be sufficiently unlikely, and to occur disjointly.

Definition 10.3.2. Let the rectangles $S \subset R$, and $\mathbf{x} \in \{0,1\}^{\mathcal{I}}$, be as described above.

(**a**) $D^{\mathbf{x}}_1(S,R)$ denotes the event that

$$\big[S \cup \big(A \cap R\backslash S^{\mathbf{x}}_{\blacksquare}\big)\big] = R\,.$$

(**b**) $D^{\mathbf{x}}_2(S,R)$ denotes the event

$$D^{\mathbf{x}}_1(S,R) \cap \big\{ A \cap S^{\mathbf{x}}_{\square} = \varnothing \big\}\,.$$

The main results of this section are the following two lemmas, which provide us with close to best possible upper bounds on the probabilities of the events $D^{\mathbf{x}}_1(S,R)$ and $D^{\mathbf{x}}_2(S,R)$. The statements are designed to facilitate a proof by induction.

Lemma 10.3.3. *Let* $S \subset R$ *be rectangles with* $\dim(R) = (a,b)$ *and* $\dim(S) = (a-s, b-t)$, *let* $\mathbf{x} \in \{0,1\}^{\mathcal{I}}$ *and set* $z = z(S,R)$. *If*

$$L_1 \leqslant \mathrm{short}(R) \leqslant \frac{B}{q} \qquad \textit{and} \qquad \mathrm{long}(R) \leqslant \frac{3e^{2B}}{q}\,, \tag{10.12}$$

and $s,t \leqslant 4\delta\sqrt{\mathrm{short}(R)}$, *then*

$$\mathbb{P}_p\big(D^{\mathbf{x}}_1(S,R)\big) \leqslant C^z \left(\frac{C}{\sqrt{a}}\right)^y \left(\frac{C}{\sqrt{b}}\right)^x \exp\big(-sg(bq) - tg(aq)\big)\,.$$

Lemma 10.3.4. *Let* $S \subset R$ *be rectangles with* $\dim(R) = (a,b)$ *and* $\dim(S) = (a-s, b-t)$*, let* $\mathbf{x} \in \{0,1\}^{\mathcal{I}}$*, and set* $z = z(S,R)$*. If*

$$\text{short}(R) > \frac{B}{q} \qquad \textit{and} \qquad \text{long}(R) \leqslant \frac{1}{2q} \log \frac{1}{q} \tag{10.13}$$

and $s,t \leqslant \dfrac{4\delta}{\sqrt{q}} \cdot \exp\big(\text{short}(R) \cdot q\big)$*, then*

$$\mathbb{P}_p\big(D_2^{\mathbf{x}}(S,R)\big) \leqslant \Big(Ce^{\text{short}(R)q}\Big)^z \big(C\sqrt{q}e^{-aq}\big)^y \Big(C\sqrt{q}e^{-bq}\Big)^x \exp\big(-sg(bq) - tg(aq)\big).$$

It will be convenient later when applying these lemmas to combine them as follows. First, the following function encodes the upper bounds on s and t:

$$f(R) := \begin{cases} \delta\sqrt{\text{short}(R)} & \text{if } \text{short}(R) \leqslant \dfrac{B}{q}, \\ \dfrac{\delta}{\sqrt{q}} \exp\big(\text{short}(R) \cdot q\big) & \text{otherwise.} \end{cases} \tag{10.14}$$

Let us say that a rectangle R is 1-*critical* if it satisfies the bounds in (10.12), and 2-*critical* if it satisfies the bounds in (10.13). Recall that if S and R have dimensions $\dim(R) = (a,b)$ and $\dim S = (a-s, b-t)$, then

$$Q(S,R) = sg\big((b-t)q\big) + tg\big((a-s)q\big),$$

and if $\mathbf{x} \in \{0,1\}^{\mathcal{I}}$ then write $\|\mathbf{x}\| := x + y = \sum_{\mathbf{d} \in Z(S,R)} x_{\mathbf{d}}$. The following corollary is an almost immediate consequence of Lemmas 10.3.3 and 10.3.4.

Corollary 10.3.5. *Let* $S \subset R$ *be rectangles such that* $\dim(R) = (a,b)$ *and* $\dim(S) = (a-s, b-t)$*, let* $\mathbf{x} \in \{0,1\}^{\mathcal{I}}$*. Let* $j \in \{1,2\}$*, and suppose that* R *is* j*-critical. If* $s,t \leqslant 4f(R)$*, then*

$$\mathbb{P}_p\big(D_j^{\mathbf{x}}(S,R)\big) \leqslant C^9 \left(\frac{\delta}{f(R)}\right)^{\|\mathbf{x}\|} \exp\big(-Q(S,R) + 4\phi(R)q\big). \tag{10.15}$$

Proof. The claimed inequality follows from those given by Lemmas 10.3.3 and 10.3.4 using the bounds on $g'(z)$ given in (10.6), and noting that $x + y + z \leqslant 8$ and $z \leqslant 4$. To spell out the details, recall from (10.6) that $g'(z) \geqslant -B/z$ if $z \leqslant B$ and $g'(z) \geqslant -3e^{-2z}$ if $z \geqslant B/2$, and note that $g'(z)$ is increasing and that $f(R) \leqslant \delta/q$. It follows that

$$\begin{aligned} &\exp\big(-sg(bq) - tg(aq)\big) \\ &\qquad \leqslant \exp\Big(-Q(S,R) - 2stq \cdot g'\big(\big(\text{short}(R) - 4f(R)\big) \cdot q\big)\Big) \\ &\qquad \leqslant \exp\big(-Q(S,R) + \delta\big) \end{aligned}$$

since $s,t \leqslant 4f(R)$ and $\delta = \delta(B)$ is sufficiently small. □

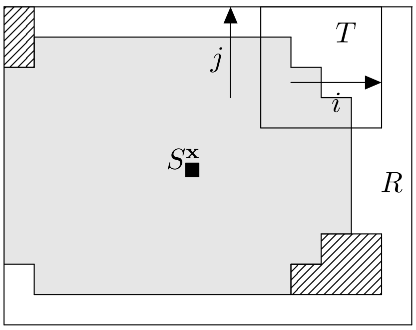

(a) Case 1: the rectangle T is internally filled outside the shaded $S^{\mathbf{x}}_{\blacksquare}$ and allows S to grow i to the right and j upwards.

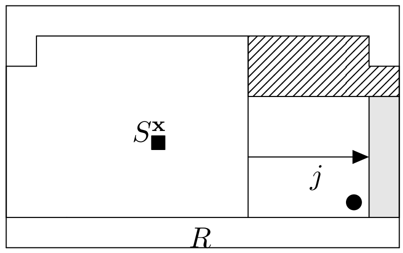

(b) Case 2: S grows j to the right until it reaches a double gap (shaded). The last column before that is necessarily occupied. In both figures, the hatched region is assumed (in the sketch proof) to be unoccupied.

Figure 10.2 – Two possible growth mechanisms.

Since the proofs of Lemmas 10.3.3 and 10.3.4 involve a significant amount of quite technical (and not especially illuminating) case analysis, we will give here only a sketch, and postpone the full details to Section 10.A.

Sketch of the proof of Lemma 10.3.3. Let R be a 1-critical rectangle with dimensions $\dim(R) = (a, b)$, and for each x, y, z and each $s, t \leqslant 4f(R)$, set

$$F^{x,y,z}(s,t) := C^z \left(\frac{C}{\sqrt{a}}\right)^y \left(\frac{C}{\sqrt{b}}\right)^x \exp\big(-sg(bq) - tg(aq)\big).$$

We will prove, by induction on the pair $(s + t, -(x + y))$, that

$$\mathbb{P}_p\big(D_1^{\mathbf{x}}(S, R)\big) \leqslant F^{x,y,z}(s,t) \tag{10.16}$$

for every $0 \leqslant s, t \leqslant 4f(R)$ and $\mathbf{x} \in \{0,1\}^{\mathcal{I}}$, and every $S \subset R$ with $\dim(S) = (a - s, b - t)$, where x and y are as defined in (10.11), and $z = z(S, R)$.

The base of the induction is the case $\min\{s,t\} = 0$. Without loss of generality suppose that $t = 0$, and note that this implies that $x = y = 0$, since otherwise $\mathbb{P}_p\big(D_1^{\mathbf{x}}(S,R)\big) = 0$. It follows that $R \backslash S$ consists of two rectangles (one of which may be empty), one of which is crossed from left to right, and the other of which is crossed from right to left. By Lemma 10.2.3, it follows that

$$\mathbb{P}_p\big(D_1^{\mathbf{x}}(S,R)\big) \leqslant \exp\big(-sg(bq)\big) \leqslant F^{0,0,z}(s,0),$$

as required. We remark that, since $\text{short}(R) \geqslant L_1$, the function $F^{x,y,z}(s,t)$ is increasing in z and decreasing in x, y, s and t.

For the induction step, fix $\mathbf{x} \in \{0,1\}^{\mathcal{I}}$ and $S \subset R$ with $\dim(S) = (a - s, b - t)$, and assume that (10.16) holds for all smaller values of the pair $(s+t, -(x+y))$ in lexicographical order. We partition into cases, depending on whether or not $z = x + y$.

Case 1: $z = x + y$, i.e., all of the non-empty buffers are included in $S^{\mathbf{x}}_{\square}$.

The key observation in this case is that if the event $D^{\mathbf{x}}_1(S, R)$ holds, then there exists a rectangle $T \subset R$ such that

$$\big[A \cap T \backslash S^{\mathbf{x}}_{\blacksquare}\big] = T \qquad \text{and} \qquad T \cap S^{\mathbf{x}}_{\blacksquare} \neq \varnothing$$

(see Figure 10.2a). For simplicity, we will assume that $\phi(T) \leqslant 36 f(R)$ (the other case is dealt with in Section 10.A), which in particular implies that $\phi(T) \leqslant \delta \cdot \mathrm{short}(S)$.

We will sum over choices of T the probability that

$$\big[A \cap T \backslash S^{\mathbf{x}}_{\blacksquare}\big] = T \qquad \text{and} \qquad \big[S \cup T \cup \big(A \cap R \backslash S^{\mathbf{x}}_{\blacksquare}\big)\big] = R\,. \tag{10.17}$$

Note that these two events depend on disjoint sets of infected sites, and are therefore independent. To bound the probabilities of these events, we will partition according to $k := \phi(T)$, and the dimensions of $[S \cup T]$,

$$\dim\big([S \cup T]\big) = (a - s + i, b - t + j)\,.$$

Note that $4 \leqslant i + j \leqslant k$, and therefore, by Lemma 10.2.1, we have

$$|A \cap T \backslash S^{\mathbf{x}}_{\blacksquare}| \geqslant \frac{k}{2} \geqslant \frac{i+j}{2}\,.$$

Note also that, given i, j and k, we have at most $4k$ choices for the rectangle T (at most k per corner of S). Therefore, given i, j and k, the expected number of rectangles T satisfying the first event in (10.17) is at most

$$4k \binom{k^2}{\lceil k/2 \rceil} p^{\lceil k/2 \rceil} \leqslant (24kp)^{k/2}\,.$$

To bound the probability of the second event in (10.17), we use the induction hypothesis. To do so, however, we need to split into cases according to whether or not the buffers of $[S \cup T]$ that are not adjacent to T contain any elements of A. In this sketch we will assume that they do not; for the full details see Section 10.A.

By the induction hypothesis (under the assumption that no additional infections are found in the buffers), it follows that[2]

$$\mathbb{P}_p\Big(\big[S \cup T \cup \big(A \cap R \backslash S^{\mathbf{x}}_{\blacksquare}\big)\big] = R\Big) \leqslant F^{x-1,y-1,z}(s - i, t - j)\,,$$

and hence, by the argument above, it will suffice (in this case) to show that

$$\sum_{i+j \geqslant 4} \sum_{k=i+j}^{36 f(R)} (24kp)^{k/2} \cdot F^{x-1,y-1,z}(s - i, t - j) \ll F^{x,y,z}(s, t)\,. \tag{10.18}$$

[2] Note that we used here the bound $z([S \cup T], R) \leqslant z(S, R)$, and the fact that F is increasing in z.

To see this, note first that $24kp \leqslant \delta$, since $k \leqslant 36f(R)$, and that we may therefore assume that $k = i + j$. Now, observe that

$$\begin{aligned} \frac{F^{x-1,y-1,z}(s-i,t-j)}{F^{x,y,z}(s,t)} &= \frac{\sqrt{ab}}{C^2}\exp\big(ig(bq)+jg(aq)\big) \\ &\leqslant \frac{\sqrt{ab}}{C^2}\left(\frac{C}{bq}\right)^{i/2}\left(\frac{C}{aq}\right)^{j/2} \end{aligned} \tag{10.19}$$

by (10.5), since $\mathrm{long}(R) \leqslant 3e^{2B}/q$. Since $i+j=k$ and $i,j \geqslant 1$, and recalling that $p \leqslant q$, we have

$$\begin{aligned} \sum_{k=4}^{36f(R)} \sum_{i+j=k} (24kp)^{k/2} \cdot \frac{\sqrt{ab}}{C^2}\left(\frac{C}{bq}\right)^{i/2}\left(\frac{C}{aq}\right)^{j/2} &\leqslant \sum_{k=4}^{36f(R)} \frac{k\cdot(C^2k)^{k/2}}{\min\{a,b\}^{(k-2)/2}} \\ &\leqslant \frac{C^5}{\mathrm{short}(R)}, \end{aligned}$$

since $\mathrm{short}(R) \geqslant L_1$. Combining this with (10.19), we obtain (10.18), as claimed.

Case 2: $z > x + y$, i.e., some non-empty buffer is not included in $S_{\square}^{\mathbf{x}}$.

Without loss of generality, let $B_{(1,0)}(S,R)$ be a non-empty buffer that is not included in $S_{\square}^{\mathbf{x}}$, so $x_{(1,0)} = 0$. The idea is to 'grow' S to the right until we find a double gap, or reach the right-hand side of R, thus leading either to an increase in $x+y$, or a decrease in $s+t$. One significant complication is that before reaching a double gap we might find an infected site in one of the other buffers, which are growing with S (see Figure 10.2b). In this sketch we will assume that this does not occur, and also that we do not reach the right-hand side of R; the other cases are dealt with in Section 10.A.

Let j be the distance to the first double gap to the right of S, that is

$$j := \min\big\{i \geqslant 0 \,:\, A \cap R \cap \big(S + (i+2,0)\big)\backslash\big(S+(i,0)\big) = \varnothing\big\},$$

and denote by $\hat{S} := \bigcup_{i=0}^{j}\big(S+(i,0)\big)$ the rectangle formed by the growth of S to the right, until it reaches that double gap. As noted above, we will assume in this sketch that

$$B_{(1,0)}(\hat{S},R) \neq \varnothing \qquad \text{and} \qquad A \cap \hat{S}_{\square}^{\hat{\mathbf{x}}}\backslash S_{\square}^{\mathbf{x}} = \varnothing\,.$$

where $\hat{\mathbf{x}} := \mathbf{x} + \mathbb{1}_{(1,0)}$ (i.e., $\hat{x}_{(1,0)} = 1$ and $\hat{x}_{\mathbf{d}} = x_{\mathbf{d}}$ for each $(1,0) \neq \mathbf{d} \in \mathcal{I}$). In other words, we found a double gap before reaching the right-hand side of R, and no new infected site was found along the way in any of the buffers. We will sum over choices of j the probability that

$$\big[S \cup \big(A \cap \hat{S}\big)\big] = \hat{S} \qquad \text{and} \qquad \big[\hat{S} \cup \big(A \cap R\backslash \hat{S}_{\blacksquare}^{\hat{\mathbf{x}}}\big)\big] = R\,. \tag{10.20}$$

Note that these two events depend on disjoint sets of infected sites, and are therefore independent; we will bound the first using Lemma 10.2.3, and the second using the induction hypothesis. Indeed, by Lemma 10.2.3 (and since $g(z)$ is decreasing) we have

$$\mathbb{P}_p\Big(\big[S \cup \big(A \cap \hat{S}\big)\big] = \hat{S}\Big) \leqslant \exp\big(-jg(bq)\big)\,,$$

and by the induction hypothesis (under the assumption that no additional infections are found in the buffers and that the right-hand side of R is not reached),

$$\mathbb{P}_p\Big(\big[\hat{S} \cup \big(A \cap R \setminus \hat{S}^{\hat{\mathbf{x}}}_{\blacksquare}\big)\big] = R\Big) \leqslant F^{x+1,y,z}(s-j,t)\,.$$

It follows that the probability that there exists $j \geqslant 0$ such that the events in (10.20) both hold is at most

$$\sum_{j=0}^{s-1} \exp\big(-jg(bq)\big) F^{x+1,y,z}(s-j,t) = \frac{Cs}{\sqrt{b}} \cdot F^{x,y,z}(s,t) \leqslant 4C\delta \cdot F^{x,y,z}(s,t)$$

since $s \leqslant 4\delta\sqrt{\text{short}\, R}$. Since $\delta = \delta(C) > 0$ was chosen sufficiently small, this bound suffices in this case. For the full details of the proof, see Section 10.A. □

The proof of Lemma 10.3.4 is very similar to that of Lemma 10.3.3, and so we shall give here only a single calculation from the proof, which illustrates the main additional technicality that arises in this setting, and shows why the term $e^{\text{short}(R)qz}$ is needed in the statement of the lemma. The full details can once again be found in Section 10.A.

Sketch of the proof of Lemma 10.3.4. Recall that $D_2^{\mathbf{x}}(S,R)$ denotes the event

$$\big[S \cup \big(A \cap R \setminus S^{\mathbf{x}}_{\blacksquare}\big)\big] = R \qquad \text{and} \qquad A \cap S^{\mathbf{x}}_{\square} = \varnothing\,.$$

Let R be a 2-critical rectangle with dimensions $\dim(R) = (a,b)$; as in the proof of Lemma 10.3.3, we use induction on the pair $(s+t, -(x+y))$, this time to prove that

$$\mathbb{P}_p\big(D_2^{\mathbf{x}}(S,R)\big) \leqslant \hat{F}^{x,y,z}(s,t)\,,$$

where

$$\hat{F}^{x,y,z}(s,t) := \Big(Ce^{\text{short}(R)q}\Big)^z \big(C\sqrt{q}e^{-aq}\big)^y \Big(C\sqrt{q}e^{-bq}\Big)^x \exp\big(-sg(bq) - tg(aq)\big)\,,$$

for every $0 \leqslant s,t \leqslant 4f(R)$ and $\mathbf{x} \in \{0,1\}^{\mathcal{I}}$, and every $S \subset R$ with $\dim(S) = (a-s, b-t)$, where x and y are as defined in (10.11), and $z = z(S,R)$.

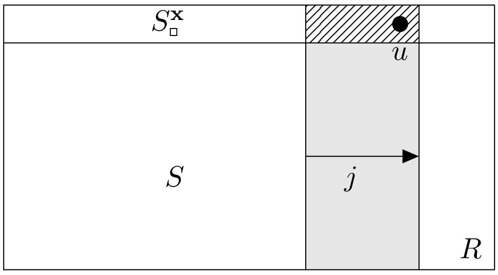

Figure 10.3 – S grows j to the right and reaches the infected site u in the hatched region before a double gap to the right. Thus, the shaded region has no vertical double gap.

In this sketch we will only consider one very particular (but instructive) configuration, which is illustrated in Figure 10.3. In this example, the top buffer is of height 1, the left and bottom buffers are empty, and we attempt to grow S to the right in search of a double gap. However, before finding one, we pass an infected site $u \in A$ in the (new part of the) top buffer, which instead causes us to grow upwards by one step.

Let j denote the $\|\cdot\|_\infty$-distance of u from S, and denote by

$$\hat{S}' := \bigcup_{i=0}^{j} \big(S + (i,0)\big) \qquad \text{and} \qquad \hat{S} := \hat{S}' \cup \big(\hat{S}' + (0,1)\big)$$

so $\hat{S}$ is the rectangle formed by the growth of S to the right, and one step upwards (using u). As noted above, we will assume in this sketch that all of the buffers of $\hat{S}$ are empty except $B_{(1,0)}(\hat{S}, R)$. We will sum over choices of j the probability that $A \cap S^{\mathbf{x}}_{\square} = \varnothing$ (as in the definition of the event $D^{\mathbf{x}}_2(S,R)$), that $u \in A$, that there is no double gap to the right of S before it reaches u, and that

$$\big[\hat{S} \cup \big(A \cap R \backslash \hat{S}^{\hat{\mathbf{x}}}_{\square}\big)\big] = R\,,$$

where $\hat{\mathbf{x}} = \mathbf{x} - \mathbb{1}_{(0,1)} = \mathbf{0}$. Note that these four events are independent, and moreover the probability of each is easy to bound. Indeed, note that $\mathbb{P}_p(u \in A) = p$, that

$$\mathbb{P}_p\big(A \cap S^{\mathbf{x}}_{\square} = \varnothing\big) = (1-p)^{a-s} = \exp\big(-(a-s)q\big) \leqslant 2 \cdot e^{-aq}$$

since $\mathrm{long}(R) \leqslant (1/(2q))\log(1/q)$, and so $s \leqslant (4\delta/\sqrt{q}) \cdot e^{\mathrm{short}(R)q} \leqslant 4\delta/q$, that

$$\mathbb{P}_p\big(\hat{S}' \backslash S \text{ has no vertical double gap}\big) \leqslant \exp\big(-(j-1)g(bq)\big)$$

by Lemma 10.2.3 (and since $g(z)$ is decreasing), and that

$$\begin{aligned}\mathbb{P}_p\Big(\big[\hat{S} \cup \big(A \cap R \backslash \hat{S}^{\hat{\mathbf{x}}}_{\square}\big)\big] = R\Big) &\leqslant \hat{F}^{0,0,1}(s-j,0) \\ &= \hat{F}^{0,1,2}(s,1)\frac{e^{aq-\mathrm{short}(R)q+jg(bq)+g(aq)}}{C^2\sqrt{q}} \cdot,\end{aligned}$$

by the induction hypothesis. It follows that the probability that there exists $j \geqslant 0$ such that the four events above all hold is at most

$$\sum_{j=1}^{s-1} 2p \cdot \frac{e^{-\operatorname{short}(R)q+g(bq)+g(aq)}}{C^2\sqrt{q}} \cdot \hat{F}^{0,1,2}(s,1) \leqslant 4s\sqrt{q} \cdot \frac{e^{-\operatorname{short}(R)q}}{C^2} \cdot \hat{F}^{0,1,2}(s,1)\,, \quad (10.21)$$

since $p \leqslant q$, and since $a, b \geqslant B/q$ implies that $e^{g(aq)+g(bq)} \leqslant 2$. Finally, recall that $s\sqrt{q} \leqslant 4\delta \cdot e^{\operatorname{short}(R)q}$, so the right-hand side of (10.21) is at most $(16\delta/C^2) \cdot \hat{F}^{0,1,2}(s,1)$, as required. Once again, see Section 10.A for the full details of the proof. □

10.4 Hierarchies

In this section we will define precisely the family of hierarchies that we will use in the proof of Theorem 10.1.1. Our definition is more complicated and restrictive than those used in [190, 225], and is designed to take advantage of the bounds proved in Section 10.3, and to allow us to prove a sufficiently strong upper bound on the number of hierarchies.

10.4.1 Good and satisfied hierarchies

Definition 10.4.1. Let R be a rectangle. A *hierarchy* $\mathcal{H}$ for R is an oriented rooted tree $G_{\mathcal{H}}$ with edges pointing away from the root ("downwards"), with edges e labelled with vectors $\mathbf{x}(e) \in \{0,1\}^{\mathcal{I}}$ and vertices u labelled with rectangles $R_u \subset R$. Let $N_{G_{\mathcal{H}}}(u)$ denote the out-neighbourhood of u in $G_{\mathcal{H}}$. We require them to satisfy the following conditions.

(a) The label of the root is R.

(b) Each vertex has out-degree at most two.

(c) If $v \in N_{G_{\mathcal{H}}}(u)$, then $R_v \subset R_u$.

(d) If $N_{G_{\mathcal{H}}}(u) = \{v, w\}$, then $[R_v \cup R_w] = R_u$.

We will write $L(\mathcal{H})$ for the set of leaves of $G_{\mathcal{H}}$, and refer to the rectangles associated with leaves $u \in L(\mathcal{H})$ as *seeds* of the hierarchy. We will also refer to vertices with out-degree two as *split vertices*.

We next define a subclass of 'good' hierarchies that is sufficiently small to allow us to use the union bound (see Lemma 10.4.9), but sufficiently large so that every internally filled rectangle R can be associated with a good hierarchy that encodes the growth of the infected sites inside R (see

Lemma 10.4.4). To do so, we will need one more piece of notation: if $S \subset R$ are rectangles with $R = [r_{(-1,0)}, r_{(1,0)}] \times [r_{(0,-1)}, r_{(0,1)}]$ and $S = [s_{(-1,0)}, s_{(1,0)}] \times [s_{(0,-1)}, s_{(0,1)}]$, then we define

$$d_{\mathbf{j}}(S,R) := |r_{\mathbf{j}} - s_{\mathbf{j}}|$$

for each $\mathbf{j} \in \mathcal{I}$, and $d(S,R) := \max\{d_{\mathbf{j}}(S,R) : \mathbf{j} \in \mathcal{I}\}$.

Definition 10.4.2. We say that a hierarchy $\mathcal{H}$ is *good* if the following conditions hold for every $u \in V(G_{\mathcal{H}})$:

(e) u is a leaf of $G_{\mathcal{H}}$ if and only if $\text{short}(R_u) \leqslant q^{-1/2}$;

(f) if $N_{G_{\mathcal{H}}}(u) = \{v\}$, then

$$d(R_v, R_u) \leqslant 2f(R_u)\,;$$

(g) if $v \in N_{G_{\mathcal{H}}}(u)$ and either $|N_{G_{\mathcal{H}}}(u)| = 2$ or $|N_{G_{\mathcal{H}}}(v)| = 1$, then

$$d(R_v, R_u) \geqslant f(R_u)\,;$$

(h) if $N_{G_{\mathcal{H}}}(u) = \{v\}$ and $|N_{G_{\mathcal{H}}}(v)| = 1$, then $\mathbf{x}(uv) \leqslant \mathbb{1}_{Z(R_v,R_u)}$, and moreover either

  (I) $\|\mathbf{x}(uv)\| = z(R_v, R_u)$, or

  (II) $\|\mathbf{x}(uv)\| = z(R_v, R_u) - 1$ and $d(R_v, R_u) \in \{f(R_u), f(R_u) + 1\}$;

(i) if $v \in N_{G_{\mathcal{H}}}(u)$ and either $|N_{G_{\mathcal{H}}}(u)| \neq 1$ or $|N_{G_{\mathcal{H}}}(v)| \neq 1$, then $\mathbf{x}(uv) = \mathbf{0}$;

where the function $f(R)$ was defined in (10.14), and

$$\|\mathbf{x}(uv)\| = \sum_{\mathbf{d} \in Z(R_v,R_u)} x(uv)_{\mathbf{d}}.$$

Finally, we need to define the family of events that we require to occur disjointly. In order to do so, let us first choose a path (the *trunk*) from the root to a leaf of a hierarchy $\mathcal{H}$ by choosing at each split vertex the out-neighbour whose associated rectangle has larger short side (if they are equal, choose arbitrarily). We will write $\text{tr}(\mathcal{H})$ for the set of edges of the trunk.

Definition 10.4.3. A hierarchy $\mathcal{H}$ is *satisfied* by A if the following events all occur *disjointly*:

(j) If $u \in L(\mathcal{H})$, then the rectangle R_u is internally filled by A;

(k) If $N_{G_{\mathcal{H}}}(u) = \{v\}$ and $uv \notin \mathrm{tr}(\mathcal{H})$, then $D_1^{\mathbf{x}}(R_v, R_u)$ holds, where $\mathbf{x} = \mathbf{x}(uv)$;

(l) If $N_{G_{\mathcal{H}}}(u) = \{v\}$ and $uv \in \mathrm{tr}(\mathcal{H})$, then $D_2^{\mathbf{x}}(R_v, R_u)$ holds, where $\mathbf{x} = \mathbf{x}(uv)$.

We remark that the purpose of the trunk is to guarantee that the unoccupied frames in the events $D_2^{\mathbf{x}}(R_v, R_u)$ occur disjointly. For the sake of brevity, we will often say that a rectangle is in the trunk of a hierarchy $\mathcal{H}$, when we really mean that the associated vertex of $G_{\mathcal{H}}$ is in the trunk, and trust that this will cause no confusion.

10.4.2 Fundamental properties

The following deterministic lemma implies that every internally filled rectangle that satisfies the condition (10.3) of Theorem 10.1.1 has a good and satisfied hierarchy.

Lemma 10.4.4. *Let R be a rectangle that is internally filled by a set A, and suppose that*

$$\mathrm{long}(R) \leqslant \frac{1}{2q} \log \frac{1}{q}\,. \tag{10.22}$$

Then there exists a good hierarchy $\mathcal{H}$ for R that is satisfied by A.

A similar lemma was proved by Holroyd in [225] using the following lemma, which is a straightforward consequence of the 'rectangles process' of Aizenman and Lebowitz [7].

Lemma 10.4.5 (Proposition 30 of [225])**.** *Let R be a rectangle such that* $\mathrm{long}(R) \geqslant 2$. *If R is internally filled by A, then there exist rectangles $S_1, S_2 \subsetneq R$, with $[S_1 \cup S_2] = R$, that are disjointly internally filled by A.*

We will need the following slight (and straightforward) strengthening of this lemma.

Lemma 10.4.6. *Let R be a rectangle with* $\mathrm{long}(R) \geqslant 2$. *If R is internally filled by A, then there exist rectangles $S_1, S_2 \subsetneq R$, with $[S_1 \cup S_2] = R$, such that*

$$\big[A \cap (S_1 \backslash S_2)\big] = S_1 \qquad \text{and} \qquad \big[A \cap (S_2 \backslash S_1)\big] = S_2\,. \tag{10.23}$$

Proof. By taking a subset if necessary, we may assume that A is a minimal percolating set for R, i.e., that A is minimal such that $[A] = R$. We claim that for such a set A, the rectangles S_1 and S_2 given by Lemma 10.4.5 in fact satisfy (10.23). Indeed, suppose that A_1 and A_2 are disjoint subsets of A such that $S_1 = [A_1]$, $S_2 = [A_2]$ and $[S_1 \cup S_2] = R$, and observe that

$$\big[A \backslash (A_1 \cap S_2)\big] = R\,,$$

since $S_2 \subset [A \backslash A_1]$ and $[A] = R$. By the minimality of A, it follows that $A_1 \cap S_2 = \varnothing$, and similarly $A_2 \cap S_1 = \varnothing$ as required. □

Now we prove that any internally filled rectangle has a good and satisfied hierarchy.

Proof of Lemma 10.4.4. The proof is similar to that of [225, Proposition 32], but since there are several slightly subtle (and important) differences, we will give the details in full.

We prove the statement by induction on $\phi(R)$. If $\text{short}(R) \leqslant q^{-1/2}$ then we can let $\mathcal{H}$ be the hierarchy with only one vertex, which is good by the bound on $\text{short}(R)$, and satisfied since R is internally filled by A. So assume that $\text{short}(R) > q^{-1/2}$, and that the lemma holds for all rectangles with semi-perimeter strictly smaller than $\phi(R)$.

We use Lemma 10.4.6 to construct a sequence of rectangles

$$R = T_0 \supsetneq T_1 \supsetneq \cdots \supsetneq T_m$$

for some $m \in \mathbb{N}$ as follows. For each $i \geqslant 0$, suppose that we have already constructed T_i, and let T_{i+1} and T_i' be the two rectangles given by Lemma 10.4.6 applied to T_i, where $d(T_{i+1}, R) \leqslant d(T_i', R)$. Now let m be minimal such that $d(T_m, R) \geqslant f(R)$, and note that m exists because (10.14) and (10.22) imply that $\text{long}(R) > 2f(R)$. We consider three cases:

Case 1: $d(T_m, R) \leqslant 2f(R)$.

In this case, instead of applying the induction hypothesis to T_m (as in, e.g., [190, 225]), we let $T_m \subset S \subset R$ be a maximal internally filled rectangle with $d(S, R) \geqslant f(R)$, and apply the induction hypothesis to S. (We remark that this is a crucial step in our proof.) Observe that, by the maximality of S, one of the following two events holds:

(I) There is no element of A within distance two of S. In this case the event $D_2^{\mathbf{x}}(S, R)$ holds for $\mathbf{x} = \mathbb{1}_{Z(S,R)}$, since $A \cap S_{\square}^{\mathbf{x}} = \varnothing$, and R is internally filled by A.

(II) $d([S \cup \{u\}], R) < f(R)$ for each $u \in A$ within distance two of S, and therefore $d(S, R) \in \{f(R), f(R) + 1\}$. Choose $\mathbf{d} \in \mathcal{I}$ such that $d_{\mathbf{d}}(S, R) \in \{f(R), f(R) + 1\}$, and set $\mathbf{x} = \mathbb{1}_{Z(S,R) \backslash \{\mathbf{d}\}}$. We claim that the event $D_2^{\mathbf{x}}(S, R)$ holds. Indeed, R is internally filled by A, and if there exists an element $u \in A \cap S_{\square}^{\mathbf{x}}$, then we have $d([S \cup \{u\}], R) \geqslant d_{\mathbf{d}}(S, R) \geqslant f(R)$, contradicting the maximality of S.

Now, let $\mathcal{H}'$ be the good and satisfied hierarchy for S given by the induction hypothesis, and form a hierarchy $\mathcal{H}$ for R by adding an edge from a vertex u corresponding to R, to the (root) vertex v of $\mathcal{H}'$ corresponding to S. If $|N_{G_{\mathcal{H}}}(v)| \neq 1$, then set $\mathbf{x}(uv) = 0$, and otherwise define $\mathbf{x}(uv)$ as above, i.e., $\mathbf{x}(uv) = \mathbb{1}_{Z(S,R)}$ in (I) and $\mathbf{x}(uv) = \mathbb{1}_{Z(S,R)\setminus\{\mathbf{d}\}}$ in (II), where $d_{\mathbf{d}}(S,R) \in \{f(R), f(R)+1\}$.

We claim that $\mathcal{H}$ is good, and satisfied by A. To see that $\mathcal{H}$ is good, recall that $\mathcal{H}'$ is good, and note that $f(R) \leqslant d(S,R) \leqslant 2f(R)$, and that if $|N_{G_{\mathcal{H}}}(v)| = 1$ then either $\|\mathbf{x}(uv)\| = z(S,R)$ (if there is no element of A within distance two of S), or $\|\mathbf{x}(uv)\| = z(S,R) - 1$ and $d(S,R) \in \{f(R), f(R)+1\}$ (otherwise).

To see that $\mathcal{H}$ is satisfied by A, recall that $\mathcal{H}'$ is satisfied by A, and note that the event $D_2^{\mathbf{x}(uv)}(S,R)$ occurs (by the observations above). Moreover, the event $D_2^{\mathbf{x}(uv)}(S,R)$ depends only on sites in $R\setminus S$, whereas the events involved in $\mathcal{H}'$ depend only on sites inside S. The events involved in $\mathcal{H}$ therefore occur disjointly, as required.

Case 2: $d(T_1,R) > 2f(R)$.

Let $\{S_1,S_2\} = \{T_1,T_0'\}$, where the labelling is chosen so that $\text{short}(S_1) \geqslant \text{short}(S_2)$, and recall that $[S_1 \cup S_2] = R$, that (10.23) holds, i.e.,

$$\big[A \cap (S_1\setminus S_2)\big] = S_1 \qquad \text{and} \qquad \big[A \cap (S_2\setminus S_1)\big] = S_2\,,$$

and that $2f(R) < \min\{d(S_1,R), d(S_2,R)\}$. Set $A_1 := A \cap S_1$ and $A_2 := A \cap (S_2\setminus S_1)$ and, applying the induction hypothesis, let $\mathcal{H}_1'$ and $\mathcal{H}_2'$ be good hierarchies for S_1 and S_2 that are satisfied by A_1 and A_2, respectively. Form a hierarchy $\mathcal{H}$ for R by adding edges from a vertex u corresponding to R, to the roots of $\mathcal{H}_1'$ and $\mathcal{H}_2'$, that is, the vertices v_1 and v_2 corresponding to S_1 and S_2 (respectively), and set $\mathbf{x}(uv_1) = \mathbf{x}(uv_2) = \mathbf{0}$.

We claim that $\mathcal{H}$ is good, and satisfied by A. To see that $\mathcal{H}$ is good, recall that $\mathcal{H}_1'$ and $\mathcal{H}_2'$ are good, and that $\min\{d(S_1,R), d(S_2,R)\} \geqslant 2f(R)$. To see that $\mathcal{H}$ is satisfied by A, recall that $\mathcal{H}_1'$ and $\mathcal{H}_2'$ are satisfied by A, and note that the trunk of $\mathcal{H}$ can be chosen to pass through S_1. Now, all of the increasing events involved in $\mathcal{H}_1'$ and $\mathcal{H}_2'$ are witnessed by disjoint subsets of A_1 and A_2, respectively, and A_1 and A_2 are disjoint sets (since $A_1 \subset S_1$ and $A_2 \cap S_1 = \varnothing$), so all of these events occur disjointly. Since S_2 is not in the trunk, the only remaining events are the decreasing events involved in $\mathcal{H}_1'$ (that the frames of rectangles in the trunk are empty), which all depend only on sites in $S_1\setminus A_1$, and therefore occur disjointly from those that depend on A_1 and A_2. The events involved in $\mathcal{H}$ therefore occur disjointly, as required.

Case 3: $d(T_m,R) > 2f(R)$, and $m > 1$.

Set $S = T_{m-1}$, and let $\{S_1,S_2\} = \{T_m, T_{m-1}'\}$, where the labelling is chosen so that $\text{short}(S_1) \geqslant \text{short}(S_2)$, and recall that $[S_1 \cup S_2] = S$, that (10.23) holds, and that

$$d(S,R) < f(R) \qquad \text{and} \qquad d(S_i,S) \geqslant d(S_i,R) - d(S,R) > f(R) \quad (10.24)$$

for each $i \in \{1,2\}$. As in Case 2, set $A_1 := A \cap S_1$ and $A_2 := A \cap (S_2\setminus S_1)$ and let $\mathcal{H}_1'$ and $\mathcal{H}_2'$ be good hierarchies for S_1 and S_2, satisfied by A_1 and

A_2, respectively, given by the induction hypothesis. Form a hierarchy $\mathcal{H}$ for R by adding an edge from a vertex u corresponding to R, to a vertex v corresponding to S, and edges from v to the roots of $\mathcal{H}'_1$ and $\mathcal{H}'_2$, that is, the vertices w_1 and w_2 corresponding to S_1 and S_2 (respectively), and set $\mathbf{x}(uv) = \mathbf{x}(vw_1) = \mathbf{x}(vw_2) = \mathbf{0}$.

It is again easy to see that $\mathcal{H}$ is good, since $\mathcal{H}'_1$ and $\mathcal{H}'_2$ are good, and using the inequalities (10.24). We claim that moreover $\mathcal{H}$ is satisfied by A; this follows almost exactly as in Case 2, but for completeness we will spell out the details. Recall that $\mathcal{H}'_1$ and $\mathcal{H}'_2$ are satisfied by A, observe that the event $D_2^{\mathbf{x}}(S,R)$ holds (since R is internally filled by A, and $\mathbf{x}(uv) = \mathbf{0}$), and note that the trunk of $\mathcal{H}$ can be chosen to pass through S_1.

Now, all of the increasing events involved in $\mathcal{H}'_1$, $\mathcal{H}'_2$ and $D_2^{\mathbf{x}}(S,R)$ are witnessed by disjoint subsets of A_1, A_2 and $A\backslash S$, respectively, and A_1, A_2 and $A\backslash S$ are disjoint sets (since $A_1 \subset S_1$, $A_2 \cap S_1 = \varnothing$, and $A_1 \cup A_2 \subset S$), so all of these events occur disjointly. Now, S_2 is not in the trunk, and $\mathbf{x}(uv) = \mathbf{0}$, so the only remaining events are the decreasing events involved in $\mathcal{H}'_1$ (that the frames of rectangles in the trunk are empty), which all depend only on sites in $S_1\backslash A_1$, and therefore occur disjointly from those that depend on A_1, A_2 and $A\backslash S$. The events involved in $\mathcal{H}$ therefore occur disjointly, as required. □

We are now ready to deduce our fundamental bound on the probability that a rectangle is internally filled, cf. [225, Section 10] or [190, Lemma 7]. Given a rectangle R, let us write $\mathcal{H}_R$ for the set of good hierarchies for R, and for each $\mathcal{H} \in \mathcal{H}_R$, set

$$G_{\mathcal{H}}^{(2)} = \mathrm{tr}(G_{\mathcal{H}}) \qquad \text{and} \qquad G_{\mathcal{H}}^{(1)} = E(G_{\mathcal{H}})\backslash G_{\mathcal{H}}^{(2)}\,.$$

Recall also that $\mathbb{P}_p$ denotes the probability space obtained by choosing A to be a p-random subset of $[n]^2$, and let us write $I(R)$ for the event that R is internally filled by A.

Lemma 10.4.7. *If R is a rectangle with $\mathrm{long}(R) \leqslant (1/(2q))\log(1/q)$, then*

$$\mathbb{P}_p\big(I(R)\big) \leqslant \sum_{\mathcal{H}\in\mathcal{H}_R} \Bigg(\prod_{j=1}^{2} \prod_{\substack{uv\in G_{\mathcal{H}}^{(j)} \\ N_{G_{\mathcal{H}}}(u)=\{v\}}} \mathbb{P}_p\big(D_j^{\mathbf{x}(uv)}(R_v,R_u)\big)\Bigg)\Bigg(\prod_{u\in L(\mathcal{H})} \mathbb{P}_p\big(I(R_u)\big)\Bigg). \tag{10.25}$$

Proof. By Lemma 10.4.4, if R is internally filled by A then there exists a hierarchy $\mathcal{H} \in \mathcal{H}_R$ that is satisfied by A. By the union bound (over $\mathcal{H}_R$), it will therefore suffice to show that for each $\mathcal{H} \in \mathcal{H}_R$, the probability that $\mathcal{H}$ is satisfied by A is bounded above by the corresponding term of the right-hand side of (10.25). But this follows immediately from Definition 10.4.3 by Reimer's Theorem, and hence (10.25) holds, as claimed. □

We remark that we did not actually need the full power of Reimer's Theorem in the proof above, since our events are particularly simple: each is the intersection of an increasing and a decreasing event, and the decreasing events are moreover primitive (i.e., a fixed set must be empty). For events of this form, the conclusion of Reimer's theorem is actually a straightforward consequence of the van den Berg–Kesten lemma.

10.4.3 Weighted counting

Recall that we can bound the probabilities on the right-hand side of (10.25) using Lemma 10.2.2 and Corollary 10.3.5. It therefore remains to control the 'size' of the set $\mathcal{H}_R$; however, since hierarchies with many empty buffers are more numerous and less likely to be satisfied, we would like to give them lower 'weight' when measuring the size of $\mathcal{H}_R$. Due to the form of the right-hand side of (10.15) (in particular, its dependence on $\|\mathbf{x}\|$), we will find the following definition useful.

Definition 10.4.8. Given a rectangle R, the *weight* of a hierarchy $\mathcal{H} \in \mathcal{H}_R$ is defined to be

$$w(\mathcal{H}) := \prod_{N_{G_{\mathcal{H}}}(u)=\{v\}} \left(\frac{1}{f(R_u)}\right)^{\|\mathbf{x}(uv)\|}.$$

Given a hierarchy $\mathcal{H}$, we will write $v(\mathcal{H})$ for the number of vertices of $G_{\mathcal{H}}$, and $s(\mathcal{H}) = |L(\mathcal{H})|$ for the number of seeds of $\mathcal{H}$. (Note that $\mathcal{H}$ has exactly $s(\mathcal{H}) - 1$ split vertices.) Given a rectangle R, let us write

$$\mathcal{H}_R(N, M) := \big\{\mathcal{H} \in \mathcal{H}_R : v(\mathcal{H}) = N,\, s(\mathcal{H}) = M\big\}. \tag{10.26}$$

The following lemma bounds the total weight of $\mathcal{H}_R(N, M)$.

Lemma 10.4.9. *Let R be a rectangle, and let $N, M \in \mathbb{N}$. Then*

$$\sum_{\mathcal{H}\in\mathcal{H}_R(N,M)} w(\mathcal{H}) \leqslant \exp\Big(16\big(N + M \log\phi(R)\big)\Big).$$

Proof. Let us first fix the tree $G_{\mathcal{H}}$ and the labels $\mathbf{x}(e)$ for each $e \in G_{\mathcal{H}}$. There are at most 3^N oriented rooted trees on N vertices with maximum out-degree at most two (and edges oriented away from the root), and at most 2^{4N} choices for the labels $\mathbf{x}(e) \in \{0,1\}^{\mathcal{I}}$. We will choose the rectangles one by one, starting at the root and working our way down the tree, counting the number of choices (given the earlier choices) at each step.

Let $u \in V(G_{\mathcal{H}})$, and suppose that we have already chosen the rectangle R_u. Suppose first that u is a split vertex, and let $N_{G_{\mathcal{H}}}(u) = \{v, w\}$. We clearly have at most $\phi(R)^4$ choices for each of R_v and R_w, and hence (recalling that there are $M - 1$ split vertices) the total number of choices for the rectangles associated with the out-neighbours of split vertices is at most $\phi(R)^{8M}$. Similarly, if $N_{G_{\mathcal{H}}}(u) = \{v\}$ and v is a split vertex or a seed, then we have at most $\phi(R)^4$ choices for R_v, so the total number of choices for the rectangles associated with such vertices is also at most $\phi(R)^{8M}$.

Suppose now that $N_{G_{\mathcal{H}}}(u) = \{v\}$ and $|N_{G_{\mathcal{H}}}(v)| = 1$, and recall from Definition 10.4.2 that $d(R_v, R_u) \leqslant 2f(R_u)$, and that either $\|\mathbf{x}(uv)\| = z(R_v, R_u)$, or

$$\|\mathbf{x}(uv)\| = z(R_v, R_u) - 1 \qquad \text{and} \qquad d(R_v, R_u) \in \{f(R_u), f(R_u) + 1\}\,.$$

In either case we have at most $2^{10} f(R_u)^{\|\mathbf{x}(uv)\|}$ choices for R_v, and it follows that

$$\sum_{\mathcal{H} \in \mathcal{H}_R(N,M)} w(\mathcal{H}) \leqslant (3 \cdot 2^{14})^N \cdot \phi(R)^{16M} \leqslant \exp\Big(16\big(N + M \log \phi(R)\big)\Big)\,,$$

as claimed. □

10.4.4 The height of a hierarchy

Let us write $h(\mathcal{H})$ for the *height* of the hierarchy $\mathcal{H}$, that is, the number of vertices in the longest path from the root to a leaf of $G_{\mathcal{H}}$. In this subsection we will prove some straightforward (though sometimes slightly technical) properties of the height of a good hierarchy.

Let us begin with a simple lower bound on the size of a seed in a good hierarchy.

Observation 10.4.10. Let R be a rectangle, and suppose that $\mathrm{long}(R) \leqslant (1/(2q)) \log(1/q)$ and $\mathrm{short}(R) \geqslant q^{-1/2}$. If $\mathcal{H} \in \mathcal{H}_R$, and $v \in V(G_{\mathcal{H}})$, then

$$\phi(R_v) \geqslant \frac{\delta}{q^{1/4}}\,.$$

Proof. It suffices to prove the claimed bound for seeds of $\mathcal{H}$, so assume that v is a seed, and that $v \in N_{G_{\mathcal{H}}}(u)$ (if $\mathcal{H}$ has only one vertex then the result is trivial). Since u is not a seed and $\mathcal{H}$ is a good hierarchy (see Definition 10.4.2), we have $\mathrm{short}(R_u) > q^{-1/2}$. Thus

$$\phi(R_v) \geqslant \min\big\{\phi(R_u) - 8f(R_u),\; f(R_u)\big\} \geqslant \frac{\delta}{q^{1/4}}\,,$$

as required, since if $|N_{G_{\mathcal{H}}}(u)| = 1$, then $d(R_v, R_u) \leqslant 2f(R_u)$, while if $|N_{G_{\mathcal{H}}}(u)| = 2$, then $d(R_v, R_u) \geqslant f(R_u)$. Note that in the first step we used the fact that if $[R_v \cup R_w] = R_u$, then $\phi(R_v) + \phi(R_w) \geqslant \phi(R_u)$, and in the second we used the definition (10.14) of $f(R)$. □

Next, let us recall a simple but key observation from [190]. Let us say that a seed S is *large* if $\text{long}(S) \geqslant 1/(3\sqrt{q})$, and denote by $m(\mathcal{H})$ the number of large seeds of a hierarchy $\mathcal{H}$. Observe (or recall from [190, Observation 17]) that every non-leaf vertex of a good hierarchy lies above a large seed.

Observation 10.4.11. Let R be a rectangle with $\text{short}(R) \geqslant q^{-1/2}$. If $\mathcal{H} \in \mathcal{H}_R$, then

$$v(\mathcal{H}) \leqslant 2 \cdot h(\mathcal{H}) \cdot m(\mathcal{H}) .$$

Proof. Since $\mathcal{H}$ is a good hierarchy for R, every non-leaf $u \in V(G_{\mathcal{H}})$ lies above a large seed. There are therefore at most $h(\mathcal{H}) \cdot m(\mathcal{H})$ vertices that are either large seeds or non-seeds. Since each small seed is adjacent to a non-seed, and each non-seed is adjacent to at most one small seed, the claimed bound follows. □

We will use Observation 10.4.11 together with the following lemma to bound the number of vertices in a 'typical' hierarchy $\mathcal{H} \in \mathcal{H}_R$.

Lemma 10.4.12. *Let R be a rectangle with* $\text{long}(R) \leqslant (1/(2q)) \log(1/q)$, *and let $\mathcal{H} \in \mathcal{H}_R$. Then either*[3]

$$h(\mathcal{H}) \leqslant \frac{L_2}{\sqrt{q}}\,, \tag{10.27}$$

or there exists a vertex $u \in V(G_{\mathcal{H}})$ such that either

$$\text{short}(R_u) \leqslant \frac{B}{q} \qquad \textit{and} \qquad \text{long}(R_u) \geqslant 2L_1 \cdot \text{short}(R_u)\,, \tag{10.28}$$

or

$$\text{short}(R_u) \geqslant \frac{B}{q} \qquad \textit{and} \qquad \text{long}(R_u) \geqslant 4 \cdot \text{short}(R_u)\,. \tag{10.29}$$

Proof. Suppose that there is no vertex $u \in V(G_{\mathcal{H}})$ satisfying either (10.28) or (10.29); we will show that $h(\mathcal{H}) \leqslant L_2/\sqrt{q}$. To do so, let v be the root of $\mathcal{H}$, let P be a longest path in $G_{\mathcal{H}}$ (from v to a seed w), and partition (the vertex set of) P into sets

$$P_1 := \{v\} \cup \big\{u \in P : \text{short}(R_u) > B/q\big\} \qquad \text{and} \qquad P_2 := P \backslash P_1\,.$$

Let u_1 be the lowest vertex of P_1, and let u_2 be the highest vertex of P_2.

We first claim that the distance (in $G_{\mathcal{H}}$) from u_2 to w is a most $L_1/(3\sqrt{q})$. To see this, note that any $w < y \leqslant u_2$ satisfies $\text{long}(R_y) < 2L_1 \cdot \text{short}(R_y)$,

[3]Recall that $L_2 = L_2(B, C, \delta, L_1)$ is a sufficiently large constant. The lemma also holds with a smaller constant in (10.27), but this particular tripartition will be convenient in Section 10.5.

and so, by Definition 10.4.2, in the next two consecutive steps up P the semi-perimeter increases by at least

$$\frac{\delta}{\sqrt{2L_1}} \cdot \sqrt{\operatorname{long}(R_y)}\,.$$

Since L_1 is large, the claimed bound follows easily.

Similarly, we claim that the distance (in $G_{\mathcal{H}}$) from u_1 to v is a most $L_1/(3\sqrt{q})$. To see this, note that any $v > y \geqslant u_1$ satisfies $4 \cdot \operatorname{short}(R_y) > \operatorname{long}(R_y)$, so in the next two consecutive steps up P, either we reach v, or the semi-perimeter increases by at least

$$\frac{\delta}{\sqrt{q}} \exp\big(q \cdot \phi(R_y)/5\big)\,.$$

It is again not difficult to see that the claimed bound holds; indeed, the semi-perimeter takes at most $x = L_1/(6\sqrt{q})$ steps to increase by $5/q$, then at most $x/2$ steps to increase by $5/q$ again, and so on, until it has increased by $(5/q)\log_2 x$ in at most $2x$ steps. □

The next lemma bounds the height of $\mathcal{H}$ in the case where only (10.28) is satisfied.

Lemma 10.4.13. *Let R be a rectangle with $\operatorname{long}(R) \leqslant (1/(2q))\log(1/q)$, and let $\mathcal{H} \in \mathcal{H}_R$. Suppose that neither (10.27) nor (10.29) holds for any vertex $u \in V(G_{\mathcal{H}})$. Then the vertex u satisfying (10.28) may be chosen so that*

$$h(\mathcal{H}) \leqslant L_1 q^{1/4} \cdot \operatorname{long}(R_u)\,. \tag{10.30}$$

Proof. Let $u \in V(G_{\mathcal{H}})$ be a vertex satisfying (10.28) with $\operatorname{long}(R_u)$ maximal, and set $c = \operatorname{short}(R_u)$ and $d = \operatorname{long}(R_u)$. Let P be the longest path in $\mathcal{H}$, and observe that P contains at most $L_1/(3\sqrt{q})$ vertices v with $\operatorname{short}(R_v) > B/q$, as in the proof of Lemma 10.4.12, since $\mathcal{H}$ contains no vertex such that (10.29) holds. Observe also that P contains at most $(L_1 q^{1/4}/2) \cdot d$ vertices v with $\operatorname{long}(R_v) \leqslant d$, since it follows from Definition 10.4.2 that in each two consecutive steps the semi-perimeter increases by at least $\delta q^{-1/4}$.

Finally, we claim that P contains at most $L_2/(3\sqrt{q})$ vertices v with $\operatorname{short}(R_v) \leqslant B/q$ and $\operatorname{long}(R_v) > d$. To see this, note that $2L_1 \cdot \operatorname{short}(R_v) > \operatorname{long}(R_v)$, by our choice of u, and therefore in each two consecutive steps up P, the semi-perimeter increases by at least

$$\frac{\delta}{\sqrt{2L_1}} \cdot \sqrt{\operatorname{long}(R_v)} \geqslant \frac{\delta}{2\sqrt{L_1}} \cdot \sqrt{\phi(R_v)}\,.$$

It now follows easily that after $L_2/(3\sqrt{q})$ steps we have $\phi(R_v) \geqslant 3L_1 B/q$, and hence $\operatorname{short}(R_v) \geqslant B/q$, as claimed.

Since (10.27) does not hold, it follows that

$$\frac{L_2}{\sqrt{q}} \leqslant h(\mathcal{H}) \leqslant \frac{L_1 q^{1/4}}{2} \cdot \mathrm{long}(R_u) + \frac{L_2}{2\sqrt{q}}\,,$$

and hence we obtain (10.30), as required. □

Define the *upper trunk* of $\mathcal{H}$ to be the following set of vertices[4] of the trunk:

$$\mathrm{up}(\mathcal{H}) := \big\{ u \in V(G_{\mathcal{H}}) : u \text{ is in the trunk of } \mathcal{H} \text{ and } \mathrm{short}(R_u) \geqslant B/q \big\}\,.$$

The final lemma of this subsection bounds the sum of the semi-perimeters of rectangles in the upper trunk when there does not exist a vertex $u \in V(G_{\mathcal{H}})$ satisfying (10.29).

Lemma 10.4.14. *Let R be a rectangle with $\mathrm{long}(R) \leqslant (1/(2q))\log(1/q)$, and let $\mathcal{H} \in \mathcal{H}_R$. Then either*

$$\sum_{u \in \mathrm{up}(\mathcal{H})} \phi(R_u) \leqslant \frac{L_2}{q^{3/2}},$$

or there exists a vertex $u \in V(G_{\mathcal{H}})$ such that

$$\mathrm{short}(R_u) \geqslant \frac{B}{q} \qquad \textit{and} \qquad \mathrm{long}(R_u) \geqslant 4 \cdot \mathrm{short}(R_u)\,. \tag{10.31}$$

Proof. If $u \in \mathrm{up}(\mathcal{H})$, and u does not satisfy (10.31), then by Definition 10.4.2 (as in the proof of Lemma 10.4.12), in the next two consecutive steps up the trunk either we reach the root v, or the semi-perimeter increases by at least

$$\frac{\delta}{\sqrt{q}} \exp\big(q \cdot \phi(R_y)/5\big)\,.$$

Set $x = L_1/\sqrt{q}$, and observe (cf. the proof of Lemma 10.4.12) that there are at most $2^{-k+1}x$ vertices $u \in \mathrm{up}(\mathcal{H})$ with $\phi(R_u) \geqslant (B + 5k)/q$, for each $0 \leqslant k \leqslant \log_2 x$. It follows that

$$\sum_{u \in \mathrm{up}(\mathcal{H})} \phi(R_u) \;\leqslant\; \sum_{k=0}^{\infty} \frac{B+5k}{q} \cdot \frac{x}{2^{k-1}} \;\leqslant\; \frac{L_2}{q^{3/2}}\,,$$

as required. □

[4]Recall that $\mathrm{tr}(\mathcal{H})$ denotes the set of *edges* of the trunk; we hope that this minor inconsistency in our notation (which will be quite convenient) will not confuse the reader.

10.4.5 The pods of a hierarchy

To finish this section, let us recall the following important lemma from [225], which is known as the 'pod lemma', and prove a generalization which we be useful in Section 10.5, below. Recall from (10.8) the definition of $U(S, R)$.

Lemma 10.4.15 (Lemma 38 of [225]). *Let $\mathcal{H} \in \mathcal{H}_R$. Then there exists a rectangle $S \subset R$ such that*

$$\dim(S) \leqslant \sum_{w \in L(\mathcal{H})} \dim(R_w)$$

and

$$\sum_{N_{G_{\mathcal{H}}}(u)=\{v\}} U(R_v, R_u) \geqslant U(S, R) - 2\big(s(\mathcal{H}) - 1\big) q g(\sqrt{q})\,.$$

Holroyd called the rectangle S the *pod* of $\mathcal{H}$. Roughly speaking, his Lemma 10.4.15 says that the 'cost' of the growth (given the size of the seeds) is minimized by placing all of the seeds near to one another, at the very bottom of the hierarchy. However, when we are in the case corresponding to (10.28) (or, more precisely, Lemma 10.4.13), we will need to make use of the special rectangle R_u, which is somewhere in the middle of $\mathcal{H}$. In order to use the fact that this rectangle appears in the hierarchy when minimizing the 'cost' of growth, we instead form two pods: one corresponding to the growth inside the special rectangle R_u, the other corresponding to the growth of this rectangle to fill R.

Lemma 10.4.16. *Let $\mathcal{H} \in \mathcal{H}_R$ and let $u \in V(G_{\mathcal{H}})$. Then there exist rectangles $S_1 \subset R_u$ and $R_u \subset S_2 \subset R$, such that*

$$\dim(S_1) + \dim(S_2) - \dim(R_u) \leqslant \sum_{w \in L(\mathcal{H})} \dim(R_w) \tag{10.32}$$

and

$$\sum_{N_{G_{\mathcal{H}}}(v)=\{w\}} U(R_w, R_v) \geqslant U(S_1, R_u) + U(S_2, R) - 2\big(s(\mathcal{H}) - 1\big) q g(\sqrt{q})\,. \tag{10.33}$$

The proof of Lemma 10.4.16 is essentially identical to Lemma 10.4.15, and so we will give only a brief sketch here, and refer the reader to [225] for the details.

Sketch proof of Lemma 10.4.16. We will use induction on the distance from u to the root. Note first that when u is the root of $\mathcal{H}$, then the claimed conclusion follows from Lemma 10.4.15 by setting $S_1 = S$ and $S_2 = R$. For the induction step, we divide into cases according to whether the root has one or two neighbours.

Indeed, suppose first that the root has one neighbour, x, and apply the induction hypothesis to the sub-hierarchy of $\mathcal{H}$ rooted at x to obtain pods S_1 and S_2. Note that these pods satisfy (10.32), and also (10.33), since the inequality

$$U(R_x, R) \geqslant U(S_2, R) - U(S_2, R_x)\,,$$

follows immediately from the definition. On the other hand, if the root has two neighbours, x and y, and u is a descendant of x, then we apply the induction hypothesis to the sub-hierarchy of $\mathcal{H}$ rooted at x, giving pods S_1' and S_2', and Lemma 10.4.15 to the sub-hierarchy of $\mathcal{H}$ rooted at y, giving a pod T. Set $S_1 := S_1'$, and choose S_2, with

$$\dim(S_2') \leqslant \dim(S_2) \leqslant \dim(S_2') + \dim(T)$$

and

$$U(S_2', R_x) + U(T, R_y) \geqslant U(S_2, R) - 2qg(\sqrt{q})$$

by applying [225, Proposition 15], exactly as in the proof of [225, Lemma 38]. Noting that $s(\mathcal{H}) = s(\mathcal{H}_x) + s(\mathcal{H}_y)$, the inequalities (10.32) and (10.33) follow. □

10.5 Proof of Theorem 10.0.1

In this section we will put the pieces together and prove Theorem 10.0.1. The main step is the proof of Theorem 10.1.1, which we restate (this time with explicit constants) for convenience. Recall that $I(R)$ denotes the event that R is internally filled by A.

Theorem 10.5.1. *Let R be a rectangle with dimensions $\dim(R) = (a, b)$, and suppose that $a \leqslant b$, and*

$$\frac{3e^{2B}}{q} \leqslant b \leqslant \frac{1}{2q}\log\frac{1}{q}\,. \tag{10.34}$$

Then

$$\mathbb{P}_p\big(I(R)\big) \leqslant \exp\left(-\min\left\{\frac{2\lambda}{q} + \frac{1}{q^{3/4}},\ (b-a)g(aq) + \frac{2}{q}\int_0^{aq} g(z)\,\mathrm{d}z - \frac{L_6}{\sqrt{q}}\right\}\right).$$

We will begin by giving an outline of the proof of Theorem 10.5.1, and proving a couple of straightforward technical lemmas. Let us fix a rectangle R as in the theorem until the end of its proof. The first step is to recall that

$$\mathbb{P}_p\big(I(R)\big) \leqslant \sum_{\mathcal{H}\in\mathcal{H}_R}\left(\prod_{j=1}^{2}\prod_{\substack{uv\in G_{\mathcal{H}}^{(j)}\\ N_{G_{\mathcal{H}}}(u)=\{v\}}}\mathbb{P}_p\big(D_j^{\mathsf{x}(uv)}(R_v,R_u)\big)\right)\left(\prod_{u\in L(\mathcal{H})}\mathbb{P}_p\big(I(R_u)\big)\right), \tag{10.35}$$

by Lemma 10.4.7, where we used the upper bound on $\text{long}(R)$ from (10.34). Recall that $g(z)$ is decreasing, and (from Definition 10.4.2) that $\text{short}(R_u) \leqslant q^{-1/2}$ for every leaf $u \in L(\mathcal{H})$ of a good hierarchy $\mathcal{H}$. Therefore, if $\mathcal{H} \in \mathcal{H}_R$, then

$$\mathbb{P}_p\big(I(R_u)\big) \leqslant 3^{\phi(R_u)} \exp\Big(- \phi(R_u) g(\sqrt{q})\Big) \tag{10.36}$$

for each leaf $u \in L(\mathcal{H})$, by Lemma 10.2.2. Moreover, since $d(R_v, R_u) \leqslant 2f(R_u)$ whenever $N_{G_{\mathcal{H}}}(u) = \{v\}$, if R_u is j-critical for some $j \in \{1,2\}$, then

$$\mathbb{P}_p\big(D_j^{\mathbf{x}(uv)}(R_v, R_u)\big) \leqslant C^9 \left(\frac{\delta}{f(R_u)} \right)^{\|\mathbf{x}(uv)\|} \exp\big(- Q(R_v, R_u) + 4\phi(R_u) q\big) \tag{10.37}$$

by Corollary 10.3.5, where $Q(R_v, R_u)$ was defined in (10.9). Unfortunately, however, some rectangles are neither 1- nor 2-critical, and we must deal with these separately.

Lemma 10.5.2. *The probability that there exists an internally filled rectangle $S \subset R$ with*

$$\text{short}(S) \leqslant \frac{B}{q} \qquad \textit{and} \qquad \text{long}(S) \geqslant \frac{3e^{2B}}{q} \tag{10.38}$$

or an internally filled rectangle $S \subset R$ with

$$\text{short}(S) \leqslant \frac{1}{q} \qquad \textit{and} \qquad \text{long}(S) \geqslant \frac{B}{2q} \tag{10.39}$$

is at most $e^{-2/q}$.

Proof. Observe that if $S \subset R$ is internally filled, then it must be crossed from left to right, and from bottom to top. By Lemma 10.2.3, it follows that if S satisfies (10.38) then

$$\mathbb{P}_p\big(I(S)\big) \leqslant \exp\Big(- \text{long}(S) \cdot g\big(q \cdot \text{short}(S)\big)\Big) \leqslant \exp\left(- \frac{3e^{2B} g(B)}{q} \right).$$

Recalling that $g(z) \sim e^{-2z}$ as $z \to \infty$ (and that B is large), and applying the union bound, it follows that the probability that there exists such a rectangle S is at most

$$\big(\text{long}(R)\big)^4 \cdot \exp\left(- \frac{3e^{2B} g(B)}{q} \right) \leqslant \frac{1}{2} \cdot e^{-2/q}.$$

The same bound (with the same proof, noting that $B > 4/g(1)$, since B is sufficiently large) holds if S satisfies (10.39). The result then follows by the union bound. □

Note that $\lambda = \pi^2/18 < 1$, so by Lemma 10.5.2 we may assume that R contains no internally filled rectangle S satisfying (10.38) or (10.39). It follows that each rectangle R_u (where $u \in V(G_{\mathcal{H}})$) either satisfies the condition (10.12) of Lemma 10.3.3 (and hence is 1-critical), or satisfies the condition (10.13) of Lemma 10.3.4 (and hence is 2-critical). Note also that, since $b \geqslant 3e^{2B}/q$, by (10.34), we may assume from now on that $a \geqslant B/q$.

The next problem is that we would like R_u to be j-critical when $uv \in G_{\mathcal{H}}^{(j)}$, and this is not necessarily the case. However, since $D_2^{\mathbf{x}(uv)}(R_v, R_u) \subset D_1^{\mathbf{x}(uv)}(R_v, R_u)$, it is not a problem if $uv \in G_{\mathcal{H}}^{(2)} = \mathrm{tr}(G_{\mathcal{H}})$ for some u with $\mathrm{short}(R_u) \leqslant B/q$. The next lemma bounds the probability that there exists $uv \in G_{\mathcal{H}}^{(1)}$ with $N_{G_{\mathcal{H}}}(u) = \{v\}$ and $\mathrm{short}(R_u) > B/q$.

Lemma 10.5.3. *The probability that there exist two disjointly internally filled rectangles $S_1, S_2 \subset R$ with*

$$\min\big\{\mathrm{short}(S_1), \mathrm{short}(S_2)\big\} \geqslant \frac{B}{q} \tag{10.40}$$

is at most $e^{-2/q+o(1/q)}$.

This lemma is an almost immediate consequence of Holroyd's theorem and the van den Berg–Kesten lemma. However, for convenience (since the version we need is not explicitly stated in [225]) we will deduce it from the following (very weak) consequence of [190, Proposition 15], which holds since $2\lambda = \pi^2/9 > 1$.

Proposition 10.5.4. *Let $S \subset R$ be a rectangle with* $\mathrm{short}(S) \geqslant B/q$. *Then*

$$\mathbb{P}_p\big(I(S)\big) \leqslant e^{-1/q}.$$

Proof of Lemma 10.5.3. By the van den Berg–Kesten inequality and Proposition 10.5.4, the probability that two given rectangles S_1 and S_2, each with short side at least B/q, are disjointly internally filled is at most $e^{-2/q}$. By the union bound, it follows that the probability that two such disjointly internally filled rectangles exist is at most

$$\big(\mathrm{long}(R)\big)^8 e^{-2/q} = e^{-2/q+o(1/q)},$$

as claimed. □

Note that if there exists a vertex u with $\mathrm{short}(R_u) \geqslant B/q$ that is not in the trunk, then there must exist a split vertex above u whose neighbours are labelled with disjointly internally filled rectangles S_1 and S_2 satisfying (10.40). Hence, by Lemma 10.5.3, and recalling that $\lambda < 1$, we may assume that every vertex $u \in V(G_{\mathcal{H}})$ with $\mathrm{short}(R_u) \geqslant B/q$ is in the trunk, and hence $uv \in \mathrm{tr}(\mathcal{H})$ whenever $N_{G_{\mathcal{H}}}(u) = \{v\}$ and $\mathrm{short}(R_u) > B/q$. We may therefore apply the inequality (10.37) to bound the probability of the event $D_j^{\mathbf{x}(uv)}(R_v, R_u)$ for each $uv \in G_{\mathcal{H}}^{(j)}$ with $N_{G_{\mathcal{H}}}(u) = \{v\}$. Setting

$$X(\mathcal{H}) := \sum_{u \in L(\mathcal{H})} \phi(R_u),$$

it follows from (10.35), (10.36) and (10.37), and Lemmas 10.5.2 and 10.5.3, that the probability that R is internally filled is bounded from above by $e^{-2/q+o(1/q)}$ plus

$$\sum_{\mathcal{H}\in\mathcal{H}_R^*} 3^{X(\mathcal{H})} e^{-X(\mathcal{H})g(\sqrt{q})} \cdot \prod_{N_{G_{\mathcal{H}}}(u)=\{v\}} C^9 \left(\frac{\delta}{f(R_u)}\right)^{\|\mathbf{x}(uv)\|} \exp\Big(- Q(R_v, R_u) + 4\phi(R_u) q\Big) \,,$$

where $\mathcal{H}_R^*$ denotes the set of hierarchies $\mathcal{H} \in \mathcal{H}_R$ that contain no rectangle satisfying either (10.38) or (10.39), and such that every vertex $u \in V(G_{\mathcal{H}})$ with $\text{short}(R_u) \geqslant B/q$ is in the trunk. By Definition 10.4.8, this is at most

$$\sum_{\mathcal{H}\in\mathcal{H}_R^*} w(\mathcal{H}) \cdot C^{9v(\mathcal{H})} \cdot 3^{X(\mathcal{H})} \cdot e^{-X(\mathcal{H})g(\sqrt{q})} \cdot \prod_{N_{G_{\mathcal{H}}}(u)=\{v\}} \exp\Big(- Q(R_v, R_u) + 4\phi(R_u) q\Big) \,. \qquad (10.41)$$

The rest of the proof of Theorem 10.5.1 is just a careful analysis of (10.41).

Proof of Theorem 10.5.1. As explained above, taking into account Lemmas 10.2.2, 10.4.7, 10.5.2 and 10.5.3, and Corollary 10.3.5, in order to prove the theorem it will suffice to bound (10.41). Let us set

$$\Lambda(\mathcal{H}) := w(\mathcal{H}) \cdot C^{9v(\mathcal{H})} \cdot 3^{X(\mathcal{H})} e^{-X(\mathcal{H})g(\sqrt{q})} \cdot \prod_{N_{G_{\mathcal{H}}}(u)=\{v\}} \exp\Big(- Q(R_v, R_u) + 4\phi(R_u) q\Big)$$

for each $\mathcal{H} \in \mathcal{H}_R^*$, and write $\mathcal{H}_R^{(1)}$ for the set of $\mathcal{H} \in \mathcal{H}_R^*$ such that

$$h(\mathcal{H}) \leqslant \frac{L_2}{\sqrt{q}} \qquad \text{and} \qquad \sum_{u\in \text{up}(\mathcal{H})} \phi(R_u) \leqslant \frac{L_2}{q^{3/2}} \,, \qquad (10.42)$$

cf. Lemmas 10.4.12 and 10.4.14. Let us note that this is the most important class of hierarchies, since it will turn out that the remaining hierarchies $\mathcal{H}_R^* \backslash \mathcal{H}_R^{(1)}$ contribute only smaller order terms to (10.41). To slightly simplify the formulae below, let us write

$$\mathcal{J}(R) := \frac{2}{q} \int_0^{aq} g(z)\, \mathrm{d}z + (b-a) g(aq) \,,$$

where we recall that $a = \text{short}(R)$ and $b = \text{long}(R)$.

Claim 1: $\displaystyle\sum_{\mathcal{H}\in\mathcal{H}_R^{(1)}} \Lambda(\mathcal{H}) \leqslant \exp\bigg(- \mathcal{J}(R) + \frac{L_5}{\sqrt{q}}\bigg) + e^{-2/q}.$

Proof of Claim 1. The proof is a fairly standard (if somewhat complicated) calculation, similar to, e.g., [190], the main new ingredient being the weighted counting of Lemma 10.4.9. The first step is to deal with hierarchies with $X(\mathcal{H}) > 1/q$, and to do so we will first show that

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H}) \cdot C^{10v(\mathcal{H})} \cdot \exp\bigg(- \frac{X(\mathcal{H})}{5} \log\frac{1}{q} + \frac{4L_2}{\sqrt{q}}\bigg) \tag{10.43}$$

for every $\mathcal{H} \in \mathcal{H}_R^{(1)}$. To see this, recall first that every vertex u with $\text{short}(R_u) \geqslant B/q$ is in the trunk, and no rectangle that appears in $\mathcal{H}$ satisfies (10.38). It follows that $\text{short}(R_u) \leqslant B/q$ and $\text{long}(R_u) \leqslant 3e^{2B}/q$ for every $u \notin \text{up}(\mathcal{H})$, and hence, by (10.42), we have

$$\prod_{N_{G_{\mathcal{H}}}(u)=\{v\}} \exp\big(4\phi(R_u)q\big) \leqslant C^{v(\mathcal{H})} \exp\big(4L_2/\sqrt{q}\big)\,, \tag{10.44}$$

since $C = C(B) > 0$ was chosen sufficiently large. Next, observe that

$$\begin{aligned} 3^{X(\mathcal{H})} e^{-X(\mathcal{H}) g(\sqrt{q})} &\leqslant \exp\bigg(-\frac{X(\mathcal{H})}{4}\bigg(\log\frac{1}{3^4 q} - 4q^{1/4}\bigg)\bigg) \\ &\leqslant \exp\bigg(-\frac{X(\mathcal{H})}{5}\log\frac{1}{q}\bigg)\,, \end{aligned} \tag{10.45}$$

since $g(\sqrt{q}) \geqslant \log(q^{-1/4}) - q^{1/4}$, by (10.4). Noting that $Q(R_v, R_u) \geqslant 0$ for every $R_v \subset R_u$, since $g(z)$ is positive, and using (10.44) and (10.45), we obtain (10.43), as claimed.

Now, recall that $m(\mathcal{H})$ denotes the number of large seeds in a hierarchy $\mathcal{H}$, and that

$$v(\mathcal{H}) \leqslant 2 \cdot h(\mathcal{H}) \cdot m(\mathcal{H})$$

by Observation 10.4.11, and observe that therefore

$$X(\mathcal{H}) \geqslant \frac{m(\mathcal{H})}{3\sqrt{q}} + \frac{\delta(s(\mathcal{H}) - m(\mathcal{H}))}{q^{1/4}} \geqslant \frac{v(\mathcal{H})}{L_3} + \frac{s(\mathcal{H})}{q^{1/5}}\,, \tag{10.46}$$

by Observation 10.4.10 and (10.42). We claim next that

$$\sum_{\mathcal{H}\in\mathcal{H}_R^{(1)} :\, X(\mathcal{H})>1/q} \Lambda(\mathcal{H}) \;\leqslant\; e^{-2/q}\,. \tag{10.47}$$

To prove (10.47), let us write

$$\mathcal{H}_R^{(1)}(N,M) := \big\{\mathcal{H} \in \mathcal{H}_R^{(1)} : v(\mathcal{H}) = N,\, s(\mathcal{H}) = M\big\}\,,$$

as in (10.26), and recall that, by Lemma 10.4.9,

$$\sum_{\mathcal{H}\in\mathcal{H}_R^{(1)}(N,M)} w(\mathcal{H}) \leqslant \exp\Big(16\big(N + M\log\phi(R)\big)\Big)\,. \tag{10.48}$$

Combining (10.43) with (10.46) and (10.48), it follows that

$$\begin{aligned}\sum_{\substack{\mathcal{H}\in\mathcal{H}_R^{(1)}\\ X(\mathcal{H})>1/q}} \Lambda(\mathcal{H}) &\leqslant \sum_{N,M} \sum_{\substack{\mathcal{H}\in\mathcal{H}_R^{(1)}(N,M)\\ X(\mathcal{H})>1/q}} w(\mathcal{H})\cdot C^{10v(\mathcal{H})}\cdot \\ &\qquad\qquad \exp\bigg(- \frac{X(\mathcal{H})}{5}\log\frac{1}{q} + \frac{4L_2}{\sqrt{q}}\bigg) \\ &\leqslant e^{-2/q}\sum_{N,M}\exp\big(-CN - Mq^{-1/5}\big) \sum_{\mathcal{H}\in\mathcal{H}_R^{(1)}(N,M)} w(\mathcal{H}) \\ &\leqslant e^{-2/q}\sum_{N,M}\exp\big(-N-M\big) \leqslant e^{-2/q},\end{aligned}$$

as claimed, where in the second step we used the bound $X(\mathcal{H}) > 1/q \gg 1$.

We will therefore assume from now on that

$$X(\mathcal{H}) \leqslant \frac{1}{q}\,. \tag{10.49}$$

We next claim that

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H})\cdot C^{10v(\mathcal{H})}\cdot \exp\bigg(- \frac{U(S,R)}{q} - \frac{X(\mathcal{H})}{4}\log\frac{1}{q} + 13L_2\cdot X(\mathcal{H})\bigg)\,. \tag{10.50}$$

To prove this we repeat the proof of (10.43), being slightly less wasteful in (10.45), and replacing the trivial bound $Q(S,R)\geqslant 0$ by a more complicated argument. To be more precise, recall that $U(R_v,R_u)\leqslant q\cdot Q(R_v,R_u)$ for every $R_v\subset R_u$, by Lemma 10.2.8, and that therefore, by Lemma 10.4.15, there exists a pod S, with $\phi(S)\leqslant X(\mathcal{H})$, such that

$$\sum_{N_{G_{\mathcal{H}}}(u)=\{v\}} Q(R_v,R_u) \geqslant \sum_{N_{G_{\mathcal{H}}}(u)=\{v\}} \frac{U(R_v,R_u)}{q} \;\geqslant\; \frac{U(S,R)}{q} - 2s(\mathcal{H})g(\sqrt{q})\,. \tag{10.51}$$

Now, recall from (10.46) that we have $X(\mathcal{H})\geqslant s(\mathcal{H})\cdot q^{-1/5} \gg s(\mathcal{H})g(\sqrt{q})$, and note that $X(\mathcal{H})\geqslant 1/(3\sqrt{q})$, since every hierarchy in $\mathcal{H}_R$ has at least one large seed. Hence, by (10.44), (10.45) and (10.51), we obtain (10.50), as claimed.

Now, by Lemma 10.2.7, we have

$$\frac{U(S,R)}{q} \geqslant \frac{2}{q}\int_0^{aq} g(z)\,\mathrm{d}z + (b-a)g(aq) - \frac{\phi(S)}{2}\log\bigg(1+\frac{1}{\phi(S)q}\bigg) - O\big(\phi(S)\big)\,,$$

since (10.49) implies[5] that $\phi(S)\leqslant X(\mathcal{H})\leqslant 1/q\leqslant a$, and hence

$$\frac{U(S,R)}{q} \geqslant \mathcal{J}(R) - \frac{X(\mathcal{H})}{2}\log\bigg(1+\frac{1}{X(\mathcal{H})q}\bigg) - O\big(X(\mathcal{H})\big)\,,$$

[5] Recall that, by Lemma 10.5.2, we may assume that $a\geqslant B/q$.

since $x \mapsto x \log\big(1 + \frac{1}{x}\big)$ is increasing. Combining this with (10.50), recalling that $v(\mathcal{H}) \leqslant L_3 \cdot X(\mathcal{H})$, by (10.46), and noting that $1 + \frac{1}{X(\mathcal{H})q} \leqslant \frac{2}{X(\mathcal{H})q}$, by (10.49), we obtain

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H}) \cdot \exp\left(-\mathcal{J}(R) - \frac{X(\mathcal{H})}{2} \log \frac{X(\mathcal{H})\sqrt{q}}{L_4}\right) \tag{10.52}$$

where $L_4 = C^{O(L_3)}$. Finally, observe that, by (10.46) and (10.48), we have

$$\sum_{\mathcal{H} \in \mathcal{H}_R^{(1)} : \, X(\mathcal{H}) = x} w(\mathcal{H}) \leqslant \exp\big(O\big(L_3 \cdot x\big)\big)$$

for any $x \in \mathbb{N}$. It follows that

$$\sum_{\substack{\mathcal{H} \in \mathcal{H}_R^{(1)} \\ X(\mathcal{H}) \leqslant 1/q}} \Lambda(\mathcal{H}) \leqslant e^{-\mathcal{J}(R)} \sum_{x = 1/(3\sqrt{q})}^{1/q} \exp\left(-\frac{x}{2} \log \frac{x\sqrt{q}}{L_4}\right) \sum_{\substack{\mathcal{H} \in \mathcal{H}_R^{(1)} \\ X(\mathcal{H}) = x}} w(\mathcal{H})$$

$$\leqslant e^{-\mathcal{J}(R)} \sum_{x = 1/(3\sqrt{q})}^{1/q} \exp\left(-\frac{x}{2} \log \frac{x\sqrt{q}}{L_4} + O\big(L_3 \cdot x\big)\right)$$
$$\leqslant \exp\left(-\mathcal{J}(R) + \frac{L_5}{\sqrt{q}}\right),$$

where $L_5 = L_4 \cdot e^{O(L_3)}$, since the summand decreases super-exponentially quickly once $x\sqrt{q}$ is larger than this. This completes the proof of Claim 1. □

If $\mathcal{H} \in \mathcal{H}_R^* \backslash \mathcal{H}_R^{(1)}$ then, by Lemmas 10.4.12 and 10.4.14, there exists a vertex $u \in V(G_{\mathcal{H}})$ satisfying either (10.28) or (10.29). The rest of the proof consists of bounding the contribution to (10.41) of hierarchies containing such a vertex. In order to simplify the argument, it will be convenient to first (in Claims 2 and 3) deal with those hierarchies in which either $v(\mathcal{H})$ or $X(\mathcal{H})$ is unusually large. We then (in Claims 4 and 5) consider the remaining hierarchies with an 'abnormal' vertex, i.e., one satisfying either (10.28) or (10.29).

We begin by considering hierarchies with unusually many vertices. Let us write $\mathcal{H}_R^{(2)}$ for the set of $\mathcal{H} \in \mathcal{H}_R^* \backslash \mathcal{H}_R^{(1)}$ such that

$$v(\mathcal{H}) \geqslant 8 \cdot s(\mathcal{H}) + \frac{4L_1}{q^{3/4}}. \tag{10.53}$$

For such hierarchies we will prove the following stronger bound.

Claim 2: $\displaystyle\sum_{\mathcal{H} \in \mathcal{H}_R^{(2)}} \Lambda(\mathcal{H}) \leqslant e^{-2/q}$.

Proof of Claim 2. The first step is to prove the following bounds,

$$|\mathrm{up}(\mathcal{H})| \leqslant \frac{L_1}{\sqrt{q}} \log\frac{1}{q} \qquad \text{and} \qquad \sum_{u \in \mathrm{up}(\mathcal{H})} \phi(R_u) \leqslant \frac{L_1}{q^{3/2}} \bigg(\log\frac{1}{q} \bigg)^2 \tag{10.54}$$

which replace those in (10.42), and hold for any $\mathcal{H} \in \mathcal{H}_R$. Both follow immediately from the upper bound on $b = \mathrm{long}(R)$ in (10.34), and the observation that in two consecutive steps of $\mathrm{up}(\mathcal{H})$, the semi-perimeter of the corresponding rectangles grows by at least $\delta/\sqrt{q}$.

Now, for each $\mathcal{H} \in \mathcal{H}_R^{(2)}$, consider the set $\mathcal{Y}(\mathcal{H})$ of edges $uv \in G_{\mathcal{H}}$ such that

$$\mathrm{short}(R_u) \leqslant \frac{B}{q}, \qquad N_{G_{\mathcal{H}}}(u) = \{v\} \qquad \text{and} \qquad |N_{G_{\mathcal{H}}}(v)| = 1\,.$$

We claim that

$$|\mathcal{Y}(\mathcal{H})| \;\geqslant\; v(\mathcal{H}) - 4s(\mathcal{H}) - |\mathrm{up}(\mathcal{H})| \;\geqslant\; \frac{v(\mathcal{H})}{2} + \frac{L_1}{q^{3/4}}\,. \tag{10.55}$$

To see this, recall that $\mathcal{H}$ has $s(\mathcal{H})$ seeds and $s(\mathcal{H}) - 1$ split vertices, and so there are at most $4s(\mathcal{H}) - 2$ vertices $u \in V(G_{\mathcal{H}})$ that are either seeds, or split-vertices, or have a single out-neighbour that is a seed or a split vertex. Moreover, $\mathcal{H} \in \mathcal{H}_R^*$ implies that every vertex $u \in V(G_{\mathcal{H}})$ with $\mathrm{short}(R_u) \geqslant B/q$ is in the trunk. The second inequality follows from (10.53) and (10.54).

We next claim that

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H}) \cdot C^{10v(\mathcal{H})} \cdot \exp\bigg(\frac{4L_1}{\sqrt{q}} \bigg(\log\frac{1}{q} \bigg)^2 - \frac{\delta^2 |\mathcal{Y}(\mathcal{H})|}{q^{1/4}} \bigg)\,. \tag{10.56}$$

The proof of this is similar to that of (10.43). Indeed, we obtain a slightly weaker bound in place of (10.44) by using (10.54) instead of (10.42), and (10.45) still holds, and the right-hand side is at most 1. Moreover, $\mathcal{H} \in \mathcal{H}_R^*$ implies that $\mathrm{long}(R_u) \leqslant 3e^{2B}/q$ for each edge $uv \in \mathcal{Y}(\mathcal{H})$, and therefore

$$Q(R_u, R_v) \geqslant \frac{\delta g(3e^{2B})}{q^{1/4}} \geqslant \frac{\delta^2}{q^{1/4}}\,,$$

for each such edge, since $d(R_u, R_v) \geqslant f(R_u) \geqslant \delta\sqrt{\mathrm{short}(R_u)} \geqslant \delta q^{-1/4}$, by Definition 10.4.2 and (10.14), and since $\delta = \delta(B)$ was chosen sufficiently small. Plugging these bounds into the definition of $\Lambda(\mathcal{H})$, we obtain (10.56).

By (10.55), it follows that

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H}) \cdot \exp\bigg(- \frac{\delta^3 |\mathcal{Y}(\mathcal{H})|}{q^{1/4}} \bigg)\,,$$

and hence, by (10.55) and Lemma 10.4.9, and since $\phi(R) \leqslant 1/q^2$, we obtain

$$\begin{aligned}\sum_{\mathcal{H}\in\mathcal{H}_R^{(2)}} \Lambda(\mathcal{H}) &\leqslant \sum_{y \geqslant L_1 q^{-3/4}} \exp\left(-\frac{\delta^3 y}{q^{1/4}}\right) \sum_{M=1}^{2y} \sum_{N=M}^{2y} \sum_{\substack{\mathcal{H}\in\mathcal{H}_R^{(2)}(N,M)\\ |\mathcal{Y}(\mathcal{H})|=y}} w(\mathcal{H}) \\ &\leqslant \sum_{y \geqslant L_1 q^{-3/4}} \exp\left(-\frac{\delta^3 y}{q^{1/4}}\right) \sum_{M=1}^{2y} \sum_{N=M}^{2y} \exp\left(16\Big(N + 2M\log\frac{1}{q}\Big)\right) \\ &\leqslant \sum_{y \geqslant L_1 q^{-3/4}} \exp\left(-\frac{\delta^4 y}{q^{1/4}}\right) \leqslant e^{-2/q},\end{aligned}$$

as required. □

We will next deal with those hierarchies for which $X(\mathcal{H})$ is unusually large. To be precise, let us define $\mathcal{H}_R^{(3)}$ to be the set of $\mathcal{H} \in \mathcal{H}_R^* \setminus \big(\mathcal{H}_R^{(1)} \cup \mathcal{H}_R^{(2)}\big)$ such that

$$X(\mathcal{H}) \geqslant \frac{1}{L_2 q^{3/4}}\,. \tag{10.57}$$

For this class of hierarchies we will prove the following bound.

Claim 3: $\displaystyle\sum_{\mathcal{H}\in\mathcal{H}_R^{(3)}} \Lambda(\mathcal{H}) \leqslant \exp\left(-\mathcal{J}(R) - \frac{1}{q^{3/4}}\right) + e^{-2/q}.$

Proof of Claim 3. Let $\mathcal{H} \in \mathcal{H}_R^{(3)}$, and observe that

$$v(\mathcal{H}) \leqslant 8 \cdot s(\mathcal{H}) + \frac{4L_1}{q^{3/4}} \qquad \text{and} \qquad \sum_{u\in \mathrm{up}(\mathcal{H})} \phi(R_u) \leqslant \frac{L_1}{q^{3/2}}\left(\log\frac{1}{q}\right)^2, \tag{10.58}$$

where the first inequality holds since $\mathcal{H} \notin \mathcal{H}_R^{(2)}$, and the second holds for any $\mathcal{H} \in \mathcal{H}_R$, by (10.54). We will repeat the proof of Claim 1, using the bounds (10.58) instead of (10.42).

Indeed, note (cf. (10.44)) that

$$\prod_{N_{G_\mathcal{H}}(u)=\{v\}} \exp\big(4\phi(R_u)q\big) \leqslant C^{v(\mathcal{H})} \exp\left(\frac{4L_1}{\sqrt{q}}\left(\log\frac{1}{q}\right)^2\right), \tag{10.59}$$

and hence, using (10.45), we obtain

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H}) \cdot C^{10v(\mathcal{H})} \cdot \exp\left(-\frac{X(\mathcal{H})}{5}\log\frac{1}{q} + \frac{4L_1}{\sqrt{q}}\left(\log\frac{1}{q}\right)^2\right).$$

Now, note that, by Observation 10.4.10 and the bounds (10.57) and (10.58), we have

$$X(\mathcal{H}) \geqslant \frac{1}{L_2} \cdot \max\left\{\frac{s(\mathcal{H})}{q^{1/4}}, \frac{1}{q^{3/4}}\right\} \geqslant \frac{v(\mathcal{H})}{L_3}\,,$$

It follows, exactly as in the proof of Claim 1 (cf. the proof of (10.47)), that

$$\sum_{\mathcal{H} \in \mathcal{H}_R^{(3)} :\, X(\mathcal{H}) > 1/q} \Lambda(\mathcal{H}) \leqslant e^{-2/q}.$$

We will therefore assume from now on that $X(\mathcal{H}) \leqslant 1/q$.

We now simply repeat the remainder of the proof of Claim 1, using the bounds $1/(L_2 q^{3/4}) \leqslant X(\mathcal{H}) \leqslant 1/q$, to obtain

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H}) \cdot \exp\left(-\mathcal{J}(R) - \frac{X(\mathcal{H})}{2} \log \frac{X(\mathcal{H})\sqrt{q}}{L_4} \right)$$

for each $\mathcal{H} \in \mathcal{H}_R^{(3)}$, and hence

$$\begin{aligned} \sum_{\substack{\mathcal{H} \in \mathcal{H}_R^{(3)} \\ X(\mathcal{H}) \leqslant 1/q}} \Lambda(\mathcal{H}) &\leqslant e^{-\mathcal{J}(R)} \sum_{x = 1/(L_2 q^{3/4})}^{1/q} \exp\left(-\frac{x}{2} \log \frac{x\sqrt{q}}{L_4} \right) \sum_{\substack{\mathcal{H} \in \mathcal{H}_R^{(3)} \\ X(\mathcal{H}) = x}} w(\mathcal{H}) \\ &\leqslant e^{-\mathcal{J}(R)} \sum_{x = 1/(L_2 q^{3/4})}^{1/q} \exp\left(-\frac{x}{2} \log \frac{x\sqrt{q}}{L_4} + O\big(L_3 \cdot x\big) \right) \\ &\leqslant \exp\left(-\mathcal{J}(R) - \frac{1}{q^{3/4}} \right), \end{aligned}$$

as claimed. □

We are now ready to deal with those hierarchies that travel 'far from the diagonal', i.e., that contain a vertex u satisfying either (10.28) or (10.29). We will first consider the (easier) case in which u is in the upper trunk, i.e.,

$$\text{short}(R_u) \geqslant \frac{B}{q} \qquad \text{and} \qquad \text{long}(R_u) \geqslant 4 \cdot \text{short}(R_u). \tag{10.60}$$

Let $\mathcal{H}_R^{(4)}$ be the set of hierarchies $\mathcal{H} \in \mathcal{H}_R^* \backslash \bigcup_{i=1}^3 \mathcal{H}_R^{(i)}$ containing a vertex u such that (10.60) holds. For these hierarchies we will prove the following bound.

Claim 4: $\displaystyle\sum_{\mathcal{H} \in \mathcal{H}_R^{(4)}} \Lambda(\mathcal{H}) \leqslant \exp\left(-\frac{2\lambda}{q} - \frac{2}{q^{3/4}} \right).$

Proof of Claim 4. Given $\mathcal{H} \in \mathcal{H}_R^{(4)}$ and $u \in V(\mathcal{H})$ satisfying (10.60), we claim that

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H}) \cdot \exp\left(-\mathcal{J}(R_u) + \frac{L_3}{q^{3/4}} \right). \tag{10.61}$$

To prove this, we will repeat the proof of Claim 1, with some minor changes. First, note that (10.49) holds, and moreover

$$\frac{s(\mathcal{H})}{L_1 q^{1/4}} \leqslant X(\mathcal{H}) \leqslant \frac{1}{L_2 q^{3/4}}\,, \tag{10.62}$$

the first holding by Observation 10.4.10, and the second since $\mathcal{H} \notin \mathcal{H}_R^{(3)}$. Now, applying Lemma 10.4.15 to $\mathcal{H}_u$, the sub-hierarchy of $\mathcal{H}$ rooted at u, and using (10.59) instead of (10.44), we obtain (cf. the proof of (10.50))

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H}) \cdot C^{10v(\mathcal{H})} \cdot \exp\left(- \frac{U(S,R_u)}{q} - \frac{X(\mathcal{H})}{5} \log\frac{1}{q} + \frac{4L_1}{\sqrt{q}} \left(\log\frac{1}{q}\right)^2 \right)$$

for some pod S with $\phi(S) \leqslant X(\mathcal{H}_u) \leqslant X(\mathcal{H})$. Using (10.58) again (this time to bound $v(\mathcal{H})$ from above), and continuing to follow the proof of Claim 1, we obtain

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H}) \cdot \exp\left(-\mathcal{J}(R_u) - \frac{X(\mathcal{H})}{2} \log\frac{X(\mathcal{H}) q^{3/5}}{L_4} + \frac{L_2}{q^{3/4}} \right)$$

instead of (10.52), which implies (10.61).

Now, let $c = \mathrm{short}(R_u)$, and observe that

$$\mathcal{J}(R_u) \;\geqslant\; \frac{2}{q}\int_0^{cq} g(z)\,\mathrm{d}z + 3cg(cq) \;\geqslant\; \frac{2}{q}\int_0^{\infty} g(z)\,\mathrm{d}z + 2cg(cq)\,,$$

by (10.60) and (10.10), and since $c \geqslant B/q$, and that

$$cg(cq) \;\geqslant\; \frac{c}{2}\cdot e^{-2cq} \;\geqslant\; \frac{1}{4q^{3/4}} \log\frac{1}{q}\,,$$

where the first inequality holds since $c \geqslant B/q$, and the second since $4c \leqslant \mathrm{long}(R) \leqslant (1/(2q))\log(1/q)$. Combining this with (10.61), it follows that

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H}) \cdot \exp\left(-\frac{2\lambda}{q} - \frac{L_3}{q^{3/4}} \right).$$

Hence, recalling that $v(\mathcal{H}) \leqslant 5L_1/q^{3/4}$ and $s(\mathcal{H}) \leqslant q^{-1/2}$ for every $\mathcal{H} \in \mathcal{H}_R^{(4)}$, by (10.58) and (10.62), and applying Lemma 10.4.9, we obtain

$$\begin{aligned} \sum_{\mathcal{H}\in\mathcal{H}_R^{(4)}} \Lambda(\mathcal{H}) \;&\leqslant\; \exp\left(-\frac{2\lambda}{q} - \frac{L_3}{q^{3/4}} \right) \sum_{M=1}^{1/q^{1/2}} \sum_{N=M}^{5L_1/q^{3/4}} \sum_{\mathcal{H}\in\mathcal{H}_R^{(4)}(N,M)} w(\mathcal{H}) \\ &\leqslant \exp\left(-\frac{2\lambda}{q} - \frac{L_3}{q^{3/4}} \right) \sum_{M=1}^{1/q^{1/2}} \sum_{N=M}^{5L_1/q^{3/4}} \exp\left(\frac{L_2}{q^{3/4}} \right) \leqslant \exp\left(-\frac{2\lambda}{q} - \frac{2}{q^{3/4}} \right), \end{aligned}$$

as claimed. □

Finally, we come to most technically challenging family of hierarchies: those which contain a vertex $u \in V(G_{\mathcal{H}})$ such that

$$\mathrm{short}(R_u) \leqslant \frac{B}{q} \qquad \text{and} \qquad \mathrm{long}(R_u) \geqslant 2L_1 \cdot \mathrm{short}(R_u)\,. \tag{10.63}$$

Let $\mathcal{H}_R^{(5)}$ denote the set of hierarchies $\mathcal{H} \in \mathcal{H}_R^* \backslash \bigcup_{i=1}^4 \mathcal{H}_R^{(i)}$ containing a vertex u such that (10.63) holds. For this final class of hierarchies we will prove the following bound.

Claim 5: $\displaystyle\sum_{\mathcal{H} \in \mathcal{H}_R^{(5)}} \Lambda(\mathcal{H}) \leqslant \exp\left(-\mathcal{J}(R) - \frac{1}{q^{3/4}}\right).$

Proof of Claim 5. Given a hierarchy $\mathcal{H} \in \mathcal{H}_R^{(5)}$, let $u \in V(\mathcal{H})$ be a vertex satisfying (10.63) with $\mathrm{long}(R_u)$ maximal, and set $c = \mathrm{short}(R_u)$ and $d = \mathrm{long}(R_u)$. We will prove that

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H}) \cdot \exp\left(-\mathcal{J}(R) - 3Cd + X(\mathcal{H}) \log \frac{1}{X(\mathcal{H}) q^{3/4}}\right), \tag{10.64}$$

from which the claim will follow easily, using Lemma 10.4.9.

In order to prove (10.64), we will need various bounds on c, d, $h(\mathcal{H})$, $v(\mathcal{H})$ and $X(\mathcal{H})$. Note first that $\mathcal{H}$ does not contain a vertex satisfying (10.60) since $\mathcal{H} \notin \mathcal{H}_R^{(4)}$. It follows, by Lemmas 10.4.13 and 10.4.14, and since $\mathcal{H} \notin \mathcal{H}_R^{(1)}$, that

$$\frac{L_2}{\sqrt{q}} \leqslant h(\mathcal{H}) \leqslant L_1 q^{1/4} d\,. \tag{10.65}$$

Indeed, the lower bound holds since $\mathcal{H} \notin \mathcal{H}_R^{(1)}$ implies that one of the inequalities in (10.42) must fail to hold, and by Lemma 10.4.14, it must be the bound on $h(\mathcal{H})$. The upper bound then follows by Lemma 10.4.13, and by our choice of u (i.e., with $\mathrm{long}(R_u)$ maximal).

We next claim that

$$v(\mathcal{H}) \leqslant d, \qquad c \leqslant \frac{1}{q} \qquad \text{and} \qquad \frac{1}{q^{3/4}} \leqslant d \leqslant \frac{B}{2q}\,. \tag{10.66}$$

Indeed, the lower bound on d follows immediately from (10.65), since $L_1 \leqslant L_2$. To prove the other bounds, recall first that (since $\mathcal{H} \in \mathcal{H}_R^*$) the rectangle R_u does not satisfy (10.38) or (10.39). Since $c \leqslant B/q$, by (10.63), it follows that $d \leqslant 3e^{2B}/q$, and hence

$$c \leqslant \frac{d}{L_1} \leqslant \frac{3e^{2B}}{L_1 q} \leqslant \frac{1}{q}\,.$$

Now, since R_u does not satisfy (10.39), it follows that $d \leqslant B/(2q)$, as claimed. Finally, to prove the bound on $v(\mathcal{H})$, recall that

$$\frac{m(\mathcal{H})}{3\sqrt{q}} \leqslant X(\mathcal{H}) \leqslant \frac{1}{L_2 q^{3/4}}, \tag{10.67}$$

by the definition of large seeds, and since $\mathcal{H} \notin \mathcal{H}_R^{(3)}$. Hence, by (10.65) and Observation 10.4.11, we obtain

$$v(\mathcal{H}) \leqslant 2 \cdot h(\mathcal{H}) \cdot m(\mathcal{H}) \leqslant 6L_1 q^{3/4} d \cdot X(\mathcal{H}) \leqslant d,$$

as claimed.

We now apply Lemma 10.4.16 to obtain two pods $S_1 \subset R_u$ and $R_u \subset S_2 \subset R$, such that

$$\phi(S_1) + \phi(S_2) - \phi(R_u) \leqslant X(\mathcal{H}) \tag{10.68}$$

and

$$\sum_{N_{G_{\mathcal{H}}}(v)=\{w\}} U(R_w, R_v) \geqslant U(S_1, R_u) + U(S_2, R) - 2s(\mathcal{H}) q g(\sqrt{q}).$$

Let $S_1 \subset S \subset S_2$ be a rectangle with

$$\dim(S) = \dim(S_1) + \dim(S_2) - \dim(R_u), \tag{10.69}$$

so $\phi(S) \leqslant X(\mathcal{H})$, by (10.68), and moreover $U(S_1, R_u) \geqslant U(S, S_2)$, since $g(z)$ is decreasing.

Recalling that $U(R_w, R_v) \leqslant q \cdot Q(R_w, R_v)$, by Lemma 10.2.8, and that $g(\sqrt{q}) \leqslant \log(1/q)$, by (10.4), it follows that

$$\sum_{N_{G_{\mathcal{H}}}(v)=\{w\}} Q(R_w, R_v) \geqslant \frac{1}{q}\Big(U(S, S_2) + U(S_2, R)\Big) - 2s(\mathcal{H}) \log\frac{1}{q}.$$

Hence, using (10.45), (10.59) and (10.62), we obtain

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H}) \cdot C^{10v(\mathcal{H})} \cdot \exp\left(-\frac{1}{q}\Big(U(S, S_2) + U(S_2, R)\Big) - \frac{X(\mathcal{H})}{4}\log\frac{1}{q} + \frac{1}{q^{3/4}}\right). \tag{10.70}$$

It only remains to bound $U(S, S_2)$ and $U(S_2, R)$; controlling $U(S, S_2)$ will take some work, but we obtain a suitable bound on $U(S_2, R)$ simply by applying Lemma 10.2.5. Indeed, by (10.66), (10.67) and (10.69), we have

$$\mathrm{long}(S_2) \leqslant \phi(S) + \mathrm{long}(R_u) \leqslant X(\mathcal{H}) + d \leqslant \frac{1}{q^{3/4}} + \frac{B}{2q} \leqslant a = \mathrm{short}(R), \tag{10.71}$$

and therefore, setting $s_2 = \text{short}(S_2)$ and $t_2 = \text{long}(S_2)$, we may apply Lemma 10.2.5, which gives

$$\frac{U(S_2,R)}{q} \geqslant (t_2 - s_2)g(t_2 q) + \frac{2}{q}\int_{t_2 q}^{aq} g(z)\,\mathrm{d}z + (b-a)g(aq)\,. \qquad (10.72)$$

As noted above, we will have to work harder to obtain a suitable bound on $U(S,S_2)$; in particular, the bound we obtain will depend on whether or not $\text{long}(S) \leqslant \text{short}(S_2)$.

Case 1: $\text{long}(S) \leqslant \text{short}(S_2)$.

This case is also straightforward, since we may apply Lemma 10.2.7, which gives

$$\frac{U(S,S_2)}{q} \geqslant \frac{2}{q}\int_0^{s_2 q} g(z)\,\mathrm{d}z + (t_2 - s_2)g(s_2 q) - \frac{X(\mathcal{H})}{2}\log\frac{2}{X(\mathcal{H})q} - O\big(X(\mathcal{H})\big)\,, \qquad (10.73)$$

where we used the inequalities $\phi(S) \leqslant X(\mathcal{H}) \leqslant 1/q$. To show that this is sufficient to deduce (10.64), we will use Lemma 10.2.10. Indeed, observe that

$$L_1 \cdot \text{short}(S_2) \leqslant L_1\big(c + X(\mathcal{H})\big) \leqslant d \leqslant \text{long}(S_2) \leqslant \frac{B}{q}\,,$$

where the first inequality follows from (10.68) (since $R_u \subset S_2$), the second follows since $2L_1 \cdot c \leqslant d$, by (10.63), and $L_2 \cdot X(\mathcal{H}) \leqslant q^{-3/4} \leqslant d$, by (10.66) and (10.67), the third since $R_u \subset S_2$, and the last by (10.71). Since $d \leqslant t_2$ it follows, by Lemma 10.2.10, that

$$\frac{2}{q}\int_{s_2 q}^{t_2 q} g(z)\,\mathrm{d}z \leqslant (t_2 - s_2)\big(g(s_2 q) + g(t_2 q)\big) - 4Cd\,.$$

Combining this with (10.72) and (10.73), it follows that

$$\frac{1}{q}\Big(U(S,S_2) + U(S_2,R)\Big) \geqslant \mathcal{J}(R) + 4Cd - \frac{X(\mathcal{H})}{2}\log\frac{2}{X(\mathcal{H})q} - O\big(X(\mathcal{H})\big)\,.$$

Hence, by (10.70), and recalling from (10.66) that $d \geqslant \max\{v(\mathcal{H}), q^{-3/4}\} \geqslant X(\mathcal{H})$, we obtain

$$\Lambda(\mathcal{H}) \leqslant w(\mathcal{H}) \cdot \exp\left(-\mathcal{J}(R) - 3Cd + \frac{X(\mathcal{H})}{2}\log\frac{1}{X(\mathcal{H})\sqrt{q}}\right)\,,$$

which is slightly stronger than (10.64).

Case 2: $\text{long}(S) > \text{short}(S_2)$.

This case is a bit more tricky, as the easiest path from S to S_2 does not reach the diagonal, see Figure 10.4. As a consequence, we cannot apply

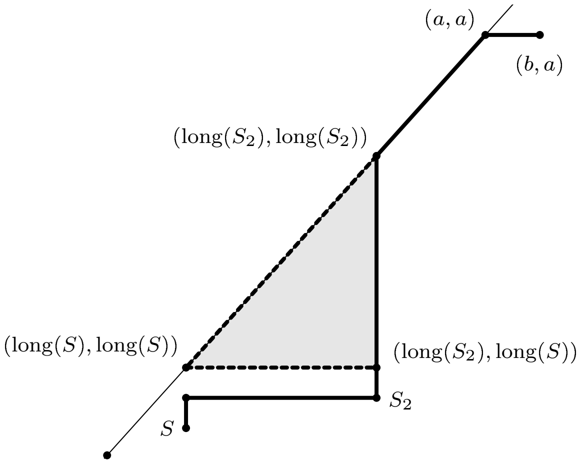

Figure 10.4 – The easiest path via which the rectangle S can grow first to S_2, and then to R, in Case 2 of Claim 5. In the proof, the lower two thick segments are replaced with the lower dashed one, and Lemma 10.2.10 is applied to the shaded triangle.

Lemma 10.2.7 directly to bound $U(S,S_2)$, nor can we apply Lemma 10.2.10 directly to the dimensions of S_2. Instead, we observe that, setting $t := \mathrm{long}(S)$, we have

$$\frac{U(S,S_2)}{q} \geqslant (t_2 - t)g(s_2 q) \geqslant (t_2 - t)g(tq)\,,$$

by Lemma 10.2.6, and since $g(z)$ is decreasing. Combining this with (10.72), we obtain

$$\frac{1}{q}\Big(U(S,S_2) + U(S_2,R)\Big) \geqslant \mathcal{J}(R) - \frac{2}{q}\int_0^{t_2 q} g(z)\,\mathrm{d}z + (t_2 - s_2)g(t_2 q) + (t_2 - t)g(tq)\,.$$

We will now apply Lemma 10.2.10 to the pair (t,t_2). Indeed, we have

$$L_1 \cdot \mathrm{long}(S) \leqslant L_1 \cdot X(\mathcal{H}) \leqslant d \leqslant \mathrm{long}(S_2) \leqslant \frac{B}{q}\,,$$

where the first inequality follows since $\phi(S) \leqslant X(\mathcal{H})$, and the others follow as in Case 1. By Lemma 10.2.10, and since $d \leqslant t_2$, it follows that

$$\frac{2}{q}\int_{tq}^{t_2 q} g(z)\,\mathrm{d}z \;\leqslant\; (t_2 - t)\big(g(tq) + g(t_2 q)\big) - 4Cd\,.$$

Note also that, by integrating (10.4), we have

$$\frac{2}{q}\int_0^{tq} g(z)\,\mathrm{d}z \;\leqslant\; t\log\frac{1}{tq} + O(t)\,.$$

Hence, recalling that $s_2 < t \leqslant \phi(S) \leqslant X(\mathcal{H})$, it follows that

$$\frac{1}{q}\Big(U(S,S_2) + U(S_2,R)\Big) \geqslant \mathcal{J}(R) + 4Cd - X(\mathcal{H})\log\frac{1}{X(\mathcal{H})q} - O\big(X(\mathcal{H})\big)\,.$$

Hence, by (10.70), and since $d \geqslant \max\{v(\mathcal{H}), q^{-3/4}\} \geqslant X(\mathcal{H})$, by (10.66), we obtain (10.64).

It now only remains to deduce the claim from (10.64); we do so using Lemma 10.4.9. Indeed, recalling that $q^{-1/5}s(\mathcal{H}) \leqslant X(\mathcal{H}) \leqslant 1/(L_2 q^{3/4}) \leqslant \max\{v(\mathcal{H}), q^{-3/4}\} \leqslant d \leqslant B/q$, by (10.62) and (10.66), we obtain

$$\begin{aligned}
&\sum_{\mathcal{H} \in \mathcal{H}_R^{(5)}} \Lambda(\mathcal{H}) \\
&\leqslant e^{-\mathcal{J}(R)} \sum_{x=1}^{1/(L_2 q^{3/4})} \left(\frac{1}{q^{3/4} x} \right)^x \sum_{d=q^{-3/4}}^{B/q} e^{-3Cd} \sum_{M=1}^{q^{1/5} d} \sum_{N=M}^{d} \sum_{\mathcal{H} \in \mathcal{H}_R^{(5)}(N,M)} w(\mathcal{H}) \\
&\leqslant e^{-\mathcal{J}(R) + q^{-3/4}} \sum_{d=q^{-3/4}}^{B/q} e^{-3Cd} \sum_{M=1}^{q^{1/5} d} \sum_{N=M}^{d} \exp\Big(O\big(N + M \log(1/q)\big)\Big) \\
&\leqslant e^{-\mathcal{J}(R) + q^{-3/4}} \sum_{d=q^{-3/4}}^{B/q} e^{-2Cd} \leqslant \exp\left(-\mathcal{J}(R) - \frac{1}{q^{3/4}} \right),
\end{aligned}$$

as required. This concludes the proof of Claim 5. $\square$

Now, combining Claims 1–5, it follows that

$$\sum_{\mathcal{H} \in \mathcal{H}_R^*} \Lambda(\mathcal{H}) \leqslant 3 \cdot \exp\left(-\min\left\{ \frac{2\lambda}{q} + \frac{2}{q^{3/4}}, \mathcal{J}(R) - \frac{L_5}{\sqrt{q}} \right\} \right).$$

As was observed before the proof (see the discussion leading up to (10.41)), this completes the proof of Theorem 10.5.1. $\square$

We are finally ready to deduce Theorem 10.0.1 from Theorem 10.5.1.

Proof of Theorem 10.0.1. Recall that the upper bound was proved in [187]; we will prove the lower bound. Let $n \in \mathbb{N}$ be sufficiently large and set

$$q := \frac{\lambda}{\log n} - \frac{4e^4 + L_6}{(\log n)^{3/2}},$$

and note that the same is satisfied by $p = q + \Theta(q^2)$ with a slightly smaller constant. Suppose that $[n]^2$ is (internally) filled. Then by the Aizenman–Lebowitz lemma there exists an internally filled rectangle $R \subset [n]^2$ with

$$\frac{1}{4q} \log \frac{1}{q} \leqslant \operatorname{long}(R) \leqslant \frac{1}{2q} \log \frac{1}{q}.$$

There are at most $n^2 (\log n)^3$ rectangles satisfying those conditions, and each one satisfies the conditions of Theorem 10.5.1. Hence, by the union bound and Lemma 10.2.9 we have

$$\mathbb{P}_p\big([A] = [n]^2\big) \leqslant n^2 (\log n)^3 \exp\left(-\frac{2\lambda}{q} + \frac{4e^4 + L_6}{\sqrt{q}} \right) \to 0$$

as $n \to \infty$, as required. □

10.6 Open problems

The most obvious problem suggested by Theorem 10.0.1 is to determine even more precise bounds on $p_c\big([n]^2, 2\big)$. By a theorem of Balogh and Bollobás [30], it is known that the 'critical window' in which the probability of percolation increases from $o(1)$ to $1 - o(1)$ has size at most $(\log n)^{-2+o(1)}$, and it is therefore natural to make the following conjecture.

Conjecture 10.6.1. *There exists a constant $\mu > 0$ such that*

$$p_c\big([n]^2, 2\big) = \frac{\pi^2}{18 \log n} - \frac{\mu + o(1)}{(\log n)^{3/2}}$$

as $n \to \infty$.

Another natural direction for future research would be to extend the results of this chapter to higher dimensions. The following conjecture was made by Uzzell [354], who also established the upper bound.

Conjecture 10.6.2 (Conjecture 7.1 of [354])**.**

$$p_c\left([n]^d, r\right) = \left(\frac{\lambda(d,r)}{\log_{(r-1)} n} - \frac{\Theta(1)}{\big(\log_{(r-1)} n\big)^{3/2}}\right)^{d-r+1}.$$

As a first step, it would be interesting to determine whether or not this conjecture holds in the case $r = 2$. In particular, one might hope that the conjecture in this case would follow from a suitable generalization of the technique used in this chapter.

However, perhaps the most interesting avenue for further research would be to prove corresponding 'sharp' and 'sharper' thresholds for other two-dimensional models, cf. the discussion of $\mathcal{U}$-bootstrap percolation in the Introduction. It would be very interesting (and, most likely, very challenging) to determine a sharp threshold for all families with polylogarithmic critical probability, or a 'sharper' threshold for either some large class of models (e.g., that studied in [125], or a corresponding class of 'unbalanced' models), or for other specific interesting examples, such as the Duarte model (see [69]). The problem in higher dimensions is also extremely interesting, but much more difficult, and proving even much weaker bounds on the critical probability for general $\mathcal{U}$-bootstrap models (see, for example, [280, Conjecture 1.6]) is an important open problem.

Appendix

10.A Proof of Lemmas 10.3.3 and 10.3.4

In this appendix we complete the proofs of Lemmas 10.3.3 and 10.3.4, by dealing with the cases that were omitted from the sketch proofs given in Section 10.3. The details are somewhat tedious but, for the convenience of the diligent (or sceptical) reader, we will attempt to spell everything out slowly and carefully.

For the entire Section 10.A we fix a rectangle $R = [(0,0),(a-1,b-1)]$. Recall from Definition 10.3.1 that if $S \subset R$ is a rectangle, then we write

$$z(S,R) = |Z(S,R)| = \left|\left\{\mathbf{d} \in \mathcal{I} \,:\, B_{\mathbf{d}}(S,R) \neq \varnothing\right\}\right|,$$

where $\mathcal{I} = \{(-1,0),(1,0),(0,-1),(0,1)\}$ is the set of directions and $B_{\mathbf{d}}(S,R)$ is the buffer in direction $\mathbf{d}$. Recall also that if $\mathbf{x} \in \{0,1\}^{\mathcal{I}}$ and $S \subset R$ is a rectangle, then

$$x = x'_{(1,0)} + x'_{(-1,0)} \qquad \text{and} \qquad y = x'_{(0,1)} + x'_{(0,-1)}\,, \tag{10.74}$$

where $\mathbf{x}' := \mathbf{x} \cdot \mathbb{1}_{Z(S,R)}$. For convenience, let us begin by restating the two key lemmas.

Lemma 10.A.1 (Lemma 10.3.3 restated)**.** *Let* $\mathbf{x} \in \{0,1\}^{\mathcal{I}}$ *and* $S \subset R$ *be a rectangle with* $\dim S = (a-s, b-t)$ *and set* $z = z(S,R)$. *If*

$$L_1 \leqslant \text{short}(R) \leqslant \frac{B}{q} \qquad \textit{and} \qquad \text{long}(R) \leqslant \frac{3e^{2B}}{q}\,, \tag{10.75}$$

and $s,t \leqslant 4\delta\sqrt{\text{short}(R)}$*, then*

$$\mathbb{P}_p\big(D_1^{\mathbf{x}}(S,R)\big) \leqslant C^z \left(\frac{C}{\sqrt{a}}\right)^y \left(\frac{C}{\sqrt{b}}\right)^x \exp\big(-sg(bq) - tg(aq)\big)\,.$$

Lemma 10.A.2 (Lemma 10.3.4 restated)**.** *Let* $\mathbf{x} \in \{0,1\}^{\mathcal{I}}$ *and* $S \subset R$ *be a rectangle with* $\dim S = (a-s, b-t)$ *and set* $z = z(S,R)$. *If*

$$\text{short}(R) > \frac{B}{q} \qquad \textit{and} \qquad \text{long}(R) \leqslant \frac{1}{2q}\log\frac{1}{q} \tag{10.76}$$

and $s,t \leqslant \dfrac{4\delta}{\sqrt{q}} \cdot \exp\big(\text{short}(R)\cdot q\big)$*, then*

$$\mathbb{P}_p\big(D_2^{\mathbf{x}}(S,R)\big) \leqslant \Big(Ce^{\text{short}(R)q}\Big)^z \big(C\sqrt{q}e^{-aq}\big)^y \Big(C\sqrt{q}e^{-bq}\Big)^x \exp\big(-sg(bq) - tg(aq)\big)\,.$$

We begin with a straightforward technical lemma, which will be required in both proofs.

Lemma 10.A.3. *Let* $15 \leqslant D \leqslant 4\delta/q$, $a, b \leqslant (1/(2q)) \log(1/q)$, *and* $s, t \leqslant D$ *with* $t \leqslant \min\{a, b\}$. *Then, for any* $x, y \in \{0, 1, 2\}$, *we have*

$$\exp\big(-3Dg(tq)\big) \leqslant \left(\frac{1}{\sqrt{a}}\right)^y \left(\frac{1}{\sqrt{b}}\right)^x \exp\big(-sg(bq) - tg(aq)\big)$$

and

$$\exp\big(-3Dg(tq)\big) \leqslant \big(\sqrt{q}e^{-aq}\big)^y \left(\sqrt{q}e^{-bq}\right)^x \exp\big(-sg(bq) - tg(aq)\big).$$

Proof. Observe first that $-3g(tq) \leqslant \log(tq) \leqslant \log(Dq)$, by (10.4) and since $t \leqslant D \leqslant 4\delta/q$. Noting that $D \log(Dq)$ is decreasing in D, it follows that

$$\exp\big(-Dg(tq)\big) \leqslant \exp\big(5\log(15q)\big) \leqslant q^4 \leqslant \min\left\{\frac{1}{ab},\, q^2 e^{-2q(a+b)}\right\},$$

where in the last step we used the bound $a, b \leqslant (1/(2q)) \log(1/q)$. Moreover, we have

$$\exp\big(-2Dg(tq)\big) \leqslant \exp\big(-sg(bq) - tg(aq)\big),$$

since $s, t \leqslant D$ and $t \leqslant \min\{a, b\}$, and recalling that g is decreasing. Combining these two inequalities we obtain the two claimed bounds. $\square$

Proof of Lemma 10.A.1. Recall that $D_1^{\mathbf{x}}(S, R)$ denotes the event that

$$\big[S \cup \big(A \cap R \backslash S_{\blacksquare}^{\mathbf{x}}\big)\big] = R.$$

Let a, b satisfy (10.75), and for each x, y, z and each $s, t \leqslant 4\delta\sqrt{\min\{a, b\}}$, set

$$F^{x,y,z}(s, t) := C^z \left(\frac{C}{\sqrt{a}}\right)^y \left(\frac{C}{\sqrt{b}}\right)^x \exp\big(-sg(bq) - tg(aq)\big).$$

We will prove, by induction on the pair $(s + t, -(x + y))$, that

$$\mathbb{P}_p\big(D_1^{\mathbf{x}}(S, R)\big) \leqslant F^{x,y,z}(s, t) \tag{10.77}$$

for every $0 \leqslant s, t \leqslant 4\delta\sqrt{\min\{a, b\}}$ and $\mathbf{x} \in \{0, 1\}^{\mathcal{I}}$, and every $S \subset R$ with $\dim(S) = (a - s, b - t)$, where x and y are as defined in (10.74), and $z = z(S, R)$.

The base of the induction, the case $\min\{s, t\} = 0$, was dealt with in Section 10.3, so let us fix $\mathbf{x} \in \{0, 1\}^{\mathcal{I}}$ and $S \subset R$ with $\dim(S) = (a - s, b - t)$, and assume that (10.77) holds for all smaller values of the pair $(s+t, -(x+y))$ in lexicographical order.

Note first that, since $\text{short}(R) \geqslant L_1$, the function $F^{x,y,z}(s, t)$ is increasing in z and decreasing in x, y, s and t. Note also that we may assume, without

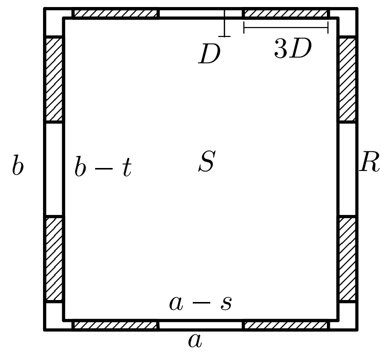

Figure 10.5 – The 8 rectangles, one of which must have a 'short' double gap if $\phi(T)$ is large. Each has width at most D and length $3D+1$. Note that $9D \leqslant \text{short}(R)$, so the rectangles do not overlap.

loss of generality, that $\mathbf{x} = \mathbf{x} \cdot \mathbb{1}_{Z(S,R)}$ (i.e., that $\mathbf{x_d} = 0$ whenever the buffer $B_{\mathbf{d}}(S,R)$ is empty), since neither side of the inequality (10.77) depends on the value of $\mathbf{x_d}$ if $\mathbf{d} \notin Z(S,R)$.

We partition into cases, depending on whether or not $z = x + y$.

Case 1: $z = x + y$, i.e., all of the non-empty buffers are included in $S^{\mathbf{x}}_{\square}$.

The key observation in this case is that if the event $D_1^{\mathbf{x}}(S,R)$ holds, then there exists a rectangle $T \subset R$ such that

$$\big[A \cap T \backslash S^{\mathbf{x}}_{\blacksquare}\big] = T \qquad \text{and} \qquad T \cap S^{\mathbf{x}}_{\blacksquare} \neq \emptyset \tag{10.78}$$

(see Figure 10.6a). In Section 10.3 we assumed that $\phi(T) \leqslant 36 f(R) = 36\delta\sqrt{\min\{a,b\}}$, so let us begin by dealing with the other case. To do so, the key observation is that, since T is internally filled by the infected sites in $T \backslash S^{\mathbf{x}}_{\blacksquare}$, one of the eight rectangles in Figure 10.5 has no double gap crossing it in the 'short' direction. To be more precise, set $D := 4\delta\sqrt{\min\{a,b\}}$ (so, in particular, $s,t \leqslant D$ and $9D \leqslant \min\{a,b\}$) and consider the $(3D+1) \times (\leqslant t)$ rectangle

$$\big([D,4D] \times [0,t-1]\big) \cap \big(R \backslash S\big)\,,$$

which is located at the bottom and to the left of Figure 10.5.

By Lemmas 10.2.3 and 10.A.3, the probability that this rectangle contains no vertical double gap is at most

$$\exp\big(-3Dg(tq)\big) \leqslant \left(\frac{1}{\sqrt{a}}\right)^{y}\left(\frac{1}{\sqrt{b}}\right)^{x} \exp\big(-sg(bq) - tg(aq)\big) \leqslant \frac{1}{C} \cdot F^{x,y,z}(s,t)\,.$$

Applying the same argument to the other seven rectangles in Figure 10.5, we may assume that each is either empty, or contains a double gap crossing it in the short direction. But in this case any rectangle satisfying (10.78) must be contained in a square of size $(4D+1) \times (4D+1)$ in one of the four corners of R, and it is therefore not possible that it has semi-perimeter greater than $9D$, as required.

We will now sum over choices of T with $\phi(T) \leqslant 9D$ the probability that

$$\big[A \cap T \backslash S^{\mathbf{x}}_{\blacksquare}\big] = T \qquad \text{and} \qquad \big[S \cup T \cup \big(A \cap R \backslash S^{\mathbf{x}}_{\blacksquare}\big)\big] = R\,. \tag{10.79}$$

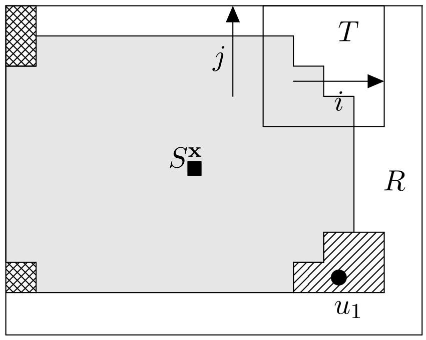

(a) The rectangle T is internally filled outside the shaded $S^{\mathbf{x}}_{\blacksquare}$ and allows S to grow i to the right and j upwards. Since there is an infected site $u_1 \in J$ (in the hatched region), but not in the double hatched set (so $J(u_1)$ is empty), in this case we have $E = \{u_1\}$.

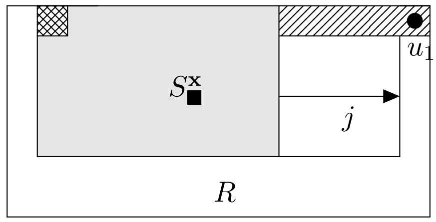

(b) In Algorithm 10.A.5, the rectangle S grows j steps to the right until it reaches either (a) the right-hand side of R, (b) a double gap, or (c) an infected site u_1 in the hatched region (here $x_{(0,1)} = x_{(-1,0)} = 1$ and $x_{(1,0)} = x_{(0,-1)} = 0$). In the figure case (c) occurs, and hence $\ell_{(0,1)} = 1$. The double hatched site in the top-left corner is not occupied, so the set $J_{(-1,0)}$ is empty and hence $\ell_{(-1,0)} = 0$. It follows that $E = \{u_1\}$ in this case.

Figure 10.6 – The two possible growth mechanisms.

Note that these two events depend on disjoint sets of infected sites, and are therefore independent. To bound the probabilities of these events, we will partition according to $k := \phi(T)$, and the dimensions of $[S \cup T]$,

$$\dim\big([S \cup T]\big) = (a - s + i, b - t + j)\,.$$

It was proved in Section 10.3 that, given i, j and k, the expected number of rectangles T satisfying the first event in (10.79) is at most $(24kp)^{k/2}$. Bounding the probability of the second event in (10.79) is unfortunately rather more complicated, and will require the induction hypothesis, and a careful analysis of the possible positions of elements of A in the buffers of $[S \cup T]$. Recall that the case in which there are no infected sites in the buffers was analyzed in Section 10.3.

In order to deal systematically with all of the possible cases, we run the following algorithm. For simplicity we assume that T contains the top right corner of $S^{\mathbf{x}}_{\square}$ (as in Figure 10.6a); the same bound follows in the other three cases by symmetry.

Algorithm 10.A.4. We define a set $E \subset R \backslash [S \cup T]$ of size 0, 1 or 2, as follows:

1. If the set

$$J := A \cap \big(B_{(-1,0)}([S \cup T], R) \cup B_{(0,-1)}([S \cup T], R)\big) \backslash S^{\mathbf{x}}_{\square}$$

   is empty, then set $E := \varnothing$.

2. Otherwise, choose $\mathbf{d} \in \{(-1,0),(0,-1)\}$ and $u_1 \in J \cap B_{\mathbf{d}}([S \cup T], R)$.[6] Now, if
$$J(u_1) := A \cap B_{\mathbf{d}'}\big([S \cup T \cup \{u_1\}], R\big) \backslash S_{\square}^{\mathbf{x}}$$
is empty, where $\{\mathbf{d}, \mathbf{d}'\} = \{(0,-1),(-1,0)\}$, then set $E := \{u_1\}$.

3. Otherwise, choose $u_2 \in J(u_1)$, and set $E := \{u_1, u_2\}$.

We now partition the second event in (10.79) according to the set E, and apply the induction hypothesis to the rectangle
$$\hat{S}(E) := [S \cup T \cup E]\,.$$
Recall that the case $E = \varnothing$ was dealt with in Section 10.3, so we may assume that $|E| \in \{1,2\}$. It is possible to deal with these two cases at the same time, but to simplify the notation we will take them one at a time.

Indeed, suppose first that $|E| = 2$. Then, by the induction hypothesis, it follows that[7]
$$\mathbb{P}_p\Big(\big[\hat{S}(E) \cup \big(A \cap R \backslash S_{\blacksquare}^{\mathbf{x}}\big)\big] = R\Big) \leqslant F^{x-2,y-2,z}(s-i-2, t-j-2)\,.$$
Now, recalling that $k = \phi(T)$, note that there are at most $(2k)^2$ choices for the set E. Hence, recalling that the expected number of rectangles T satisfying the first event in (10.79) is at most $(24kp)^{k/2}$, it will suffice (in this case) to show (cf. (10.18)) that
$$\sum_{i+j \geqslant 4} \sum_{k=i+j}^{36f(R)} 4k^2p^2 \cdot (24kp)^{k/2} \cdot F^{x-2,y-2,z}(s-i-2, t-j-2) \ll F^{x,y,z}(s,t)\,. \tag{10.80}$$
To see this, note first that $24kp \leqslant \delta$, since $k \leqslant 36\delta\sqrt{\min\{a,b\}} \leqslant q^{-1/2}$, and that we may therefore assume that $k = i + j$. Now, observe that
$$\frac{F^{x-2,y-2,z}(s-i-2, t-j-2)}{F^{x,y,z}(s,t)} = \frac{ab}{C^4} \cdot e^{(i+2)g(bq)+(j+2)g(aq)} \leqslant \frac{ab}{C^4}\left(\frac{C}{bq}\right)^{(i+2)/2}\left(\frac{C}{aq}\right)^{(j+2)/2}$$
by (10.5), since $\max\{a,b\} \leqslant 3e^{2B}/q$. Since $i+j = k$, and recalling that $p \leqslant q$ and that $\min\{a,b\} \geqslant \max\{L_1, C^3k\}$, we obtain
$$\sum_{k=4}^{36f(R)} \sum_{i+j=k} 4k^2 \cdot (24kp)^{k/2} \cdot \frac{1}{C^2}\left(\frac{C}{bq}\right)^{i/2}\left(\frac{C}{aq}\right)^{j/2} \leqslant \sum_{k=4}^{36f(R)} \frac{4k^3 \cdot (C^2k)^{k/2}}{\min\{a,b\}^{k/2}} \leqslant \frac{C^5}{\min\{a,b\}^2}\,,$$

[6]Whenever we have a choice to make in either algorithm, we choose the first direction / site of A in some (arbitrary) pre-defined order on $\mathcal{I}$ / the sites in R.

[7]Note that we used here the bound $z(\hat{S}, R) \leqslant z(S, R)$, and the monotonicity of F in z, s, t.

which implies (10.80).

The argument for the case $|E| = 1$ is almost the same. Observe first that, by symmetry, we may assume that $u_1 \in B_{(0,-1)}([S \cup T], R)$, as in Figure 10.6a. Noting that the set $J(u_1)$ (in Algorithm 10.A.4) is empty, it follows from the induction hypothesis that

$$\mathbb{P}_p\Big(\big[\hat{S}(E) \cup \big(A \cap R \backslash S^{\mathbf{x}}_{\blacksquare}\big)\big] = R\Big) \leqslant F^{x-1,y-2,z}(s-i, t-j-2)\,.$$

There are at most $2k$ choices for the vertex u_1, so it will suffice (in this case) to show that

$$\sum_{i+j \geqslant 4} \sum_{k=i+j}^{36f(R)} 2kp \cdot (24kp)^{k/2} \cdot F^{x-1,y-2,z}(s-i, t-j-2) \ll F^{x,y,z}(s,t)\,. \quad (10.81)$$

Since $24kp \leqslant \delta$, we may again assume that $k = i + j$. Now, observe that

$$\begin{aligned} \frac{F^{x-1,y-2,z}(s-i, t-j-2)}{F^{x,y,z}(s,t)} &= \frac{a\sqrt{b}}{C^3} \cdot e^{ig(bq)+(j+2)g(aq)} \\ &\leqslant \frac{a\sqrt{b}}{C^3} \left(\frac{C}{bq}\right)^{i/2} \left(\frac{C}{aq}\right)^{(j+2)/2} \end{aligned}$$

by (10.5), since $\max\{a,b\} \leqslant 3e^{2B}/q$. Since $i \geqslant 1$, we obtain

$$\begin{aligned} \sum_{k=4}^{36f(R)} \sum_{i+j=k} 2k \cdot (24kp)^{k/2} \cdot \frac{\sqrt{b}}{C^2} \left(\frac{C}{bq}\right)^{i/2} \left(\frac{C}{aq}\right)^{j/2} &\leqslant \sum_{k=4}^{36f(R)} \frac{2k^2 \cdot (C^2k)^{k/2}}{\min\{a,b\}^{(k-1)/2}} \\ &\leqslant \frac{C^5}{\min\{a,b\}^{3/2}}\,, \end{aligned}$$

which implies (10.81), since $\min\{a,b\} \geqslant \max\{L_1, C^3k\}$. This completes the proof in Case 1.

Case 2: $z > x + y$, i.e., some non-empty buffer is not included in $S^{\mathbf{x}}_{\square}$.

Without loss of generality, let $B_{(1,0)}(S,R)$ be a non-empty buffer that is not included in $S^{\mathbf{x}}_{\square}$, so $x_{(1,0)} = 0$. As explained in the sketch proof, the idea is to 'grow' S to the right until we find a double gap, an infected site in one of the buffers, or reach the right-hand side of R, thus leading either to an increase in $x + y$, or a decrease in $s + t$. In Section 10.3 we dealt with the cases in which we find a double gap before reaching the right-hand side, and that before doing so we do not find any infected sites in the buffers above or below S. Here we will deal with the other cases.

Algorithm 10.A.5. We define a set $E \subset R \backslash S^{\mathbf{x}}_{\square}$ of size 0, 1, 2 or 3, and also an integer $j \in \{0, \ldots, s\}$ and a variables $\ell_{\mathbf{d}} \in \{0,1\}$ for each direction $\mathbf{d} \in \mathcal{I}$, as follows:[8]

[8]Initially $E = \varnothing$ and $\ell_{\mathbf{d}} = 0$ for every direction $\mathbf{d} \in \mathcal{I}$.

1. Set $\tilde{S} := \bigcup_{i=0}^{j} \big(S + (i,0)\big)$ and $\tilde{\mathbf{x}} := \mathbf{x} + \mathbb{1}_{(1,0)}$, where
$$j := \min\big\{i \geqslant 0 \,:\, A \cap R \cap \big(S + (i+2,0)\big)\backslash\big(S + (i,0)\big) = \varnothing\big\}\,.$$
If the set $A \cap \tilde{S}^{\tilde{\mathbf{x}}}_{\square}\backslash S^{\mathbf{x}}_{\square}$ is empty, then go to Step 8.

2. Set $\tilde{S} := \bigcup_{i=0}^{j} \big(S + (i,0)\big)$, where j is minimal such that $\tilde{S}\backslash S$ is crossed from left to right, and
$$\big(A \cap \tilde{S}^{\tilde{\mathbf{x}}}_{\square}\big)\backslash\big(S^{\mathbf{x}}_{\square} \cup B_{(1,0)}(\tilde{S},R)\big) \neq \varnothing\,.$$

3. Now, if $x_{(0,1)} = 1$ and the set[9]
$$J_{(0,1)} := A \cap \Big(\big(B_{(0,1)}(\tilde{S},R)\backslash B_{(0,1)}(S,R)\big) \cup \{v_{(1,1)}\}\Big)$$
is non-empty, then add a site $u_{(0,1)} \in J_{(0,1)}$ to E, and set $\ell_{(0,1)} := 1$.

4. Similarly, if $x_{(0,-1)} = 1$ and the set
$$J_{(0,-1)} := A \cap \Big(\big(B_{(0,-1)}([\tilde{S} \cup E],R)\backslash B_{(0,-1)}(S,R)\big) \cup \{v_{(1,-1)}\}\Big)$$
is non-empty, then add a site $u_{(0,-1)} \in J_{(0,-1)}$ to E, and set $\ell_{(0,-1)} := 1$.

5. If $x_{(0,1)} = 1$ and $E = \{v_{(1,-1)}\}$, and the set[10]
$$J'_{(0,1)} := A \cap \Big(\big(B_{(0,1)}([\tilde{S} \cup E],R)\backslash B_{(0,1)}(S,R)\big)\Big)$$
is non-empty, then add a site $u_{(0,1)} \in J'_{(0,1)}$ to E, and set $\ell_{(0,1)} := 1$.

6. If $x_{(-1,0)} = 0$ then go to Step 8. Otherwise, if the set
$$J_{(-1,0)} := A \cap B_{(-1,0)}([\tilde{S} \cup E],R)\backslash S^{\mathbf{x}}_{\square}$$
is non-empty, then add a site $u_{(-1,0)} \in J_{(-1,0)}$ to E, and set $\ell_{(-1,0)} := 1$.

7. If either $\ell_{(0,1)} + \ell_{(0,-1)} = x_{(0,1)} + x_{(0,-1)}$ or $\ell_{(-1,0)} = 0$ then go to Step 8. Otherwise:

   (a) If $\ell_{(0,1)} = 1$ and $\ell_{(0,-1)} = 0$, and the set
   $$J_{(-1,-1)} := A \cap B_{(0,-1)}([\tilde{S} \cup E],R)\backslash S^{\mathbf{x}}_{\square}$$
   is non-empty, then add a site $u_{(-1,-1)} \in J_{(-1,-1)}$ to E and set $\ell_{(0,-1)} := 1$.

---

[9]For each $i \in \{1,-1\}$, let $v_{(1,i)}$ denote the unique 'corner' site in $R\backslash\tilde{S}$ with a neighbour in each of the buffers $B_{(1,0)}(\tilde{S},R)$ and $B_{(0,i)}(\tilde{S},R)$.

[10]Note that this set has at most one element.

(*b*) If $\ell_{(0,1)} = 0$ and $\ell_{(0,-1)} = 1$, and the set

$$J_{(-1,1)} := A \cap B_{(0,1)}([\tilde{S} \cup E], R) \setminus S_{\square}^{\mathbf{x}}$$

is non-empty, then add a site $u_{(-1,1)} \in J_{(-1,1)}$ to E and set $\ell_{(0,1)} := 1$.

8. Set $\hat{S} := [\tilde{S} \cup E]$ and STOP.

Observe that Algorithm 10.A.5 outputs a set E and an integer j (which together determine the sets $\tilde{S}$ and $\hat{S}$, and the variable $\ell_{\mathbf{d}}$ for each $\mathbf{d} \in \mathcal{I}$) with the following properties:

(*a*) $E \subset A$;

(*b*) $[S \cup (A \cap \tilde{S})] = \tilde{S}$;

(*c*) On the event $D_1^{\mathbf{x}}(S, R)$ the event $D_1^{\hat{\mathbf{x}}}(\hat{S}, R)$ occurs, i.e.,

$$\big[\hat{S} \cup \big(A \cap R \setminus \hat{S}_{\blacksquare}^{\hat{\mathbf{x}}}\big)\big] = R \,,$$

where $\hat{\mathbf{x}} = \mathbf{x} - \boldsymbol{\ell}$, except in the case treated in Section 10.3, in which $\hat{\mathbf{x}} = \mathbf{x} + \mathbb{1}_{(1,0)}$.

We will analyse each case individually, and sum over all possible sets E and $j \in \{0, \dots, s\}$.

Suppose first that $E = \varnothing$. In Section 10.3 we dealt with this case, under the additional assumption that $B_{(1,0)}(\hat{S}, R) \neq \varnothing$ (i.e., we found a double gap before reaching the right-hand side of R); here we will deal with the other case (i.e., that $\hat{S}$ reaches the right-hand side of R). To do so, we need to bound the probability that

$$[S \cup (A \cap \tilde{S})] = \tilde{S} \qquad \text{and} \qquad D_1^{\hat{\mathbf{x}}}(\hat{S}, R) \,,$$

where in this case $\hat{S} = \tilde{S}$ and $\hat{\mathbf{x}} = \mathbf{x}$. Note that these two events depend on disjoint sets of sites, and are therefore independent; we will bound the first using Lemma 10.2.3, and the second using the induction hypothesis. Indeed, by Lemma 10.2.3 (and since $g(z)$ is decreasing) we have

$$\mathbb{P}_p\big([S \cup (A \cap \tilde{S})] = \tilde{S}\big) \leqslant \exp\big(-jg(bq)\big) \,,$$

and by the induction hypothesis (since $B_{(1,0)}(\hat{S}, R) = \varnothing$ and $E = \varnothing$),

$$\mathbb{P}_p\big(D_1^{\mathbf{x}}(\hat{S}, R)\big) \leqslant F^{x,y,z-1}(s-j, t) \,.$$

Thus, the probability that $E = \varnothing$, $B_{(1,0)}(\hat{S}, R) = \varnothing$ and $D_1^{\mathbf{x}}(S, R)$ is at most[11]

$$\exp\big(-jg(bq)\big) F^{x,y,z-1}(s-j, t) = \frac{1}{C} \cdot F^{x,y,z}(s, t)$$

[11] Note that there is a unique possible value of j for which $B_{(1,0)}(\hat{S}, R) = \varnothing$.

which suffices since C is sufficiently large.

We will therefore assume from now on that $E \neq \varnothing$, which means that before reaching a double gap (when trying to grow S rightwards, forming $\tilde{S}$), we find an infected site in either the buffer above or below $\tilde{S}$. More precisely, recall that $\tilde{S} := \bigcup_{i=0}^{j} \big(S + (i,0)\big)$, where j is minimal such that $\tilde{S} \backslash S$ is crossed from left to right, and

$$\big(A \cap \tilde{S}_{\square}^{\tilde{\mathbf{x}}}\big) \backslash \big(S_{\square}^{\mathbf{x}} \cup B_{(1,0)}(\tilde{S}, R)\big) \neq \varnothing\,,$$

where $\tilde{\mathbf{x}} = \mathbf{x} + \mathbb{1}_{(1,0)}$. Note that $x_{(0,1)} + x_{(0,-1)} \geqslant 1$, and that moreover $\ell_{(0,1)} + \ell_{(0,-1)} \geqslant 1$ (since the set above is non-empty).

Suppose first that $|E| = 1$, and therefore (without loss of generality) we have $\ell_{(0,1)} = 1$, $E = \{u_{(0,1)}\}$ and $\ell_{(0,-1)} = \ell_{(-1,0)} = 0$. This implies that the (independent) events

$$u_{(0,1)} \in A, \qquad [S \cup (A \cap \tilde{S})] = \tilde{S} \qquad \text{and} \qquad D_1^{\hat{\mathbf{x}}}(\hat{S}, R)$$

occur, where $\hat{\mathbf{x}} = \mathbf{x} - \mathbb{1}_{(0,1)}$. Given j, it follows from the minimality of j that there are at most four choices for $u_{(0,1)}$, so

$$\mathbb{P}_p\Big(\big\{u_{(0,1)} \in A\big\} \cap \big\{[S \cup (A \cap \tilde{S})] = \tilde{S}\big\}\Big) \leqslant 4p \exp\big(-jg(bq)\big)\,.$$

Suppose first that $u_{(0,1)} \neq v_{(1,1)}$. In this case, by the induction hypothesis, we have

$$\mathbb{P}_p\big(D_1^{\hat{\mathbf{x}}}(\hat{S}, R)\big) \leqslant F^{x,y-1,z}(s-j, t-2)\,,$$

and so the probability in this case can be bounded by

$$\sum_{j=0}^{s} 4p \exp\big(-jg(bq)\big) F^{x,y-1,z}(s-j, t-2) \leqslant \frac{4(s+1)p\sqrt{a}}{C} e^{2g(aq)} \cdot F^{x,y,z}(s,t)\,.$$

Recalling that $s \leqslant 4\delta\sqrt{a}$ and that $e^{2g(aq)} \leqslant C/aq$, by (10.5), and recalling that δ is sufficiently small, we obtain a suitably strong bound in this case. Similarly, if $u_{(0,1)} = v_{(1,1)}$, then the induction hypothesis gives

$$\mathbb{P}_p\big(D_1^{\hat{\mathbf{x}}}(\hat{S}, R)\big) \leqslant F^{x,y-1,z}(s-j-1, t-1)\,,$$

and so the probability in this case can be bounded by

$$\sum_{j=0}^{s} 4p \exp\big(-jg(bq)\big) F^{x,y-1,z}(s-j-1, t-1) \leqslant$$
$$\frac{4(s+1)p\sqrt{a}}{C} e^{g(aq)+g(bq)} \cdot F^{x,y,z}(s,t)\,.$$

Since $s \leqslant 4\delta\sqrt{b}$ and $e^{g(aq)+g(bq)} \leqslant C/(q\sqrt{ab})$, we again obtain a suitable bound.

In all remaining cases we win easily, since each extra infected site is extremely expensive. Nevertheless, we will go carefully through each case. Indeed, suppose next that $|E| = 2$, so (without loss of generality) either

(a) $\ell_{(0,1)} = \ell_{(0,-1)} = 1$, $\ell_{(-1,0)} = 0$, and $E = \{u_{(0,1)}, u_{(0,-1)}\}$, or

(b) $\ell_{(0,1)} = \ell_{(-1,0)} = 1$, $\ell_{(0,-1)} = 0$, and $E = \{u_{(0,1)}, u_{(-1,0)}\}$.

In case (a), the (independent) events

$$E \subset A, \qquad [S \cup (A \cap \tilde{S})] = \tilde{S} \qquad \text{and} \qquad D_1^{\hat{\mathbf{x}}}(\hat{S}, R)$$

occur, where $\hat{\mathbf{x}} = \mathbf{x} - \mathbb{1}_{(0,1)} - \mathbb{1}_{(0,-1)}$. Given j, there are at most five choices for each element of E, so

$$\mathbb{P}_p\big(E \subset A \text{ and } [S \cup (A \cap \tilde{S})] = \tilde{S}\big) \leqslant (5p)^2 \exp\big(-jg(bq)\big).$$

Suppose first that the set $E \cap \{v_{(1,1)}, v_{(1,-1)}\}$ is empty. In this case, by the induction hypothesis, we have

$$\mathbb{P}_p\big(D_1^{\hat{\mathbf{x}}}(\hat{S}, R)\big) \leqslant F^{x,y-2,z}(s-j, t-4),$$

and so the probability in this case can be bounded by

$$\sum_{j=0}^{s} (5p)^2 e^{-jg(bq)} \cdot F^{x,y-2,z}(s-j, t-4) \leqslant \frac{25(s+1)p^2 a}{C^2} e^{4g(aq)} \cdot F^{x,y,z}(s,t).$$

Since $s \leqslant 4\delta\sqrt{a}$ and $e^{2g(aq)} \leqslant C/aq$, we win (easily) in this case. Similarly, if the set $E \cap \{v_{(1,1)}, v_{(1,-1)}\}$ is non-empty, then the induction hypothesis gives

$$\mathbb{P}_p\big(D_1^{\hat{\mathbf{x}}}(\hat{S}, R)\big) \leqslant F^{x,y-2,z}(s-j-1, t-3),$$

and so the probability in this case can be bounded by

$$\sum_{j=0}^{s} (5p)^2 e^{-jg(bq)} \cdot F^{x,y-2,z}(s-j-1, t-3) \leqslant \frac{25(s+1)p^2 a}{C^2} e^{3g(aq)+g(bq)} \cdot F^{x,y,z}(s,t).$$

Since $s \leqslant 4\delta\sqrt{b}$ and $e^{3g(aq)+g(bq)} \leqslant C^2/(q^2\sqrt{a^3 b})$, we again obtain a suitable bound.

In case (b), the (independent) events

$$E \subset A, \qquad [S \cup (A \cap \tilde{S})] = \tilde{S} \qquad \text{and} \qquad D_1^{\hat{\mathbf{x}}}(\hat{S}, R)$$

occur, where $\hat{\mathbf{x}} = \mathbf{x} - \mathbb{1}_{(0,1)} - \mathbb{1}_{(-1,0)}$. Given j, there are at most four choices for each element of E, so

$$\mathbb{P}_p\big(E \subset A \text{ and } [S \cup (A \cap \tilde{S})] = \tilde{S}\big) \leqslant (4p)^2 \exp\big(-jg(bq)\big)\,.$$

Suppose first that $u_{(0,1)} \neq v_{(1,1)}$. In this case, by the induction hypothesis, we have

$$\mathbb{P}_p\big(D_1^{\hat{\mathbf{x}}}(\hat{S}, R)\big) \leqslant F^{x-1,y-1,z}(s-j-2, t-2)\,,$$

and so the probability in this case can be bounded by

$$\sum_{j=0}^{s} (4p)^2 e^{-jg(bq)} \cdot F^{x-1,y-1,z}(s-j-2,t-2) \leqslant \frac{16(s+1)p^2\sqrt{ab}}{C^2} e^{2g(aq)+2g(bq)} \cdot F^{x,y,z}(s,t)\,.$$

Since $s \leqslant 4\delta\sqrt{a}$ and $e^{2g(aq)+2g(bq)} \leqslant C^2/(abq^2)$, we are done in this case, as before. Similarly, if $u_{(0,1)} = v_{(1,1)}$, then the induction hypothesis gives

$$\mathbb{P}_p\big(D_1^{\hat{\mathbf{x}}}(\hat{S}, R)\big) \leqslant F^{x-1,y-1,z}(s-j-3, t-1)\,,$$

and so the probability in this case can be bounded by

$$\sum_{j=0}^{s} (4p)^2 e^{-jg(bq)} \cdot F^{x-1,y-1,z}(s-j-3,t-1) \leqslant \frac{16(s+1)p^2\sqrt{ab}}{C^2} e^{g(aq)+3g(bq)} \cdot F^{x,y,z}(s,t)\,.$$

Since $s \leqslant 4\delta\sqrt{b}$ and $e^{g(aq)+3g(bq)} \leqslant C^2/(q^2\sqrt{ab^3})$, we again obtain a suitable bound.

Finally, suppose that $|E| = 3$, and observe that $\ell_{(0,1)} = \ell_{(0,-1)} = \ell_{(-1,0)} = 1$ and, without loss of generality, either

(*a*) $E = \big\{u_{(0,1)}, u_{(0,-1)}, u_{(-1,0)}\big\}$, or

(*b*) $E = \big\{u_{(0,1)}, u_{(-1,0)}, u_{(-1,-1)}\big\}$.

In either case, the (independent) events

$$E \subset A, \qquad [S \cup (A \cap \tilde{S})] = \tilde{S} \qquad \text{and} \qquad D_1^{\hat{\mathbf{x}}}(\hat{S}, R)$$

occur, where $\hat{\mathbf{x}} = \mathbf{x} - \mathbb{1}_{(0,1)} - \mathbb{1}_{(0,-1)} - \mathbb{1}_{(-1,0)} = \mathbf{0}$. Given j, there are at most six choices for each element of E, so

$$\mathbb{P}_p\big(E \subset A \text{ and } [S \cup (A \cap \tilde{S})] = \tilde{S}\big) \leqslant (6p)^3 \exp\big(-jg(bq)\big)\,.$$

Suppose first that the set $E \cap \{v_{(1,1)}, v_{(1,-1)}\}$ is empty. In this case, by the induction hypothesis, we have

$$\mathbb{P}_p\big(D_1^{\hat{\mathbf{x}}}(\hat{S}, R)\big) \leqslant F^{x-1,y-2,z}(s-j-2, t-4)\,,$$

and so the probability in this case can be bounded by

$$\sum_{j=0}^{s-1} (6p)^3 e^{-jg(bq)} \cdot F^{x-1,y-2,z}(s-j-2, t-4) \leqslant \frac{6^3 sp^3 a\sqrt{b}}{C^3} e^{4g(aq)+2g(bq)} \cdot F^{x,y,z}(s,t)\,.$$

Since $s \leqslant 4\delta\sqrt{a}$ and $e^{4g(aq)+2g(bq)} \leqslant C^3/(q^3a^2b)$, we again win easily in this case. Similarly, if $E \cap \{v_{(1,1)}, v_{(1,-1)}\}$ is non-empty, then the induction hypothesis gives

$$\mathbb{P}_p\big(D_1^{\hat{\mathbf{x}}}(\hat{S}, R)\big) \leqslant F^{x-1,y-2,z}(s-j-3, t-3)\,,$$

and so the probability in this case can be bounded by

$$\sum_{j=0}^{s-1} (6p)^3 e^{-jg(bq)} \cdot F^{x-1,y-2,z}(s-j-3, t-3) \leqslant \frac{6^3 sp^3 a\sqrt{b}}{C^3} e^{3g(aq)+3g(bq)} \cdot F^{x,y,z}(s,t)\,.$$

Since $s \leqslant 4\delta\sqrt{b}$ and $e^{3g(aq)+3g(bq)} \leqslant C^3/(q^3\sqrt{a^3b^3})$, we again obtain a suitable bound.

Summing over the various cases completes the proof of Lemma 10.A.1. □

The proof of Lemma 10.A.2 is very similar to that of Lemma 10.A.1, and we will be able to reuse large parts of the proof (in particular Algorithms 10.A.4 and 10.A.5).

Proof of Lemma 10.A.2. Recall that $D_2^{\mathbf{x}}(S, R)$ denotes the event that

$$\big[S \cup \big(A \cap R \backslash S^{\mathbf{x}}_{\blacksquare}\big)\big] = R \qquad \text{and} \qquad A \cap S^{\mathbf{x}}_{\square} = \varnothing\,.$$

As in the proof of Lemma 10.A.1, we use induction on the pair $(s+t, -(x+y))$, this time to prove that

$$\mathbb{P}_p\big(D_2^{\mathbf{x}}(S, R)\big) \leqslant \hat{F}^{x,y,z}(s,t)\,,$$

where

$$\hat{F}^{x,y,z}(s,t) := \Big(Ce^{\text{short}(R)q}\Big)^z \big(C\sqrt{q}e^{-aq}\big)^y \Big(C\sqrt{q}e^{-bq}\Big)^x \exp\big(-sg(bq) - tg(aq)\big)\,,$$

for every $0 \leqslant s, t \leqslant 4\delta \cdot q^{-1/2} \cdot \exp\big(\min\{a,b\} \cdot q\big)$ and $\mathbf{x} \in \{0,1\}^{\mathcal{I}}$, and every $S \subset R$ with $\dim(S) = (a-s, b-t)$, where x and y are as defined in (10.74), and $z = z(S,R)$. The base of the induction remains unchanged from Lemma 10.A.1.

As in the proof of Lemma 10.A.1, we partition into cases depending on whether or not $z = x + y$, the function $\hat{F}^{x,y,z}(s,t)$ is increasing in z and decreasing in x, y, s and t, and we may assume without loss of generality that $\mathbf{x} = \mathbf{x} \cdot \mathbb{1}_{Z(S,R)}$.

Case 1: $z = x + y$, i.e., all of the non-empty buffers are included in $S^{\mathbf{x}}_{\square}$.

As in Lemma 10.A.1, the event $D_2^{\mathbf{x}}(S,R)$ requires the existence of a rectangle T such that

$$\big[A \cap T \backslash S^{\mathbf{x}}_{\blacksquare}\big] = T \qquad \text{and} \qquad T \cap S^{\mathbf{x}}_{\blacksquare} \neq \varnothing\,.$$

The first step is to apply Lemma 10.A.3, as in the proof of Lemma 10.A.1, to exclude rectangles T with $\phi(T) > 9D$, where this time we set $D := 4\delta \cdot q^{-1/2} \cdot \exp\big(\min\{a,b\} \cdot q\big)$. It follows from (10.76) that $s, t \leqslant D \leqslant 4\delta/q$ and $9D \leqslant \min\{a,b\}$, and we may therefore argue exactly as before, except using the second inequality in Lemma 10.A.3, which gives

$$\begin{aligned} \exp\big(-3Dg(tq)\big) &\leqslant \big(\sqrt{q}e^{-aq}\big)^y \Big(\sqrt{q}e^{-bq}\Big)^x \exp\big(-sg(bq) - tg(aq)\big) \\ &\leqslant \frac{1}{C} \cdot \hat{F}^{x,y,z}(s,t)\,. \end{aligned}$$

We will therefore assume from now on that $\phi(T) \leqslant 9D$, and sum over choices of T with $\phi(T) \leqslant 9D$ the probability that

$$\big[A \cap T \backslash S^{\mathbf{x}}_{\blacksquare}\big] = T, \quad \big[S \cup T \cup \big(A \cap R \backslash S^{\mathbf{x}}_{\blacksquare}\big)\big] = R \quad \text{and} \quad A \cap S^{\mathbf{x}}_{\square} = \varnothing\,. \tag{10.82}$$

Note that these events depend on disjoint sets of sites and are therefore independent. It was proved in Section 10.3 that, given $k := \phi(T)$ and the dimensions of $[S \cup T]$, the expected number of rectangles T satisfying the first event in (10.82) is at most $(24kp)^{k/2}$. For the intersection of the second and third events, we will partition the space according to the set E given by Algorithm 10.A.4, and apply the induction hypothesis to the set

$$\hat{S}(E) := [S \cup T \cup E]\,.$$

Suppose first that $E = \varnothing$, and recall that this means that the set

$$\big(B_{(-1,0)}([S \cup T], R) \cup B_{(0,-1)}([S \cup T], R)\big) \backslash S^{\mathbf{x}}_{\square}$$

contains no element of A. Together with the second and third events appearing in (10.82), this implies that the event $D_2^{\hat{\mathbf{x}}}(\hat{S}(E), R)$ occurs, where $\hat{\mathbf{x}} = \mathbf{x} - \mathbb{1}_{(1,0)} - \mathbb{1}_{(0,1)}$. By the induction hypothesis, we have

$$\mathbb{P}_p\big(D_2^{\hat{\mathbf{x}}}(\hat{S}(E), R)\big) \leqslant \hat{F}^{x-1,y-1,z}(s-i, t-j)\,,$$

where $\dim([S \cup T]) = (a - s + i, b - t + j)$. Now, since $B/q \leqslant a, b \leqslant (1/(2q)) \log(1/q)$, we have

$$\frac{\hat{F}^{x-1,y-1,z}(s-i,t-j)}{\hat{F}^{x,y,z}(s,t)} = \frac{e^{(a+b)q}}{C^2 q} \exp\big(ig(bq) + jg(aq)\big) \leqslant \frac{2^{i+j}}{C^2 q^2}$$

since $e^{g(aq)+g(bq)} \leqslant 2$. Recalling that $k \leqslant 9D \leqslant 36\delta/q$, it follows that the probability of this case is at most

$$\sum_{i+j\geqslant 4} \sum_{k=i+j}^{9D} (24kp)^{k/2} \cdot \frac{2^{i+j}}{C^2 q^2} \cdot \hat{F}^{x,y,z}(s,t) \leqslant \frac{1}{C} \cdot \hat{F}^{x,y,z}(s,t),$$

as required.

Suppose next that $|E| = 1$, and observe that, by symmetry, we may assume that $u_1 \in B_{(0,-1)}([S \cup T], R)$, as in Figure 10.6a. Recalling that the set $J(u_1)$ (in Algorithm 10.A.4) is empty, it follows that the event $D_2^{\hat{\mathbf{x}}}(\hat{S}(E), R)$ occurs, where $\hat{\mathbf{x}} = \mathbf{x} - \mathbb{1}_{(1,0)} - \mathbb{1}_{(0,1)} - \mathbb{1}_{(0,-1)}$. By the induction hypothesis, we have

$$\mathbb{P}_p\big(D_2^{\hat{\mathbf{x}}}(\hat{S}(E), R)\big) \leqslant \hat{F}^{x-1,y-2,z}(s-i, t-j-2)$$

and, since $B/q \leqslant a, b \leqslant (1/(2q)) \log(1/q)$, we have (as before)

$$\frac{\hat{F}^{x-1,y-2,z}(s-i,t-j-2)}{\hat{F}^{x,y,z}(s,t)} = \frac{e^{(2a+b)q}}{C^3 q^{3/2}} \exp\big(ig(bq) + (j+2)g(aq)\big) \leqslant \frac{2^{i+j+2}}{C^3 q^3}.$$

Noting that there are at most $2k$ choices for the vertex u_1, and recalling that $k \leqslant 9D \leqslant 36\delta/q$, it follows that the probability of this case is at most

$$\sum_{i+j\geqslant 4} \sum_{k=i+j}^{9D} 2kp \cdot (24kp)^{k/2} \cdot \frac{2^{i+j+2}}{C^3 q^3} \cdot \hat{F}^{x,y,z}(s,t) \leqslant \frac{1}{C^2} \cdot \hat{F}^{x,y,z}(s,t),$$

as required.

Finally, suppose that $|E| = 2$, and observe that in this case the event $D_2^{\hat{\mathbf{x}}}(\hat{S}(E), R)$ occurs, where $\hat{\mathbf{x}} = \mathbf{0}$, and that $\mathbf{x} = \mathbf{1}$. By the induction hypothesis, we have

$$\mathbb{P}_p\big(D_2^{\hat{\mathbf{x}}}(\hat{S}(E), R)\big) \leqslant \hat{F}^{x-2,y-2,z}(s-i-2, t-j-2)$$

and, since $B/q \leqslant a, b \leqslant (1/(2q)) \log(1/q)$, we have (as before)

$$\frac{\hat{F}^{x-2,y-2,z}(s-i-2,t-j-2)}{\hat{F}^{x,y,z}(s,t)} = \frac{e^{(2a+2b)q}}{C^4 q^2} \exp\big((i+2)g(bq) + (j+2)g(aq)\big) \leqslant \frac{2^{i+j+4}}{C^4 q^4}.$$

Noting that there are at most $4k^2$ choices for E, and since $k \leqslant 9D \leqslant 36\delta/q$, it follows that the probability of this case is at most

$$\sum_{i+j\geqslant 4} \sum_{k=i+j}^{9D} (2kp)^2 \cdot (24kp)^{k/2} \cdot \frac{2^{i+j+4}}{C^4 q^4} \cdot \hat{F}^{x,y,z}(s,t) \leqslant \frac{1}{C^3} \cdot \hat{F}^{x,y,z}(s,t)\,,$$

as required. This completes the proof in Case 1.

Case 2: $z > x + y$.

As in the proof of Lemma 10.A.1, let $B_{(1,0)}(S,R)$ be a non-empty buffer that is not included in $S_{\square}^{\mathbf{x}}$, so $x_{(1,0)} = 0$, and define a set E using Algorithm 10.A.5.

Suppose first that $E = \varnothing$, and recall that $\tilde{S} = \bigcup_{i=0}^{j} \big(S + (i,0)\big)$, where

$$j = \min\big\{i \geqslant 0 \,:\, A \cap R \cap \big(S + (i+2,0)\big)\backslash\big(S + (i,0)\big) = \varnothing\big\},$$

and that $\tilde{S}_{\square}^{\tilde{\mathbf{x}}}\backslash S_{\square}^{\mathbf{x}}$ contains no elements of A, where $\tilde{\mathbf{x}} = \mathbf{x} + \mathbb{1}_{(1,0)}$. There are two sub-cases, depending on whether or not $B_{(1,0)}(\tilde{S},R) = \varnothing$, that is, whether or not we reached the right-hand side without finding a double gap. Suppose first that we did find a double gap (i.e., $B_{(1,0)}(\tilde{S},R) \neq \varnothing$). We will sum over choices of j the probability that

$$[S \cup (A \cap \tilde{S})] = \tilde{S}, \qquad \big[\tilde{S} \cup \big(A \cap R\backslash \tilde{S}_{\blacksquare}^{\tilde{\mathbf{x}}}\big)\big] = R \qquad \text{and} \qquad A \cap \tilde{S}_{\square}^{\tilde{\mathbf{x}}} = \varnothing\,. \tag{10.83}$$

Note that these three events depend on disjoint sets of sites, and are therefore independent; we will bound the first using Lemma 10.2.3, and the intersection of the second and third using the induction hypothesis. Indeed, by Lemma 10.2.3 (and since $g(z)$ is decreasing) we have

$$\mathbb{P}_p\big([S \cup (A \cap \tilde{S})] = \tilde{S}\big) \leqslant \exp\big(-jg(bq)\big)\,.$$

Moreover, the second and third events imply that the event $D_2^{\tilde{\mathbf{x}}}(\tilde{S},R)$ occurs, and by the induction hypothesis we have

$$\mathbb{P}_p\big(D_2^{\tilde{\mathbf{x}}}(\tilde{S},R)\big) \leqslant \hat{F}^{x+1,y,z}(s-j,t)\,.$$

It follows that the probability that there exists $j \geqslant 0$ such that the events in (10.83) all hold is at most

$$\sum_{j=0}^{s-1} e^{-jg(bq)} \cdot \hat{F}^{x+1,y,z}(s-j,t) = Cs\sqrt{q}e^{-bq} \cdot \hat{F}^{x,y,z}(s,t) \leqslant 4C\delta \cdot \hat{F}^{x,y,z}(s,t)$$

as required, since $s \leqslant 4\delta \cdot q^{-1/2} \cdot \exp\big(\min\{a,b\} \cdot q\big)$ and $\delta = \delta(C) > 0$ is sufficiently small.

We next deal with the case $E = \varnothing$ and $B_{(1,0)}(\tilde{S}, R) = \varnothing$ (i.e., we reached the right-hand side without finding a double gap). In this case, the event $D_2^{\mathbf{x}}(\tilde{S}, R)$ occurs, and by the induction hypothesis we have

$$\mathbb{P}_p\big(D_2^{\mathbf{x}}(\tilde{S}, R)\big) \leqslant \hat{F}^{x,y,z-1}(s-j,t)\,,$$

since $z(\tilde{S}, R) = z(S, R) - 1$. It follows that the probability that the events in (10.83) all hold, and $B_{(1,0)}(\tilde{S}, R) = \varnothing$, is at most

$$\exp\big(-jg(bq)\big)\hat{F}^{x,y,z-1}(s-j,t) \leqslant \frac{1}{C}\cdot \hat{F}^{x,y,z}(s,t)$$

which suffices since C is sufficiently large.

We will therefore assume from now on that $E \neq \varnothing$, so $\tilde{S} = \bigcup_{i=0}^{j}\big(S + (i,0)\big)$, where j is minimal such that $\tilde{S}\backslash S$ is crossed from left to right, and

$$\big(A \cap \tilde{S}_{\square}^{\tilde{\mathbf{x}}}\big)\backslash\big(S_{\square}^{\mathbf{x}} \cup B_{(1,0)}(\tilde{S}, R)\big) \neq \varnothing\,.$$

In other words, before reaching a double gap we found an infected site in either the buffer above or below $\tilde{S}$. In particular, note that $\ell_{(0,1)} + \ell_{(0,-1)} \geqslant 1$.

Suppose first that $|E| = 1$, and therefore that (without loss of generality) we have $\ell_{(0,1)} = 1$, $E = \{u_{(0,1)}\}$ and $\ell_{(0,-1)} = \ell_{(-1,0)} = 0$. Then the events

$$u_{(0,1)} \in A, \quad [S \cup (A \cap \tilde{S})] = \tilde{S}, \quad A \cap B_{(0,1)}(S,R) = \varnothing \quad \text{and} \quad D_2^{\hat{\mathbf{x}}}(\hat{S}, R)$$

occur, where $\hat{S} = [\tilde{S} \cup E]$ and $\hat{\mathbf{x}} = \mathbf{x} - \mathbb{1}_{(0,1)}$. There is an important subtlety in this case, since these events might not be independent: the buffer $B_{(0,1)}(S, R)$ might 'stick out' of the top of $\hat{S}$, and therefore intersect the set of sites that the event $D_2^{\hat{\mathbf{x}}}(\hat{S}, R)$ depends on. However, the only dependence is between the decreasing event $\{A \cap B_{(0,1)}(S, R) = \varnothing\}$ and the increasing part of the event $D_2^{\hat{\mathbf{x}}}(\hat{S}, R)$ (since $\hat{x}_{(0,1)} = 0$), so by Harris' inequality [202][12] the probability that all four events occur is at most the product of their probabilities.

Given j, there are at most four choices for $u_{(0,1)}$, so

$$\mathbb{P}_p\Big(\big\{u_{(0,1)} \in A\big\} \cap \big\{[S \cup (A \cap \tilde{S})] = \tilde{S}\big\}\Big) \leqslant 4p\exp\big(-jg(bq)\big)\,.$$

Suppose first that $z(\hat{S}, R) < z(S, R)$. In this case, by the induction hypothesis, we have

$$\mathbb{P}_p\Big(\big\{A \cap B_{(0,1)}(S,R) = \varnothing\big\} \cap D_2^{\hat{\mathbf{x}}}(\hat{S}, R)\Big) \leqslant \\ (1-p)^{a-s}\cdot \hat{F}^{x,y-1,z-1}(s-j-1,t-2)\,,$$

[12] Harris' inequality states that increasing events in a product space are positively correlated. It is often referred to as the FKG inequality, which is a generalization that was proved somewhat later.

and so the probability in this case can be bounded by[13]

$$\begin{aligned}\sum_{j=0}^{s} 4p(1-p)^{a-s}e^{-jg(bq)} \cdot \hat{F}^{x,y-1,z-1}(s-j-1,t-2) \\ \leqslant 4(s+1)p \cdot e^{-(a-s)q} \cdot \frac{e^{aq-\min\{a,b\}q}}{C^2\sqrt{q}} \cdot e^{2g(aq)+g(bq)} \cdot \hat{F}^{x,y,z}(s,t) \\ \leqslant \frac{\delta}{C} \cdot \hat{F}^{x,y,z}(s,t)\,,\end{aligned}$$

since $s \leqslant 4\delta \cdot q^{-1/2} \cdot \exp\big(\min\{a,b\} \cdot q\big) \leqslant 4\delta/q$, and hence $e^{sq+2g(aq)+g(bq)} \leqslant 2$.

On the other hand, if $z(\hat{S},R) = z(S,R)$ then the buffer $B_{(0,1)}(S,R)$ must have height at least two, and hence

$$\mathbb{P}_p\Big(\big\{A \cap B_{(0,1)}(S,R) = \varnothing\big\} \cap D_2^{\mathbf{x}}(\hat{S},R)\Big) \leqslant \\ (1-p)^{2(a-s)} \cdot \hat{F}^{x,y-1,z}(s-j-1,t-2)\,.$$

The probability in this case can therefore be bounded, as above, by

$$\begin{aligned}&\sum_{j=0}^{s} 4p(1-p)^{2(a-s)}e^{-jg(bq)} \cdot \hat{F}^{x,y-1,z}(s-j-1,t-2) \\ &\leqslant 4(s+1)p \cdot e^{-2(a-s)q} \cdot \frac{e^{aq}}{C\sqrt{q}} \cdot e^{2g(aq)+g(bq)} \cdot \hat{F}^{x,y,z}(s,t) \leqslant \delta \cdot \hat{F}^{x,y,z}(s,t)\,,\end{aligned}$$

as required.

The remaining cases are similar but easier, since each extra infected site is extremely expensive. We will therefore be able to be use slightly weaker bounds, which simplifies the analysis somewhat. Suppose next that $|E| = 2$, so either

(a) $\ell_{(0,1)} = \ell_{(0,-1)} = 1$, $\ell_{(-1,0)} = 0$, and $E = \{u_{(0,1)}, u_{(0,-1)}\}$, or

(b) $\ell_{(0,1)} = \ell_{(-1,0)} = 1$, $\ell_{(0,-1)} = 0$, and $E = \{u_{(0,1)}, u_{(-1,0)}\}$.

In either case, given j there are at most five choices for each element of E, so

$$\mathbb{P}_p\big(E \subset A \;\text{ and }\; [S \cup (A \cap \tilde{S})] = \tilde{S}\big) \leqslant (5p)^2 \exp\big(-jg(bq)\big)\,.$$

Moreover, in case (a), by the induction hypothesis and Harris' inequality, we have

$$\mathbb{P}_p\Big(\big\{A \cap B = \varnothing\big\} \cap D_2^{\mathbf{x}}(\hat{S},R)\Big) \leqslant (1-p)^{2(a-s)} \cdot \hat{F}^{x,y-2,z}(s-j-1,t-4)\,,$$

[13]We remark that this is the only point in the proof where we will need the term $e^{\text{short}(R)qz}$ in the bound in Lemma 10.3.4. This term gives rise to the term $\text{up}(\mathcal{H})$ in the proof of Theorem 10.1.1 and the corresponding precision needed in Lemma 10.4.14.

where $B = B_{(0,1)}(S,R) \cup B_{(0,-1)}(S,R)$ and $\hat{\mathbf{x}} = \mathbf{x} - \mathbb{1}_{(0,1)} - \mathbb{1}_{(0,-1)}$.[14] The probability in this case can therefore be bounded by

$$\sum_{j=0}^{s} (5p)^2 (1-p)^{2(a-s)} e^{-jg(bq)} \cdot \hat{F}^{x,y-2,z}(s-j-1,t-4) \leqslant \frac{\delta}{C} \cdot \hat{F}^{x,y,z}(s,t)\,,$$

since $s \leqslant 4\delta/q$ and $e^{2sq+4g(aq)+g(bq)} \leqslant 2$. Similarly, in case (b) we have

$$\mathbb{P}_p\Big(\big\{A \cap B = \varnothing\big\} \cap D_2^{\hat{\mathbf{x}}}(\hat{S},R)\Big) \leqslant (1-p)^{a-s+b-t} \cdot \hat{F}^{x-1,y-1,z}(s-j-3,t-2)\,,$$

where $B = B_{(0,1)}(S,R) \cup B_{(-1,0)}(S,R)$ and $\hat{\mathbf{x}} = \mathbf{x} - \mathbb{1}_{(0,1)} - \mathbb{1}_{(-1,0)}$, which allows us to bound the probability in this case by

$$\sum_{j=0}^{s} (5p)^2 (1-p)^{a-s+b-t} e^{-jg(bq)} \cdot \hat{F}^{x-1,y-1,z}(s-j-3,t-2) \leqslant \frac{\delta}{C} \cdot \hat{F}^{x,y,z}(s,t)\,,$$

exactly as before. The calculation when $|E| = 3$ is almost the same. Recall that $\ell_{(0,1)} = \ell_{(0,-1)} = \ell_{(-1,0)} = 1$ and, without loss of generality, either

(a) $E = \{u_{(0,1)}, u_{(0,-1)}, u_{(-1,0)}\}$, or

(b) $E = \{u_{(0,1)}, u_{(-1,0)}, u_{(-1,-1)}\}$.

In either case, given j there are at most six choices for each element of E, so

$$\mathbb{P}_p\big(E \subset A \text{ and } [S \cup (A \cap \tilde{S})] = \tilde{S}\big) \leqslant (6p)^3 \exp\big(-jg(bq)\big)\,.$$

Moreover, in either case, by the induction hypothesis and Harris' inequality, we have

$$\mathbb{P}_p\Big(\big\{A \cap B = \varnothing\big\} \cap D_2^{\hat{\mathbf{x}}}(\hat{S},R)\Big) \leqslant (1-p)^{2(a-s)+b-t} \cdot \hat{F}^{x-1,y-2,z}(s-j-3,t-4)\,,$$

where $B = B_{(0,1)}(S,R) \cup B_{(0,-1)}(S,R) \cup B_{(-1,0)}(S,R)$ and $\hat{\mathbf{x}} = \mathbf{0}$. The probability in this case can therefore be bounded by

$$\sum_{j=1}^{s} (6p)^3 (1-p)^{2(a-s)+b-t} e^{-jg(bq)} \cdot \hat{F}^{x-1,y-2,z}(s-j-3,t-4) \leqslant \frac{\delta}{C^2} \cdot \hat{F}^{x,y,z}(s,t)\,,$$

since $s \leqslant 4\delta/q$ and $e^{2sq+tq+4g(aq)+3g(bq)} \leqslant 2$. This completes the proof of Lemma 10.A.2. □

[14] Note that here (and also below) we could have gained substantially by dividing into cases, as above, depending on whether or not $z(\hat{S},R) < z(S,R)$. In this case, however, this weaker bound will suffice.

Part III

Percolation

Chapter 11

Generalised oriented site percolation

This chapter is based on joint work with Réka Szabó [222]. It concerns GOSP discussed in Section 1.5.4, though its relation to bootstrap percolation is unimportant for the present chapter.

11.1 Introduction

Oriented percolation on $\mathbb{Z}^d$ is a classical model in probability theory and statistical physics, whose behaviour is relatively well understood with many of the main advances on the subject dating back to the 1980s (see [131, 224, 257, 258] for comprehensive expositions). It is also essentially equivalent to the well-known contact process, but also linked to many other models and often used as a tool in proofs.

In this chapter we study the supercritical phase of a natural generalisation of oriented site percolation on $\mathbb{Z}^d$ with arbitrary finite neighbourhood, which we define next. Our goal is to examine the importance of symmetry and planarity to the qualitative behaviour of oriented percolation. The generalisation is further motivated by its relations with probabilistic cellular automata and bootstrap percolation, as discussed in Section 11.2.

11.1.1 Model

Our model of interest is *generalised oriented site percolation* (GOSP) on $\mathbb{Z}^d$ for $d \geqslant 2$. The model is defined by a *neighbourhood*—a finite set $X \subset \mathbb{Z}^d \backslash \{o\}$ (o shall denote the origin of $\mathbb{Z}^d$) with $|X| \geqslant 2$ such that

$$\exists u \in \mathbb{R}^d, \forall x \in X: \quad \langle x, u \rangle > 0, \tag{11.1}$$

which ensures the orientation of the model, and a *parameter* $p \in [0, 1]$. For convenience we will always assume that $u = e_d$, where $(e_i)_{i=1}^d$ denotes the canonical basis of $\mathbb{R}^d$ and that the group generated by X is $\mathbb{Z}^d$. This can be achieved by an invertible linear transformation of $\mathbb{Z}^d$ and, possibly, a restriction to a sublattice. We denote by $\mathbb{P}_p$ the product Bernoulli measure of parameter p on $\mathbb{Z}^d$. The *configuration* $\omega \in \Omega = \{0,1\}^{\mathbb{Z}^d}$ is assumed to be distributed according to this measure. We endow the vertex set $\mathbb{Z}^d$ with the locally finite translation-invariant oriented graph structure with edge set $\{(a, a+x) : a \in \mathbb{Z}^d, x \in X\}$ generated by X (see Fig. 11.1). We refer to this graph as $\mathbb{Z}^d$ when X is clear from the context and $\mathcal{G}_X$ otherwise.

One can naturally identify $\omega \in \Omega$ with the set of *open sites* $\{x \in \mathbb{Z}^d : \omega_x = 1\} \subset \mathbb{Z}^d$, all other sites being *closed*. The open sites induce a subgraph of $\mathcal{G}_X$ by keeping all edges between open sites. We can then introduce the following variant of the natural notion of being connected in this graph. For any $a, b \in \mathbb{Z}^d$ we say that *a infects b* (there is a path from a to b) and write $a \to b$ for the event that there exists a sequence of open vertices $a_1, \dots, a_m = b$ such that $a_1 - a \in X$ and $a_i - a_{i-1} \in X$ for all $i \in [2, m]$. Note that we do not require for a to be open in order for $a \to b$ to occur. We make this choice so that $a \to b$ and $b \to c$ are independent for all $a, b, c \in \mathbb{Z}^d$.

For any $B \subset \mathbb{Z}^d$ we further define $a \xrightarrow{B} b$ as $a \to b$ but with $a_i \in B$ for $i \in [1, m]$. We write $a \xrightarrow{B} \infty$ for the existence of infinitely many b such that $a \xrightarrow{B} b$ and similarly for $\infty \xrightarrow{B} b$. We further extend the notation by defining the event $C \xrightarrow{B} D$ for $B, C, D \subset \mathbb{Z}^d$ as $\exists c \in C, \exists d \in D$ such that $c \xrightarrow{B} d$. We say that *C percolates* in B if $C \xrightarrow{B} \infty$. We define the *order parameter*

$$\theta(p) = \mathbb{P}_p(o \to \infty),$$

the *critical probability*

$$p_{\mathrm{c}} = p_{\mathrm{c}}(X) = \inf\{p > 0 : \theta(p) > 0\}$$

and say that *there is percolation* at p if $\theta(p) > 0$ (by ergodicity this is equivalent to the a.s. existence of an infinite open path). Depending on the value of p, we may speak of *subcritical*, *critical* and *supercritical* regimes. We focus on the study of the supercritical phase, where $\theta(p) > 0$.

It is convenient to view the last coordinate of $\mathbb{Z}^d$ as the time in an interacting particle system. We therefore usually denote points in $\mathbb{Z}^d$ by (x, t) with $x \in \mathbb{Z}^{d-1}$ and $t \in \mathbb{Z}$. Let $R = \max\{t \in \mathbb{Z} : (x, t) \in X\}$ be the *range* of X. Consider the slab $S_t = \mathbb{Z}^{d-1} \times (\mathbb{Z} \cap [t, t+R))$ of width R with normal vector e_d, which we call *time slab*, and denote $S = S_0$. Given an *initial condition* $A \subset S$ and a domain $B \subset \mathbb{Z}^d$, which we omit if $B = \mathbb{Z}^{d-1} \times [R, \infty)$, the *state at time* $t \in \mathbb{N}$ is

$${}_B\xi_t^A = \left\{ b \in S : \exists a \in A, a \xrightarrow{B} b + te_d \right\},$$

so $({}_{\mathbb{Z}^{d-1} \times [R,\infty)}\xi_t^A)_{t=0}^{\infty} = (\xi_t^A)_{t=0}^{\infty}$ is a Markov chain with state space $\{0,1\}^S$. For simplicity if $A = \{o\} \subset S$, we write simply o instead of A. Finally, in

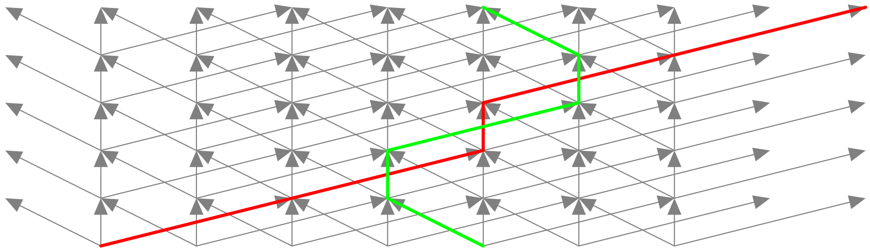

Figure 11.1 – The graph $\mathcal{G}_X$ on $\mathbb{Z}^2$ for $X = \{(-1,1),(0,1),(2,1)\}$. The two thickened paths cross although there is no common vertex nor an edge pointing from one to the other. Here $S = \mathbb{Z} \times \{0\}$, since $R = 1$.

the supercritical phase it is useful to define

$$\bar{\mathbb{P}}_p = \mathbb{P}_p(\cdot | \forall t \geqslant 0, \xi_t^o \neq \varnothing).$$

11.1.2 Examples

Standard *oriented percolation* in 2 dimensions (2dOP) can be defined by $X = \{(0,1),(1,1)\}$. However, we more customarily consider $X = \{(-1,1),(1,1)\}$ instead, which only spans half of $\mathbb{Z}^2$, but we will mostly disregard this minor detail. We denote by $p_{\mathrm{c}}^{\mathrm{OP}}$ the critical probability of 2dOP. In higher dimensions the situation is more ambiguous and at least the choices $X = \{e_i : i \in \{1,\dots,d\}\}$; $X' = \{e_d + \varepsilon e_i : i \in \{1,\dots,d-1\}, \varepsilon \in \{-1,1\}\}$ and $X'' = X' \cup \{e_d\}$ for the neighbourhood could be legitimately called *d-dimensional oriented percolation* (ddOP). For concreteness, we will use ddOP to refer to X'' and simply OP for generic statements.

As a prototype example of neighbourhood which is not covered by the classical approach, but handled here, we retain the two-dimensional GOSP defined by $X = \{(-1,1),(0,1),(2,1)\}$ (see Fig. 11.1). It exhibits the two main additional difficulties of GOSP w.r.t. 2dOP: lack of symmetry w.r.t. the vertical axis and the non-planarity. The latter property is witnessed by the fact that paths may jump over each other without intersecting (see Fig. 11.1).

11.1.3 Results

Denote by $t^A(x)$ the *hitting time* of $x \in \mathbb{Z}^{d-1}$ from A:

$$t^A(x) := \min\left\{t : (x,0) \in \xi_t^A\right\}, \tag{11.2}$$

and define the following subsets of S:

$$H_t^A := \left\{(x,s) \in S : t^A(x) \leqslant t - s\right\}, \tag{11.3}$$

$$K_t^A := \left\{(x,s) \in S : \xi_t^A(x,s) = \xi_t^S(x,s)\right\}, \tag{11.4}$$

which we refer to as *hit* and *coupled regions* with initial condition A respectively. We omit A if it is o. Our main result is the following.

Theorem 11.1.1. *Consider a GOSP in any dimension $d \geqslant 2$. For any $p > p_c$ there exists a deterministic convex compact set $U = U(p) \subset \mathbb{R}^{d-1}$ with non-empty interior such that for all $\varepsilon > 0$, $\bar{\mathbb{P}}_p$-a.s., for every t large enough it holds that*

$$H_t \cap K_t \supset (((1-\varepsilon)tU) \times [0,R)) \cap \mathbb{Z}^d, \tag{11.5}$$

$$\xi_t^o \subset (((1+\varepsilon)tU) \times [0,R)) \cap \mathbb{Z}^d. \tag{11.6}$$

The function $p \mapsto U(p)$ is continuous on $(p_c, 1]$ for the Hausdorff distance on non-empty compact subsets of $\mathbb{R}^{d-1}$. Furthermore, for any open set $O \subset U$, considering the cone $C = \bigcup_{t>0}(tO \times \{t\})$, we have

$$\mathbb{P}_p\left(\exists x \in C, x \xrightarrow{C} \infty\right) = 1. \tag{11.7}$$

Our second result provides more precise information in the near-critical regime in two dimensions.

Theorem 11.1.2. *For GOSP in two dimensions there exists $v \in \mathbb{R}$ such that*

$$\bigcap_{p>p_c} \mathring{U}(p) = \{v\},$$

where $\mathring{U}(p)$ is the interior of the limit shape from Theorem 11.1.1.

11.1.4 Organisation

The chapter is structured as follows. In Section 11.2 we provide the background for this chapter. In Section 11.3 we gather some preliminaries and notation. Sections 11.4 and 11.5 contain the proofs of Theorems 11.1.1 and 11.1.2 respectively.

The proofs are quite long and involve numerous intermediate results of independent interest. Inevitably, some of the steps are already known or require little or no new input as compared to existing arguments for OP or for the contact process. Nevertheless, we choose to also present these steps (without their proofs), so that the new ingredients we provide can be fitted into the global strategy and the reader is not obliged to scour the vast and entangled literature for all the "well-known" ingredients necessary. Moreover, in order not to disturb the flow of reasoning and to single out the novel contributions, we gather them in Section 11.A. Hence, specialists aware of classical results in two and more dimensions and of more recent developments around shape theorems may be able to directly consult Section 11.A.

11.2 Background

We only discuss the supercritical regime, which is the focus of this chapter. Let us begin by emphasising that Theorems 11.1.1 and 11.1.2 are both known for OP and so are all intermediate results featuring in their proofs. More precisely, in the case of ddOP Eqs. (11.5) and (11.6) are due to [132–134]; the continuity in Theorem 11.1.1 was only recently established in [168], based on [166, 167]; Eq. (11.7) was proved in [113, Chapter 5]. Correspondingly, Theorem 11.1.2 for 2dOP was established in [134] (see also [131, 258]).

Following progress on OP, natural generalisations similar to GOSP have often been considered. For the sake of comparability, in the present discussion, we focus on the most restrictive interesting case: GOSP with $X \subset \{(a, 1) : a \in \mathbb{Z}^{d-1}\}$, like the example of Fig. 11.1. These models exhibit the main difficulties inherent to GOSP and are known as *percolation probabilistic cellular automata* (PPCA), 2dOP being called *Stavskaya's PCA* in this context [331, 335, 336, 348–350].

Bezuidenhout and Gray [50] adapted the well-known renormalisation scheme of Bezuidenhout and Grimmett [51] to show that in any dimension PPCA (and more general models) do not percolate at criticality. Their renormalisation will be the starting point of the proof of Theorem 11.1.1. In two dimensions an attempt at proving Theorem 11.1.2 and related results for PPCA (and more general models) was made by Durrett and Schonmann [136], themselves building on [134]. Unfortunately, they imposed a restrictive technical assumption amounting to assuming that X consists of consecutive sites. These neighbourhoods precisely lack the two main obstacles of GOSP—asymmetry and paths jumping over each other (see Fig. 11.1). Furthermore, unaware of their work, Taggi [335, 336] claimed results for PPCA in two dimensions based on [134], as outlined in [131]. Owing to non-planarity, his proof is only correct for neighbourhoods of consecutive sites. As we will see, [335, Theorem 2.2] does indeed hold for all PPCA (and, more generally, GOSP), but requires a different treatment either based on higher dimensional techniques or on our enhancement of the approach of Durrett–Schonmann used to prove Theorems 11.1.1 and 11.1.2 respectively.

Let us note that GOSP are a particular case of *bootstrap percolation* (see Chapter 12 and [315]). As established in Chapter 12 (see particularly Remark 12.5.7 there), Theorems 11.1.1 and 11.1.2 on GOSP can be used to obtain results for more general bootstrap percolation models, particularly in conjunction with quantitative bounds on the limit shape U, as discussed in the first arXiv version of the present chapter [221, Section 5.7]. Other related models and generalisations of OP, to which much of the present approach applies can be found in [50, 130, 133, 364] (also see [121]).

11.3 Preliminaries

11.3.1 Duality

An important property of GOSP is that it is "nearly" self-dual (see [258, 334] for background on duality). The dual of GOSP with neighbourhood X can be thought of as a GOSP with paths moving "backwards" in time. More precisely, write $a \rightsquigarrow b$ if there exist $m \geqslant 0$ and $(a_i)_{i=0}^m$ with $a_0 = a$ and $a_m = b$ such that for all $0 \leqslant i < m$ we have $a_i \in \omega$ and $a_i - a_{i+1} \in X$. In other words, $a \to b$ iff $b \rightsquigarrow a$. Note that there are two differences with $b \to a$. Firstly, the steps are reversed: $a_{i+1} - a_i \in -X$. Secondly, for the dual connections we require that the initial site is open instead of the final one. Based on this notion we define the *dual process* $(\tilde{\xi}_t^A)$ again with state space S but time coordinate $-e_d$. We draw the reader's attention to the fact that this process does not have the same law as the primal process (ξ_t^A), for instance

$$\tilde{\theta}(p) = \mathbb{P}_p(o \rightsquigarrow \infty) = p\mathbb{P}_p(o \to \infty) = p\theta(p).$$

However, up to such minor amendments all our results apply equally well to the dual process and we will use them as needed without systematically stating them.

11.3.2 The contact process

OP is closely related to the *contact process* (CP) [203, 257, 258]. The latter is often used to model epidemics on a graph: vertices are individuals, which can be healthy or infected. In this continuous time Markov dynamics infected individuals recover with rate 1 and transmit the infection to each neighbour with rate $\lambda > 0$ (*infection rate*). The CP admits a well-known graphical construction that is a space-time representation [258]. We assign to each vertex v and ordered pair (u, v) of neighbours independent Poisson point processes D_v with rate 1 and $D_{(u,v)}$ with rate λ respectively. For each atom t of D_v we place a "recovery mark" at (v, t) and for each atom of $D_{(u,v)}$ we draw an "infection arrow" from (u, t) to (v, t). An *infection path* is a connected path moving in the increasing time direction without crossing recovery marks, but possibly jumping along infection arrows in the direction of the arrow. Starting from a set of initially infected vertices A, the set of *infected* vertices at time t is the set of vertices v such that (v, t) can be reached by an infection path from some $(u, 0)$ with $u \in A$.

This representation can be thought of as a continuous time version of OP with infection paths in CP corresponding to paths in OP. Several of the results presented below are originally stated for CP but their proofs transfer to discrete models with the following very minor adaptations.

Firstly, setting $\gamma = \max(\|x\|/t : (x, t) \in X)$, we clearly have that $o \to (x, t)$ implies $\|x\| \leqslant \gamma t$, so, just like for CP, influence can spread at most linearly in time.

Secondly, since the group generated by X is $\mathbb{Z}^d$, for all $n > 0$ there exist a time t and $v \in \mathbb{Z}^{d-1}$ such that

$$\mathbb{P}_p\left(\xi_t^o \supset v + B_n\right) > 0, \tag{11.8}$$

where $B_n = ([-n, n)^{d-1} \times [0, R)) \cap \mathbb{Z}^d$. This is the analogue of the fact that with positive probability the CP infects an arbitrarily large box in unit time.

Finally, for the CP one often needs to control the time an infection path spends at a vertex: either to ensure that it does not stay long at a vertex before jumping or that the path spends at least δt time at a vertex during a time interval of length t. The first assertion is trivial in discrete time as a path "jumps immediately" to the next vertex, and the discrete-analogue of the second assertion is visiting a vertex at least $\lceil \delta t \rceil$ times in a time interval of length t.

11.4 Proof of Theorem 11.1.1

Throughout Sections 11.4 and 11.5, proofs are usually omitted altogether when they only require minor changes (including those of Section 11.3.2). Nevertheless, we provide a sketch or at least a vague idea, whenever possible. The proofs requiring new ideas are gathered in Section 11.A.

The present section is structured as follows. In Section 11.4.1 we recall the Bezuidenhout–Grimmett renormalisation and its extension. We next derive several exponential bounds in Section 11.4.2 obtained based on restart arguments for later use. Section 11.4.3 then completes the proof of the asymptotic shape result and its continuity from Theorem 11.1.1. Finally, Section 11.4.4 puts together all of the above with some further large deviation results to prove the percolation in restricted regions of Theorem 11.1.1. New ingredients needed in Sections 11.4.2 to 11.4.4 are deferred to Sections 11.A.1 to 11.A.3.

11.4.1 Bezuidenhout–Grimmett renormalisation

We begin the study of the supercritical phase by briefly describing the well-known Bezuidenhout–Grimmett (BG) renormalisation. It was first introduced in [51] for the CP on $\mathbb{Z}^d$ (see also [257, Sec. I.2]) and later generalised for translation-invariant finite-range attractive[1] spin systems on $\mathbb{Z}^d$ by Bezuidenhout and Gray [50], the latter reference being the more relevant for us. It is a construction that allows us to compare GOSP and 2dOP. The main idea is to show that GOSP when restricted to a sufficiently thick two-dimensional space-time slab dominates a supercritical 2dOP, which in turn

[1] A spin system is *attractive* if adding extra sites in the initial condition only makes more sites infected by it at any later time (see e.g. [257, Sec. I.1.]), as in the case of GOSP.

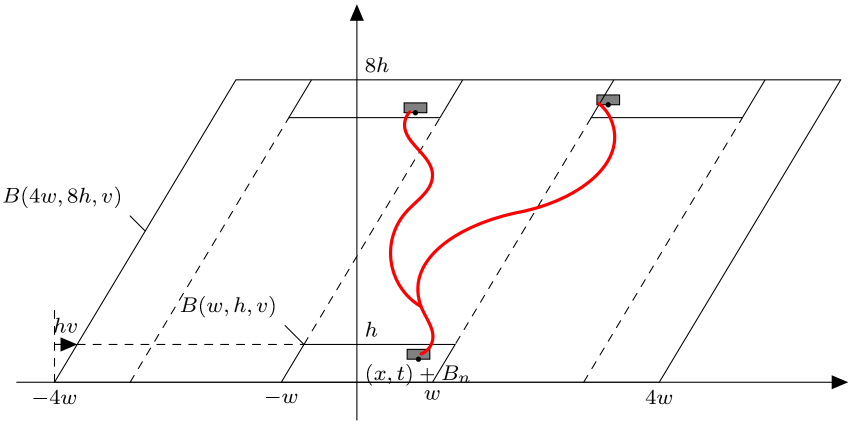

Figure 11.2 – The event described in Theorem 11.4.1 for $d = 2$. Note that in this case w and v are one-dimensional.

implies that if there is percolation in 2dOP, then there is percolation in this restricted region. As all the results below are already known for 2dOP, this will entail numerous consequences for GOSP.

Before proceeding to the renormalisation we will need a few geometric definitions. Our basic *box* is $B_n = [-n, n)^{d-1} \times [0, R)$ for natural n, recalling R from Section 11.1.1 (although many of our regions will be defined as subsets of $\mathbb{R}^d$, we systematically refer to the integer points in them). For $w \in \mathbb{Z}^{d-1}$, $h \in \mathbb{Z}$ and $v \in \mathbb{R}^{d-1}$ we further introduce the *block* (see Fig. 11.2)

$$B(w, h, v) = [0, h)(v, 1) + \prod_{i=1}^{d-1}[-w_i, w_i) \times \{0\} \subset \mathbb{R}^d, \qquad (11.9)$$

so that $B_n = B(\{n\}^{d-1}, R, 0)$. Note that if the model is symmetric we can always assume $v = 0$.

The key theorem allowing comparison between the two models is as follows.

Theorem 11.4.1. *If p is such that $\theta(p) > 0$, then for all $\varepsilon > 0$ there exist positive integers n, h and vectors w, v with $n < w_i$ for all i and $h > R$ such that if $(x, t) \in B(w, h, v)$, then*

$$\mathbb{P}_p\big(\exists (y, s) \in B(w, h, v) + 7h(v, 1) \pm 2we_{d-1} \text{ such that } (x, t) + B_n \text{ infects } (y, s) + B_n \text{ in } B(4w, 8h, v)\big) > 1 - \varepsilon. \qquad (11.10)$$

In other words we can choose parameters such that when considering the truncated process in $B(4w, 8h, v)$ with high probability a box B_n centered at some (x, t) inside the block $B(w, h, v)$ infects a copy of itself centered at either of the target blocks that are translates of $B(w, h, v)$ (see Fig. 11.2). The proof of this theorem being quite long and technical, we direct the interested reader to [50], where it is established in a setting essentially including GOSP.

Recall that 2dOP is defined by $X = \{(-1,1),(1,1)\}$. We denote by ζ_k^A the set of (even) sites in $\mathbb{Z}^2$ with second coordinate k which are infected by 2dOP with initial condition A. We are now ready for the 2dOP comparison of [51].

Theorem 11.4.2 (BG renormalisation)**.** *Fix $q < 1$ and assume that p is such that Eq. (11.10) holds for $\varepsilon > 0$ sufficiently small depending on q and some n, h, w, v as in Theorem 11.4.1. Then the following holds for some n, h, w, v. For any initial condition $A \subset S$ we denote*

$$A' = \{j \in 2\mathbb{Z} : \exists (x,s) \in B(w,h,v) + 2w_{d-1}je_{d-1}, (x,s) + B_n \subset A\}.$$

Then there exists a coupling of 2dOP $\zeta^{A'}$ of parameter q and GOSP ξ^A such that for all $j \in \mathbb{Z}$ and k

$$j \in \zeta_k^{A'} \text{ implies that } (x,0) + B_n \subset \xi_t^A \quad \text{for some } (x,t) \in B(w,h,v) + 7hk(v,1) + 2w_{d-1}je_{d-1}.$$

In particular, $\theta(p) > 0$.

Informally, each site of 2dOP corresponds to a translate of the block $B(w,h,v)$ in GOSP. We can couple the two processes so that if a site in 2dOP is in the cluster of a vertex in ζ^A, then there is a box infected by A in the block corresponding to that site in the GOSP.

The proof of Theorem 11.4.2 is as in [257, Theorem I.2.23] (also see [51]). Indeed, one may construct the coupling by induction as follows. If $j \in \zeta_k$, then there is an infected copy of the box B_n in the block corresponding to j, k, so we may apply the result of Theorem 11.4.1 to get that with probability $1-\varepsilon$ there will also be such infected boxes in the blocks corresponding to $j+1, k+1$ and $j-1, k+1$. It is easily seen (as the GOSP configuration is composed of independent variables) that the resulting process is a 1-dependent 2dOP with parameter at least $1-\varepsilon$, so by a standard comparison between 1-dependent and independent percolation [259] we obtain Theorem 11.4.2 as desired.

It is useful to note that the BG renormalisation concerns only certain translates of the block $B(w,h,v)$. However, we may tile $\mathbb{Z}^d$ with disjoint blocks so that each tilted space-time slab of the form

$$\bigcup_{j,k \in \mathbb{Z}} B(w,h,v) + w_{d-1}je_{d-1} + kh(v,1)$$

is formed by 14 disjoint 2dOP lattices of blocks. We may perform the couplings of all the corresponding 2dOP processes with the same GOSP simultaneously as above so that sites in different 2dOP have a finite range dependence, hence they may be made independent by [259]. In total, for $A \subset S$

we can couple ξ^A with independent 2dOP processes naturally indexed by $\mathbb{Z}^{d-2} \times \{1, \ldots, 14\}$ with initial conditions corresponding to the parts of A in each of the 2dOP lattices.

11.4.2 Restart arguments

Recall that ξ_t^A is the set of sites infected by A at time t. The *extinction time of A* is the absorption time of the chain started at A, that is

$$\tau^A = \min\left\{t \geqslant 0 : \xi_t^A = \varnothing\right\} \in \{0, 1, \ldots\} \cup \{\infty\}. \tag{11.11}$$

We say that the process started from A *dies out* if $\tau^A < \infty$ and *survives* otherwise.

The BG renormalisation allows us to use a so-called *restart argument* that can be described as follows. We let GOSP (ξ_t^A) evolve until we find an infected translate of the box B_n (which by Eq. (11.8) has a strictly positive probability, thus it happens after at most a geometrically distributed number of steps) or the process dies out. If we infect a box, we start the 2dOP process (ζ_k) of Theorem 11.4.2 (with appropriate initial condition) from the corresponding block coupled with (ξ_t^A). If (ζ_k) percolates, then (ξ_t^A) percolates as well. If (ζ_k) dies out and (ξ_t^A) still survives, we restart the procedure. We repeat this until either the GOSP dies out or the renormalisation yields a percolating 2dOP (since the parameter q of (ζ_k) is supercritical, $q > p_c^{\mathrm{OP}}$, this will happen after at most a geometric number of trials).

We will use this technique to transfer properties from 2dOP to GOSP. The exponential bounds we prove next in this section (like all other results) are already established for 2dOP [134] and the d-dimensional CP [132]. Recall that τ^A is the extinction time of the set A.

Theorem 11.4.3 (Exponential death bounds)**.** *For every $p > p_c$ there exists a constant $\varepsilon = \varepsilon(p) > 0$ such that for all $A \subset S$ and $t \geqslant R$ it holds that*

$$\mathbb{P}_p\left(t \leqslant \tau^A < \infty\right) \leqslant e^{-\varepsilon t}, \tag{11.12}$$

$$\mathbb{P}_p\left(\tau^A < \infty\right) \leqslant e^{-\varepsilon |A|}. \tag{11.13}$$

The proof of Theorem 11.4.3 goes along the same lines as the proof of [257, Theorem I.2.30], using the restart argument. For Eq. (11.12), on $\{\tau^A < \infty\}$, we can bound τ^A by the sum of the number of steps until we find an infected box and the survival time of the coupled 2dOP in each round. As 2dOP satisfies Eq. (11.12) and the required quantities of the restart argument are bounded by geometric random variables, we get the desired exponential decay.

For Eq. (11.13) first tile $\mathbb{Z}^d$ with the disjoint translates of the space-time slab

$$T = \bigcup_{j,k \in \mathbb{Z}} B(4w, 7h, v) + 4w_{d-1} j e_{d-1} + 7hk(v, 1)$$

and consider the processes restricted to these slabs with initial conditions corresponding to the parts of A in each slab. Observe that these processes are independent and τ^A can be bounded from below by the maximum of their extinction times. Therefore, it is enough to show the analogue of Eq. (11.13) for $A \subset T$. As in Eq. (11.8) we can show that there exists a t depending on w and h, but not on A, such that with strictly positive probability every vertex in $A \subset T$ can infect a box B_n in

$$t(v,1) + \bigcup_{j\in\mathbb{Z}} B(w,h,v) + w_{d-1} j e_{d-1}.$$

Thus, (save for an exponentially unlikely event) for some $\delta > 0$ at least $\delta|A|$ disjoint blocks at the same time contain an infected box. We start the 2dOP process of Theorem 11.4.2 from all these blocks. Observe that $\{\tau^A < \infty\}$ can only happen if all the coupled 2dOP process dies, but since Eq. (11.13) holds for 2dOP, this has exponentially small probability in the size of A.

Remark 11.4.4. Equation (11.12) implies that the law of ξ_t^S converges to the upper invariant measure of the process (corresponding to the distribution of sites $x \in S$ such that $x \rightsquigarrow \infty$) exponentially fast in t.

The next result is the analogue of condition (a) of Lemma 5.1 in [132].

Theorem 11.4.5. *Let ξ and $\tilde{\xi}$ be independent primal and dual GOSP. Then for every $p > p_c$ there exist constants $\varepsilon, c, C > 0$ and a vector $v \in \mathbb{R}^{d-1}$ depending on p such that for all integer $t > 0$ and $A, B \subset S$ satisfying $\max_{a\in A, b\in B} \|a-b\| < ct$ we have*

$$\mathbb{P}_p\left(\xi_t^A \cap \tilde{\xi}_t^{B+(2tv+z_t,0)} = \varnothing, \xi_t^A \neq \varnothing, \tilde{\xi}_t^{B+(2tv+z_t,0)} \neq \varnothing\right) \leqslant Ce^{-\varepsilon t}, \quad (11.14)$$

where $z_t \in \mathbb{R}^{d-1} \times \{0\}$ is such that $2tv + z_t \in \mathbb{Z}^{d-1}$ and $\|z_t\|^2 \leqslant (d-1)/4$.

It is important to note that due to the lack of symmetry this result is more technical for GOSP than for ddOP. We leave the proof to Section 11.A.1 and only indicate that it relies mainly on the BG renormalisation, restart argument, Eq. (11.12) and several properties known for 2dOP, which will be established below for GOSP.

Remark 11.4.6. We can use Theorem 11.4.5 to prove that the infinite cluster is unique. Together with $\theta(p_c) = 0$ following from Theorems 11.4.1 and 11.4.2, this customarily yields that $\theta : p \mapsto \theta(p)$ is continuous on $[0,1]$. This was first established for ddOP in [197].

Recall the hit and coupled regions of Eqs. (11.3) and (11.4).

Theorem 11.4.7 (At least linear growth). *For every $p > p_c$ there exist constants $\varepsilon, c > 0$ and a vector $v \in \mathbb{R}^{d-1}$ depending on p such that for all $t > 0$ and $x \in \mathbb{Z}^{d-1}$ such that $\|x - vt\| < ct$ it holds that*

$$\mathbb{P}_p\left((x,0) \notin H_t, \tau^o = \infty\right) \leqslant e^{-\varepsilon t}, \tag{11.15}$$

$$\mathbb{P}_p\left((x,0) \notin K_t, \tau^o = \infty\right) \leqslant e^{-\varepsilon t}. \tag{11.16}$$

This result is also a consequence of Theorem 11.4.5. The proof is an adaptation of the proof of conditions (c) and (d) of Theorem 5.2 in [132] for the CP (see also [133]). For Eq. (11.15) recall the definition of the dual process $\tilde{\xi}$ from Section 11.3.1 and note that for any $0 \leqslant s \leqslant t$ the event $\{\exists y \in S, o \to y + se_d, (x,t) \rightsquigarrow y + se_d\}$ is equivalent to $o \to (x,t)$. By a restart argument we can find a time r such that $(x,r) \rightsquigarrow \infty$ and up to an exponentially unlikely event we can take $0 \leqslant t - r < ct$. Then the conclusion follows from Eq. (11.14) with $A = \{o\}$, $B = \{(x - 2rv, 0)\}$ and t in the theorem equal to $r/2$.

For Eq. (11.16) observe that for $\omega \subset \mathbb{Z}^{d-1} \times \{R, R+1, \ldots\}$

$$\{(x,0) \notin K_{2t}, \tau^o = \infty\} = \{o \not\to (x,2t), (x,2t) \rightsquigarrow S, o \to \infty\}.$$

The conclusion then follows again from Eq. (11.14) with $A = \{o\}$ and $B = \{(x - 4tv, 0)\}$.

Finally, for completeness let us mention a consequence of the restart arguments concerning GOSP on tori. Recall that we assume that the direction $u = e_d$, and let $\mathbb{T}_n^{d-1} = (\mathbb{Z}/n\mathbb{Z})^{d-1}$ denote the $(d-1)$-dimensional discrete torus of side n. Consider GOSP on the graph with vertex set $\mathbb{T}_n^{d-1} \times \mathbb{Z}$ obtained as the quotient of $\mathcal{G}_X$. The *extinction time* is defined as

$$\tau^{\mathbb{T}} = \sup\left\{t \geqslant 0 : \mathbb{T}_n^{d-1} \times \{0, -1, -2, \ldots\} \to \mathbb{T}_n^{d-1} \times \{t\}\right\}.$$

Corollary 11.4.8. *For all $p < p_c$ there exists $c(p)$ such that*

$$\frac{\tau^{\mathbb{T}}}{\log n} \xrightarrow[n\to\infty]{\mathbb{P}_p} \frac{d-1}{c(p)}, \tag{11.17}$$

and for all $p > p_c$ there exist $c, C \in (0, \infty)$ such that

$$\frac{\tau^{\mathbb{T}}}{\mathbb{E}_p[\tau^{\mathbb{T}}]} \xrightarrow[n\to\infty]{(\mathrm{d})} \mathcal{E}, \tag{11.18}$$

$$e^{cn^{d-1}} < \mathbb{E}_p\left[\tau^{\mathbb{T}}\right] < e^{Cn^{d-1}}, \tag{11.19}$$

for n large enough, where $\mathcal{E}$ is the standard exponential distribution.

Equation (11.17) follows as in [135] (see also [257, Theorem I.3.3]) from the subcritical result established in [6, 276]: for all $p < p_c$ there exists $c(p)$ such that

$$-\lim_{t\to\infty} \frac{1}{t} \log \mathbb{P}_p(\tau^o \geqslant t) = c(p) > 0. \tag{11.20}$$

Equation (11.18) was proved for the CP in two dimensions in [312] (see also [137]), while in d dimensions this was done in [285] (also see [327] for subsequent development). Equation (11.19) was proved in [135] for the two-dimensional CP and in [103] in d dimensions. The proofs rely on Theorems 11.4.2, 11.4.3 and 11.4.7 (see [221, App. A.4] for a sketch following [285]).

Let us note that for the CP on a finite box (in our setting this corresponds to cutting the bonds crossing the boundary of a fundamental domain of the torus) in [137, 287] it was established that in fact $\log \mathbb{E}_p[\tau]/n^{d-1}$ converges as $n \to \infty$. However, for GOSP considering a box is either inappropriate or requires tiling the lattice first, making the result somewhat unnatural and unhandy due to the implicit definition of the tilting direction, which may even depend on p as $p \to p_{\mathrm{c}}$. It would appear that proving the existence of the above limit on the torus is unknown even for OP and CP.

11.4.3 Asymptotic shape

With the results of Section 11.4.2 at hand we are ready to prove the asymptotic shape theorem and the continuity of the limit shape, that is Eqs. (11.5) and (11.6) and the continuity result of Theorem 11.1.1. These results are known for the CP. However, certain issues arise due to the possibility that the model may have a "drift," e.g. if the convex envelope of the neighbourhood X does not intersect the line $\mathbb{R}e_d$. This problem is absent if we can take $v = 0$ in Theorem 11.4.7. For simplicity, in Section 11.4.3 we only briefly recall the arguments used to prove the desired results under this additional assumption, leaving out the minor changes described in Section 11.3.2. Thus, we leave the new input needed for removing the assumption $v = 0$ to Section 11.A.2.

It was proved in [133] for permanent one-site growth procesess (translation invariant, attractive processes with local rules, with $\varnothing$ absorbing state and positive probability of survival) that the exponential estimates from Eqs. (11.12), (11.15) and (11.16) with $v = 0$ imply the shape theorem: Eqs. (11.5) and (11.6). The idea is that, given these estimates, the hitting times are subadditive, stationary and integrable. Then, using subadditive ergodic theory [243], one can prove that for $x \in \mathbb{Z}^{d-1}$

$$\frac{t(nx)}{n} \to \mu(x) \quad \bar{\mathbb{P}}_p\text{-a.s.} \tag{11.21}$$

The *time constant* $\mu(x)$ can be extended into a norm on $\mathbb{R}^{d-1}$ with unit ball U, yielding the result for the hit region. Then we can argue that there are a lot of vertices around the boundary of the cone defined by U that are reached from the origin and by Eq. (11.12) survive. Using Eq. (11.16) we can conclude that the union of the coupled regions of these vertices eventually covers $(1 - \varepsilon)tU$.

Our next goal is to prove that the limit shape U is continuous in p. For this, we will require a quantity called essential hitting time. It was first introduced by Garet and Marchand in [166], inspired by [248], to prove shape theorems in a more difficult setting. Using this notion, they later proved large deviation inequalities [167] and continuity of the asymptotic shape [168]. We next discuss these results still under the assumption that $v = 0$ in Theorem 11.4.7.

Roughly speaking (see [166] for the correct definition), under $\bar{\mathbb{P}}_p$ the *essential hitting time* $\sigma(x)$ of $x \in \mathbb{Z}^{d-1}$ is a time such that $o \to (x, \sigma(x)) \to \infty$. Crucially, the essential hitting time is nearly subadditive [166, Theorem 2]. Using this property, one can show that $\bar{\mathbb{P}}_p$-a.s., as $n \to \infty$, $\sigma(xn)/n$ converges. Controlling the discrepancy between the essential hitting time and the hitting time [166, Proposition 17], we can conclude that the limit is also the one of $t(xn)/n$, whose existence is a byproduct of [133]. This control of $\sigma(x) - t(x)$ further allows us to bound the moments of $\sigma(x)$ under $\bar{\mathbb{P}}_p$ and to get exponential estimates for the essentially hit region analogous to Eq. (11.15) with $v = 0$ [166, Corollary 20 and 21].

Relying on the almost subadditivity of $\sigma(x)$, one may establish large deviation results corresponding to [167, Theorems 1.1 and 1.4], still under the assumption $v = 0$ to be removed in Section 11.A.2.

Theorem 11.4.9. *For every $p > p_c$ and every $\varepsilon > 0$ there exist constants $c, C > 0$ such that for any $x \in \mathbb{Z}^{d-1}$ and $t \geqslant 1$*

$$\bar{\mathbb{P}}_p\left(K_t \cap H_t \supset (((1-\varepsilon)tU) \times [0,R)) \cap \mathbb{Z}^d\right) \geqslant 1 - Ce^{-ct},$$
$$\mathbb{P}_p\left(\xi_t^o \subset (((1+\varepsilon)tU) \times [0,R)) \cap \mathbb{Z}^d\right) \geqslant 1 - Ce^{-ct},$$

where U is as in Theorem 11.1.1.

Finally, one can show the continuity of the limit shape in Theorem 11.1.1 as in [168, Theorem 1], recycling much of the proof of Theorem 11.4.9.

11.4.4 Percolation in restricted regions

Relying on the results of Section 11.4.2, we next establish large deviations for the infinite cluster density, which we then use to prove Eq. (11.7) of Theorem 11.1.1.

Theorem 11.4.10. *Let*

$$Y_n := |\{(x,t) \in B_n : (x,t) \rightsquigarrow \infty\}| / |B_n|. \tag{11.22}$$

For all $p > p_c$ there exists a convex function $\varphi : [0,1] \to [0,\infty)$ such that $\varphi(a) = 0$ if and only if $a = p\theta(p) = \tilde{\theta}(p)$ and for all $a < b$ in $[0,1]$,

$$\lim_{n\to\infty} \frac{\log \mathbb{P}_p(Y_n \in [a,b])}{n^{d-1}} = -\inf_{x\in[a,b]} \varphi(x).$$

Most of this result is very general and holds for any translation invariant attractive spin system, as established in [251]. Roughly speaking, the existence of the limit follows from the fact that if several boxes have $Y_n \geqslant a$, then so does their union; the convexity follows similarly, asking for one box with $Y_n \geqslant x$ and one with $Y_n \geqslant y$ and considering their union. The relevance of $\tilde{\theta}(p)$ comes from cutting a large box into smaller ones and using attractiveness to replace boundary conditions by maximal ones, in order to enable the use of large deviations results for i.i.d. random variables. Indeed, the invariant measure of GOSP in a large box with infected boundary condition still infects sites "far from the boundary" with probability close to $\tilde{\theta}(p)$. This follows from the fact that the upper invariant measure of the infinite volume attractive process must dominate the (decreasing) limit of these invariant measures as the size of the box diverges (see [258, Theorem III.2.7]).

The only somewhat model-specific property is the fact that for any $a < \tilde{\theta}(p)$ there exists $c > 0$ such that for all n we have

$$\mathbb{P}_p(Y_n \leqslant a) \leqslant e^{-cn^{d-1}}. \tag{11.23}$$

This was established for 2dOP in [138]. Unfortunately, the argument is 2-dimensional, so we provide a proof for GOSP in any dimension (and in particular ddOP), which appears to be novel. This is done via a new renormalisation in Section 11.A.3, relying on Theorems 11.4.3 and 11.4.7.

Remark 11.4.11. One can further study fluctuations of the density. For translation invariant attractive spin systems on $\mathbb{Z}^d$ [250] examined when, starting from the upper invariant measure, we reach a value of Y_n smaller than $\tilde{\theta}(p)$. Later this result was extended to upper fluctuations in [161] for the CP. These proofs rely on Theorem 11.4.10 and can be adapted to GOSP. We also direct the reader to [113, Chapter 5] for information regarding the properties and shape of large finite clusters and more large deviations.

Now we are ready to prove a more geometric property of the infinite cluster, Eq. (11.7) of Theorem 11.1.1, establishing that percolation occurs in restricted regions. This result was proved for ddOP in [113, Theorem 1.3 of Chapter 5], but given the results available to us, we may directly retrieve it (for GOSP) from Theorems 11.4.3, 11.4.9 and 11.4.10 as follows. Fixing some $u \in O$ and $\delta > 0$ small, for a site (x,t) at distance at most $\delta^2 t$ from (tu,t) surviving for time δt, by Eq. (11.12) and Theorem 11.4.9 it is likely that its coupled region contains a box of side $\delta^2 t$ centered at $((1+\delta)tu,(1+\delta)t)$. By Theorem 11.4.10, it is likely that at least $\delta^{2d}t^{d-1}$ sites in that box are infected by (x,t) and, since δ is small and (x,t) is at distance of order t from the boundary of C, this has to happen inside C. Finally, by Eq. (11.13), some of those sites is likely to survive. Repeating this procedure to infinity and recalling that the probability of failing at each step is exponentially small in t, we obtain Eq. (11.7) of Theorem 11.1.1, as desired.

11.5 Proof of Theorem 11.1.2

In this section we assume $d = 2$. In this case one can say more about GOSP based on techniques for 2dOP, for which [258, Chapter VI] and [131] are excellent references. In Section 11.5.1 we gather some standard preliminaries. In Section 11.5.2 we introduce an alternative renormalisation technique, whose refinement enables us to prove Theorem 11.1.2.

11.5.1 Edge speed

Define the *right edge* of the process as

$$r_t = \max\left\{x \in \mathbb{Z} : \exists y \in \{0, \ldots, R-1\}, (x, y) \in \xi_t^{S^-}\right\},$$

where $S^- = ((-\infty, 0] \times [0, R)) \cap \mathbb{Z}^2$ is the left half of S. Similarly define the *left edge* l_t as the minimum of x above with $S^+ = ([0, \infty) \times [0, R)) \cap \mathbb{Z}^2$ instead of S^-. It is important to note that, as discussed in Section 11.1.2, although the model is two-dimensional, it is *not* planar and paths may jump over each other without crossing (recall Fig. 11.1). Nevertheless, the right and left edges do retain some of their properties from the 2dOP case.

Theorem 11.5.1 (Edge speed)**.** *For any $p \in [0, 1]$ there exists*

$$\alpha = \lim_{t\to\infty} \frac{\mathbb{E}_p[r_t]}{t} = \inf_{t \geqslant 1} \frac{\mathbb{E}_p[r_t]}{t} \in [-\infty, \infty).$$

Moreover, $r_t/t \xrightarrow{t\to\infty} \alpha$ a.s. and if $\alpha > -\infty$, then $\mathbb{E}_p\left[\left|\frac{r_t}{t} - \alpha\right|\right] \xrightarrow{t\to\infty} 0$. Similar statements hold for $\beta = \lim \mathbb{E}_p[l_t]/t$.

The proof (and statement) is identical to [258, Theorem VI.2.19] and is a consequence of a subadditive ergodic theorem due to Durrett [130] (see particularly Theorem 6.1 thereof). The idea is to introduce a version of the right edge between time s and t which, contrary to r_t, is subadditive in an appropriate sense.

We next show that the two-dimensional approach coincides with the more general one from the previous section.

Theorem 11.5.2. *For any $p > p_{\mathrm{c}}$ the limit shape U from Theorem 11.1.1 and the edge speeds α, β from Theorem 11.5.1 satisfy $U = [\beta, \alpha]$.*

To see this, note that by Theorem 11.4.7 with positive probability $o \to \infty$ and at all times the coupled region is large enough to ensure that the right and left edges are infected by o. We can then conclude, since Theorems 11.1.1 and 11.5.1 are almost sure statements.

A notable advantage of having the edge representation of the limit shape is the following result.

Theorem 11.5.3. *The right edge speed α is strictly increasing on $(p_c, 1)$.*

The proof is very similar to [131, Eq. (12)] and was reiterated in [136] in a setting including PPCA. It proceeds in two steps. Firstly, one shows by a clever but simple algebraic manipulation that adding a vertical column of sites to any initial condition entirely on its right increases $\mathbb{E}_p[r_t]$ by at least 1 for all t. This property only relies on the fact that the process is additive in the sense that $\xi_t^{A\cup B} = \xi_t^A \cup \xi_t^B$ for all A, B, t. Secondly, one observes that if p is increased by a small amount δ, it may happen that the additional vertices opened by increasing it lead precisely to adding such a vertical column in ξ_t to the right of r_t (corresponding to parameter p).

11.5.2 Alternative renormalisation

In two dimensions it is possible to study the supercritical phase via a more elementary renormalisation scheme than the BG one. For 2dOP this approach due to Durrett and Griffeath [134] is used classically to derive most of the results stated above in that setting. However, applying this renormalisation to GOSP (in two dimensions) turns out to be quite tricky. Let us first give the rough lines of the renormalisation before explaining what goes wrong for GOSP and how to address it.

For 2dOP, let us assume that p satisfies $\alpha(p) > \beta(p)$. By Theorem 11.5.1 we have that $r_t/t \to \alpha$ a.s. Moreover, one can show (see [131, 134]) that for all $\varepsilon > 0$ there exists $c > 0$ such that for all $t > 0$

$$\mathbb{P}_p(r_t > (\alpha + \varepsilon)t) \leqslant e^{-ct}. \tag{11.24}$$

We may then establish (see Fig. 11.4) that for L large enough depending on $\varepsilon > 0$, the box $B(\varepsilon L, L, \alpha)$ is crossed from bottom to top by an open path with high probability, namely for $\varepsilon > 0$

$$\lim_{L\to\infty} \mathbb{P}_p\left({}_{B(\varepsilon L,L,\alpha)}\xi_L^{S^-} = \varnothing\right) = 0. \tag{11.25}$$

Indeed, by Eq. (11.24), it is forbidden for the right edge to leave the box on one side; by Theorem 11.5.1 the right edge at time L is likely to be in the middle of the top side of the box; while if the path reaching the right edge at time L leaves the box on the other side, that would imply that the path necessarily went faster than allowed by Eq. (11.24), in order to make up for the delay (the last idea is due to Gray [136]).

Hence, we have that with probability close to 1 long thin boxes with tilting α (and similarly for β) are crossed. This reasoning is perfectly valid for GOSP. In order use such boxes to construct a renormalisation, one places around each renormalised vertex two of them directed by α and β and says the vertex is open if they are crossed by paths (see [131, Fig. 7] or [134, Fig. 1]). For 2dOP it is then clear that if the resulting renormalised 2dOP percolates, then so does the original one. Indeed, one can switch from the path in one box to another as soon as they intersect, which is necessarily the case for planar graphs such as the one associated to 2dOP.

It is not hard to see that the argument remains valid for PPCA with neighbourhood consisting of consecutive sites of the form $(x, 1)$. However, for GOSP with arbitrary neighbourhood X it is no longer true that two paths which "cross" have to intersect in an open point. An attempt to remedy this was made by Durrett and Schonmann [136], whose approach will be of use to us. Yet, when restricted to PPCA, their result only applies to the ones with neighbourhood of consectuive sites as above, making it trivial (their main idea is not needed for those models). As their work is somewhat informal, we indicate that this follows from the restrictive hypothesis (H3) located at the end of Sec. 4 of [136].

Improving the approach of [136] and using Theorem 11.5.3, we outline how to obtain the following result in Section 11.A.4.

Theorem 11.5.4. *If for some $p \in [0, 1]$ We have $\alpha(p) > \beta(p)$, then $\theta(p) > 0$ and*

$$\lim_{p' \to p-} \alpha(p') = \alpha(p) \qquad\qquad \lim_{p' \to p-} \beta(p') = \beta(p).$$

In particular, this implies $\alpha(p_c) \leqslant \beta(p_c)$. On the other hand, Theorem 11.5.1 readily implies that $\alpha(p) \geqslant \beta(p)$ for $p > p_c$. The final ingredient for proving Theorem 11.1.2 is the continuity to the right of α (and β), which also follows from Theorem 11.5.1, since α is the decreasing limit of the continuous non-decreasing functions $\mathbb{E}_p[r_t]$.[2] Combining these properties, we get

$$\lim_{p \to p_c+} \alpha(p) - \beta(p) = 0,$$

which, together with Theorems 11.5.2 and 11.5.3, implies Theorem 11.1.2. We note that in higher dimensions it is unknown whether $\bigcap_{p>p_c} \mathring{U}(p)$ is empty, a singleton or a larger set.

Appendix

11.A Proofs

In this appendix we gather the proofs of the novel steps in the proof of the main results. The following basic result for 2dOP proved by contour arguments (see [131, 138]) will be used several times.

---

[2]Indeed, $\mathbb{E}_p[r_t]$ is the limit of the polynomials $\mathbb{E}_p[\max(r_t, -M)]$ as $M \to \infty$. The limit is uniform for $p \in [p', 1]$ for $p' > 0$, since the negative tail of r_t is bounded by a geometric variable with success rate $(p')^t$. Recalling Eq. (11.20) and $p_c > 0$ (by comparison with branching), we further have $\alpha(p) = -\infty$ and $\beta(p) = \infty$ for $p < p_c$.

Proposition 11.A.1. *For every $\varepsilon > 0$ there exist $c, \delta > 0$ such that for 2dOP with parameter $p \geqslant 1-\delta$ it holds that for all finite $A \subset \mathbb{Z}$ and integer t*

$$\mathbb{P}_p(|\{a \in A, a \to \infty\}|/|A| \leqslant 1-\varepsilon) < e^{-c|A|}$$
$$\bar{\mathbb{P}}_p(|\xi_t^o|/t \leqslant 1-\varepsilon) < e^{-ct}.$$

11.A.1 Primal-dual intersection—proof of Theorem 11.4.5

Recall the notation of Theorem 11.4.5 and Section 11.4.2.

Observe that the BG renormalisation restricts the process to a space-time slab in which all but one space dimension are suppressed. Throughout the chapter we assumed this to be the $d-1^{\text{st}}$ dimension, but we can replace e_{d-1} by any e_i in Theorem 11.4.1 for $i \in \{1,\dots,d-1\}$. If the model is symmetric, the parameters n, w, h, v can be chosen to be the same in all directions, however this is not necessarily the case in general. Let us fix ε and denote the values of n, w, h, v in Theorem 11.4.1 corresponding to e_i by n^i, w^i, h^i and v^i. We then set $v = \sum_{i=1}^{d-1} v^i/(d-1)$. For simplicity we will disregard the offset z_t.

We start by noticing that by Eq. (11.12) we may assume that $\tau^A = \infty$ in ω and $\tilde{\tau}^{B+2tv} = \infty$ in ω translated by $-2te_d$. We can then choose $(x,s) \in A$ and $(\tilde{x},\tilde{s}) \in B + 2t(v,1)$ such that $\tau^{\{(x,s)\}} = \tilde{\tau}^{\{(\tilde{x},\tilde{s}-2t)\}} = \infty$.

We can perform a restart argument starting from (x,s) and $(\tilde{x},\tilde{s})$ until they simultaneously infect a (translate of the) box $B_{n^{d-1}}$ each and that the two boxes give rise to a percolating 2dOP in their respective e_{d-1}-space-time slabs. As in Section 11.4.2 the restart argument is exponentially unlikely to require more than δt steps for some small $\delta > 0$, the positions of the two boxes differ by $2t(v,1) + \Delta$ with $\|\Delta\| = O(c+\delta)t$ (here asymptotic notation is w.r.t. $t \to \infty$).

Fix a large integer K so that Eq. (11.8) holds for $n = n^{d-2}$ and $t = \lfloor Kh^{d-1}/2 \rfloor$. Then by Proposition 11.A.1 we get that at time $t/(7h^{d-1}(d-1)) - K$ both in ζ and $\tilde{\zeta}$ (the renormalised 2dOP corresponding to ξ and $\tilde{\xi}$) infect at least 2/3 of the (renormalised) sites that can be reached from the sites corresponding to the initial two boxes $B_{n^{d-1}}$. By the pigeonhole principle, as c and δ are sufficiently small, there are at least $t/(22h^{d-1}(d-1))$ sites $(z, t/(7h^{d-1}(d-1)) - K)$ which are infected in ζ and such that $(\tilde{z}, t/(7h^{d-1}(d-1)) - K)$ is infected in $\tilde{\zeta}$ with $\tilde{z} - z = -\lfloor \Delta_{d-1}/w_{d-1}^{d-1} \rfloor$. It then follows fromEq. (11.8), Proposition 11.A.1 and the pigeonhole principle that up to an exponentially unlikely event at least one such couple $z, \tilde{z}$ gives rise to two boxes $B_{n^{d-2}} + (y, t/(d-1))$ and $B_{n^{d-2}} + (\tilde{y}, t(2 - 1/(d-1)))$ infected in ξ and $\tilde{\xi}$ respectively such that $\tilde{y} - y = 2tv - 2tv^{d-1}/(d-1) + \sum_{i=1}^{d-2} \Delta_i e_i$ and such that the 2dOP renormalisations in direction e_{d-2} of each of the boxes percolate.

Repeating the same reasoning for each direction and recalling the definition of v, we obtain the desired conclusion.

11.A.2 Tilting

Recall the setting of Section 11.4.3. In this section we show how to remove the additional assumption $v = 0$ used there in the proofs of Eqs. (11.5) and (11.6) and the continuity of U in Theorem 11.1.1, as well as Theorem 11.4.9. The reasoning for Theorem 11.4.9 and the continuity being identical to the one for Eqs. (11.5) and (11.6), we only address the latter.

Indeed, we can assume w.l.o.g. that the vector v in Theorem 11.4.7 is in $\mathbb{Q}^{d-1}$ and then apply the linear map $(x,t) \mapsto (x - vt, t)$ to the lattice. We will refer to the resulting lattice $\hat{\mathbb{Z}}^d$ as the *tilted lattice* and define its *period* $\hat{R} := \min\{t \in \mathbb{Z} : tv \in \mathbb{Z}^{d-1}, t \geqslant R\}$ and *base* $\hat{B} = (\mathbb{R}^{d-1} \times [0, \hat{R})) \cap \hat{\mathbb{Z}}^d$. For $A \subset S$ we define the tilted process, hitting time, hit and coupled regions

$$
\begin{aligned}
\hat{\xi}_t^A &:= \left\{(x,s) \in \hat{B} : (x + v(t+s), 0) \in \xi_{t+s}^A\right\},\\
\hat{K}_t^A &:= \left\{(x,s) \in \hat{B} : \hat{\xi}_t^A(x,s) = \hat{\xi}_t^S(x,s)\right\},\\
\hat{t}^A(x,s) &:= \min\left\{t \in s + \hat{R}\mathbb{Z} : t \geqslant 0, (x + vt, 0) \in \xi_t^A\right\} \quad (x,s) \in \hat{B},\\
\hat{H}_t^A &:= \left\{(x,s) \in \hat{B} : \hat{t}^A(x,s) \leqslant t\right\}
\end{aligned}
$$

The proof of Section 11.4.3 applies to GOSP in the tilted setting, yielding Eqs. (11.5) and (11.6) in $\hat{\mathbb{Z}}^d$ for some convex compact limit shape $\hat{U} \subset \mathbb{R}^{d-1}$ containing o in its interior. We then need to transfer the result back to the original lattice with $U = \hat{U} + v$. By the definition of $\hat{\xi}_t$ and $\hat{K}_t$, Eq. (11.6) and the inclusion in the coupled region in Eq. (11.5) are immediate. It remains to show that for every $\varepsilon > 0$, $\bar{\mathbb{P}}_p$-a.s. for every t large enough $H_t \supset (((1-\varepsilon)tU) \times [0,R)) \cap \mathbb{Z}^d$. Our strategy, somewhat similar to [133], is as follows. Fixing $x \in \mathbb{Z}^{d-1}$ such that $(x,0)$ should belong to H_t, we trace the line of slope v from (x,t) and determine when it intersects the boundary of the cone $\bigcup_{t' \geqslant 0}(t'U) \times \{t'\}$ (see Fig. 11.3). Someone close to the intersection point should be infected around time t' by the result available in $\hat{\mathbb{Z}}^d$. But then at a time $t' - \varepsilon t$ some site close to the intersection has survived for time εt, so, applying Eq. (11.12) and Theorem 11.4.7 to this site we manage to reach (x,t) as desired. With this in mind, let us spell out the details.

Equations (11.5) and (11.6) of Theorem 11.1.1 on $\hat{\mathbb{Z}}^d$ imply that for every $\varepsilon > 0$ $\bar{\mathbb{P}}_p$-a.s. there exists a constant C such that for every $(x,s) \in \hat{B}$ with $\|x\| \geqslant C$

$$
\left(x, \hat{t}(x,s)\right) \in \bigcup_{t>0} \left(t\partial\hat{U}\right) \times \{t\}, \tag{11.26}
$$

setting $\partial\hat{U} := ((1+\varepsilon)\hat{U}) \backslash ((1-\varepsilon)\hat{U})$. Observe that this event implies that in the original lattice there is at least one vertex infected by the origin in the

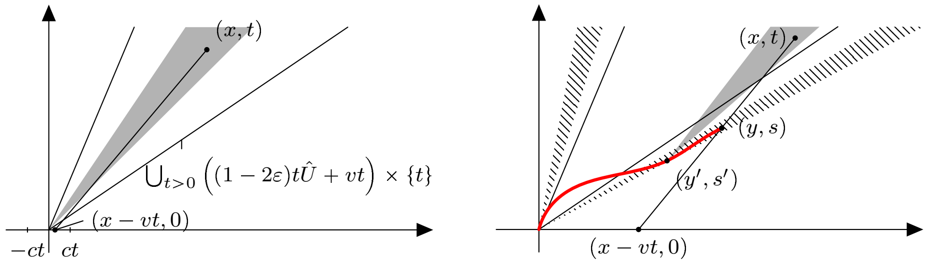

(a) Case $\|x-vt\| \leqslant ct$. (b) Case $\|x-vt\| > ct$.

Figure 11.3 – The original lattice for $d = 2$. Shaded areas represent the cone $\bigcup_{t>0}[t(v-c), t(v+c)] \times \{t\}$ rooted at o and (y', s') respectively. The hatched region is Δ.

intersection of $\Delta := \bigcup_{t>0}(t\partial\hat{U}_1 + vt) \times \{t\}$ and the ray $(x + vt, t)_{t\geqslant s}$ for all $(x, s) \in \hat{B}$ such that $\|x\| \geqslant C$ (see Fig. 11.3b).

Fix c and ε so that Theorems 11.4.3 and 11.4.7 hold for the original lattice. We now argue that for any $t > C/c$ and for any $x \in ((1-2\varepsilon)t\hat{U} + vt) \cap \mathbb{Z}^{d-1}$, we have $(x, 0) \in H_t$ except with probability exponentially small in t. If $\|x - vt\| \leqslant ct$ (see Fig. 11.3a), then Theorem 11.4.7 directly gives the desired result.

Assuming $\|x - vt\| > ct$, let $\delta > 0$ be small enough depending on $X, c, \hat{U}, v, \varepsilon$, but not t, x, C. Equation (11.26) implies that $\bar{\mathbb{P}}_p$-a.s. there is a vertex (y, s) in the intersection of Δ and the segment from $(x - vt, 0)$ to (x, t), such that $(y, 0) \in \xi^o_s$. As $s > \delta t$, we can take a site along an associated infection path at time closest to $s - \delta t$, and denote it by (y', s') (see Fig. 11.3b).

We then have that $(y', 0) \in \xi^o_{s'}$ and (y', s') survives for time at least $\delta t/2$, so we use Eq. (11.12) to conclude that (y', s') survives with probability exponentially high in t. Once we have survival, we can use Theorem 11.4.7 to show that x is in the hit region of (y', s') at time t with high probability, thus it is also in the hit region of the origin. Indeed,

$$\left\|x - y' - v(t - s')\right\| = \left\|y - y' - v(s - s')\right\| \leqslant t\sqrt{\delta} < c(t - s'),$$

since (x, t) is at distance at least κt from Δ (and thus from (y, s)) for some $\kappa > 0$ depending only on $\hat{U}, v, \varepsilon$.

This completes the proof of Eqs. (11.5) and (11.6) of Theorem 11.1.1 as stated for the original lattice.

11.A.3 Density large deviations—proof of Eq. (11.23)

Recall the notation of Theorem 11.4.10 and Section 11.4.4. We will use a similar argument to [138] but based on a completely different renormalisation.

Let us fix $p > p_c$ and $a < \tilde{\theta}(p)$, let C be large enough depending on p, let L be large enough depending on p, a, C and define $w = (L, \ldots, L) \in \mathbb{Z}^{d-1}$ and $s = CL + L/C$. Recalling Eq. (11.9), let $B = B(w, CL, v)$ with v as in Theorem 11.4.7. We say that B is *good* if the following events all occur.

(1) For each site $(x,t) \in B \cap S = B(w, R, v)$ we have either $\tau^{(x,t)} < L/C$ or $\tau^{(x,t)} \geqslant s$, where $\tau^{(x,t)}$ is defined as $\tau^{\{(x,R-1)\}}$ for the configuration ω translated by $-(t-R+1)e_d$.

(2) For each site $(x,t) \in B \cap S$ such that $\tau^{(x,t)} \geqslant L/C$ we have $K_s^{(x,t)} \supset B(3w, R, v) + sv$ and $K_{CL}^{(x,t)} \supset B(3w, R, v) + CLv$ with $K_u^{(x,t)}$ defined as $K_{u-t+R-1}^{\{(x,R-1)\}}$ for the configuration ω translated by $-(t-R+1)e_d$.

(3)
$$\xi_s^S \cap (B(w/C, R, v) + s(v,0) + Le_{d-1}) \neq \varnothing,$$
$$\xi_s^S \cap (B(w/C, R, v) + s(v,0) - Le_{d-1}) \neq \varnothing.$$

In words, each site which does not die quickly survives well beyond the top of B and infects the same set of sites at the top of $B \pm Le_{d-1}$, at least one of which does not die quickly. Indeed, the neighbourhood X being finite, the only way to reach $B(w/C, R, v) + s(v,1) \pm Le_{d-1}$ is to go through $B(w, R, v) + CL(v,1) \pm Le_{d-1}$. Therefore, considering a renormalised two-dimensional lattice with sites corresponding to disjoint translates of B, the resulting 2dOP is C^2-dependent, as B being good only depends on the configuration in $B(C^2 w, 2CL, v)$.

We next show that the parameter of the 2dOP is close to 1 when L is large enough, so that by [259] it stochastically dominates an independent 2dOP with parameter close to 1. Indeed, Event (1) fails with exponentially small (in L) probability by Eq. (11.12); Event (2) fails with exponentially small probability by Theorem 11.4.7 and Eq. (11.12); Event (3) fails with stretched exponentially small probability by Eq. (11.13) applied to the dual process.

It is easily checked that if a renormalised site B percolates in 2dOP, then each site in $B \cap S$ either dies in time at most L/C or also percolates. Recalling Proposition 11.A.1, the rest of the proof is essentially as in [138]. Taking n much larger than L, we may cut a box B_n into $(n/(2L))^{d-2}$ strips each giving rise to a different renormalised 2dOP. It is then standard to show that the total proportion of percolating renormalised sites is not close to 1 with probability at most $\exp(-\varepsilon n^{d-1})$ for some $\varepsilon > 0$ depending on L but not on n. Moreover, by standard large deviations for independent random variables, the proportion of sites, which survive at least L/C steps in B_n is smaller than $\theta(p)$ with probability at most $\exp(-\varepsilon' n^{d-1})$ for some $\varepsilon' > 0$ depending on L but not on n. We may then conclude by discarding the renormalised sites which do not percolate. Finally, performing the same reasoning for the dual process rather than the primal one, we obtain the desired conclusion (with $\tilde{\theta}(p)$ instead of $\theta(p)$).

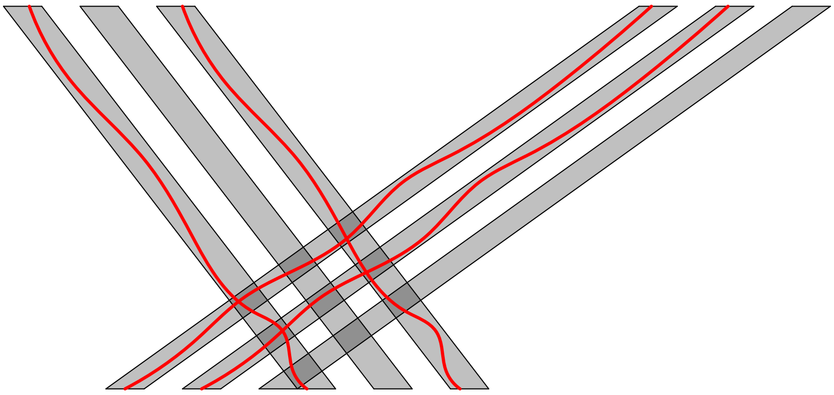

Figure 11.4 – The shaded boxes are likely to contain paths crossing them. In order to transition from one path to another, we use additional infections as illustrated in Fig. 11.5.

11.A.4 Enhanced two-dimensional renormalisation—proof of Theorem 11.5.4

Recall the notation of Theorem 11.5.4 and Section 11.5.

The idea of [136] is to introduce several translates of the original long boxes from Eq. (11.25) and hope that their paths have positive probability of intersecting and do so independently (see Fig. 11.4). More precisely, increasing p by a small amount ϵ, they require that the additional infections suffice to transition with positive probability from one path to the other at the place where they cross. As we shall see, although it is not possible to do this, as intended, in one step at the crossing point, it is possible to find a place where to do it in several steps.

More precisely, using Eq. (11.25), fix $\varepsilon > 0$ and $\delta > 0$ small and L large enough so that

$$\mathbb{P}_p\left(S \xrightarrow{B} S + (0, L)\right) > 1 - \delta$$

with $B = B(\varepsilon L, L, \alpha) - 2\varepsilon L e_1$ and similarly for $B' = B(\varepsilon L, L, \beta) + 2\varepsilon L e_1$ (see Fig. 11.5). Fix two open paths $\gamma = (a_0, a_1, \dots, a_n)$ and $\gamma' = (a'_0, a'_1, \dots, a'_m)$ crossing B and B' respectively. Fix v and large n and t as in Eq. (11.8) independent of all other constants. Let η be a set of additional infections, with each site at distance at most $O(t)$ from $B \cap B'$ infected independently with probability $\epsilon > 0$. Then we claim that

$$\mathbb{P}_{p+\epsilon}\left(a_0 \xrightarrow{\gamma\cup\gamma'\cup\eta} a'_m\right) > \epsilon^{O(t^2)}.$$

To see this, simply consider the region

$$\hat{\gamma} = \bigcup_{a\in\gamma}(a + B_n + (v, t)),$$

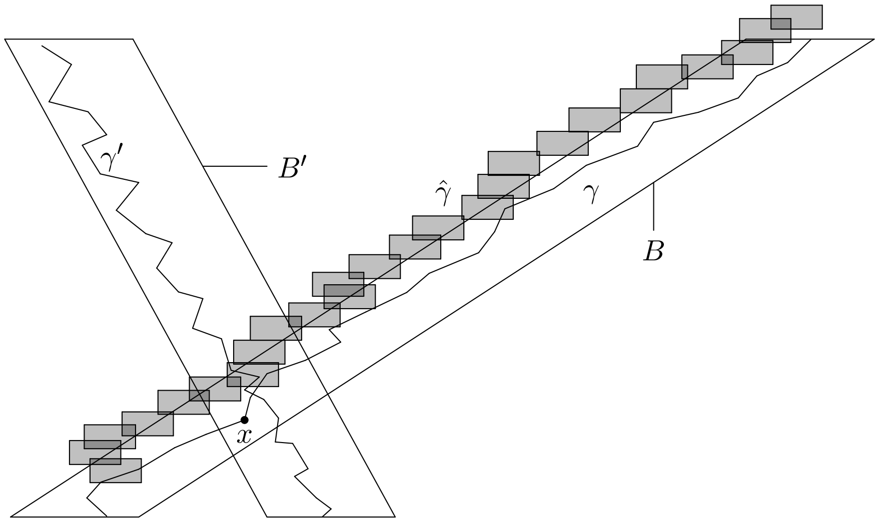

Figure 11.5 – The paths γ and γ'. The shifted and thickened version $\hat{\gamma}$ is the union of the shaded boxes. Due to its thickness it necessarily intersects γ' and does so close to the intersection of the boxes B and B'. The only x yielding the intersection in the example is indicated by a dot.

which is a shifted and thickened version of γ (see Fig. 11.5). It is clear that $\hat{\gamma} \cap \gamma' \neq \varnothing$, so it suffices for one $a \in \gamma$ such that $(a + B_n + (v,t)) \cap \gamma' \neq \varnothing$ to infect all possible sites for a time interval t in η. Since there are $O(t^2)$ of them, we obtain the desired result.

Hence, with positive probability we can go from a_0 to a'_m and, similarly, from a'_0 to a_n, which is just as good as having $\gamma \cap \gamma' \neq \varnothing$ for our purposes (except that the latter cannot be achieved by sprinkling). With this at hand, the approach of [136] works without the annoying hypothesis (H3) to renormalise GOSP to 2dOP with parameter close to 1. Consequently, $\alpha(p) > \beta(p)$ does imply $\theta(p) > 0$, but also, since the probability of a renormalised site being open is continuous in p and arbitrarily close to 1, Theorem 11.5.4 follows (see [131] for more details).

Chapter 12

Subcritical bootstrap percolation

This chapter is based on [209]. Recall Section 1.5.4.

Remark 12.0.1. We acknowledge that since the completion of the original [209] of the present chapter it has come to our attention that Schonmann [315, Theorem 5.1] has proved by independent means a result roughly equivalent to part of Theorem 12.3.5 (see also Remark 12.4.8).

In this chapter we write $\mathbb{P}_q$ for the product Bernoulli measure μ and $\mathbb{E}_q$ for its expectation. We say that bootstrap percolation (*BP*) occurs if the the closure $[A]$ of the initial set A of infections with law $\mathbb{P}_q$ is the entire lattice $\mathbb{Z}^2$. We further use the term *percolation* to refer to a random subset of $\mathbb{Z}^2$ with law $\mathbb{P}_p$ for any p, not necessarily referring to a BP process.

12.1 Models

Oriented site percolation *OP* can be viewed as the BP model defined by $\mathcal{U} = \{\{(1,1),(-1,1)\}\}$. It is easy to check (and was noticed already by Schonmann [315]) that $x \notin [A]$ if and only if there exists an infinite oriented path (with North-East and North-West steps) starting at x of initially healthy sites. In particular, q_c for this model is equal to $1-p_\mathrm{c}^{\mathrm{OP}}$, where $p_\mathrm{c}^{\mathrm{OP}}$ is the usual critical probability of OP parametrised in terms of the density of healthy sites (this is one of the reasons for denoting our parameter q). Up to applying an invertible linear transformation to $\mathbb{Z}^2$, any family with one rule consisting of two non-collinear sites is equivalent to OP, so we will abusively also call them OP. Furthermore, one may consider *bidirectional OP* with $\mathcal{U}' = \{\{(1,1),(-1,1)\},\{(-1,-1),(1,-1)\}\}$, for which the surviving healthy sites are those initially belonging to a bi-infinite oriented path, so that the critical probability is again $1-p_\mathrm{c}^{\mathrm{OP}}$. OP is a very classical and well-understood model, for background on which we direct the reader to [131] in addition to Section 12.5.

Spiral model The *Spiral model* of Toninelli, Biroli and Fisher [347] is defined by $\mathcal{U} = \{U_1, U_2, U_3, U_4\}$, where

$$\begin{aligned} U_1 &= \{(1,-1),(1,0),(1,1),(0,1)\} & U_2 &= \{(1,-1),(1,0),(-1,-1),(0,-1)\} \\ U_3 &= -U_1, & U_4 &= -U_2. \end{aligned} \tag{12.1}$$

This model was introduced to witness the somewhat surprising fact that subcritical BP can have a discontinuous phase transition in the sense that $\theta(q_c) = \mathbb{P}_{q_c}(0 \notin [A]) > 0$. This was established rigorously by Toninelli and Biroli [344] based on a close relationship with OP, which we will discuss further in Section 12.6.

Directed triangular bootstrap percolation *DTBP* was introduced by Balister, Bollobás, Przykucki and Smith [28] as an example of a simple, but somewhat generic, subcritical model. Its main feature is its lack of symmetry and it should be viewed as a benchmarking example. It can be defined as 2-neighbour BP on a directed triangular lattice, but can also be embedded in $\mathbb{Z}^2$ by

$$\mathcal{U} = \{\{(1,0),(0,1)\},\{(1,0),(-1,-1)\},\{(0,1),(-1,-1)\}\}. \tag{12.2}$$

As for most subcritical models not much is known about it. As a quantitative illustration of their result, the authors of [28] established that for DTBP

$$10^{-101} \leqslant q_c \leqslant 1 - p_c^{\mathrm{OP}} \leqslant 0.3118,$$

invoking Gray, Weirman and Smythe [193] for the last inequality.

Ordinary site percolation Finally, *SP* is one of the most classical percolation models (see [196]), which will also be useful for us, although it is not a particular case of BP. Similarly to OP, it consists in declaring each site of $\mathbb{Z}^2$ open independently with probability p and looking for infinite paths of open sites with respect to the usual nearest neighbour graph structure of $\mathbb{Z}^2$ instead of the oriented one for OP. We denote p_c^{SP} the critical probability of appearance of such infinite paths.

12.2 Definitions and notation

In this section we gather most of the notation used throughout the chapter. We invite the reader familiar with percolation to skip ahead to Section 12.3 and go back to this section as needed. As some of the notions will be used relatively locally, let us indicate that the central notion of the present chapter is the one in Definition 12.2.1.

Critical probability Recall that $0 < q < 1$ is the density of infected sites and $\mathbb{P}_q$ is the associated Bernoulli product law of the random set $A \subset \mathbb{Z}^2$ and that $[\cdot]$ denotes the closure with respect to the BP process defined by a non-trivial update family $\mathcal{U}$, that we keep implicit when there is no risk of confusion. Also, $B_x = [-x, x]^2 \cap \mathbb{Z}^2$ for all $x \in [0, \infty)$. Define

$$\begin{aligned} \theta_n(q) &= \mathbb{P}_q(0 \notin [A \cap B_n]), \\ \theta(q) &= \lim_n \theta_n(q) = \mathbb{P}_q(0 \notin [A]). \end{aligned}$$

The critical probability is given by

$$q_{\mathrm{c}} = \inf \left\{ q \in [0,1], \mathbb{P}_q([A] = \mathbb{Z}^2) = 1 \right\} = \sup \left\{ q, \theta(q) > 0 \right\},$$

the first equality following from ergodicity and the second one resulting from invariance by translation as for SP (see e.g. [196]). We also introduce another critical probability

$$\tilde{q}_{\mathrm{c}} = \inf \left\{ q \in [0,1], \sum_n n\theta_n(q) < \infty \right\}, \tag{12.3}$$

which is actually the only relevant one for our proofs, only noting that $\tilde{q}_{\mathrm{c}} \geqslant q_{\mathrm{c}}$. Several other equivalent definitions will be proved in Theorem 12.3.5, so that $\tilde{q}_{\mathrm{c}}$ is in particular the critical probability of exponential decay of $\theta_n(q)$. We emphasise that working with $\tilde{q}_{\mathrm{c}}$ instead of q_{c} will only lead to stronger results in applications.

Directions and half-planes In order to define the central notion of this chapter, critical densities, we will need some conventions and notation concerning directions and half-planes, which will mostly follow previous authors. We identify the unit circle $S^1 \subset \mathbb{R}^2$ with the torus $\mathbb{R}/2\pi\mathbb{Z}$ via

$$(\cos\theta, \sin\theta) \longleftrightarrow \theta \mod 2\pi.$$

Despite the identification we shall preferentially use the letters u, v for directions and θ for angles. For $n \in \mathbb{N}$ directions $u_1, \dots u_n \in S^1$ we write $u_1 < \dots < u_n$ if one can find $\theta_1 < \dots < \theta_n < \theta_1 + 2\pi$ and θ in $\mathbb{R}$ such that for each i we have $u_i \longleftrightarrow (\theta_i - \theta) \mod 2\pi$.

Recall that $\langle \cdot, \cdot \rangle$ and S^1 are the canonical scalar product on $\mathbb{R}^2$ and its unit sphere (circle). Furthermore, for $u \in S^1$ and $a \in \mathbb{R}$ set

$$\mathbb{H}_u^a = \{v \in \mathbb{R}^2, \langle v, u \rangle < a\},$$

$\mathbb{H}_u = \mathbb{H}_u^0$ and $\mathbb{H}_u^{a+} = \{v \in \mathbb{R}^2, \langle v, u \rangle \leqslant a\}$. We denote by $V_{u,v} = \mathbb{H}_u \cap \mathbb{H}_v$ the cone defined by the directions $u, v \in S^1$. We also recall the standard notation $a \vee b = \max(a, b)$ and $a \wedge b = \min(a, b)$.

Critical densities We are now ready to introduce the new notion of 'critical densities' adapted to subcritical BP (for critical and supercritical ones they will turn out to be identically 0). Let us note that this is not an extension, but rather a complement, of the 'difficulties' of [70], which are trivial for subcritical models.

Before we frighten the reader with the definition, let us say that the critical density in a direction u is morally the critical probability of the model with infected boundary condition in $\mathbb{H}_u$. The definition we give differs from this one in two ways—it concerns the critical probability for certain decay of $\theta_n(q)$ and it is defined in a region whose shape approaches a half-plane. Nevertheless, this distinction will only be of major importance for Section 12.4.2. That is because in applications we will always rely on simple OP-like models, in which we know that there is exponential decay above criticality and that the critical density is continuous in the shape of the region, so that the two notions coincide. Finally, we actually conjecture that they are *always* equal. With this in mind, let us state the definition we shall use.

Definition 12.2.1. For $u \in S^1$ and $\theta \in [-\pi, \pi]$ define

$$d_u^\theta = \inf \left\{ q \in [0,1], \sum_n n\mathbb{P}_q \left(0 \notin [((A \cup V_{u,u+\theta}) \cap B_n)]\right) < \infty \right\}.$$

Taking the (monotone) limit of this quantity, we set

$$d_u^\pm = \lim_{\theta \to 0\pm} d_u^\theta$$

and we call d_u^- and d_u^+ the *left and right critical densities* of u respectively. The *critical density* of u is then given by $d_u = d_u^+ \vee d_u^-$. We call $u \mapsto d_u$ the *critical density function* of the model (of $\mathcal{U}$).

It is clear from the definition that this quantity is somewhat of the same complexity as q_c, so that it is not feasible to be able to compute the critical densities for all u even for the simplest of subcritical models—OP.

The next observation directly follows from Definition 12.2.1, but will be the base for our upper bounds on q_c.

Observation 12.2.2. Let $\mathcal{U}$ be an update family. Let $u \in S^1$ be a direction and $\mathcal{U}' \subset \mathcal{U}$ be a subfamily of rules. Then

$$d_u(\mathcal{U}) \leqslant d_u(\mathcal{U}').$$

One-arm events Generally in percolation theory, a one-arm event is an event corresponding to 'a point being connected to infinity' or its finite-size truncations. In BP there is one very natural infinite volume one-arm event—$\{0 \notin [A]\}$, which corresponds to the presence of an infinite cluster (set) of healthy sites ensuring the occurrence of the event. There are several natural ways to truncate this event. In particular, we have

$$\{0 \notin [A]\} = \bigcap_n \{\tau_0 \geqslant n\} = \bigcap_n \{0 \notin [A \cap B_n]\},$$

etc., where τ_0 is the infection time of the origin. We interpret this event as $0 \to \infty$ (0 'looks at' infinity) and its truncated version $\{0 \notin [A \cap B_n]\}$ as $0 \to \partial B_n$ (∂ stands for the boundary). In models involving some kind of directionality, like BP, one may need to distinguish between 'point-to-infinity' and 'infinity-to-point' and similarly for truncated versions. The second one, which we define next, turns out to be more tractable, albeit less natural.

For $n \in \mathbb{N}$ and $x \in B_n$ we denote the infection time of x in B_n with healthy boundary condition by

$$\tau_x^{B_n} = \inf\left\{t,\, x \in (A \cap B_n)_t^{B_n}\right\},$$

where the dynamics only affects the configuration in B_n. More formally, for any sets $X \subset \mathbb{Z}^2$ and $A_0 \subset \mathbb{Z}^2$, we inductively define

$$A_{t+1}^X = A_t^X \cup \left\{x \in X, \exists U \in \mathcal{U}, x + U \subset A_t^X\right\}.$$

Definition 12.2.3. Fix a large constant $C > 0$ depending on $\mathcal{U}$. Denote by $E_n \subset \{0,1\}^{B_n}$ the event that there exists an integer N and a sequence $(x_i)_0^N$ of sites in B_n such that

- x_N is at distance at most C from the boundary ∂B_n of B_n.
- $x_0 = 0$
- $x_{i-1} \in x_i + X$ for all $1 \leqslant i \leqslant N$, where $X = \bigcup_{U \in \mathcal{U}} U$
- $\tau_{x_i}^{B_n} \geqslant i$.

Also set $\tilde{\theta}_n(q) = \mathbb{P}_q(E_n)$ and $\tilde{\theta}(q) = \lim_n \tilde{\theta}_n(q)$.

Note that the healthy boundary condition does not influence this event too much. Indeed, it is clear that some x_i is close to $\partial B_{n/2}$, so the occurrence of E_n implies the existence of a site 'in the bulk' (far from the boundary) with large infection time. We will use this observation to obtain information on the distribution of the infection time τ_0 below $\tilde{q}_c$.

The events E_n, which we interpret as $\partial B_n \to 0$, have the notable advantage of being 'reflexive' in the sense that, when exploring a configuration to check if E_n holds, looking back at the explored region from its boundary, one sees the event itself occurring in a smaller domain, which is crucial for the argument of Duminil-Copin, Raoufi and Tassion [127] that we will use.

Also very importantly, this event is defined in terms of a path rather than a 'cluster', although it does require the existence of 'clusters' of healthy sites. Of course, the main disappointment is that although very closely related to (and only differing by at most polynomial factors from) the natural events $\{0 \notin [A \cap B_n]\}$ or $\{\tau_0 \geqslant n\}$, it does not allow us to prove that $\tilde{q}_c = q_c$, but only provides additional constraints on the phase $[q_c, \tilde{q}_c)$. The reason is that we may have $\bigcap_n E_n \neq \{0 \notin [A]\}$, meaning that in BP the '$0 \to \infty$' and '$\infty \to 0$' events are different.

Randomised algorithms and revealment We will need the natural notion of *algorithm determining* a random variable Y on $\Omega_0 = \{0,1\}^{B_n}$ endowed e.g. with the measure $\mathbb{P}_q$. Roughly speaking, this is an algorithm which reveals the state of one bit (the value of $\omega_0 \in \Omega_0$ on one site $x \in B_n$) at a time possibly depending on knowledge of the configuration already explored. It keeps exploring bits one at a time until the value of Y is witnessed by the explored sites (determined regardless of the state of the remaining unexplored sites).

More formally, an algorithm is a rooted strict binary tree T directed away from the root. Its internal nodes are labelled by sites of B_n indicating the state of which site is being revealed. For each such internal node labelled by x, the two out-edges are labelled by the two possible values of the corresponding bit, so that given $\omega_0 \in \Omega_0$, the algorithm with input ω_0 continues along the edge labelled by $\omega_0(x)$. The leaves of the tree are labelled by the possible values of Y (with repetition) indicating which value of Y is witnessed (guaranteed) by the states indicated by the edges from the root to the leaf. More precisely, let P_l denote the path from the root to a leaf l labelled by a possible value y of Y. Then the vertices of P_l all have distinct labels (each site is revealed at most once) and for any $\omega_0 \in \Omega_0$ such that for all internal nodes $v \in P_l$ we have $\omega_0(x_v) = \epsilon_v$ it holds that $Y(\omega_0) = y$, where x_v is the label of v and ϵ_v is the label of the out-edge of v belonging to P_l. Clearly, given an algorithm and an input $\omega_0 \in \Omega_0$, there exists a unique leaf l_{ω_0} such that for every internal node in $v \in P_{l_{\omega_0}}$ we have $\omega_0(x_v) = \epsilon_v$. This simply corresponds to what the algorithm actually does for the specific realisation of the random input—which sites it checks, in what order, what values it finds for their states and, finally, what value of the random variable Y it determines based on those states.

A *randomised algorithm* is an algorithm-valued random variable. As we will apply these algorithms to inputs which are random themselves, we need to define them on a common probability space $(\Omega, \mathbb{P})$, so that the random algorithm is independent from the random input. For a randomised algorithm define its *maximal revealment*

$$\delta = \max_{x \in B_n} \mathbb{P}(\exists v \in P_{l_{\omega_0}}, x_v = x),$$

i.e. the maximal probability that any fixed site is explored by the algorithm.

Noise sensitivity We next define noise sensitivity, although our proofs will mostly use black-box theorems based on Fourier analysis instead of the definition.

Definition 12.2.4. Let $G_n \subset \{0,1\}^{B_n}$ be a sequence of events. For every $\omega_0 \in \{0,1\}^{B_n}$ let $N_\varepsilon(\omega_0)$ be the configuration obtained when each bit of ω_0 is resampled independently with probability ε and unchanged otherwise. Resampled bits are taken to be independently infected with probability q as originally.

We say that the sequence G_n is *noise sensitive*, if for every $\varepsilon > 0$

$$\lim_{n\to\infty} \frac{Cov\left(\mathbb{1}_{\omega_0 \in G_n}, \mathbb{1}_{N_\varepsilon(\omega_0)\in G_n}\right)}{Var(\mathbb{1}_{G_n})} = 0.$$

Let us note that this definition following [43] is stronger than the original one from [47], which is trivial for events with probabilities tending to 0 and equivalent, if the probabilities are bounded away from 0.

12.3 Results

Our goal is to provide a toolbox for studying subcritical models in full generality. Although our results will apply also to supercritical and critical models, most of them are either empty or relatively easy for such families. Unless explicitly mentioned we do not consider trivial subcritical models.

Critical densities and upper bounds on q_c Let $\mathcal{C} = \{[u, u+\pi], u \in S^1\}$ be the set of closed semi-circles of S^1. The most central result of our chapter is the following directional decomposition of the critical probability.

Theorem 12.3.1. *Let $\mathcal{U}$ be any update family. Then*

$$\tilde{q}_c = \sup_{u\in S^1} d_u = \inf_{C\in\mathcal{C}} \sup_{u\in C} d_u. \tag{12.4}$$

If $\mathcal{U}$ is not subcritical, then $\tilde{q}_c = 0$.

Combining Theorem 12.3.1 with Observation 12.2.2, we obtain the following upper bound on q_c.

Corollary 12.3.2. *Let $\mathcal{U}$ be an update family. Then for any set of subfamilies $\mathcal{U}_i \subset \mathcal{U}$ we have*

$$q_c(\mathcal{U}) \leqslant \tilde{q}_c(\mathcal{U}) \leqslant \inf_{C\in\mathcal{C}} \sup_{u\in C} \min_i d_u(\mathcal{U}_i).$$

Critical densities of OP In order to make use of Corollary 12.3.2 and obtain a concrete non-trivial upper bound in relative generality, we express the critical densities of OP in terms of a classical quantity called 'edge speed'. This is done in Section 12.5 by combining many standard facts about OP recalled there together with the definition of the 'edge speed'.

Application to DTBP Though simple, the bound in Corollary 12.3.2 is very versatile and can lead to non-trivial results for the right choice of subfamilies we have information for. Of course, in some cases it will reduce to the trivial bound $q_c(\mathcal{U}) \leqslant \min_{U\in\mathcal{U}} q_c(\{U\})$ (since it is sometimes sharp already), which has not been brought up explicitly in the literature, but was mentioned for DTBP in [28], taking only $\mathcal{U}_1 = \{U\}$ for some rule $U \in \mathcal{U}$ (they are all isomorphic). There it was observed that $q_c \leqslant 1 - p_c^{\mathrm{OP}} < 0.312$, the second inequality being due to Gray, Weirman and Smythe [193].

As an exemplary application of our result, we improve this bound on DTBP, answering Question 17 of [28] (of course, the question may now be reiterated). We prove the following by combining Corollary 12.3.2, the expression of critical densities of OP and a variant of the argument from [193].

Theorem 12.3.3. *For DTBP*

$$q_c \leqslant \tilde{q}_c \leqslant d^{OP}_{\arctan(-1/3)} < 0.2452,$$

where d^{OP} is the critical density of OP.

Application to Spiral Another application concerns the Spiral model. For that model Toninelli and Biroli [344] proved that $q_c = 1 - p_c^{\mathrm{OP}}$, there is exponential decay for $q > q_c$ and its transition is discontinuous, as well as providing bounds on the exponentially diverging correlation length. It turns out that our method exactly recovers the first two assertions, giving a new proof of the following.

Theorem 12.3.4 (Theorem 3.3. of [344])**.** *For the Spiral model $q_c = \tilde{q}_c = 1 - p_c^{\mathrm{OP}}$.*

This is a consequence of Corollary 12.3.2 together with an adaptation of a straightforward but fundamental lemma from [344], which inputs a crucial feature of the model identified by Jeng and Schwarz [235].

Exponential decay In the proof of Theorem 12.3.1 we actually prove that $\theta_n(q)$ decays exponentially fast in n for $q > \tilde{q}_c$. We provide a second proof of this fact, which also gives additional information on the phase $q < \tilde{q}_c$.

Theorem 12.3.5. *Recalling Definition 12.2.3, for any update family the following holds.*

- *If $q > \tilde{q}_c$, then there exists $c(q) > 0$ such that*

$$\max\left(\theta_n(q), \tilde{\theta}_n(q)\right) \leqslant \exp(-c(q) \cdot n).$$

- *There exists $c > 0$ such that for $q < \tilde{q}_c$*

$$\tilde{\theta}(q) \geqslant c \cdot (\tilde{q}_c - q) > 0.$$

- *If $q < \tilde{q}_c$, then there exists $c(q) > 0$ such that*

$$\mathbb{P}_q(\tau_0 > n) \geqslant c(q)/n$$

and in particular $\mathbb{E}_q[\tau_0] = \infty$.

Although we expect that $q_c = \tilde{q}_c$, this implies that if $q_c \neq \tilde{q}_c$, then the expected infection time is infinite at q_c (Question 11 of [28]).

The proof relies heavily on the new simple but powerful method of Duminil-Copin, Raoufi and Tassion [127] based on randomised algorithms. With some additional work on their only model-dependent Lemma 3.2, somewhat surprisingly the technique applies to BP, which is a rather unconventional setting for such arguments from SP.

Finally, we answer Question 12 of [28] on exponential decay for $q < q_c$ in the negative and provide satisfactory information concerning Question 14 of the same paper on the relationship between BP and SP.

Noise sensitivity Exploiting the algorithm we devise in order to prove Theorem 12.3.5, we obtain the following relatively complete information about noise sensitivity.

Theorem 12.3.6. *Recalling Definition 12.2.3, for any update family and any $q \in (0, 1)$ the following hold.*

- *$\tilde{\theta}(q) = 0$ if and only if the events E_n are noise sensitive and if and only if there is an algorithm with vanishing revealment determining their occurrence.*
- *If $\theta(q) > 0$, then the events $\{0 \notin [A \cap B_n]\}$ are not noise sensitive.*
- *If $\theta(q) = \tilde{\theta}(q) = 0$, then the events $\{0 \notin [A \cap B_n]\}$ are noise sensitive and there is an algorithm with vanishing revealment determining their occurrence.*

The proof relies on fundamental results of Benjamini, Kalai and Schramm [47] and Schramm and Steif [318].

In particular, this proves that Spiral is not noise sensitive at criticality, while OP is, so that the conditions on continuity of the transition are indeed relevant for noise sensitivity. Let us also mention that proving that the missing case—$\tilde{\theta}(q) > 0 = \theta(q)$—never occurs is only slightly stronger than proving Conjecture 12.8.1 stating that $q_c = \tilde{q}_c$. If it indeed does not occur, then Theorem 12.3.6 provides the final answer to Question 13 of [28] as far as one-arm events are concerned. Furthermore, Theorem 12.3.6 suggests some limitations for the intuition given by Bartha and Pete [43] (see Question 1.3 therein). Namely, Theorem 12.3.6 indicates that noise sensitivity non-trivially depends on the continuity of the transition, while [43] suggests that it should only depend on whether the model is subcritical or not, though for a more restrained class of models. Therefore, if a variant of Question 1.3 of [43] is to hold in general, additional ramifications should be needed.

Spectral gap and mean infection time of KCM Another application of our exponential decay results concerns KCM. We extend to full generality the scope of the main result of Cancrini, Martinelli, Roberto and Toninelli [88] using their method together with exponential decay.

Theorem 12.3.7. *Consider any KCM. If $q < \tilde{q}_c$, then the spectral gap of its generator is 0 and the mean infection time of the origin in the stationary process (with initial law $\mathbb{P}_q$) is infinite. If $q > \tilde{q}_c$, then the spectral gap is strictly positive and the mean infection time of the origin in the stationary process is finite.*

In other words, $\tilde{q}_c$ is the phase transition of the spectral gap of the associated KCM, so that it can be directly read off the associated BP as is the case of the non-ergodicity transition occurring at q_c [88].

We should note that the statement in the case of supercritical and critical models (for which $\tilde{q}_c = 0$ by Theorem 12.3.1) is also a trivial consequence of the quantitative result of [269]. We are particularly indebted to Cristina Toninelli for discussions around this theorem and its proof.

12.4 Critical densities

In this section, after some short preparatory work of establishing basic properties of critical densities, we characterise $\tilde{q}_c$ in terms of them, which can be viewed as the most central result of the chapter.

12.4.1 Preliminaries

We start with a few observations which follow trivially from Definition 12.2.1, but are essential nonetheless.

Observation 12.4.1. For all $u, \theta \in S^1$ one has

$$d_u^\theta \leqslant \tilde{q}_c$$

and therefore the same holds for d_u and $d_u^{\pm}$. Moreover, $\theta \mapsto d_u^{\theta}$ is non-decreasing for $\theta \in [0, \pi]$ and non-increasing for $\theta \in [-\pi, 0]$ and $d_u^{\pm\pi} = \tilde{q}_c$.

Observation 12.4.2. For all $u, \theta \in S^1$ one has

$$d_u^{\theta} = d_{u+\theta}^{-\theta}.$$

The following fundamental lemma is based on a classical topological trick.

Lemma 12.4.3. *Let $\varepsilon > 0$ and $I \neq S^1$ be a closed interval of S^1, which we identify with an interval $[u, v]$ of $\mathbb{R}$. Then there exists $n \in \mathbb{N}$ and a finite sequence $u = u_0 < u_1 < \ldots < u_n = v$ of directions in I such that*

$$\forall i \in [1, n],\ 0 \leqslant d_{u_{i-1}}^{u_i - u_{i-1}} - \left(d_{u_{i-1}}^{+} \vee d_{u_i}^{-}\right) < \varepsilon. \tag{12.5}$$

Proof. Recall that by Observation 12.4.2 for $u', v' \in S^1$ with $0 < v' - u' < \pi$ we have $d_{u'}^{v'-u'} = d_{v'}^{-(v'-u')}$. Then by Observation 12.4.1 one always has $d_{u'}^{v'-u'} \geqslant d_{u'}^{+} \vee d_{v'}^{-}$, so we need only establish the second inequality.

Set

$$I_0 = \big\{v' \in [u, v],\ \exists n \exists (u_i) \in (S^1)^{n+1},\\ u = u_0 < \ldots < u_n = v', \text{ satisfying (12.5)}\big\},$$

and $v_0 = \sup I_0$, which we shall prove to be v. To do this we prove that I_0 is open to the right:

$$\forall v' \in I_0\, \exists \delta > 0,\ [v', v' + \delta] \cap I \subset I_0$$

and closed to the right:

$$\exists v' \in I, (v_i) \in I_0^{\mathbb{N}},\ v_i \nearrow v' \Rightarrow v' \in I_0,$$

which suffices as I is an interval and $u \in I_0$.

For the first part, fix $v' \in I_0 \backslash \{v\}$, n and $(u_i)_0^n$, $u_n = v'$ as provided by the definition of I_0. By Observation 12.4.1 there exists $(v - v') \wedge \pi > \delta > 0$ small enough so that $d_{v'}^{\delta} - d_{v'}^{+} < \varepsilon$, which proves that $[v', v' + \delta] \subset I_0$.

The proof of I_0 being closed goes along the same lines looking to the left instead of to the right. More precisely, let v_i form an increasing sequence of elements of I_0 converging to $v' \in I$. By definition for i sufficiently large $v' - v_i < \delta$, where $0 < \delta < (v' - u) \wedge \pi$ is such that $d_{v'}^{-\delta} - d_{v'}^{-} < \varepsilon$. Hence, taking a sequence given by the definition of $v_i \in I_0$ and appending v' to it, we obtain $v' \in I_0$, which concludes the proof. ☐

Remark 12.4.4. One can use the technique of quasi-stable directions [70] to deal more easily with intervals of unstable and isolated stable directions. We do not do this as our construction works for the more difficult stable intervals and trivially also applies to unstable ones.

Also notice that if one knew that $(u, \theta) \mapsto d_u^{\theta}$ is continuous, this would follow by uniform continuity on a compact set.

We shall in fact need the following variant which follows immediately.

Corollary 12.4.5. *With the notation of Lemma 12.4.3 there also exist two directions such that* $v < v' < u' < u$ *and*

$$d_u^{u'-u} - d_u^- < \varepsilon,$$
$$d_v^{v'-v} - d_v^+ < \varepsilon.$$

Proof. Given a sequence as in Lemma 12.4.3 we apply one step of the reasoning to the right of v, obtaining v' sufficiently close to v and one step to the left of u. We simply observe that the inequalities we obtained in the proof of the Lemma were in fact the stronger ones in the statement of the corollary. □

12.4.2 Critical density characterisation of $\tilde{q}_c$—proof of Theorem 12.3.1

In order to prove Theorem 12.3.1 we will first need to show that above the maximal critical density in a semi-circle a certain well-chosen big droplet of infection grows indefinitely in that direction with high probability. We thus start by defining our droplets (see Figure 12.1).

Definition 12.4.6. Let $n \geqslant 3$, $u = u_0 < \ldots < u_{n+1} = v$ be directions with $u_n = u_1 + \pi$ and $u_n < v < u < u_1$ and let L be in $\mathbb{R}_+$. We then define the *droplet* of size L by

$$D_L = \bigcap_{i=0}^{n+1} \mathbb{H}_{u_i}^L - x_L, \quad D_{L+} = \bigcap_{L'>L} D_{L'} = \left(\bigcap_{i=1}^{n} \mathbb{H}_{u_i}^{L+} - x_L \right) \cap V_{u,v}, \quad (12.6)$$

where $x_L \in \mathbb{R}^2$ is such that $\langle x_L, u \rangle = \langle x_L, v \rangle = L$, so that droplets are inscribed in $V_{u,v}$.

It is crucial for the reasoning to follow that all sides of this droplet are of length $\Theta(L)$ for large L when the directions are fixed.

The growth mechanism is, of course, quite different from the one encountered for critical and supercritical models (finding an infection somewhere on the side of a droplet and relying on quasi-stable directions to make sure that the sides expand to fill the corners as well). Our strategy is to infect sites one by one by inspecting an area of size $\Omega(L)$ to have sufficiently small probability that the site remains uninfected in that zone. We can then use the union bound to infect a new row on one side of the droplet. We use this procedure to make the droplet grow, making sure that each side grows linearly, so that we can finally sum the probabilities using the decay provided by the definition of critical densities.

The next lemma roughly tells us that once a set of directions is fixed as in Corollary 12.4.5, a large infected droplet is highly likely to grow to infect

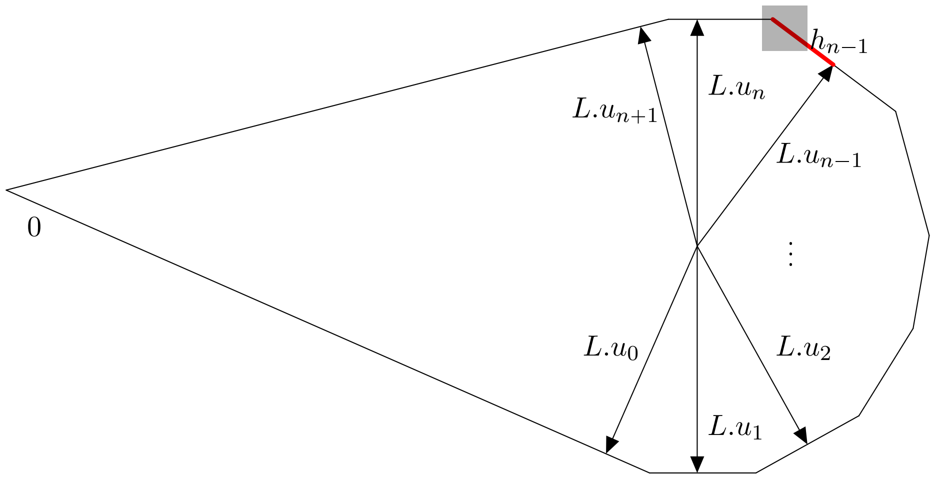

Figure 12.1 – The droplet D_L of size L for the directions $u_0, \ldots u_{n+1}$ defined in (12.6). The left 'half-side', h_{n-1} of l_{n-1} is thickened. The shaded box is $x + B_{L/C}$ for some $x \in h_{n-1}$.

the cone it is inscribed in if given a sufficiently high (compared to the critical densities) additional density of infections.

Lemma 12.4.7. *Let $n > 2$ and let $(u_i)_0^{n+1}$ be directions such that*

$$u = u_0 < u_1 < \ldots < u_n < u_{n+1} = v,$$

and $u_1 + \pi = u_n < u_{n+1} < u_0 < u_1$. Fix C large enough depending on the directions. Let $q > \max d_{u_{i-1}}^{u_i - u_{i-1}}$ for all $1 \leqslant i \leqslant n+1$ and let $\delta > 0$. Then for L large enough and for any $\Lambda \geqslant CL$

$$\mathbb{P}_q\left([D_L \cup (A \cap B_{C\Lambda})] \supset V_{u,v} \cap B_{\Lambda/C}\right) > 1 - \delta.$$

Proof. Let $(u_i)_{i=0}^n$, C, q and δ be as in the statement of the lemma.

Consider L such that $\mathbb{Z}^2 \cap (D_{L+} \backslash D_L) \neq \varnothing$ and let $L' = \sup\{l,\ D_l \cap \mathbb{Z}^2 = D_{L+} \cap \mathbb{Z}^2\}$. Consider the (possibly empty) new line of $D_{L'} \backslash D_L$ in direction u_i, $l_i = \mathbb{Z}^2 \cap D_{L'} \cap \left(\left(\mathbb{H}_{u_i}^{L'} \backslash \mathbb{H}_{u_i}^{L}\right) - x_L\right)$, for $1 \leqslant i \leqslant n$. Let $h_i = \{x \in l_i, \langle u_i + \pi/2, x + x_L\rangle \geqslant 0\}$ be the left *half-side* of l_i (looking from inside the droplet), see Figure 12.1. For each site $x \in h_i$ and $\Lambda \geqslant CL$ we have

$$\begin{aligned}\mathbb{P}_q\left(x \notin [D_L \cup (A \cap B_{C\Lambda})]\right) &\leqslant \mathbb{P}_q\left(x \notin \left[(A \cup D_L) \cap (x + B_{L/C})\right]\right)\\ &\leqslant \mathbb{P}_q\left(0 \notin \left[\left(A \cup V_{u_i,u_{i+1}}\right) \cap B_{L/C}\right]\right),\end{aligned}$$

since inside a box of size L/C around x the droplet locally looks like (at least) $V_{u_i,u_{i+1}}$, see Figure 12.1. Then the union bound over all sites in all half-sides gives

$$\mathbb{P}_q\left([D_L \cup (A \cap B_{C\Lambda})] \not\supset D_{L'}\right) \leqslant \sum_{i=1}^{n} |l_i| \Big(\mathbb{P}_q\left(0 \notin \left[\left(A \cup V_{u_i,u_{i+1}}\right) \cap B_{L/C}\right]\right) + \mathbb{P}_q\left(0 \notin \left[\left(A \cup V_{u_{i-1},u_i}\right) \cap B_{L/C}\right]\right) \Big).$$

We now iterate this bound. Let L_0 be large enough (depending on C, δ and $(u_i)_{i=0}^{n+1}$) and such that such that $\mathbb{Z}^2 \cap (D_{L_0+} \backslash D_{L_0}) \neq \varnothing$. Define $L_{j+1} = \sup\{l,\, D_l \cap \mathbb{Z}^2 = D_{L_j+} \cap \mathbb{Z}^2\}$ for all $j \geqslant 0$. Again by the union bound for any $L \geqslant L_0$ and $\Lambda \geqslant CL$ we have

$$\begin{aligned}\mathbb{P}_q&\left([D_L \cup (A \cap B_{C\Lambda})] \not\supset D_{\Lambda}\right) \\ &\leqslant \sum_{i=1}^{n}\sum_{j=0}^{\infty} |l_i^j| \Big(\mathbb{P}_q\left(0 \notin \left[\left(A \cup V_{u_i,u_{i+1}}\right) \cap B_{L_j/C}\right]\right) + \mathbb{P}_q\left(0 \notin \left[\left(A \cup V_{u_{i-1},u_i}\right) \cap B_{L_j/C}\right]\right) \Big),\end{aligned}$$

where $l_i^j = \mathbb{Z}^2 \cap D_{L_{j+1}} \cap \left(\left(\mathbb{H}_{u_i}^{L_{j+1}} \backslash \mathbb{H}_{u_i}^{L_j}\right) - x_{L_j}\right)$.

Let us upper bound the first term for $i = 1$ for concreteness. Let $j_k = \min\{j,\, L_j \geqslant Ck\}$. Then for any $k \geqslant \lfloor L_0/C \rfloor$

$$\sum_{j=j_k}^{j_{k+1}-1} |l_1^j| \mathbb{P}_q\left(0 \notin \left[(A \cup V_{u_1,u_2}) \cap B_{L_j/C}\right]\right) \leqslant \mathbb{P}_q\left(0 \notin [(A \cup V_{u_1,u_2}) \cap B_k]\right) \sum_{j=j_k}^{j_{k+1}-1} |l_1^j|.$$

Finally, the last sum is easily seen to be at most $C^3 k$ (it is essentially equal to the area covered by the u_i side while growing from D_{Ck} to $D_{C(k+1)}$), so in total we get

$$\begin{aligned}\mathbb{P}_q&\left([D_L \cup (A \cap B_{C\Lambda})] \not\supset D_{\Lambda}\right) \\ &\leqslant \sum_{k=\lfloor L_0/C \rfloor}^{\infty} C^3 k \sum_{i=0}^{n} \mathbb{P}_q\left(0 \notin \left[\left(A \cup V_{u_i,u_{i+1}}\right) \cap B_k\right]\right) \leqslant \delta\end{aligned}$$

by Definition 12.2.1 and the choice of q. This concludes the proof, since $D_\Lambda \supset V_{u,v} \cap B_{\Lambda/2}$ (by construction the u, v-sector of the Euclidean ball of radius Λ/C is contained in D_Λ). □

We are now ready to prove Theorem 12.3.1.

Proof of Theorem 12.3.1. By Observation 12.4.1 we have

$$\tilde{q}_c \geqslant \sup_{u\in S^1} d_u \geqslant \inf_{C\in\mathcal{C}} \sup_{u\in C} d_u,$$

so we are left with proving $\tilde{q}_c \leqslant \inf_{C\in\mathcal{C}} \sup_{u\in C} d_u$.

Fix $\varepsilon > 0$ sufficiently small and $C \in \mathcal{C}$ such that

$$\varepsilon + \inf_{C'\in\mathcal{C}} \sup_{u\in C'} d_u > \sup_{u\in C} d_u.$$

Also fix a set of directions as required by Lemma 12.4.7 with $C = [u_1, u_n]$ and satisfying

$$\begin{aligned} \forall i \in [2,n],\ d_{u_{i-1}}^{u_i - u_{i-1}} - (d^+_{u_{i-1}} \vee d^-_{u_i}) &< \varepsilon \\ d_{u_1}^{-(u_1-u_0)} - d^-_{u_1} &< \varepsilon \\ d_{u_n}^{u_{n+1}-u_n} - d^+_{u_n} &< \varepsilon, \end{aligned}$$

as provided by Corollary 12.4.5. Without loss of generality (after rotating the lattice) we assume $u_n = (0,1)$. Fix $\delta > 0$ sufficiently small depending on the directions (u_i) and ε. Let $q' = 2\varepsilon + \sup_{u'\in C} d_{u'}$, so that $q = q' - \varepsilon$ satisfies the condition $q > \max d_{u_{i-1}}^{u_i - u_{i-1}}$ of Lemma 12.4.7.

We sample (a part of) the infected sites as the union of two independent percolations—one with probability ε and another one with probability q. At this point one can easily obtain $q' \geqslant q_c$ using Lemma 12.4.7 to prove that a droplet of size L grows with high probability in the second percolation and find such a large droplet in the first one. However, in order to avoid using $q_c = \tilde{q}_c$, we give a slightly more involved but fairly standard renormalisation procedure to prove the desired inequality for $\tilde{q}_c$. Furthermore, we will be able to deduce that $\tilde{q}_c$ is also the critical probability of exponential decay.

Let L be large enough for the assertion of Lemma 12.4.7 to hold. Also fix N sufficiently large depending on L such that $\mathbb{P}_\varepsilon(\exists x \in B_N,\, A \cap B_N \supset D_L + x) \geqslant 1 - \delta$. Finally, let $c \in \mathbb{N}$ be large enough depending only on the directions (u_i) (and on the constant C in Lemma 12.4.7), but not on δ. Consider a renormalised lattice $\mathcal{L} = \mathbb{Z}^2$ and say $X \in \mathcal{L}$ is open if $N.X + B_N \subset [A \cap (N.X + B_{cN})]$. This process is clearly only $2c$-dependent[1] and we claim that each site is open with probability at least $1 - 2\delta$. Indeed, $N(X - (\lfloor\sqrt{c}\rfloor, 0)) + B_N$ contains a droplet of size L in the percolation process with parameter ε with probability at least $1 - \delta$ and this droplet grows to infect $NX + B_N$ with probability at least $1 - \delta$ in the percolation process with parameter q only using infections inside $NX + B_{cN}$ by Lemma 12.4.7.

Hence, by the Liggett–Schonmann–Stacey theorem [259] the renormalised process stochastically dominates an independent site percolation with parameter $1 - \delta'$ with δ' which can be made arbitrarily small by choosing

[1]Each site is independent from the states of sites at distance more than $2c$ from it.

δ sufficiently small. In particular, it is known (from the standard Peierls argument, see e.g. [196]) that the probability that there is no contour (self-avoiding closed path) of open sites around 0 decays exponentially. Yet, if such a contour exists in a renormalised box of size $a > c$, we know that $0 \in [A \cap B_{2aN}]$. Indeed, since the family is not trivial subcritical, the renormalised site $NX + B_N$ for X in the contour becomes infected using $A \cap (NX + B_{cN})$ and the union of these sets for all X in the contour is enough to infect the origin. To see this, simply use the fact that there exists an unstable direction and that the BP process inside the infected contour behaves as though everything outside the contour is infected. Thus, $\theta_m(q')$ decays exponentially in m, since N is a constant. Hence, $q' \geqslant \tilde{q}_c$, concluding the proof of (12.4).

Let us now consider a non-subcritical family and show that $\tilde{q}_c = 0$. Fix $q > 2\varepsilon$. It is not hard to see (e.g. by repeating the proof from [74]) that a sufficiently large droplet is very likely to grow using a density ε of infections to infect an entire cone of fixed opening depending only on ε and $\mathcal{U}$ (see Figure 7 of [74]). We can then repeat the renormalisation above using this input instead of Lemma 12.4.7 to obtain that there is exponential decay at q and thereby $\tilde{q}_c = 0$. □

Remark 12.4.8. Note that we also proved that $\tilde{q}_c$ is the critical probability of exponential decay: for each $q > \tilde{q}_c$

$$\liminf_n \frac{-\log \theta_n(q)}{n} > 0,$$

while this fails for $q < \tilde{q}_c$. Moreover, since the family is not trivial, the exponential decay of the absence of a renormalised contour of radius n implies also exponential decay of $\mathbb{P}_q(\tau_0 \geqslant n)$ for $q > \tilde{q}_c$.

Remark 12.4.9. In fact, using droplets contained between two parallel lines (see Figures 5 and 7 of [74]) instead of a cone with strictly positive opening one can obtain a slightly stronger characterisation of $\tilde{q}_c$ only involving one of the left or right critical densities at each endpoint of the semi-circle.

12.5 Critical densities of oriented percolation

In this section we determine the critical densities of the simplest subcritical BP model—OP. This is established in order to be used in conjunction with Theorem 12.3.1 in the next section to deduce information about other models. Interestingly, although determining critical densities corresponds to studying the phase transition of OP with an absorbing boundary condition (in a restricted region), this problem does not seem to have been thoroughly studied. The only case which we are aware of that has been considered [160] is the symmetric one—$u = \pi$, for which the result, as we shall see, is that the transition is the same as on the entire plane.

Let us recall a few classical results from OP theory all of which can be found up to minor modifications in Durrett's review [131] (see also [130,134, 193,258]). We will not redo most of the proofs, as they are discussed in more detail for GOSP in Chapter 11 and since they have appeared numerous times in the literature in slightly modified forms.

Recall that OP is defined by $\mathcal{U} = \{U\} = \{\{(-1,1),(1,1)\}\}$. For the sake of convenience, in this section we parametrise in terms of $p = 1 - q$—the density of healthy (open) sites, so that $\mathbb{P}_p$ still denotes the product Bernoulli measure such that each site is open with probability p. For the rest of this section we consider only the sublattice of $\mathbb{Z}^2$ generated by U without further mention. Denote by $x \to y$ for x and y in $\mathbb{Z}^2$ the event that there exist $x_0, \dots, x_N$ with $x_0 = x$, $x_N = y$, $x_i - x_{i-1} \in U$ and x_i open for $0 < i \leqslant N$, that we call an *OP path* (from x to y). Let

$$r_n = \sup\{x \in \mathbb{Z}, \exists y \leqslant 0, (y,0) \to (x,n)\}$$

be the *right edge* with the convention $\sup \varnothing = -\infty$.

Lemma 12.5.1. *There exists a function $\alpha : [0,1] \to [-\infty, 1]$ called* edge speed *with the following properties.*

(1) For any p we have $\mathbb{P}_p$-a.s.

$$r_n/n \to \alpha(p) = \inf_n \mathbb{E}_p[r_n/n] = \lim_n \mathbb{E}_p[r_n/n].$$

(2) α is strictly increasing on $\left[p_{\mathrm{c}}^{\mathrm{OP}}, 1\right]$.

(3) α and continuous on $\left[p_{\mathrm{c}}^{\mathrm{OP}}, 1\right]$ with $\alpha\left(p_{\mathrm{c}}^{\mathrm{OP}}\right) = 0$, $\alpha(1) = 1$ and $\alpha(p) = -\infty$ for $p < p_{\mathrm{c}}^{\mathrm{OP}}$.

The first equalities and the a.s. limit are proved as in [258], following [130, 131]. The other assertions are proved exactly as in [131]. We will use this definition of α in the remainder of the chapter. The contour argument used in [131] to prove the continuity of α (together with the Borel-Cantelli lemma) actually gives the following.

Lemma 12.5.2. *For all $p > p_{\mathrm{c}}^{\mathrm{OP}}$ and $a < \alpha(p)$ we have that with positive probability there exists an infinite OP path $((a_i, i))_{i\in\mathbb{N}}$ with $a_0 = 0$ and $\inf_n a_n/n \geqslant a$.*

The next Lemma can be proved exactly like Theorem 7 of [195] (see also [131]).

Lemma 12.5.3. *If $a > \alpha(p)$, then for some $\gamma > 0$*

$$\mathbb{P}_p(r_n \geqslant an) \leqslant e^{-\gamma n}.$$

The following bound on α will only be used in the next section.

Lemma 12.5.4. *For all $p \in [0,1]$ we have*

$$\alpha(p) \leqslant \frac{p^3+p-1}{p^3-2p^2+3p-1}.$$

Proof. The two-paragraph argument of Section 2 of [193] adapts immediately to give that $\alpha^{-1}(a)$ is larger than the root of the equation

$$(p^3-p^2+2p-1)/(p-p^2) = \frac{1+a}{1-a}.$$

Rephrasing this we obtain exactly the desired inequality. □

Let ψ be the composition of the tangent, the inverse of α and finally $1-\cdot$

$$\psi : [-\pi,-3\pi/4] \cup [-\pi/4,0] \xrightarrow{|\tan|} [0,1] \xrightarrow{\alpha^{-1}} \left[p_c^{\mathrm{OP}},1\right] \xrightarrow{1-\cdot} [0,q_c].$$

Putting the preceding facts together we obtain the critical densities of OP.

Theorem 12.5.5. *The critical density of $\mathcal{U} = \{U\} = \{\{(1,1),(-1,1)\}\}$ is given by*

$$d_u(\mathcal{U}) = \begin{cases} 0, & u \in [-3\pi/4,-\pi/4] \\ 1-p_c^{\mathrm{OP}} = q_c, & u \in [0,\pi] \\ \psi(u), & \text{otherwise.} \end{cases}$$

For bidirectional OP $\mathcal{U}' = \{U,-U\}$, where $-U = \{(-1,-1),(1,-1)\}$, the critical densities are $d_u(\mathcal{U}') = d_u(\mathcal{U}) \wedge d_{-u}(\mathcal{U})$. One also has $d_u^0 = d_u^{\pm} = d_u$ for all u in both cases.

Remark 12.5.6. If the OP rule is rather $\tilde{U} = \{(x,y),(z,t)\}$ with the two linearly independent vectors (sites), let $L \in GL_2(\mathbb{R})$ be such that $L \cdot \tilde{U} = U = \{(-1,1),(1,1)\}$ and $\det L > 0$. Then the critical densities are also transformed via $d_u^{\{\tilde{U}\}} = d_{u'}^{\{U\}}$, where u' is the direction of $(L(u-\pi/2))+\pi/2$.

Proof of Theorem 12.5.5. If $u \in (-3\pi/4,-\pi/4)$ we have nothing to prove, as the directions are unstable. By symmetry it suffices to treat $u \in [-\pi/4,\pi/2]$, so fix one such direction and let $\tilde{q} = q_c$ if $u \in (0,\pi/2]$ and $\psi(u)$ otherwise. Notice that $\alpha(1-\tilde{q}) = -\tan(u)$ in the latter case and 0 in the former one.

Let $q < \tilde{q}$. By Lemmas 12.5.1 and 12.5.2 we know that with positive probability there exists an infinite OP path of healthy sites starting at 0 not intersecting $\mathbb{H}_u$. This proves that $q \leqslant d_u^\theta$ for all θ, so $q \leqslant d_u^0 \leqslant d_u^{\pm} \leqslant d_u$ and the same inequalities hold for $\tilde{q}$.

Conversely, let $q > \tilde{q}$. Then by Lemmas 12.5.1 and 12.5.3

$$\mathbb{P}_q(0 \notin [(A \cap B_n) \cup V_{u-\theta,u+\theta}])$$

decays exponentially for $\theta > 0$ small enough, so that $d_u^0 \leqslant d_u^\pm \leqslant d_u \leqslant q$. Thus, with the inequalities from the previous case we obtain

$$d_u = d_u^\pm = d_u^0 = \tilde{q}.$$

Now consider bidirectional OP. It is clear that 0 remaining healthy for this process is equivalent to 0 remaining healthy for the family $\{U\}$ and for the family $\{-U\}$, both of which are simply OP. Moreover, these two events are independent conditionally on the state of 0 (as the oriented paths occur in the upper and lower half-planes respectively). Thus, the critical densities are indeed obtained as claimed. □

Remark 12.5.7. In order to be able to usefully apply Corollary 12.3.2 in full generality to any subcritical model, we require a generalisation of Theorem 12.5.5 to GOSP. Indeed, every non-trivial subcritical model contains rules corresponding to GOSP as explained in Section 12.1. The proof of Theorem 12.5.5 remains unchanged for GOSP, provided we have all the ingredients needed, Lemmas 12.5.1–12.5.3. In Chapter 11 we explained how those are established.

12.6 Applications of the upper bound

The most natural and easy way to use Corollary 12.3.2, which we call *basic bound*, is for subfamilies consisting of only one rule:

$$q_c(\mathcal{U}) \leqslant \tilde{q}_c(\mathcal{U}) \leqslant \inf_{C\in\mathcal{C}} \sup \min_{U\in\mathcal{U}} d_u(\{U\}), \tag{12.7}$$

since the r.h.s. terms correspond to OP treated in the previous paragraph or similarly behaved GOSP. In principle this approach includes the trivial one consisting of using $q_c(\mathcal{U}) \leqslant \min_{U\in\mathcal{U}} q_c^{\{U\}}$, but also allows better estimates.

We give two illustrative applications of the general bound of Corollary 12.3.2. The first one follows from the basic bound given by single rule subfamilies as outlined above, while the second one is more subtle.

12.6.1 The basic bound—the DTBP model

Our first example is DTBP. We improve the upper bound of [28] as asked in their Question 17 by proving Theorem 12.3.3.

Proof of Theorem 12.3.3. Our starting point is (12.7). Let U_i be the three rules in the update family $\mathcal{U}$ of DTBP defined in (12.2). We can then use Theorem 12.5.5 and Remark 12.5.6 to determine the r.h.s. We spare the reader the tedious details, but it is elementary to see (see Figure 12.2) that by symmetry there are three local maxima of $u \mapsto \min_i d_u(\{U_i\})$—the one

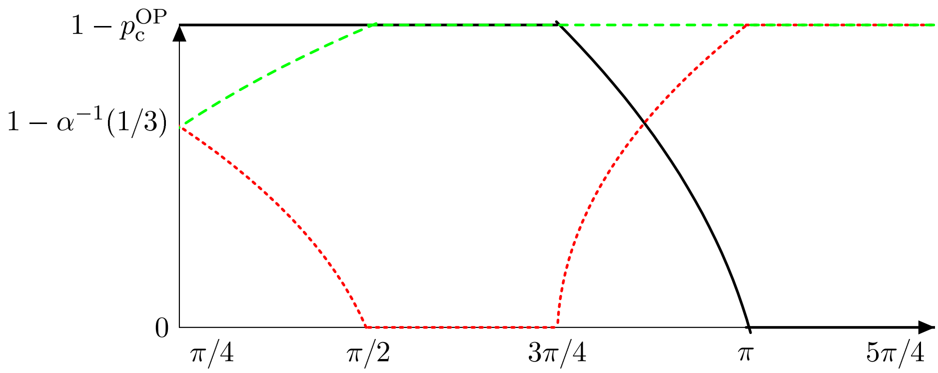

Figure 12.2 – A schematic representation of the critical densities of the three OP rules in DTBP. For symmetry reasons we only depict the domain $u \in [\pi/4, 5\pi/4]$.

at $\pi/4$ being the global maximum in $[-\pi/4, 3\pi/4]$. Hence, Theorem 12.5.5 and Remark 12.5.6 give

$$q_c(\mathcal{U}) \leqslant d^{OP}_{L(-\pi/4)+\pi/4} = d^{OP}_{\arctan(-1/3)} = 1 - \alpha^{-1}(1/3),$$

where $L(x,y) = (x, y-x)$ maps the DTBP rule $\{(-1,-1),(0,1)\}$ into $\{(-1,0),(0,1)\}$, which is OP rotated by $\pi/4$.

In fact, the other two maxima are also easily determined to be at $\pi - \arctan(1/2)$ and $\arctan(1/2) - \pi/2$. They turn out to give the same value as the one at $\pi/4$, but we did not need that for establishing the upper bound. Finally, Lemma 12.5.4 provides the desired bound $\alpha^{-1}(1/3) > 0.7548$. □

It should be noted that the numerical bound is not optimised, but merely given to testify that the gain is significant. For comparison, based on a refinement of the same method in [193] in conjunction with the trivial bound $q_c(\mathcal{U}) \leqslant 1 - p_c^{OP} = 1 - \alpha^{-1}(0)$ the authors of [28] obtain $q_c(\mathcal{U}) < 0.312$. Even if the exact value of p_c^{OP} were known, it follows from rigorous upper bounds that the trivial bound cannot go beyond 0.274 [29]. Numerical studies indicate that in fact $1 - p_c^{OP} \approx 0.2945$ [294]. Unfortunately, we have been unable to find appropriate numerical estimates for α for values far from q_c in the literature, so we cannot provide a corresponding result for our bound $1 - \alpha^{-1}(1/3)$. Finally, all these values are also to be compared with the numerical estimate $q_c(\mathcal{U}) \approx 0.118$ suggested in [28], which indicates that there is much room for further improvements.

12.6.2 Motivation of the second-level bound

Unfortunately, the basic bound (12.7) is not tight. Something more, it is possible to find two rules U_1 and U_2, such that $d(\{U_1, U_2\})$ is nowhere equal to $d(\{U_1\}) \wedge d(\{U_2\})$. Even worse, changing U_2 may lead to a change in $d(\{U_1, U_2\})$ while $d(\{U_2\})$ remains the same. We give the following instructive counterexample, along whose lines many can be constructed.

Proposition 12.6.1. *Let $\mathcal{U}_n = \{U_1, U_n\}$ with $U_n = \{(n,n), (-n,n)\}\}$ for $n \geqslant 1$. Then as $n \to \infty$*

$$q_c(\mathcal{U}_n) \leqslant 1 - \inf\left\{p,\, p_c^{\mathrm{OP}} \leqslant \theta^{OP}(p)\right\} + o(1),$$

where $\theta^{\mathrm{OP}}(p) = \mathbb{P}_{1-p}\left(0 \notin [A]_{\{U_1\}}\right)$ is the probability that 0 is never infected in OP.

Proof. Let $B'_n = (-n,n] \times (0,n)$ and denote by $\mathcal{L} = \{n.(m-k, m+k),\ m,k \in \mathbb{N}\}$ the sites concerned by the second rule. Note that for all $x \in \mathcal{L}$ the boxes $x + B'_n$ are disjoint and disjoint from $\mathcal{L}$.

Fix $\varepsilon > 0$ and $p = 1-q$ such that $\theta^{OP}(p) < p_c^{\mathrm{OP}} - \varepsilon$. Let n be large enough so that

$$\mathbb{P}_q\left(x \notin \left[A \cap \left(x + B'_n\right)\right]\right) \leqslant \frac{\theta^{OP}(p) + \varepsilon}{p}.$$

Such an n exists, because the process with initial infection in $x + B'_n$ is identical to the one under the family $\{U_1\}$, which is OP and for which we know that the probability converges to $\theta^{OP}(p)/p$.

Then we can associate to each site of $x \in \mathcal{L}$ an independent Bernoulli random variable with parameter $\theta^{OP}(p)+\varepsilon$—the indicator of the event $G_x = \{x \notin A;\, x \notin [A \cap (x + B'_n)]\}$. Furthermore, $\{x \notin [A]\} \subset G_x$ for all x. But then in order for 0 to remain uninfected at all times it is necessary to have an infinite path with steps in U_n starting at 0 of sites x such that G_x occurs and the probability of this event is $\theta^{OP}(\theta^{OP}(p) + \varepsilon) = 0$, since $\theta^{OP}(p) \leqslant p_c^{\mathrm{OP}} - \varepsilon$. □

This example shows where the main difficulty of the subcritical models resides once GOSP is well understood. The division into three universality classes is based on the unstable directions of a model, which can be directly obtained by superimposing the ones for each rule, which are very easy to determine [28, 74]. In the refined result based on 'difficulties' for critical models [70] Bollobás, Duminil-Copin, Morris and Smith only require information in the finitely many isolated stable directions—their difficulty. In their case, like here, there is no easy way of calculating the difficulty of an isolated stable direction without looking at the entire update family. However, in the simple case of critical models the difficulty happens to be a finite discrete quantity, which invites direct exhaustive computation (which for simple models is readily done by hand), and indeed [70] does not provide a recipe for determining difficulties (it turns out that determining them is NP-hard—see Chapter 9). This is essentially the same problem that we are

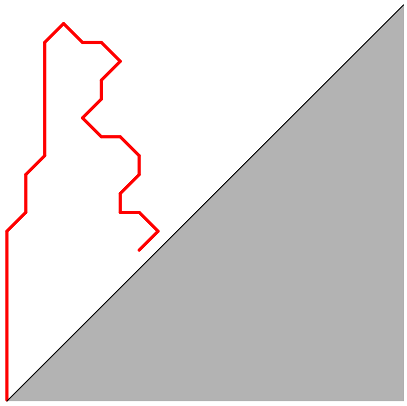

Figure 12.3 – An example of the healthy path used in the proof of Lemma 12.6.2. The shaded region is entirely infected.

facing here, but the critical densities of subcritical models being much richer, they are even harder to decompose and analyse.

On the bright side the bound from Corollary 12.3.2 need not be applied to subfamilies with a single rule. Hence, if we have information on the joint critical densities of, say, all pairs of rules in the family $\mathcal{U}$, then we can extract a (better) upper bound for $q_c(\mathcal{U})$. We next turn our attention to an example where this approach works brilliantly, while to apply the basic bound (and obtain worse results) we would need an understanding of GOSP.

12.6.3 Spiral model

Indeed, in the Spiral model the subfamilies with two rules happen to be simpler than the single-rule ones when restricted to appropriate half-planes. Recall the definition of its update family $\mathcal{U} = \{U_1, U_2, U_3, U_4\}$ from (12.1). We will use Corollary 12.3.2 to provide a new proof of one of the main results of [344]—Theorem 12.3.4.

The proof is nearly complete at this point, but we need one last ingredient, a variant of Lemma 4.11 of [344], which is actually more naturally expressed in the language of critical densities. This is where one uses the "no parallel crossing" property, which Jeng and Schwarz [235] identified as essential, as without it the pairs of rules do not simplify to OP.

Lemma 12.6.2 (Adaptation of Lemma 4.11 of [344])**.** *Let $u \in (\pi/2, 5\pi/4)$. Then*

$$d_u(\{U_1, U_2\}) = d_u(\mathcal{U}'),$$

where $\mathcal{U}' = \{\{(0,1),(1,1)\},\{(0,-1),(-1,-1)\}\}$ is a bidirectional OP.

Since there are a few additional technicalities, we give the proof, focusing on the new parts, so the reader is also invited to consult [344] for more details.

Proof of Lemma 12.6.2. Let $u \in I = (\pi/2, 5\pi/4)$ and $\pi/2 - u < \theta < 5\pi/4 - u$. We claim that $d_u^\theta(\{U_1, U_2\}) = d_u^\theta(\mathcal{U}')$, which clearly implies the desired result.

Let $B = [-n, n] \times [0, cn]$ for some fixed $n \in \mathbb{N}$ sufficiently large and $0 \leqslant c \leqslant 1$ sufficiently small ($c < \tan(u - \pi/2)$ if $u \in (\pi/2, \pi)$ and the same with u replaced with $u + \theta$) and define the events

$$\begin{aligned} \mathcal{E}_1 &= \left\{0 \notin [(A \cup V_{u,u+\theta}) \cap B]_{\mathcal{U}'}\right\} \\ \mathcal{E}_2 &= \left\{0 \notin [(A \cup V_{u,u+\theta}) \cap B]_{\{U_1,U_2\}}\right\}. \end{aligned}$$

We argue that $\mathcal{E}_1 \supset \mathcal{E}_2$. Fix a realisation of A such that $\mathcal{E}_2 \backslash \mathcal{E}_1$ holds and call the sites in

$$B \backslash [(A \cup V_{u,u+\theta}) \cap B]_{\{U_1,U_2\}}$$

survivors. Let P be the rightmost path of survivors from 0 with steps in $\{(0,1),(1,1)\}$ (performing the step $(1,1)$ whenever possible and $(0,1)$ only when $(1,1)$ is not possible) and denote x its endpoint. Indeed, P cannot reach the (top) boundary ∂B of B, since $\mathcal{E}_1$ does not hold (survivors are necessarily initially healthy). Since x is a survivor and both $x + (0,1)$ and $x + (1,1)$ are not (otherwise x is not the end of the path), there needs to be a survivor y among $x+(1,0)$ and $x+(1,-1)$ (see Figure 12.3). In particular, $x \neq 0$, as both $(0,1)$ and $(1,-1)$ are in $\mathbb{H}_u \cap \mathbb{H}_{u+\theta}$.

Since y is a survivor, there has to exist a path of survivors starting at y with steps in U_2 reaching ∂B. However, it is easy to see (see Figure 12.3) that such a path cannot reach ∂B without intersecting $V_{u,u+\theta}$ or P. The former possibility is excluded, since $V_{u,u+\theta}$ are not survivors and the latter one contradicts the choice of P to be the rightmost path of survivors from 0.

Hence, $\mathcal{E}_2 \subset \mathcal{E}_1$. A similar reasoning applies with B tilted by $3\pi/4$. Finally, recalling that the region $V_{\pi/2,5\pi/4}$ is entirely infected for all values of (u,θ) considered, we obtain that

$$0 \notin [(A \cup V_{u,u+\theta}) \cap B_n]_{\{U_1,U_2\}} \Longrightarrow 0 \notin [(A \cup V_{u,u+\theta}) \cap B_{cn}]_{\mathcal{U}'}.$$

The same implication with $\mathcal{U}'$ and $\{U_1, U_2\}$ swapped is clear from the fact that $U_1 \supset \{(0,1),(1,1)\}$ and $U_2 \supset \{(0,-1),(-1,-1)\}$, so we are done by Definition 12.2.1. □

Proof of Theorem 12.3.4. First note that if $q < 1 - p_c^{\mathrm{OP}}$, then with probability 1 there exists a bidirectional $\mathcal{U}'$ path of healthy sites, which remains healthy also for $\mathcal{U}$. Therefore, $q_c(\mathcal{U}) \geqslant 1 - p_c^{\mathrm{OP}}$.

We apply Corollary 12.3.2 to $\mathcal{U}$ and the families $\mathcal{U}_1 = \{U_1, U_2\}$, $\mathcal{U}_2 = \{U_2, U_3\}$, $\mathcal{U}_3 = \{U_3, U_4\}$ and $\mathcal{U}_4 = \{U_4, U_1\}$. We simply bound $d_u(\mathcal{U}_1)$ by 1 for $u \in (-\pi, \pi/2]$ and apply Lemma 12.6.2 and Theorem 12.5.5 with Remark 12.5.6 to obtain a bound on $d_u(\mathcal{U}_1)$ for all u. By symmetry the same applies to the other three families up to rotation by $\pi/2$. Hence,

$$\begin{aligned} q_{\mathrm{c}}(\mathcal{U}) \leqslant \tilde{q}_{\mathrm{c}}(\mathcal{U}) &= \sup_{u\in S^1} d_u \leqslant \sup_{u\in(\pi/2,\pi]} d_u(\mathcal{U}_1) \\ &= \sup_{u\in(\pi/2,\pi]} d_u(\mathcal{U}') \leqslant \sup_{u\in S^1} d_u(\mathcal{U}') = 1-p_{\mathrm{c}}^{\mathrm{OP}}. \end{aligned} \qquad \square$$

Remark 12.6.3. It is important to note that Lemma 12.6.2 does not hold for all directions u. It is clear, for example, that when $u = 0$ it suffices to have an infinite uni-directional healthy path with steps $\{(1,0),(1,-1)\}$ starting at 0, which occurs for $q < 1-p_{\mathrm{c}}^{\mathrm{OP}} \neq 0 = d_u(\mathcal{U}')$. Moreover, the complete Spiral model is *not* equivalent to any (uni- or bi-directional) OP, as it is clear from the fact that it has a discontinuous phase transition [344], while the phase transition of OP is continuous [51]—BP occurs for both bidirectional OP involved, but not for Spiral. Thus, it is crucial to restrict the process to half-planes where it is equivalent to OP. This idea also underlies the reasoning of [344].

12.7 Exponential decay and applications

In Section 12.4 we characterised $\tilde{q}_{\mathrm{c}}$ in terms of critical densities and proved that it is the critical probability of exponential decay. We now give a second proof of the latter, which makes the conclusions slightly stronger and more manipulable. For instance, if we assume that $\theta_n(q)$ decays like a power law, (12.3) gives that for $q < \tilde{q}_{\mathrm{c}}$ the exponent is at least -2, which is what we will prove here without assuming that the decay is a power law. Moreover, this method will grant us access to noise sensitivity as well as proving that a one-arm event has strictly positive probability below $\tilde{q}_{\mathrm{c}}$, so that this is indeed a phase transition regardless of whether $q_{\mathrm{c}} = \tilde{q}_{\mathrm{c}}$ or not. Finally, we give a straightforward but important application of exponential decay to the spectral gap and mean infection time of KCM.

As a motivation we start by answering Questions 12 and 14 of Balister, Bollobás, Przykucki and Smith [28]. We then reprove exponential decay and all the results gathered in Theorem 12.3.5 using the method developed by Duminil-Copin, Raoufi and Tassion [127] and then use a modification of the algorithm we made for the proof of exponential decay to also deduce the results concerning noise sensitivity in Theorem 12.3.6.

12.7.1 Answers to Questions 12 and 14 of [28]

Let us begin this section by explaining why, contrary to the expectations of the authors of [28], one should expect exponential decay *above* criticality rather than below it, thus answering Question 12 of that paper. As the reasoning will be identical, we also answer Question 14, but before that we will need to establish the following straightforward fact that will serve as a source of examples.

Proposition 12.7.1. *For every $\varepsilon > 0$ there exists a GOSP model with $q_c \geqslant 1 - \varepsilon$.*

Proof. Fix $1 - q = \varepsilon > 0$ and let $N = N(\varepsilon) \in \mathbb{N}$ be large enough. Consider the following GOSP update family

$$\mathcal{U} = \{U\} = \{\mathbb{H}_{-\pi/2} \cap B_{8N}\}.$$

We perform the following renormalisation. We call a renormalised site $X \in \mathbb{Z}^2$ good if there is a healthy site in $4N.X + B_N$. The renormalised process clearly yields a percolation with parameter larger than p_c^{OP} for N large enough. Indeed, sites are good independently (as $(4NX + B_N) \cap (4NY + B_N) = \varnothing$ for $X \neq Y \in \mathbb{Z}^2$) with probability $1 - q^{|B_N|}$. In particular, for N large enough there is a positive probability that the renormalised site 0 belongs to an infinite OP path of good renormalised sites. But this implies that the ordinary site 0 belongs to an infinite oriented path of healthy vertices in the graph structure on $\mathbb{Z}^2$ defined by U, i.e. 0 remains healthy forever with positive probability. Hence, BP does not occur a.s. and $1 - \varepsilon = q \leqslant q_c$ as desired. □

Question 14

The authors of [28] ask for which subcritical models below criticality there is no infinite path (non-oriented with nearest neighbour steps) of sites in $[A]$ and seem to be in favour of a positive answer for all subcritical BP models. On the one hand, it is indeed *possible* for this scenario to occur and that is the case for the simplest subcritical model—OP.

Proposition 12.7.2. *Consider OP and let $q < q_c$. Then a.s. there is no infinite path in $[A]$.*

Proof. Let $q < q_c$. Recall that the edge speed from Lemma 12.5.1 satisfies $\alpha(1 - q) > \varepsilon$ for some $\varepsilon > 0$. It then follows from Lemma 12.5.2 that with positive probability there exists an infinite initially healthy oriented path $(a_i, i)_{i \in \mathbb{N}}$ (i.e. with $|a_{i+1} - a_i| = 1$ for all i) starting at 0 with $\inf a_i/i \geqslant \varepsilon$. Reflecting this event, we see that with positive probability there exists a bi-infinite oriented path $(a_i, i)_{i \in \mathbb{Z}}$ containing 0 such that $\inf_{i \neq 0} a_i/|i| \geqslant \varepsilon$. By ergodicity and symmetry a.s. there exist two bi-infinite oriented paths of initially healthy vertices $(a_i)_{i \in \mathbb{Z}}$ and $(b_i)_{i \in \mathbb{Z}}$ such that $a_0 < 0$, $b_0 > 0$, $\liminf_{|i| \to \infty} a_i/|i| \geqslant \varepsilon$ and $\limsup_{|i| \to \infty} b_i/|i| \leqslant -\varepsilon$. As these are oriented paths of healthy sites, they never become infected in the BP process. Moreover, the two paths intersect both in the upper and lower half-planes, $\mathbb{H}_{-\pi/2}$ and $\mathbb{H}_{\pi/2}$, forming a contour of sites in $\mathbb{Z}^2 \backslash [A]$ around the origin. In particular, a.s. there is no infinite non-oriented path with nearest neighbour steps in $[A]$ containing the origin, which concludes the proof by ergodicity. □

On the other hand, it is obvious that *any* subcritical model with $q_c > p_c^{SP}$ is an example of the opposite behaviour. Minimal such examples are provided by large enough GOSP as in Proposition 12.7.1, but also by any trivial subcritical model. Indeed, for any $p_c^{SP} < q < q_c$ we a.s. have an infinite non-oriented path of initially infected sites.

As we do not give the characterisation asked for in [28], let us explain why we believe the question to be somewhat extrinsic in the light of the above example and counter-examples. Indeed, the graph structure of $\mathbb{Z}^2$, which defines the infinite path in $[A]$ that [28] asks for, is not relevant to the model itself, defined only by $\mathcal{U}$. For example if one is to replace $\mathcal{U}$ by $2\mathcal{U}$ (e.g. in the above examples) the problem is changed non-trivially, while the bootstrap process is really the same. Finally, let us note that we do not expect that $q_c > p_c^{SP}$ (or $q_c \geqslant p_c^{SP}$) is a necessary condition.

Question 12

With the previous reasoning in mind, let us go back to Question 12 of [28] about exponential decay. The question is whether at $q < q_c$ there would be exponential decay in n of the probability of 0 being connected by sites in $[A]$ to the boundary of B_n, to quote [28] "Here we mean 'connected' in the site percolation sense, although other notions of connectedness are also interesting".

This is not the case, since in many models there is even no decay at all (the probability of being connected in the non-oriented nearest neighbour sense by sites in $[A]$ to the boundary of B_n may remain bounded away from 0 as $n \to \infty$ for some $q < q_c$), let alone exponential one. For example consider any subcritical model with $q_c > p_c^{SP}$. Obviously, for $p_c^{SP} < q < q_c$ there is a positive probability for 0 to be initially connected to infinity by an infected non-oriented nearest neighbour path, but also with probability 1 BP does not occur, so some (positive density of) sites remain healthy forever. This is by no means contradictory, since, e.g. in the example of Proposition 12.7.1, a path, in the graph sense given by the GOSP rule and not the non-oriented nearest neighbour one, of healthy sites witnessing that 0 never becomes infected can easily jump over an infinite infected non-oriented nearest neighbour path in the usual $\mathbb{Z}^2$ sense.

12.7.2 Exponential decay—proof of Theorem 12.3.5

Even though exponential decay below q_c is not always present, we prove that there is exponential decay *above* q_c, as it is well known to be the case for OP (this follows e.g. from Lemmas 12.5.1 and 12.5.3). We shall use the recent method of Duminil-Copin, Raoufi and Tassion [127] in order to prove the exponential decay of the one-arm events E_n from Definition 12.2.3. In fact, much of the proof of [127] calls for no modification.[2] We will only need the following replacement for their Lemma 3.2.

Lemma 12.7.3. *There exists a randomised algorithm determining $\mathbb{1}_{E_n}$ with maximal revealment*

$$\delta \leqslant \frac{3}{n-1}\sum_{k=0}^{n-1}\tilde{\theta}_k(p).$$

Proof. The algorithm is as follows. First pick k uniformly at random in $[1, n-1]$. Let $S \subset B_n$ denote the current set of sites whose state has been checked by the algorithm. We start by revealing (in an arbitrary order) all sites at distance at most C from ∂B_k, the boundary of B_k, and adding them to S. Afterwards we repeat the following. As long as there exists a site $x_0 \in B_n \backslash S$ for which there exist an integer $N \geqslant 1$ and a sequence $x_1, \dots x_N$ of sites in S verifying the following conditions, the algorithm picks one of the possible x_0 arbitrarily and checks its state.

- x_N is at distance at most C from ∂B_k.
- For all $0 < i \leqslant N$ we have $x_{i-1} \in x_i + X$.
- For all $0 < i \leqslant N$ we have that S is a witness of the event $\tau_{x_i}^{B_n} \geqslant i$.

When no such sites remain, the first stage of the algorithm terminates.

If at this point $0 \notin S$, then the algorithm stops. Otherwise, we directly reveal all remaining sites in B_n (in an arbitrary order) and stop.

It is clear that this algorithm does determine $\mathbb{1}_{E_n}$. Indeed, if all sites were revealed, this is vacuously true for any function, while if at the end of the first stage we had $0 \notin S$, we know that E_k does not occur (by definition) and therefore neither does $E_n \subset E_k$ (by extraction of a shorter path from a longer one).

We now proceed to bound its revealment. Fix the value of k and consider a site $x \in \partial B_l$ for some $0 \leqslant l \leqslant n$. The events E_n are such that when x is revealed, we are certain that either $E_{|k-l|}$ translated by x occurs or the original event E_k occurs. Hence, its revealment is at most $\tilde{\theta}_{|k-l|}(q) + \tilde{\theta}_k(q)$. Taking the average on k this gives a maximal revealment bounded by

$$\frac{3}{n-1}\sum_{0}^{n-1}\tilde{\theta}_l(p). \qquad \square$$

With this Lemma we are ready to apply the method of [127] to prove Theorem 12.3.5.

[2] We encourage the reader unfamiliar with that paper to see the second half of the course recording [124], which gives precisely the part we need and precisely in the simpler form we use here adapted to product measures, except for Lemma 12.7.3 we prove.

Proof of Theorem 12.3.5. Let us start by proving the theorem for subcritical models. For the first two items, using Lemma 3.1 of [127] we can repeat the proof of their Theorem 1.2, using the result of [291] (instead of its more general form, Theorem 1.1 of [127]) together with our replacement for their Lemma 3.2—Lemma 12.7.3—and Russo's formula. Setting

$$\hat{q}_{\mathrm{c}} = \sup\left\{q,\ \limsup \frac{\log \sum_0^{n-1} \tilde{\theta}_k(q)}{\log n} \geqslant 1\right\},$$

this yields the following.

- If $q > \hat{q}_{\mathrm{c}}$, then there exists $c(q) > 0$ such that
$$\tilde{\theta}_n(q) \leqslant \exp(-c(q).n).$$
- There exists $c > 0$ such that for $q < \hat{q}_{\mathrm{c}}$
$$\tilde{\theta}(q) \geqslant c.(\hat{q}_{\mathrm{c}} - q) > 0.$$

We next prove that $\hat{q}_{\mathrm{c}} = \tilde{q}_{\mathrm{c}}$.

First notice that $0 \notin [A \cap B_n]$ implies the existence of a path, in the sense of Definition 12.2.3, of sites x_i with $\tau_{x_i}^{B_n} = \infty$ from 0 to ∂B_n (since there are no finite stable healthy sets) with $x_{i+1} \in x_i + X$ and $x_0 = 0$. But such a path needs to come at distance less than $C/4$ of $\partial B_{n/2}$ at some point x_k, so $E_{n/3}$ translated by x_k occurs. Thus, by the union bound

$$\theta_n(q) \leqslant Cn\tilde{\theta}_{n/3}(q).$$

Therefore, exponential decay for $\tilde{\theta}_n$ implies exponential decay for θ_n and thereby $\tilde{q}_{\mathrm{c}} \leqslant \hat{q}_{\mathrm{c}}$ and for $q > \hat{q}_{\mathrm{c}}$ we have (for some other $c(q)$)

$$\theta_n(q) \leqslant \exp(-c(q).n).$$

Conversely, we know that for $q < \hat{q}_{\mathrm{c}}$ the sequence $\tilde{\theta}_n(q)$ converges to $\tilde{\theta}(q) > 0$. Note that on the event E_n there exists a site x with $\tau_x^{B_n} \geqslant n/C$ at distance at most $C/4$ from $\partial B_{n/2}$ in the path in Definition 12.2.3. Then by the union bound we obtain

$$Cn\theta_{\sqrt{n}/(2C)}(q) \geqslant \tilde{\theta}_n(q) \to \tilde{\theta}(q) > 0,$$

since $\tau_0^{B_{C^2n}} \geqslant 4Cn \Rightarrow 0 \notin [A \cap B_{\sqrt{n}}]$. Indeed, since $\mathcal{U}$ is not supercritical, we can find three or four stable directions containing the origin in their convex envelope, which guarantees that $[B_{\sqrt{n}}] \subset B_{\sqrt{Cn}}$ and inside this box sites will become infected at least one at a time. This proves that $\theta_n(q) \geqslant c/n^2$ for some $c > 0$ and thus $q \leqslant \tilde{q}_{\mathrm{c}}$ by (12.3). Hence, $\tilde{q}_{\mathrm{c}} = \hat{q}_{\mathrm{c}}$ and the proof of the first two items is complete.

Let us turn to the third one. As we already observed the occurrence of E_n implies the existence of a site x within distance $C/4$ of $\partial B_{n/2}$ with $\tau_x^{B_n} \geqslant n/C$. However, the event $\tau_x \geqslant n/C$ does not depend on sites outside B_n, so that it is the same as $\tau_x^{B_n} \geqslant n/C$ and the first one's probability is independent of $x \in B_{2n/3}$. Then the union bound gives

$$Cn\mathbb{P}_q(\tau_0 \geqslant n/C) \geqslant \tilde{\theta}_n(q) \to \tilde{\theta}(q) > 0.$$

Thus, for $q < \tilde{q}_c$ we have $\mathbb{P}_q(\tau_0 > n) \geqslant c/n$ for some $c > 0$ and in particular the first moment of τ_0 is infinite, which completes the proof for subcritical models.

For $\mathcal{U}$ critical or supercritical and $q > 0$ it suffices to recall from Remark 12.4.8 that $\mathbb{P}_q(\tau_0 \geqslant n)$ decays exponentially, which immediately implies the exponential decay of $\tilde{\theta}_n(q)$ by the union bound as above and thus completes the proof (the second and third items being void for $\tilde{q}_c = 0$). □

12.7.3 Noise sensitivity—proof of Theorem 12.3.6

We next use the algorithm we have to study noise sensitivity and prove Theorem 12.3.6.

The harder part of the proof of Theorem 12.3.6 relies on the following easy consequence of Theorem 1.8 of Schramm and Steif [318] and Theorem 1.9 of Benjamini, Kalai and Schramm [47].[3]

Theorem 12.7.4 ([47, 318]). *Let G_n be a sequence of cylinder events (depending on finitely many sites). If there exists a randomised algorithm determining the occurrence of G_n with maximal revealment $\delta_n \to 0$, then the sequence is noise sensitive.*

The straightforward converses in Theorem 12.3.6, stated for completeness, follow from the next easy lemma.

Lemma 12.7.5. *Let G_n be a nested sequence of cylinder events such that $\bigcap_n G_n = G_\infty$ and $0 < \mathbb{P}_q(G_\infty) < 1$. Then G_n are not noise sensitive.*

Proof. Firstly, $Var(\mathbb{1}_{G_n}) \to Var(\mathbb{1}_{G_\infty}) \in (0, 1/4]$. Secondly, $\mathbb{1}_{G_n} \xrightarrow{L^2} \mathbb{1}_{G_\infty}$, so that for any $\delta > 0$ there exists n_δ such that for all $n \geqslant n_\delta$ we have $\|\mathbb{1}_{G_n} - \mathbb{1}_{G_{n_\delta}}\|_{L^2} < \delta$. Finally, for any $\varepsilon > 0$ the function $f \mapsto (x \mapsto \mathbb{E}[f(N_\varepsilon(x))|x])$ is an L^2 contraction, so that for all $n \geqslant n_\delta$ we also have $\|\mathbb{1}_{N_\varepsilon(x)\in G_n} - \mathbb{1}_{N_\varepsilon(x)\in G_{n_\delta}}\|_{L^2} < \delta$. These three facts combined imply that it is sufficient to show that for any $\delta > 0$ small enough and any $\varepsilon > 0$ small enough depending on δ it holds that $Var(\mathbb{1}_{G_{n_\delta}}) - Cov(\mathbb{1}_{x\in G_{n_\delta}}, \mathbb{1}_{N_\varepsilon(x)\in G_{n_\delta}}) < \delta$.

[3] The results of these papers are stated for $q = 1/2$, but they are also valid for any fixed value of $0 < q < 1$. Moreover, the result does hold for the stronger Definition 12.2.4.

But this is the case, as G_{n_δ} is a cylinder event, so that for ε small enough $\mathbb{P}_q(\mathbb{1}_{x\in G_{n_\delta}} \neq \mathbb{1}_{N_\varepsilon(x)\in G_{n_\delta}}) < \delta$. Hence,

$$\lim_{\varepsilon\to 0}\liminf_{n\to\infty} \frac{Cov(\mathbb{1}_{x\in G_n}, \mathbb{1}_{N_\varepsilon(x)\in G_n})}{Var(\mathbb{1}_{G_n})} = 1,$$

which concludes the proof by Definition 12.2.4. □

Remark 12.7.6. The consequences of Lemma 12.7.5 can also be deduced easily from [47, Theorem 1.4].

Proof of Theorem 12.3.6. Fix $0 < q < 1$. First assume that $\theta(q) > 0$. Then by Lemma 12.7.5 we have that the events $0 \notin [A \cap B_n]$ are not noise sensitive and then Theorem 12.7.4 proves that no low-revealment algorithm exists. The proof in the case $\tilde{\theta}(q) > 0$ that the events E_n are not noise sensitive is analogous. Assume, on the contrary, that $\tilde{\theta}(q) = 0$. Then Lemma 12.7.3 provides an algorithm with revealment $\delta_n \to 0$, which completes the proof of the first two items of Theorem 12.3.6.

Finally, assume that $\theta(q) = \tilde{\theta}(q) = 0$. Since $\theta(q) = 0$ we also have $\mathbb{P}_q(\tau_0 \geqslant n) \to 0$. Fix $\varepsilon > 0$ and let n be large enough so that we can find $n/C > k_0 > C$ with $k_0 < \varepsilon/(64C\mathbb{P}_q(\tau_0 \geqslant n/C))$ and $\frac{2}{k_0}\sum_{m=0}^{2k_0}\tilde{\theta}_m(q) < \varepsilon$. Denote by H_k the event that there exists x at distance at most C from ∂B_k such that $\tau_x^{B_n} > n/C$. Then by the union bound $\mathbb{P}_q(H_k) < 16Ck\mathbb{P}_q(\tau_0 \geqslant n/C) < \varepsilon$ for $k < 4k_0$.

We perform the same algorithm as in the proof of Lemma 12.7.3, but with k chosen uniformly in $[3k_0, 4k_0)$. When the first stage (exploration) of the algorithm stops we check if H_k occurs, which is indeed known (witnessed by the set of inspected sites S). If it does, then we simply check all the remaining sites to determine if $0 \in [A \cap B_n]$. The probability that this last step occurs is exactly $\mathbb{P}_q(H_k) < \varepsilon$. If H_k does not occur, we know that $0 \in [A \cap B_n]$ (since there are no finite stable healthy sets). We can then bound the revealment similarly to what we did in Lemma 12.7.3—we consider a site $y \in \partial B_l$ and take cases depending on its position. If $l \geqslant 5k_0$, the revealment is at most $\varepsilon + \tilde{\theta}_{l-4k_0}(q) \leqslant \varepsilon + \tilde{\theta}_{k_0}(q) < 2\varepsilon$ and similarly for $l < 2k_0$. For $2k_0 \leqslant l < 5k_0$ we average on k as before to obtain a revealment bounded by $\varepsilon + \frac{2}{k_0}\sum_{m=0}^{2k_0}\tilde{\theta}_m(q)$. Hence, the maximal revealment is indeed bounded by 2ε. Then, as previously, Theorem 12.7.4 gives that $0 \in [A \cap B_n]$ is noise sensitive, which concludes the proof. □

12.7.4 Spectral gap and mean infection time of KCM

To conclude our discussion of exponential decay, we turn to its applications to the KCM defined at the end of the introduction. Cancrini, Martinelli, Roberto and Toninelli [88] proved the positivity of the spectral gap above q_c for several specific models including OP, whose KCM counterpart is known

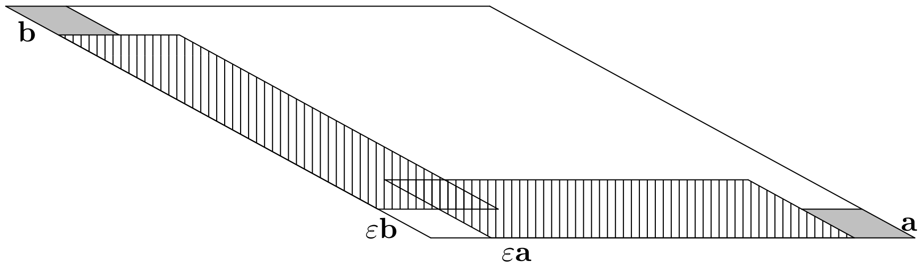

Figure 12.4 – Illustration of the definition of a renormalised site being good. The two hatched parallelograms become infected by the first condition, while the second one concerns the two shaded rhombi.

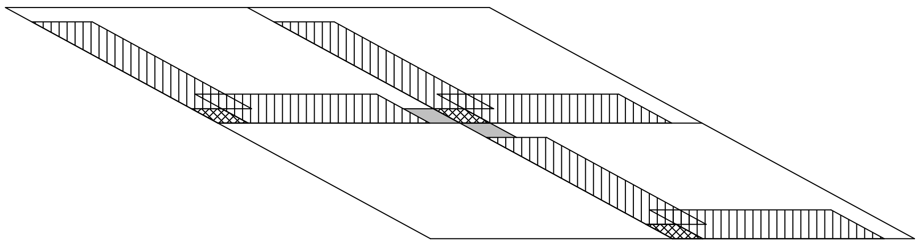

Figure 12.5 – Infection procedure used to prove that if the top-right, bottom-right and top-left renormalised sites are good, the bottom-left one becomes entirely infected.

as the North-East model. They also proved that the result holds for any model under an unhandy additional condition. We now use Theorem 12.3.5 together with their results to prove that for all KCM the gap is positive above $\tilde{q}_{\mathrm{c}}$ and 0 below and the mean infection time of the origin is finite and infinite respectively. It is very interesting to note that we will use the exponential decay of $\tilde{\theta}_n$ and not θ_n, which does not suffice.

In order to link the spectral gap and the mean infection times we need the following simple facts from [272] and [89].

Lemma 12.7.7 (Lemma 4.3 [272], Theorem 4.7 [89]). *For all $0 < q < 1$ the mean infection time of the origin in the BP and the corresponding stationary KCM processes satisfy*

$$\delta\mu\left(\tau_0^{\mathrm{BP}}\right) \leqslant \mathbb{E}_\mu(\tau_0) \leqslant T_{\mathrm{rel}}/q,$$

where T_{rel} is the inverse spectral gap of the KCM and $\delta > 0$ is a sufficiently small constant.

Proof of Theorem 12.3.7. Let $\mathcal{U}$ be a (non-trivial) update family and without loss of generality assume that it contains a rule $U_0 \subset \mathbb{H}_{-\pi/2+\delta} \cap \mathbb{H}_{-\pi/2-2\delta}$ for some $\delta > 0$ sufficiently small such that $-\pi/2 - \delta$ is a rational direction. Fix $q > \tilde{q}_c$ and $\varepsilon(\delta) > 0$ and $\eta(\delta, \varepsilon) > 0$ sufficiently small. The positivity of the gap is implied by Theorem 3.3 of [88] if we can find a suitable renormalisation satisfying the following (see Definition 3.1 [88]).[4]

(a) Each renormalised site is good with probability at least $1 - \varepsilon$.

(b) If the renormalised sites $(0, 1)$, $(1, 0)$ and $(1, 1)$ are all good, then

$$[A \cap (\{\mathbf{a}, \mathbf{b}, \mathbf{a} + \mathbf{b}\} + B')] \supset B',$$

where $\mathbf{a}$ and $\mathbf{b}$ are the two base vectors of the renormalisation and B' is the renormalisation box—the parallelogram generated by $\mathbf{a}$ and $\mathbf{b}$ i.e.

$$B' = ([0, 1) \cdot \mathbf{a}) + ([0, 1) \cdot \mathbf{b}),$$

where we use the notation $C + D = \{c + d,\ c \in C, d \in D\}$.

Set $\mathbf{a} = (n, 0)$ and $\mathbf{b} = n(\cos(-\pi + \delta), \sin(-\pi + \delta))$ for $n(\eta)$ sufficiently large. We call the renormalised site 0 *good* if the following all hold (see Figure 12.4) and we extend the definition to any site by translation.

- For all x in the parallelograms $[\varepsilon, 1 - \varepsilon] \cdot \mathbf{a} + [0, 2\varepsilon] \cdot \mathbf{b}$ and $[\varepsilon, 1 - \varepsilon] \cdot \mathbf{b} + [0, 2\varepsilon] \cdot \mathbf{a}$ it holds that $\tau_x^{B'} < \eta n$.
- For all x in the rhombus $[1 - \varepsilon, 1) \cdot \mathbf{a} + [0, \varepsilon] \cdot \mathbf{b}$ it holds that $\tau_x^{B'} < \eta n$ if we impose infected boundary condition on $[1, 1 + 2\varepsilon] \cdot \mathbf{a} + [0, 1 - \varepsilon] \cdot \mathbf{b}$ and healthy on the rest of $\mathbb{Z}^2 \backslash B'$. Also the symmetric condition holds for the rhombus $[1 - \varepsilon, 1) \cdot \mathbf{b} + [0, \varepsilon] \cdot \mathbf{a}$.

Condition (b) on the renormalisation is easily checked from this definition, using only the rule U_0 (see Figure 12.5). Indeed, all hatched regions become infected by the first condition, so that the double hatched rhombi are infected by U_0. Finally, the shaded rhombi become infected by the second condition, since the infected boundary condition is already met. The renormalised site considered is then entirely infected using U_0. Thus, we only need to check that a renormalised site is good with probability at least $1 - \varepsilon$.

Since the conditions concern $O(n^2)$ sites, by symmetry and monotonicity it suffices to observe that

$$\mathbb{P}_q\left(\tau_0^{[-C\eta n, C\eta n] \times [0, C\eta n]} \geqslant \eta n\right)$$

decays exponentially with n. Indeed, for this event to occur, there must exist a path of sites $x_0, \ldots, x_{\lceil n\eta \rceil} = 0$ with $x_i - x_{i+1} \in U_0$ and $\tau_{x_i}^{[-C\eta n, C\eta n] \times [0, C\eta n]} \geqslant i$ for all $0 \leqslant i < \eta n$, which in particular means that $E_{\eta^2 n}$ translated by x_0 occurs. Hence, using the first item of Theorem 12.3.5 and the union bound

[4] The statement in [88] is given for square boxes, but generalises without change.

we obtain the desired result and thereby the spectral gap is strictly positive. By Lemma 12.7.7 this implies that the mean infection time of the KCM is finite.

Finally, by Theorem 12.3.5 for $q < \tilde{q}_c$ the mean infection time of BP is infinite, so Lemma 12.7.7 shows that in this regime the spectral gap is 0 and the mean infection time of the KCM is infinite. □

12.8 Open problems

To conclude, let us mention some interesting open problems related to this chapter besides its direct extensions based on GOSP.

12.8.1 Simplifications

We next mention the two prime conjectures which would greatly simplify the statements of our results besides being interesting on their own. We start with the uniqueness of the transition.

Conjecture 12.8.1. *For all update families we have*

$$q_c = \tilde{q}_c.$$

We should note that, the Kahn–Kalai–Linial theorem [238] tells us that (up to replacing the box by the torus as in [30] or adapting the technique of [126]) $\theta_n(q)$ decays at least like $n^{-\varepsilon(q-q_c)}$ above criticality and Theorem 12.3.5 establishes that below $\tilde{q}_c$ it decays at most like n^{-2}. As it is commonly the case, it is likely that breaching this gap will prove difficult.

As mentioned earlier if one proves the slightly stronger property

$$\tilde{\theta}(q) > 0 \Rightarrow \theta(q) > 0, \tag{12.8}$$

which implies Conjecture 12.8.1, then Theorem 12.3.6 exhausts the noise sensitivity problem for subcritical BP at least for the most natural event $0 \in [A \cap B_n]$, which we consider since there is no obvious choice of "crossing" event. Indeed, in view of Question 12.8.3 below, it is not clear whether it is relevant to consider the event of complete infection on the torus. Also in the light of Theorem 12.3.6 the converse implication of (12.8) is not uninteresting at $\tilde{q}_c$.

Secondly, it would be practical to know if the complication of taking limits in Definition 12.2.1 is necessary. We suspect that this is never the case.

Question 12.8.2. *What are the continuity properties of d_u^θ as a function of (u, θ)?*

12.8.2 Torus

Although the most natural setting for subcritical models is the infinite volume quantity θ, which is approximated by its restriction to boxes θ_n, another common choice in order to avoid boundary issues is to consider the torus $\mathbb{T}_n = (\mathbb{Z}/n\mathbb{Z})^2$. Indeed, results for critical and supercritical models are meaningful in this setting and are essentially equivalent to the law of the infection time in infinite volume [74]. Yet, for subcritical models the mechanism of infection is rather different—instead of rare large droplets that grow easily we have common droplets which only manage to grow with a lot of help. Owing to this it is not clear how quantities on the torus relate to those on the entire grid. We should mention that most of our results carry through if all is defined on the torus, but it is interesting to note that not even the next question seems to have been answered yet.

Question 12.8.3. *Does one have that for all subcritical families*

$$q_{\mathrm{c}} = \lim_n \inf\{q,\, \mathbb{P}_q([A]_{\mathbb{T}_n} = \mathbb{T}_n) \geqslant 1/2\},$$

where the closure is taken with respect to the BP process on the torus and A is a random subset of $\mathbb{T}_n$ of density q?

Bibliography

[1] M. A. Abdullah and N. Fountoulakis, *A phase transition in the evolution of bootstrap percolation processes on preferential attachment graphs*, Random Structures Algorithms **52** (2018), no. 3, 379–418 pp. MR3783205

[2] J. Adler, *Bootstrap percolation*, Phys. A **171** (1991), no. 3, 453–470 pp.

[3] J. Adler and A. Aharony, *Diffusion percolation. I. Infinite time limit and bootstrap percolation*, J. Phys. A **21** (1988), no. 6, 1387–1404 pp. MR939745

[4] J. Adler and U. Lev, *Bootstrap percolation: visualizations and applications*, Braz. J. Phys. **33** (2003), no. 3, 641–644 pp.

[5] J. Adler, D. Stauffer, and A. Aharony, *Comparison of bootstrap percolation models*, J. Phys. A **22** (1989), no. 7, L297–L301 pp.

[6] M. Aizenman and D. J. Barsky, *Sharpness of the phase transition in percolation models*, Comm. Math. Phys. **108** (1987), no. 3, 489–526 pp. MR874906

[7] M. Aizenman and J. L. Lebowitz, *Metastability effects in bootstrap percolation*, J. Phys. A **21** (1988), no. 19, 3801–3813 pp. MR968311

[8] D. Aldous and P. Diaconis, *The asymmetric one-dimensional constrained Ising model: rigorous results*, J. Stat. Phys. **107** (2002), no. 5-6, 945–975 pp. MR1901508

[9] D. Aldous and J. A. Fill, *Reversible markov chains and random walks on graphs*, 2002. Unfinished monograph, recompiled 2014, available at `http://www.stat.berkeley.edu/~aldous/RWG/book.html`.

[10] D. J. Aldous, *Meeting times for independent Markov chains*, Stochastic Process. Appl. **38** (1991), no. 2, 185–193 pp. MR1119980

[11] G. B. Alers, M. B. Weissman, A. Kinzig, and N. Israeloff, *Noise simulations of an Ising model of glassy kinetics*, Phys. Rev. B **36** (1987), no. 16, 8429–8434 pp.

[12] G. Alon and G. Kozma, *Comparing with octopi*, Ann. Inst. Henri Poincaré Probab. Stat. **56** (2020), no. 4, 2672–2685 pp. MR4164852

[13] N. Alon, *An extremal problem for sets with applications to graph theory*, J. Combin. Theory Ser. A **40** (1985), no. 1, 82–89 pp. MR804870

[14] C. Alves, G. Amir, R. Baldasso, and A. Teixeira, *Percolation phase transition on planar spin systems*, arXiv e-prints (2021), available at arXiv:2105.13314.

[15] C. Alves and R. Ribeiro, *Spread of infection over P.A. random graphs with edge insertion*, arXiv e-prints (2021), available at arXiv:2103.16708.

[16] H. Amini, *Bootstrap percolation in living neural networks*, J. Stat. Phys. **141** (2010), no. 3, 459–475 pp. MR2728841

[17] H. Amini and N. Fountoulakis, *Bootstrap percolation in power-law random graphs*, J. Stat. Phys. **155** (2014), no. 1, 72–92 pp. MR3180970

[18] E. D. Andjel, *Characteristic exponents for two-dimensional bootstrap percolation*, Ann. Probab. **21** (1993), no. 2, 926–935 pp. MR1217573

[19] E. D. Andjel, T. S. Mountford, and R. H. Schonmann, *Equivalence of exponential decay rates for bootstrap-percolation-like cellular automata*, Ann. Inst. Henri Poincaré Probab. Stat. **31** (1995), no. 1, 13–25 pp. MR1340029

[20] C. Ané, S. Blachère, D. Chafaï, P. Fougères, I. Gentil, F. Malrieu, C. Roberto, and G. Scheffer, *Sur les inégalités de Sobolev logarithmiques*, Panoramas et synthèses, vol. 10, Société Mathématique de France, Paris, 2000. MR1845806

[21] O. Angel and B. Kolesnik, *Large deviations for subcritical bootstrap percolation on the random graph*, arXiv e-prints (2017), available at arXiv:1705.06815.

[22] O. Angel and B. Kolesnik, *Sharp thresholds for contagious sets in random graphs*, Ann. Appl. Probab. **28** (2018), no. 2, 1052–1098 pp. MR3784495

[23] F. Arceri, F. P. Landes, L. Berthier, and G. Biroli, *Glasses and aging: a statistical mechanics perspective*, arXiv e-prints (2020), available at arXiv:2006.09725.

[24] D. J. Ashton, *Fast and slow dynamics in kinetically constrained models of glasses*, Ph.D. Thesis, University of Nottingham, 2008.

[25] A. Asselah and P. Dai Pra, *Quasi-stationary measures for conservative dynamics in the infinite lattice*, Ann. Probab. **29** (2001), no. 4, 1733–1754 pp. MR1880240

[26] L. Avena, O. Blondel, and A. Faggionato, *A class of random walks in reversible dynamic environments: antisymmetry and applications to the East model*, J. Stat. Phys. **165** (2016), no. 1, 1–23 pp. MR3547832

[27] P. Balister, B. Bollobás, R. Morris, and P. Smith, *Uncomputability of the percolation threshold for monotone cellular automata* (In preparation).

[28] P. Balister, B. Bollobás, M. Przykucki, and P. Smith, *Subcritical $\mathcal{U}$-bootstrap percolation models have non-trivial phase transitions*, Trans. Amer. Math. Soc. **368** (2016), no. 10, 7385–7411 pp. MR3471095

[29] P. Balister, B. Bollobás, and A. Stacey, *Upper bounds for the critical probability of oriented percolation in two dimensions*, Proc. Roy. Soc. London Ser. A **440** (1993), no. 1908, 201–220 pp. MR1232833

[30] J. Balogh and B. Bollobás, *Sharp thresholds in bootstrap percolation*, Phys. A **326** (2003), no. 3-4, 305–312 pp.

[31] J. Balogh and B. Bollobás, *Bootstrap percolation on the hypercube*, Probab. Theory Related Fields **134** (2006), no. 4, 624–648 pp. MR2214907

[32] J. Balogh, B. Bollobás, H. Duminil-Copin, and R. Morris, *The sharp threshold for bootstrap percolation in all dimensions*, Trans. Amer. Math. Soc. **364** (2012), no. 5, 2667–2701 pp. MR2888224

[33] J. Balogh, B. Bollobás, and R. Morris, *Bootstrap percolation in three dimensions*, Ann. Probab. **37** (2009), no. 4, 1329–1380 pp. MR2546747

[34] J. Balogh, B. Bollobás, and R. Morris, *Majority bootstrap percolation on the hypercube*, Combin. Probab. Comput. **18** (2009), no. 1-2, 17–51 pp. MR2497373

[35] J. Balogh, B. Bollobás, and R. Morris, *Bootstrap percolation in high dimensions*, Combin. Probab. Comput. **19** (2010), no. 5-6, 643–692 pp. MR2726074

[36] J. Balogh, B. Bollobás, and R. Morris, *Graph bootstrap percolation*, Random Structures Algorithms **41** (2012), no. 4, 413–440 pp. MR2993128

[37] J. Balogh, B. Bollobás, R. Morris, and O. Riordan, *Linear algebra and bootstrap percolation*, J. Combin. Theory Ser. A **119** (2012), no. 6, 1328–1335 pp. MR2915649

[38] J. Balogh, G. Kronenberg, A. Pokrovskiy, and T. Szabó, *The maximum length of K_r-bootstrap percolation*, arXiv e-prints (2019), available at arXiv:1907.04559.

[39] J. Balogh, Y. Peres, and G. Pete, *Bootstrap percolation on infinite trees and non-amenable groups*, Combin. Probab. Comput. **15** (2006), no. 5, 715–730 pp. MR2248323

[40] J. Balogh and G. Pete, *Random disease on the square grid*, Random Structures Algorithms **13** (1998), no. 3-4, 409–422 pp. MR1662792

[41] J. Balogh and B. G. Pittel, *Bootstrap percolation on the random regular graph*, Random Structures Algorithms **30** (2007), no. 1-2, 257–286 pp. MR2283230

[42] S. Banerjee, M. Jenamani, and D. K. Pratihar, *A survey on influence maximization in a social network*, Knowl. Inf. Syst. **62** (2020), no. 9, 3417–3455 pp.

[43] Z. Bartha and G. Pete, *Noise sensitivity in bootstrap percolation*, arXiv e-prints (2015), available at arXiv:1509.08454.

[44] E. Bayraktar and S. Chakraborty, *$K_{r,s}$ graph bootstrap percolation*, arXiv e-prints (2019), available at arXiv:1904.12764.

[45] F. Benevides and M. Przykucki, *On slowly percolating sets of minimal size in bootstrap percolation*, Electron. J. Combin. **20** (2013), no. 2, Paper No. 46, 20 pp. MR3084588

[46] F. Benevides and M. Przykucki, *Maximum percolation time in two-dimensional bootstrap percolation*, SIAM J. Discrete Math. **29** (2015), no. 1, 224–251 pp. MR3304258

[47] I. Benjamini, G. Kalai, and O. Schramm, *Noise sensitivity of Boolean functions and applications to percolation*, Inst. Hautes Études Sci. Publ. Math. **90** (1999), 5–43 pp. MR1813223

[48] L. Berthier and J. P. Garrahan, *Nontopographic description of inherent structure dynamics in glassformers*, J. Chem. Phys. **119** (2003), no. 8, 4367–4371 pp.

[49] L. Bertini and C. Toninelli, *Exclusion processes with degenerate rates: convergence to equilibrium and tagged particle*, J. Stat. Phys. **117** (2004), no. 3-4, 549–580 pp. MR2099727

[50] C. Bezuidenhout and L. Gray, *Critical attractive spin systems*, Ann. Probab. **22** (1994), no. 3, 1160–1194 pp. MR1303641

[51] C. Bezuidenhout and G. Grimmett, *The critical contact process dies out*, Ann. Probab. **18** (1990), no. 4, 1462–1482 pp. MR1071804

[52] N. Bhatnagar, P. Caputo, P. Tetali, and E. Vigoda, *Analysis of top-swap shuffling for genome rearrangements*, Ann. Appl. Probab. **17** (2007), no. 4, 1424–1445 pp. MR2344312

[53] G. Biroli and J. P. Garrahan, *Perspective: the glass transition*, J. Chem. Phys. **138** (2013), no. 12, 12A301, 13 pp.

[54] M. Biskup and R. H. Schonmann, *Metastable behavior for bootstrap percolation on regular trees*, J. Stat. Phys. **136** (2009), no. 4, 667–676 pp. MR2540158

[55] D. Blanquicett, *Fixation for two-dimensional $\mathcal{U}$-Ising and $\mathcal{U}$-voter dynamics*, J. Stat. Phys. **182** (2021), no. 2, Paper No. 21, 24 pp. MR4205282

[56] O. Blondel, *Dynamiques de particules sur réseaux avec contraintes cinétiques*, Ph.D. Thesis, Université Paris Diderot, 2013.

[57] O. Blondel, *Front progression in the East model*, Stochastic Process. Appl. **123** (2013), no. 9, 3430–3465 pp. MR3071385

[58] O. Blondel, *Tracer diffusion at low temperature in kinetically constrained models*, Ann. Appl. Probab. **25** (2015), no. 3, 1079–1107 pp. MR3325269

[59] O. Blondel, N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Fredrickson-Andersen one spin facilitated model out of equilibrium*, Markov Process. Related Fields **19** (2013), no. 3, 383–406 pp. MR3156958

[60] O. Blondel, A. Deshayes, and C. Toninelli, *Front evolution of the Fredrickson-Andersen one spin facilitated model*, Electron. J. Probab. **24** (2019), Paper No. 1, 32 pp. MR3903501

[61] O. Blondel, C. Erignoux, M. Sasada, and M. Simon, *Hydrodynamic limit for a facilitated exclusion process*, Ann. Inst. Henri Poincaré Probab. Stat. **56** (2020), no. 1, 667–714 pp. MR4059004

[62] O. Blondel, C. Erignoux, and M. Simon, *Stefan problem for a nonergodic facilitated exclusion process*, Prob. Math. Phys. **2** (2021), no. 1, 127–178 pp.

[63] O. Blondel and C. Toninelli, *Is there a fractional breakdown of the Stokes-Einstein relation in kinetically constrained models at low temperature?*, EPL **107** (2014), no. 2, 26005, 6 pp.

[64] O. Blondel and C. Toninelli, *Kinetically constrained lattice gases: tagged particle diffusion*, Ann. Inst. Henri Poincaré Probab. Stat. **54** (2018), no. 4, 2335–2348 pp. MR3865675

[65] T. Bodineau, V. Lecomte, and C. Toninelli, *Finite size scaling of the dynamical free-energy in a kinetically constrained model*, J. Stat. Phys. **147** (2012), no. 1, 1–17 pp. MR2922756

[66] T. Bodineau and C. Toninelli, *Activity phase transition for constrained dynamics*, Comm. Math. Phys. **311** (2012), no. 2, 357–396 pp. MR2902193

[67] B. Bollobás, *Weakly k-saturated graphs*, Beiträge zur Graphentheorie (Kolloquium, Manebach, 1967), 1968, 25–31 pp. MR0244077

[68] B. Bollobás, *The art of mathematics: Coffee time in Memphis*, Cambridge University Press, New York, 2006. MR2285090

[69] B. Bollobás, H. Duminil-Copin, R. Morris, and P. Smith, *The sharp threshold for the Duarte model*, Ann. Probab. **45** (2017), no. 6B, 4222–4272 pp. MR3737910

[70] B. Bollobás, H. Duminil-Copin, R. Morris, and P. Smith, *Universality of two-dimensional critical cellular automata*, Proc. Lond. Math. Soc. (To appear).

[71] B. Bollobás, S. Griffiths, R. Morris, L. Rolla, and P. Smith, *Nucleation and growth in two dimensions*, Random Structures Algorithms **56** (2020), no. 1, 63–96 pp. MR4052846

[72] B. Bollobás, K. Gunderson, C. Holmgren, S. Janson, and M. Przykucki, *Bootstrap percolation on Galton-Watson trees*, Electron. J. Probab. **19** (2014), Paper No. 13, 27 pp. MR3164766

[73] B. Bollobás, M. Przykucki, O. Riordan, and J. Sahasrabudhe, *On the maximum running time in graph bootstrap percolation*, Electron. J. Combin. **24** (2017), no. 2, Paper No. 2.16, 20 pp. MR3650265

[74] B. Bollobás, P. Smith, and A. Uzzell, *Monotone cellular automata in a random environment*, Combin. Probab. Comput. **24** (2015), no. 4, 687–722 pp. MR3350030

[75] A. Bonato, J. Janssen, and E. Roshanbin, *How to burn a graph*, Internet Math. **12** (2016), no. 1-2, 85–100 pp. MR3474055

[76] L. Bonorino, R. de Paula, P. Gonçalves, and A. Neumann, *Hydrodynamics of porous medium model with slow reservoirs*, J. Stat. Phys. **179** (2020), no. 3, 748–788 pp. MR4099999

[77] A. Bovier and F. den Hollander, *Metastability: A potential-theoretic approach*, Grundlehren der mathematischen Wissenschaften, vol. 351, Springer, Cham, 2015. MR3445787

[78] M. Bradonjić and I. Saniee, *Bootstrap percolation on random geometric graphs*, Probab. Engrg. Inform. Sci. **28** (2014), no. 2, 169–181 pp. MR3244950

[79] M. Bradonjić and I. Saniee, *Bootstrap percolation on periodic trees*, 2015 Proceedings of the Twelfth Workshop on Analytic Algorithmics and Combinatorics (ANALCO), 2015, 89–96 pp. MR3448631

[80] M. Bramson and D. Griffeath, *On the Williams-Bjerknes tumour growth model: II*, Math. Proc. Camb. Phil. Soc. **88** (1980), no. 2, 339–357 pp. MR578279

[81] M. Bramson and D. Griffeath, *On the Williams-Bjerknes tumour growth model. I*, Ann. Probab. **9** (1981), no. 2, 173–185 pp. MR606980

[82] N. S. Branco, *Probabilistic bootstrap percolation*, J. Stat. Phys. **70** (1993), no. 3-4, 1035–1044 pp.

[83] N. S. Branco, R. R. dos Santos, and S. L. A. de Queiroz, *Bootstrap percolation: a renormalisation group approach*, J. Phys. C **17** (1984), no. 14, L373–L377 pp.

[84] K. Bringmann and K. Mahlburg, *Improved bounds on metastability thresholds and probabilities for generalized bootstrap percolation*, Trans. Amer. Math. Soc. **364** (2012), no. 7, 3829–3859 pp. MR2901236

[85] A. Broder and E. Shamir, *On the second eigenvalue of random regular graphs*, 28th Annual Symposium on Foundations of Computer Science, 1987, 286–294 pp.

[86] S. Butler and P. Harrowell, *The origin of glassy dynamics in the 2D facilitated kinetic Ising model*, J. Chem. Phys. **95** (1991), no. 6, 4454–4465 pp.

[87] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Relaxation times of kinetically constrained spin models with glassy dynamics*, J. Stat. Mech. Theory Exp. **03** (2007), L03001, 8 pp.

[88] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Kinetically constrained spin models*, Probab. Theory Related Fields **140** (2008), no. 3-4, 459–504 pp. MR2365481

[89] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Facilitated spin models: recent and new results*, Methods of contemporary mathematical statistical physics, 2009, 307–340 pp. MR2581609

[90] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Kinetically constrained lattice gases*, Comm. Math. Phys. **297** (2010), no. 2, 299–344 pp. MR2651901

[91] N. Cancrini, F. Martinelli, C. Roberto, and C. Toninelli, *Mixing time of a kinetically constrained spin model on trees: power law scaling at criticality*, Probab. Theory Related Fields **161** (2015), no. 1-2, 247–266 pp. MR3304752

[92] N. Cancrini, F. Martinelli, R. Schonmann, and C. Toninelli, *Facilitated oriented spin models: some non equilibrium results*, J. Stat. Phys. **138** (2010), no. 6, 1109–1123 pp. MR2601425

[93] R. Candelier, O. Dauchot, and G. Biroli, *Building blocks of dynamical heterogeneities in dense granular media*, Phys. Rev. Lett. **102** (2009), no. 8, 088001, 4 pp.

[94] R. Candelier, A. Widmer-Cooper, J. K. Kummerfeld, O. Dauchot, G. Biroli, P. Harrowell, and D. R. Reichman, *Spatiotemporal hierarchy of relaxation events, dynamical heterogeneities, and structural reorganization in a supercooled liquid*, Phys. Rev. Lett. **105** (2010), no. 13, 135702, 4 pp.

[95] E. Candellero and N. Fountoulakis, *Bootstrap percolation and the geometry of complex networks*, Stochastic Process. Appl. **126** (2016), no. 1, 234–264 pp. MR3426518

[96] P. Caputo, F. Martinelli, and F. L. Toninelli, *Mixing times of monotone surfaces and SOS interfaces: a mean curvature approach*, Comm. Math. Phys. **311** (2012), no. 1, 157–189 pp. MR2892467

[97] R. Cerf and E. N. M. Cirillo, *Finite size scaling in three-dimensional bootstrap percolation*, Ann. Probab. **27** (1999), no. 4, 1837–1850 pp. MR1742890

[98] R. Cerf and F. Manzo, *The threshold regime of finite volume bootstrap percolation*, Stochastic Process. Appl. **101** (2002), no. 1, 69–82 pp. MR1921442

[99] R. Cerf and F. Manzo, *Nucleation and growth for the Ising model in d dimensions at very low temperatures*, Ann. Probab. **41** (2013), no. 6, 3697–3785 pp. MR3161463

[100] J. Chalupa, P. L. Leath, and G. R. Reich, *Bootstrap percolation on a Bethe lattice*, J. Phys. C **12** (1979), no. 1, L31–L35 pp.

[101] A. K. Chandra, P. Raghavan, W. L. Ruzzo, and R. Smolensky, *The electrical resistance of a graph captures its commute and cover times*, Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing, 1989, 574–586 pp.

[102] C. M. Chaves and B. Koiller, *Universality, thresholds and critical exponents in correlated percolation*, Phys. A **218** (1995), no. 3-4, 271–278 pp.

[103] J. W. Chen, *The contact process on a finite set in higher dimensions*, Chinese J. Contemp. Math. **15** (1994), no. 1, 13–20 pp.

[104] P. Chleboun, A. Faggionato, and F. Martinelli, *Time scale separation and dynamic heterogeneity in the low temperature East model*, Comm. Math. Phys. **328** (2014), no. 3, 955–993 pp. MR3201217

[105] P. Chleboun, A. Faggionato, and F. Martinelli, *Mixing time and local exponential ergodicity of the East-like process in* $\mathbb{Z}^d$, Ann. Fac. Sci. Toulouse Math **24** (2015), no. 4, 717–743 pp. MR3434253

[106] P. Chleboun, A. Faggionato, and F. Martinelli, *Relaxation to equilibrium of generalized East processes on* $\mathbb{Z}^d$*: renormalization group analysis and energy-entropy competition*, Ann. Probab. **44** (2016), no. 3, 1817–1863 pp. MR3502595

[107] P. Chleboun, A. Faggionato, F. Martinelli, and C. Toninelli, *Mixing length scales of low temperature spin plaquettes models*, J. Stat. Phys. **169** (2017), no. 3, 441–471 pp. MR3711606

[108] P. Chleboun and F. Martinelli, *Mixing time bounds for oriented kinetically constrained spin models*, Electron. Commun. Probab. **18** (2013), Paper No. 60, 9 pp. MR3084571

[109] P. Chleboun and A. Smith, *Mixing of the square plaquette model on a critical length scale*, Electron. J. Probab. **25** (2020), Paper No. 89, 53 pp. MR4136469

[110] P. Chleboun and A. Smith, *Cutoff for the square plaquette model on a critical length scale*, Ann. Appl. Probab. **31** (2021), no. 2, 668–702 pp. MR4254492

[111] J.-O. Choi and U. Yu, *Bootstrap and diffusion percolation transitions in three-dimensional lattices*, J. Stat. Mech. Theory Exp. **6** (2020), 063218, 11 pp. MR4161627

[112] F. Chung, P. Diaconis, and R. Graham, *Combinatorics for the East model*, Adv. Appl. Math. **27** (2001), no. 1, 192–206 pp. MR1835679

[113] O. Couronné, *Sur les grands clusters en percolation*, Ph.D. Thesis, Université Paris-Sud, 2004.

[114] Y. Couzinié and F. Martinelli, *On a front evolution problem for the multidimensional East model*, arXiv e-prints (2021), available at arXiv:2112.14693.

[115] S. David, I. Hartarsky, and M. Tiba, *Strong Ramsey games in unbounded time*, European J. Combin. **86** (2020), 103096, 26 pp. MR4073378

[116] P. de Gregorio, A. Lawlor, P. Bradley, and K. A. Dawson, *Clarification of the bootstrap percolation paradox*, Phys. Rev. Lett. **93** (2004), no. 2, 025501, 4 pp.

[117] P. de Gregorio, A. Lawlor, and K. A. Dawson, *Bootstrap percolation*, Encyclopedia of complexity and systems science, 2009, 608–626 pp.

[118] R. de Paula, P. Gonçalves, and A. Neumann, *Energy estimates and convergence of weak solutions of the porous medium equation*, arXiv e-prints (2021), available at arXiv:2101.08900.

[119] P. Dehghanpour and R. H. Schonmann, *Metropolis dynamics relaxation via nucleation and growth*, Comm. Math. Phys. **188** (1997), no. 1, 89–119 pp. MR1471333

[120] P. Dehghanpour and R. H. Schonmann, *A nucleation-and-growth model*, Probab. Theory Related Fields **107** (1997), no. 1, 123–135 pp. MR1427719

[121] A. Deshayes, *An asymptotic shape theorem for random linear growth models*, arXiv e-prints (2015), available at arXiv:1505.05000.

[122] P. Diaconis and L. Saloff-Coste, *Logarithmic Sobolev inequalities for finite Markov chains*, Ann. Appl. Probab. **6** (1996), no. 3, 695–750 pp. MR1410112

[123] J. A. M. S. Duarte, *Simulation of a cellular automaton with an oriented bootstrap rule*, Phys. A **157** (1989), no. 3, 1075–1079 pp.

[124] H. Duminil-Copin, *Sharp threshold phenomena in statistical physics 2/4*, 2017. Course recording available at `https://www.youtube.com/watch?v=I0lSoWH4GIM`.

[125] H. Duminil-Copin and A. Holroyd, *Finite volume bootstrap percolation with balanced threshold rules on* $\mathbb{Z}^2$ (2012). Preliminary preprint.

[126] H. Duminil-Copin, A. Raoufi, and V. Tassion, *A new computation of the critical point for the planar random-cluster model with* $q \geqslant 1$, Ann. Inst. Henri Poincaré Probab. Stat. **54** (2018), no. 1, 422–436 pp. MR3765895

[127] H. Duminil-Copin, A. Raoufi, and V. Tassion, *Sharp phase transition for the random-cluster and Potts models via decision trees*, Ann. of Math. (2) **189** (2019), no. 1, 75–99 pp. MR3898174

[128] H. Duminil-Copin and A. van Enter, *Erratum to "Sharp metastability threshold for an anisotropic bootstrap percolation model"*, Ann. Probab. **44** (2016), no. 2, 1599–1599 pp. MR3474478

[129] H. Duminil-Copin and A. C. D. van Enter, *Sharp metastability threshold for an anisotropic bootstrap percolation model*, Ann. Probab. **41** (2013), no. 3A, 1218–1242 pp. MR3098677

[130] R. Durrett, *On the growth of one-dimensional contact processes*, Ann. Probab. **8** (1980), no. 5, 890–907 pp. MR586774

[131] R. Durrett, *Oriented percolation in two dimensions*, Ann. Probab. **12** (1984), no. 4, 999–1040 pp. MR757768

[132] R. Durrett, *The contact process, 1974–1989*, Mathematics of random media (Blacksburg, VA, 1989), 1991, 1–18 pp. MR1117232

[133] R. Durrett and D. Griffeath, *Contact processes in several dimensions*, Z. Wahrsch. Verw. Gebiete **59** (1982), no. 4, 535–552 pp. MR656515

[134] R. Durrett and D. Griffeath, *Supercritical contact processes on Z*, Ann. Probab. **11** (1983), no. 1, 1–15 pp. MR682796

[135] R. Durrett and X.-F. Liu, *The contact process on a finite set*, Ann. Probab. **16** (1988), no. 3, 1158–1173 pp. MR942760

[136] R. Durrett and R. H. Schonmann, *Stochastic growth models*, Percolation theory and ergodic theory of infinite particle systems (Minneapolis, Minn., 1984–1985), 1987, 85–119 pp. MR894544

[137] R. Durrett and R. H. Schonmann, *The contact process on a finite set. II*, Ann. Probab. **16** (1988), no. 4, 1570–1583 pp. MR958203

[138] R. Durrett and R. H. Schonmann, *Large deviations for the contact process and two-dimensional percolation*, Probab. Theory Related Fields **77** (1988), no. 4, 583–603 pp. MR933991

[139] W. Ertel, K. Froböse, and J. Jäckle, *Constrained diffusion dynamics in the hard-square lattice gas at high density*, J. Chem. Phys. **88** (1988), no. 8, 5027–5034 pp.

[140] A. Ertul, *Cutoff for the Fredrickson-Andersen one spin facilitated model*, arXiv e-prints (2021), available at arXiv:2103.00019.

[141] A. Ertul, *Diffusion et relaxation pour des systèmes de particules avec contraintes cinétiques*, Ph.D. Thesis, Université Lyon 1, 2021.

[142] A. Ertul and A. Shapira, *Self-diffusion coefficient in the Kob-Andersen model*, Electron. Commun. Probab. **26** (2021), Paper No. 3, 12 pp. MR4218031

[143] A. Faggionato, F. Martinelli, C. Roberto, and C. Toninelli, *Aging through hierarchical coalescence in the East model*, Comm. Math. Phys. **309** (2012), no. 2, 459–495 pp. MR2864800

[144] A. Faggionato, F. Martinelli, C. Roberto, and C. Toninelli, *Universality in one-dimensional hierarchical coalescence processes*, Ann. Probab. **40** (2012), no. 4, 1377–1435 pp. MR2978129

[145] A. Faggionato, F. Martinelli, C. Roberto, and C. Toninelli, *The East model: recent results and new progresses*, Markov Process. Related Fields **19** (2013), no. 3, 407–452 pp. MR3156959

[146] A. Faggionato, C. Roberto, and C. Toninelli, *Universality for one-dimensional hierarchical coalescence processes with double and triple merges*, Ann. Appl. Probab. **24** (2014), no. 2, 476–525 pp. MR3178489

[147] V. Falgas-Ravry and A. Sarkar, *Bootstrap percolation in random geometric graphs*, arXiv e-prints (2021), available at arXiv:2110.12166.

[148] A. Fey, L. Levine, and Y. Peres, *Growth rates and explosions in sandpiles*, J. Stat. Phys. **138** (2010), no. 1-3, 143–159 pp. MR2594895

[149] L. R. Fontes, R. H. Schonmann, and V. Sidoravicius, *Stretched exponential fixation in stochastic Ising models at zero temperature*, Comm. Math. Phys. **228** (2002), no. 3, 495–518 pp. MR1918786

[150] L. R. G. Fontes and R. H. Schonmann, *Bootstrap percolation on homogeneous trees has 2 phase transitions*, J. Stat. Phys. **132** (2008), no. 5, 839–861 pp. MR2430783

[151] N. Fountoulakis, M. Kang, C. Koch, and T. Makai, *A phase transition regarding the evolution of bootstrap processes in inhomogeneous random graphs*, Ann. Appl. Probab. **28** (2018), no. 2, 990–1051 pp. MR3784494

[152] P. Frankl, *An extremal problem for two families of sets*, European J. Combin. **3** (1982), no. 2, 125–127 pp. MR670845

[153] G. H. Fredrickson, *Linear and nonlinear experiments for a spin model with cooperative dynamics*, Ann. N. Y. Acad. Sci. **484** (1986), no. 1, 185–205 pp.

[154] G. H. Fredrickson, *Recent developments in dynamical theories of the liquid-glass transition*, Annu. Rev. Phys. Chem. **39** (1988), no. 1, 149–180 pp.

[155] G. H. Fredrickson and H. C. Andersen, *Kinetic Ising model of the glass transition*, Phys. Rev. Lett. **53** (1984), no. 13, 1244–1247 pp.

[156] G. H. Fredrickson and H. C. Andersen, *Facilitated kinetic Ising models and the glass transition*, J. Chem. Phys. **83** (1985), no. 11, 5822–5831 pp.

[157] G. H. Fredrickson and S. A. Brawer, *Monte Carlo investigation of a kinetic Ising model of the glass transition*, J. Chem. Phys. **84** (1986), no. 6, 3351–3366 pp.

[158] E. Friedgut and G. Kalai, *Every monotone graph property has a sharp threshold*, Proc. Amer. Math. Soc. **124** (1996), no. 10, 2993–3002 pp. MR1371123

[159] J. Friedman, J. Kahn, and E. Szemerédi, *On the second eigenvalue of random regular graphs*, Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing, 1989, 587–598 pp.

[160] P. Fröjdh, M. Howard, and K. B. Lauritsen, *Directed percolation and other systems with absorbing states: impact of boundaries*, Int. J. Modern Phys. B **15** (2001), no. 12, 1761–1797 pp.

[161] A. Galves, F. Martinelli, and E. Olivieri, *Large-density fluctuations for the one-dimensional supercritical contact process*, J. Stat. Phys. **55** (1989), no. 3-4, 639–648 pp. MR1003532

[162] S. Ganguly, E. Lubetzky, and F. Martinelli, *Cutoff for the East process*, Comm. Math. Phys. **335** (2015), no. 3, 1287–1322 pp. MR3320314

[163] S. Ganguly and F. Martinelli, *Upper triangular matrix walk: cutoff for finitely many columns*, Random Structures Algorithms **55** (2019), no. 2, 313–341 pp. MR3983784

[164] J. Gao, T. Zhou, and Y. Hu, *Bootstrap percolation on spatial networks*, Sci. Rep. **5** (2015), no. 1, 14662, 10 pp.

[165] O. Garet and R. Marchand, *Asymptotic shape for the chemical distance and first-passage percolation on the infinite Bernoulli cluster*, ESAIM Probab. Stat. **8** (2004), 169–199 pp. MR2085613

[166] O. Garet and R. Marchand, *Asymptotic shape for the contact process in random environment*, Ann. Appl. Probab. **22** (2012), no. 4, 1362–1410 pp. MR2985164

[167] O. Garet and R. Marchand, *Large deviations for the contact process in random environment*, Ann. Probab. **42** (2014), no. 4, 1438–1479 pp. MR3262483

[168] O. Garet, R. Marchand, and M. Théret, *Continuity of the asymptotic shape of the supercritical contact process*, Electron. Commun. Probab. **20** (2015), Paper No. 92, 11 pp. MR3434209

[169] J. P. Garrahan, *Glassiness through the emergence of effective dynamical constraints in interacting systems*, J. Phys.: Condens. Matter **14** (2002), no. 7, 1571–1579 pp.

[170] J. P. Garrahan, *Aspects of non-equilibrium in classical and quantum systems: slow relaxation and glasses, dynamical large deviations, quantum non-ergodicity, and open quantum dynamics*, Phys. A **504** (2018), 130–154 pp. MR3805503

[171] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, *Dynamical first-order phase transition in kinetically constrained models of glasses*, Phys. Rev. Lett. **98** (2007), no. 19, 195702, 4 pp.

[172] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, *First-order dynamical phase transition in models of glasses: an approach based on ensembles of histories*, J. Phys. A **42** (2009), no. 7, 075007, 34 pp. MR2525457

[173] P. Garrahan, P. Sollich, and C. Toninelli, *Kinetically constrained models*, Dynamical heterogeneities in glasses, colloids and granular media and jamming transitions, 2011, 341–369 pp.

[174] R. J. Glauber, *Time-dependent statistics of the Ising model*, J. Mathematical Phys. **4** (1963), no. 2, 294–307 pp. MR148410

[175] S. Goel, R. Montenegro, and P. Tetali, *Mixing time bounds via the spectral profile*, Electron. J. Probab. **11** (2006), Paper No. 1, 1–26 pp. MR2199053

[176] S. Goldstein, J. L. Lebowitz, and E. R. Speer, *The discrete-time facilitated totally asymmetric simple exclusion process*, Pure Appl. Funct. Anal. **6** (2021), no. 1, 177–203 pp. MR4213301

[177] P. Gonçalves, *Microscopic dynamics for the porous medium equation*, Dynamics, games and science. II, 2011, 387–392 pp. MR2883293

[178] P. Gonçalves, C. Landim, and C. Toninelli, *Hydrodynamic limit for a particle system with degenerate rates*, Ann. Inst. Henri Poincaré Probab. Stat. **45** (2009), no. 4, 887–909 pp. MR2572156

[179] I. S. Graham, L. Piché, and M. Grant, *Dynamics and effective thermodynamics of a model structural glass*, J. Phys.: Condens. Matter **5** (1993), no. 29, L349–L354 pp.

[180] I. S. Graham, L. Piché, and M. Grant, *Model for dynamics of structural glasses*, Phys. Rev. E **55** (1997), no. 3, 2132–2144 pp.

[181] J. Gravner and D. Griffeath, *First passage times for threshold growth dynamics on* $\mathbf{Z}^2$, Ann. Probab. **24** (1996), no. 4, 1752–1778 pp. MR1415228

[182] J. Gravner and D. Griffeath, *Nucleation parameters for discrete threshold growth on* $\mathbf{Z}^2$, Experiment. Math. **6** (1997), no. 3, 207–220 pp. MR1481590

[183] J. Gravner and D. Griffeath, *Cellular automaton growth on* $\mathbf{Z}^2$*: theorems, examples, and problems*, Adv. Appl. Math. **21** (1998), no. 2, 241–304 pp. MR1634709

[184] J. Gravner and D. Griffeath, *Scaling laws for a class of critical cellular automaton growth rules*, Random walks (Budapest, 1998), 1999, 167–186 pp. MR1752894

[185] J. Gravner and D. Griffeath, *Random growth models with polygonal shapes*, Ann. Probab. **34** (2006), no. 1, 181–218 pp. MR2206346

[186] J. Gravner, C. Hoffman, J. Pfeiffer, and D. Sivakoff, *Bootstrap percolation on the Hamming torus*, Ann. Appl. Probab. **25** (2015), no. 1, 287–323 pp. MR3297774

[187] J. Gravner and A. E. Holroyd, *Slow convergence in bootstrap percolation*, Ann. Appl. Probab. **18** (2008), no. 3, 909–928 pp. MR2418233

[188] J. Gravner and A. E. Holroyd, *Local bootstrap percolation*, Electron. J. Probab. **14** (2009), Paper No. 14, 385–399 pp. MR2480546

[189] J. Gravner and A. E. Holroyd, *Polluted bootstrap percolation with threshold two in all dimensions*, Probab. Theory Related Fields **175** (2019), no. 1-2, 467–486 pp. MR4009714

[190] J. Gravner, A. E. Holroyd, and R. Morris, *A sharper threshold for bootstrap percolation in two dimensions*, Probab. Theory Related Fields **153** (2012), no. 1-2, 1–23 pp. MR2925568

[191] J. Gravner, A. E. Holroyd, and D. Sivakoff, *Polluted bootstrap percolation in three dimensions*, Ann. Appl. Probab. **31** (2021), no. 1, 218–246 pp. MR4254479

[192] J. Gravner and E. McDonald, *Bootstrap percolation in a polluted environment*, J. Stat. Phys. **87** (1997), no. 3-4, 915–927 pp. MR1459046

[193] L. Gray, J. C. Wierman, and R. T. Smythe, *Lower bounds for the critical probability in percolation models with oriented bonds*, J. Appl. Probab. **17** (1980), no. 4, 979–986 pp. MR587198

[194] D. Griffeath, *Additive and cancellative interacting particle systems*, Lecture Notes in Mathematics, vol. 724, Springer, Berlin, Heidelberg, 1979. MR538077

[195] D. Griffeath, *The basic contact processes*, Stochastic Process. Appl. **11** (1981), no. 2, 151–185 pp. MR616064

[196] G. Grimmett, *Percolation*, Second edition, Grundlehren der mathematischen Wissenschaften, Springer, Berlin, Heidelberg, 1999. Originally published by Springer, New York (1989). MR1707339

[197] G. Grimmett and P. Hiemer, *Directed percolation and random walk*, In and out of equilibrium, 2002, 273–297 pp. MR1901958

[198] A. Guggiola and G. Semerjian, *Minimal contagious sets in random regular graphs*, J. Stat. Phys. **158** (2015), no. 2, 300–358 pp. MR3299880

[199] K. Gunderson, S. Koch, and M. Przykucki, *The time of graph bootstrap percolation*, Random Structures Algorithms **51** (2017), no. 1, 143–168 pp. MR3668849

[200] K. Gunderson and M. Przykucki, *Lower bounds for bootstrap percolation on Galton-Watson trees*, Electron. Commun. Probab. **19** (2014), Paper No. 42, 7 pp. MR3233204

[201] L. Hambardzumyan, H. Hatami, and Y. Qian, *Lower bounds for graph bootstrap percolation via properties of polynomials*, J. Combin. Theory Ser. A **174** (2020), 105253, 12 pp. MR4082061

[202] T. E. Harris, *A lower bound for the critical probability in a certain percolation process*, Math. Proc. Camb. Phil. Soc. **56** (1960), no. 1, 13–20 pp. MR115221

[203] T. E. Harris, *Contact interactions on a lattice*, Ann. Probab. **2** (1974), no. 6, 969–988 pp. MR356292

[204] T. E. Harris, *On a class of set-valued Markov processes*, Ann. Probab. **4** (1976), no. 2, 175–194 pp. MR400468

[205] P. Harrowell, *Visualizing the collective motions responsible for the α and β relaxations in a model glass*, Phys. Rev. E **48** (1993), no. 6, 4359–4363 pp.

[206] I. Hartarsky, *Maximal bootstrap percolation time on the hypercube via generalised snake-in-the-box*, Electron. J. Combin. **25** (2018), no. 3, Paper No. 3.10, 12 pp. MR3853862

[207] I. Hartarsky, *Bisection for kinetically constrained models revisited*, Electron. Commun. Probab. **26** (2021), Paper No. 60, 10 pp.

[208] I. Hartarsky, *Bootstrap percolation, probabilistic cellular automata and sharpness*, arXiv e-prints (2021), available at arXiv:2112.01778.

[209] I. Hartarsky, *$\mathcal{U}$-bootstrap percolation: critical probability, exponential decay and applications*, Ann. Inst. Henri Poincaré Probab. Stat. **57** (2021), no. 3, 1255–1280 pp. MR4291442

[210] I. Hartarsky, *Refined universality for critical KCM: upper bounds*, arXiv e-prints (2021), available at arXiv:2104.02329.

[211] I. Hartarsky, M. Baity-Jesi, R. Ravasio, A. Billoire, and G. Biroli, *Maximum-energy records in glassy energy landscapes*, J. Stat. Mech. Theory Exp. **9** (2019), 093302, 26 pp. MR4021480

[212] I. Hartarsky and B. N. B. de Lima, *Weakly constrained-degree percolation on the hypercubic lattice*, arXiv e-prints (2020), available at arXiv:2010.08955.

[213] I. Hartarsky and L. Marêché, *Refined universality for critical KCM: lower bounds*, Combin. Probab. Comput. (To appear).

[214] I. Hartarsky, L. Marêché, and C. Toninelli, *Universality for critical KCM: infinite number of stable directions*, Probab. Theory Related Fields **178** (2020), no. 1, 289–326 pp. MR4146539

[215] I. Hartarsky, F. Martinelli, and C. Toninelli, *Sharp threshold for the FA-2f kinetically constrained model*, arXiv e-prints (2020), available at arXiv:2012.02557.

[216] I. Hartarsky, F. Martinelli, and C. Toninelli, *Universality for critical KCM: finite number of stable directions*, Ann. Probab. **49** (2021), no. 5, 2141–2174 pp. MR4317702

[217] I. Hartarsky, F. Martinelli, and C. Toninelli, *Sharp threshold for the FA-jf kinetically constrained model* (In preparation).

[218] I. Hartarsky, F. Martinelli, and C. Toninelli, *Coalescing and branching simple symmetric exclusion process*, Ann. Appl. Probab. (To appear).

[219] I. Hartarsky and T. R. Mezei, *Complexity of two-dimensional bootstrap percolation difficulty: algorithm and NP-hardness*, SIAM J. Discrete Math. **34** (2020), no. 2, 1444–1459 pp. MR4117299

[220] I. Hartarsky and R. Morris, *The second term for two-neighbour bootstrap percolation in two dimensions*, Trans. Amer. Math. Soc. **372** (2019), no. 9, 6465–6505 pp. MR4024528

[221] I. Hartarsky and R. Szabó, *Generalised oriented site percolation, probabilistic cellular automata and bootstrap percolation*, arXiv e-prints (2021), available at arXiv:2103.15621v1.

[222] I. Hartarsky and R. Szabó, *Generalised oriented site percolation*, Markov Process. Related Fields (To appear).

[223] J. Hermon and Y. Peres, *A characterization of L_2 mixing and hypercontractivity via hitting times and maximal inequalities*, Probab. Theory Related Fields **170** (2018), no. 3-4, 769–800 pp. MR3773799

[224] H. Hinrichsen, *Non-equilibrium critical phenomena and phase transitions into absorbing states*, Adv. Phys. **49** (2000), no. 7, 815–958 pp.

[225] A. E. Holroyd, *Sharp metastability threshold for two-dimensional bootstrap percolation*, Probab. Theory Related Fields **125** (2003), no. 2, 195–224 pp. MR1961342

[226] A. E. Holroyd, *The metastability threshold for modified bootstrap percolation in d dimensions*, Electron. J. Probab. **11** (2006), Paper No. 17, 418–433 pp. MR2223042

[227] A. E. Holroyd, T. M. Liggett, and D. Romik, *Integrals, partitions, and cellular automata*, Trans. Amer. Math. Soc. **356** (2003), no. 8, 3349–3368 pp. MR2052953

[228] E. Ising, *Beitrag zur Theorie des Ferromagnetismus*, Z. Phys. **31** (1925), no. 1, 253–258 pp.

[229] R. L. Jack, P. Mayer, and P. Sollich, *Mappings between reaction–diffusion and kinetically constrained systems: $A + A \leftrightarrow A$ and the Fredrickson–Andersen model have upper critical dimension $d_c = 2$*, J. Stat. Mech. Theory Exp. **03** (2006), P03006, 33 pp.

[230] R. L. Jack, T. Nemoto, and V. Lecomte, *Dynamical phase coexistence in the Fredrickson–Andersen model*, J. Stat. Mech. Theory Exp. **5** (2020), 053204, 35 pp. MR4153014

[231] J. Jäckle and S. Eisinger, *A hierarchically constrained kinetic Ising model*, Z. Phys. B Con. Mat. **84** (1991), no. 1, 115–124 pp.

[232] J. Jäckle, K. Froböse, and D. Knödler, *Size dependence of self-diffusion in the hard-square lattice gas*, J. Stat. Phys. **63** (1991), no. 1-2, 249–260 pp.

[233] S. Janson, *On percolation in random graphs with given vertex degrees*, Electron. J. Probab. **14** (2009), Paper No. 5, 87–118 pp. MR2471661

[234] S. Janson, T. Łuczak, T. Turova, and T. Vallier, *Bootstrap percolation on the random graph $G_{n,p}$*, Ann. Appl. Probab. **22** (2012), no. 5, 1989–2047 pp. MR3025687

[235] M. Jeng and J. M. Schwarz, *On the study of jamming percolation*, J. Stat. Phys. **131** (2008), no. 4, 575–595 pp. MR2398943

[236] M. Jeng and J. M. Schwarz, *Force-balance percolation*, Phys. Rev. E **81** (2010), no. 1, 011134, 14 pp.

[237] M. Jerrum and A. Sinclair, *Approximating the permanent*, SIAM J. Comput. **18** (1989), no. 6, 1149–1178 pp. MR1025467

[238] J. Kahn, G. Kalai, and N. Linial, *The influence of variables on Boolean functions*, 29th Annual Symposium on Foundations of Computer Science, 1988, 68–80 pp.

[239] G. Kalai, *Weakly saturated graphs are rigid*, Convexity and graph theory (Jerusalem, 1981), 1984, 189–190 pp. MR791030

[240] V. Kanade, F. Mallmann-Trenn, and T. Sauerwald, *On coalescence time in graphs: when is coalescing as fast as meeting?*, Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, 2019, 956–965 pp. MR3909528

[241] M. Kang and T. Makai, *Bootstrap percolation on $G(n,p)$ revisited*, Proceedings of the 27th International Conference on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms—AofA'16, 2016, 12 pp. MR3817532

[242] R. M. Karp, *Reducibility among combinatorial problems*, Complexity of computer computations, 1972, 85–103 pp. MR0378476

[243] J. F. C. Kingman, *Subadditive ergodic theory*, Ann. Probab. **1** (1973), no. 6, 883–899 pp. MR356192

[244] W. Kob and H. C. Andersen, *Kinetic lattice-gas model of cage effects in high-density liquids and a test of mode-coupling theory of the ideal-glass transition*, Phys. Rev. E **48** (1993), no. 6, 4364–4377 pp.

[245] P. M. Kogut and P. L. Leath, *Bootstrap percolation transitions on real lattices*, J. Phys. C **14** (1981), no. 22, 3187–3194 pp.

[246] G. Kordzakhia and S. P. Lalley, *Ergodicity and mixing properties of the northeast model*, J. Appl. Probab. **43** (2006), no. 3, 782–792 pp. MR2274800

[247] G. Kronenberg, T. Martins, and N. Morrison, *Weak saturation numbers of complete bipartite graphs in the clique*, J. Combin. Theory Ser. A **178** (2021), 105357, 15 pp. MR4169833

[248] T. Kuczek, *The central limit theorem for the right edge of supercritical oriented percolation*, Ann. Probab. **17** (1989), no. 4, 1322–1332 pp. MR1048929

[249] H. Lacoin, *The scaling limit for zero-temperature planar Ising droplets: with and without magnetic fields*, Topics in percolative and disordered systems, 2014, 85–120 pp. MR3229287

[250] J. L. Lebowitz and R. H. Schonmann, *On the asymptotics of occurrence times of rare events for stochastic spin systems*, J. Stat. Phys. **48** (1987), no. 3-4, 727–751 pp. MR914904

[251] J. L. Lebowitz and R. H. Schonmann, *Pseudo-free energies and large deviations for non-Gibbsian FKG measures*, Probab. Theory Related Fields **77** (1988), no. 1, 49–64 pp. MR921818

[252] T.-Y. Lee and H.-T. Yau, *Logarithmic Sobolev inequality for some models of random walks*, Ann. Probab. **26** (1998), no. 4, 1855–1873 pp. MR1675008

[253] R. Lenormand and C. Zarcone, *Growth of clusters during imbibition in a network of capillaries*, Kinetics of aggregation and gelation, 1984, 177–180 pp.

[254] W. Lenz, *Beitrag zum Verständnis der magnetischen Erscheinungen in festen Körpern*, Physik. Z **21** (1920), 613–615 pp.

[255] S. Léonard, *Hétérogénéité dynamique et échelles de longueur dans les systèmes vitreux hors-équilibre*, Ph.D. Thesis, Université Montpellier II, 2007.

[256] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*, American Mathematical Society, Providence, RI, 2009. With a chapter by J. G. Propp and D. B. Wilson. MR2466937

[257] T. M. Liggett, *Stochastic interacting systems: contact, voter and exclusion processes*, Grundlehren der mathematischen Wissenschaften, vol. 324, Springer, Berlin, Heidelberg, 1999. MR1717346

[258] T. M. Liggett, *Interacting particle systems*, Classics in mathematics, Springer, Berlin, Heidelberg, 2005. Originally published by Springer, New York (1985). MR2108619

[259] T. M. Liggett, R. H. Schonmann, and A. M. Stacey, *Domination by product measures*, Ann. Probab. **25** (1997), no. 1, 71–95 pp. MR1428500

[260] O. Louidor, R. Tessler, and A. Vandenberg-Rodes, *The Williams–Bjerknes model on regular trees*, Ann. Appl. Probab. **24** (2014), no. 5, 1889–1917 pp. MR3226167

[261] R. Lyons and Y. Peres, *Probability on trees and networks*, Cambridge series in statistical and probabilistic mathematics, vol. 42, Cambridge University Press, Cambridge, 2016. MR3616205

[262] D. Maldonado, *Universalité et complexité des automates cellulaires coagulants*, Ph.D. Thesis, Université d'Orléans, 2018.

[263] L. Marêché, *Exponential convergence to equilibrium for the d-dimensional East model*, Electron. Commun. Probab. **24** (2019), Paper No. 55, 10 pp. MR4003129

[264] L. Marêché, *Kinetically constrained models : relaxation to equilibrium and universality results*, Ph.D. Thesis, Université Paris Diderot, Sorbonne Paris Cité, 2019.

[265] L. Marêché, *Combinatorics for general kinetically constrained spin models*, SIAM J. Discrete Math. **34** (2020), no. 1, 370–384 pp. MR4062795

[266] L. Marêché, *Exponential convergence to equilibrium in supercritical kinetically constrained models at high temperature*, Markov Process. Related Fields **26** (2020), no. 5, 861–884 pp. MR4246011

[267] L. Marêché, F. Martinelli, and C. Toninelli, *Exact asymptotics for Duarte and supercritical rooted kinetically constrained models*, Ann. Probab. **48** (2020), no. 1, 317–342 pp. MR4079438

[268] F. Martinelli, *Lectures on Glauber dynamics for discrete spin models*, Lectures on probability theory and statistics (Saint-Flour, 1997), 1999, 93–191 pp. MR1746301

[269] F. Martinelli, R. Morris, and C. Toninelli, *Universality results for kinetically constrained spin models in two dimensions*, Comm. Math. Phys. **369** (2019), no. 2, 761–809 pp. MR3962008

[270] F. Martinelli, A. Shapira, and C. Toninelli, *Diffusive scaling of the Kob–Andersen model in* $\mathbb{Z}^d$, Ann. Inst. Henri Poincaré Probab. Stat. **56** (2020), no. 3, 2189–2210 pp. MR4116722

[271] F. Martinelli and C. Toninelli, *Kinetically constrained spin models on trees*, Ann. Appl. Probab. **23** (2013), no. 5, 1967–1987 pp. MR3134727

[272] F. Martinelli and C. Toninelli, *Towards a universality picture for the relaxation to equilibrium of kinetically constrained models*, Ann. Probab. **47** (2019), no. 1, 324–361 pp. MR3909971

[273] K. Matzke, *The saturation time of graph bootstrap percolation*, arXiv e-prints (2015), available at arXiv:1510.06156.

[274] F. Mauch and J. Jäckle, *Recursive dynamics in an asymmetrically constrained kinetic Ising chain*, Phys. A **262** (1999), no. 1-2, 98–117 pp.

[275] M. C. Medeiros and C. M. Chaves, *Universality in bootstrap and diffusion percolation*, Phys. A **234** (1997), no. 3, 604–610 pp.

[276] M. V. Menshikov, *Coincidence of critical points in percolation problems*, Dokl. Akad. Nauk SSSR **288** (1986), no. 6, 1308–1311 pp. MR852458

[277] L. Miclo, *An example of application of discrete Hardy's inequalities*, Markov Process. Related Fields **5** (1999), no. 3, 319–330 pp. MR1710983

[278] R. Morris, *Minimal percolating sets in bootstrap percolation*, Electron. J. Combin. **16** (2009), no. 1, Paper No. 2, 20 pp. MR2475525

[279] R. Morris, *Zero-temperature Glauber dynamics on* $\mathbb{Z}^d$, Probab. Theory Related Fields **149** (2011), no. 3-4, 417–434 pp. MR2776621

[280] R. Morris, *Bootstrap percolation, and other automata*, European J. Combin. **66** (2017), 250–263 pp. MR3692148

[281] R. Morris, *Monotone cellular automata*, Surveys in combinatorics 2017, 2017, 312–371 pp. MR3728111

[282] N. Morrison and J. A. Noel, *Extremal bounds for bootstrap percolation in the hypercube*, J. Combin. Theory Ser. A **156** (2018), 61–84 pp. MR3762103

[283] T. Mountford and G. Valle, *Exponential convergence for the Fredrickson-Andersen one-spin facilitated model*, J. Theoret. Probab. **32** (2019), no. 1, 282–302 pp. MR3908915

[284] T. S. Mountford, *Rates for the probability of large cubes being non-internally spanned in modified bootstrap percolation*, Probab. Theory Related Fields **93** (1992), no. 2, 159–167 pp. MR1176723

[285] T. S. Mountford, *A metastable result for the finite multidimensional contact process*, Canad. Math. Bull. **36** (1993), no. 2, 216–226 pp. MR1222537

[286] T. S. Mountford, *Critical length for semi-oriented bootstrap percolation*, Stochastic Process. Appl. **56** (1995), no. 2, 185–205 pp. MR1325219

[287] T. S. Mountford, *Existence of a constant for finite system extinction*, J. Stat. Phys. **96** (1999), no. 5-6, 1331–1341 pp. MR1722992

[288] Y. Nagahata, *Lower bound estimate of the spectral gap for simple exclusion process with degenerate rates*, Electron. J. Probab. **17** (2012), Paper No. 92, 19 pp. MR2994840

[289] H. Nakanishi and H. Takano, *Numerical study on the kinetic Ising model for glass transition*, Phys. Lett. A **115** (1986), no. 3, 117–121 pp.

[290] M. E. J. Newman and C. Moore, *Glassy dynamics and aging in an exactly solvable spin model*, Phys. Rev. E **60** (1999), no. 5, 5068–5072 pp.

[291] R. O'Donnell, M. Saks, O. Schramm, and R. A. Servedio, *Every decision tree has an influential variable*, 46th Annual IEEE Symposium on Foundations of Computer Science, 2005, 31–39 pp.

[292] R. I. Oliveira and Y. Peres, *Random walks on graphs: new bounds on hitting, meeting, coalescing and returning*, 2019 Proceedings of the Sixteenth Workshop on Analytic Algorithmics and Combinatorics (ANALCO), 2019, 119–126 pp. MR3909448

[293] E. Olivieri and M. E. Vares, *Large deviations and metastability*, Encyclopedia of mathematics and its applications, vol. 100, Cambridge University Press, Cambridge, 2005. MR2123364

[294] R. N. Onody and U. P. C. Neves, *Series expansion of the directed percolation probability*, J. Phys. A **25** (1992), no. 24, 6609–6615 pp. MR1210881

[295] N. H. Packard and S. Wolfram, *Two-dimensional cellular automata*, J. Stat. Phys. **38** (1985), no. 5-6, 901–946 pp. MR802562

[296] N. Pancotti, G. Giudice, J. I. Cirac, J. P. Garrahan, and M. C. Bañuls, *Quantum East model: localization, nonthermal eigenstates, and slow dynamics*, Phys. Rev. X **10** (2020), no. 2, 021051, 21 pp.

[297] Y. Peres and A. Sly, *Mixing of the upper triangular matrix walk*, Probab. Theory Related Fields **156** (2013), no. 3-4, 581–591 pp. MR3078280

[298] Y. Peres and P. Winkler, *Can extra updates delay mixing?*, Comm. Math. Phys. **323** (2013), no. 3, 1007–1016 pp. MR3106501

[299] N. S. Pillai and A. Smith, *Mixing times for a constrained Ising process on the torus at low density*, Ann. Probab. **45** (2017), no. 2, 1003–1070 pp. MR3630292

[300] N. S. Pillai and A. Smith, *Mixing times for a constrained Ising process on the two-dimensional torus at low density*, Ann. Inst. Henri Poincaré Probab. Stat. **55** (2019), no. 3, 1649–1678 pp. MR4010947

[301] N. Pippenger, *The hypercube of resistors, asymptotic expansions, and preferential arrangements*, Math. Mag. **83** (2010), no. 5, 331–346 pp. MR2762645

[302] B. Pittel, J. Spencer, and N. Wormald, *Sudden emergence of a giant k-core in a random graph*, J. Combin. Theory Ser. B **67** (1996), no. 1, 111–151 pp. MR1385386

[303] M. Pollak and I. Riess, *Application of percolation theory to 2D-3D Heisenberg ferromagnets*, Phys. stat. sol. (b) **69** (1975), no. 1, K15–K18 pp.

[304] M. Przykucki, *Maximal percolation time in hypercubes under 2-bootstrap percolation*, Electron. J. Combin. **19** (2012), no. 2, Paper No. 41, 13 pp. MR2946099

[305] D. Reimer, *Proof of the van den Berg-Kesten conjecture*, Combin. Probab. Comput. **9** (2000), no. 1, 27–32 pp. MR1751301

[306] J. Reiter, *Statics and dynamics of the two-spin–facilitated kinetic Ising model*, J. Chem. Phys. **95** (1991), no. 1, 544–554 pp.

[307] J. Reiter, F. Mauch, and J. Jäckle, *Blocking transitions in lattice spin models with directed kinetic constraints*, Phys. A **184** (1992), no. 3, 458–476 pp.

[308] E. Riedl, *Largest minimal percolating sets in hypercubes under 2-bootstrap percolation*, Electron. J. Combin. **17** (2010), no. 1, Paper No. 80, 13 pp. MR2651733

[309] F. Ritort and P. Sollich, *Glassy dynamics of kinetically constrained models*, Adv. Phys. **52** (2003), no. 4, 219–342 pp.

[310] L. Saloff-Coste, *Lectures on finite Markov chains*, Lectures on probability theory and statistics (Saint-Flour, 1996), 1997, 301–413 pp. MR1490046

[311] F. Sausset, C. Toninelli, G. Biroli, and G. Tarjus, *Bootstrap percolation and kinetically constrained models on hyperbolic lattices*, J. Stat. Phys. **138** (2010), no. 1-3, 411–430 pp. MR2594903

[312] R. H. Schonmann, *Metastability for the contact process*, J. Stat. Phys. **41** (1985), no. 3-4, 445–464 pp. MR814841

[313] R. H. Schonmann, *Critical points of two-dimensional bootstrap percolation-like cellular automata*, J. Stat. Phys. **58** (1990), no. 5-6, 1239–1244 pp. MR1049067

[314] R. H. Schonmann, *Finite size scaling behavior of a biased majority rule cellular automaton*, Phys. A **167** (1990), no. 3, 619–627 pp. MR1075567

[315] R. H. Schonmann, *On the behavior of some cellular automata related to bootstrap percolation*, Ann. Probab. **20** (1992), no. 1, 174–193 pp. MR1143417

[316] R. H. Schonmann, *Slow droplet-driven relaxation of stochastic Ising models in the vicinity of the phase coexistence region*, Comm. Math. Phys. **161** (1994), no. 1, 1–49 pp. MR1266068

[317] R. H. Schonmann and S. B. Shlosman, *Wulff droplets and the metastable relaxation of kinetic Ising models*, Comm. Math. Phys. **194** (1998), no. 2, 389–462 pp. MR1627669

[318] O. Schramm and J. E. Steif, *Quantitative noise sensitivity and exceptional times for percolation*, Ann. of Math. (2) **171** (2010), no. 2, 619–672 pp. MR2630053

[319] D. L. Schwartz, *Applications of duality to a class of Markov processes*, Ann. Probab. **5** (1977), no. 4, 522–532 pp. MR448631

[320] J. M Schwarz, A. J Liu, and L. Q Chayes, *The onset of jamming as the sudden emergence of an infinite k-core cluster*, EPL **73** (2006), no. 4, 560–566 pp.

[321] A. Shapira, *Bootstrap percolation and kinetically constrained models in homogeneous and random environments*, Ph.D. Thesis, Université Paris Diderot, Sorbonne Paris Cité, 2019.

[322] A. Shapira, *Metastable behavior of bootstrap percolation on Galton-Watson trees*, ALEA Lat. Am. J. Probab. Math. Stat. **16** (2019), no. 2, 1089–1104 pp. MR4030529

[323] A. Shapira, *Hydrodynamic limit of the Kob-Andersen model*, arXiv e-prints (2020), available at arXiv:2003.08495.

[324] A. Shapira, *Kinetically constrained models with random constraints*, Ann. Appl. Probab. **30** (2020), no. 2, 987–1006 pp. MR4108128

[325] A. Shapira, *A note on the spectral gap of the Fredrickson–Andersen one spin facilitated model*, J. Stat. Phys. **181** (2020), no. 6, 2346–2352 pp. MR4179809

[326] A. Shapira and E. Slivken, *Time scales of the Fredrickson-Andersen model on polluted $\mathbb{Z}^2$ and $\mathbb{Z}^3$*, arXiv e-prints (2019), available at arXiv:1906.09949.

[327] A. Simonis, *Metastability of the d-dimensional contact process*, J. Stat. Phys. **83** (1996), no. 5-6, 1225–1239 pp. MR1392426

[328] E. Slivken, *Bootstrap percolation on the Hamming torus with threshold 2*, arXiv e-prints (2014), available at arXiv:1407.2317.

[329] P. Sollich and M. R. Evans, *Glassy time-scale divergence and anomalous coarsening in a kinetically constrained spin chain*, Phys. Rev. Lett. **83** (1999), no. 16, 3238–3241 pp.

[330] P. Sollich and M. R. Evans, *Glassy dynamics in the asymmetrically constrained kinetic Ising chain*, Phys. Rev. E **68** (2003), no. 3, 031504, 16 pp.

[331] O. N. Stavskaya and I. I. Piatetski-Shapiro, *On homogeneous nets of spontaneously active elements*, Syst. Theory Res **20** (1971), 75–88 pp.

[332] A. Sudbury and P. Lloyd, *Quantum operators in classical probability theory. IV. Quasi-duality and thinnings of interacting particle systems*, Ann. Probab. **25** (1997), no. 1, 96–114 pp. MR1428501

[333] R. Sun and J. M. Swart, *The Brownian net*, Ann. Probab. **36** (2008), no. 3, 1153–1208 pp. MR2408586

[334] J. M. Swart, *Duality and intertwining of Markov chains*, 2013. Lecture notes for the ALEA in Europe School October 21-25 2013, Luminy (Marseille). Available at `http://staff.utia.cas.cz/swart`.

[335] L. Taggi, *Critical probabilities and convergence time of percolation probabilistic cellular automata*, J. Stat. Phys. **159** (2015), no. 4, 853–892 pp. MR3336983

[336] L. Taggi, *Convergence time of probabilistic cellular automata on the torus*, Probabilistic cellular automata, 2018, 179–196 pp. MR3793691

[337] E. Teomy and Y. Shokef, *Finite-density effects in the Fredrickson-Andersen and Kob-Andersen kinetically-constrained models*, J. Chem. Phys. **141** (2014), no. 6, 064110, 7 pp.

[338] E. Teomy and Y. Shokef, *Relation between structure of blocked clusters and relaxation dynamics in kinetically constrained models*, Phys. Rev. E **92** (2015), no. 3, 032133, 10 pp.

[339] P. Tetali, *Random walks and the effective resistance of networks*, J. Theoret. Probab. **4** (1991), no. 1, 101–109 pp. MR1088395

[340] W. Thomson and P. Tait, *Treatise on natural philosophy*, Oxford University Press, Oxford, 1867.

[341] C. Toninelli, *Kinetically constrained models and glassy dynamics*, Ph.D. Thesis, Università degli studi di Roma "La Sapienza", 2003.

[342] C. Toninelli, *Interacting particle systems with kinetic constraints*, In preparation.

[343] C. Toninelli and G. Biroli, *Dynamical arrest, tracer diffusion and kinetically constrained lattice gases*, J. Stat. Phys. **117** (2004), no. 1-2, 27–54 pp. MR2098557

[344] C. Toninelli and G. Biroli, *A new class of cellular automata with a discontinuous glass transition*, J. Stat. Phys. **130** (2008), no. 1, 83–112 pp. MR2375957

[345] C. Toninelli, G. Biroli, and D. S. Fisher, *Cooperative behavior of kinetically constrained lattice gas models of glassy dynamics*, J. Stat. Phys. **120** (2005), no. 1-2, 167–238 pp. MR2165529

[346] C. Toninelli, G. Biroli, and D. S. Fisher, *Jamming percolation and glass transitions in lattice models*, Phys. Rev. Lett. **96** (2006), no. 3, 035702, 4 pp.

[347] C. Toninelli, G. Biroli, and D. S. Fisher, *Toninelli, Biroli, and Fisher reply:*, Phys. Rev. Lett. **98** (2007), no. 12, 129602, 1 pp.

[348] A. Toom, *Cellular automata with errors: problems for students of probability*, Topics in contemporary probability and its applications, 1995, 117–157 pp. MR1410537

[349] A. L. Toom, N. B. Vasilyev, O. N. Stavskaya, L. G. Mityushin, G. L. Kurdyumov, and S. A. Pirogov, *Discrete local Markov systems*, Stochastic cellular systems: ergodicity, memory, morphogenesis, 1990, 1–182 pp.

[350] A. Toom, *Contornos, conjuntos convexos e autômatos celulares*, Publicações matemáticas do IMPA. [IMPA mathematical publications], Instituto de Matemática Pura e Aplicada (IMPA), Rio de Janeiro, 2001. 23-o Colóquio Brasileiro de Matemática. [23rd Brazilian Mathematics Colloquium]. MR1849134

[351] G. L. Torrisi, M. Garetto, and E. Leonardi, *A large deviation approach to supercritical bootstrap percolation on the random graph $G_{n,p}$*, Stochastic Process. Appl. **129** (2019), no. 6, 1873–1902 pp. MR3958417

[352] G. L. Torrisi, M. Garetto, and E. Leonardi, *Bootstrap percolation on the stochastic block model with k communities*, arXiv e-prints (2020), available at arXiv:1812.09107.

[353] R. M. Turner, R. L. Jack, and J. P. Garrahan, *Overlap and activity glass transitions in plaquette spin models with hierarchical dynamics*, Phys. Rev. E **92** (2015), no. 2, 022115, 13 pp. MR3531148

[354] A. J. Uzzell, *An improved upper bound for bootstrap percolation in all dimensions*, Combin. Probab. Comput. **28** (2019), no. 6, 936–960 pp. MR4015663

[355] P. Valiant, *Linear bounds on the North-East model and higher-dimensional analogs*, Adv. Appl. Math. **33** (2004), no. 1, 40–50 pp. MR2064356

[356] J. van den Berg and U. Fiebig, *On a combinatorial conjecture concerning disjoint occurrences of events*, Ann. Probab. **15** (1987), no. 1, 354–374 pp. MR877608

[357] J. van den Berg and H. Kesten, *Inequalities with applications to percolation and reliability*, J. Appl. Probab. **22** (1985), no. 3, 556–569 pp. MR799280

[358] A. C. D. van Enter, *Proof of Straley's argument for bootstrap percolation*, J. Stat. Phys. **48** (1987), no. 3-4, 943–945 pp. MR914911

[359] A. C. D. van Enter, *Scaling and inverse scaling in anisotropic bootstrap percolation*, Probabilistic cellular automata, 2018, 69–77 pp. MR3793683

[360] A. C. D. van Enter, J. Adler, and J. A. M. S. Duarte, *Finite-size effects for some bootstrap percolation models*, J. Stat. Phys. **60** (1990), no. 3-4, 323–332 pp. MR1069635

[361] A. C. D. van Enter and T. Hulshof, *Finite-size effects for anisotropic bootstrap percolation: logarithmic corrections*, J. Stat. Phys. **128** (2007), no. 6, 1383–1389 pp. MR2352015

[362] S. Whitelam, L. Berthier, and J. P. Garrahan, *Renormalization group study of a kinetically constrained model for strong glasses*, Phys. Rev. E **71** (2005), no. 2, 026128, 14 pp.

[363] A. Whittemore, *Bootstrap percolation on random geometric graphs*, Ph.D. Thesis, 2021. MR4315377

[364] J. C. Wierman, *Duality for directed site percolation*, Particle systems, random media and large deviations, 1985, 363–380 pp. MR814726

[365] T. Williams and R. Bjerknes, *Stochastic model for abnormal clone spread through epithelial basal layer*, Nature **236** (1972), no. 5340, 19–21 pp.

[366] S. Wolfram, *Universality and complexity in cellular automata*, Phys. D **10** (1984), 1–35 pp. MR762650

[367] Q. Yang, H. Liu, and F. Gao, *Logarithmic Sobolev inequalities for two-sided birth-death processes*, Wuhan Univ. J. Nat. Sci. **13** (2008), no. 2, 133–136 pp. MR2491203

[368] A. N. Zehmakan, *On the spread of information through graphs*, Ph.D. Thesis, ETH Zurich, 2019.

RÉSUMÉ

On étudie deux classes de modèles étroitement liées de physique statistique sur le réseau carré bidimensionnel – les modèles cinétiquement contraints et la percolation bootstrap. Les premiers sont apparus pour modéliser la dynamique des liquides surfondus près de leur transition vitreuse, tandis que la percolation bootstrap modélise de nombreux cadres tels que certains aimants ou encore des phénomènes sociaux. Nous considérons les modèles cinétiquement contraints et la percolation bootstrap d'un point de vue rigoureux probabiliste. On s'intéresse à leur comportement lorsque leur paramètre tend vers sa valeur critique (possiblement dégénérée). Plus concrètement, nous étudions le taux de divergence de certains temps caractéristiques tels que le temps d'infection d'un site fixé et le temps de relaxation.
Parmi les résultats les plus conséquents de la thèse est la détermination des classes d'universalité de modèles cinétiquement contraints ainsi que leurs échelles de temps caractéristiques à l'équilibre en basse température. C'est-à-dire, on établit une partition de tous les modèles possibles en groupes à comportement similaire et fournit une recette pour déterminer ce comportement à partir de la définition du modèle. Des contributions sont apportées à tout le spectre de classes d'universalité de modèles cinétiquement contraints, mais dans certains cas aussi à la percolation bootstrap plus simple et mieux comprise. En supplément des résultats universels, nous donnons des asymptotiques exactes à la fois en percolation bootstrap et en modèle cinétiquement contraint pour le modèle le plus classique à deux voisins. De plus, nous marquons des progrès sur le modèle cinétiquement contraint à un voisin appelé modèle de Fredrickson–Andersen 1-spin facilité.

La thèse est constituée de trois parties principales, basées sur des techniques provenant de domaines différents. La première relève de la dynamique de systèmes de particules en interaction. La deuxième emploie des arguments de combinatoire. La troisième et dernière partie prend un point de vue de percolation.

MOTS CLÉS

Percolation bootstrap, modèles cinétiquement contraints, percolation orientée, systèmes de particules en interaction, dynamique de Glauber, universalité, classification, seuil aigu, inégalité de Poincaré, trou spectral

ABSTRACT

We study two tightly related classes of statistical mechanics models on the two-dimensional square lattice—kinetically constrained models and bootstrap percolation. The former arose as models of the dynamics of supercooled liquids close to the glass transition, while the latter are used to model a number of settings including magnets and social phenomena. We consider both kinetically constrained models and bootstrap percolation from a rigorous probabilistic perspective. We are interested in their behaviour as their parameter approaches its (possibly degenerate) critical value. More specifically, we investigate the rate of divergence of certain characteristic time scales, such as the infection time of a fixed site and the relaxation time.
Among the highlights of the thesis is determining the universality classes of kinetically constrained models together with their characteristic equilibrium time scales at low temperature. That is, we establish a partition of all possible models into groups with similar behaviour and provide a recipe for determining the behaviour from the definition of the model. Contributions are made to the full spectrum of universality classes of kinetically constrained models, but in some cases also to the simpler and better understood bootstrap percolation. In addition to universal results, we provide sharp asymptotics in both bootstrap percolation and kinetically constrained models for the most classical, 2-neighbour, model, as well as advances on the 1-neighbour kinetically constrained model known as the Fredrickson–Andersen 1-spin facilitated model.

The thesis consists of three main parts based on techniques from different domains. The first one relates to dynamics of interacting particle systems. The second one relies on combinatorial arguments. The third and final part takes a percolation viewpoint.

KEYWORDS

Bootstrap percolation, kinetically constrained models, oriented percolation, interacting particle systems, Glauber dynamics, universality, classification, sharp threshold, Poincaré inequality, spectral gap