



HAL
open science

Performance of spatial Multi-LRU caching under traffic with temporal locality

Apostolos Avranas, Anastasios Giovanidis

► **To cite this version:**

Apostolos Avranas, Anastasios Giovanidis. Performance of spatial Multi-LRU caching under traffic with temporal locality. 9th International Symposium on Turbo Codes & Iterative Information Processing 2016, Sep 2016, Brest, France. 10.1109/ISTC.2016.7593134 . hal-01339972

HAL Id: hal-01339972

<https://hal.science/hal-01339972>

Submitted on 30 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.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License

Performance of spatial Multi-LRU caching under traffic with temporal locality

Apostolos Avranas[†] and Anastasios Giovanidis*

Abstract—In this work a novel family of decentralised caching policies for wireless networks is introduced, referred to as *spatial multi-LRU*. These improve cache-hit probability by exploiting multi-coverage. Two variations are proposed, the *multi-LRU-One* and *-All*, which differ in the number of replicas inserted in the covering edge-caches. The evaluation is done under spatial traffic that exhibits temporal locality, with varying content catalogue and dependent demands. The performance metric is hit probability and the policies are compared to (1) the single-LRU and (2) an upper bound for all centralised policies with periodic popularity updates. Numerical results show the multi-LRU policies outperform both comparison policies. The reason is their passive adaptability to popularity changes. Between the *-One* and *-All* variation, which one is preferable strongly depends on the available storage space and on traffic characteristics. The performance also depends on the popularity shape.

Keywords—Wireless; Cache; LRU; Information Centric Networking; Multi-Coverage; Hit Probability; Popularity

I. INTRODUCTION

The design of today’s and future networks is characterised by a paradigm shift, from a host-centric communication architecture, towards an Information Centric Networking (ICN) one. Following this novel concept, network nodes are equipped with storage capacity where data objects can be temporarily cached and retrieved [15]. In this way information can be made available closer to the user, it can be accessed reliably [19] with minimum delay, and possibly with a quality adaptable to the users’ preferences, as envisioned in the case of multimedia files. The principal benefit is the reduction of traffic flow at the core network by serving demands from intermediate nodes [18]. The edge-nodes constitute a very important part of the ICN architecture, since it is where the wireless users directly have access to. When these nodes are equipped with storage capability, download path length is minimised [6].

In this work, we consider the wireless edge of a content centric network, which consists of a set of transmitting nodes taking fixed positions on a planar area, and a set of users dynamically arriving at this area and asking for service. The set of transmitters can refer to base stations (BSs) of a cellular network, small stations of heterogeneous networks, WIFI hotspots, or any other type of wireless nodes that can provide access to users. A user can be covered by multiple of these nodes, but she/he will be served by only one.

An important question is how to best manage the available edge-memories, in order to maximise the hit probability of

user-demands. We define the hit probability as the probability that a user will find her/his demand cached in the memory of one of the cells she/he is covered from. By managing, we mean to choose a policy that decides which objects to install in each cache and how each cache inventory is updated over time. For the edge-network with storage capability, there exists a considerable number of policies proposed in the recent literature. To find the optimal content placement that minimises downlink delay, Golrezaei et al [8] formulate a binary optimisation problem and propose approximation and greedy algorithms for its solution. Błaszczyszyn and Giovanidis [3] provide a randomised strategy that maximises the hit probability using reduced knowledge over the network topology, i.e. just the coverage probability. Poularakis et al [16] formulate a joint content placement and routing problem that maximises the fraction of requests served locally by the deployed small BSs and propose an approximation algorithm for its solution. Other proposals can be found in [5], [12].

In the above works, the proposed caching policies are centralised and result as solutions of optimisation problems. These use as input considerable system information, including exact knowledge over an assumed static catalogue of files and its popularities, as well as the network topology. Since such knowledge is not trivial to be obtained, while centralised optimisation is not very practical, Giovanidis and Avranas have proposed in [9] a novel family of distributed caching policies that profit from multi-coverage, without the need of the above information. These policies are named *spatial multi-LRU* and constitute an extension of the standard single-cache Least Recently Used (single-LRU) principle, to the case with several inter-dependent caches. The multi-LRU policies were evaluated in [9] assuming time-stationary traffic of demands, as that in [8], [3], [16]. The Independent Reference Model (IRM) [7] describes such traffic, based on which (a) there is a fixed catalogue of requested objects, (b) the popularity of each object is known and constant over time and, (c) every request is independent of previous ones. With IRM input, the authors of [9] derive approximate Che-like analytical formulas for the hit probability. Evaluation of the novel policies shows that their performance, without use of excess system information, approaches that of centralised ones.

Although the IRM offers tractability, it is not enough to describe real traffic aspects. In real networks new objects (never requested before) keep on appearing, while older ones become obsolete after some time. Furthermore, the popularity of a content does not remain constant but varies over time, and there is dependence between requests of the same object within some time horizon. All these characteristics are described by the term *temporal locality* [10], [20], [14]. Generators of such traffic have been proposed in the literature [1]. A so-called

[†]Mathematical and Algorithmic Sciences Lab, France Research Center, Huawei Technologies Co. Ltd., Arcs de Seine Bt. A, 20 quai du Point du Jour 92100 Boulogne Billancourt, France; apostolos.avranas@huawei.com

*Université Paris-Saclay, CNRS & Télécom ParisTech, LTCI lab, 23 avenue d’Italie, 75013 Paris, France; anastasios.giovanidis@telecom-paristech.fr

Shot Noise Model (SNM) is presented in [20], [11], which we make use here as a basis of our own traffic model. Under SNM the demand process is a superposition of independent Poisson processes (not necessarily homogeneous), one for each content.

Altogether, in this paper we study the behaviour of spatial multi-LRU policies under a traffic model with temporal locality. We evaluate their performance for hit probability metrics. The importance of this study lies in the significant performance gains exhibited for the novel policies under such traffic with evolving catalogue and varying file popularities. These gains are impressive not only when compared to the *single-LRU* policy, but also to the family of static policies that periodically need to track popularity changes and appropriately update their catalogue. Since, the newly proposed policies do not use any a-priori knowledge (exact or estimated) of traffic-related information, they can adapt to any type of traffic input and (when multi-coverage is available) they always achieve considerable benefits, something not possible under the same input from static/semi-static policies ([8], [3], [16], [5], [12]).

In Section II the spatial multi-LRU policies are introduced. Section III describes the wireless multi-coverage model. A significant part of the work is dedicated to the generation of spatial traffic with temporal locality and its characteristics (Section IV). The numerical evaluation is given in Section V, where a comparison with the single-LRU, and also with semi-static policies that periodically update their catalogue is provided. Final conclusions are drawn in Section VI. Calculations not included in the paper can be found in the extended ArXiv version of the work, with the same title.

II. SPATIAL MULTI-LRU

In the *single-LRU* policy each user can be connected to a single station (the one with the strongest signal) and can have access to its cache. For a given cache, we index the files placed in the inventory. The first position is called *Most Recently Used (MRU)*, and the last *Least Recently Used (LRU)*. A new request for an object triggers one out of two options. (a. Update) If the object is already in the cache it is moved to the MRU position and is downloaded immediately to the user. (b. Insertion) If the object is not in the cache it is downloaded from the core network and inserted as new at the MRU position, while the object in the LRU position is evicted.

The multi-LRU policies make use of multi-coverage in the following way. A user can check all caches of covering stations for the demanded object and can download it from anyone that has it in its inventory. Hence, they exploit the fact that by searching, the user indirectly informs each covering station over the inventory content of its neighbours. Then, cache updates and object insertions can be done in an efficient way for hit performance. Let \mathcal{M} , with cardinality m , be the set of caches (belonging to transmitters) that cover the requesting user. We introduce the following policy variations.

- **multi-LRU-One:** Action is taken only in one of the m caches. (a. Update) If \mathcal{M} is not empty and the content is found in a subset of \mathcal{M} , only one cache from the subset is used for download and, for this, the content is moved to the MRU position. (b. Insertion) If the object is not found in any of the m caches, it is inserted in only one cache from \mathcal{M} and its LRU object is evicted. This one cache can be chosen as the

closest to the user, a random one, or by some other criterion. In this work we choose the closest node.

- **multi-LRU-All:** (a. Update) If the content is found in a non-empty subset of \mathcal{M} , all caches of that subset are updated. (b. Insertion) If the object is not found in any cache of \mathcal{M} the object is inserted in all m .

III. WIRELESS MULTI-COVERAGE

We consider the following network model. The transmitter positions form a 2-dimensional square lattice (grid). To every station we allocate storage space which can contain at most K objects (irrespective of the sizes of different files. Equivalently for the model, unit size per object is assumed). Each transmitter has a possibly random area of wireless coverage associated such that, when users arrive within this area they can be served by the transmitter. Neighbouring coverage areas can overlap. Hence, a user arriving at a random location may be covered by multiple BSs, or may not be covered at all. Let p_m be the probability of a randomly located user to be covered by m stations. Of course it holds $\sum_{m=0}^{\infty} p_m = 1$. The expected number of covering stations is denoted by $\overline{N}_{bs} := \sum_{m=0}^{\infty} m p_m$.

Different coverage models define differently the shapes of coverage cells and as a consequence give different values for the coverage probabilities p_m . Specific models to be considered are the SINR model, and the SNR/Boolean model [2]. For both, the coverage cell of a station is the set of planar points for which the received signal quality exceeds some given signal-quality threshold T . The difference between them is that the SINR model refers to networks with interference (e.g. when BSs serve on the same OFDMA frequency sub-slot), whereas the SNR model, to networks that are noise-limited (e.g. by use of frequency reuse, neighbouring stations do not operate on the same bandwidth). For the Boolean model the covering cell is a disk having center its station and radius R_b . It coincides with the SNR model, when no randomness of signal fading over the wireless channel is considered.

IV. SPATIAL TRAFFIC WITH TEMPORAL LOCALITY

In what follows we give a detailed description of the traffic model. A central attribute of real traffic is the temporal evolution of the content catalogue. New objects are constantly born while others become obsolete. We denote by \mathcal{C}_t the catalogue (set) of active objects at time t , with cardinality $C(t) := |\mathcal{C}_t|$. The evolution of the catalogue size is a random process. We assume that the arrival of a new object c_i coincides with the time of its first request t_i . The time instants of first requests (arrivals) are modelled as a homogeneous Poisson Point Process (PPP) Φ on \mathbb{R} with intensity $\lambda_c > 0$ [$\frac{\text{objects}}{\text{unit-time}}$]. Here we take as unit-time 1 *day*.

A pair of variables is related to each content as an independent mark on the arrival process: (a) The first random variable denoted by T_i is the i 'th content's *lifespan*, which gives the length of time period during which it is requested by users, and after the period's end it becomes obsolete. We could allow for the realisation τ_i to take infinite values but in such option the size of the catalogue would grow indefinitely, unless the popularity of different objects tends to zero fast enough. To avoid treating such question for now, we let $\tau_i < \infty$. With such

option the catalogue size $C(t)$ fluctuates over time and remains finite. The time interval of an object is $\Delta t_i := [t_i, t_i + \tau_i)$. (b) The second random variable attached to the object c_i is the *volume* V_i i.e. the total number of requests during its lifespan. The pair of values (τ_i, v_i) per object is chosen independently of other objects and in the general case should be drawn from a joint probability distribution with a given density $f_{(T,V)}(\tau, v)$, where T and V are the generic variables. This means that in general the two variables should be dependent, to describe the correlation between the longevity of user interest towards an object and the number of times this object is requested.

To simplify the traffic model it is assumed here that T and V are random variables independent of each other, i.e. $f_{(T,V)}(\tau, v) = f_T(\tau)f_V(v)$. This simplification can help with the calculation of main traffic quantities and can allow for easier manipulation without obvious impact on the performance of the caching policies. Both variables (lifespan and volume) follow a Power-law, i.e both T and V are Pareto distributed. This is in coherence with statistical analysis of traffic measurements in the literature for the lifespan [4] and the volume V [13]. The Pareto distribution in both cases has parameter $\beta > 1$ (for the expected value to be finite), and its p.d.f. is given by (here for V) $f_V(v) = \frac{\beta V_{\min}^{\beta}}{v^{\beta+1}}$. Its expected value depends on the values of β and V_{\min} through the expression $\mathbb{E}[V] = \frac{\beta V_{\min}}{\beta-1}$. So, for the traffic generation when one of the values β or $\mathbb{E}[V]$ is determined, the other one follows. To guarantee $V \in \mathbb{N}_+$ for the samples, we choose $V_{\min} = 0.5$ and we round to get discrete values. Sampling from a Pareto distribution, generates Zipf-like distributed sizes of objects due to the Power-law behaviour.

Having sampled the lifespan and number of requests (τ_i, v_i) for a specific object arriving at t_i , it remains to determine how these requests are positioned within Δt_i . To include additional attributes of temporal locality in the traffic model, we let requests be distributed according to a finite point process (given $V < \infty$) and more specifically a (non-homogeneous) binomial point process (BPP) Ψ_i on \mathbb{R}^{v_i-1} with density function $g_i(t, t_i, \tau_i)$ over t ,

$$\Psi_i \sim \text{Binomial}(\Delta t_i, v_i - 1, g_i(t, t_i, \tau_i)). \quad (1)$$

We randomly position only $v_i - 1$ requests, because the first request always coincides with the time of content arrival t_i . The choice of the Binomial distribution further implies that requests take position independently of each other. Since $g_i(t, t_i, \tau_i)$ describes how each of the $v_i - 1$ requests is distributed within Δt_i according to the function's *shape*, the higher the value of g_i for some t , the more probable it is that a request will appear at that point. From now on, the triplet $(\Delta t_i, v_i - 1, g_i)$ will be referred to as the *popularity profile*, and the probability density function (p.d.f.) $g_i(t, t_i, \tau_i)$ as the *popularity shape* of content c_i . The popularity shape is an important aspect of the model. For some c_i it holds

$$g_i(t, t_i, \tau_i) = 0 \quad \text{for } t \notin \Delta t_i, \quad (2)$$

and

$$\int_{t_i}^{t_i + \tau_i} g_i(t, t_i, \tau_i) dt = 1. \quad (3)$$

The cumulative distribution function (c.d.f.) is

$$G_i(s, t_i, \tau_i) := \mathbb{P}(t \leq s | t_i, \tau_i) = \int_{-\infty}^s g_i(t, t_i, \tau_i) dt, \quad (4)$$

with $G_i(t_i, t_i, \tau_i) = 0$ and $G_i(t_i + \tau_i, t_i, \tau_i) = 1$. The joint p.d.f. of the $v_i - 1$ requests is their product, due to independence, $g_i^{(v_i-1)}(x_1, \dots, x_{v_i-1}, t_i, \tau_i) = g(x_1, t_i, \tau_i) \cdots g(x_{v_i-1}, t_i, \tau_i)$.

When the requests follow a homogeneous BPP for some c_i with a given Δt_i , the shape function is *uniform* and takes the expression $g_i(t, t_i, \tau_i) = \tau_i^{-1} \mathbf{1}_{\{t \in \Delta t_i\}}$. For further shape options we refer the reader to [17]. In this reference work, for finite volume per object three shapes are proposed, namely (i) the *logistic*, (ii) the *Gompertz*, and (iii) the *negative exponential*¹. Applying this to our traffic generator, when a new object arrives it is assigned a shape of index k with probability a_k , the exact value of which is a tuneable parameter.

The spatial (geographical) aspect of traffic plays an important role in influencing the performance of the policies studied here. In our work requests are uniformly positioned on a finite 2D plane. However, the traffic model can be easily extended to incorporate spatial locality e.g. by enforcing per content requests to follow a spatial 2-dimensional Gaussian distribution around some specified center.

A. Characteristic Quantities of Traffic

Based on the above description, characteristic quantities of the generated traffic can be derived (see also arXiv version).

- *Mean Catalogue Size* $\mathbb{E}[C]$. Because of the stationarity of the arrival PPP the expected number of active contents (hence catalogue size) does not depend on time $s \in \mathbb{R}_+$,

$$\mathbb{E}[C] = \mathbb{E}[N_{act}(s)] = \lambda_c \mathbb{P}(V > 1) \mathbb{E}[T]. \quad (5)$$

- *Mean Total Number of Requests within* $[0, s]$ [days],

$$N_{req}([0, s]) = s \lambda_c \mathbb{E}[V]. \quad (6)$$

- *Cache-to-(Mean)-Catalogue-Size-Ratio (CCSR)* where we omit $\mathbb{P}(V > 1)$, which is just a scaling constant,

$$\rho := \frac{K}{\lambda_c \mathbb{E}[T]}. \quad (7)$$

V. NUMERICAL EVALUATION

For the performance evaluation of the multi-LRU policies extensive simulations are performed. We consider a rectangular window where $N_{st} = 20$ stations are placed on a square lattice configuration. Assuming a Boolean coverage model every station can serve users within distance R_b from itself. The larger the R_b , the bigger the coverage area of each station, and the stronger the multi-coverage effects. Obviously R_b determines the expected number of stations covering a random user, i.e. \overline{N}_{bs} . In the figures to follow, it is \overline{N}_{bs} shown to vary rather than R_b , as the former variable is sensitive to both the choice of BS density and the coverage model characteristics.

If not mentioned otherwise, $\lambda_c = 2400$ [objects/day], with expected request volume per content $\mathbb{E}[V] = 2.1$. The distribution lifespan p.d.f. is truncated Pareto, with $T \in [\tau_{min}, \tau_{max}] = [0.1, 96]$ [days], and $\mathbb{E}[T] = 35$. The overall duration of the simulations is 15 [months]. The popularity

¹ By assuming that the lifespan of every content ends when it reaches its $1 - \varepsilon$ of its total views v_i, τ_i can be mapped to the curve parameter λ in [17], which parametrises the speed of a content's popularity change. In this paper we chose $\varepsilon = 0.02$.

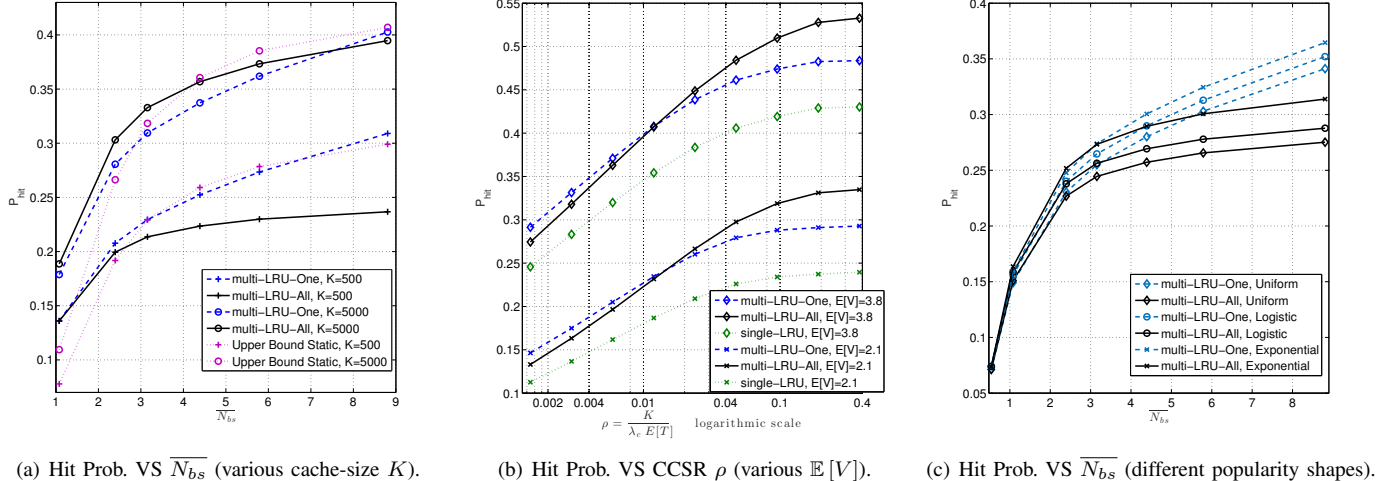


Fig. 1: Evaluation of the hit probability of multi-LRU policies for different system variables. Parameter values in (b) $\overline{N}_{bs} = 2.4$. In (c) $K = 1500$.

shape g_i of content c_i is chosen with probability $a_1 = 0.06$ as Logistic, $a_2 = 0.38$ as Gompertz and with $a_3 = 0.56$ as Exponential. Later we include the Uniform shape.

In each simulation, we keep some variables fixed and let others vary, to evaluate the policy performance and understand their influence. The variables to be varied are categorised into (1) *network* and (2) *traffic* variables. The network variables are the mean coverage number \overline{N}_{bs} and the cache size K . The traffic variables are the average request volume per content $\mathbb{E}[V]$, the mean lifespan $\mathbb{E}[T]$ and the probability vector for the shape (a_1, a_2, a_3) .

1) Network influence: The studied system has both spatial and temporal dimensions. For a hit to happen, two consecutive requests of the same content should not only be close in time, but should also occur in the coverage area of the same station. So, to increase the hit probability, data-objects should both: stay long-enough in each cache, and be inserted in as many caches as possible. But since storage space per station is limited, a trade-off arises. The multi-LRU-All which for every missed request inserts the object in all user-covering stations, spreads the object geographically wider at the cost of occupying several memory slots. For each object inserted in a station another one is evicted. As a result, in the -All variation contents stay stored for a shorter time period. On the other hand, multi-LRU-One does not spread the object to neighbouring caches, but the objects stay longer in each memory after insertion.

This trade-off between geographical expansion and preservation of a content in caches is depicted in Fig. 1(a) for different values of the memory size. The larger the storage space, the longer it takes for an object to be evicted. So, for large memory the geographical expansion of an object is beneficial. This is shown for $K = 5000$ in which case multi-LRU-All surpasses multi-LRU-One. But for each case K there is a value of \overline{N}_{bs} after which the performance of the -All is less than -One, because after this value further increase of content diversity is at the cost of content variety. The smaller the cache size, the more valuable storage space becomes because an object stays less time in the cache before eviction. Hence, for smaller K -One shows better performance and exceeds the -All variation even for small values of multi-coverage \overline{N}_{bs} .

In Fig. 1(b) the performance of multi-LRU policies versus the CCSR ratio ρ is illustrated. This ratio is equal to the mean number of memory slots per active content and is a measure of the system's storage capability, because the smaller it is than one, the less storage resources are available. Of course if it is close or larger than one, the average catalogue size can be cached entirely in each station, and hence the hit probability should be close to the "cacheability" limit (the hit probability for all objects that appear at least twice). Keeping the denominator of ρ constant, Fig. 1(b) shows the impact of the memory size on the policy performance. Obviously, hit probability increases with K , but for smaller K , the -One variation is preferable to the -All, as explained also previously. There is again a critical point in K after which the -All variation is preferable (for large storage). If the ratio ρ is further increased, the performance gains are diminishing for both variations, and saturation occurs.

2) Traffic influence: The qualitative impact of the mean lifespan value on the hit probability can be understood by reading Fig. 1(b) in the opposite direction of the x-axis, from right to left. Keeping the numerator constant, as ρ decreases $\mathbb{E}[T]$ increases. This means that the same storage capacity serves a larger active catalogue size $\mathbb{E}[C]$. Consequently the overall performance drops. Moreover, Fig. 1(b) illustrates that the hit probability improves as the average number of requests per content $\mathbb{E}[V]$ increases. The reason is that for higher $\mathbb{E}[V]$ a content put in storage is requested and hit more times.

In Fig. 1(c) each curve corresponds to a scenario where contents are assumed to follow only one particular popularity shape. Specifically, either the logistic, or the negative exponential, or the uniform shape is used. In the negative exponential shape popularity takes big values in a short time period and then drops abruptly. A steep popularity shape implies that consecutive requests of the same content appear close to each other in time. This makes more probable the event that the content is not evicted before its next request happens. With this in mind, it can be understood that the negative exponential can lead to higher hit probabilities than the uniform shape. Interestingly, for isolated caches the authors in [20], [14] state that the shape does not affect significantly the hit probability of LRU. In our model, this can be observed when \overline{N}_{bs} takes

small values so every user can connect to at most one station. This observation is confirmed in Fig 1(c). But as $\overline{N_{bs}}$ increases and multi-coverage effects appear, *the multi-LRU performance depends considerably on the correlation between requests of the same content, and thus the shape of the popularity curves.*

3) *Comparison with the single-LRU:* Under single-LRU a user can access only one (the closest in this work) station's memory even when covered by more than one. As a result hit performance is independent of $\overline{N_{bs}}$ (provided coverage is enough so that a user is always covered by at least one station). Depriving the user of the ability to retrieve its content from all covering stations, strongly reduces the overall hit probability. That's why in Fig. 1(b) where $\overline{N_{bs}} = 2.4$ is relatively small, both multi-LRU policies show relative gains compared to the single-LRU, for every value of ρ . The maximum gains reach 30% when $\mathbb{E}[V] = 2.1$ and 20% when $\mathbb{E}[V] = 3.8$.

4) *Comparison with centralised Policies with periodic Popularity updates and prefetching (POP):* As mentioned in the introduction, most caching policies proposed in the literature distribute contents to memories in a centralised way [8], [3]. These caching algorithms (POP) use as input the popularity of the contents, assumed known and constant over time. But, when the traffic exhibits temporal locality, popularities change perpetually and should be estimated regularly. Based on every new estimation, memories should be updated. Specifically we can assume that at time $t_n = n\Delta t_{ev}$, $n \in \mathbb{Z}$, the caches are updated by the POP using the estimated popularities during the time interval $[t_n - \Delta t_{pop}, t_n]$. Let \mathcal{F}_{x,t_n} be the set of the x most requested contents in $[t_n - \Delta t_{pop}, t_n]$.

We propose an upper bound. Let a user requesting content c connect to m stations with probability p_m . Then the user "sees" mK memory slots. The most favorable scenario is when all memory slots are filled at time t_n by the mK most popular files until that time, i.e. the content set \mathcal{F}_{mK,t_n} . Then the maximum hit probability within the time interval $[t_n, t_n + \Delta t_{ev}]$ is

$$P_{hit}^{(POP)}[t_n, t_n + \Delta t_{ev}] \leq \sum_{m=1}^{\infty} p_m \mathbb{P}(c \in \mathcal{F}_{mK,t_n}). \quad (8)$$

Clearly this upper bound depends on the intervals Δt_{ev} and Δt_{pop} . The more regularly the algorithm updates the caches the better performance it achieves. But considering that a caching update will use backhaul and computational resources by the controller, Δt_{ev} cannot be too small. We can assume that $\Delta t_{ev} = 1$ day and the caching policy runs every night when the request load is low. As far as Δt_{pop} is concerned there is a "crisp" optimal choice. If it is too big, \mathcal{F}_{mK,t_n} will possibly include outdated contents. On the other hand, small Δt_{pop} can result in excluding even the popular objects from \mathcal{F}_{mK,t_n} because they have not been sufficiently requested. The bound in (8) was evaluated by Monte Carlo simulations. The optimum Δt_{pop} number of days was found by increasing it until $P_{hit}^{(POP)}$ starts decreasing. Fig. 1(a) depicts $P_{hit}^{(POP)}$ for two memory sizes, where Δt_{pop} was chosen (among a large set of possibilities) equal to 5 and 10 days, for $K = 500$ and 5000 respectively.

Fig. 1(a) shows that even the upper bound for POP does not surpass in performance the appropriate multi-LRU policy, except maybe for a small range of $\overline{N_{bs}}$. In [9] where the traffic was assumed static (IRM), multi-LRU performed lower than the centralised policies POP. But under a temporal traffic

model (which is also more realistic), the ability of multi-LRU policies to update at each request the caches, without the need to estimate the content popularities, results in a considerable performance boost.

VI. CONCLUSIONS

In this work, the proposed Multi-LRU policies, that update each cache content in a per-demand basis, are evaluated under traffic with temporal locality. Their hit probability is compared to the single-LRU, as well as the family of centralised policies with periodic cache updates based on popularity estimates. It is shown that Multi-LRU outperforms single-LRU as well as the centralised policies for a large range of parameter values. The reason is the fast passive adaptability of the policies to traffic changes. The performance strongly depends on coverage and traffic parameters, and notably the content popularity shape.

REFERENCES

- [1] V. Almeida, A. Bestavros, M. Crovella, and A. de Oliveira. Characterizing reference locality in the WWW. *IEEE PDIS*, 1996.
- [2] F. Baccelli, and B. Błaszczyszyn. On a Coverage Process Ranging from the Boolean Model to the Poisson-Voronoi Tessellation with Applications to Wireless Communications. *Advances in Applied Probability*, Vol. 33, No. 2, pp. 293-323, 2001.
- [3] B. Błaszczyszyn and A. Giovanidis. Optimal geographic caching in cellular networks. *IEEE ICC*, 2015.
- [4] X. Cheng, J. Liu, and C. Dale. Understanding the Characteristics of Internet Short Video Sharing: A YouTube-Based Measurement Study. *IEEE Trans. on Multimedia*, vol. 15, no. 5, pp. 1184-1194, Aug. 2013.
- [5] M. Dehghan, A. Seetharam, B. Jiang, T. He, T. Salonidis, J. Kurose, D. Towsley, and R. Sitaraman. On the complexity of optimal routing and content caching in heterogeneous networks. *INFOCOM*, 2015.
- [6] S. E. Elayoubi and J. Roberts. Performance and cost effectiveness of caching in the mobile access network. *ACM SIGCOMM ICN*, 2015.
- [7] R. Fagin and T.G. Price. Efficient calculation of expected miss ratios in the independent reference model. *SIAM J. Comput.*, 7:288-297, 1978.
- [8] N. Golrezaei, K. Shanmugam, A. Dimakis, A. Molisch, and G. Caire. Femtocaching: Wireless video content delivery through distributed caching helpers. *INFOCOM*, 2012.
- [9] A. Giovanidis, and A. Avranas. Spatial Multi-LRU Caching for Wireless Networks with Coverage Overlaps. *ACM SIGMETRICS / IFIP Performance*, 2016.
- [10] S. Jin and A. Bestavros. Sources and characteristics of web temporal locality. *IEEE MASCOTS*, 2000.
- [11] E. Leonardi, and G.L. Torrisi. Least recently used caches under the Shot Noise Model. *INFOCOM*, 2015.
- [12] K. Naveen, L. Massoulie, E. Baccelli, A. C. Viana, and D. Towsley. On the interaction between Content Caching and Request Assignment in Cellular Cache Networks. *AllThingsCellular'15*, 2015.
- [13] M. E. J. Newman. Power laws, Pareto distributions and Zipf's law. *Contemporary Physics*, 46:323-351, 2005.
- [14] F. Olmos, B. Kauffmann, A. Simonian, and Y. Carlinet. Catalog dynamics: Impact of content publishing and perishing on the performance of a LRU cache. *26th International Teletraffic Congress*, 2014.
- [15] G. Paschos, E. Bastug, I. Land, G. Caire, and M. Debbah. Wireless Caching: Technical Misconceptions and Business Barriers. *arXiv:1602.00173*, 2016.
- [16] K. Poularakis, G. Iosifidis, and L. Tassioulas. Approximation algorithms for mobile data caching in small cell networks. *IEEE Trans. on Communications*, 62(10), October 2014.
- [17] C. Richier, E. Altman, R. Elazouzi, T. Altman, G. Linares, and Y. Portilla. Modelling View-count Dynamics in YouTube. *arXiv:1404.2570*.
- [18] V. Sourlas, P. Flegkas, G.S. Paschos, D. Katsaros, and L. Tassioulas. Storage planning and replica assignment in content-centric publish/subscribe networks. *Elsevier Computer Networks*, 55(18):4021-4032, Dec. 2011.
- [19] R. Timo, and M. Wigger. Joint cache-channel coding over erasure broadcast channels. *IEEE ISWCS*, Brussels, Belgium, 2015.
- [20] S. Traverso, M. Ahmed, M. Garetto, P. Giaccone, E. Leonardi, and S. Niccolini. Unravelling the impact of temporal and geographic locality in content caching systems. *IEEE Trans. on Multimedia*, 17(10):1839-1854, 2015.