Last passage percolation on the complete graph and particle systems

Sanjay Ramassamy
Institut de Physique Théorique

Joint works with Sergey Foss (Sobolev Institute / Heriot-Watt), Takis Konstantopoulos (Liverpool), Bastien Mallein (Paris Nord / ENS) and Arvind Singh (Paris-Saclay)

Séminaire Matrices Et Graphes Aléatoires
Institut Henri Poincaré, June 17 2022

- Barak-Erdős graphs (BEGs) are the directed acyclic version of Erdős-Rényi random graphs.

- They arise as a special case of last passage percolation on the complete graph.

- The geometric infinite-bin model (IBM) is an interacting particle system which can be coupled to BEGs.

- This coupling relates the speed of the front of the geometric IBM to the length of the longest path of BEGs.

Outline :

1. Barak-Erdős graphs

2. The infinite-bin model (IBM)

3. Properties of Barak-Erdős graphs via the IBM

4. Extensions

1 Barak-Erdős graphs

Construction of Barak-Erdős graphs

- Fix $n \geq 1$ integer and $0 \leq p \leq 1$.
- Vertex set is $\{1, 2, \ldots, n\}$.
- For each pair $i < j$, add an edge directed from i to j with probability p, independently for each pair (i, j).

Construction of Barak-Erdős graphs

- Fix $n \geq 1$ integer and $0 \leq p \leq 1$.
- Vertex set is $\{1, 2, \ldots, n\}$.
- For each pair $i < j$, add an edge directed from i to j with probability p, independently for each pair (i, j).

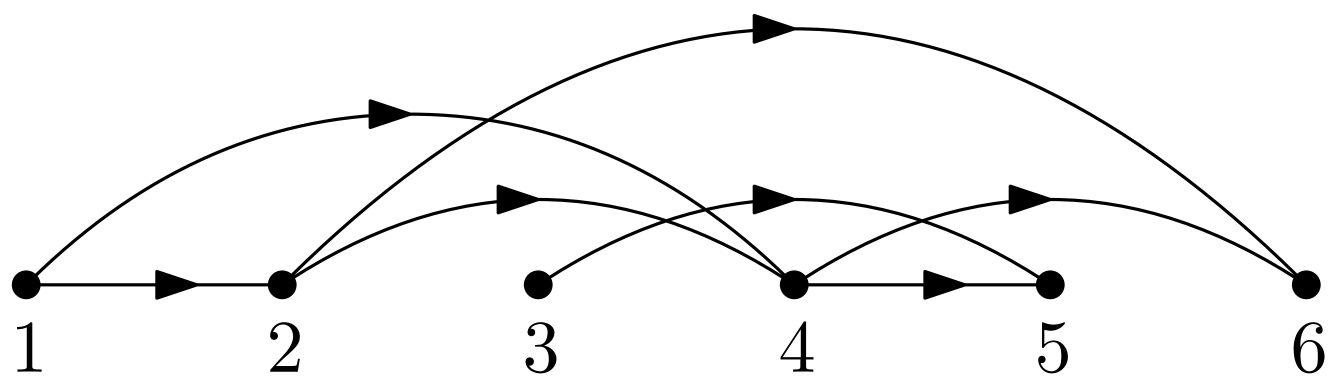

- Introduced by Barak and Erdős in 1984.
- The most studied feature is the length of the longest path $L_n(p)$.

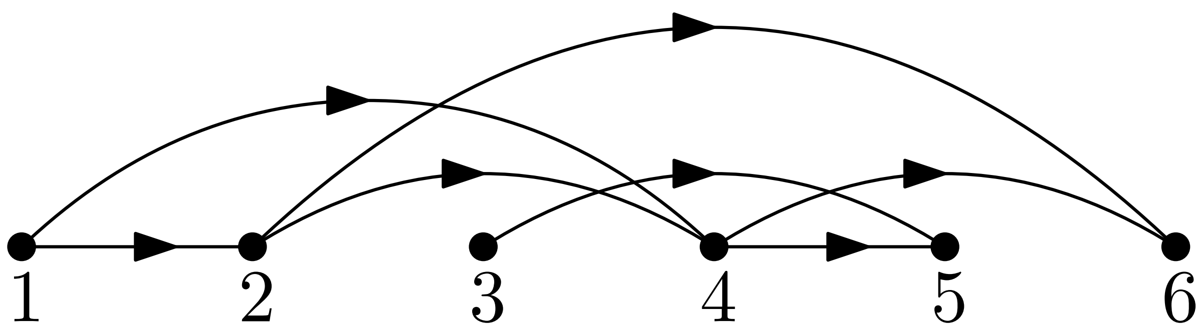

L=3

$1 \to 2 \to 4 \to 5$
$1 \to 2 \to 4 \to 6$

- Applications to performance evalution of computer systems (Gelenbe-Nelson-Philips-Tantawi, Isopi-Newman), mathematical ecology (Cohen-Newman) and queuing systems (Foss-Konstantopoulos).

- The length of the longest path grows linearly in the number of vertices:

$$\frac{1}{n} L_n(p) \xrightarrow[n\to\infty]{} C(p) \text{ in probability (Newman '92)}$$

- The growth rate C is a function of p:

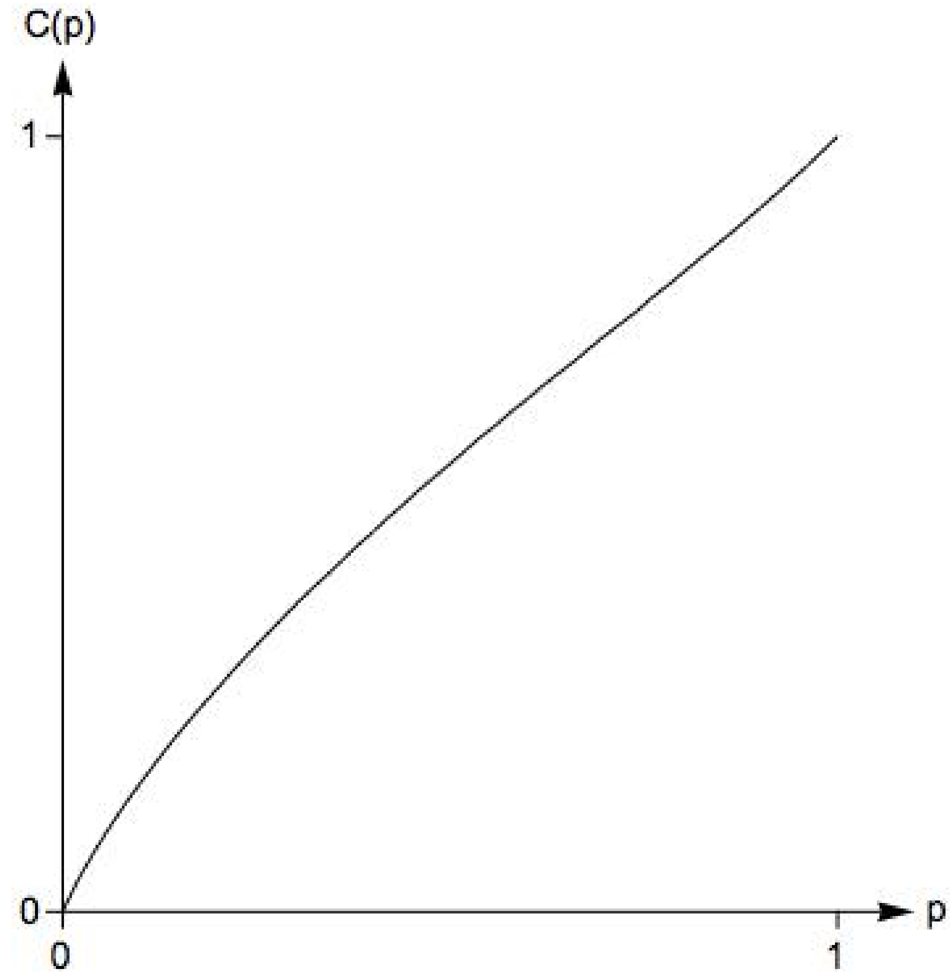

Properties of $C(p)$

- $C(p)$ is continuous and $C'(0) = e$ (Newman '92).

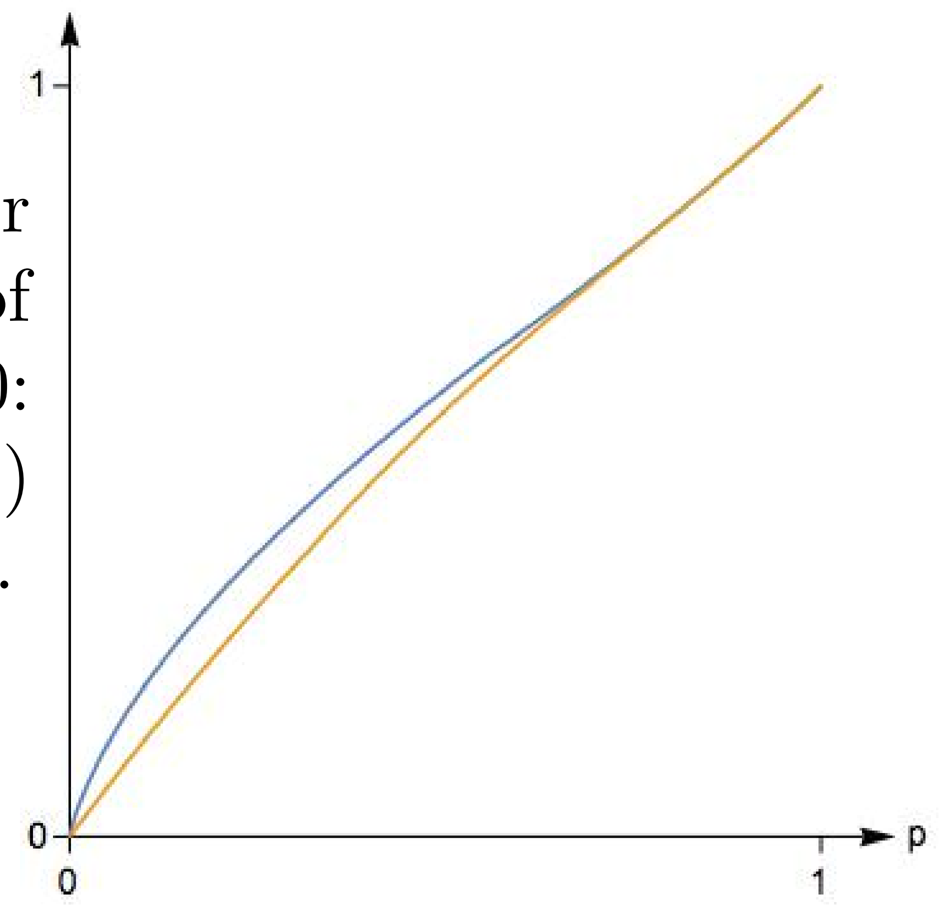

- Upper and lower bound for $C(p)$, yielding expansion of $C(1-q)$ for q tending to 0: $1-q+q^2-3q^3+7q^4+O(q^5)$ (Foss-Konstantopoulos '03).

Results by Mallein-R. (2019-2021)

- For $p > 0$, $C(p)$ is analytic and can be obtained as the sum of a series.

- The power series expansion of $C(p)$ centered at 1 has integer coefficients.

- $C(p)$ has no second derivative at $p = 0$:

$$C(p) = ep - \frac{e\pi^2 p}{2(\log p)^2} + o\left(\frac{p}{(\log p)^2}\right) \text{ when } p \to 0.$$

2 The infinite-bin model (IBM)

Configurations

Infinitely many balls placed inside bins indexed by $\mathbb{Z}$, such that the set of indices of nonempty bins has a maximal element, the *front*.

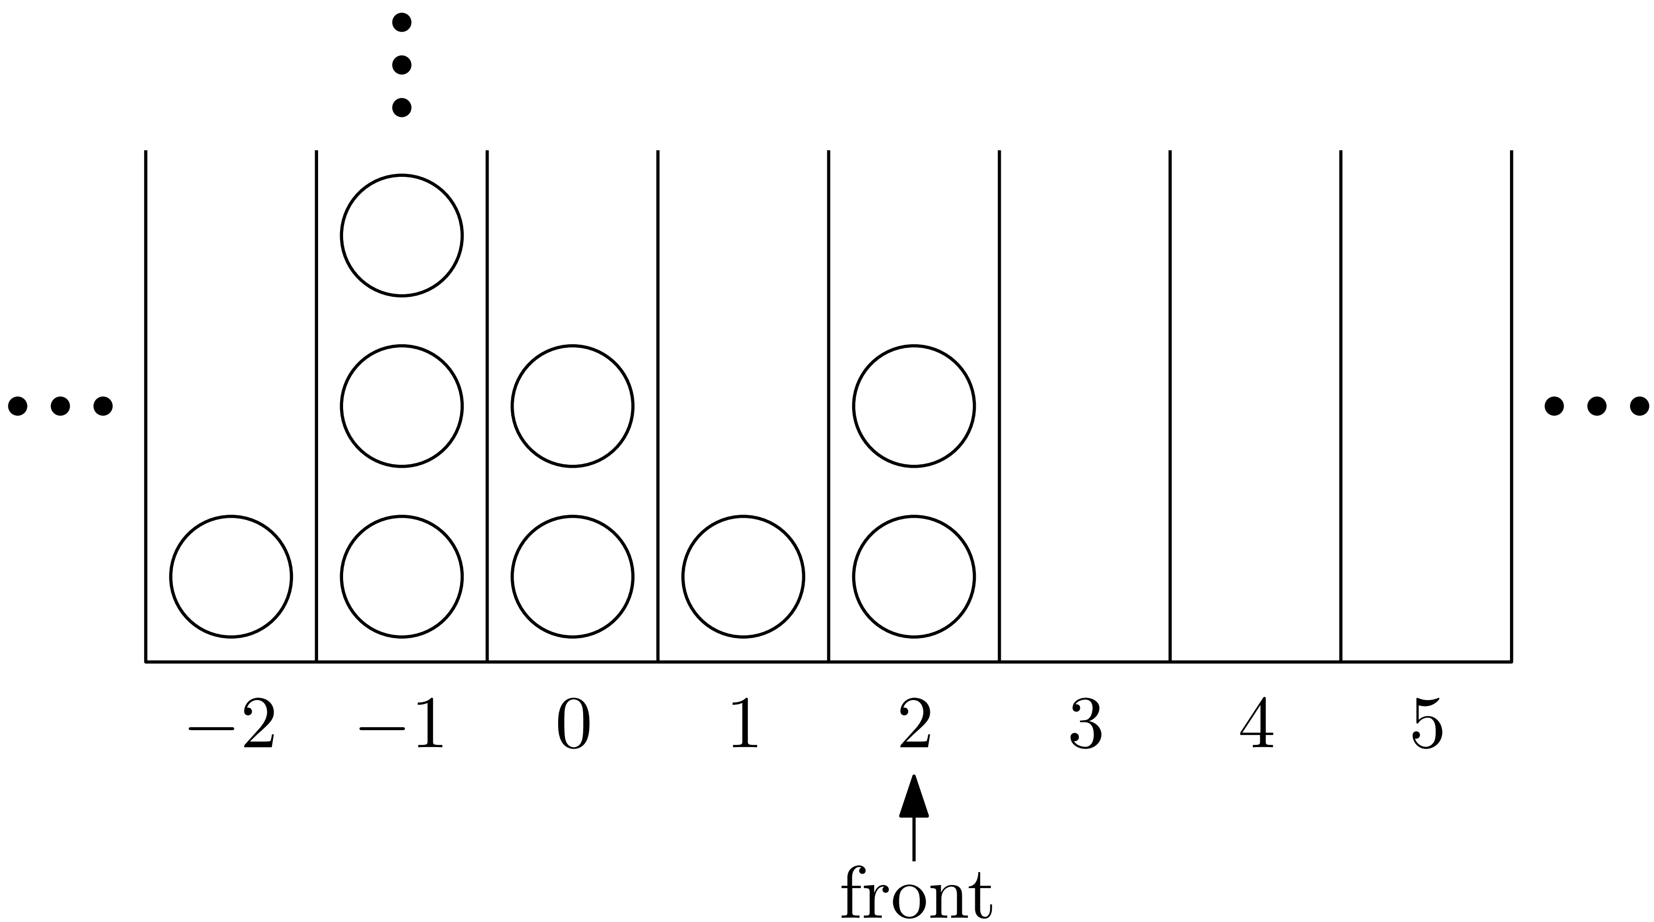

Move of type k

Add a ball in the bin immediately to the right of the bin containing the k-th ball, where balls are counted from right to left.

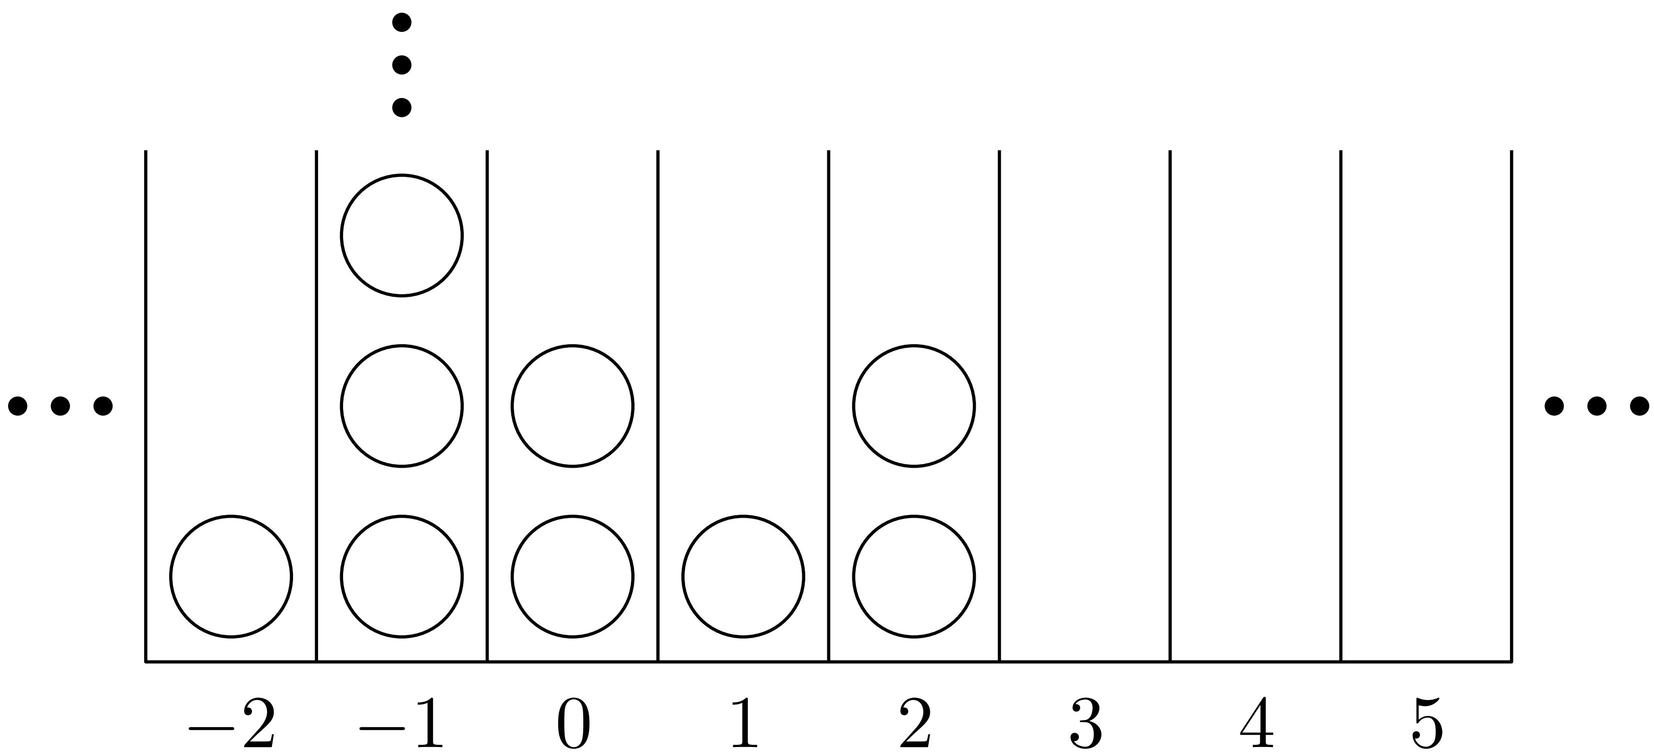

Move of type k

Add a ball in the bin immediately to the right of the bin containing the k-th ball, where balls are counted from right to left.

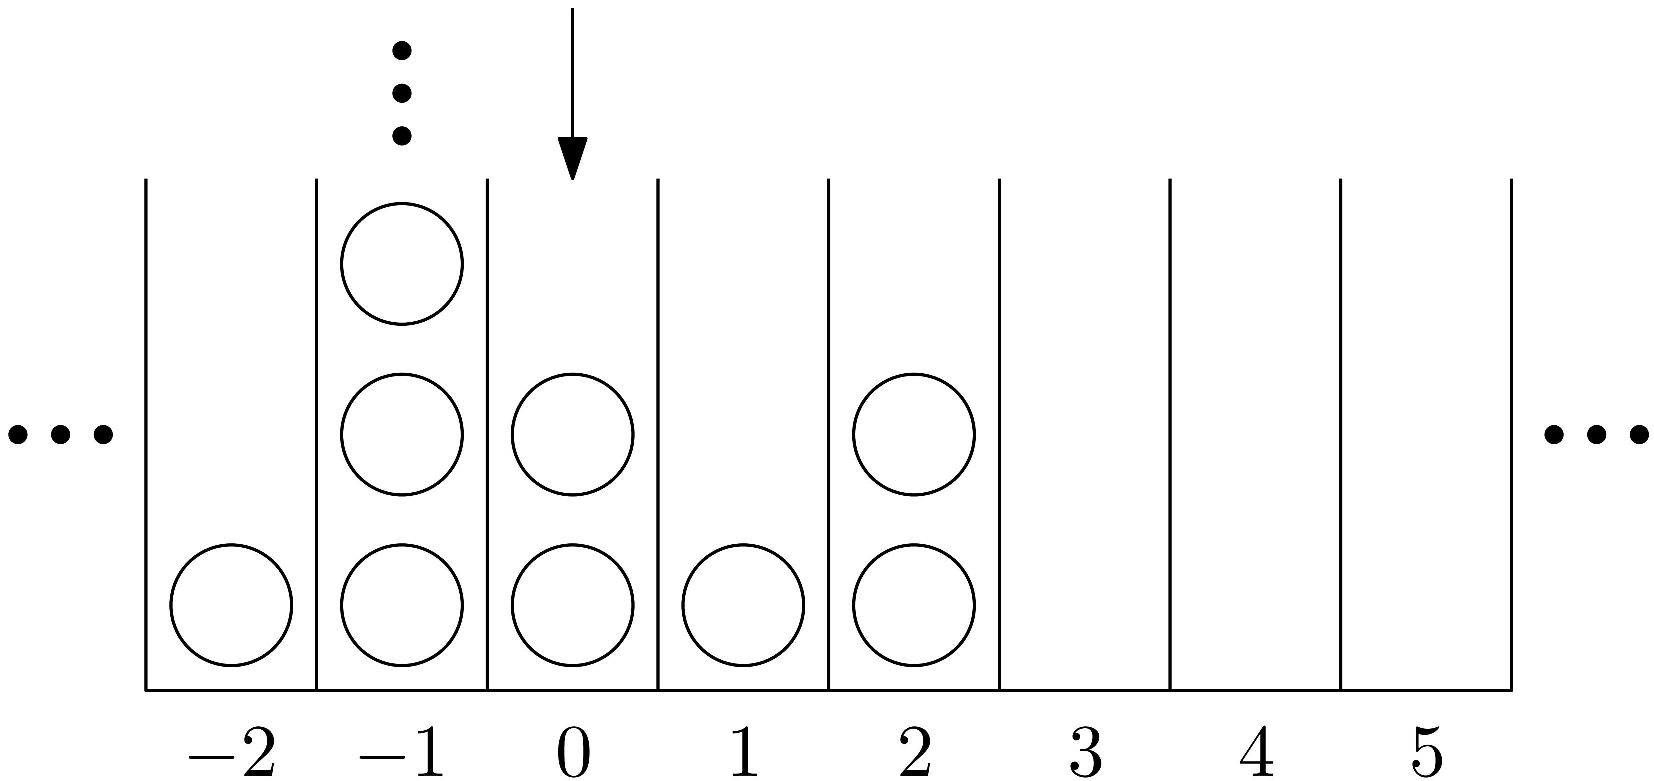

Move of type 5

Move of type k

Add a ball in the bin immediately to the right of the bin containing the k-th ball, where balls are counted from right to left.

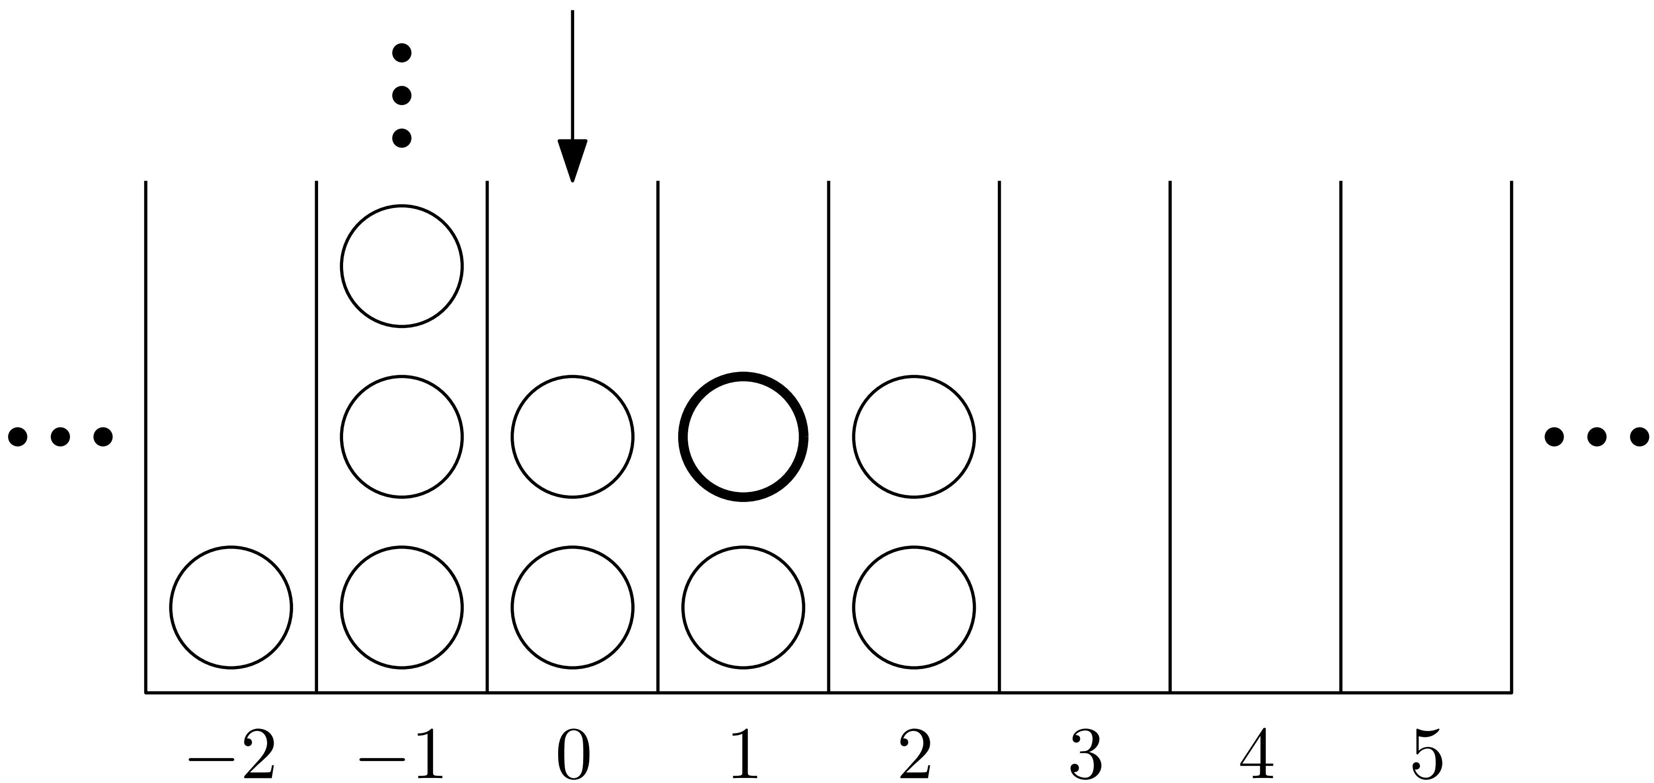

Move of type 5

ϕ_k : move of type k

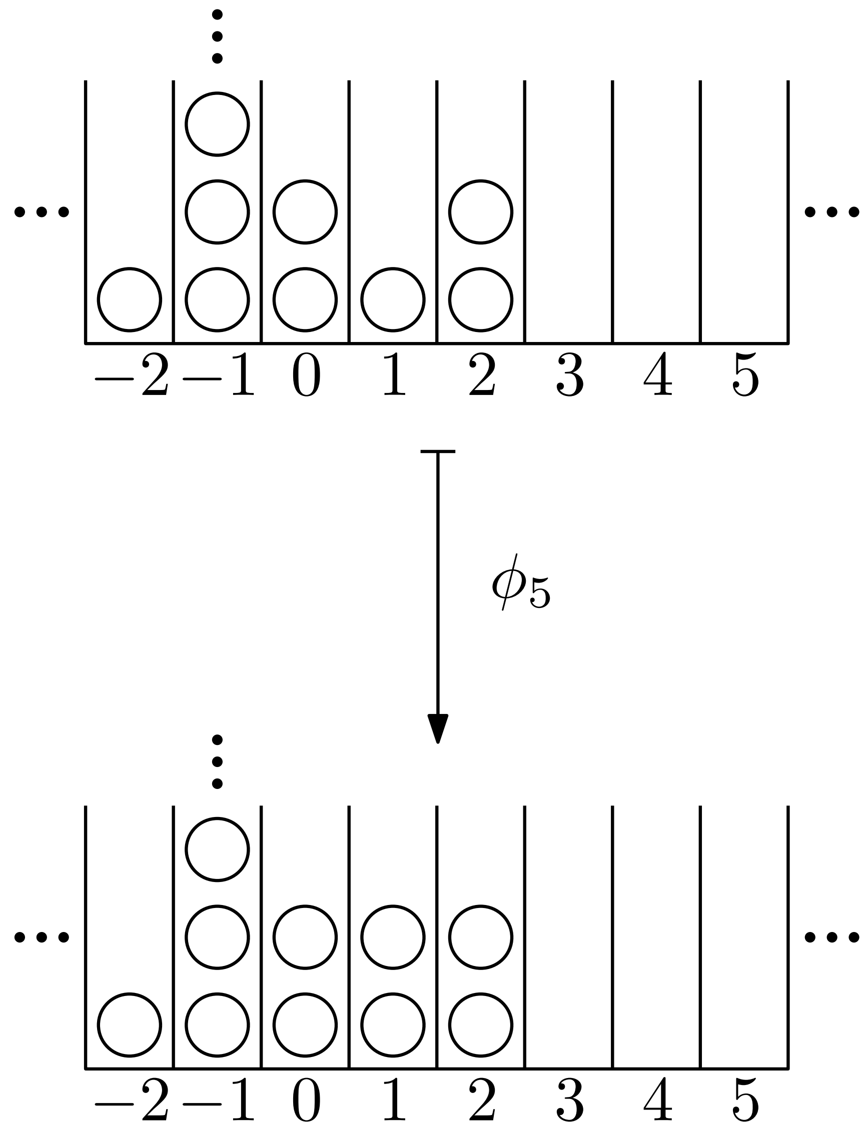

ϕ_k : move of type k

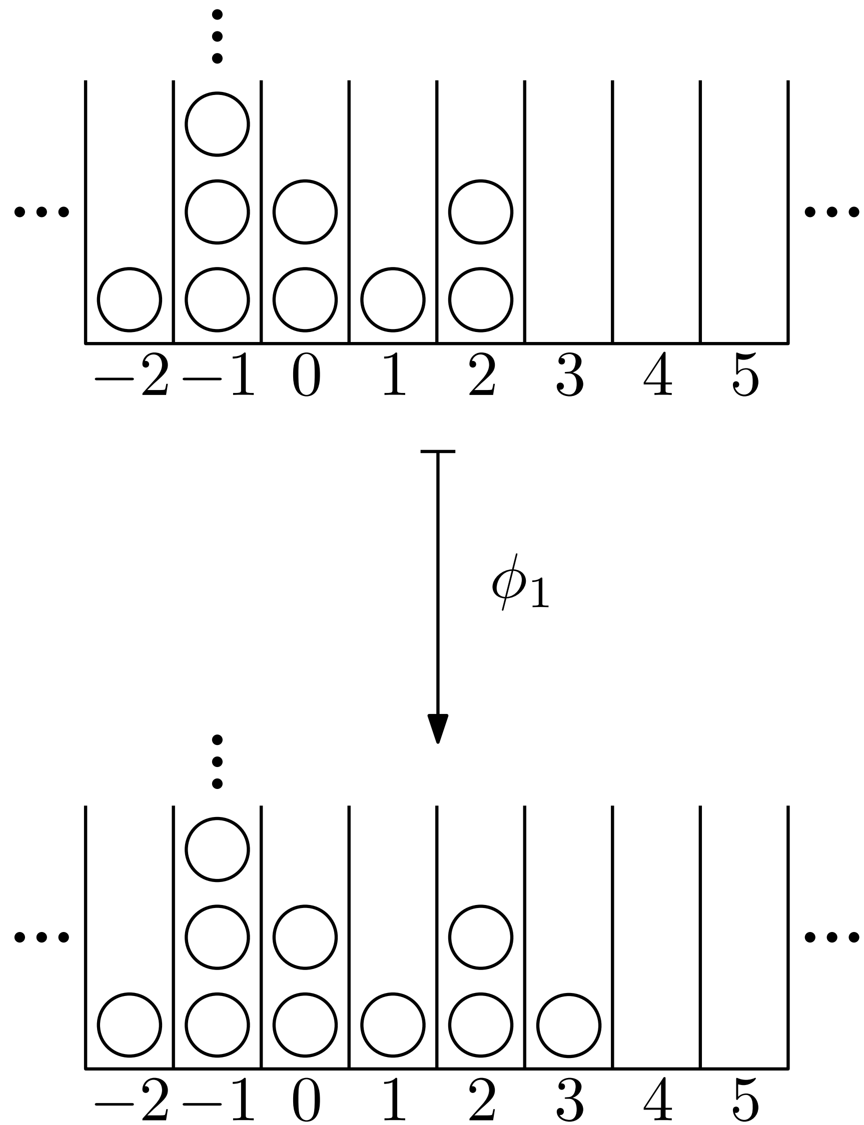

Markovian evolution

- Fix an initial configuration X_0 and a probability distribution μ on $\mathbb{N} = \{1, 2, 3, \ldots\}$.

- The infinite-bin model with move distribution μ (IBM(μ) for short) and initial configuration X_0 is the Markov chain $(X_n)_{n \geq 0}$ satisfying

$$X_n = \phi_{\xi_n}(X_{n-1}),$$

where the sequence $(\xi_n)_{n \geq 1}$ is i.i.d. distributed like μ.

- Introduced by Foss and Konstantopoulos in 2003 to study the longest paths of Barak-Erdős graphs.

- Special case when μ is the uniform measure on $\{1, \ldots, n\}$ already appeared in Aldous-Pitman '83.

Speed of the front

Consider the IBM(μ) with initial configuration X_0. Denote by F_n the position of the front at time n.

Theorem (Foss-Konstantopoulos, Mallein-R.). *There exists* $v_\mu \in (0, 1]$ *such that*

$$\lim_{n\to\infty} \frac{F_n - F_0}{n} = v_\mu \ a.s.$$

v_μ *(independent of* X_0*) is called the* speed *of the IBM(*μ*).*

Coupling with Barak-Erdős (Foss-Konstantopoulos ’03)

- μ_p : geometric distribution on $\{1, 2, \ldots\}$ with parameter p, *i.e.* $\mu_p(k) = p(1-p)^{k-1}$ for $k \geq 1$.

- The speed of the IBM(μ_p) equals the growth rate of the length of the longest path in Barak-Erdős graphs with edge probability p :

$$v_{\mu_p} = C(p).$$

- Grow Barak-Erdős graph one vertex after the other.

- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

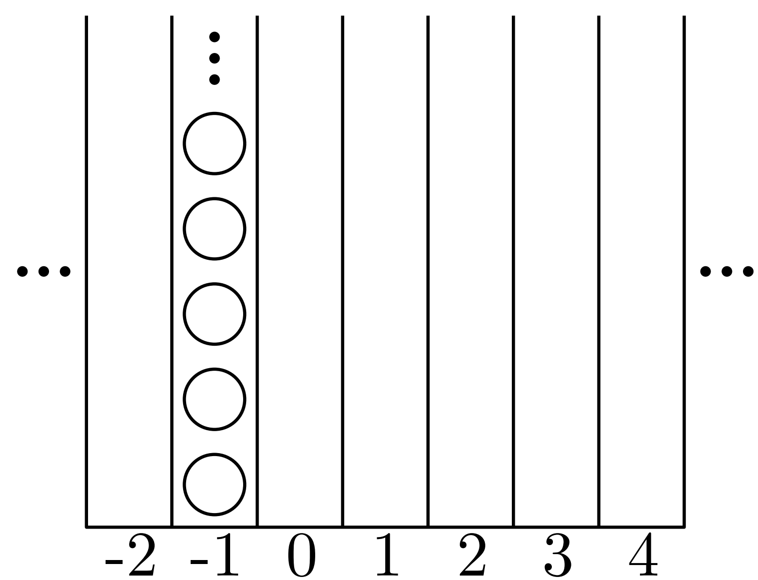

- Grow Barak-Erdős graph one vertex after the other.
- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

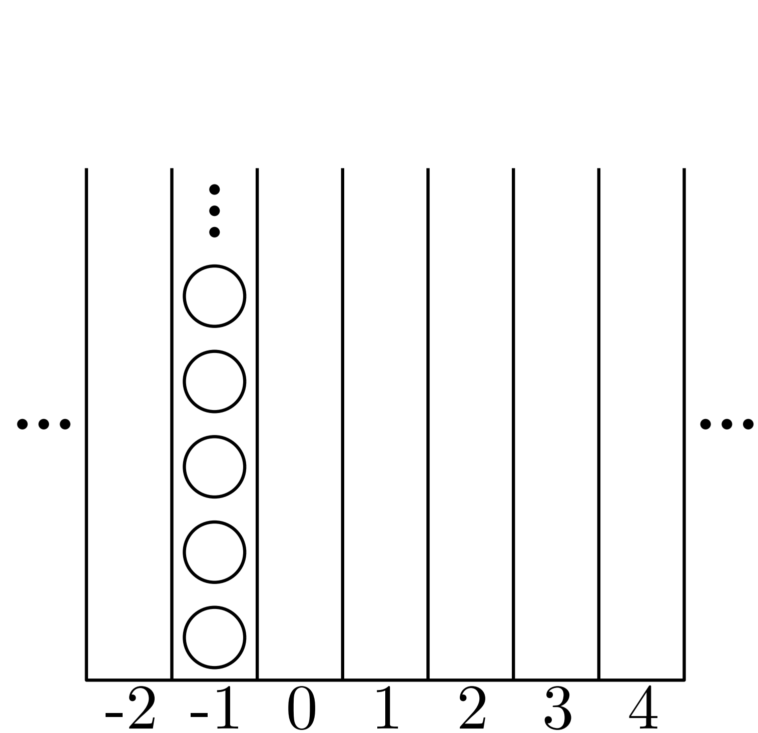

- Grow Barak-Erdős graph one vertex after the other.
- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

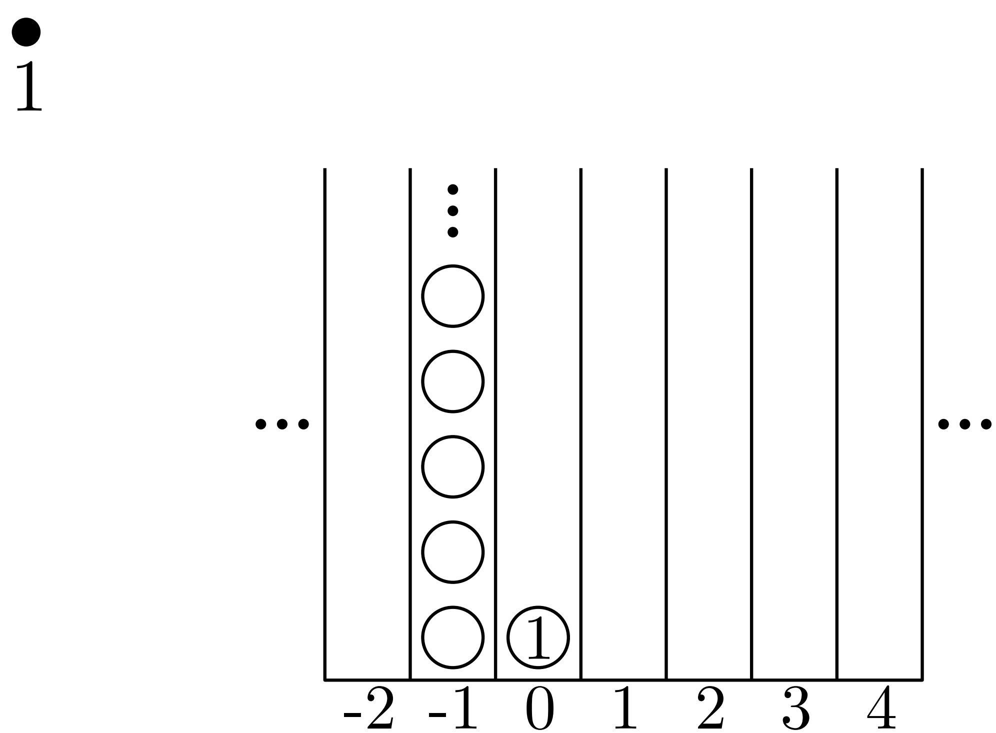

- Grow Barak-Erdős graph one vertex after the other.

- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

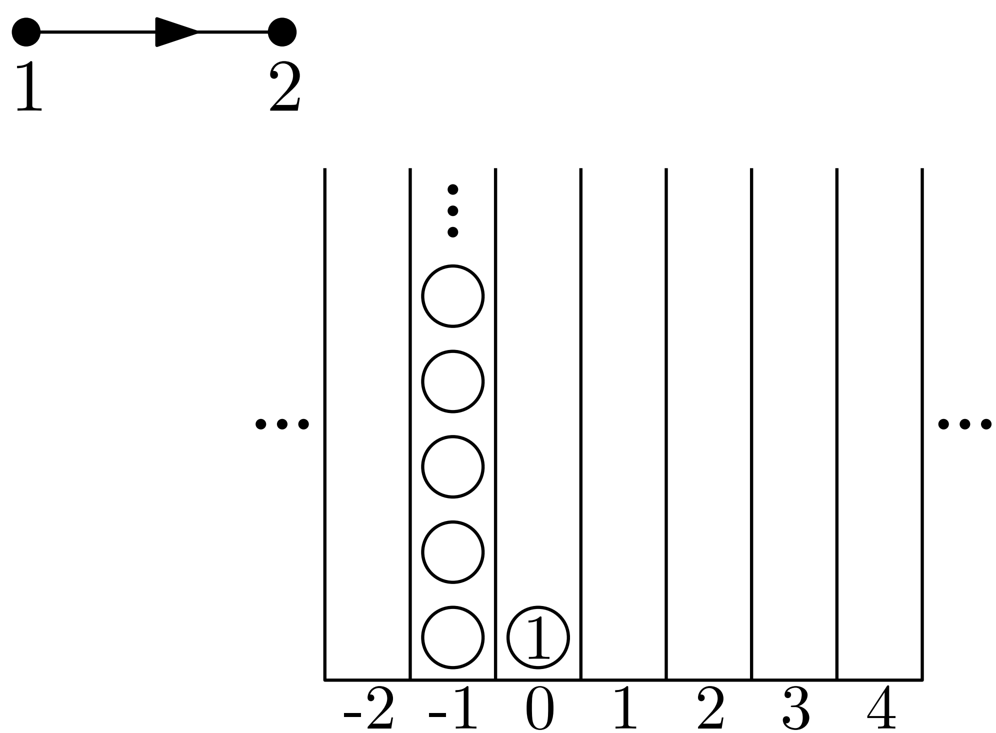

- Grow Barak-Erdős graph one vertex after the other.
- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

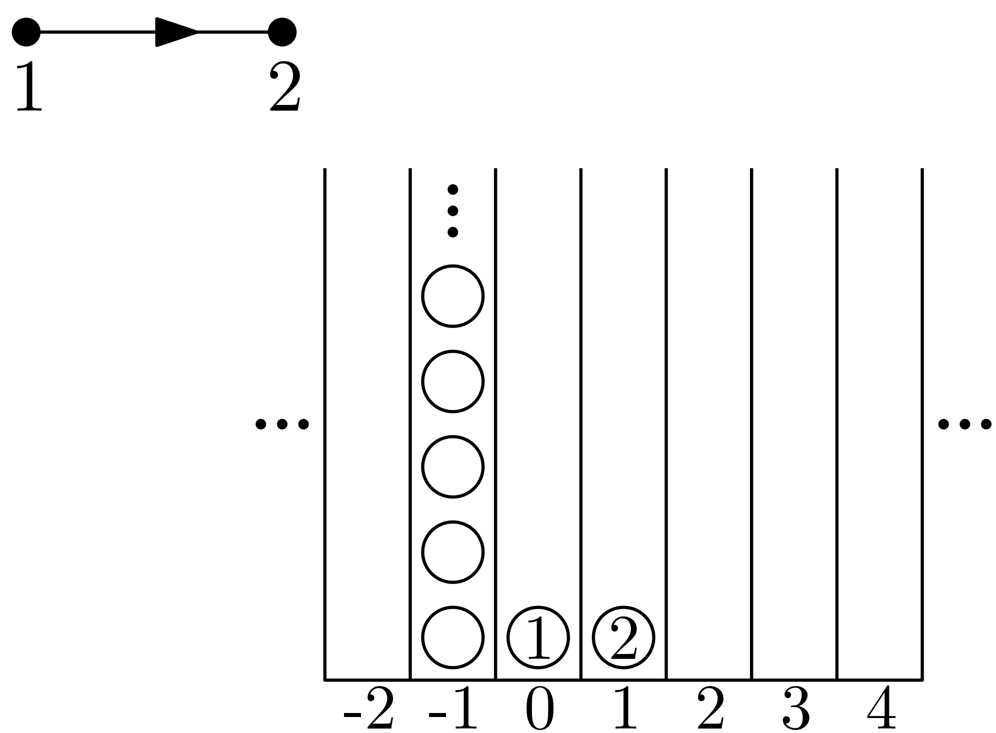

- Grow Barak-Erdős graph one vertex after the other.
- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

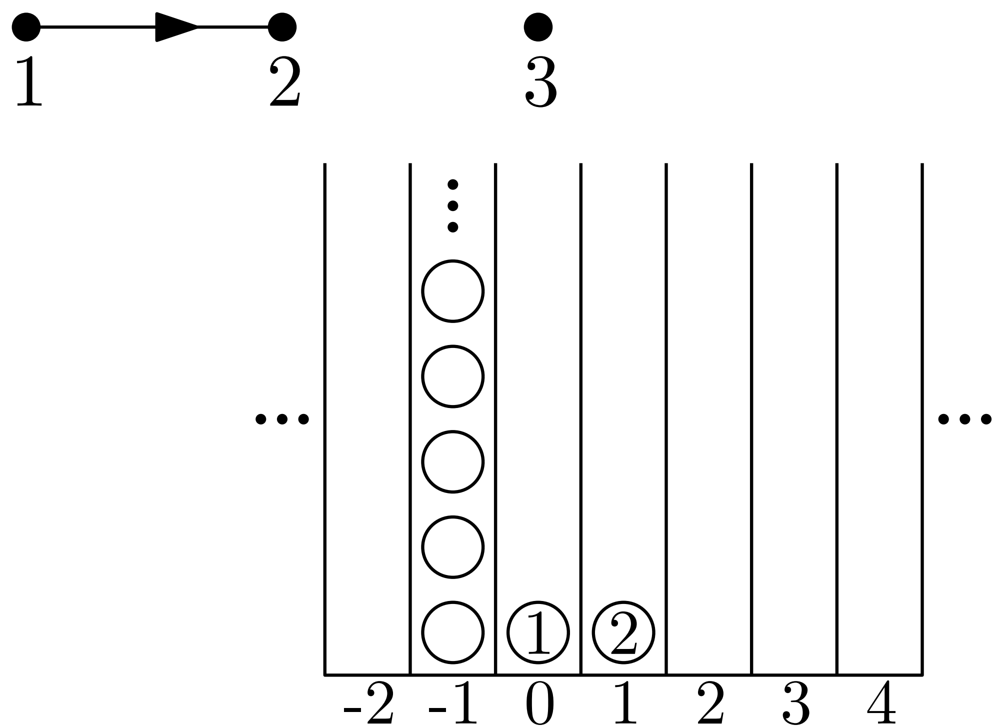

- Grow Barak-Erdős graph one vertex after the other.

- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

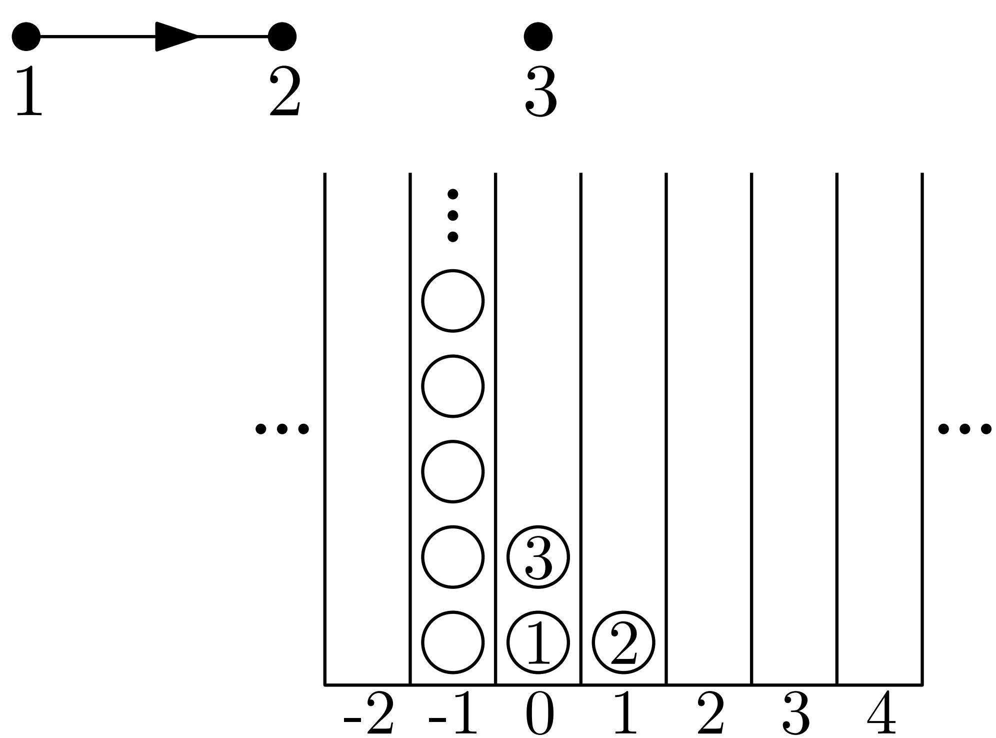

- Grow Barak-Erdős graph one vertex after the other.
- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

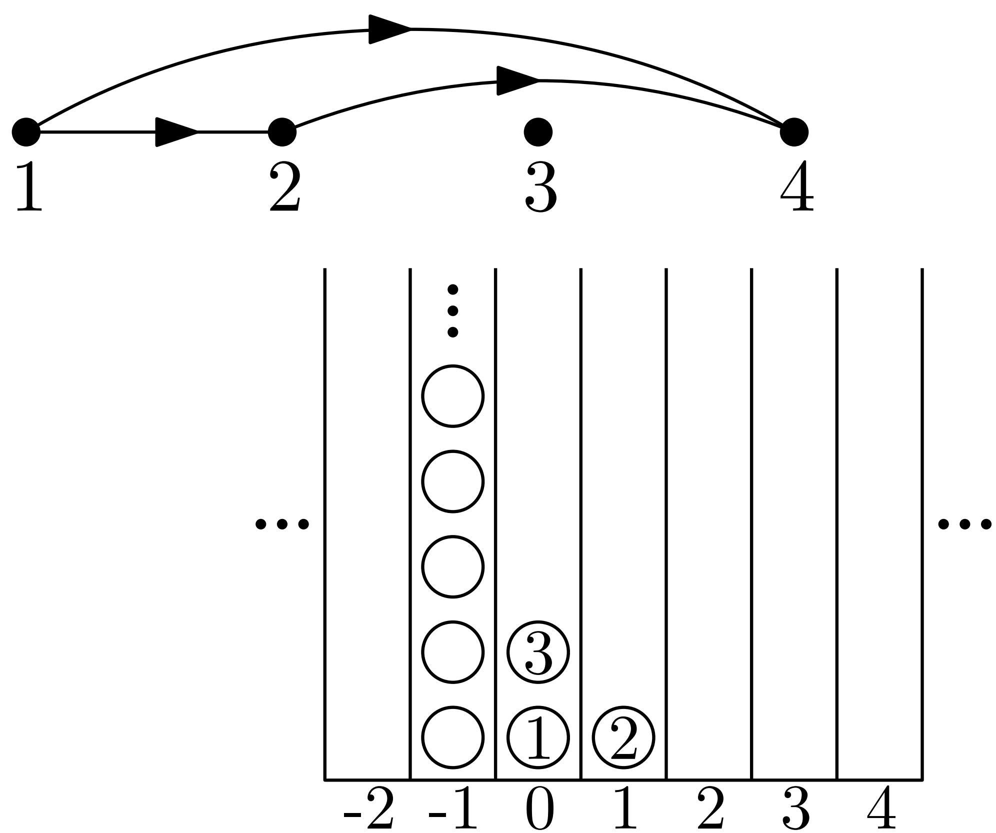

- Grow Barak-Erdős graph one vertex after the other.

- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

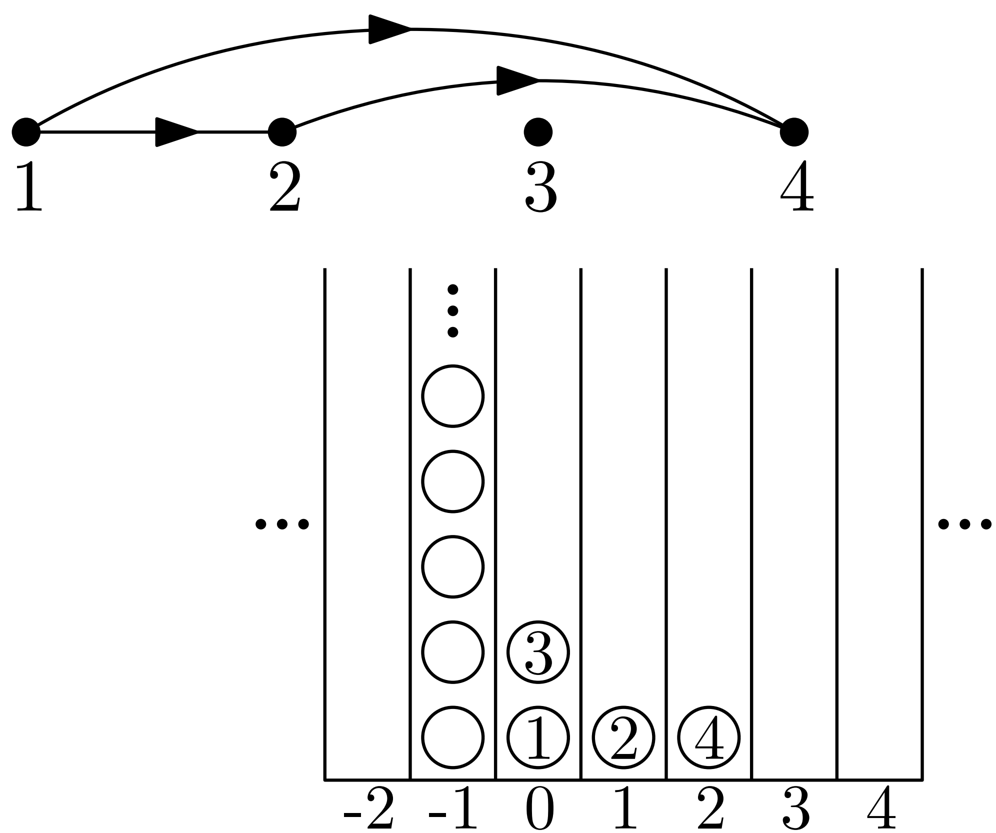

- Grow Barak-Erdős graph one vertex after the other.

- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

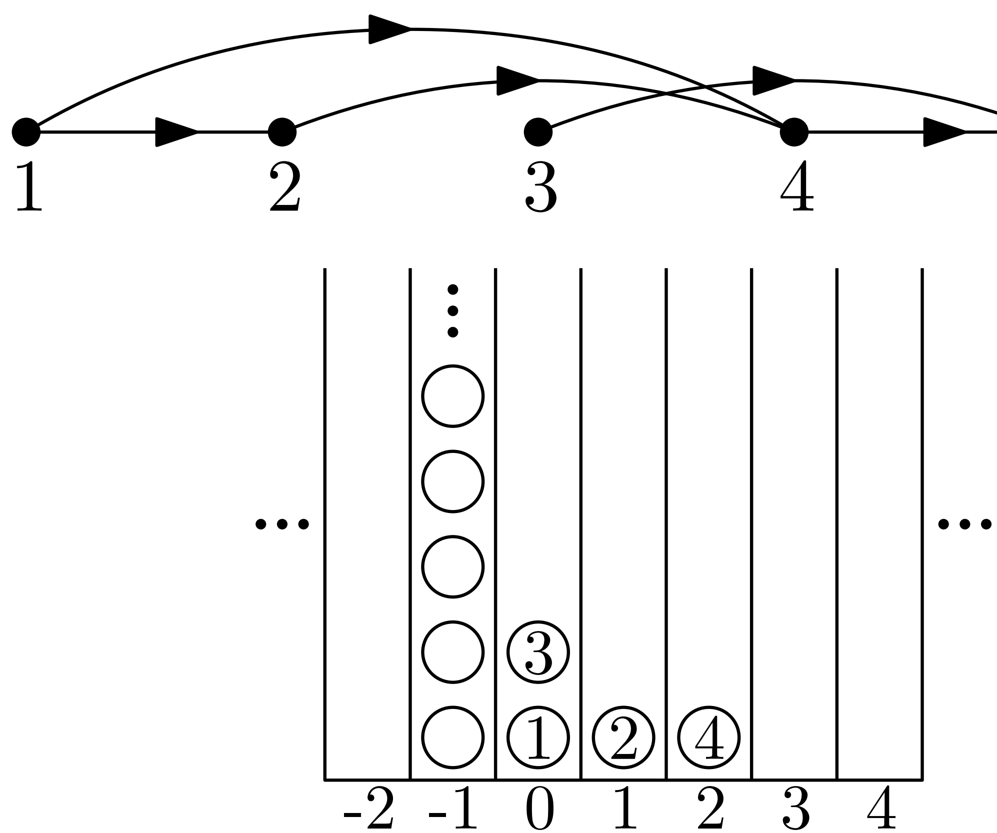

- Grow Barak-Erdős graph one vertex after the other.

- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

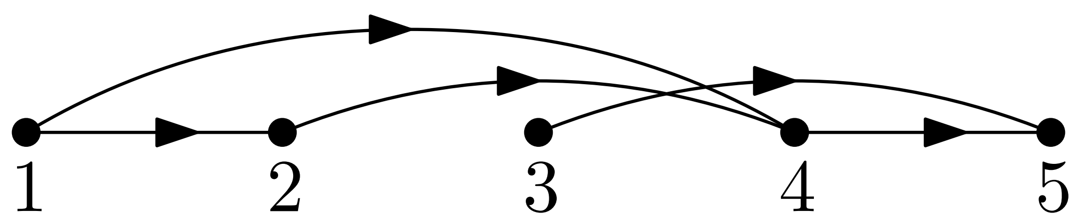

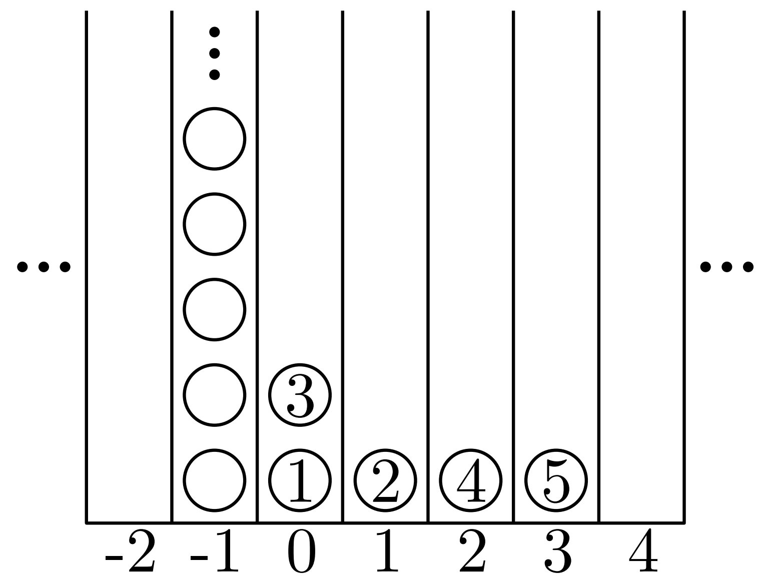

- Grow Barak-Erdős graph one vertex after the other.

- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

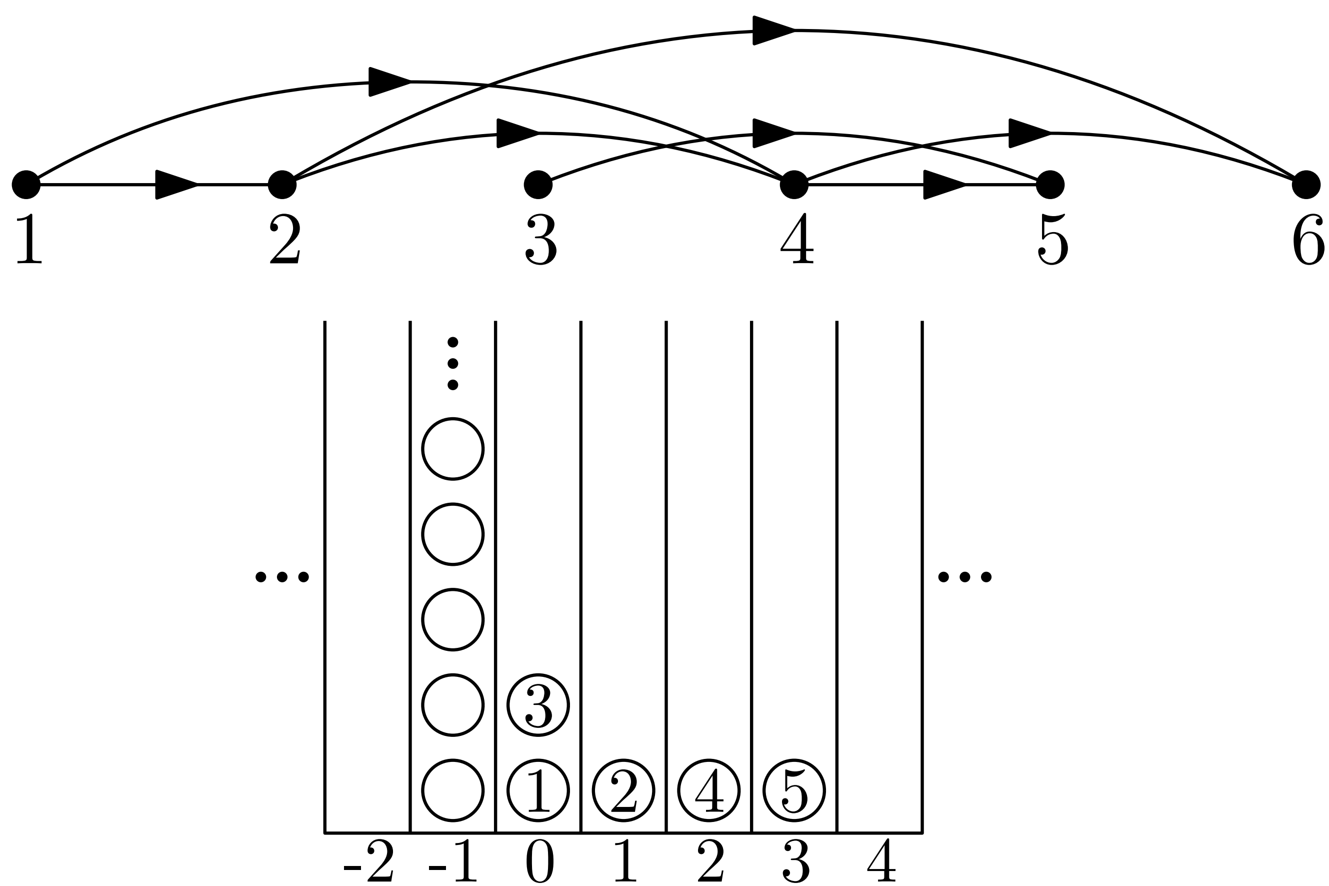

- Grow Barak-Erdős graph one vertex after the other.

- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

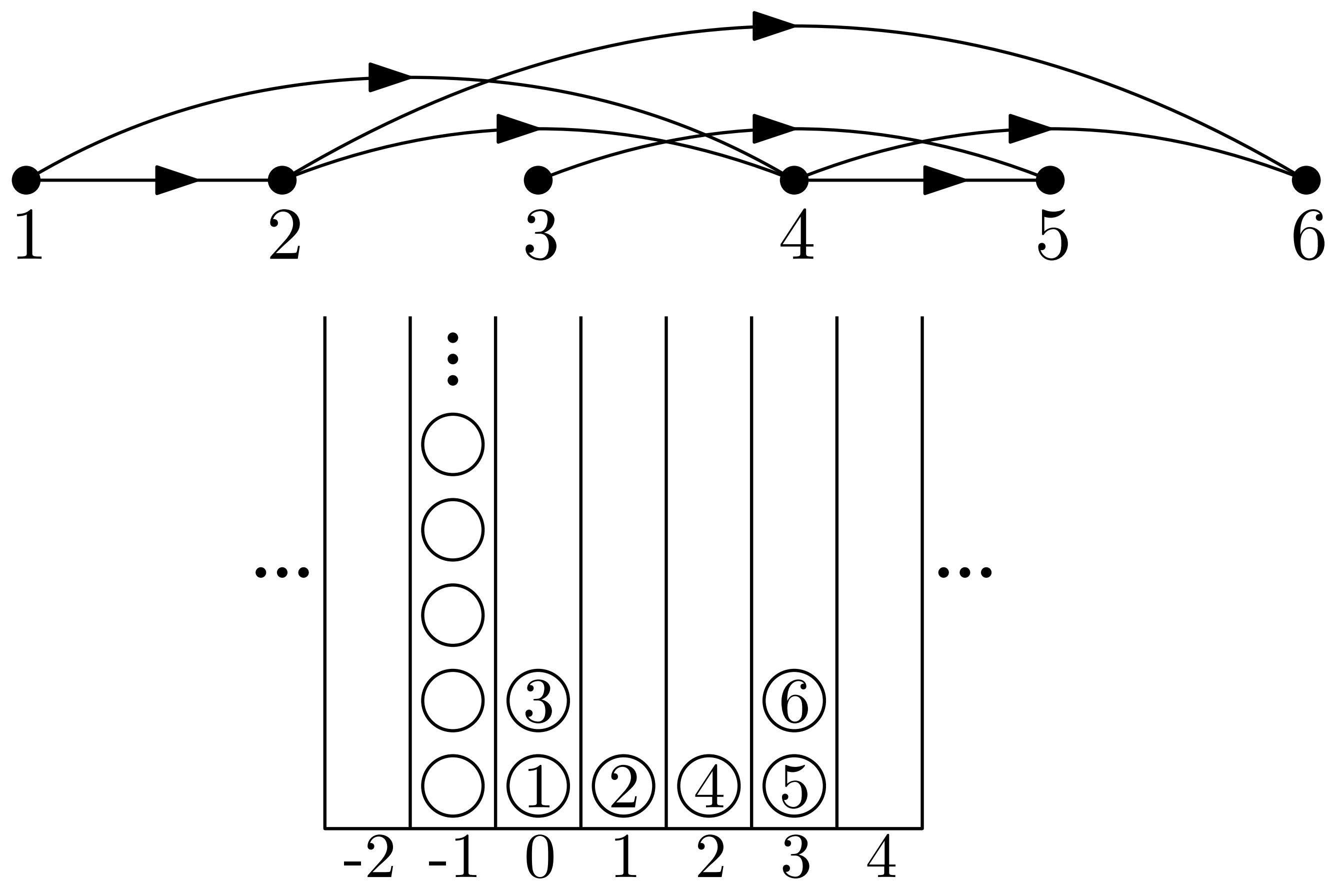

- Grow Barak-Erdős graph one vertex after the other.

- For each vertex n, call l_n the length of the longest path ending at n. Place a ball with label n in the bin indexed by l_n.

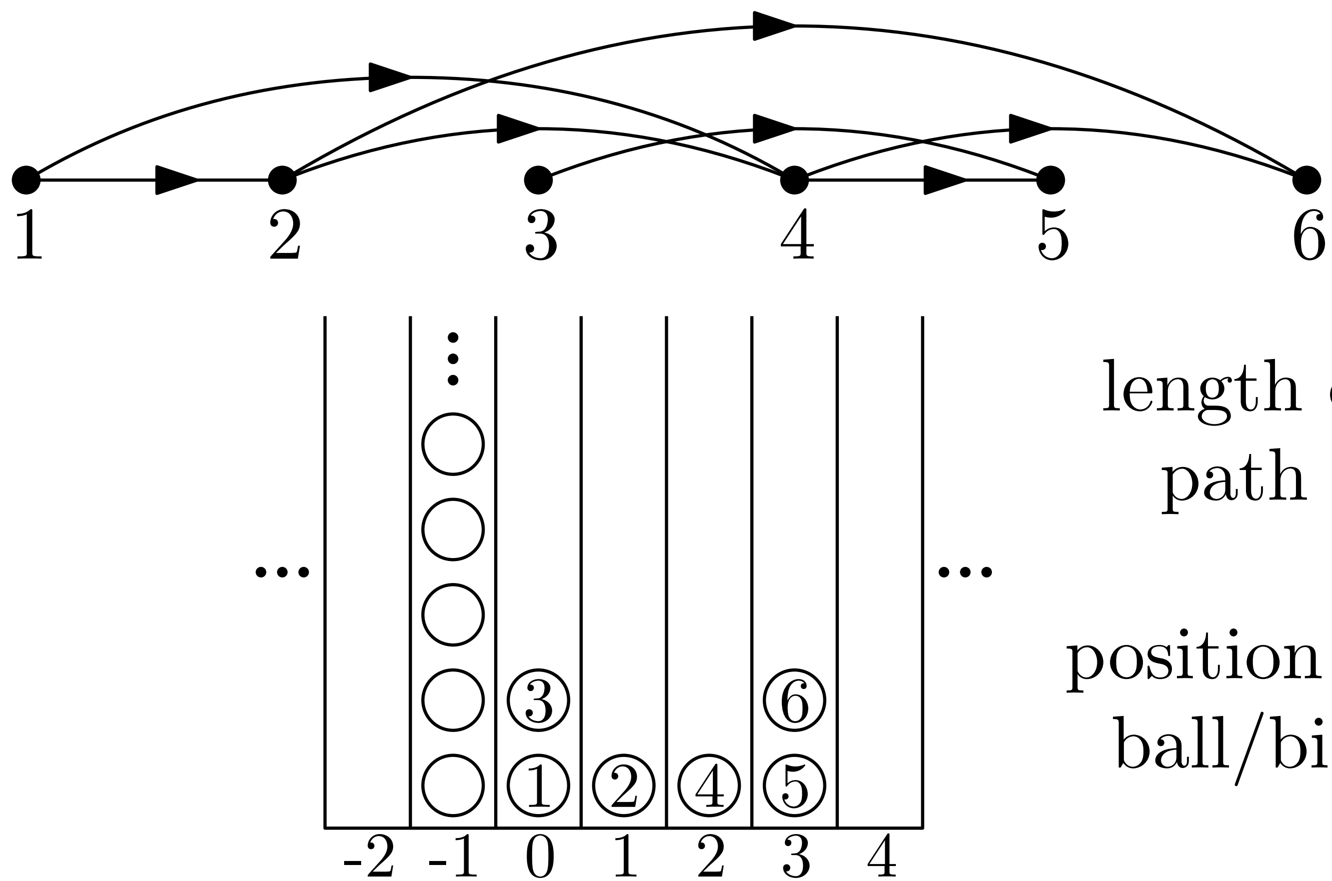

length of longest path in BEG = position of front in ball/bin process

Why is this random evolution of balls in bins an IBM(μ_p) ?

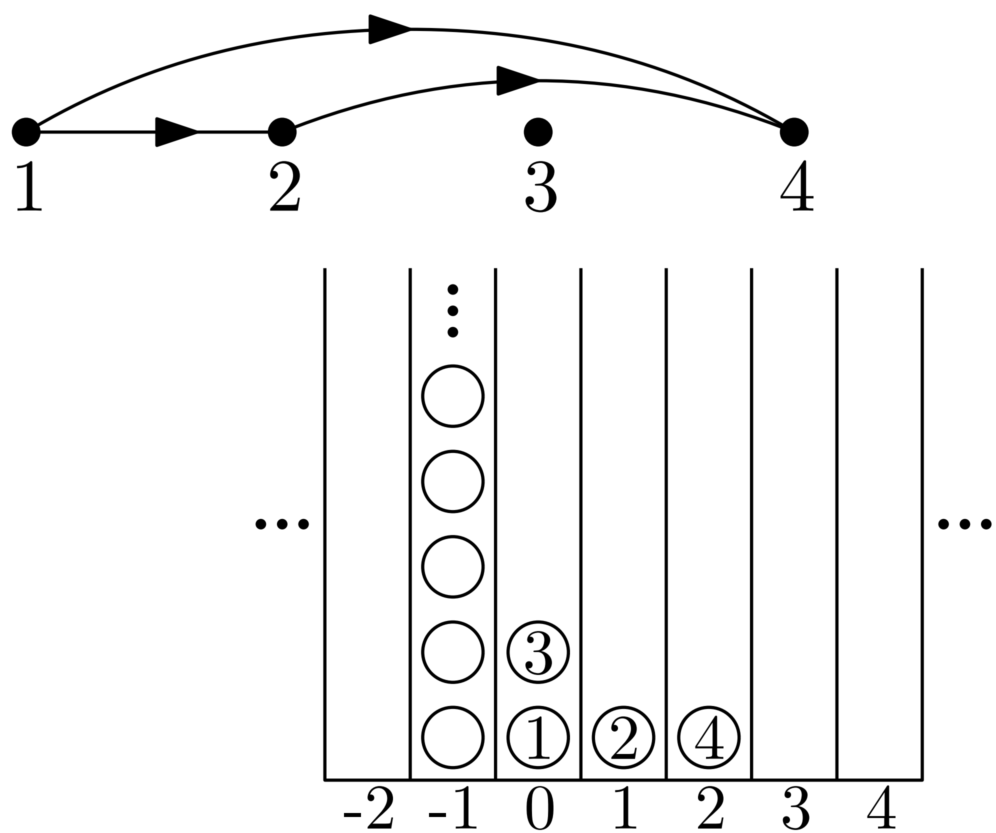

Why is this random evolution of balls in bins an IBM(μ_p) ?

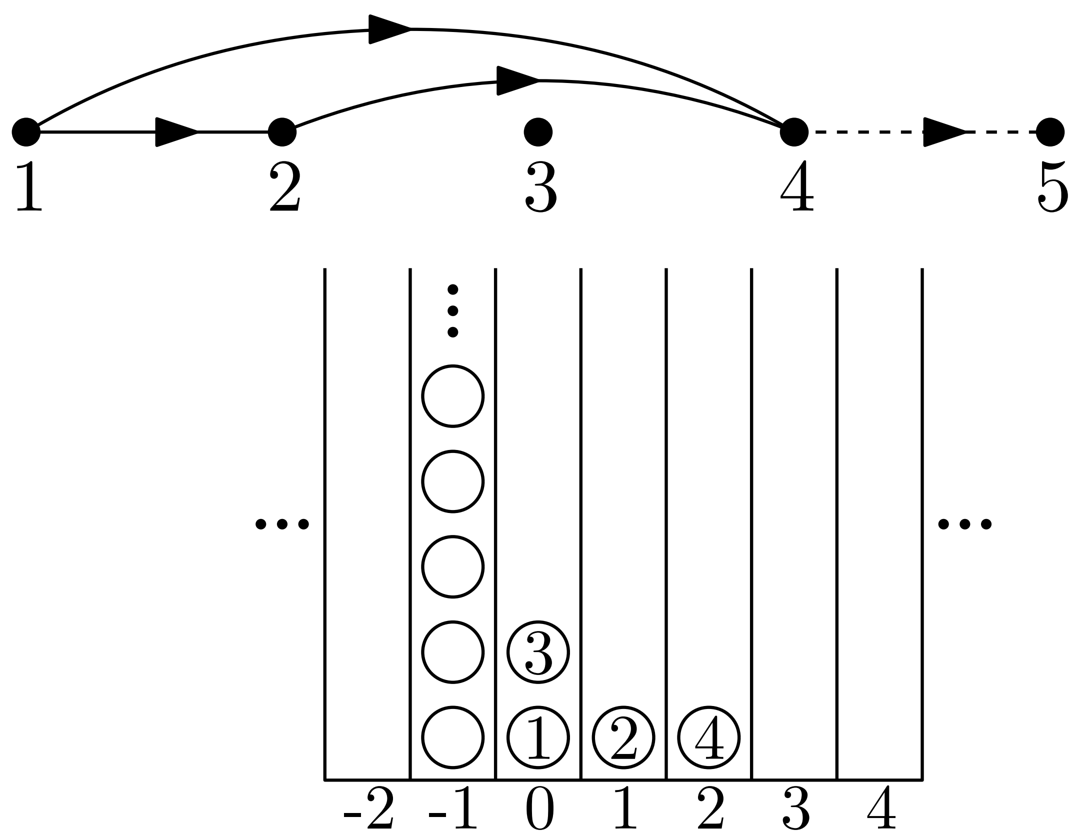

Why is this random evolution of balls in bins an IBM(μ_p) ?

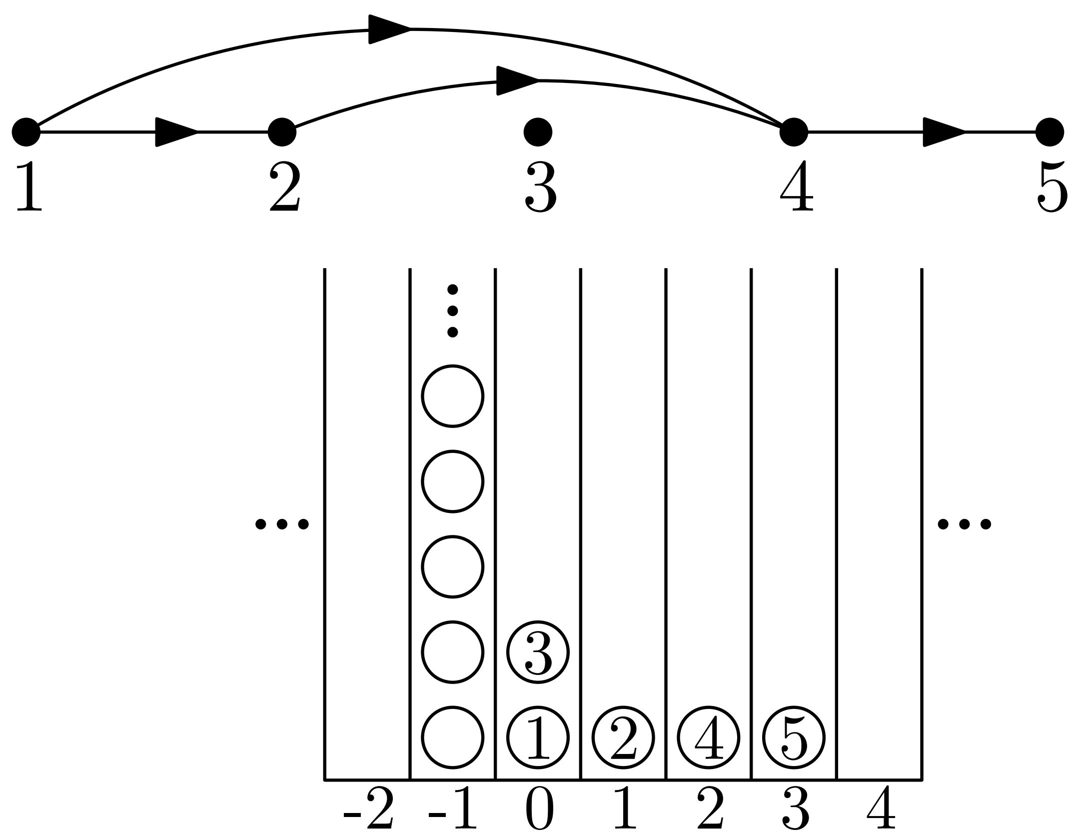

Why is this random evolution of balls in bins an IBM(μ_p) ?

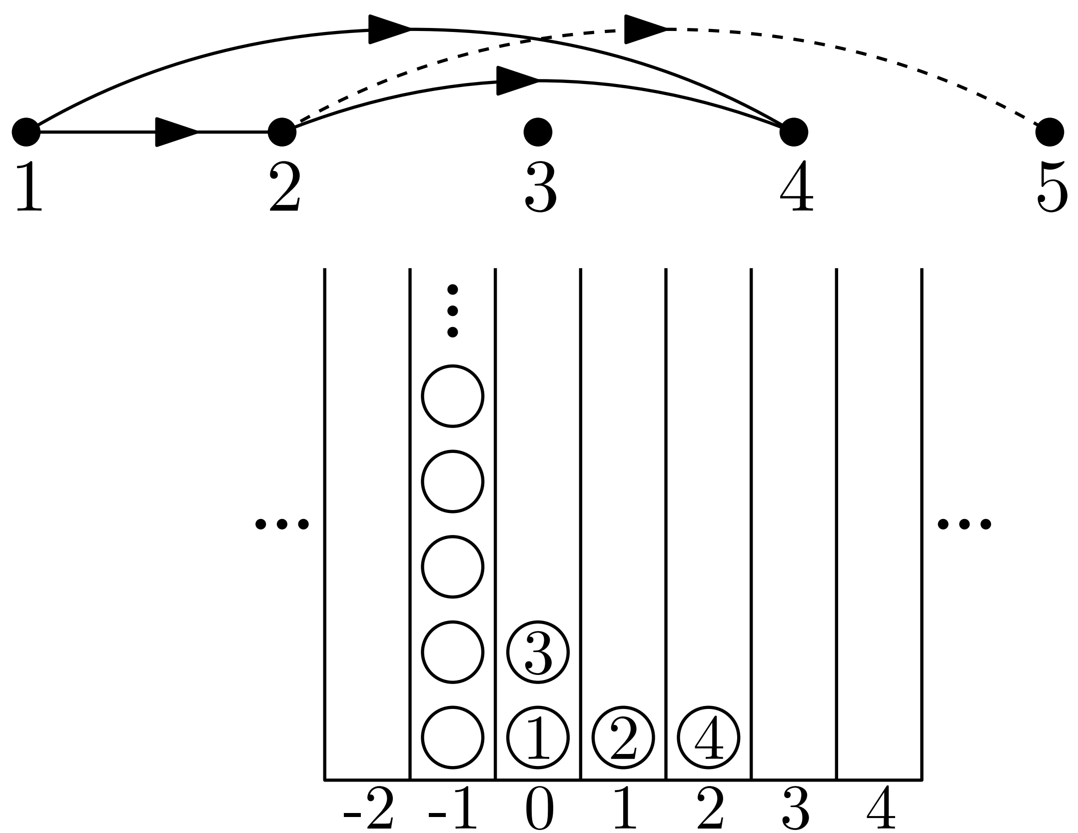

Why is this random evolution of balls in bins an IBM(μ_p) ?

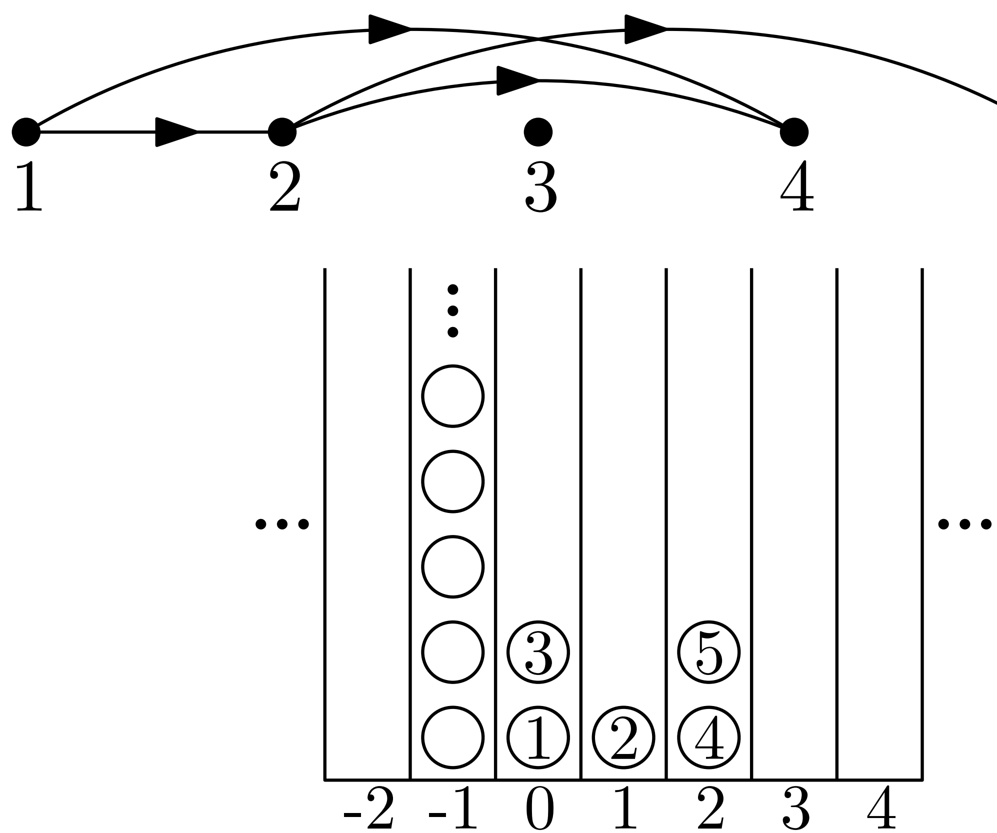

Why is this random evolution of balls in bins an IBM(μ_p) ?

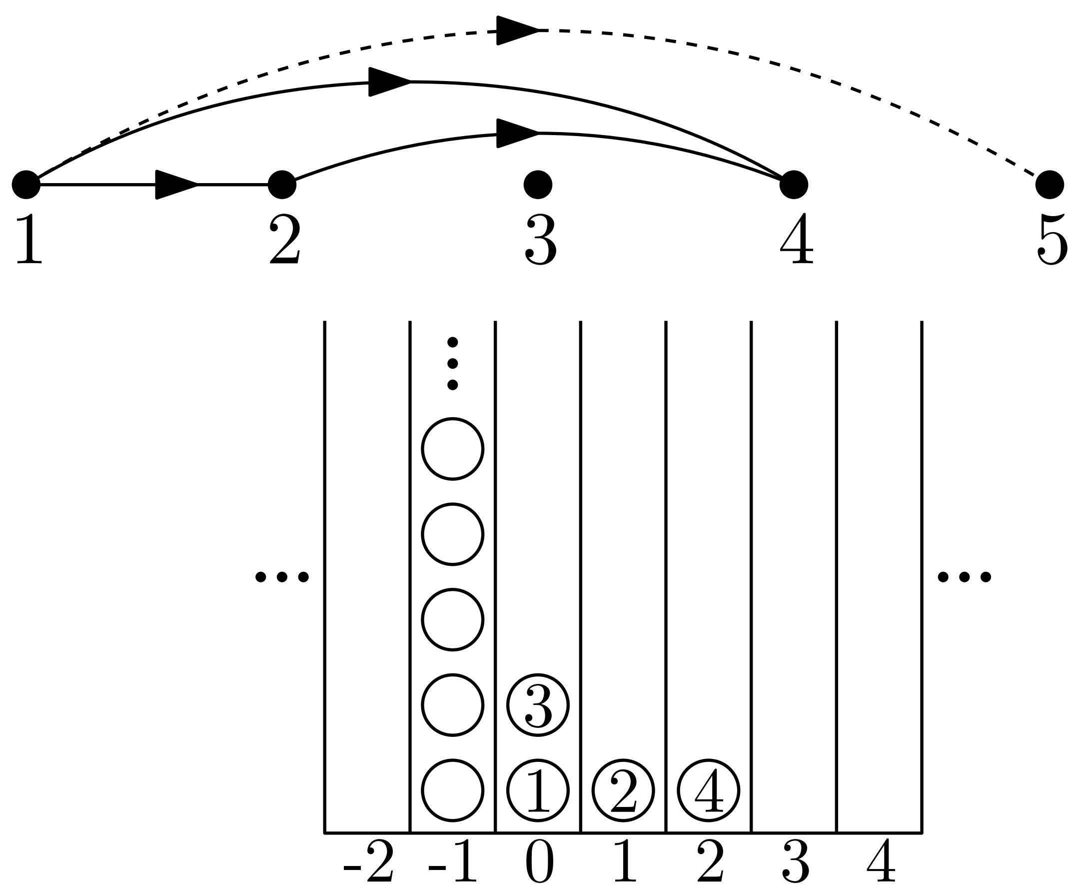

Why is this random evolution of balls in bins an IBM(μ_p) ?

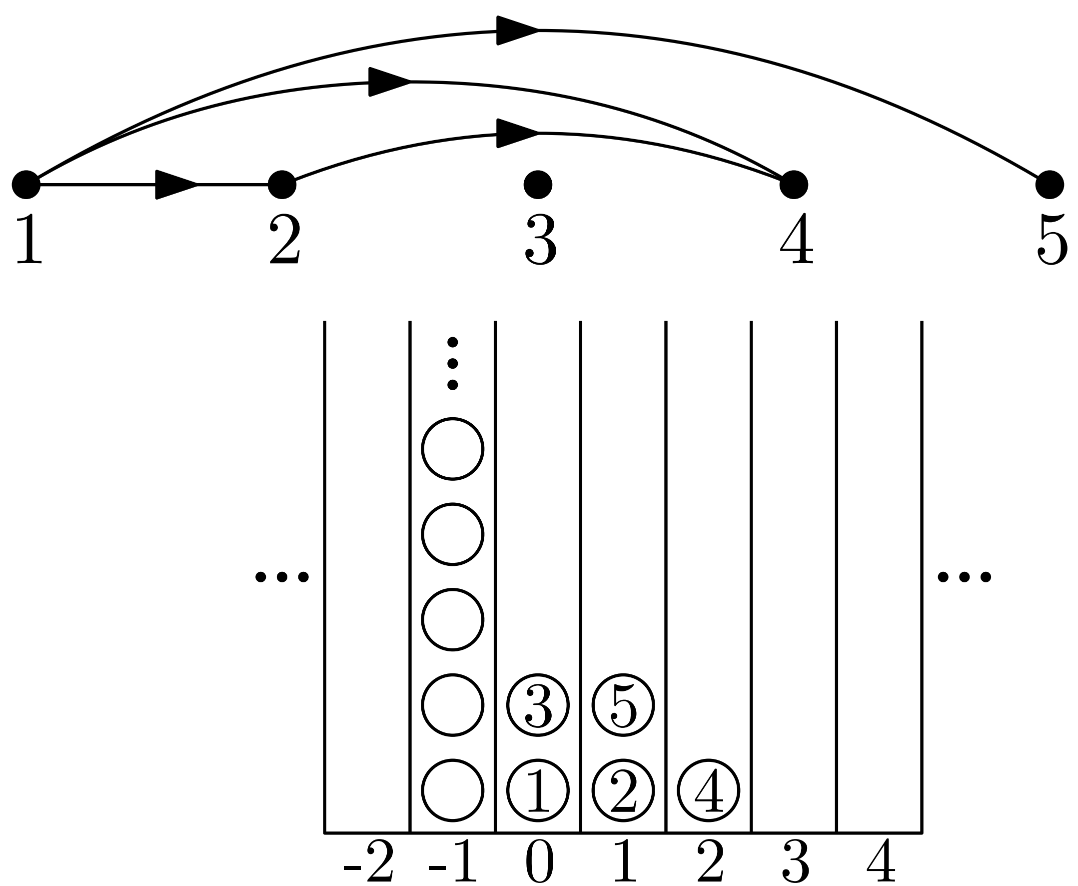

Why is this random evolution of balls in bins an IBM(μ_p) ?

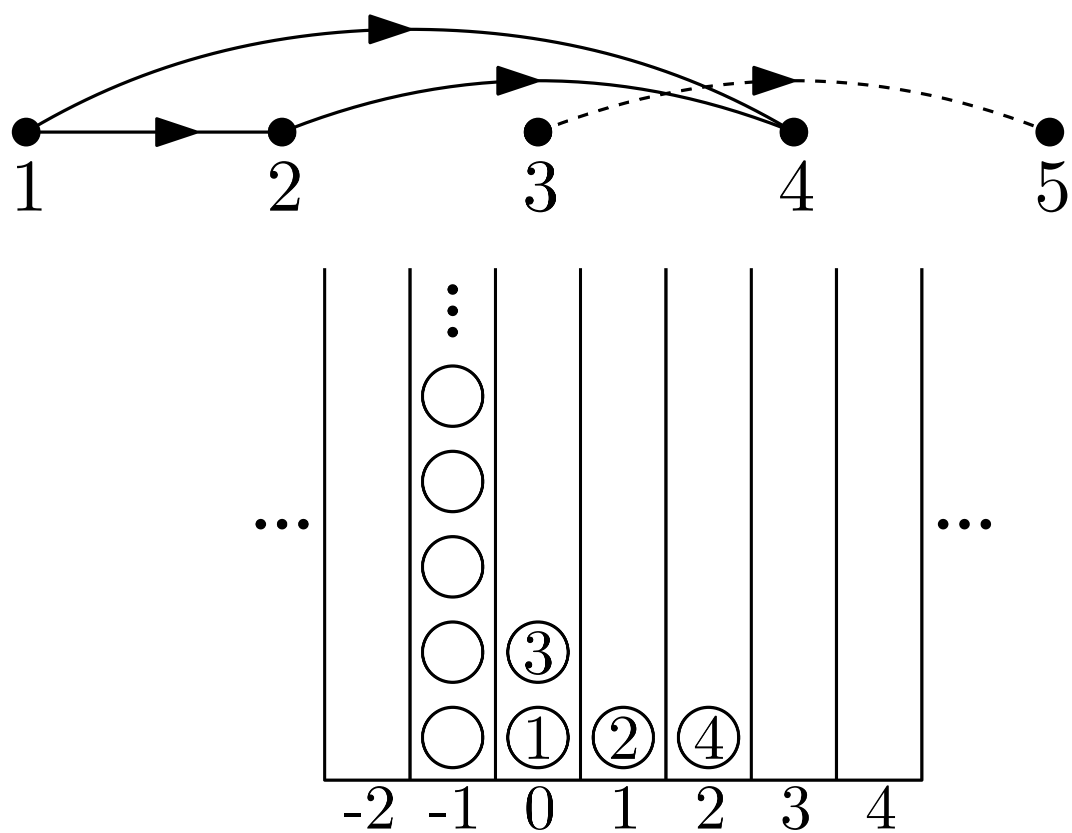

Why is this random evolution of balls in bins an IBM(μ_p) ?

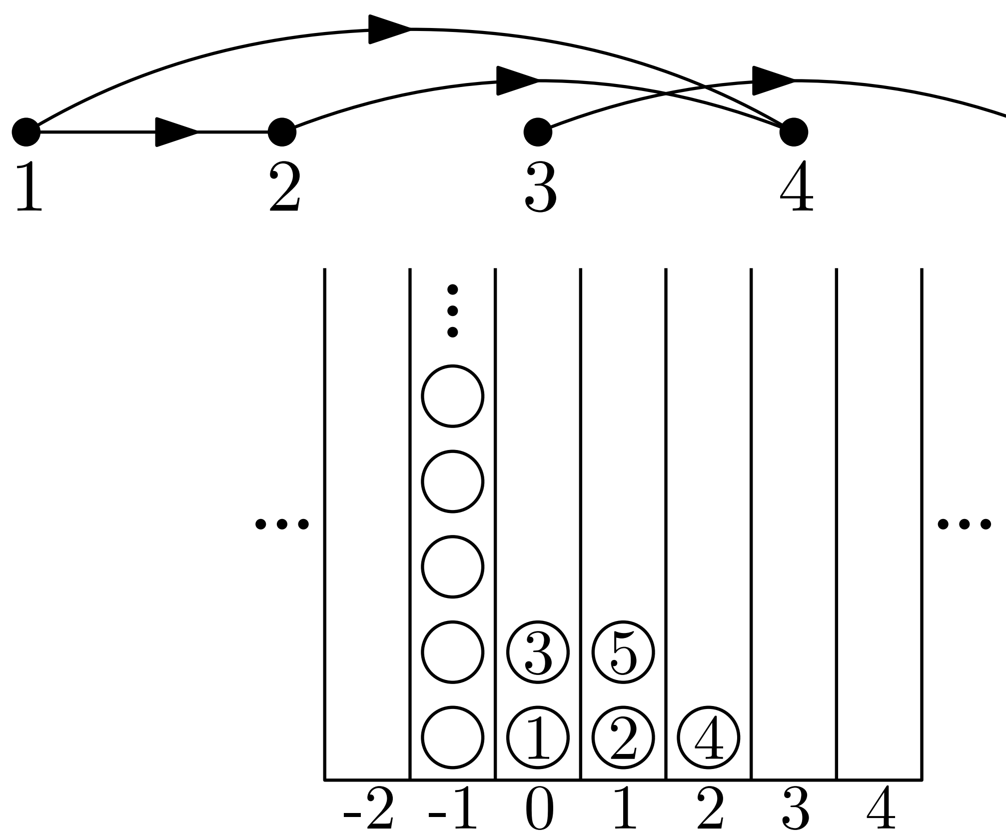

Why is this random evolution of balls in bins an IBM(μ_p) ?

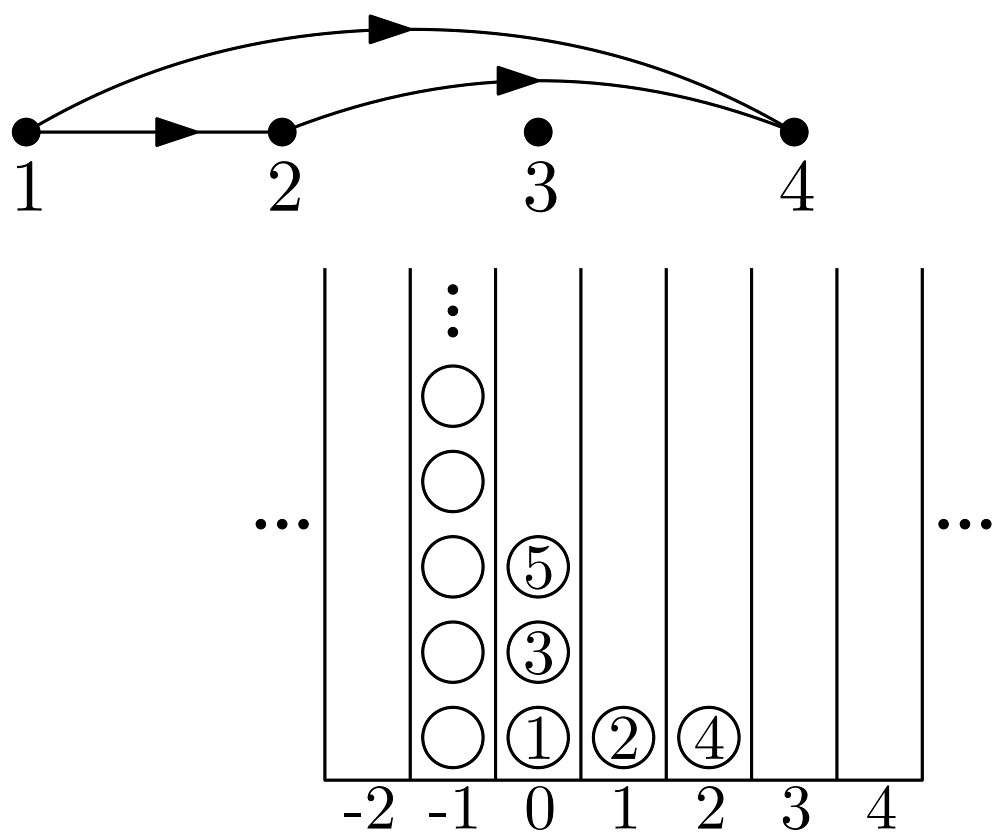

Why is this random evolution of balls in bins an IBM(μ_p) ?

IBM(μ_p): pick a ball whose rank follows the geometric distribution of parameter p and add a ball to its right.

Geometric distribution: number of trials until first success.

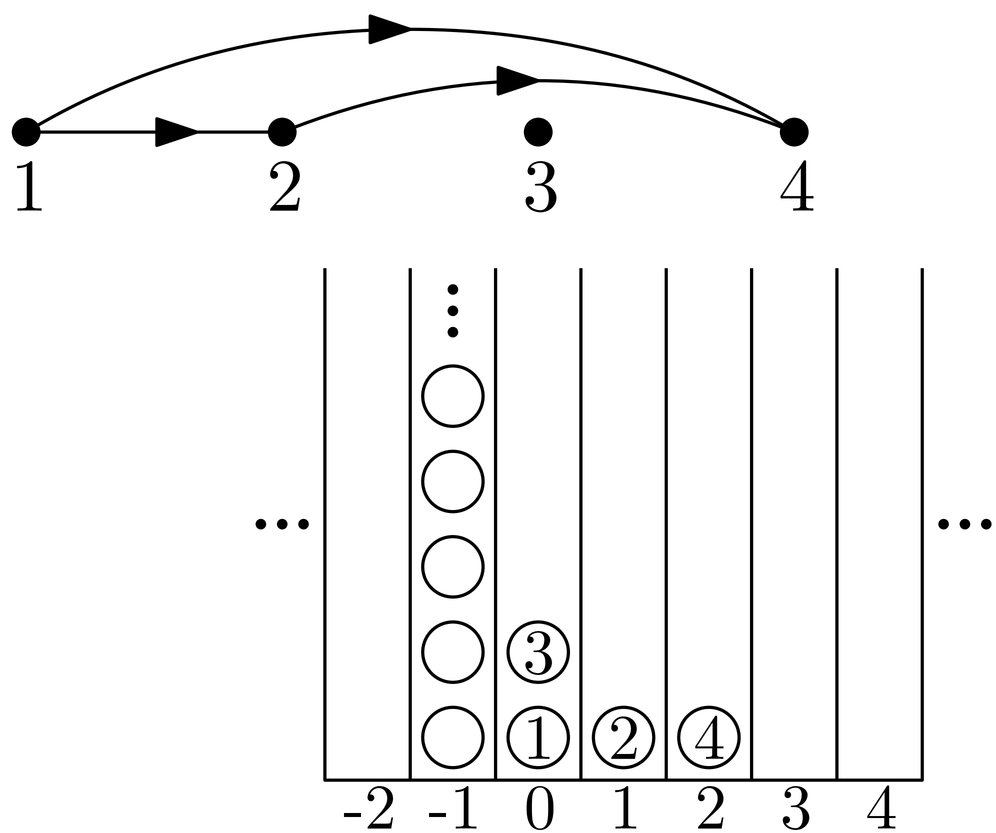

Stationary version of the IBM

- How to construct a stationary process, with time indexed by $\mathbb{Z}$ rather than $\mathbb{Z}_+$?

- Given a sequence $(\xi_n)_{n\in\mathbb{Z}}$, we want to construct a process $(X_n)_{n\in\mathbb{Z}}$ satisfying

$$\forall n \in \mathbb{Z}, X_n = \phi_{\xi_n}\left(X_{n-1}\right).$$

- Given the value of ξ_n for every $n \leq 0$, one can a.s. reconstruct X_0 up to a global shift (Foss-Konstantopoulos, Mallein-R.).

- Assume first that $\mu(1) > 0$ and μ has finite expectation.

- A finite sequence of positive integers $(u_1, \ldots, u_n)$ is called *triangular* if $u_i \leq i$ for all $1 \leq i \leq n$.

- E.g. 1132 is a triangular sequence, while 1324 is not.

- If one applies a triangular sequence of moves to an initial configuration Y_0, the only information needed to place the new balls is the position of the front of Y_0. All the new balls will be placed to the right of $F(Y_0)$.

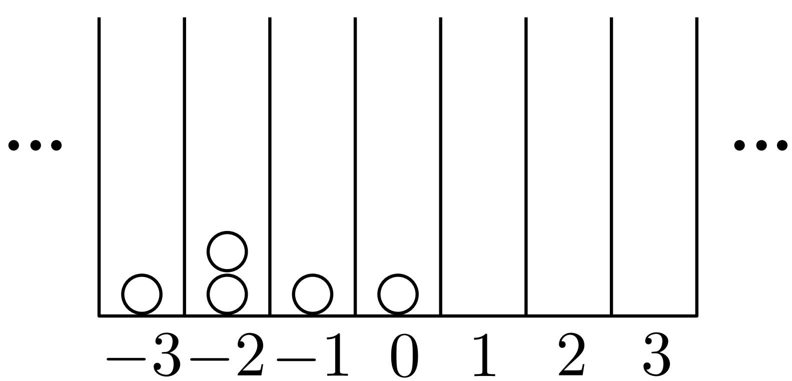

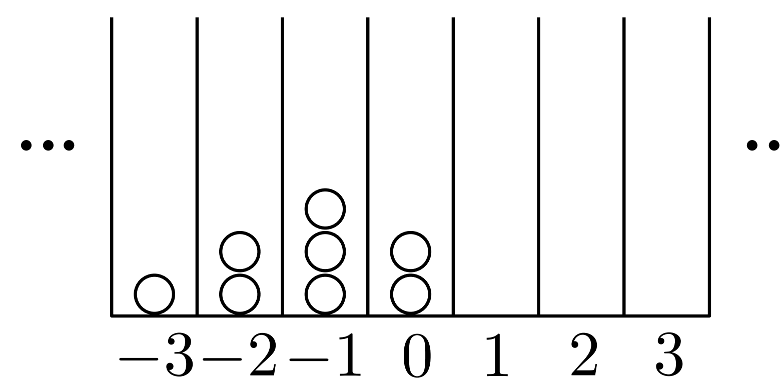

- A finite sequence of positive integers $(u_1, \ldots, u_n)$ is called *triangular* if $u_i \leq i$ for all $1 \leq i \leq n$.

- E.g. 1132 is a triangular sequence, while 1324 is not.

- If one applies a triangular sequence of moves to an initial configuration Y_0, the only information needed to place the new balls is the position of the front of Y_0. All the new balls will be placed to the right of $F(Y_0)$.

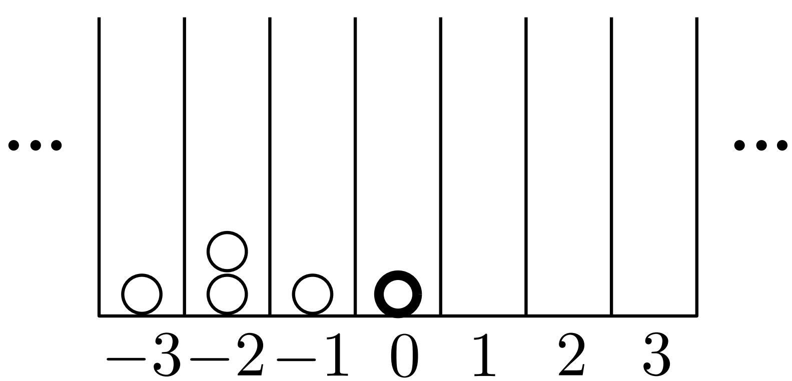

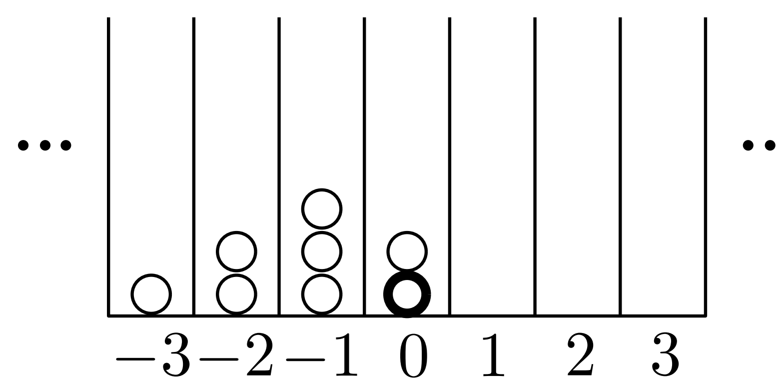

- A finite sequence of positive integers $(u_1, \ldots, u_n)$ is called *triangular* if $u_i \leq i$ for all $1 \leq i \leq n$.

- E.g. 1132 is a triangular sequence, while 1324 is not.

- If one applies a triangular sequence of moves to an initial configuration Y_0, the only information needed to place the new balls is the position of the front of Y_0. All the new balls will be placed to the right of $F(Y_0)$.

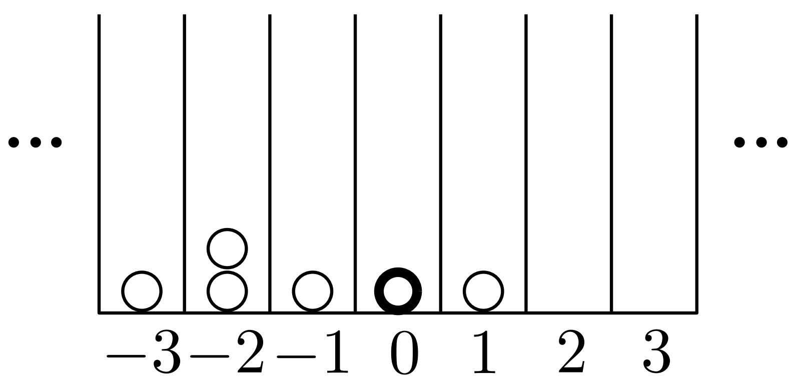

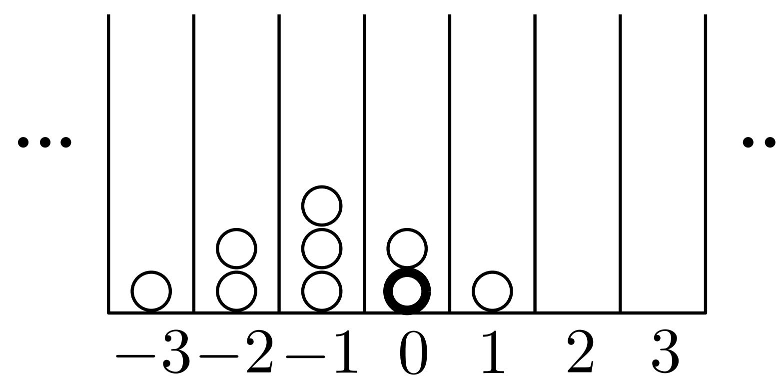

- A finite sequence of positive integers $(u_1, \ldots, u_n)$ is called *triangular* if $u_i \leq i$ for all $1 \leq i \leq n$.

- E.g. 1132 is a triangular sequence, while 1324 is not.

- If one applies a triangular sequence of moves to an initial configuration Y_0, the only information needed to place the new balls is the position of the front of Y_0. All the new balls will be placed to the right of $F(Y_0)$.

- A finite sequence of positive integers $(u_1, \ldots, u_n)$ is called *triangular* if $u_i \leq i$ for all $1 \leq i \leq n$.

- E.g. 1132 is a triangular sequence, while 1324 is not.

- If one applies a triangular sequence of moves to an initial configuration Y_0, the only information needed to place the new balls is the position of the front of Y_0. All the new balls will be placed to the right of $F(Y_0)$.

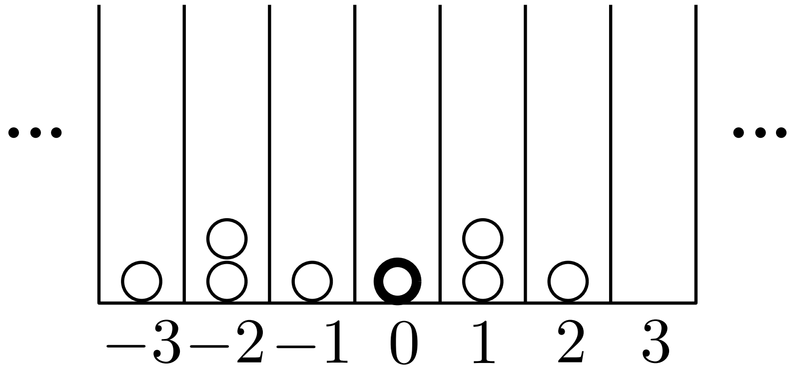

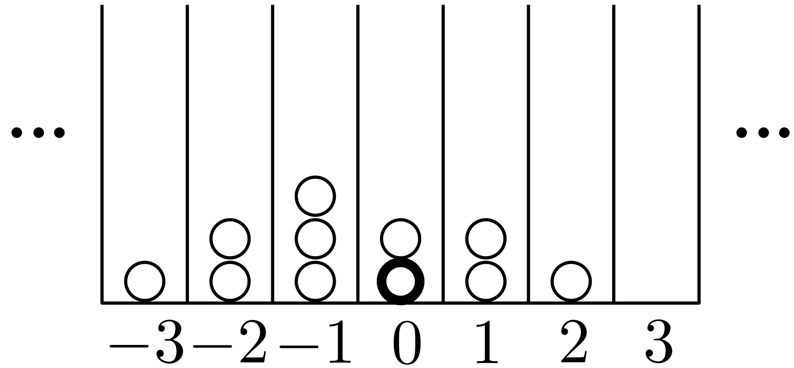

- A finite sequence of positive integers $(u_1, \ldots, u_n)$ is called *triangular* if $u_i \leq i$ for all $1 \leq i \leq n$.

- E.g. 1132 is a triangular sequence, while 1324 is not.

- If one applies a triangular sequence of moves to an initial configuration Y_0, the only information needed to place the new balls is the position of the front of Y_0. All the new balls will be placed to the right of $F(Y_0)$.

- Looking at the sequence $(\xi_n)_{n\leq 0}$, go back in time starting from 0 until you find the first triangular sequence ending at 0. This happens a.s. in finite time.

- Denote by $T_1 \leq 0$ this time. For $k \geq 2$, let T_k be the largest $t < T_{k-1}$ such that the sequence between t and 0 is triangular.

- Let $k \geq 1$. Whatever the initial configuration is, after application of the sequence of moves between times T_k and 0, the content of the rightmost k non-empty bins is the same.

1 2 2 4 1 3 2 1 2 1 2 2

↳ time 0

- Looking at the sequence $(\xi_n)_{n \leq 0}$, go back in time starting from 0 until you find the first triangular sequence ending at 0. This happens a.s. in finite time.

- Denote by $T_1 \leq 0$ this time. For $k \geq 2$, let T_k be the largest $t < T_{k-1}$ such that the sequence between t and 0 is triangular.

- Let $k \geq 1$. Whatever the initial configuration is, after application of the sequence of moves between times T_k and 0, the content of the rightmost k non-empty bins is the same.

1 2 2 4 1 3 2 1 2 $\boxed{1}$ 2 2

T_1 ↑ time 0

- Looking at the sequence $(\xi_n)_{n\leq 0}$, go back in time starting from 0 until you find the first triangular sequence ending at 0. This happens a.s. in finite time.

- Denote by $T_1 \leq 0$ this time. For $k \geq 2$, let T_k be the largest $t < T_{k-1}$ such that the sequence between t and 0 is triangular.

- Let $k \geq 1$. Whatever the initial configuration is, after application of the sequence of moves between times T_k and 0, the content of the rightmost k non-empty bins is the same.

$$1\ 2\ 2\ 4\ 1\ 3\ 2\ \underset{T_2}{\boxed{1}}\ 2\ \underset{T_1}{\boxed{1}}\ 2\ \underset{\uparrow \text{time } 0}{2}$$

- Looking at the sequence $(\xi_n)_{n \leq 0}$, go back in time starting from 0 until you find the first triangular sequence ending at 0. This happens a.s. in finite time.

- Denote by $T_1 \leq 0$ this time. For $k \geq 2$, let T_k be the largest $t < T_{k-1}$ such that the sequence between t and 0 is triangular.

- Let $k \geq 1$. Whatever the initial configuration is, after application of the sequence of moves between times T_k and 0, the content of the rightmost k non-empty bins is the same.

$$\boxed{1}\ 2\ 2\ 4\ 1\ 3\ 2\ \boxed{1}\ 2\ \boxed{1}\ 2\ 2$$

T_3 T_2 T_1 ↑ time 0

- Looking at the sequence $(\xi_n)_{n\leq 0}$, go back in time starting from 0 until you find the first triangular sequence ending at 0. This happens a.s. in finite time.

- Denote by $T_1 \leq 0$ this time. For $k \geq 2$, let T_k be the largest $t < T_{k-1}$ such that the sequence between t and 0 is triangular.

- Let $k \geq 1$. Whatever the initial configuration is, after application of the sequence of moves between times T_k and 0, the content of the rightmost k non-empty bins is the same.

- Define the content of the rightmost k non-empty bins of $X(0)$ to be that common content. This definition is compatible for different values of k.

- A finite sequence is called triangular minimal if it is triangular and it has no strict suffix which is triangular.

- E.g. 12 and 122 are triangular minimal, but not 12122.

- The subsequence of $(\xi_n)_{n \leq 0}$ between positions T_1 and 0 is triangular minimal.

Speed formula for the IBM

- For the bi-infinite stationary version of the IBM(μ), the speed of the front is the probability that the front advances at time 0.

- Can read this information from the sequence $(\xi_n)_{n \leq 0}$.

- Let $\mathcal{T}_m^+$ be the set of triangular minimal sequences such that the last move makes the front advance. E.g. 122 is in $\mathcal{T}_m^+$ but 12 is not.

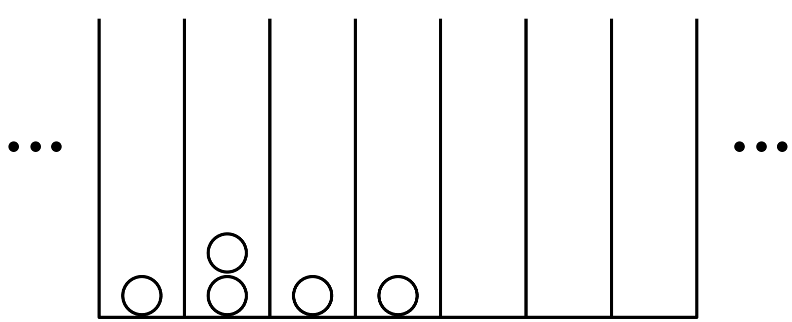

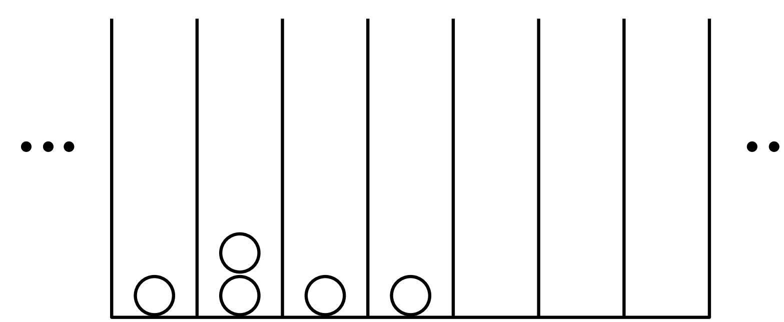

Speed formula for the IBM

- For the bi-infinite stationary version of the IBM(μ), the speed of the front is the probability that the front advances at time 0.

- Can read this information from the sequence $(\xi_n)_{n \leq 0}$.

- Let $\mathcal{T}_m^+$ be the set of triangular minimal sequences such that the last move makes the front advance. E.g. 122 is in $\mathcal{T}_m^+$ but 12 is not.

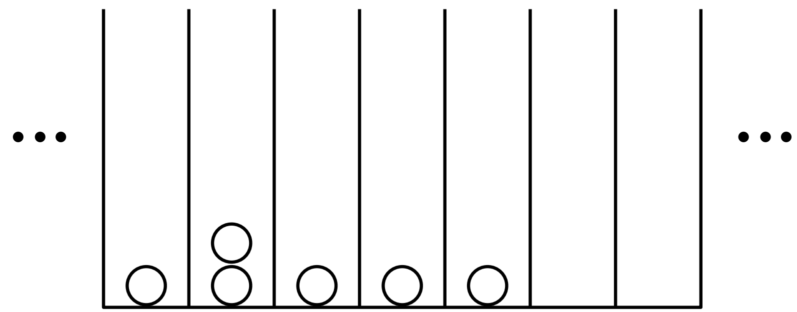

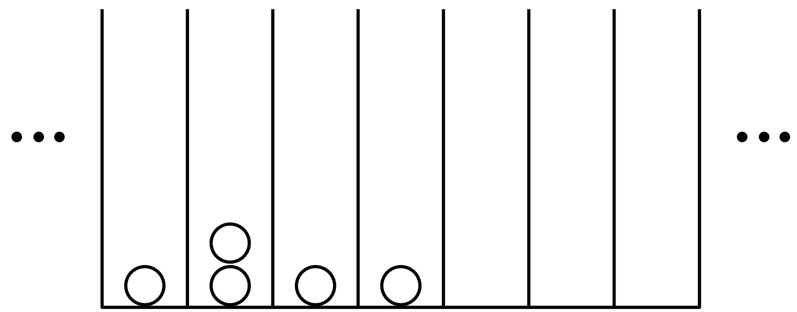

Speed formula for the IBM

- For the bi-infinite stationary version of the IBM(μ), the speed of the front is the probability that the front advances at time 0.
- Can read this information from the sequence $(\xi_n)_{n\leq 0}$.
- Let $\mathcal{T}_m^+$ be the set of triangular minimal sequences such that the last move makes the front advance. E.g. 122 is in $\mathcal{T}_m^+$ but 12 is not.

Speed formula for the IBM

- For the bi-infinite stationary version of the IBM(μ), the speed of the front is the probability that the front advances at time 0.

- Can read this information from the sequence $(\xi_n)_{n \leq 0}$.

- Let $\mathcal{T}_m^+$ be the set of triangular minimal sequences such that the last move makes the front advance. E.g. 122 is in $\mathcal{T}_m^+$ but 12 is not.

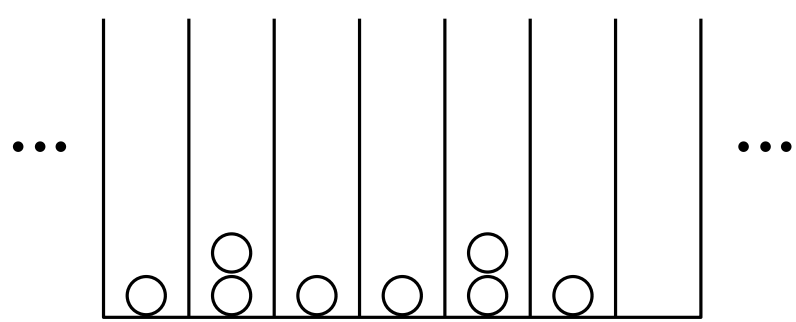

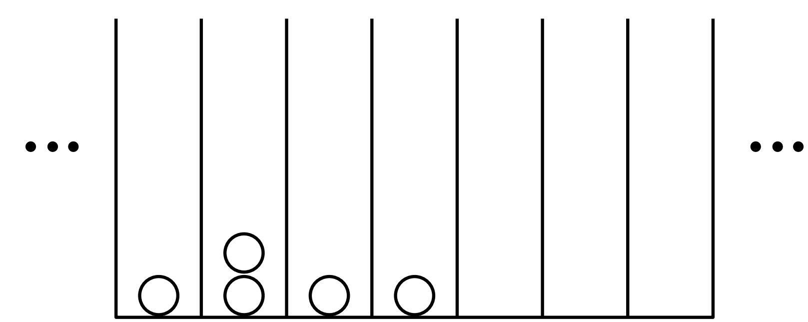

Speed formula for the IBM

- For the bi-infinite stationary version of the IBM(μ), the speed of the front is the probability that the front advances at time 0.
- Can read this information from the sequence $(\xi_n)_{n\leq 0}$.
- Let $\mathcal{T}_m^+$ be the set of triangular minimal sequences such that the last move makes the front advance. E.g. 122 is in $\mathcal{T}_m^+$ but 12 is not.

Speed formula for the IBM

- For the bi-infinite stationary version of the IBM(μ), the speed of the front is the probability that the front advances at time 0.

- Can read this information from the sequence $(\xi_n)_{n \leq 0}$.

- Let $\mathcal{T}_m^+$ be the set of triangular minimal sequences such that the last move makes the front advance. E.g. 122 is in $\mathcal{T}_m^+$ but 12 is not.

- Given a word α, the weight $w_\mu(\alpha)$ is the product of the probabilities under μ of the letters of α.

  E.g. $w_\mu(2,1,3) = \mu(2) \times \mu(1) \times \mu(3)$.

Theorem (Mallein-R.). *For any μ with finite expectation and such that $\mu(1) > 0$, the speed of the IBM(μ) is*

$$v_\mu = \sum_{\alpha \in \mathcal{T}_m^+} w_\mu(\alpha)$$

- There exists a more technical construction for an arbitrary non-Dirac probability measure μ, with a sum over a different class of words (Mallein-R.).

Perfect simulation

- Sampling exactly from the stationary distribution of a Markov process, unlike MCMC methods which sample from a distribution close to the stationary distribution.

- Perfect simulation is possible for any finite-dimensional marginal of the IBM. This produces an unbiased estimation for $C(p)$ (Foss-Konstantopoulos, Mallein-R.).

3 Properties of Barak-Erdős graphs via the IBM

Analyticity of $C(p)$

- The special case when μ is μ_p, the geometric distribution of parameter p, gives a formula for $v_{\mu_p} = C(p)$:

$$C(p) = \sum_{\alpha \in \mathcal{T}_m^+} p^{L(\alpha)} (1-p)^{H(\alpha)},$$

  where the height $H(\alpha)$ (resp. length $L(\alpha)$) of a sequence $\alpha = (\alpha_1, \ldots, \alpha_n)$ is defined to be $\alpha_1 + \cdots + \alpha_n - n$ (resp. n).

- Proving the existence of finite exponential moments for the time T_1 one has to wait before discovering a triangular sequence implies that $C(p)$ is analytic for $p > 0$.

$p \to 0$ limit

Theorem (Mallein-R.)**.**

$$C(p) = ep - \frac{e\pi^2 p}{2(\log p)^2} + o\left(\frac{p}{(\log p)^2}\right) \quad \textit{when } p \to 0.$$

In particular, $C(p)$ has no second derivative at $p = 0$.

Proof strategy

1. Compare to an IBM(μ) where μ is the uniform distribution on $\{1, \ldots, n\}$ and n is large.

2. The IBM with uniform distribution is coupled with a branching random walk with selection.

3. Use known estimates on branching random walks.

Step 1 : reduction to uniform case

- Want the speed of the IBM(μ_p), where μ_p is geometric with parameter p small.

- If $p = \frac{1}{n}$:

k	1	2	3	$\cdots$	n	$\cdots$
$\mu_p(k)$	$\frac{1}{n}$	$\frac{1}{n}\left(1-\frac{1}{n}\right)$	$\frac{1}{n}\left(1-\frac{1}{n}\right)^2$	$\cdots$	$\frac{1}{n}\left(1-\frac{1}{n}\right)^{n-1}$	$\cdots$

- Roughly equal to ν_n, the uniform distribution on $\{1,\ldots,n\}$.

Step 2 : branching random walk analogy

- IBM(ν_n) first studied by Aldous and Pitman in 1983 : speed behaves like $\frac{e}{n}$ when $n \to \infty$.

- Use a coupling with a continuous-time branching random walk on $\mathbb{Z}$ with selection of the rightmost n individuals :

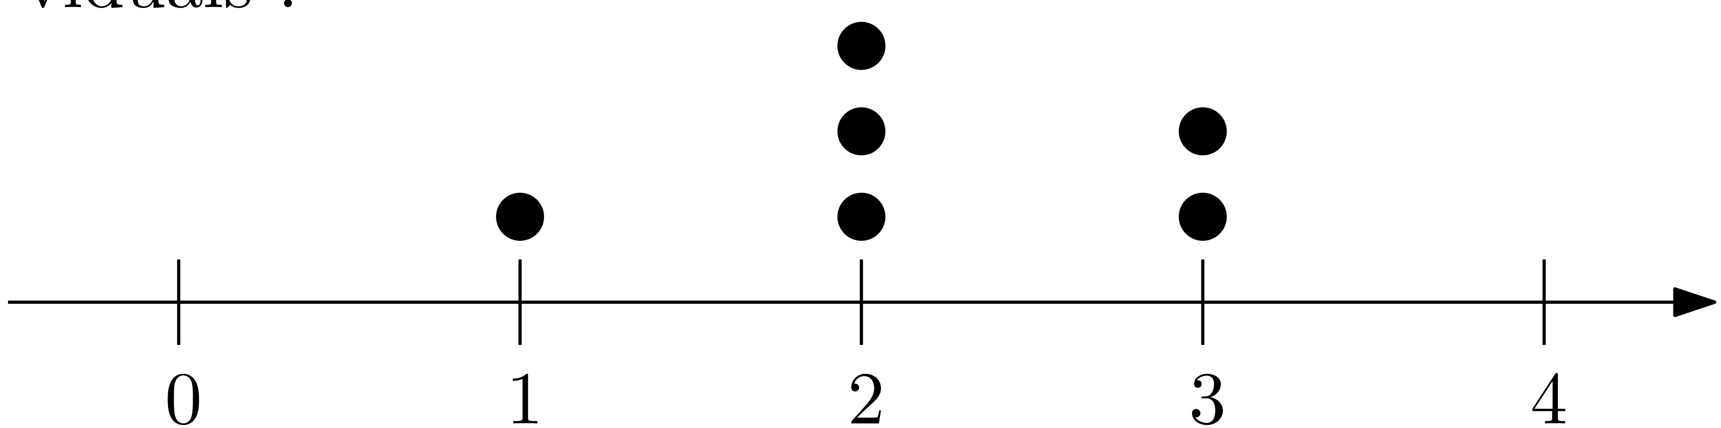

Step 2 : branching random walk analogy

- IBM(ν_n) first studied by Aldous and Pitman in 1983 : speed behaves like $\frac{e}{n}$ when $n \to \infty$.

- Use a coupling with a continuous-time branching random walk on $\mathbb{Z}$ with selection of the rightmost n individuals :

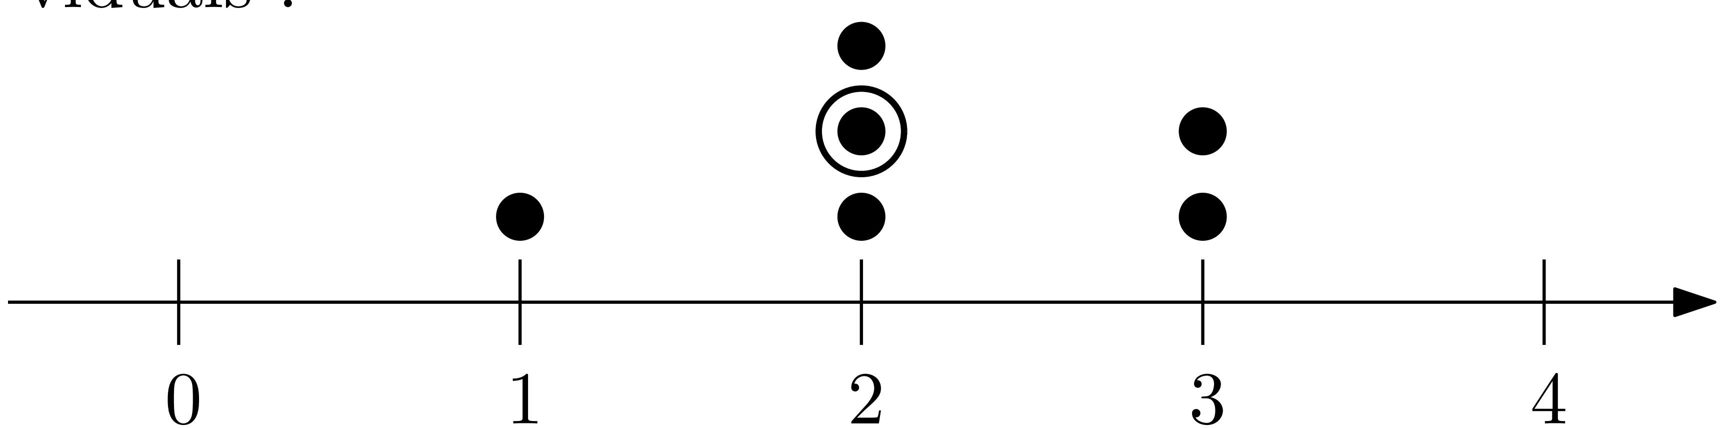

Step 2 : branching random walk analogy

- IBM(ν_n) first studied by Aldous and Pitman in 1983 : speed behaves like $\frac{e}{n}$ when $n \to \infty$.

- Use a coupling with a continuous-time branching random walk on $\mathbb{Z}$ with selection of the rightmost n individuals :

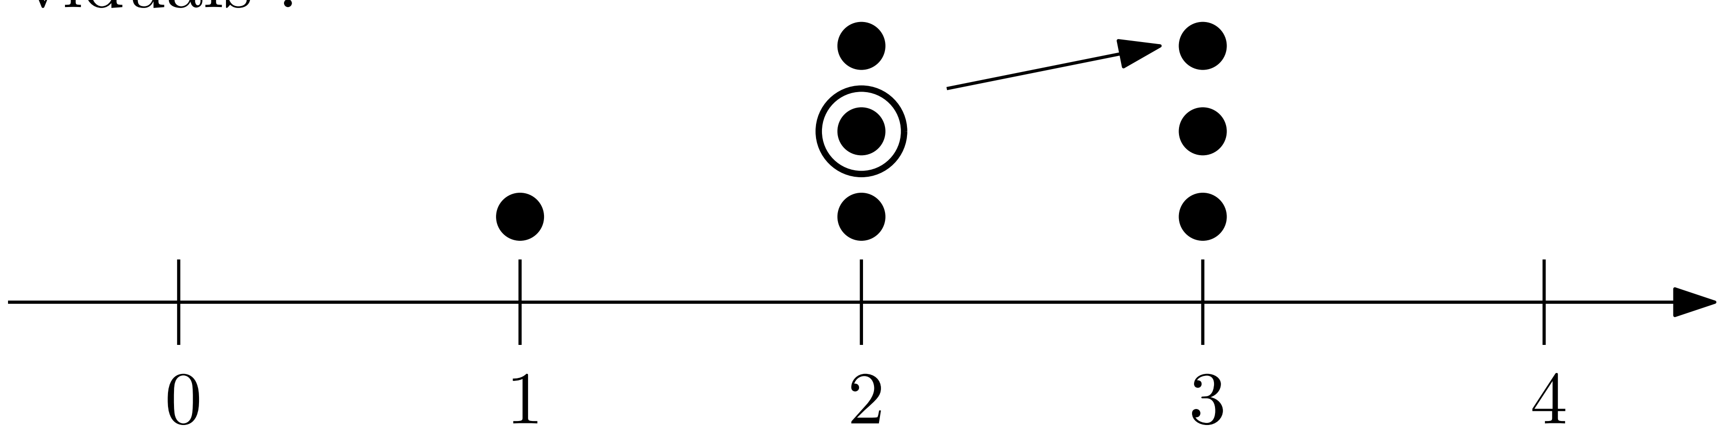

Step 2 : branching random walk analogy

- IBM(ν_n) first studied by Aldous and Pitman in 1983 : speed behaves like $\frac{e}{n}$ when $n \to \infty$.

- Use a coupling with a continuous-time branching random walk on $\mathbb{Z}$ with selection of the rightmost n individuals :

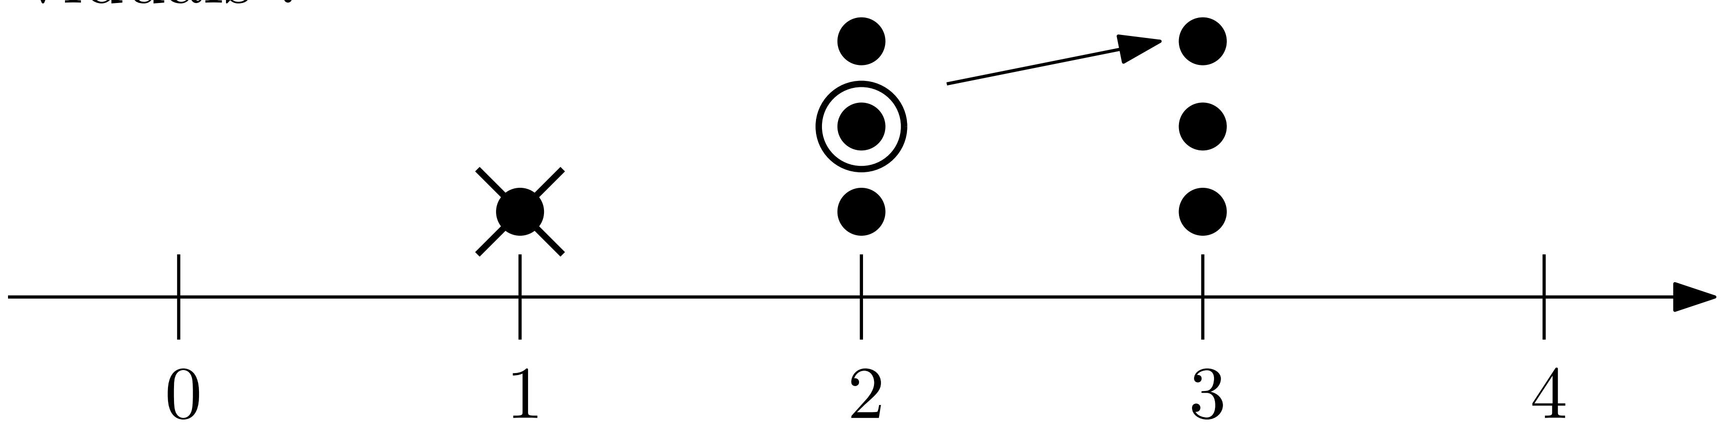

Step 2 : branching random walk analogy

- IBM(ν_n) first studied by Aldous and Pitman in 1983 : speed behaves like $\frac{e}{n}$ when $n \to \infty$.

- Use a coupling with a continuous-time branching random walk on $\mathbb{Z}$ with selection of the rightmost n individuals :

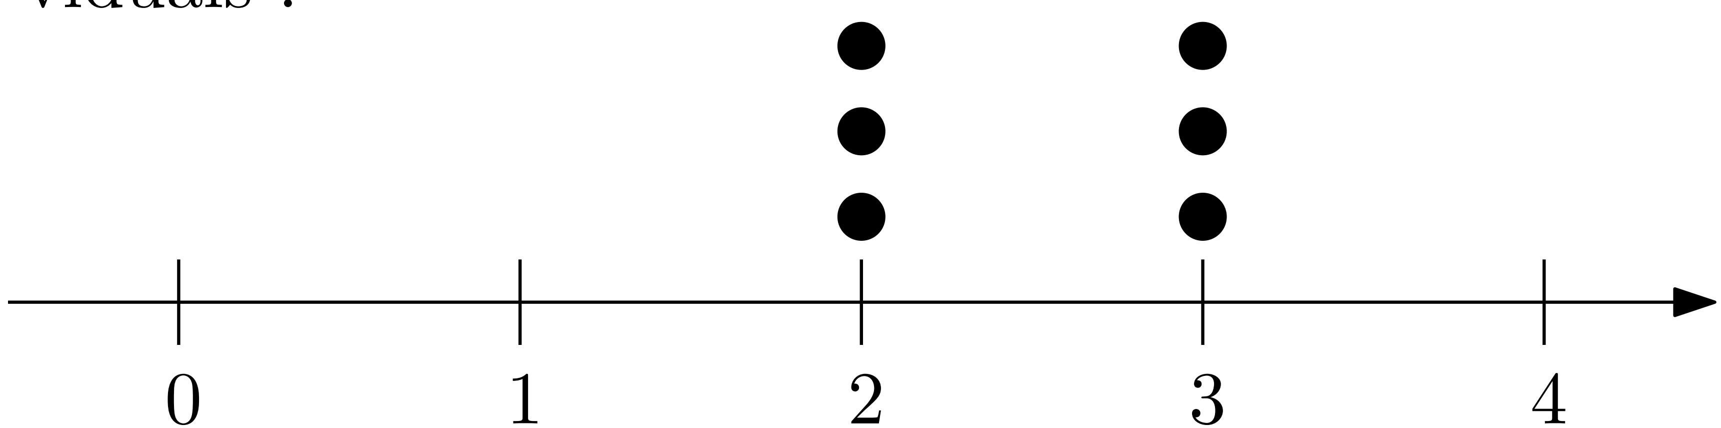

Step 3 : use branching random walk estimates

- Discrete-time branching random walk with selection widely studied: Brunet-Derrida '97, Bérard-Gouéré '10, Bérard-Maillard '14, Mallein '15...
- Translate these discrete-time results to the continuous setting.
- Conjecture for next term (after Brunet-Derrida):

$$\frac{3e\pi^2 p \log(-\log p)}{(-\log p)^3}$$

Power series expansion around $p = 1$

Theorem (Mallein-R.). $C(1-q)$ *can be expanded as a power series in* q *with radius at least* $\frac{\sqrt{2}-1}{2}$ *and its coefficients are integers.*

Fix a positive integer h. Then

$$C(1-q) = \sum_{\alpha \in \mathcal{T}_m^+} (1-q)^{L(\alpha)} q^{H(\alpha)}$$

$$= \sum_{\substack{\alpha \in \mathcal{T}_m^+ \\ H(\alpha) \le h}} (1-q)^{L(\alpha)} q^{H(\alpha)} + \sum_{\substack{\alpha \in \mathcal{T}_m^+ \\ H(\alpha) > h}} (1-q)^{L(\alpha)} q^{H(\alpha)}$$

Power series expansion around $p = 1$

$$C(1-q) = 1 - q + q^2 - 3q^3 + 7q^4 - 15q^5 + \cdots$$
$$= \sum_{k\geq 0} (-1)^k a_k q^k.$$

Foss-Konstantopoulos '03: first 5 terms
Mallein-R. '16: first 17 terms
Terlat '21: first 24 terms

Was not in the Online Encyclopedia of Integer Sequences before I added it (sequence A321309).

- Open problem: show that (a_k) forms an increasing sequence of positive integers and find a class of objects counted by them.

- Can one find some more or less explicit formula for $C(p)$?

If u_k is an integer sequence, write its generating function as

$$F(q) = \sum_{k \geq 0} u_k q^k .$$

$F(q)$ is *rational* if it is a quotient of two polynomials in q.

$F(q)$ is *algebraic* if it satisfies $G(q, F(q)) = 0$ for some bivariate polynomial G.

$F(q)$ is *D-finite* if it satisfies $H(q, F(q), F'(q), \cdots, F^{(m)}(q)) = 0$ for some H which is polynomial in its first variable and linear in its other variables.

$F(q)$ is *D-algebraic* if it satisfies a similar equation with H polynomial in every variable.

- Because of the singularity in $p/(\log p)^2$ we found around $p = 0$, $C(p)$ cannot be neither rational, nor algebraic, nor D-finite.

- Is $C(p)$ D-algebraic ?

- Maple package called *gfun* developed by Salvy: if you enter "enough" terms of the sequence u_k, it will guess the generating function if it falls in one of the four categories.

4 Extensions

Last passage percolation (LPP)

- Consider a deterministic directed acyclic graph and attach i.i.d. random lengths to each edge. The length of a path is the sum of the lengths of its edges.

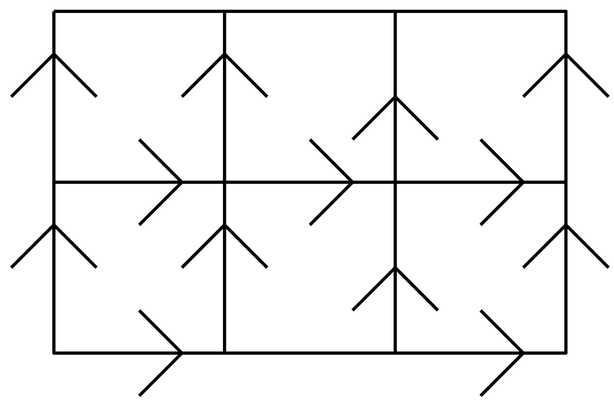

Last passage percolation (LPP)

- Consider a deterministic directed acyclic graph and attach i.i.d. random lengths to each edge. The length of a path is the sum of the lengths of its edges.

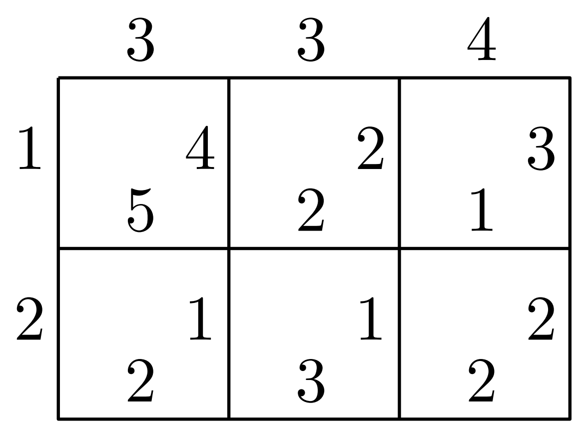

Last passage percolation (LPP)

- Consider a deterministic directed acyclic graph and attach i.i.d. random lengths to each edge. The length of a path is the sum of the lengths of its edges.

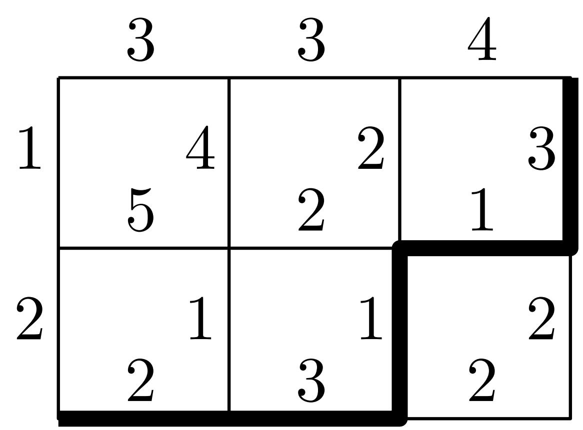

Path of length 10

Last passage percolation (LPP)

- Consider a deterministic directed acyclic graph and attach i.i.d. random lengths to each edge. The length of a path is the sum of the lengths of its edges.

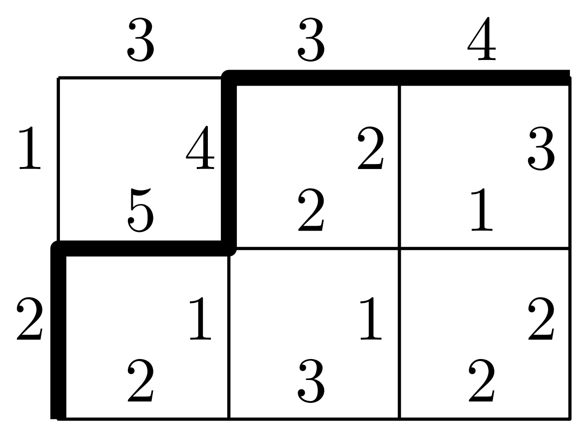

- How does the length of the longest path grow as the size of the graph grows ?

LPP on the complete graph on $\mathbb{Z}$

- Let ν be a probability distribution on $\mathbb{R} \cup \{-\infty\}$ such that its support has a finite upper bound.

- To each edge of the complete graph on $\mathbb{Z}$, attach an i.i.d. weight/length distributed like ν.

- Denote by C_ν the linear growth rate of the "length" of the "longest" path.

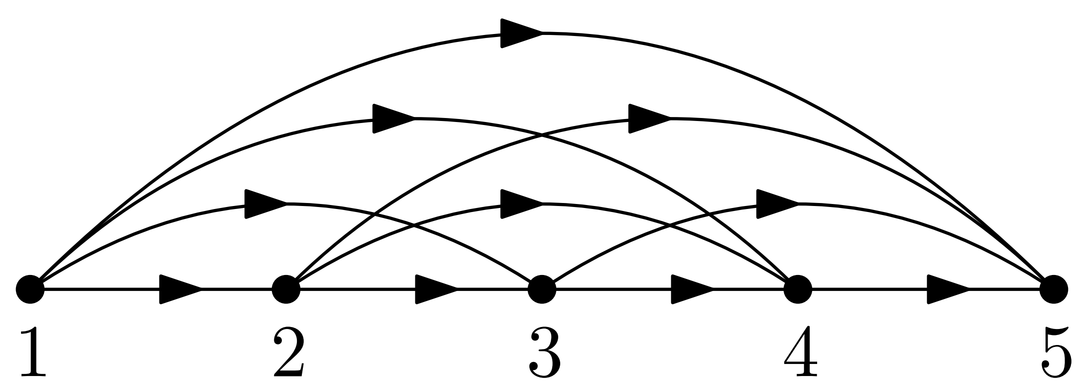

- The case when ν takes the value 1 with probability p and the value $-\infty$ with probability $1-p$ is the Barak-Erdős case. Then $C_\nu = C(p)$.

- Foss-Konstantopoulos-Mallein-R. '21: perfect simulation for C_ν via the introduction of an auxiliary particle system generalizing the geometric infinite-bin model.

Two-valued distribution ν

- Let $x \in \mathbb{R} \cup \{-\infty\}$. If $\nu = p\delta_1 + (1-p)\delta_x$, denote C_ν by $C(p,x)$. Then $C(p) = C(p, -\infty)$.

- Foss-Konstantopoulos-Pyatkin '20:
  The function $x \mapsto C(p,x)$ is non-differentiable exactly at values of x which are either 0 or a negative rational or of the form n or $1/n$, where n is an integer ≥ 2.

- Terlat '22: For fixed x, the function $p \mapsto C(p,x)$ is analytic on $(0,1]$. If $x > 0$, that function is actually rational on $[0,1]$.

Freezing (Mallein-R.-Singh ’22)

- Consider $(X_n)_{n\geq 0}$ an IBM(μ). Fix $k \in \mathbb{Z}$ and look at the total number of balls added to the kth bin during the entire process. Is it finite or infinite ?

- We say that the kth bin *freezes* if $X_\infty(k) < +\infty$ a.s.

- If there exists $M \geq 1$ such that $X_0(k) \leq M$ for every k (*uniformly bounded* starting configuration), then every bin freezes whatever μ is.

- If the starting configuration X_0 is *locally finite* (i.e. $X_0(k)$ is finite for every k), then the freezing depends on μ.

- Let (Y_n) be an IBM(μ) started from an infinite column at bin 0 (and nothing else).

- Let $k \geq 1$ be an integer. The measure μ is said to be *of type* k if a.s. $Y_\infty(1), \ldots, Y_\infty(k-1)$ are infinite and $Y_\infty(k)$ is finite.

- In other words, a measure of type k produces a cascade of $k-1$ infinite bins starting from a single infinite bin.

- The measure μ is of type 1 if and only if it has a finite first moment.

- If μ is of finite type k and the sequence $\mu(j)_{j\geq 1}$ is non-increasing, then every bin freezes for any choice of a locally finite X_0.

- If μ is of infinite type (i.e. a single infinite bin produces an infinite cascade of infinite bins), then there exists a locally finite initial configuration X_0 such that the IBM(μ) (X_n) verifies $X_\infty(k) = \infty$ for every $k \in \mathbb{Z}$ a.s.

THANK YOU !