

Geographically Fair In-Network Caching for Mobile Data Offloading

Mahmoud El Chamie, Chadi Barakat, Giovanni Neglia

► **To cite this version:**

Mahmoud El Chamie, Chadi Barakat, Giovanni Neglia. Geographically Fair In-Network Caching for Mobile Data Offloading. IFIP Networking 2015, May 2015, Toulouse, France. <10.1109/IFIPNetworking.2015.7145318>. <hal-01136402>

HAL Id: hal-01136402

<https://hal.inria.fr/hal-01136402>

Submitted on 29 Jun 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Geographically Fair In-Network Caching for Mobile Data Offloading

Mahmoud El Chamie

Chadi Barakat

Giovanni Neglia

Abstract—Data offloading from the cellular network to low-cost WiFi has been the subject of several research works in the last years. In-network caching has also been studied as an efficient means to further reduce cellular network traffic. In this paper we consider a scenario where mobile users can download popular contents (e.g., maps of a city, shopping information, social media, etc.) from WiFi-enabled caches deployed in an urban area. We study the optimal distribution of contents among the caches (i.e., what contents to put in each cache) to minimize users’ access cost in the whole network. We argue that this optimal distribution does not necessarily provide geographic fairness, i.e., users at different locations can experience highly variable performance. In order to mitigate this problem, we propose two different cache coordination algorithms based on gossiping. These algorithms achieve geographic fairness while preserving the minimum access cost for end users.

I. INTRODUCTION

The reduction of the load on the 3G/4G infrastructure by offloading mobile traffic to low-cost WiFi has been the subject of several research works in the last years [1], [2]. The main idea is to redirect content requests, when it is possible, from the mobile infrastructure to WiFi hotspots. If requests are delay tolerant (videos, images, documents, emails, etc.), it is indeed possible to delay them until a WiFi opportunity arises. In [3], the authors study the performance of delaying mobile data offloading with user-defined deadlines in the framework of queuing models. In [4] the authors propose schemes for the prediction of future WiFi opportunities and advocate the gain in terms of 3G traffic reduction, while in [5] requests are redirected to neighboring mobile devices where the operators push the content.

In addition to traffic redirection via low-cost WiFi, in-network caching has also been studied to further reduce cellular network traffic [6], [7]. Indeed, because of the skewed content popularity, even limited-size caches at the network edges can reduce the load on the core. The authors in [8] illustrate this benefit by the help of a measurement-based study of 3G HTTP traffic. The idea of in-network caching is the key principle in Information-Centric Networks (ICN), like the Content-Centric Network (CCN) paradigm [9].

Our present work shares the same objective of mobile traffic offloading with this set of related work taking advantage of in-network caching. In particular we consider a scenario where WiFi-enabled caches are deployed at different locations in a




Fig. 1. The mobility-network model for a user.

city to provide mobile users (e.g., tourists) access to a set of relatively static popular contents (e.g., maps of the city, shopping information, social media, etc.). Caches are assumed to be only loosely connected to each other and to the Internet, e.g., through some low-rate or intermittent communication infrastructure. For example they could take advantage of a Delay Tolerant Network (DTN) backbone leveraging the public transportation system [10]. Fig. 1 shows our user-caches scenario.

Users can retrieve the contents at any time through their 3G data connection or wait to enter in the transmission range of a cache storing the content. The user is assumed to incur a higher cost when s/he accesses the content through the cellular network than when s/he exploits the WiFi connection, e.g., because of direct data charges, data roaming, reduction of the monthly residual amount of data available, larger battery depletion, or lower download rates. At the same time, the user is assumed to be *impatient*, i.e., s/he may not be willing to delay content fruition more than a certain time and can then resort to use her/his cellular data connection.

Two questions arise in such scenario. First, how should the set of contents be distributed across the caches in order to minimize users’ cost to access the contents? Second, can the system benefit from some form of cooperation among the caches despite their loose connectivity? In this paper we address both these issues. We start formulating an optimization problem for contents’ placement in caches (i.e., which files to put in which cache) in the presence of impatient users, and then propose an efficient algorithm to find the placements that minimize user’s content retrieval cost and maximize the amount of offloaded traffic, thus relieving the cellular network. We highlight the fact that the cost-optimal distribution of con-

tents does not guarantee in general geographical fairness, i.e., users experience high variability in performance at different locations in the network. We propose then a new principle of “content shuffling by gossiping” that takes advantage of the limited communication infrastructure connecting the caches. We believe the idea of local content shuffling can have applications well beyond the scenario considered in this paper.

The rest of the paper is organized as follows. In Section II we describe how we modeled the scenario. Section III studies the optimal in-network caching distribution to minimize content retrieval costs for users. We then introduce the concept of geographic fairness in Section IV and propose two caches’ coordination algorithms to improve it. Section V presents some simulation results, and Section VI concludes the paper.

II. MODEL

A. Scenario

Consider a network of WiFi-enabled caches that are deployed in a region where a set of popular contents (maps of the city, shops information, etc.) is made available to mobile users (e.g., tourists). The users can use their WiFi supported devices (e.g., tablets, smartphones) while moving in that area to connect to the closest cache. When the user requests a content, s/he has two ways to get it: either to download the content using her/his 3G data connection with a cost C_{3G} , or to wait to meet a cache storing the content and download it through WiFi with a cost C_{WiFi} . We assume that the cost per content using the WiFi is smaller than that of 3G, i.e.,

$$C_{WiFi}^{(i)} < C_{3G}^{(i)} \quad \text{for every content } i.$$

A single cache cannot store the full content catalog, hence, while moving around the city, the user connects to different caches that may or may not have the requested content. Note that if a cache does not contain the content, it will not retrieve the content for the user, but s/he can either search for the content from another cache or download it using her/his 3G connection. This is coherent with our assumption that caches are poorly connected among themselves and to the Internet. We consider impatient users, i.e., each user has a maximum time (we call it the *patience time*) s/he is willing to wait to download a content from a cache. If a cache with the content has not been found by such time, the user will use the 3G to get it and will then incur a higher cost.

The choice of which content to store at each cache is fundamental to satisfy a large number of users’ queries and reduce their costs. Intuitively, if the patience time is the same for all the contents, highly requested contents should be stored in more caches than the less requested ones. What happens if patience times are heterogeneous is not clear (even qualitatively), and in any case a model is required to determine the exact number of replicas.

B. Contents

We suppose that the catalog size for users’ requests is K , i.e., the total number of different contents or files requested by the users is K (along this paper, we use the two terms contents and files interchangeably). A request for one of the

catalog contents is a random variable $R \in \{1, \dots, K\}$ with the following discrete probability distribution:

$$q_i = \text{Prob}(R = i) \quad i = 1, \dots, K. \quad (1)$$

Without loss of generality, we suppose that the contents are ordered in decreasing order of popularity, i.e., the first content has the highest popularity and the K -th content has the lowest one, i.e., $q_1 \geq q_2 \geq \dots \geq q_K$. Contents are assumed to have the same storage size.

C. Caches

The service provider has a fixed number N of identical caches. Each cache can store up to B contents and provide them to nearby users upon request. The caches are distributed in different geographic places in the city.

Caches have the possibility to establish a connection with each other from time to time using a low cost network (a dedicated wired network, an ad-hoc network using buses or any other mobile devices as data ferries [11], etc.). Once two caches are connected, we assume that they can exchange contents. Therefore, the caches form a network that can be represented by an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the caches are its vertices ($\mathcal{V} = 1, \dots, N$) and each edge $(i, j) \in \mathcal{E}$ identifies that there is a recurrent communication opportunity between cache i and j (a formal communication model is introduced in Section IV).

In the first part of this work, we will study the optimal content distribution among caches, in the sense that this distribution minimizes the users’ cost to retrieve the content. To this regard, the cache network topology does not play any role, because the optimization problem can be solved in a centralized way and the caches do not need to exchange any content. However, when we study the problem of geographic fairness later in the paper, we will propose algorithms that rely on contents’ shuffling among the caches. Here, the inter-caches communication infrastructure is important in determining the performance of the algorithms.

D. Users

We consider a homogeneous set of users, so in this section we will refer to one of them. As mentioned earlier, the user generates requests for contents to download according to the probability distribution introduced in Section II-B. The user keeps moving inside the city and can get in contact with a given WiFi cache multiple times. We say that a user is *associated* to a cache if her/his WiFi device is actually associated to the cache’s Access Point and can download contents. During an association, the user will complete the download of all the contents s/he is interested in and which are stored at the cache. This holds if we assume that the user stops moving until the download is complete, or that the download time is smaller than the user’s association time to the cache.

Let us consider a request generated by the user at time t_0 . We define the *residual time* between the user and cache j as the first time after t_0 when the user associates to cache j and we denote it by X_j . Through the whole paper, we consider valid the following assumption:

Assumption 1. The residual times X_1, \dots, X_N between the user and the caches in the network are independent and identically distributed (i.i.d.) random variables.

To highlight the generality of this mobility model for the user, we give the following mobility model example that satisfies our assumption.

Example: Assume that for each cache j the point process of the instants when the user associates to cache j is a renewal process, where the random time Z_j between two consecutive associations has Cumulative Distribution Function (CDF) $F_{Z_j}(t)$. If the N renewal processes are i.i.d. and stationary, then Assumption 1 is satisfied. The residual time X_j is the residual time for $t = 0$ of the corresponding renewal process [12, p. 116] and its CDF can then be calculated as follows: $F_{X_j}(x) = \int_0^x \bar{F}_{Z_j}(\tau) d\tau / \mathbb{E}[Z_j]$, where $\bar{F}_{Z_j}(\tau) = 1 - F_{Z_j}(\tau)$ is the complementary CDF. ■

While our results below hold under Assumption 1 for generic distributions for the variables X_1, \dots, X_N , we often consider the exponential distribution as an example in order to derive closed form expressions and provide further insights about the problem.

The user is impatient, i.e., he is willing to wait a time at most T_i before being able to download content i . If the user finds content i at one cache by time T_i , s/he downloads it from the cache using the Wifi connection at a cost $C_{WiFi}^{(i)}$. If the timer expires and the user has not found the content, s/he downloads it using 3G at a cost $C_{3G}^{(i)}$. We refer to T_i as the user's *patience time*.

Let R be the (random) content requested at time t_0 , then the cost C_R incurred to download it is:

$$C_R = C_{3G}^{(R)} \mathbb{1}_{\text{file downloaded by 3G}} + C_{WiFi}^{(R)} \mathbb{1}_{\text{file downloaded by WiFi}},$$

where $\mathbb{1}$ is the indicator function (equal to 1 if the condition is satisfied and 0 otherwise). The randomness in C_R comes from two factors: 1) the identity of the requested content and 2) the instants when the user associates to the different caches after $t = t_0$.

III. TUNING CONTENT DISTRIBUTION FOR OPTIMAL NETWORK PERFORMANCE

Let C be the expected cost for the request C_R , i.e., $C = \mathbb{E}[C_R]$. C depends obviously on the distribution of contents among the caches. On one hand, placing more replicas of popular contents in the network of caches reduces the costs because it allows more frequent downloads from WiFi, which is less costly than 3G. On the other hand, placing popular contents everywhere reduces the remaining available space for other, less popular, contents to be stored in the network without bringing necessarily the same level of cost reduction. Therefore, due to caching space constraints of the caching network, the number of copies should be selected wisely to reduce the costs. Let us start with the following proposition that gives the expression of C .

Proposition 1. The expected cost per content request for any user is

$$C = \alpha + \sum_{i=1}^K \gamma_i q_i (1 - F_{Y^{(i)}}(T_i)), \quad (2)$$

where α and γ_i are positive constants, $Y^{(i)}$ is the minimum of the residual times between the user and any of the caches having content i , and $F_{Y^{(i)}}(T_i) = \text{Prob}(Y^{(i)} \leq T_i)$ is its cumulative distribution function.

Proof: Suppose that there are n_i copies of a content i in the network (i.e., n_i is the number of caches storing content i). Then the minimum of the residual times is

$$Y^{(i)} = \min\{X_1, X_2, \dots, X_{n_i}\},$$

where X_1, \dots, X_{n_i} are i.i.d. random variables having the same distribution (due to Assumption 1). The probability that the user connects to a cache having content i within time delay T_i (and downloads it using WiFi) is:

$$\begin{aligned} & \text{Prob}[\text{content downloaded by WiFi} \mid \text{it is content } i] \\ &= \text{Prob}[Y^{(i)} \leq T_i] = F_{Y^{(i)}}(T_i), \end{aligned} \quad (3)$$

and the probability it is downloaded using 3G is:

$$\begin{aligned} & \text{Prob}[\text{content downloaded by 3G} \mid \text{it is content } i] \\ &= \text{Prob}[Y^{(i)} > T_i] = 1 - F_{Y^{(i)}}(T_i). \end{aligned} \quad (4)$$

Let A_1 be the event that the content is downloaded by WiFi and A_2 the event that the content is downloaded by 3G. Then,

$$\begin{aligned} \mathbb{E}[C_R] &= \sum_{i=1}^K C_{WiFi}^{(i)} \text{Prob}[A_1 \mid R = i] \times \text{Prob}[R = i] \\ &+ \sum_{i=1}^K C_{3G}^{(i)} \text{Prob}[A_2 \mid R = i] \times \text{Prob}[R = i] \\ &= \sum_{i=1}^K C_{WiFi}^{(i)} (F_{Y^{(i)}}(T_i)) q_i + \sum_{i=1}^K C_{3G}^{(i)} (1 - F_{Y^{(i)}}(T_i)) q_i \\ &= \sum_{i=1}^K q_i C_{WiFi}^{(i)} + \sum_{i=1}^K q_i (C_{3G}^{(i)} - C_{WiFi}^{(i)}) (1 - F_{Y^{(i)}}(T_i)), \end{aligned}$$

then the proposition follows taking $\alpha = \sum_{i=1}^K q_i C_{WiFi}^{(i)}$ and $\gamma_i = C_{3G}^{(i)} - C_{WiFi}^{(i)}$. ■

For the special case when the residual time between the user and a cache is exponentially distributed with rate λ , then $Y^{(i)} = \min\{X_1, X_2, \dots, X_{n_i}\}$ is also exponentially distributed with rate $n_i \lambda$. Therefore the cost C can be written in closed form as follows:

$$C = \alpha + \sum_{i=1}^K \gamma_i q_i \exp(-\lambda n_i T_i). \quad (5)$$

Proposition 1 shows that the average cost depends on the number of replicas each content has in the network. Notice that the average cost lies between $C_{WiFi} = \sum_{i=1}^K q_i C_{WiFi}^{(i)}$ and $C_{3G} = \sum_{i=1}^K q_i C_{3G}^{(i)}$, i.e., $C_{WiFi} < C \leq C_{3G}$, where $C = C_{3G}$ if the WiFi is not used or the caches do not store any content (i.e., $n_i = 0$ for all i), and C can get arbitrary close to C_{WiFi} if there are enough copies of every content (i.e., $\lim_{n_1, \dots, n_K \rightarrow \infty} C = C_{WiFi}$).

A. Optimizing Content Distribution

Due to the storage size constraints at caches, the service provider can optimize the performance of the network by selecting the number of replicas of each content to be placed in the caches so that the cost is minimized. The optimal number of replicas for each content can be derived solving the following optimization problem:

$$\begin{aligned} & \text{minimize} && f(n_1, n_2, \dots, n_K) = C \text{ given in (2)} \\ & \text{subject to} && \sum_{i=1}^K n_i \leq B.N, \\ & && n_i \in \{0, 1, \dots, N\}, \text{ for } i = 1, \dots, K. \end{aligned} \quad (6)$$

The first constraint states that the total number of replicas in the network should not exceed the available space in the caches. The second set of constraints limits the number of replicas of a content to the number of available caches and restricts this number to a non-negative integer.

The result of this optimization is the number of replicas of each content to be placed in the caches. Note that any placement of that number of contents under the condition that the same content is not placed more than once in the same cache (otherwise it just consumes space without adding any value) would provide the optimal cost. This placement is clearly not unique in general. We will get back to this issue later in the paper when talking about *geographically fair placements* in Section IV. Before providing an efficient algorithm to solve (6), we prove the following important property of its objective function:

Proposition 2. *The objective function $f(n_1, \dots, n_K)$ is a separable convex function in the variables of its argument.*

Proof: The function f is clearly separable in its variables because $f(n_1, \dots, n_K) = \alpha + \sum_{i=1}^K \gamma_i q_i (1 - F_{Y^{(i)}}(T_i))$, and $F_{Y^{(i)}}(T_i)$ is a function of n_i only. Then it is sufficient to prove that each of the functions $1 - F_{Y^{(i)}}(T_i)$ is convex in n_i . Denote $Y(n_i) = \min\{X_1, \dots, X_{n_i}\}$ where X_1, \dots, X_{n_i} are all i.i.d. (due to Assumption 1). Let us define $\bar{F}_{Y(n_i)}(T) = 1 - F_{Y(n_i)}(T) = \text{Prob}(Y(n_i) > T)$. Then we have for every $n_i > 0$,

$$\begin{aligned} & \bar{F}_{Y(n_i+1)}(T) - \bar{F}_{Y(n_i+2)}(T) \\ &= \text{Pr}(\min\{Y(n_i), X\} > T) - \text{Pr}(\min\{Y(n_i+1), X\} > T) \\ &= \text{Prob}(Y(n_i) > T) \times \text{Prob}(X > T) \\ &\quad - \text{Prob}(Y(n_i+1) > T) \times \text{Prob}(X > T) \\ &\leq \text{Prob}(Y(n_i) > T) - \text{Prob}(Y(n_i+1) > T) \\ &= \bar{F}_{Y(n_i)}(T) - \bar{F}_{Y(n_i+1)}(T), \end{aligned}$$

thus f is convex in the variables of its arguments. ■

B. A Greedy Optimal Algorithm

The optimization problem (6) gives the number of copies of every content in the network guaranteeing minimal cost. Let $n_1^*, n_2^*, \dots, n_K^*$ be this optimal allocation and let $C^* = f(n_1^*, n_2^*, \dots, n_K^*)$ be the corresponding minimum cost value. Let us define the following function for a content i :

$$D_i(n_i) = \gamma_i q_i (F_{Y(n_i+1)}(T_i) - F_{Y(n_i)}(T_i))$$

where $Y(n_i) = \min\{X_1, \dots, X_{n_i}\}$ and, by convention, $F_{\min\{X_1, \dots, X_{n_i}\}}(T_i) = 0$ if $n_i = 0$. In the particular case of an exponential mobility model with rate λ , then $D_i(n_i)$ has the following closed form:

$$D_i(n_i) = \gamma_i q_i (\exp(-n_i \lambda T_i) - \exp(-(n_i + 1) \lambda T_i)).$$

The optimal number of replicas of each content can be given by the following greedy algorithm:

Algorithm 1 Calculate $n_1^*, n_2^*, \dots, n_K^*$ of optimization (6)

```

1:  $n_1 = n_2 = \dots = n_K = 0$ 
2:  $S = \{1, 2, \dots, K\}$ 
3:  $V_i = D_i(n_i) \quad \forall i \in S$ 
4: while  $\sum_i n_i < B.N$  do
5:    $j = \text{Argmax}_{i \in S} V_i$ 
6:    $n_j \leftarrow n_j + 1$ 
7:    $V_j \leftarrow D_j(n_j)$ 
8:   if  $n_j = N$  then
9:     Remove  $j$  from  $S$ 
10:  end if
11: end while
12: return  $n_1, n_2, \dots, n_K$ 

```

Proposition 3. *Algorithm 1 returns the optimal solution $n_1^*, n_2^*, \dots, n_K^*$ of optimization (6).*

Proof: At every iteration, the algorithm is adding one content having the maximum difference V_i (steps 5 and 6 in the algorithm) and removing the content if $n_i > N$ (step 9). Thus Algorithm 1 is eventually equivalent to calculating the elements $D_i(m)$, for $i = 1, \dots, K$ and $m = 0, \dots, N - 1$, and then considering the contents corresponding to the largest BN values among those elements. The latter can be shown to provide the optimal solution by adapting the proof for a similar algorithm in the context of the *discrete resource allocation* problem with a separable convex function given in [13, Chapter 4]. Also in our case, what is fundamental for the proof is that the objective function is separable convex as we showed in Proposition 2. Additionally, this algorithm can be implemented efficiently using the *max heap* data structure [14] with a computational complexity $O(K + BN \log K)$, and memory complexity $O(K)$. ■

We give an example of optimal content distribution in Fig. 2. This numerical result corresponds to a mobility model with exponentially distributed residual times with rate λ , a Zipf law for content popularity, and two different cases of the patience time. For the case when T_i is also assumed to follow Zipf,¹ the figure shows that the replicas are not proportional to their popularity. Less popular contents (contents 8 and 9 in the figure) have the highest number of replicas in the network. The reason for the different number of replicas is that it is not just the popularity that plays a role in selecting the items, but there are other factors. On one hand, the cost functions are convex

¹A Zipf distribution for T_i may be justified if the user knows contents popularity. For a more popular content, the user would rather be more patient with the hope of finding it later in the network (because it is popular), but s/he would not be as patient for a less popular one as s/he knows that s/he is less likely to find it.




Fig. 2. The optimal distribution of contents as function of their popularity given by Algorithm 1 for two different cases of the patience, T_i being zipf and T_i being constant ($\alpha = 0$, $\gamma_i = 1$, $K = 10000$ files, $N = 50$ caches, $B = 10$ buffer size, q_i is zipf distribution with parameter $s = 1$, and $\lambda = 5$ is constant).

decreasing in the number of replicas, so adding one replica does not reduce the cost as much as it did a previous replica of the same content. On the other hand, the cost also depends on the patience time (T_i , for $i = 1, \dots, K$) by the user for the contents. In the first example, the patience time by users for the more popular contents is assumed to be larger than for the other contents which explains why less replicas are sufficient. In the case when T_i is constant ($T_i = 0.0067$ in this case), the first 4 contents have the same optimal number of replicas (50 in this case) and then that number starts decreasing again with the popularity. Note that contents that are not shown in the figure (from 23 up to 10000) are not cached in the network (i.e., $n_i^* = 0$).

IV. FAIRNESS

In the previous section we derived the optimal number of copies for each content to minimize the global content retrieval cost for the user. Any placement of this number of copies (as far as the placement ensures that the same cache does not contain more than one replica of the same content) guarantees the minimal cost. This is due to the implicit assumption of homogeneous user mobility, so that residual times between the user and each cache are identically distributed. However, different placements could lead to the user experiencing different quality-of-service at different caches. For example, this would happen for constant patience times, if one cache would store only copies of the most popular contents and another only copies of the least popular contents. We denote this problem as *geographical unfairness*. The global optimization problem considered until now provides minimum global cost but not geographical fairness.

It appears natural to formulate then a different optimization problem, where we try to optimally allocate the replicas among the caches, so that the expected cost at each cache is almost the same. To argue formally, let us quantify the performance of our proposed WiFi caches network. Since the aim of this network is to reduce the costs for users, we can define the




Fig. 3. Variability in caches' utilities shows that, even with optimal number of contents, the placement of these copies in the network can cause variable geographic performances ($N = 50$ caches randomly deployed, $K = 10000$ files, $B = 10$ buffer size, q_i and T_i are zipf distribution with parameter $s = 1$, $\lambda = 0.04$ is constant).

gain P of the network as the expected reduction of retrieval cost per request due to in-network caching. Since γ_i is the reduction in the cost when content i is downloaded from the caches and $F_{Y^{(i)}}(T_i)$ is the probability of finding this content, then $P = \sum_{i=1}^K q_i \gamma_i F_{Y^{(i)}}(T_i)$. The utility of a cache would then be its contribution to the performance,

$$V_j = \sum_{\text{contents } i \text{ in cache } j} q_i \gamma_i \frac{F_{Y^{(i)}}(T_i)}{n_i^*},$$

where we divided by n_i^* because when a request for a content i is satisfied, it is equally likely to be satisfied at any of the n_i^* caches having a replica of this content. Therefore, the utility (or value) U_i of a replica of content i is given by

$$U_i = \begin{cases} q_i \gamma_i \frac{F_{Y^{(i)}}(T_i)}{n_i^*} & \text{if } n_i^* > 0, \\ 0 & \text{if } n_i^* = 0. \end{cases} \quad (7)$$

For exponentially distributed residual times it is

$$U_i = \begin{cases} \gamma_i q_i \frac{1 - \exp(-n_i^* \lambda T_i)}{n_i^*} & \text{if } n_i^* > 0, \\ 0 & \text{if } n_i^* = 0, \end{cases} \quad (8)$$

so U_i can be considered as a measure for the importance of a replica of a content.

A numerical example for the potential high variability of caches' utilities is given in Fig. 3. The figure shows that the optimal number of copies does not necessarily provide geographic fairness. In this specific example, the caches having the most popular contents happen to be placed at the edges of the geographic area, and thus users moving close to these areas would get higher satisfaction compared to users moving near the regions where caches have less popular contents. The question that arises here is the following: how can we preserve optimal performance (having always the optimal number of copies in the caches) and at the same time get geographic fairness (same geographic performance everywhere in the area)?

So our goal is to achieve lower performance variability across the area. Since there are different ways of distributing

the replicas in the network, and each replica has a different performance value, for the purpose of fairness, it is desirable that this distribution of replicas minimizes the variability of values in different caches. Let us define

$$V_{max} = \max_{j \in \mathcal{V}} V_j,$$

which is the maximum utility of a caches. Therefore, to achieve fairness, we should distribute the contents in the caches in a way that minimizes V_{max} while still guaranteeing the minimum cost. This problem can be formulated as the following Mixed Integer Linear Program (MILP):

$$\begin{aligned} & \text{minimize} && y \\ & \text{subject to} && \sum_{i=1}^K U_i x_{ij} \leq y, \text{ for } j = 1, \dots, N, \\ & && \sum_{i=1}^K x_{ij} = B, \text{ for } j = 1, \dots, N, \\ & && \sum_{j=1}^N x_{ij} = n_i^*, \text{ for } i = 1, \dots, K, \\ & && x_{ij} \in \{0, 1\}, y \in \mathbb{R}, \end{aligned} \quad (9)$$

where the variable x_{ij} indicates whether a content i is placed in cache j or not. The newly introduced real variable y together with the first set of constraints guarantee that V_{max} is minimized. The second set of constraints guarantees that each router is not exceeding its buffer capacity. And the third set of constraints is to ensure that the number of replicas for a content does not exceed its optimal value. However, this optimization turns out to be computationally hard,

Proposition 4. *The fairness problem (9) is NP-complete.*

Proof: The proof is due to a reduction from the 3-partition problem (known to be NP-complete). The 3-partition problem is to decide whether a given set of S integers m_1, \dots, m_S can be partitioned into triples that all have the same sum. We show that given any 3-partition problem, it can be solved by the formulation (9) with a particular choices $T_i, q_i, \gamma_i, B, N,$ and K . There are different possibilities for these choices, one of them is to take $B = 3, N = S/3, K = S$ to be the number of contents, and $U_i = m_i$ (obtained by taking $T_i = \infty$ so that $n_i^* = 1$ for all $i = 1, \dots, K, \gamma_i = \sum_i m_i$ for all $i,$ and $q_i = m_i / \sum_i m_i$). Thus by solving this problem by (9) and checking, in polynomial time, if the values of the caches are equal, we can solve the original 3-partition problem. Therefore, the optimal placement of contents in the caches is at least as hard as the 3-partition problem and this ends the proof. ■

Given the difficulty to determine the static content distribution which maximizes fairness, an alternative approach could be to achieve it on the long run by dynamically changing content distribution among the caches. A naive idea is to periodically reshuffle the copies among the caches, by calculating in a centralized way a random distribution of the copies and transmit them to each cache accordingly. This solution guarantees that all the caches have the same time-average utility on the long run (because of the strong law of large numbers), but has significant communication cost, given that NB copies should be re-transmitted at every iteration. In order to improve on top of this idea by reducing the amount of communication required, we consider shuffling algorithms

based only on local exchanges among caches as in gossiping protocols [15] [16], without the need to download B copies at each cache.

The movement of contents between caches is done by exchanging files through the communication network connecting these caches. We consider an asynchronous operation, where pairs of caches can autonomously decide to exchange some contents taking advantage of their mutual connectivity, which, as we mentioned before, may also be intermittent. We model then the network as discrete-time system where at every time slot two caches exchange files, then k exchanges occur in the network by slot k . Let $G(k)$ be the graph representing the interaction at time k : this graph has the N caches as vertices and only one edge connecting the two caches that exchange contents at time k . We consider that the interaction among caches can be modeled probabilistically as follows:

Assumption 2. *At a given communication iteration, the probability that caches i and j exchange contents is given by the following expression,*

$$\text{Prob}((ij) \in G(k)) = \begin{cases} \sigma_{ij} > 0 & \text{if } (i, j) \in \mathcal{E} \\ 0 & \text{else;} \end{cases}$$

The probabilities σ_{ij} take into account both the probability there is a communication opportunity between caches i and j (as we said the network could be only intermittently connected) and the decision of the caches to actually exchange some contents. Given this model, the union of all possible graphs $G(k)$ is the connected graph \mathcal{G} with probability one.

Once the connection between two caches is established, these caches should choose which files (contents) to exchange. We propose here two exchange rules for the files:

- **Exchange Rule 1:** Common files are not exchanged, only different files are randomly divided between the two caches.
- **Exchange Rule 2:** Common files are not exchanged. The connected caches exchange one file (if it exists) that gives the maximum decrease in the local objective function $V_{max}^{(ij)} = \max\{V_i, V_j\}$.

Let $V_j(k)$ be the value of the cache j after the k -th exchange such that $V_j(0) = V_j^0$ are the values of the caches due to an initial distribution of the optimal copies in the network (the initial values vector is denoted by \mathbf{V}^0).

For both the exchange rules the total utility of the caches ($\sum_j V_j(k)$) does not vary neither with k nor with the initial placement, i.e.,

$$\sum_{j=1}^N V_j(k) = \sum_{i=1}^K n_i^* U_i, \quad \forall k, \forall \mathbf{V}^0.$$

This is evident if we think that the total number of copies per content does not change or if we consider that at every iteration, local utilities change as follows:

$$\begin{cases} V_i(k) + V_j(k) = V_i(k-1) + V_j(k-1) & \text{if } (i, j) \in G(k) \\ V_i(k) = V_i(k-1) & \text{else.} \end{cases} \quad (10)$$

Utilities of caches evolve randomly under both the exchange rules. For exchange rule 2 the randomness is only due to the

random selection of the caches which interact. For exchange rule 1 the randomness is also due to the random reallocation of the files among the caches. For this reason it is useful to consider the following quantities

$$x_j(k) = \begin{cases} V_j^0 & \text{if } k = 0 \\ \mathbb{E}(V_j(k)) & \text{if } k > 1, \end{cases}$$

where the expected value is calculated over the possible sequence of pairs of interacting caches for exchange rule 2 and over the sequence of pairs of interacting caches *and* possible re-allocations of the files for exchange rule 1. We can now give the following definition:

Definition 1. We say that the caches achieve asymptotic geographic fairness under a given exchange rule if for any initial optimal distribution of contents (i.e., any \mathbf{V}^0)

- 1) $\lim_{k \rightarrow \infty} x_j(k) = \bar{x}$ for $j = 1, \dots, N$,
- 2) $\bar{x} = \frac{1}{N} \sum_{i=1}^K n_i^* U_i$.

The first condition says that all caches would have on average the same value, i.e., it is indifferent to be in an area or in another because the caches have the same value. The other condition says that the distribution of contents remains optimal and its value is independent from the initial placement of contents (actually this corresponds to averaging the initial utilities of caches in the network).

Although the similarities illustrated above for the two exchange rules, there are some evident differences. For exchange rule 2 we have $|V_i(k) - V_j(k)| \leq |V_i(k+1) - V_j(k+1)|$ for every iteration k . Given that the sum of the utilities is constant, this inequality implies that the maximum over all the utilities ($\max_i V_i(k)$) is a non-increasing sequence and then converges to a given value. This value is a local minimum that the greedy local exchange rule 2 cannot improve. On the other side, utilities in exchange rule 1 do not converge in general, because at any time slot some contents are randomly reallocated. Nevertheless, exchange rule 1 has the following desirable property:

Proposition 5. Under Assumption 2 and the Exchange Rule 1, caches achieve asymptotic fairness.

Proof: The proof is provided in the Appendix. ■

V. PERFORMANCE STUDY

A. Comparison of Optimal Placement with LRU Caches

An alternative approach to content placement is to use LRU (Least Recently Used) replacement policy in a non-coordinated way where each cache applies its own policy as a function of the traffic it sees. For our particular scenario where caches see the same traffic distribution, they would have the same probability distribution of having a content (mainly most popular contents are found everywhere). To compare with our optimal placement, we assume that when a user does not find a requested file in an LRU cache, it continues searching in nearby caches (according to the mobility model) and does not wait for the cache to download the content (even if the cache will eventually replace one of its existing file with the recently missed file using 3G synchronization, we don't include this




Fig. 4. Cost comparison between having LRU replacement policy in caches (\underline{C}_{LRU}) and the cost due to optimal replicas from problem (6) (C^*) while varying the number of caches in the network from 1 to 100 ($\alpha = 0$, $\gamma_i = 1$, $K = 10000$ files, $B = 10$ buffer size, q_i and T_i are zipf distribution with parameter $s = 1$, $\lambda = 5$ is constant).

latter cost for the sake of comparison). Having said that, similar to the derivation of the expected cost in Proposition 1, we can write the expected cost per request C_{LRU} for using LRU caches with exponential mobility model of rate λ as follows:

$$C_{LRU} = \alpha + \sum_{i=1}^K \gamma_i q_i \mathbb{E}[\exp(-Z_i \lambda T_i)], \quad (11)$$

where Z_i is a random variable for the number of caches having content i . Since each of the caches are considered as independent LRU cache, then we have $\mathbb{E}[Z_i] = h_i N$ where h_i is the probability that a file i is present in one LRU cache (hit probability for file i). By applying Jensen's inequality for the convex functions $\phi_i(Z_i) = \exp(-Z_i \lambda T_i)$, we can derive a lower bound \underline{C}_{LRU} on C_{LRU} :²

$$C_{LRU} \geq \alpha + \sum_{i=1}^K \gamma_i q_i \exp(-h_i N \lambda T_i). \quad (12)$$

Notice that by comparing equations (5) and (12), n_i is now replaced by $h_i N$, which is no other than the mean number of replicas of content i across the network of LRU caches. To find h_i , we apply "Che approximation" [17], which is proved to be very accurate [18], given as follows:

$$h_i \approx 1 - \exp(-q_i t_C) \text{ for } i = 1, \dots, K, \quad (13)$$

where t_C is the unique root of the following equation $\sum_{i=1}^K (1 - \exp(-q_i t)) = B$.

Fig. 4 shows that the average cost per request when caches run optimization problem (6) is less than the one in the case caches use LRU replacement policy. Moreover, when the number of caches is small, adding one more cache in the

²Alternatively, another approach can be used to derive the same lower bound, we observe that $\mathbb{E}[\phi_i(Z_i)]$ is the probability of not finding content i in the network (i.e., probability of a miss) which can be calculated precisely by the following formula $\mathbb{E}[\phi_i(Z_i)] = (h_i \exp(-\lambda T_i) + (1 - h_i))^N$. Therefore, the mentioned lower bound can be also easily derived since $h_i \exp(-\lambda T_i) + (1 - h_i) \geq \exp(-h_i \lambda T_i)$.




Fig. 5. Utility of caches by considering different exchange rules.

network reduces quickly the cost for the optimal placement, which does not hold for the lower bound cost \underline{C}_{LRU} for LRU caches. At one point, it also seems from the figure that the difference between the two costs tends to be constant. It is important to recall that the cost given in the figure for the LRU does not include the cost of synchronization using the 3G (i.e., practically speaking we should add to the LRU cost curve in the figure the cost of downloading a new content after every miss). Having said that, the LRU has the advantage that it will automatically adapt to changes in the popularity distribution (q_i 's) while in our case we should rerun our optimization problem and perform a new placement (that can be done incrementally starting from the existing one). We leave the study of dynamic traffic to our future research.

B. Fairness Exchange Rules

As we have seen so far, fairness is not granted, and the caches should run an additional gossiping protocol to guarantee fairness while preserving optimal performance. The two exchange rules for gossiping the content provide two different (orthogonal) approaches to the problem.

Fig. 5 shows the resulting utility of caches after applying 500 iterations of gossiping between caches. We consider a random geometric graph (RGG)³ topology with $n = 50$ caches that communicate according to Assumption 2. We consider that all the edges in the network have the same rate $\sigma_{ij} = cte$ for all $(ij) \in \mathcal{E}$. Thus, each edge can be selected with the same probability and the expression in equation (16) for the weight matrix $W = \mathbb{E}[W(k)]$ reduces to:

$$W = I + \frac{1}{2m}(A - D),$$

where I is the identity matrix, A is the adjacency matrix of the graph (having entries $a_{ij} = 1$ if $(ij) \in \mathcal{E}$ and 0 otherwise), D is a diagonal matrix having the degrees of the nodes, and

³RGGs are random graphs where n nodes are placed uniformly at random in a 1 by 1 unit square area, and any two nodes within a connectivity threshold forms an edge in the graph.

m is the number of edges in the graph (i.e., $m = |\mathcal{E}|$). We just remind the reader that in our fairness model, the lower V_{max} achieved, the fairer is the content distribution.

Considering Exchange Rule 1, the expected utility curve (in red) in the figure shows the expected value (or utility) of a given cache in the network, that can be calculated using the expected weight matrix W . The dotted (light blue) curve shows instead the evolution of the actual utility of a given node. If we imagine to repeat an infinite number of times the process, and average all the value curves for the same node, then we would get exactly the red one. The solid line curve (dark blue) shows the maximum value in the network (V_{max}). This curve shows that Exchange Rule 1 does not provide fairness at any time slot because V_{max} always has high values. At the same time each cache is changing its contents regularly over time, and thus experiences different utilities such that, on average, all caches are equivalent. Fairness is only guaranteed by Exchange Rule 1 in this sense.

On the contrary, for Exchange Rule 2 the maximum value V_{max} (green curve) keeps decreasing. In fact, as we have explained in the previous section, caches are performing a sort of distributed local search, by exchanging contents only when the total objective function V_{max} can be reduced. This gives better fairness as time goes on in comparison to the initial placement (reduced from $V_{max} = 1.6$ at iteration 0 till $V_{max} = 1.05$ at iteration 500). Nevertheless, the final value to which V_{max} converges is just a local optimum. Notice that the static placement that optimizes V_{max} should give a value that lies between the red and the green curve (i.e., the optimal value in this case is $V_{max}^* \in [0.7, 1.05]$).

VI. CONCLUSION

We have presented in this paper an optimization problem for the optimal caching in mobile networks. Caching is supposed to be done at WiFi hotspots that are spread across a city. Caching is meant to offload traffic of the 3G/4G infrastructure, and hence to reduce costs for both the operator and the users. We showed that by an optimal replication and placement of contents in caches accounting for their popularity, the average cost per requested content is minimized. Replicas are then dynamically shuffled between caches to provide a balanced performance for users in different geographic places while preserving optimality for global network performance. We leveraged the theory of gossiping to realize content shuffling between caches and developed algorithms to ensure that the user cost is indifferent with respect to her/his place in the network. Up to our knowledge, this work is a first attempt to explore the problem of unfairness that might result from content placement, and that proposes shuffling of the content of caches to achieve fairness without compromising global performance. For future work, we will keep investigating this direction by focusing in particular on the effect of traffic dynamicity and the way it interacts with our developed algorithms.

ACKNOWLEDGMENT

The authors would like to thank Nicolas Nisse for his feedback regarding the NP-completeness of the MILP problem (9).

APPENDIX

PROOF OF PROPOSITION 5

To show the fairness property, let us study the dynamics and the convergence properties of the vector $\mathbf{x}(k)$ where $x_i(k)$ is its i th element. Notice that with every iteration (exchange), we have

$$\begin{cases} x_i(k+1) = x_j(k+1) = \frac{x_i(k)+x_j(k)}{2} & \text{if } (i, j) \in G(k) \\ x_i(k+1) = x_i(k) & \text{else,} \end{cases} \quad (14)$$

then we can write in matrix form

$$\mathbf{x}(k+1) = W(k)\mathbf{x}(k), \quad (15)$$

where $W(k)$ is a matrix having off-diagonal elements equal to zero except for $w_{ij} = w_{ji} = 0.5$ such that $(i, j) \in G(k)$, and all diagonal elements equal to 1 except for $w_{ii} = w_{jj} = 0.5$. Then the matrix $W(k)$ satisfies the following properties for every k : $W(k)\mathbf{1} = \mathbf{1}$, $W(k) = W(k)^T$, and $W(k)^2 = W(k)$ (where $\mathbf{1}$ is the vector of all ones). These properties are important for studying the stability of gossip algorithms (see [15], [16]). Notice that due to Assumption 2, $W(0), W(1), \dots$ are all i.i.d. matrices. Let $W = \mathbb{E}[W(k)]$, then we have

$$W = \sum_{(i,j) \in \mathcal{E}} \text{Prob}[(i, j) \in G(k)]W(k) = \sum_{(i,j) \in \mathcal{E}} \sigma_{ij}W^{(ij)}, \quad (16)$$

where $W^{(ij)}$ is a doubly stochastic matrix having off-diagonal elements equal to zero except for $w_{ij} = w_{ji} = 0.5$. Notice that W is a positive definite matrix (having dominant diagonal entries) and it is irreducible because we assume that the union graph \mathcal{G} is connected. Thus we can say that the second largest eigenvalue (in magnitude) of W , $\lambda_2(W)$, satisfies

$$\begin{cases} 0 \leq \lambda_2(W) < 1, \\ \lambda_2(W) = \lambda_1(W - J), \end{cases} \quad (17)$$

where $\lambda_1(\cdot)$ is the largest eigenvalue (in magnitude) of a matrix, and $J = \frac{1}{N}\mathbf{1}\mathbf{1}^T$. To show the convergence of equation (15) under Assumption 2 and the Exchange Rule, we define $\mathbf{y}(k) = \mathbf{x}(k) - \bar{x}\mathbf{1}$ such that $\bar{x} = \frac{1}{N}\sum_{j=1}^N V_j^0 = \frac{1}{N}\sum_{i=1}^K n_i^* U_i$. We will thus show the two properties for geographic fairness of Definition 1 by showing that $\mathbf{y}(k)$ converges to $\mathbf{0}$ almost surely. Let $\alpha_k = \|\mathbf{y}(k)\|_2^2 = \mathbf{y}(k)^T \mathbf{y}(k)$, then we have

$$\begin{aligned} \mathbb{E}[\alpha_{k+1} | \mathbf{y}(k)] &= \mathbb{E}[\mathbf{y}(k)^T W(k)^T W(k) \mathbf{y}(k) | \mathbf{y}(k)] \\ &= \mathbb{E}[\mathbf{y}(k)^T W(k) \mathbf{y}(k) | \mathbf{y}(k)] \\ &= \mathbf{y}(k)^T W \mathbf{y}(k) = \mathbf{y}(k)^T (W - J) \mathbf{y}(k) \\ &\leq \lambda_2(W) \mathbf{y}(k)^T \mathbf{y}(k). \end{aligned}$$

Since $\mathbb{E}[\alpha_{k+1}] = \mathbb{E}[\mathbb{E}[\alpha_{k+1} | \mathbf{y}(k)]]$, then by a simple recursive equation we get that

$$\mathbb{E}[\alpha_k] \leq A\lambda_2^k, \quad (18)$$

where $A = \|\mathbf{y}(0)\|_2^2$ is just a constant. Using the Markovian inequality, we get that for any $\epsilon > 0$,

$$\text{Prob}[\alpha_k > \epsilon] \leq \frac{\mathbb{E}[\alpha_k]}{\epsilon} \leq \frac{A\lambda_2^k}{\epsilon},$$

and thus

$$\sum_{k=0}^{\infty} \text{Prob}[\alpha_k > \epsilon] \leq \frac{A}{\epsilon} \sum_{k=0}^{\infty} \lambda_2^k = \frac{A}{\epsilon(1-\lambda_2)} < \infty,$$

then since the rate of convergence of $\mathbb{E}[\alpha_k]$ is fast enough ($\sum_{k=0}^{\infty} \text{Prob}[\alpha_k > \epsilon] < \infty$), then α_k converges almost surely to 0 [19, p. 138]. Therefore, $x_j(k)$ converges almost surely to \bar{x} for $j = 1, \dots, N$ and this ends the proof because all properties of geographic fairness with optimality in Definition 1 are satisfied.

REFERENCES

- [1] S. Dimatteo, P. Hui, B. Han, and V. Li, "Cellular Traffic Offloading through WiFi Networks," in *Mobile Adhoc and Sensor Systems (MASS), 2011 IEEE 8th International Conference on*, Oct 2011, pp. 192–201.
- [2] P. Deshpande, X. Hou, and S. R. Das, "Performance Comparison of 3G and Metro-scale WiFi for Vehicular Network Access," in *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, ser. IMC '10. ACM, 2010, pp. 301–307.
- [3] F. Mehmeti and T. Spyropoulos, "Is it worth to be patient? analysis and optimization of delayed mobile data offloading," in *INFOCOM, 2014 Proceedings IEEE*, April 2014, pp. 2364–2372.
- [4] A. Balasubramanian, R. Mahajan, and A. Venkataramani, "Augmenting Mobile 3G Using WiFi," in *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys '10. ACM, 2010, pp. 209–222.
- [5] J. Whitbeck, Y. Lopez, J. Leguay, V. Conan, and M. D. De Amorim, "Fast Track Article: Push-and-track: Saving Infrastructure Bandwidth Through Opportunistic Forwarding," *Pervasive Mob. Comput.*, vol. 8, no. 5, pp. 682–697, Oct. 2012.
- [6] F. Neves dos Santos, B. Ertl, C. Barakat, T. Spyropoulos, and T. Turletti, "CEDO: Content-centric Dissemination Algorithm for Delay-tolerant Networks," in *Proceedings of the 16th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, ser. MSWiM '13. ACM, 2013, pp. 377–386.
- [7] R. Bhatia, G. Narlikar, I. Rimaq, and A. Beck, "UNAP: User-Centric Network-Aware Push for Mobile Content Delivery," in *INFOCOM 2009, IEEE*, April 2009, pp. 2034–2042.
- [8] J. Erman, A. Gerber, M. Hajiaghayy, D. Pei, S. Sen, and O. Spatscheck, "To Cache or Not to Cache: The 3G Case," *IEEE Internet Computing*, vol. 15, no. 2, pp. 27–34, Mar. 2011.
- [9] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking Named Content," in *Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies (CoNEXT '09)*. ACM, 2009, pp. 1–12.
- [10] S. Gaito, D. Maggiorini, C. Quadri, and G. P. Rossi, "On the impact of a road-side infrastructure for a dtn deployed on a public transportation system," in *Proceedings of the 11th International IFIP TC 6 Conference on Networking - Volume Part II*, ser. IFIP'12, 2012, pp. 265–276.
- [11] W. Zhao, M. Ammar, and E. Zegura, "A message ferrying approach for data delivery in sparse mobile ad hoc networks," in *Proceedings of the 5th ACM MobiHoc '04*. ACM, 2004, pp. 187–198.
- [12] S. M. Ross, *Stochastic Processes (Wiley Series in Probability and Statistics)*, 2nd ed. Wiley, Feb. 1995.
- [13] T. Ibaraki and N. Katoh, *Resource Allocation Problems: Algorithmic Approaches*. Cambridge, MA, USA: MIT Press, 1988.
- [14] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms, Third Edition*, 3rd ed. The MIT Press, 2009.
- [15] F. Bénézit, A. G. Dimakis, P. Thiran, and M. Vetterli, "Order-optimal consensus through randomized path averaging," *IEEE Trans. Inf. Theor.*, vol. 56, no. 10, pp. 5150–5167, oct 2010.
- [16] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE/ACM Trans. Netw.*, vol. 14, no. SI, pp. 2508–2530, jun 2006.
- [17] H. Che, Y. Tung, and Z. Wang, "Hierarchical Web caching systems: modeling, design and experimental results," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 7, pp. 1305–1314, Sep 2002.
- [18] C. Fricker, P. Robert, and J. Roberts, "A Versatile and Accurate Approximation for LRU Cache Performance," in *Proceedings of the 24th International Teletraffic Congress*, ser. ITC '12. International Teletraffic Congress, 2012, pp. 8:1–8:8.
- [19] A. Karr, *Probability*, ser. Springer texts in statistics. Springer-Verlag, 1993.