Performance Evaluation of the Random Replacement Policy for Networks of Caches

Massimo Gallo[a], Bruno Kauffmann[b], Luca Muscariello[b], Alain Simonian[b], Christian Tanguy[b]

[a] *Alcatel-Lucent Bell Labs, Centre de Villarceaux, Route de Villejust, 91620 Nozay, France*
[b] *Orange Labs, 38–40 rue du Gnral Leclerc, 92794 Issy-Les-Moulineaux, France*

Abstract

The overall performance of content distribution networks as well as recently proposed information-centric networks rely on both memory and bandwidth capacities. The hit ratio is the key performance indicator which captures the bandwidth / memory tradeoff for a given global performance.
This paper focuses on the estimation of the hit ratio in a network of caches that employ the Random replacement policy (RND). Assuming that requests are independent and identically distributed, general expressions of miss probabilities for a single RND cache are provided as well as exact results for specific popularity distributions (such results also hold for the FIFO replacement policy). Moreover, for any Zipf popularity distribution with exponent $\alpha > 1$, we obtain asymptotic equivalents for the miss probability in the case of large cache size.
We extend the analysis to networks of RND caches, when the topology is either a line or a homogeneous tree. In that case, approximations for miss probabilities across the network are derived by neglecting time correlations between miss events at any node; the obtained results are compared to the same network using the Least-Recently-Used discipline, already addressed in the literature. We further analyze the case of a mixed tandem cache network where the two nodes employ either Random or Least-Recently-Used policies. In all scenarios, asymptotic formulas and approximations are extensively compared to simulation results and shown to be very accurate. Finally, our results enable us to propose recommendations for cache replacement disciplines in a network dedicated to content distribution.

1. Introduction

Communication networks use an ever increasing amount of data storage to cache information in the aim of performance improvement. Data caching con-

Email addresses: massimo.gallo@alcatel-lucent.com (Massimo Gallo), bruno.kauffmann@orange.com (Bruno Kauffmann), luca.muscariello@orange.com (Luca Muscariello), alain.simonian@orange.com (Alain Simonian), christian.tanguy@orange.com (Christian Tanguy)

sists in temporarily storing pieces of data into a memory, so as to directly provide the data upon possible forthcoming requests. The performance gain stems from the Round-Trip-Time reduction and the increase in network capacity when the cache is located downstream a bandwidth bottleneck, *e.g.*, a communication link with limited bandwidth or a shared bus in a network of chips.

As a major application, the increasing amount of content delivered to Internet users has pushed the use of Web caching into communication models based on distributed caching such as Content Delivery Networks (CDNs) or Peer-To-Peer (P2P) Networks. Additionally, new information-centric network architectures ([1], [2], [3]) have been recently proposed, that include built-in network storage as a central feature of the underlying communication model. Content storage then becomes a primary resource in such networks, aiming at minimizing content delivery time under an ever increasing demand that cannot simply be satisfied by increasing link bandwidth. On the economic side, the use of network storage to bypass bandwidth bottlenecks appears cost effective as memory turns out to be cheaper than transmission capacity.

One of the fundamental operations of a cache is defined by its replacement policy which determines the object to be removed from the cache when the latter is full. Many replacement policies are based on content popularity, with significant cost for managing the sorted lists. This is the case, in particular, for the Least Frequently Used (LFU) policy and more sophisticated variants of it. On the contrary, Most Recently Used (MRU), Least Recently Used (LRU), First-In-First-Out (FIFO) and Random (RND) policies have the compelling feature to replace cached objects with constant delay. In-network storage, as envisaged in the new architectures mentioned above, may require packet-level caching at line rate; current routers running complex replacement policies might not, however, sustain such high rates [4]. In this framework, RND or FIFO policies can therefore be seen as presenting the least possible complexity; in fact, the latter require less memory access per packet than LRU or MRU, and it has been shown [4] that this advantage is critical for sustaining high-speed caching with current memory technology.

In this paper, we address performance issues of caching networks running the RND replacement policy; our analysis also holds for the FIFO policy whose performance is known to be equivalent to that of RND. We mainly focus on the analytical characterization of the miss probabilities under the Independent Reference Model (IRM) assumptions when the number of available objects is infinite. Exploiting the Markovian properties of the cache occupancy and its associated product-form distribution, we first express the miss rate as a ratio of normalizing constants. This enables us to provide exact formulae for the miss rate in case of either geometric or specific Zipf content popularity distributions. On the basis of Large Deviations results for discrete probability distributions, Proposition 3.9 then asserts our main result: When the popularity distribution follows a general power-law with decay exponent $\alpha > 1$, the miss probability is asymptotic to $A\rho_\alpha/C^{\alpha-1}$ for large cache size C, where constants A and ρ_α depend on α only. In Proposition 3.10, we extend that result to miss probabilities conditioned by the object popularity rank.

A second major contribution of the paper is given by Proposition 5.1, where we evaluate the performance of networks of caches under the RND policy, for both linear and homogeneous tree networks and asymptotically Zipf popularity distributions. An approximate closed formula for the miss probability across the network is provided and compared to corresponding estimates for LRU cache networks. The analysis is also extended to the mixed tandem cache network where one cache employs LRU and the other uses RND.

The specific focus on Zipf distributions or, more generally, power-law distributions is motivated by numerous studies on Internet object popularity, starting from the late 90's experiments on World Wide Web documents ([5],[6]) to the content stored in enterprises media servers ([7],[8]) and recent studies on Internet media content ([9], [10]). While other content popularity distributions might be considered, we do not provide here a complete review of the literature on Internet content popularity characterization; the above references confirm the relevance of Zipf distributions for studying caching performance.

The remainder of the paper is organized as follows. Section 2 presents related work on the analytic performance evaluation of caching systems. Section 3 analyzes the RND cache replacement policy and its comparison to LRU for a single cache; these analytic results are compared to exact numerical and simulation results in Section 4. Section 5 reports the approximate analysis of the network of RND caches for two topologies, namely the line and the tree. Numerical and simulation results for the network case are reported in Section 6. Section 7 further evaluates the tandem cache system where one cache implements LRU and the other RND. Section 8 concludes the paper.

2. Related Work

There is a significant body of work on caching systems and their associated replacement policies; we here only report the literature focusing on the analytic characterization of the performance of such systems.

The replacement policy most often analyzed is LRU whose performance is evaluated considering the move-to-front rule, consisting in putting the latest requested object in front of a list; a miss event for a LRU cache with finite size takes place when the position (also referred to as search cost) of an object in the list is larger than that size. Under the Independent Reference Model, [11] calculates the expected search cost and its variance for finite lists. An explicit formula is given in [12, 13] for the probability distribution of that cost; such a formula is, however, impractical for numerical evaluation in case of large object population and large cache size. Integral representations obtained in [14, 15] using the Laplace transform of the search cost function reduce the problem to numerical integration.

An asymptotic analysis of LRU miss probabilities for Zipf and Weibull popularity distributions is derived in [16] and provides simple closed formulas. Extensions to correlated requests are obtained in [17, 18], showing that short-term correlation does not impact the asymptotic results derived in [16]; the case of variable object sizes is also considered in [19]. The average miss probability for a LRU cache when requests are issued according to a general, possibly non-stationary, stochastic point process is obtained in [20].

The analytic evaluation of the RND policy has first been initiated in [21] for a single cache where the miss probability is given a general expression for any popularity distribution. To the best of our knowledge, its application to specific popularity distributions has, however, not yet been envisaged together with its numerical tractability for large object population and cache size. Besides, a fluid limit analysis is performed in [22] by considering a sequence of single cache systems indexed by a integer n; for any popularity distribution, the average miss probability is then shown to tend a limit when n tends to infinity. The latter limit verifies an implicit equation and is related to the heuristic approximation proposed in [23] for the FIFO replacement policy. Finally, the so-called "Che approximation" has been applied in [24] to the RND replacement scheme; in that approach, the sojourn time for any object r in the cache is assumed to be approximately inversely proportional to the request rate for objects other than r. The miss rate is similarly shown to be given by an implicit equation and the validatity of the approach is validated numerically.

We are aware of a few papers that address the issue of networks of caches. Significant work is devoted to systems where a document is copied to all crossed caches. As we also assume in Section 5, most papers use the approximation that the output of any cache is also IRM, with filtered popularity. A network of LRU caches, in particular, has been analyzed in [25], using the above-mentioned approximation obtained in [23] for the miss probability at each node; miss probabilities can then be obtained as the solution of an iterative algorithm that is proved to converge. Using the same assumption, the results of [16] are also extended in [26] to a two-level request process where objects are segmented into packets, when assuming that the LRU policy applies to packets; the analysis is applied to line and tree topologies with in-path caching. Moreover, reference [27] extends [26] when network links have finite bandwidth. An even stronger assumption is used in [28], where hierarchical networks are analyzed assuming that, at each level, the requests form an independent process and have Zipf popularity. As noted in [28] and [29], the output of a cache is generally not IRM; reference [30] provides exact results for tree topologies of caches, when each cache keeps documents according to an exponentially distributed timer.

Besides, other papers focus on networks of caches where documents are not systematically copied to all caches along their path. In particular, the interconnection of LRU caches using the Leave a Copy Down mechanism (each cache of a given level stores the transmitted object) is analyzed in [31]. Mixed replacement policies are analyzed in the context of a hierarchical cache in [32], and hybrid policies with good performance and implementation cost are proposed.

3. Single cache model

In this section, we address the single cache system with RND replacement policy, deriving analytic expressions of the miss probability together with asymptotics for large cache size. To avoid technical results at first reading, the reader may quickly go through the notation of Section 3.1 and directly skip to main results given in Propositions 3.9 and 3.10.

3.1. Basic results

Consider a cache memory with finite size which is offered requests for objects. If the cache size is attained and a request for an object cannot be satisfied (corresponding to a *miss* event), the requested object is fetched from the repository server and cached locally at the expense of replacing some other object in the cache. The object replacement policy is assumed to follow the RND discipline, *i.e.*whenever a miss occurs, the object to be replaced is chosen at random, uniformly among the objects present in cache.

We consider a discrete time system: at time $t \in \mathbb{N}$, the t-th object requested from the cache is denoted by $R(t) \in \{1, 2, ..., N\}$, where $N \in \mathbb{N} \cup \{+\infty\}$ is the total number of objects which can be possibly requested (in the present analysis, the set of all possible objects is considered to be invariant in time). We assume that all N objects are ordered with decreasing popularity, the probability that the object with popularity rank r is requested being q_r, $1 \leq r \leq N$. In the following, we consider the *Independent Reference Model*, where variables $R(t)$, $t \in \mathbb{N}$, are mutually independent and identically distributed with distribution defined by

$$\mathbb{P}(R = r) = q_r, \quad 1 \leq r \leq N.$$

Let $C \leq N$ be the cache size. We denote by $\mathcal{N}_C$ the set of all ordered vectors $(j_1, ..., j_C) \in \{1, ..., N\}^C$ with $j_k < j_\ell$ for $k < \ell$. Define the cache state at time $t \in \mathbb{N}$ by the vector $\mathbf{S}(t) \in \mathcal{N}_C$, where component $S_k(t)$, $1 \leq k \leq C$, is the rank of the object with the kth highest popularity in the cache. Define

$$G(C) = \sum_{(j_1,...,j_C) \in \mathcal{N}_C} q_{j_1} ... q_{j_C} \tag{1}$$

with $G(0) = 1$, and let $M(C)$ finally denote the stationary miss probability calculated over all possibly requested objects.

It is shown [21] that the cache configurations $(\mathbf{S}(t))_{t \in \mathbb{N}}$ define a reversible Markov process with stationary probability distribution given by

$$\mathbb{P}(\mathbf{S} = \mathbf{s}) = \frac{1}{G(C)} \prod_{j \in \mathbf{s}} q_j, \quad \mathbf{s} \in \mathcal{N}_C; \tag{2}$$

moreover ([21], Theorem 4), miss probability $M(C)$ equals

$$M(C) = \frac{\displaystyle\sum_{(j_1,...,j_C) \in \mathcal{N}_C} q_{j_1} ... q_{j_C} \sum_{r \notin (j_1,...,j_C)} q_r}{\displaystyle\sum_{(j_1,...,j_C) \in \mathcal{N}_C} q_{j_1} ... q_{j_C}}. \tag{3}$$

The latter formula is easily understood as the sum, for each possible cache configuration, of all probabilities to request an object which is not available in the cache, weighted by the probability to be in that configuration. It is also shown [21] that the latter stationary distribution and miss rate probability for RND are identical to that of a cache using the FIFO replacement policy, so that our results apply to that policy as well.

We now show that expression (3) can be written in terms of normalizing constants $G(C)$ and $G(C+1)$ only; this will give formula (3) a compact form suitable for the derivation of both exact and asymptotic expressions for $M(C)$.

Lemma 3.1. *The miss rate $M(C)$ is given by*

$$M(C) = (C+1)\frac{G(C+1)}{G(C)} \tag{4}$$

with $G(C)$ defined in (1).

Proof. The denominator in (3) equals $G(C)$ by definition. The numerator can be expressed as

$$\begin{aligned}
&\sum_{1 \leq j_1 < ... < j_C \leq N} q_{j_1} ... q_{j_C} \sum_{r \notin \{j_1, ..., j_C\}} q_r \\
&= \sum_{1 \leq j_1 < ... < j_C \leq N} q_{j_1} ... q_{j_C} \times \left(\sum_{1 \leq r < j_1} q_r + ... + \sum_{j_{C-1} < r < j_C} q_r + \sum_{j_C < r \leq N} q_r \right) \\
&= (C+1) \sum_{1 \leq r < j_1 < ... < j_C \leq N} q_r q_{j_1} ... q_{j_C} = (C+1) G(C+1)
\end{aligned}$$

and expression (4) of $M(C)$ follows. □

The latter results readily extend to the case when the total number N of objects is infinite, since the sum $\Sigma_{j \geq 1} q_j$ is finite. The calculation of coefficients $G(C)$, $0 \leq C \leq N$, is now performed through their associated generating function F defined by

$$F(z) = \sum_{0 \leq C \leq N} G(C) z^C, \quad z \in \mathbb{C}, \tag{5}$$

for either finite or infinite population size N (as $M(C) \leq 1$, Lemma 3.1 entails that $G(C+1)/G(C) \leq 1/(C+1)$ and the ratio test implies that the power series defining $F(z)$ has infinite convergence radius). We easily obtain the second preliminary result.

Lemma 3.2. *The generating function F is given by*

$$F(z) = \prod_{1 \leq r \leq N} (1 + q_r z) \tag{6}$$

for all $z \in \mathbb{C}$.

Proof. Expanding the latter product and using definition (1) readily provide the claimed result. □

To further study the single cache properties, let $M_r(C)$ denote the per-object miss probability, given that the requested object is precisely $r \in \{1, ..., N\}$. Defining

$$G_r(C) = \sum_{1 \leq j_1 < ... < j_C \leq N,\ r \notin \{j_1, ..., j_C\}} q_{j_1} ... q_{j_C} \tag{7}$$

with $G_r(0) = 1$, we then have $M_r(C) = \mathbb{P}(r \notin \mathbf{S})$ so that (2) and (7) yield

$$M_r(C) = \frac{G_r(C)}{G(C)}. \tag{8}$$

Lemma 3.3. *For given $r \in \{1, ..., N\}$, the per-object miss probability $M_r(C)$ can be expressed by*

$$M_r(C) = 1 + \sum_{\ell=1}^{C} (-1)^\ell q_r^\ell \frac{G(C - \ell)}{G(C)}. \tag{9}$$

The stationary probability $q_r(2)$, $r \in \{1, ..., N\}$, that a miss event occurs for object r is given by

$$q_r(2) = \frac{M_r(C)}{M(C)} q_r \tag{10}$$

where $M(C)$ is the averaged miss probability.

Proof. By definition (7), the generating function $F_r(z)$ of coefficients $G_r(C)$, $0 \leq C \leq N$, is given by

$$F_r(z) = \frac{F(z)}{1 + q_r z}, \quad z \in \mathbb{C}. \tag{11}$$

Expanding the latter ratio as a powers series of z gives

$$G_r(C) = \sum_{\ell=0}^{C} (-1)^{C-\ell} q_r^{C-\ell} G(\ell)$$

and provides (9) after using definition (8) for $M_r(C)$. Besides, letting $\mathcal{M}$ denote a miss event, the Bayes formula entails

$$q_r(2) = \mathbb{P}(R = r \mid \mathcal{M}) = \mathbb{P}(R = r) \frac{\mathbb{P}(\mathcal{M} \mid R = r)}{\mathbb{P}(\mathcal{M})} = q_r \frac{M_r(C)}{M(C)}$$

hence relation (10). □

If the popularity distribution has unbounded support, then $\lim_{r \uparrow +\infty} q_r = 0$ and formula (9) implies that the per-object probability $M_r(C)$ tends to 1 as $r \uparrow +\infty$ for fixed C; (10) consequently entails

$$q_r(2) \sim \frac{q_r}{M(C)}, \quad r \uparrow +\infty. \tag{12}$$

For given C, asymptotic (12) shows that the tail of distribution $(q_r(2))_{r\in\mathbb{N}}$ at infinity is proportional to that of distribution $(q_r)_{r\in\mathbb{N}}$. The distribution $(q_r(2))_{r\in\mathbb{N}}$ describes the output process of the single cache, which process is generated by consecutive missed requests. It will serve as an essential ingredient to the further extension of the single cache model to network cache configurations considered in Section 5.

3.2. Some exact results

Coefficients $G(C)$, $C \geq 0$, and associated miss probability $M(C)$ can be explicitly derived for some specific popularity distributions. In the following, the total population N of objects is always assumed to be infinite.

Corollary 3.4. *Assume a geometric popularity distribution* $q_r = (1-\kappa)\kappa^{r-1}$, $r \geq 1$, *with given* $\kappa \in]0,1[$. *For all* $C \geq 0$, *the miss rate equals*

$$M(C) = \frac{1-\kappa}{1-\kappa^{C+1}}(C+1)\kappa^C. \tag{13}$$

Proof. Using Lemma 3.2, F is readily shown to verify the functional identity $F(z) = (1+(1-\kappa)z)F(\kappa z)$ for all $z \in \mathbb{C}$. Expanding each side of that identity in power series of z and identifying identical powers provides the value of the ratio $G(C+1)/G(C)$, hence result (13) by (4). □

Let us now assume that the popularity distribution follows a Zipf distribution defined by

$$q_r = \frac{A}{r^\alpha}, \quad r \geq 1, \tag{14}$$

with exponent $\alpha > 1$ and normalization constant $A = 1/\zeta(\alpha)$, where ζ is the Riemann's Zeta function. We now show that explicit rational expressions for miss rate $M(C)$ can be obtained for some integer values of α.

Corollary 3.5. *Assume a Zipf popularity distribution with exponent* α. *The miss probability then equals*

$$M(C) = \begin{cases} \dfrac{3}{2C+3} & \text{if} \quad \alpha = 2, \\[2ex] \dfrac{45}{(4C+3)(4C+5)(2C+3)} & \text{if} \quad \alpha = 4, \\[2ex] \dfrac{840}{(6C+7)(6C+5)(3C+4)(3C+2)(2C+3)} & \text{if} \quad \alpha = 6 \end{cases}$$

for all $C \geq 0$.

Proof. When $\alpha = 2$, the normalization constant equals $A = 1/\zeta(2) = 6/\pi^2$. Using the infinite product formula ([33], p.85, formula 4.5.68)

$$F(z) = \prod_{j\geq 1}\left(1 + \frac{u^2}{\pi^2 j^2}\right) = \frac{\sinh u}{u}$$

and expanding the left hand side into powers of $u^2 = A\pi^2 z$ gives the expansion $F(z) = \Sigma_{C \geq 0} G(C) z^C$ where $G(C) = (\pi^2 A)^C/(2C+1)!$ for all $C \geq 0$. Computing ratio (4) with the above expression of $G(C)$ then provides the miss probability $M(C) = 3/(2C+3)$, as claimed.

The subsequent cases follow a similar derivation pattern: for $\alpha = 4$ (resp. $\alpha = 6$), the infinite product $F(z) = f(u)$ is given a simple closed form as a finite product of functions $u \mapsto \sinh(\omega u)/u$ where $\omega^2 = -1$ (resp. $\omega^3 = -1$) and $\Re(\omega) > 0$ and with variable change $u^4 = A\pi^4 z$ (resp. $u^6 = A\pi^6 z$). Expanding then $f(u)$ in power series of z gives a rational expression for coefficient $G(C)$, from which the rational expression for miss probability $M(C)$ is derived. □

Rational expressions of Corollary 3.5 do not seem, however, to generalize for integer values $\alpha = 2p$ with $p \geq 4$; upper bounds can nevertheless be envisaged and are the object of further study.

As also suggested by Corollary 3.5, the cache size corresponding to a target miss probability should be a decreasing function of α. This property is generalized in **Section 3.3** where an asymptotic evaluation of $M(C)$ is provided for large cache size C and any Zipf distribution with real exponent $\alpha > 1$.

3.3. Large cache approximation

The specific popularity distributions considered in Corollaries 3.4 and 3.5 show that $M(C)$ is of order Cq_C for large C. In the present section, we derive general asymptotics for probabilities $M(C)$ and $M_r(C)$ with large cache size and show that such a magnitude order for $M(C)$ is generally valid for any Zipf popularity distribution with exponent $\alpha > 1$.

We first start by formulating a general Large Deviations result for evaluating coefficients $G(C)$ for large C (Theorem 3.6 below). The derivation of that result is essentially based on introducing random variables X_C, $C \geq 0$, related to the sequence $(G(C))_{C \geq 0}$ through their "shifted" distribution

$$\mathbb{P}(X_C = x) = \frac{G(x)}{F(\theta_C)} \theta_C^x, \quad x \geq 0,$$

for some relevant argument θ_C; using the local estimates stated in [34] for densities of arbitrary sequences of random variables then provides general estimate (18). To apply that estimate to the Zipf distribution (14), we then state two preliminary results (Lemmas 3.7 and 3.8) on the behavior of the associated generating function F at infinity. This finally enables us to claim our central result (Proposition 3.9) for the behavior of $M(C)$ for large C.

Theorem 3.6. *(See Proof in* ***Appendix A****)*

(i) Given the generating function F defined in (6) and $C \geq 0$, equation

$$zF'(z) = CF(z) \tag{15}$$

has a unique real positive solution $z = \theta_C$;

(ii) Assume there exists some constant $\sigma > 0$ such that the limit

$$\lim_{C \uparrow +\infty} e^{s\sqrt{C}} \frac{F(\theta_C e^{-s/\sqrt{C}})}{F(\theta_C)} = e^{\sigma^2 s^2/2} \tag{16}$$

holds for any given $s \in \mathbb{C}$ *with* $\Re(s) = 0$ *and, given any* $\delta > 0$, *there exists* $\eta \in]0,1[$ *and an integer* C_δ *such that*

$$\sup_{\delta \leq |y| \leq \pi} \left| \frac{F(\theta_C e^{iy})}{F(\theta_C)} \right|^{1/C} \leq \eta \tag{17}$$

for $C \geq C_\delta$. *We then have*

$$G(C) \sim \frac{\exp(H_C)}{\sigma\sqrt{2\pi C}} \tag{18}$$

as C *tends to infinity, with* $H_C = \log F(\theta_C) - C \log \theta_C$.

Following Theorem 3.6, the asymptotic behavior of $M(C)$ can then be derived from (18) together with identity (4). This approach is now applied to the Zipf popularity distribution (14); in this aim, the behavior of the associated generating function F is specified as follows.

Lemma 3.7. *(See Proof in* ***Appendix B****)*
For $\alpha > 1$ *and large* $z \in \mathbb{C} \setminus \mathbb{R}^-$, $\log F(z)$ *expands as*

$$\log F(z) = \alpha(\rho_\alpha A z)^{1/\alpha} - \frac{1}{2}\log(Az) + S_\alpha + o(1) \tag{19}$$

where $A = 1/\zeta(\alpha)$, *with constants*

$$\rho_\alpha = \left(\frac{\pi/\alpha}{\sin(\pi/\alpha)} \right)^\alpha \tag{20}$$

and S_α *depending on* α *only.*

The next lemma specifies in turn the behaviour of the solution θ_C to equation (15) for large C.

Lemma 3.8. *(See Proof in* ***Appendix C****)*
For $\alpha > 1$ *and large* C, *the unique real positive solution* θ_C *to (15) verifies*

$$\theta_C = \frac{C^\alpha}{A\rho_\alpha} + C^{\alpha-1} r_C \tag{21}$$

with $r_C = A_1 + O(C^{-1})$ *with* $A_1 = \alpha/2\rho_\alpha A$ *if* $\alpha \neq 2$, *and* $r_C = O(\log C)$ *if* $\alpha = 2$.

We can now state our central result.

Proposition 3.9. *For a Zipf popularity distribution with exponent* $\alpha > 1$, *the miss probability* $M(C)$ *is asymptotic to*

$$M(C) \sim \frac{A\rho_\alpha}{C^{\alpha-1}} \tag{22}$$

for large C, *with prefactor* ρ_α *defined in (20).*

Proof. As verified in **Appendix D**, the conditions of Theorem 3.6 are satisfied for a Zipf distribution. Using asymptotics (19) and (21) of Lemmas 3.7 and 3.8 to make the argument H_C explicit in (18), we then have

$$\begin{aligned} H_C &= \log F(\theta_C) - C\log\theta_C \\ &= \alpha(\rho_\alpha A\theta_C)^{1/\alpha} - \frac{1}{2}\log(A\theta_C) + S_\alpha + o(1) - C\log\left[\frac{C^\alpha}{A\rho_\alpha} + C^{\alpha-1}r_C\right] \end{aligned}$$

so that $H_{C+1} - H_C = -\alpha\log C + k + o(1)$ with $k = \log(A\rho_\alpha)$. By definition (4) of $M(C)$, the latter estimates enable us to derive that

$$\begin{aligned} M(C) = (C+1)\frac{G(C+1)}{G(C)} &\sim C\exp[H_{C+1} - H_C] \\ \sim C\frac{e^k}{C^\alpha} &= \frac{A\rho_\alpha}{C^{\alpha-1}} \end{aligned}$$

as claimed. □

We therefore conclude with Proposition 3.9 that $M(C) = O(Cq_C)$ for any Zipf popularity with exponent $\alpha > 1$.

Remark 3.1. *Using elementary asymptotics for* $\sin(\pi/\alpha)$*, expression (20) for factor* ρ_α *readily shows that* $\lim_{\alpha\uparrow+\infty}\rho_\alpha = 1$ *and* $\rho_\alpha \sim 1/(\alpha-1)$ *as* $\alpha\downarrow 1$*, respectively. Note also that coefficient* $A\rho_\alpha = \rho_\alpha/\zeta(\alpha)$ *in (22) has bounded variations for* $\alpha > 1$ *as it is lower bounded by 1 (attained for* $\alpha = 1$ *and* $\alpha = +\infty$*) and upper bounded by its maximum* $1.503...$ *(attained for* $\alpha = 2.172...$*).*

Remark 3.2. *Proposition 3.9 also provides asymptotic (22) for* $M(C)$ *under the weaker assumption that the popularity distribution* $(q_r)_{r\geq 1}$ *has a heavy tail of order* $r^{-\alpha}$ *for large* r *and some* $\alpha > 1$*, without being precisely Zipf as in (14). In fact, all necessary properties for deriving Lemmas 3.7 and 3.8 are based on that tail behavior only.*

To close this section, we now address the asymptotic behavior of $M_r(C)$ defined in (8).

Proposition 3.10. *For any Zipf popularity distribution with exponent* $\alpha > 1$ *and given the object rank* r*, the per-object miss probability* $M_r(C)$ *is estimated by*

$$M_r(C) \sim \frac{\rho_\alpha r^\alpha}{C^\alpha + \rho_\alpha r^\alpha} \tag{23}$$

for large C*, with prefactor* ρ_α *defined in (20).*

Proof. The generating function F_r of the sequence $G_r(C)$, $C \geq 0$, being given by (11), apply Theorem 3.6 to estimate coefficients $G_r(C)$ for large C. Concerning condition *(i)*, the solution $\eta = \eta_C$ to equation $\eta F_r'(\eta) = CF_r(\eta)$ reduces to equation (15) for $\theta = \theta_C$ where the term $q_r\theta/(1+q_r\theta)$ has been suppressed; but suppressing that term does not modify the estimate for θ_C with large C so that

$\eta_C \sim \theta_C$. On the other hand, condition *(ii)* is readily verified by generating function F_r and we then obtain

$$G_r(C) \sim \frac{G(C)}{1 + q_r \theta_C}. \tag{24}$$

By Lemma 3.8, we have $\theta_C \sim C^\alpha / A\rho_\alpha$ for large C; definition (8) of $M_r(C)$ and estimate (24) with $q_r = A/r^\alpha$ give

$$M_r(C) = \frac{G_r(C)}{G(C)} \sim \frac{1}{1 + q_r \theta_C} \sim \frac{\rho_\alpha r^\alpha}{C^\alpha + \rho_\alpha r^\alpha} \tag{25}$$

and result (23) follows. □

For any value $\alpha > 1$, (23) is consistent with the fact that $M_r(C)$ is an increasing function of object rank r and a decreasing function of cache size C.

3.4. Comparing RND to LRU

Let us now compare the latter results with the LRU replacement policy investigated in [15, 16]. Recall that, for a Zipf popularity distribution with exponent $\alpha > 1$, the miss probability $M(C)$ for LRU can be estimated by $M(C) \sim A\lambda_\alpha / C^{\alpha-1}$ for large C ([16], Theorem 3), with

$$\lambda_\alpha = \frac{1}{\alpha}\left[\Gamma\left(1 - \frac{1}{\alpha}\right)\right]^\alpha \tag{26}$$

where Γ is Gamma function; prefactor λ_α can also be estimated by $\lambda_\alpha \sim e^\gamma/\alpha$ as $\alpha \uparrow +\infty$ (where $\gamma \approx 0.57721$ denotes Euler's constant with $e^\gamma = 1,781...$) and $\lambda_\alpha \sim 1/(\alpha - 1)$ as $\alpha \downarrow 1$. From functional properties of the Γ function, it can be shown (see **Appendix E**) that prefactors ρ_α and λ_α respectively associated with RND and LRU performance are such that

$$\forall\, \alpha > 1, \quad \rho_\alpha > \lambda_\alpha, \tag{27}$$

thus confirming the fact that LRU discipline performs better than RND in general. More specifically, Remark 3.1 and the latter estimates for λ_α show that the difference $\rho_\alpha - \lambda_\alpha$ tends to 1 as $\alpha \uparrow +\infty$. This difference decreases, however, for smaller values of α since ρ_α and λ_α behave similarly as α is close to 1 (see Figure 1); in fact, we can easily show that

$$\rho_\alpha = \frac{e^{-h \log h}}{h} + O(h), \quad \lambda_\alpha = \frac{e^{-h \log h}}{h} - \gamma + O(h \log h)$$

for small $h = \alpha - 1$ and where γ is Euler's constant, so that $\lim_{\alpha \downarrow 1}(\rho_\alpha - \lambda_\alpha) = \gamma$. For small enough values of α, say $1 < \alpha \leq 2$, we can therefore consider that both disciplines provide essentially similar performance levels in terms of miss probabilities for large cache sizes.

In contrast to heavy-tailed popularity distributions, we can consider a light-tailed distribution where $q_r = A\exp(-Br^\beta)$, $r \geq 1$, with positive parameters

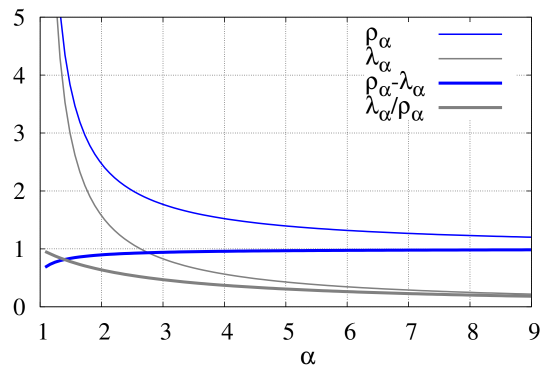

Figure 1: *Respective prefactors ρ_α and λ_α for RND and LRU policies, and variations of $\rho_\alpha - \lambda_\alpha$ and $\lambda_\alpha/\rho_\alpha$.*

A, B, β. It is shown in this case [16] that the miss probability for LRU is asymptotic to

$$M(C) \sim \frac{e^\gamma}{\beta B} C^{1-\beta} q_C$$

for large C. For a geometric popularity distribution (with $\beta = 1$), the latter estimate shows that $M(C) = O(q_C)$; on the other hand, formula (13) of Corollary 3.4 shows that $M(C) = O(Cq_C)$ for RND discipline. This illustrates the fact that RND and LRU replacements provide significantly different performance levels if the popularity distribution is highly concentrated on a relatively small number of objects.

4. Numerical results: single cache

In this section, we present numerical and simulation results to validate the preceding estimates for a single RND policy cache. In the following, when considering a finite object population with total size N, the Zipf popularity distribution is normalized accordingly. We also mention that the content popularity distribution obviously refers to document classes instead of individual documents. For comparison purpose with the existing LRU analysis, we represent these classes by a single index, as if they were a single document. In the following, cache sizes must accordingly be scaled up to the typical class size.

Simulations are performed using CCNPL-sim simulator [35]. In every simulation run, performance measures are collected once the system has reached the stationary state, thus ignoring the transient behavior.

The most critical parameter in our simulation setting is the numerical value of α. As the Zipf distribution flattens when α gets closer to 1, much longer simulation runs are necessary to have good estimates of the miss rate. Small enough values of α must, nevertheless, be considered as they are more realistic.

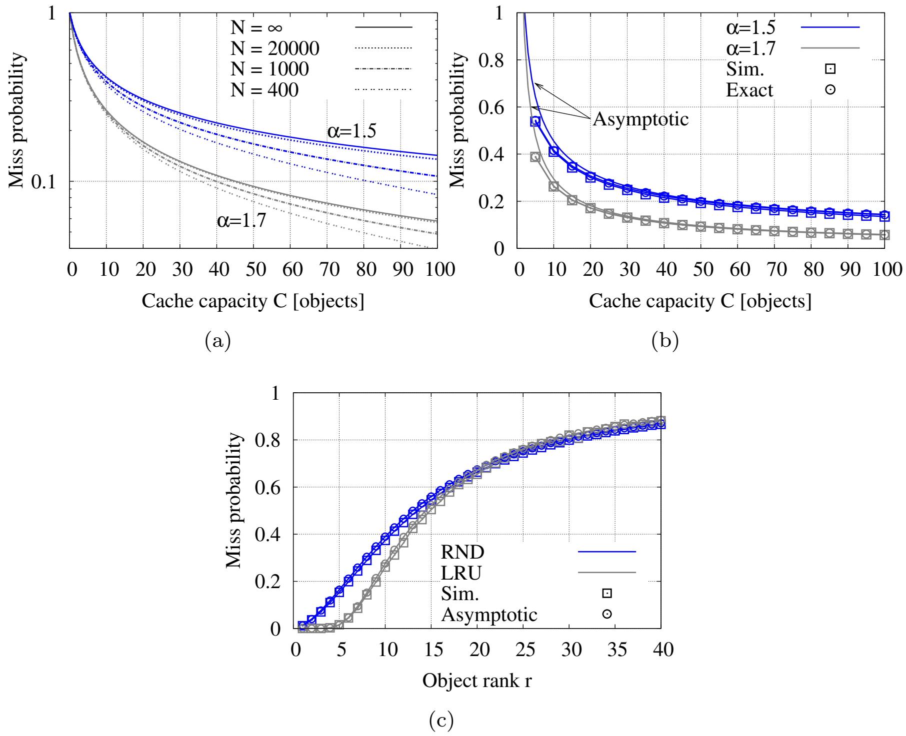

Figure 2: Single cache results: (a) exact formula for $M(C)$ with RND policy (b) asymptotic of $M(C)$ with RND policy (c) asymptotic of $M_r(C)$ with $C = 25$, $N = 20\ 000$ $\alpha = 1.7$for RND and LRU policies.

Estimates of α have been reported in [10] for web sites providing access to video content like `www.metacafe.com` for which $\alpha = 1.43$, `www.dailymotion.com` and `www.veoh.com` for which $\alpha = 1.72$ and $\alpha = 1.76$, respectively. In the following, we hence fix $\alpha = 1.5$ or $\alpha = 1.7$ in our numerical experiments.

Fig. 2(a) first reports exact formula (3) for $M(C) = M(C; N)$ as a function of cache size C and for increasing values of total population N, where $M(C; N)$ measures the total miss probability for a cache of size C when the number of objects N is finite. As expected, the convergence speed of $M(C; N)$ to $M(C; \infty)$ as $N \uparrow +\infty$ increases with α. In the case when $\alpha = 1.5$ for instance, a population of $N = 20\ 000$ can be considered a good approximation for an infinite object population ($N = \infty$), while there is almost no difference between values $M(C; N = 20\ 000)$ and $M(C; N = \infty)$ when $\alpha = 1.7$.

In Fig. 2(b), we compare the exact formula (3) for $M(C)$, asymptotic (22) and simulation results for the above scenario. Formula (3) for $N = \infty$ is computed with arbitrary precision and we used $N = 20\ 000$ for simulation as a good approximation for an infinite object population. Simulation and exact results are very close (especially for $\alpha = 1.7$), while asymptotic (22) gives a very good estimation of the miss probability as soon as cache size C is above 20.

Fig. 2(c) presents the miss probability $M_r(C)$ as a function of the object rank r for both RND and LRU policies with fixed $C = 25$, $N = 20\ 000$ and $\alpha = 1.7$. Results are reported for the most popular classes and confirm the asymptotic accuracy of estimate (23) for RND and the corresponding one for LRU policy [29]. Beside the good approximation provided by the asymptotics, it is important to remark that RND and LRU performance are very close when object rank $r \geq 15$, while there is a slight difference for the most popular objects (say $r < 15$). Moreover, comparing $M(C = 25)$ for RND and LRU (respectively equal to 0.147 and 0.108), we observe only a 4% smaller miss probability using LRU with respect to RND policy. This may suggest RND as a good candidate for caches working at very high speed, where LRU may become too expensive in terms of computation due to its relative complexity.

5. In-network cache model

In order to generalize the single-cache model, networks of caches with various topologies can be considered.

5.1. Line topology

We first consider the tandem system defined as follows. Any request is addressed to a first cache $\sharp 1$ with size C_1; if it is not satisfied, it is addressed to a second cache $\sharp 2$ with size C_2:

- if this request is satisfied at cache $\sharp 2$, the object is copied to cache $\sharp 1$, with replacement performed according to the RND discipline;

- if this request is not satisfied at cache $\sharp 2$, the object is retrieved from a repository server and copied in caches $\sharp 1$ **and** $\sharp 2$ according to the RND discipline. Note that this replacement scheme, hereafter denoted by **IPC** for **In-Path Caching**, ignores any collaboration between the two caches and blindly copies objects in all caches along the path towards the requesting source.

We now fix some notation and properties for the above defined tandem model. Let $R_1(t) \in \{1, 2, ..., N\}$ denote the object requested at cache $\sharp 1$ at time t; we still assume that variables $R_1(t)$, $t \in \mathbb{N}$, describe an IRM process with distribution defined by $\mathbb{P}(R_1 = r) = q_r$, $1 \leq r \leq N$. Denoting by $\mathbf{S}_1(t)$ (resp. $\mathbf{S}_2(t)$) the state vector of cache $\sharp 1$ (resp. $\sharp 2$) at time t, the bivariate process $(\mathbf{S}_1(t), \mathbf{S}_2(t))_{t \in \mathbb{N}}$ is easily shown to define a Markov process that, however, is not reversible. It is therefore unlikely that the stationary distribution of process $(\mathbf{S}_1, \mathbf{S}_2)$ can be derived in a simple product-form.

Alternatively, we here follow an approach based on the approximation of the request process to cache $\sharp 2$. Let t_n, $n \in \mathbb{N}$, denote the successive instants when a miss occurs at first cache $\sharp 1$, and $R_2(n)$ be the object corresponding to that miss event at time t_n. First note that the common distribution of variables $R_2(n)$ is the stationary distribution $(q_r(2))_{r \in \mathbb{N}}$ introduced in Lemma 3.3, equation (10), with cache size C replaced by C_1. In the following, we will further assume that

(H) the request process for cache $\sharp 2$ is an IRM, that is, all variables $R_2(n)$, $n \in \mathbb{N}$, are independent with common distribution

$$\mathbb{P}(R_2 = r) = q_r(2), \quad r \in \mathbb{N}.$$

The simplifying assumption **(H)** neglects any correlation structure for the output process of cache $\sharp 1$ (that is, the input to cache $\sharp 2$) produced by consecutive missed requests. Recall also that the tail of distribution $(q_r(2))_{r\in\mathbb{N}}$, defined by (12), is proportional to that of distribution $(q_r)_{r\in\mathbb{N}}$.

The latter two-stage tandem model can be easily extended to a tandem network consisting in a series of K caches ($K > 2$) where any request dismissed at caches $\sharp 1, ..., \sharp \ell$, $\ell \geq 1$, is addressed to cache $\sharp(\ell+1)$. As an immediate generalization of the **IPC** scheme, we assume that any requested document which experiences a miss at cache $\sharp j$, $1 \leq j \leq \ell$, and an object hit at cache $\sharp(\ell+1)$ is copied backwards at all downstream caches $\sharp 1, ..., \sharp \ell$. A request miss therefore corresponds to a miss event at each cache $1, 2, ..., K$. Furthermore, assumption **(H)** is generalized by saying that any cache $\sharp \ell$ considered in isolation behaves as a single cache with IRM input produced by consecutive missed requests at cache $\sharp(\ell - 1)$. The size of cache $\sharp \ell$ is denoted by C_ℓ.

In the following, the "global" miss probability $M_r(C_1, \ldots, C_\ell)$ (resp. "local" miss probability $M_r^*(C_1, \ldots, C_\ell)$) for request r at cache ℓ is the miss probability for object r over all caches $1, ..., \ell$ (resp. the miss probability for object r at cache ℓ) so that

$$M_r(C_1, \ldots, C_\ell) = \prod_{j=1}^{\ell} M_r^*(C_1, \ldots, C_j) \tag{28}$$

(note that for a single cache, we have $M_r(C_1) = M_r^*(C_1)$). To simplify notation, we abusively write $M_r(\ell)$ (resp. $M_r^*(\ell)$) instead of $M_r(C_1, \ldots, C_\ell)$ (resp. instead of $M_r^*(C_1, \ldots, C_\ell)$). Finally, if $q_r(\ell)$, $r \geq 1$, defines the distribution of the input process at cache $\sharp \ell$, the averaged local miss probability $M^*(\ell)$ at cache $\sharp \ell$ is given by

$$M^*(\ell) = \sum_{r \geq 1} M_r^*(\ell) q_r(\ell) \tag{29}$$

for any $\ell \in \{1, ..., K\}$.

Proposition 5.1. *(See Proof in **Appendix F**)*

*For the K-cache tandem system with **IPC** scheme, suppose that the request process at cache $\sharp 1$ is IRM with Zipf popularity distribution with exponent $\alpha > 1$, and that assumption **(H)** holds for all caches $\sharp 2, ..., \sharp K$.*

For any $\ell \in \{1, ..., K\}$ and large cache sizes $C_1, ..., C_\ell$ [1]*, the global miss*

[1]For $\ell > 1$, a formal definition of these large sizes is to set the scaling $C_j = c_j x$ for all $j \in \{1, ..., \ell\}$ with bounded c_j, and let the parameter x tend to infinity.

probability $M_r(\ell)$ (resp. local miss probability $M_r^*(\ell)$) is given by

$$M_r(\ell) \sim \frac{\rho_\alpha r^\alpha}{\rho_\alpha r^\alpha + \sum_{j=1}^{\ell} C_j^\alpha}, \quad M_r^*(\ell) \sim \frac{\rho_\alpha r^\alpha + \sum_{j=1}^{\ell-1} C_j^\alpha}{\rho_\alpha r^\alpha + \sum_{j=1}^{\ell} C_j^\alpha}. \tag{30}$$

The average global miss probability $M(\ell) = \Sigma_{r\geq 1} M_r(\ell) q_r$ for all objects requested along the cache network is given by

$$M(\ell) \sim \frac{A\rho_\alpha}{\left(\sum_{j=1}^{\ell} C_j^\alpha\right)^{1-\frac{1}{\alpha}}} \tag{31}$$

for any $\ell \in \{1, ..., K\}$ and large cache sizes C_1, ..., C_ℓ.

Proposition 5.1 shows how the K-stage tandem system with **IPC** scheme improves the performance in terms of miss probability by adding a term C_j^α when the j-th cache is added to the path.

5.2. Tree topology

The previous linear network model can be easily extended to the homogeneous tree topology with Zipf distributed requests. By homogeneous, we mean that all leaves of the tree are located at a common depth of the root, and that the cache size for each node at a given level i is equal to C_i (where C_1 is the cache size of the leaves). An example of such a tree is a complete binary tree of given height.

Let $\Lambda_1, \dots, \Lambda_J$ be the J leaves of the tree. We assume that all requests arrive at the leaves, following an IRM, that is, $\mathbb{P}(R(t) = r, \Lambda(t) = j) = p_j q_r$ for all $1 \leq r \leq N$, $1 \leq j \leq J$, where $(p_1, \dots, p_J)$ are positive values such that $\Sigma_{1\leq j\leq J}\, p_j = 1$ and $\Lambda(t)$ denotes the leaf where the request t arrives at time t. Requests are served according to the IPC rule, *i.e.*, are forwarded upwards until the content is found, and the content is then copied in each cache between this location and the addressed leaf.

Corollary 5.2. *Consider a homogeneous tree with **IPC** scheme and suppose that assumption **(H)** holds for all its internal nodes. The results of Proposition 5.1 then extend to that tree with IRM request process at leaves and Zipf popularity distribution with exponent $\alpha > 1$.*

Proof. Only the order of requests in time matters since their precise timing is irrelevant; we can consequently assume that the requests arrive according to a Poisson process with intensity 1. From the property of independent thinning and merging of Poisson processes, it follows that the requests for a given object

r at leaf j is also a Poisson process with intensity $p_j q_r$ (with a Zipf distribution $q_r = A/r^\alpha$, $r \geq 1$) and that the request process at leaf j is a Poisson process with intensity p_j. Now, using assumption **(H)** and applying the previous results for a single cache to each leaf, we deduce that at any leaf j, the miss sequence for object r is a Poisson process with intensity $p_j q_r M_r^*(1)$. Merging these miss sequences from all children of a given second-level node, we deduce that the requests at this node build up a Poisson process and that the probability of requesting an object r is $q_r M_r^*(1)/M(1) = q_r(2)$. This process has the same properties as the IRM process with distribution $(q_r(2))_{r\in\mathbb{N}}$ used in the proof of Proposition 5.1, which therefore applies. Repeating recursively this reasoning at each level, we conclude that Proposition 5.1 holds in this context. □

Remark 5.1. *Corollary 5.2 is also valid for a homogeneous tree where different replacement policies are used at different levels i (e.g. Random at first level and LRU at second one).*

6. Numerical results: network of caches

In this section, we report numerical and simulation results to show the accuracy of the approximations presented in Section 5.

Fig. 3(a) first depicts estimate (30) of $M_r^*(1)$ and $M_r^*(2)$ for both RND and LRU with $C_1 = C_2 = 25$, $N = 20\ 000$ and $\alpha = 1.7$ (the approximation for the tandem LRU is taken from [26]). We focus on the second cache, as the performance of the first one has been analyzed in previous sections. We note a good agreement between the approximations given in Section 5 and simulation results.

Moreover, while less popular objects are affected in the same way when employing either RND or LRU (in our specific example, $r \geq 15$), a significantly different behavior is detectable for popular objects ($r < 15$). Local miss probabilities $M_r^*(1)$ and $M_r^*(2)$ help understanding where an object has been cached, conditioned on its rank. The combination of LRU and IPC clearly tends to favor stationary configurations where popular objects are likely to be stored in the first cache (see [26] for a similar discussion). When using RND instead of LRU, however, the distribution of the content across the two caches is fairly different; as illustrated in Fig. 3(a), while the most popular objects are likely to be retrieved at the first cache when using either LRU or RND, such objects can also be found in the second cache only when using RND. It therefore appears that while both LRU and RND tend to store objects proportionally to their popularity, RND more evenly distributes objects across the whole path.

Fig. 3(b) reports the global miss probability $M_r(2) = M_r^*(1)M_r^*(2)$ at the second cache (the probability to query an object of rank r at the repository server) with the same setting than that described for Fig. 3(a). In this example, we see that objects with rank r less than about 15 are slightly more frequently requested at the server when using RND rather than LRU, but RND is more favorable than LRU for objects with higher rank $r \geq 15$. In average, the global miss probability at the second cache $M(2)$, reported in Fig. 3(c) for $C_1 = C_2 \leq$

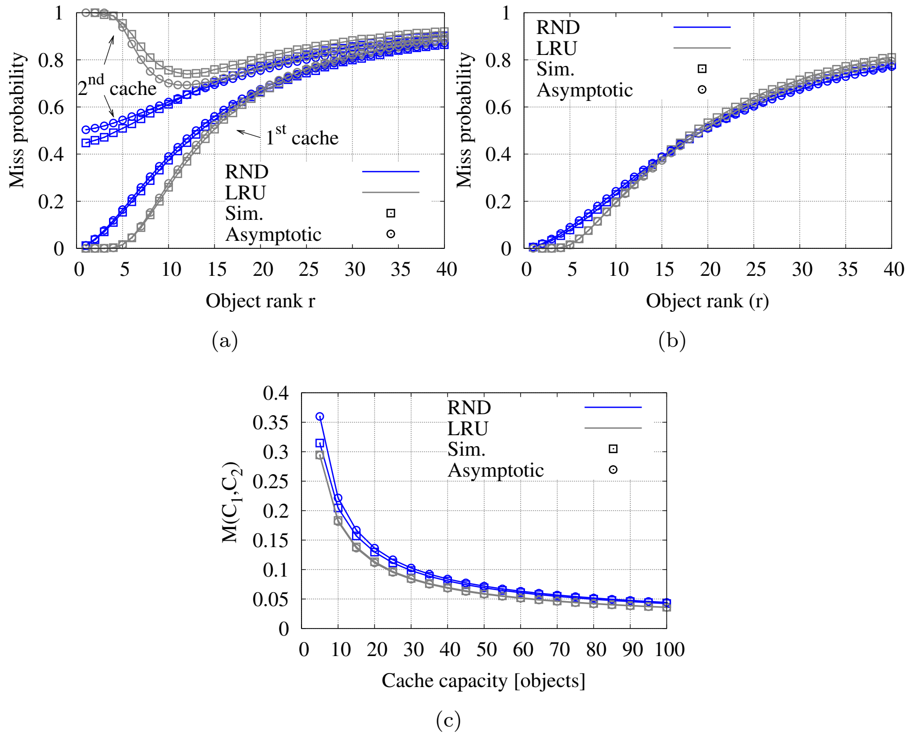

Figure 3: Tandem cache results: (a) asymptotics of $M_r^*(1)$, $M_r^*(2)$ (b) and of $M_r(2)$ for RND and LRU policies compared to simulation with $C_1 = C_2 = 25$, $N = 20\ 000$, $\alpha = 1.7$ (c) asymptotic for $M(2) = M(C_1, C_2)$ with $C_1 = C_2 \leq 100$, $N = 20\ 000$, $\alpha = 1.7$.

100, $N = 20\ 000$ and $\alpha = 1.7$, is very similar using either RND or LRU, with a slight advantage to LRU. The average miss probability $M(2)$ indicates the amount of data that is to be requested from the server.

We finally observe that the approximations calculated in Section 5 for RND prove very accurate (as in [27], [26] for LRU). Furthermore, these approximations work well in a large number of scenarios that we do not report here for the sake of conciseness.

7. Mixture of RND and LRU

We have considered so far networks of caches where all caches use the RND replacement policy. In practice, it is feasible to use different replacement algorithms in the same network. This section addresses the case of a two-cache tandem network, where one cache uses the RND replacement algorithm while the other uses the LRU algorithm. As in Section 5.2, these results also hold in the case of an homogeneous tree.

7.1. Large cache size estimates

We provide estimates for miss probabilities in the case when cache sizes C_1 and C_2 are large.

Proposition 7.1. *For the two-cache tandem system with **IPC** scheme, suppose the request process at cache* $\sharp 1$ *is IRM with Zipf popularity distribution with exponent* $\alpha > 1$ *and that assumption **(H)** for cache* $\sharp 2$ *holds.*

I) *When cache* $\sharp 1$ *(resp. cache* $\sharp 2$*) uses the RND (resp. LRU) replacement policy, the global (resp. local) miss probability* $M_r(2)$ *(resp.* $M_r^*(2)$*) on cache* $\sharp 2$ *is given by*

$$\left\{ \begin{array}{ll} M_r(2) & \sim \dfrac{\rho_\alpha r^\alpha}{\rho_\alpha r^\alpha + C_1^\alpha} \exp\left(-\dfrac{\rho_\alpha C_2^\alpha}{\alpha \lambda_\alpha \left(\rho_\alpha r^\alpha + C_1^\alpha \right)} \right), \\ M_r^*(2) & \sim \exp\left(-\dfrac{\rho_\alpha C_2^\alpha}{\alpha \lambda_\alpha \left(\rho_\alpha r^\alpha + C_1^\alpha \right)} \right) \end{array} \right. \tag{32}$$

for large cache sizes C_1*,* C_2 *and constants* ρ_α*,* λ_α *introduced in (20) and (26), respectively.*

II) *When cache* $\sharp 1$ *(resp. cache* $\sharp 2$*) uses the LRU (resp. RND) replacement policy, the global (resp. local) miss probability* $M_r(2)$ *(resp.* $M_r^*(2)$*) on cache* $\sharp 2$ *is given by*

$$\left\{ \begin{array}{ll} M_r(2) & \sim \dfrac{\rho_\alpha r^\alpha}{\rho_\alpha r^\alpha \exp\left(\dfrac{C_1^\alpha}{\alpha \lambda_\alpha r^\alpha} \right) + C_2^\alpha}, \\ M_r^*(2) & \sim \dfrac{\rho_\alpha r^\alpha}{\rho_\alpha r^\alpha + C_2^\alpha \exp\left(-\dfrac{C_1^\alpha}{\alpha \lambda_\alpha r^\alpha} \right)} \end{array} \right. \tag{33}$$

for large cache sizes C_1*,* C_2*.*

Proof. We follow the same derivation pattern as the proof of Proposition 5.1 detailed in **Appendix F**.

I) When cache $\sharp 1$ uses the RND replacement policy, we know from **Appendix F** that the request process at cache $\sharp 2$ is IRM with popularity distribution

$$q_r(2) = q_r \frac{M_r^*(1)}{M^*(1)} \sim \frac{C_1^{\alpha - 1}}{C_1^\alpha + \rho_\alpha r^\alpha}, \quad r \geq 1,$$

and is asymptotically Zipf for large r. We then follow the proof of Proposition 6.2 of [26]. Let $S_2(0, t)$ be the number of different objects requested at cache $\sharp 2$ in the time interval $[0, t]$; it verifies

$$\mathbb{E}\left[S_2(0, t) \right] = \sum_{r \geq 1} \left(1 - e^{-q_r(2) t} \right).$$

We then first deduce that

$$\mathbb{E}\left[S_2(0,t)\right] \geq \int_1^{+\infty} \left(1 - e^{-q_u(2)t}\right) \mathrm{d}u. \tag{34}$$

Using the variable change $v = C_1^{\alpha-1}t/(C_1^\alpha + \rho_\alpha u^\alpha)$ in the latter integral, we further obtain

$$\mathbb{E}\left[S_2(0,t)\right] \geq \left(\frac{C_1^{\alpha-1}t}{\rho_\alpha}\right)^{\frac{1}{\alpha}} \times \int_0^{\frac{C_1^{\alpha-1}t}{C_1^\alpha+\rho_\alpha}} \frac{1}{\alpha}\left(1-e^{-v}\right) v^{-1-\frac{1}{\alpha}} \left(1 - \frac{C_1 v}{t}\right)^{\frac{1}{\alpha}-1} \mathrm{d}v.$$

Letting $t \uparrow +\infty$, the monotone convergence theorem applied to the family of functions $v \mapsto (1 - C_1 v/t)^{-1+1/\alpha}$, $t > 0$, together with a further integration by parts yield

$$\lim_{t\uparrow+\infty} \frac{\mathbb{E}\left[S_2(0,t)\right]^\alpha}{t} \geq \left(\frac{C_1^{\alpha-1}}{\rho_\alpha}\right)\left[\Gamma\left(1-\frac{1}{\alpha}\right)\right]^\alpha. \tag{35}$$

Starting integral (34) from $u = 0$ instead of $u = 1$, the latter asymptotic bound is seen to hold also as an upper bound of $\mathbb{E}\left[S_2(0,t)\right]^\alpha/t$, thus showing that (35) actually holds as an equality. The local per-object miss rate on the second cache for an LRU cache is then

$$M_r^*(2) \sim \exp\left[-q_r(2)C_2^\alpha \left(\lim_{t\uparrow+\infty}\frac{\mathbb{E}\left[S_2(0,t)\right]^\alpha}{t}\right)^{-1}\right]$$

which proves expressions (32).

II) When cache ♯1 applies the LRU replacement policy, the local per-object miss rate at cache ♯1 is known asymptotically [29] to equal

$$M_r^*(1) \sim \exp\left[-\frac{C_1^\alpha}{r^\alpha\left[\Gamma\left(1-\frac{1}{\alpha}\right)\right]^\alpha}\right] = \exp\left[-\frac{C_1^\alpha}{\alpha\lambda_\alpha r^\alpha}\right]$$

and the local average miss rate is

$$M^*(1) \sim \frac{1}{\alpha}\left[\Gamma\left(1-\frac{1}{\alpha}\right)\right]^\alpha \frac{A}{C_1^{\alpha-1}} = \frac{\lambda_\alpha A}{C_1^{\alpha-1}}.$$

Using assumption **(H)**, it then follows that the input process at cache ♯2 is IRM with popularity distribution given by

$$q_r(2) = q_r\frac{M_r^*(1)}{M^*(1)} \sim \frac{C_1^{\alpha-1}}{\lambda_\alpha r^\alpha}\exp\left[-\frac{C_1^\alpha}{\alpha\lambda_\alpha r^\alpha}\right], \quad r \geq 1.$$

Note that this distribution is asymptotically Zipf, that is, $q_r(2) \sim A'(2)/r^\alpha$ for large r with coefficient $A'(2) = A/M^*(1)$. Applying estimate (25) to the above defined distribution $q_r(2)$, $r \geq 1$, it then follows that

$$M_r^*(2) \sim \frac{1}{1 + q_r(2)\theta'(2)}$$

where the associated root $\theta'(2)$ is easily estimated by $\theta'(2) = C_2^\alpha / A'(2)\rho_\alpha$ by using Lemma 3.8. We then derive that

$$M_r^*(2) \sim \left(1 + \frac{C_2^\alpha}{A'(2)\rho_\alpha}\frac{AM_r^*(1)}{r^\alpha M^*(1)}\right)^{-1}$$

which finally leads to expressions (33). □

7.2. Numerical results

We here report numerical and simulation results for mixed homogeneous tree topologies to show the accuracy of the approximations presented in Section 7.1, so as to derive some more general observations about the mixture of RND and LRU in a network of caches. Following Section 5.2, we actually simulate tree topologies with 2 leaves with cache size C_1 and one root with cache size C_2.

Fig. 4(a) reports $M_r^*(2)$ for RND-LRU and LRU-RND homogeneous tree networks with asymptotics (32) and (33), respectively. We first note that the

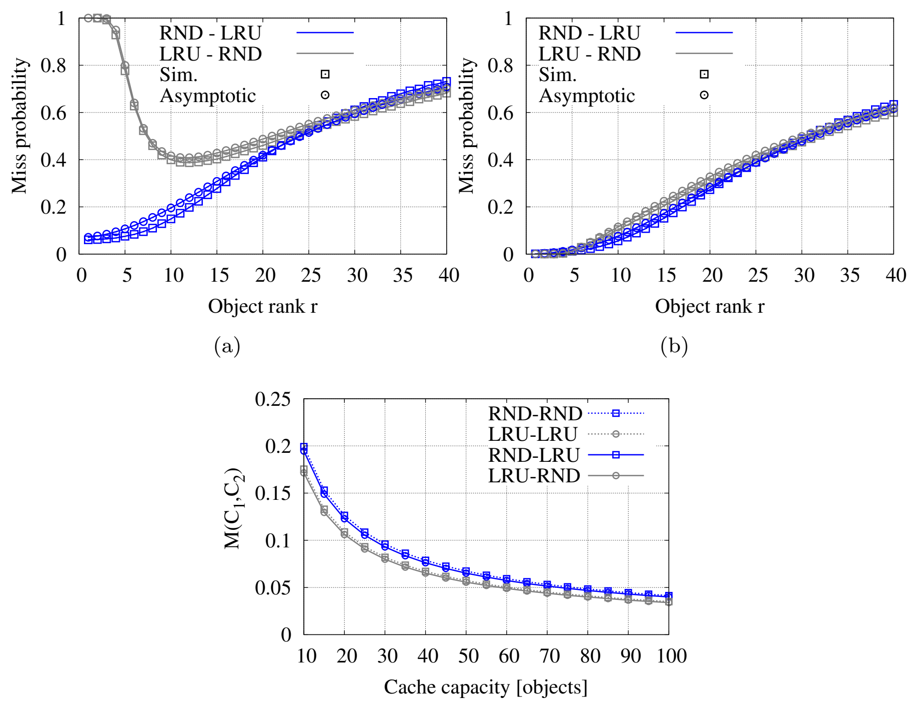

(a) (b) (c)

Figure 4: Mixed tree cache networks with $C_1 = 25, C_2 = 50$, $\alpha = 1.7$: (a) asymptotic of $M_r^*(2)$ (b) asymptotic of $M_r(2)$ (c) Simulation for $M(2) = M(C_1, C_2)$ with $C_1 = C_2 \leq 100$, $\alpha = 1.7$.

latter provide estimates with reasonable accuracy. Besides, we observe that the behavior of $M_r^*(2)$ is in strict relation to the policy used for cache $\sharp 1$. If cache $\sharp 1$ is RND, then $M_r^*(2)$ has a behavior similar to that observed in RND caches in tandem; similarly, if cache $\sharp 1$ is LRU, $M_r^*(2)$ behaves as in the case of LRU caches in tandem.

This phenomenon has a simple explanation. RND and LRU policies act similarly on objects ranked in the tail of the Zipf popularity distribution. However, the two replacement policies manage popular objects in a rather different way, as already observed in Section 6. The "local" popularity distribution seen by 2nd level caches results from the shaping around the mean of the popularity distribution given at 1st level, while the tail remains unchanged. The portion of the distribution that is affected by such a shaping process is determined by the cache size C_1 at first level. In the analysis reported in Fig. 4, the 1st level cache significantly determines the performance of the overall tandem system. In Fig. 4(b), we observe that the global miss probability $M_r(2)$ in the two mixed tandem caches is similar, while the distribution of the objects across the two nodes varies considerably.

Finally, Fig. 4(c) depicts simulation results for $M(2) = M(C_1, C_2)$, for all possible configurations in an homogeneous tree network where $C_1 = C_2 \leq 100$. We observe that the LRU-RND tree cache network achieves slightly better performance than the LRU-LRU system. This behaviour would suggest to prefer LRU at 1st level since it performs better in terms of miss probability, while using RND at 2nd level in order to save significant processing time.

8. Conclusion

The recent technological evolution of memory capacities, as illustrated by the deployment of CDNs and the proposition of new information-centric architectures where caching becomes an intrinsic network property, raise new interests on cache studies.

In this paper, we have studied the RND replacement policy where objects to be removed are chosen uniformly at random. Assuming that the content popularity follows a Zipf law with parameter $\alpha > 1$, and that request processes are IRM, we prove (Proposition 3.9) that the miss rate is asymptotically equivalent to $A\rho_\alpha / C^{\alpha-1}$ for large cache size C. This shows that the difference between LRU and RND caches is asymptotically independent of the cache size and depends on coefficient $A\rho_\alpha$ only. These results are extended to typical network topologies, namely tandems and homogeneous trees, under the approximation that request processes are IRM at any node. The case of mixed policies (RND at one network level and LRU at the other one) is also considered. Simulation results show that the IRM assumption applied to several network topologies is efficient and provides accurate estimates.

Our results suggest that the performance of RND is reasonably close to that of LRU in terms of average miss rate. As a consequence, RND is a good candidate for high-speed caching when the complexity of the replacement policy becomes critical. In the presence of a hierarchy of caches, caches at deep levels (*i.e.* access networks) typically serve a relatively small number of requests per second which can be easily sustained by a cache running LRU; LRU policy should consequently be implemented at the bottom level since it provides the best performance. Meanwhile, higher-level caches see many aggregated requests and should therefore use the RND policy which yields similar performance while being less computationally expensive. Recall finally that the FIFO policy, the alternative high-speed candidate, has theoretical performance similar to that of RND.

We have assumed in this paper that the parameter α of the Zipf distribution is larger than 1, and the total number of available objects is infinite. For further study, we first intend to explore the case $\alpha \leq 1$ with a finite number of objects. Besides, since Zipf popularity distributions do not represent all types of Internet traffic, we also intend to analyze the performance of RND caches when the popularity has a light-tailed (e.g. Weibull) distribution. Finally, all results derived in this paper hold for i.i.d. requests processes; actual traces show, however, that consecutive requests are correlated, both in time and space. The definition of an accurate and realistic model which can take these correlations into account, as well as the extension of the present results to such a request model, is also on our research agenda.

Acknowledgements

This work has been partially funded by the French National Research Agency (ANR), CONNECT project, under grant number ANR-10-VERS-001, and by the European FP7 IP project SAIL under grant number 257448. We thank Philippe Olivier for his useful comments.

References

[1] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, R. L. Braynard, Networking named content, in: Proc. of ACM CoNEXT, 2009.

[2] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, I. Stoica, A data-oriented (and beyond) network architecture, in: Proc. of ACM SIGCOMM, 2007.

[3] Network of Information (NetInf), `http://www.sail-project.eu/`.

[4] D. Perino, M. Varvello, A reality check for content-centric networking, in: Proc. of ACM SIGCOMM ICN, 2011.

[5] P. Barford, A. Bestavros, A. Bradley, M. Crovella, Changes in web client access patterns: Characteristics and caching implications, World Wide Web 2 (1999) 15–28.

[6] L. Breslau, P. Cao, L. Fan, G. Phillips, S. Shenker, Web caching and Zipf-like distributions: evidence and implications, in: Proc. of IEEE INFOCOM, 1999.

[7] L. Cherkasova, M. Gupta, Analysis of enterprise media server workloads: access patterns, locality, content evolution, and rates of change, IEEE/ACM Transactions on Networking 12 (5) (2004) 781 – 794.

[8] C. P. Costa, I. S. Cunha, A. Borges, C. V. Ramos, M. M. Rocha, J. M. Almeida, B. Ribeiro-Neto, Analyzing client interactivity in streaming media, in: Proc. of ACM WWW, 2004.

[9] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, S. Moon, I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system, in: Proc. of ACM IMC, 2007.

[10] S. Mitra, M. Agrawal, A. Yadav, N. Carlsson, D. Eager, A. Mahanti, Characterizing web-based video sharing workloads, ACM Transactions on the Web 5 (2011) 8:1–8:27.

[11] J. McCabe, On serial files with relocatable records, Operations Research 12 (1965) 609–618.

[12] J. Burville, J. F. C. Kingman, On a model for storage and search, Journal of Applied Probability 10 (1973) 697–701.

[13] W. J. Hendricks, The stationary distribution of an interesting Markov chain, Journal of Applied Probability 9 (1972) 231–233.

[14] J. A. Fill, L. Holst, On the distribution of search cost for the move-to-front rule, Random Structures and Algorithms 8 (3) (1996) 179–186.

[15] P. Flajolet, D. Gardy, L. Thimonier, Birthday paradox, coupon collectors, caching algorithms and self-organizing search, Discrete Applied Mathematics 39 (3) (1992) 207–229.

[16] P. R. Jelenković, Asymptotic approximation of the move-to-front search cost distribution and least-recently-used caching fault probabilities, The Annals of Applied Probability 9 (2) (1999) 430–464.

[17] E. G. Coffman, P. R. Jelenković, Performance of the move-to-front algorithm with Markov-modulated request sequences, Operation Research Letters 25 (3) (1999) 109–118.

[18] P. R. Jelenković, A. Radovanović, Least-recently-used caching with dependent requests, Theoretical Computer Science 326 (1–3) (2004) 293–327.

[19] P. R. Jelenković, A. Radovanović, Optimizing LRU caching for variable document sizes, Journal of Combinatorics, Probability and Computing 13 (4–5) (2004) 627–643.

[20] T. Osogami, A fluid limit for cache algorithms with general request processes., in: Proc. of IEEE INFOCOM, 2009.

[21] E. Gelenbe, A unified approach to the evaluation of a class of replacement algorithms, IEEE Transactions on Computers 22 (6) (1973) 611–618.

[22] N. Tsukada, R. Hirade, N. Miyoshi, Fluid limit analysis of FIFO and RR caching for independent reference models 69 (2012) 403–412.

[23] A. Dan, D. Towsley, An approximate analysis of the LRU and FIFO buffer replacement schemes, in: Proc. of ACM SIGMETRICS, 18, pp.143-152, 1990.

[24] C. Fricker, P. Robert, J. Roberts, A versatile and acurrate approximation for LRU cache performance, in: Proceedings of ITC24, 2012.

[25] E. J. Rosensweig, J. Kurose, D. Towsley, Approximate models for general cache networks, in: Proc. of IEEE INFOCOM, 2010.

[26] G. Carofiglio, M. Gallo, L. Muscariello, D. Perino, Modeling data transfer in content-centric networking, in: Proc. of ITC23, 2011.

[27] G. Carofiglio, M. Gallo, L. Muscariello, Bandwidth and storage sharing performance in information-centric networking, in: Proc. of ACM SIGCOMM ICN, 2011.

[28] H. Che, Y. Tung, Z. Wang, Hierarchical web caching systems: modeling, design and experimental results, Selected Areas in Communications, IEEE Journal on 20 (7) (2002) 1305–1314.

[29] P. R. Jelenković, X. Kang, Characterizing the miss sequence of the LRU cache, ACM SIGMETRICS Performance Evaluation Review 36 (2008) 119–121.

[30] N. Fofack, P. Nain, G. Neglia, D. Towsley, Analysis of ttl-based cache networks, in: Performance Evaluation Methodologies and Tools (VALUETOOLS), International Conference on, 2012.

[31] N. Laoutaris, H. Che, I. Stavrakakis, The LCD interconnection of LRU caches and its analysis, Performance Evaluation 63 (2006) 609–634.

[32] O. Babaoğlu, Hierarchical replacement decisions in hierarchical stores, in: Proc. of ACM SIGMETRICS, 1982.

[33] M. Abramowitz, I. Stegun, Handbook of Mathematical Functions with formulas, graphs and mathematical tables, 10th Edition, Dover, 1972.

[34] N. R. Chaganty, J. Sethuraman, Multidimensional large deviations local limit theorems, Journal of Multivariate analysis 20 (1986) 190–204.

[35] M. Gallo, L. Muscariello, M. Papalini, CCNPL-Sim: Content-Centric Networking Packet level simulator `http://code.google.com/p/ccnpl-sim/wiki/Installation`, 2012.

[36] M. M. Rao, R. J. Swift, Probability Theory with Applications, 2nd Edition, Springer, 2006.

[37] R. Godement, Analyse Mathématique II, 2nd Edition, Springer, 2003.

Appendix A. Proof of Theorem 3.6

(i) Using (6), equation (15) reduces to $g(z) = C$ where

$$g(z) = z\frac{F'(z)}{F(z)} = \sum_{j\geq 1} \frac{q_j z}{1+q_j z}. \tag{A.1}$$

Continuous function $g : z \in [0, +\infty[\rightarrow g(z) \in [0, +\infty[$ vanishes at $z = 0$, is strictly increasing on $[0, +\infty[$ and tends to $+\infty$ when $z \uparrow +\infty$. Equation (15) has consequently a unique positive solution θ_C. Note that θ_C tends to $+\infty$ as $C \uparrow +\infty$, since $g(z) \leq \Sigma_{j\geq 1}(q_j z) = z$, hence $\theta_C \geq C$.

(ii) Consider the random variable X_C with support $\mathbb{N}$ and distribution

$$\mathbb{P}(X_C = x) = \frac{G(x)}{F(\theta_C)}\theta_C^x, \quad x \geq 0, \tag{A.2}$$

where θ_C satisfies (15); note that definition (A.2) for X_C is equivalent to

$$G(x) = \frac{F(\theta_C)}{\theta_C^x}\mathbb{P}(X_C = x), \quad x \geq 0. \tag{A.3}$$

By definitions (5) and (A.2), the generating function of random variable X_C is $z \mapsto F(z\theta_C)/F(\theta_C)$; in view of (15), the expectation of variable X_C is then

$$\mathbb{E}(X_C) = \frac{\mathrm{d}}{\mathrm{d}z}\frac{F(z\theta_C)}{F(\theta_C)}\Big|_{z=1} = \theta_C\frac{F'(\theta_C)}{F(\theta_C)} = C$$

so that random variables $Y_C = (X_C - C)/\sqrt{C}$, $C \geq 0$, are all centered. Besides, the Laplace transform of variable Y_C is given by

$$\mathbb{E}(e^{-sY_C}) = e^{s\sqrt{C}}\mathbb{E}(e^{-sX_C/\sqrt{C}}) = e^{s\sqrt{C}}\frac{F(\theta_C e^{-s/\sqrt{C}})}{F(\theta_C)}$$

for all $s \in \mathbb{C}$. By Lévy's continuity theorem ([36], Theorem 4.2.4), assumption (16) entails that variables Y_C converge in distribution when $C \uparrow +\infty$ towards a centered Gaussian variable with variance σ^2; moreover, assumption (17) ensures that the conditions of Chaganty-Sethuraman's theorem ([34], Theorem 4.1) hold so that $\mathbb{P}(X_C = C) = \mathbb{P}(Y_C = 0)$ is asymptotic to

$$\mathbb{P}(X_C = C) \sim \frac{1}{\sigma\sqrt{2\pi C}} \tag{A.4}$$

as $C \uparrow +\infty$. Equation (A.3) for $x = C$ and asymptotic (A.4) together provide estimate (18) for $G(C)$. □

Appendix B. Proof of Lemma 3.7

Let $q(x) = A/x^\alpha$ and $f_z(x) = \log(1 + q(x)z)$ for any real $x \geq 1$ and $z \in \mathbb{C}$; definitions (6), (14) and the above notation then entail that

$$\log F(z) = \sum_{r \geq 1} f_z(r);$$

function $\log F$ is analytic in the domain $\mathbb{C} \setminus \mathbb{R}^-$. For given $z \in \mathbb{C} \setminus \mathbb{R}^-$ and integer $J \geq 1$, the Euler-Maclaurin summation formula ([37], Chap.VI, Sect.2, formula (16.4)) reads

$$\sum_{r=1}^{J} f_z(r) = \int_1^J f_z(x)\mathrm{d}x + \frac{1}{2}\left[f_z(J) + f_z(1)\right] + \frac{1}{12}\left[f_z'(J) - f_z'(1)\right] + \frac{T_z(J)}{6} \quad \text{(B.1)}$$

with

$$T_z(J) = \int_1^J B_3(\{x\}) f_z^{(3)}(x)\mathrm{d}x,$$

where $B_3(x) = x(x-1)(2x-1)/2$ is the third Bernoulli polynomial and $\{x\}$ denotes the fractional part of real x; derivatives of f_z are taken with respect to x. Consider the behavior of the r.h.s. of (B.1) as J tends to infinity. We first have $f_z(1) = \log(1 + Az)$ and $f_z(J) = O(J^{-\alpha})$ for large J; differentiation with respect to x entails

$$f_z'(1) = -\frac{\alpha A z}{1 + Az}$$

and $f_z'(J) = O(J^{-\alpha-1})$ for large J. Differentiating twice again with respect to x shows that the third derivative of f_z is $O(x^{-\alpha-3})$ for large positive x, and is consequently integrable at infinity. Letting J tend to infinity in (B.1) and using the above observations together with the boundedness of periodic function $x \geq 1 \mapsto B_3(\{x\})$, we obtain

$$\log F(z) = \int_1^{+\infty} f_z(x)\mathrm{d}x + \frac{1}{2}\log(1 + Az) + \frac{\alpha A z}{12(1 + Az)} + \frac{T_z}{6} \quad \text{(B.2)}$$

with remainder term

$$T_z = \int_1^{+\infty} B_3(\{x\}) f_z^{(3)}(x)\mathrm{d}x. \quad \text{(B.3)}$$

Now, considering the first integral in the r.h.s. of (B.2), the variable change $x = t(Az)^{1/\alpha}$ gives

$$\int_1^{+\infty} f_z(x)\mathrm{d}x = (Az)^{1/\alpha}\Big[L - \int_0^{1/(Az)^{1/\alpha}} \log\left(1 + \frac{1}{t^\alpha}\right)\mathrm{d}t\Big] \quad \text{(B.4)}$$

where L is the finite integral

$$L = \int_0^{+\infty} \log\left(1 + \frac{1}{t^\alpha}\right)\mathrm{d}t = \alpha \int_0^{+\infty} \frac{\mathrm{d}t}{1 + t^\alpha} = \alpha \frac{\pi/\alpha}{\sin(\pi/\alpha)} = \alpha \rho_\alpha^{1/\alpha} \quad \text{(B.5)}$$

(use an integration by parts and see ([33], p.256, formula 6.1.17) for the second to last equality) with ρ_α introduced in (20) for $\alpha > 1$, and where

$$\int_0^{1/(Az)^{1/\alpha}} \log\left(1+\frac{1}{t^\alpha}\right)\mathrm{d}t = \frac{\log(Az)}{(Az)^{1/\alpha}} + \frac{\alpha}{(Az)^{1/\alpha}} + o(1); \qquad \text{(B.6)}$$

gathering (B.5) and (B.6) in expression (B.4) then provides expansion

$$\int_1^{+\infty} f_z(x)\mathrm{d}x = L(Az)^{1/\alpha} - \log(Az) - \alpha + o(1). \qquad \text{(B.7)}$$

We finally show that remainder term (B.3) tends to a finite limit t_α when $z \uparrow +\infty$. In fact, making the third derivative $f_z^{(3)}$ explicit gives

$$T_z = \alpha A z\Big[\alpha(1+\alpha)U_z - 2V_z\Big] \qquad \text{(B.8)}$$

where

$$U_z = \int_1^{+\infty} B_3(\{x\})\frac{x^{\alpha-1}}{x^2(x^\alpha+Az)^2}\mathrm{d}x,\ V_z = \int_1^{+\infty} B_3(\{x\})\frac{(Az+(\alpha+1)x^\alpha)^2}{x^3(x^\alpha+Az)^3}\mathrm{d}x.$$

First estimate U_z for large z:

• if $1 < \alpha \leq 2$, variable change $x = t(Az)^{1/\alpha}$ yields

$$U_z = (Az)^{-2-2/\alpha}\int_{1/(Az)^{1/\alpha}}^{+\infty} B_3\left(\left\{\frac{t}{(Az)^{1/\alpha}}\right\}\right)u(t)\mathrm{d}t$$

where $u(t) = t^{\alpha-1}/(t^2(t^\alpha+1)^2)$. Polynomial B_3 is upper bounded by some positive constant b on interval $[0,1]$ so that

$$\mid U_z \mid \leq b(Az)^{-2-2/\alpha}\int_{1/(Az)^{1/\alpha}}^{+\infty} u(t)\mathrm{d}t \qquad \text{(B.9)}$$

where the latter integral diverges for the lower bound $t = 0$ and is $O(z^{-1+2/\alpha})$ for $1 < \alpha < 2$ or $O(\log z)$ for $\alpha = 2$. As a consequence, U_z is of order $O(z^{-3})$ for $1 < \alpha < 2$ or $O(z^{-3}\log z)$ for $\alpha = 2$ and therefore always tends to 0 when z tends to infinity;

• if $\alpha > 2$, the integral in (B.9) converges for lower bound $t = 0$ and therefore $U_z = O(z^{-2-2/\alpha})$ tends to 0.

We thus deduce that for all $\alpha > 1$, zU_z tends to 0 when z tends to infinity. Addressing now V_z, write

$$V_z = (Az)^{-1-3/\alpha}\int_1^{+\infty} B_3(\{x\})v\left(\frac{x}{(Az)^{1/\alpha}}\right)\mathrm{d}x$$

with

$$v(t) = \frac{(1+(\alpha+1)t^\alpha)^2}{t^3(1+t^\alpha)^3}.$$

We have $t^3 v(t) \leq M$ for $t \geq 0$ and some constant $M > 0$ so that

$$(Az)^{-3/\alpha} v\left(x(Az)^{-1/\alpha}\right) \leq \frac{M}{x^3}$$

for all $x \geq 1$, where function $x \mapsto M/x^3$ is integrable on $[1, +\infty[$; besides, $v(t) \sim 1/t^3$ as $t \downarrow 0$ so that $(Az)^{-3/\alpha} v(x(Az)^{-1/\alpha})$ tends to $1/x^3$ for given $x \geq 1$. Applying the dominated convergence theorem to the family of functions $x \geq 1 \mapsto (Az)^{-3/\alpha} v(x(Az)^{-1/\alpha})$, we therefore deduce that $(Az)V_z$ has the finite limit

$$v_\alpha = \int_1^{+\infty} B_3(\{x\}) \frac{\mathrm{d}x}{x^3}.$$

We conclude from (B.8) and the above discussion that T_z tends to the finite limit $t_\alpha = -2\alpha v_\alpha$ when $z \uparrow +\infty$.

Gathering terms in (B.2)-(B.7), we are finally left with expansion (19) where constants A and $S_\alpha = -\alpha + \alpha/12 + t_\alpha/6$ depend on α only (some further calculations can provide the actual value $S_\alpha = -\alpha \log(2\pi)/2$, although that value is not necessary in our discussion). $\square$

Appendix C. Proof of Proposition 3.8

Recall definition (A.1) of function g and write equivalently

$$g(z) = \sum_{r \geq 1} g_z(r).$$

where we let $g_z(x) = Az(x^\alpha + Az)^{-1}$. The Euler-Maclaurin summation formula ([37], Chap.VI, Sect.2, formula (16.4)) applies again in the form

$$\sum_{r=1}^{J} g_z(r) = \int_1^J g_z(x)\mathrm{d}x + \frac{1}{2}\left[g_z(J) + g_z(1)\right] + \frac{1}{12}\left[g_z'(J) - g_z'(1)\right] + \frac{W_z(J)}{6} \quad \text{(C.1)}$$

for given $z \in \mathbb{C} \setminus \mathbb{R}^-$, integer $J \geq 1$ and where

$$\mid W_z(J) \mid \leq \frac{12}{(2\pi)^2} \int_1^J \mid g_z^{(3)}(x) \mid \mathrm{d}x$$

(derivatives of g_z are taken with respect to variable x). Consider the behavior of the r.h.s. of (C.1) as J tends to infinity. Firstly, $g_z(1) = Az/(Az+1)$ and $g_z(J) = O(J^{-\alpha})$ as $J \uparrow +\infty$; secondly, $g_z'(1) = -A\alpha z(1+Az)^{-2}$ together with $g_z'(J) = O(J^{-\alpha-1})$ for large J. Differentiating twice again shows that the third derivative of g_z is $O(x^{-\alpha-3})$ for large positive x and is consequently integrable at infinity. Letting J tend to infinity in (C.1) therefore implies equality

$$g(z) = \int_1^{+\infty} g_z(x)\mathrm{d}x + \frac{1}{2}\frac{Az}{Az+1} + \frac{1}{12}\frac{A\alpha z}{(1+Az)^2} + \frac{W_z}{6} \quad \text{(C.2)}$$

where

$$| W_z | \leq \frac{12}{(2\pi)^2} \int_1^{+\infty} | g_z^{(3)}(x) | \, \mathrm{d}x.$$

Using the explicit expression of the derivative $g_z^{(3)}$, it can be simply shown that

$$| W_z | = O(z^{-2/\alpha}), \; | W_z | = O(z^{-1} \log z), \; | W_z | = O(z^{-1}) \tag{C.3}$$

if $\alpha > 2$, $\alpha = 2$ and $1 < \alpha < 2$, respectively. Now, considering the first integral in the r.h.s. of (C.2), the variable change $x = t(Az)^{1/\alpha}$ gives

$$\int_1^{+\infty} g_z(x) \mathrm{d}x = I(Az)^{1/\alpha} - 1 + O\left(\frac{1}{z}\right) \tag{C.4}$$

where $I = L/\alpha = \rho_\alpha^{1/\alpha}$, with integral L introduced in (B.5) for $\alpha > 1$. Expanding all terms in powers of z for large z, it therefore follows from (C.2) and (C.4) that

$$g(z) = \rho_\alpha^{1/\alpha}(Az)^{1/\alpha} - \frac{1}{2} + W_z$$

with W_z estimated in (C.3). For large C, equation (A.1), *i.e.*$g(\theta_C) = C$, then reads

$$A\theta_C = \left[\frac{C}{I} + \frac{1}{2I} + O(W_{\theta_C})\right]^\alpha = \left(\frac{C}{I}\right)^\alpha + \frac{\alpha}{2I^\alpha} C^{\alpha-1} + O(C^{\alpha-2}) \tag{C.5}$$

for $\alpha > 2$ since (C.3) implies $W_{\theta_C} = O(\theta_C^{-2/\alpha}) = O(C^{-2})$ in this case. The case $1 < \alpha < 2$ gives a similar expansion since the remainder is then of order $W_{\theta_C}/C = O(C^{-\alpha}/C) = O(C^{-\alpha-1})$. Finally, the case $\alpha = 2$ yields

$$A\theta_C = \left[\frac{C}{I} + \frac{1}{2I} + O(W_{\theta_C})\right]^2 = \left(\frac{C}{I}\right)^2 + O\left(\frac{\log C}{C}\right). \tag{C.6}$$

Gathering results (C.5)-(C.6) finally provides expansions (21) for θ_C. □

Appendix D. Proof of Proposition 3.9

We here verify that conditions (16) and (17) of Theorem 3.6 are satisfied in the case of a Zipf popularity distribution with exponent $\alpha > 1$. Let us first establish convergence result (16). Using Lemma 3.7, we readily calculate

$$\begin{aligned} e^{s\sqrt{C}} \frac{F(\theta_C e^{-s/\sqrt{C}})}{F(\theta_C)} = \; & e^{s\sqrt{C}} \exp\Big[\alpha(\rho_\alpha A\theta_C)^{1/\alpha}\left(e^{-s/\alpha\sqrt{C}} - 1\right) \\ & + \frac{s}{2\sqrt{C}} + \varepsilon(\theta_C e^{-s/\sqrt{C}}) - \varepsilon(\theta_C)\Big] \end{aligned} \tag{D.1}$$

for any given $s \in \mathbb{C}$ with $\Re(s) = 0$, $| \Im(s) | \leq a$ and where $\varepsilon(\theta) \to 0$ as $\theta \uparrow +\infty$. By Lemma 3.8, we further obtain $\alpha(\rho_\alpha A\theta_C)^{1/\alpha} = \alpha C + \alpha\rho_\alpha r_C + o(r_C)$ and the expansion of $e^{-s/\alpha\sqrt{C}} - 1$ at first order in $1/C$ entails that

$$\alpha(\rho_\alpha A\theta_C)^{1/\alpha}\left(e^{-s/\alpha\sqrt{C}} - 1\right) = -s\sqrt{C} + \frac{s^2}{2\alpha} + O\left(\frac{1}{\sqrt{C}}\right);$$

letting C tend to infinity, we then derive from (D.1) and the previous expansions that

$$e^{s\sqrt{C}}\frac{F(\theta_C e^{-s/\sqrt{C}})}{F(\theta_C)} \to \exp\left(\frac{s^2}{2\alpha}\right)$$

so that assumption (16) is satisfied with $\sigma^2 = 1/\alpha$.

Let us finally verify boundedness condition (17). Rephrasing (D.1) for the value $s = -iy\sqrt{C}$, we have

$$\left(\frac{F(\theta_C e^{iy})}{F(\theta_C)}\right)^{1/C} = \exp\Big[\frac{\alpha(\rho_\alpha A\theta_C)^{1/\alpha}}{C}\left(e^{iy/\alpha}-1\right) - \frac{iy}{2C} + \frac{\varepsilon(\theta_C e^{iy}) - \varepsilon(\theta_C)}{C}\Big] \quad \text{(D.2)}$$

for any $y \in \mathbb{R}$. But as above, $\alpha(\rho_\alpha A\theta_C)^{1/\alpha}/C$ tends to the constant α when $C \uparrow +\infty$ so that

$$\left|\frac{F(\theta_C e^{iy})}{F(\theta_C)}\right|^{1/C} \leq \mid E(y) \mid^{\beta} \times \left|\exp\Big[\frac{\varepsilon(\theta_C e^{iy}) - \varepsilon(\theta_C)}{C}\Big]\right| \quad \text{(D.3)}$$

for some positive constant β and where

$$E(y) = \left|\exp\left(e^{iy/\alpha}-1\right)\right| = \exp\left(\cos\left(\frac{y}{\alpha}\right)-1\right).$$

Function E is continuous, even and given $\delta > 0$, h is decreasing on interval $[\delta, \pi]$ since $\alpha > 1$, hence $E(y) \leq E(\delta) = \eta_\delta < 1$ for $\delta \leq y \leq \pi$. Using the estimates of remainder T_z derived in **Appendix B**, it is further verified that, given any compact $\mathcal{K} \subset \mathbb{C}$ not containing the origin $u = 0$, we have $\lim_{C\uparrow+\infty} \varepsilon(\theta_C u) = 0$ uniformly with respect to $u \in \mathcal{K}$; this entails that the exponential term in the right-hand side of (D.3) tends to 1 when $C \uparrow +\infty$ uniformly with respect to $u = e^{iy}$, $y \in [\delta, \pi]$. We finally conclude that condition (17) is verified. □

Appendix E. Proof of inequality (27)

Using functional properties $\Gamma(z)\Gamma(1-z) = \pi/\sin(\pi z)$, $0 < z < 1$, and $\Gamma(z) = \Gamma(1+z)/z$ verified by Γ function ([33], p.256, formulae 6.1.17 and 6.1.15), expressions (20) and (26) for ρ_α and λ_α give

$$\log\left(\frac{\rho_\alpha}{\lambda_\alpha}\right) = \frac{\log\Gamma(1+x) - x\log x}{x}$$

where $x = 1/\alpha$. Letting $h(x) = \log\Gamma(1+x) - x\log x$, we now show that $h(x) > 0$ for $0 < x < 1$, which will provide the conclusion.

We first note that $h(0) = h(1) = 0$ and $h'(x) = \psi(x+1) - \log x - 1$ where $\psi(z) = \Gamma'(z)/\Gamma(z)$ is defined by

$$\psi(z) = -\gamma - \frac{1}{z} + \sum_{n\geq 1}\frac{z}{(z+n)n},$$

γ denoting Euler's constant; ψ is known [33] to satisfy the functional identity $\psi(z+1) = \psi(z)+1/z$. The above expressions of h' and ψ successively imply that $h'(0) = +\infty$ and $h'(1) = -\gamma$ since $\psi(2) = \psi(1) + 1 = -\gamma + 1$. Differentiating again, we obtain $h''(x) = (k(x) - 1)/x$ for $0 < x < 1$ where

$$k(x) = \sum_{n \geq 1} \frac{x}{(x+n)^2};$$

a final differentiation readily shows that $k'(x) > 0$ for $0 < x < 1$.

Function k is therefore strictly increasing on interval $[0, 1]$ from $k(0) = 0$ to $k(1) = \pi^2/6 - 1 < 1$, and therefore $k(x) < 1$ for $0 < x < 1$. As a consequence, h'' is always negative on $[0, 1]$, which implies that h' is strictly decreasing on that interval from $h'(0) = +\infty$ and $h'(1) = -\gamma$; h' thus vanishes at a unique value $x^* \in]0, 1[$ which is the unique extremum of h in $[0, 1]$. Function h is therefore increasing on $[0, x^*]$ from $h(0) = 0$ to $h(x^*) > 0$ and decreasing on $[x^*, 1]$ from $h(x^*) > 0$ to $h(1) = 0$. This implies that $h(x) > 0$ for $0 < x < 1$, as claimed. □

Appendix F. Proof of Proposition 5.1

Let $q_r(\ell + 1)$, $r \geq 1$, denote the distribution of the input process at cache $\sharp(\ell + 1)$, $\ell \geq 1$. By the same reasoning than that performed in Lemma 3.3, we can write

$$q_r(\ell + 1) = q_r(\ell) \frac{M_r^*(\ell)}{M^*(\ell)} = q_r \frac{M_r^*(\ell)...M_r^*(1)}{M^*(\ell)...M^*(1)} \tag{F.1}$$

for all $r \in \mathbb{N}$, where $M_r^*(\ell)$ (resp. $M^*(\ell)$) is the local miss probability of a request for object r at cache $\sharp\ell$ (resp. the averaged local miss probability for all objects requested at cache $\sharp\ell$) introduced in (29) and with notation $q_r = q_r(1)$. As $M_r^*(\ell') \to 1$ for all $\ell' \leq \ell$ when $r \uparrow +\infty$, we deduce from (F.1) that

$$q_r(\ell + 1) \sim \frac{A(\ell)}{r^\alpha}$$

when $r \uparrow +\infty$, where $A(\ell) = A/M^*(1)M^*(2)...M^*(\ell)$. Apply then estimate (25) to obtain

$$M_r^*(\ell + 1) \sim \frac{1}{1 + \theta(\ell + 1)q_r(\ell + 1)} \tag{F.2}$$

where $\theta(\ell+1) \sim C_{\ell+1}^\alpha / A(\ell)\rho_\alpha$ and with $q_r(\ell+1)$ given by (F.1); using the value of $A(\ell)$ above and the definition $q_r = A/r^\alpha$, the product $\theta(\ell+1)q_r(\ell+1)$ reduces to

$$\begin{aligned}\theta(\ell + 1)q_r(\ell + 1) &\sim \frac{C_{\ell+1}^\alpha}{A(\ell)\rho_\alpha} \cdot q_r \frac{M_r^*(\ell)...M_r^*(1)}{M^*(\ell)...M^*(1)} \\ &= \frac{C_{\ell+1}^\alpha}{\rho_\alpha r^\alpha} M_r^*(\ell)...M_r^*(1).\end{aligned} \tag{F.3}$$

Writing then $M_r(\ell+1) = M_r(\ell)M_r^*(\ell+1)$, asymptotics (F.2) and (F.3) together yield

$$M_r(\ell+1) \sim M_r(\ell)\left[1 + \frac{C_{\ell+1}^\alpha}{\rho_\alpha r^\alpha} M_r^*(\ell)...M_r^*(1)\right]^{-1}$$

so that

$$\frac{1}{M_r(\ell+1)} \sim \frac{1}{M_r(\ell)} + \frac{C_{\ell+1}^\alpha}{\rho_\alpha r^\alpha}$$

since $\Pi_{j=1}^{\ell} M_r^*(j) = M_r(\ell)$ after (28); the latter recursion readily provides expression (30) for $M_r(\ell)$, $1 \leq \ell \leq K$.

Using relation $M_r(\ell+1) = M_r(\ell)M_r^*(\ell+1)$ again together with expression (30) of $M_r(\ell)$ provides in turn expression (30) for $M_r^*(\ell)$, $1 \leq \ell \leq K$.

We finally justify estimate (31) for the average miss probability $M(\ell)$. Writing by definition $M(\ell) = \Sigma_{r \geq 1} q_r M_r(\ell)$, we estimate the latter sum by a Riemann integral; using asymptotic (30) for $M_r(\ell)$ then gives

$$\begin{aligned} M(\ell) = \sum_{r=1}^{+\infty} q_r M_r(\ell) &\sim \sum_{r=1}^{+\infty} \frac{A}{r^\alpha}\left(\frac{\rho_\alpha r^\alpha}{\rho_\alpha r^\alpha + C_1^\alpha + ... + C_\ell^\alpha}\right) \\ &\sim A\rho_\alpha \int_0^{+\infty} C_1 \mathrm{d}x \frac{1}{\rho_\alpha C_1^\alpha x^\alpha + C_1^\alpha + ... + C_\ell^\alpha} \\ &= \frac{A\rho_\alpha}{C_1^{\alpha-1}} \int_0^{+\infty} \frac{\mathrm{d}x}{1 + v^\alpha + \rho_\alpha x^\alpha} \end{aligned} \tag{F.4}$$

after setting $v^\alpha = (C_2^\alpha + ... + C_\ell^\alpha)/C_1^\alpha$. Variable change $t = x\rho_\alpha^{1/\alpha}/(1+v^\alpha)^{1/\alpha}$ reduces the latter integral to

$$\frac{1}{v^\alpha + 1} \times \left(\frac{v^\alpha + 1}{\rho_\alpha}\right)^{1/\alpha} \int_0^{+\infty} \frac{\mathrm{d}t}{1 + t^\alpha} = \frac{1}{(v^\alpha + 1)^{1-1/\alpha}}$$

so that estimate (F.4) finally provides (31). □

Massimo Gallo is a member of technical staff at Alcatel-Lucent Bell Labs, Paris, France since May 2013. He received the Bachelor and the Master of Science degrees in Computer and Communication Networks from Politecnico di Torino in 2006 and 2008. He obtained the PhD in Networks and computer science from Telecom ParisTech, Paris, France in 2012, performing his graduate research at Orange Labs, France Telecom, Paris, France. His main research interests are on performance evaluation, simulation, system design and experimentation in communication Networks.

Bruno Kauffmann graduated from École Normale Supérieure in 2007. He received his PhD from cole Normale Suprieure and Universit Pierre et Marie Curie in 2011, under the supervision of Prof. Franois Baccelli. During his thesis, he worked on inverse problems in networks, leveraging queuing theory results and statistical techniques to infer the state of a network. He is with Orange Labs since 2011. Bruno's research focus lies in stochastic modeling and performance analysis, applied to current CDNs solution and to ICN as a future Internet architecture.

Luca Muscariello is a senior expert on future networks at Orange Labs, former France Telecom R&D. He graduated in Telecommunications Engineering from Politecnico di Torino, Italy, in 2002. He holds a PhD in Telecommunications Engineering from Politecnico di Torino awarded as European Doctorate by the Conference of European Rectors. His graduate research was performed in the area of Internet traffic measurement, characterization and modeling at Politecnico di Torino, France Telecom R&D and at VTT in Helsinki. His current research interests include performance evaluation of communication networks, protocol design and optimization as well as network measurements and experimentation. He has filed 10 patents and published more than 30 technical papers in major conferences and journals on communication networks like IEEE INFOCOM, ACM SIGMETRICS, IEEE ICNP and IEEE/ACM Transactions on Networking. He has been part of several program committees like IEEE INFOCOM and is program co-chair of Valuetools 2013. He is a senior member of the IEEE and SEE and a member of the ACM.

Alain Simonian an graduated from the École Polytechnique, France, in 1982 and Telecom ParisTech in 1984, respectively, and received his PdD in Mathematics from the Université Rennes in 1993. His research interests address the modeling and the performance evaluation for the design of new telecommunications networks (Information-Based Networking ICN , 3G and 4G Mobile networks), and probabilistic risk models and optimal investments plans (Option theory). After working with Strategic Planning and Finance departments in Orange (France Telecom), he is now with Orange Labs and presently studies Caching policies in the ICN, along with traffic control rules for the next-generation LTE networks. On the more fundamental side, he is also active in developing new mathematical methods for the resolution of queuing systems non amenable to classical resolution techniques.

Christian Tanguy is an alumnus of the École Normale Supérieure. He graduated from the École Nationale Supérieure des Télécommunications in 1987, obtained the *Thèse de troisième cycle* in atomic physics (University Paris 6, 1983; supervisor: Prof. Claude Cohen-Tannoudji), and the *Habilitation à diriger des recherches* in solid-state physics (University Paris 7, 1996). He has been working in optoelectronics until 2001, mainly as a theoretical physicist. He is currently working in reliability theory and its applications to telecommunications networks, and is involved in the modeling of efficient, third-generation photovoltaic solar cells.