Caching Heterogeneous Size Content in Small Cell Networks with CoMP Joint Transmissions

Guilherme Iecker Ricardo, Giovanni Neglia, Thrasyvoulos Spyropoulos

► To cite this version:

Guilherme Iecker Ricardo, Giovanni Neglia, Thrasyvoulos Spyropoulos. Caching Heterogeneous Size Content in Small Cell Networks with CoMP Joint Transmissions. GLOBECOM 2021 - IEEE Global Communications Conference, Dec 2021, Madrid, Spain. hal-03411870

HAL Id: hal-03411870

<https://hal.science/hal-03411870>

Submitted on 2 Nov 2021

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.

Caching Heterogeneous Size Content in Small Cell Networks with CoMP Joint Transmissions

Guilherme Iecker Ricardo^{1,2}, Giovanni Neglia², and Thrasyvoulos Spyropoulos¹

¹EURECOM, France, guilherme.ricardo@eurecom.fr, thrasyvoulos.spyropoulos@eurecom.fr

²Inria, Université Côte d’Azur, France, giovanni.neglia@inria.fr

Abstract—In 5G and beyond network architectures, operators and content providers base their content distribution strategies on small cell networks. On top of such networks, edge caching and Coordinated Multi-Point (CoMP) Joint Transmissions are used to improve performance. Online solutions for average delay minimization problem have been studied in the related literature, although only under the strong assumption that files have equal sizes. In this paper we aim to fill this gap and propose an online caching policy, *q*LRU-HS, that takes into account heterogeneous sizes and asymptotically converges to the optimal cache allocation under the Independent Reference Model. Our experiments confirm such convergence in practice and reveal that *q*LRU-HS outperforms other state-of-the-art solutions.

Index Terms—Edge caching, CoMP, joint transmission, heterogeneous cellular networks, optimization, distributed algorithms.

I. INTRODUCTION

Cellular data consumption has experienced an unprecedented increase. According to recent CISCO’s forecast [1], by 2023 there will be 13 billion mobile connections, showing an increase of nearly 50% over 2018. Network densification is considered a key strategy to cope with traffic increase [2]. Specifically, 3G/4G macro-cell architecture will be incremented with a large number of overlapping small cells (e.g., femto, pico), in order to improve both coverage and capacity. In a dense cellular network, each user is in general in the transmission range of many BSs and has access to the content of their caches. Cellular networks with this architecture are often called *small cell* networks.

On top of such a densified network, two additional techniques have been proposed to provide higher Quality of Experience (QoE). Assuming that every small base station (BS) has limited data storage capacity, the first technique is *caching* relevant content, e.g., the most popular content (with a higher probability of being requested). It allows users to directly access their desired content from the nearby BSs. As a consequence, the access latency as well as the backhaul congestion and servers’ load can be drastically reduced. The second technique is Coordinated Multi-Point (CoMP) *joint transmissions* [3]. The idea is that two or more BSs jointly transmit the requested file to the user. By doing so, users

experience higher rates and, consequently, smaller delays to obtain the content. The problem is to define a caching management strategy that is able to optimize the QoE taking CoMP joint transmissions into consideration.

Some related work considers *offline solutions*, where there is a centralized entity aware of the files popularities (assumed to be constant over time) and of the whole network topology [4]–[6]. With this information, it is possible to decide which files should be cached at each BS to optimize a given performance metric, e.g., the probability of finding requested files in the cache, bandwidth usage, etc. However, having all this information available is a very strong assumption and is hardly satisfied in real systems. Moreover, if files have different sizes, the problem complexity increases. Although some works propose interesting solutions to take into account heterogeneous sizes, e.g., through dynamic programming [7], bound-and-bound enumerate search [8], or greedy-based policies [9], the solution is often costly and most of the related literature lacks theoretical bounds and guarantees.

Alternatively, in the *online caching* framework, every BS employs a local caching policy that dynamically updates the local set of files reacting to the request process. Since the BSs take decisions on-the-fly, online policies are more reactive to files’ popularity short-term variability [10] and, in comparison to offline solutions, each BS needs to know and exchange much less information. For these reasons, online caching policies are more appropriate to be deployed in real systems. Some related work proposes online caching policies for small cell networks, e.g., to maximize the hit ratio [11]–[13], or minimize the average delay [14], etc. However, these studies are all based on the strong assumption that files have equal size. To the best of our knowledge, online caching policies for different file sizes have been considered only in the single cache setup [15]–[17].

In this paper, we propose an online caching policy that is able to minimize the average delay in a small cell network considering heterogeneous file sizes. In our proposed policy, BSs estimate the marginal gain per byte for keeping a copy of a cached content, calculated as the delay reduction due to the copy divided by the file size. The marginal gain per byte is used to drive the (probabilistic) caching decisions towards the optimal performance.

The main contributions of this paper are summarized as follows:

This work has been supported by the French government, through the EUR DS4H Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-17-EURE-0004 as well as the “5C-for-5G” JCJC project with the reference number ANR-17-CE25-0001.

- In Section II, we present a model that captures file retrieval delay under CoMP joint transmissions opportunities and heterogeneous files sizes.
- We define the offline optimization problem that we use as a baseline for our proposed solution in Section III. To this purpose we consider a greedy algorithm that computes a possibly infeasible caching allocation but with desirable approximation guarantees.
- We introduce our caching policy q LRU-HS in Section IV. q LRU-HS is designed for delay minimization with heterogeneous file sizes and is asymptotically optimal as its parameter q converges to 0.
- In Section V, we provide numerical results based on simulations that show our policy's convergence to the optimum when q vanishes. Then, we evaluate its performance against other policies from the related literature.

II. SYSTEM MODEL AND OPERATION

We consider a set $[B]$ of base stations (BSs) arbitrarily located in a given area $A \subseteq \mathbb{R}^2$, where $[n]$ denotes the set $\{1, \dots, n\}$, for $n \in \mathbb{Z}_+$. There is a set $[U]$ of user equipments (UEs) spread across area A . Because of the high density of BSs, each UE u will, in general, be within communication range of multiple BSs. We denote by I_u the set of UE u 's *neighboring* BSs, i.e., all BSs that have UE u within their coverage area and are able to receive requests and transmit content back to u . In order to simplify our analysis, we consider that the Signal-to-Noise Ratio (SNR) $h_u^{(b)}$ of the wireless channel between BS b and UE u is constant, i.e., $h_u^{(b)} = \bar{h} \in \mathbb{R}^+$, if u and b are connected ($b \in I_u$), and $h_u^{(b)} = 0$, otherwise.

Each BS b is equipped with a cache that can store up to $C^{(b)}$ bytes. We consider a catalog of files $[F]$, where file $f \in [F]$ has size equal to s_f bytes and is requested with probability λ_f over the area A (λ_f quantifies then the popularity of content f). In particular we assume that the request process follows the Independent Reference Model (IRM): each request is for file f with probability λ_f , independently from the past. Let $X_f^{(b)} \in \{0, 1\}$ be a variable indicating whether BS b caches file f ($X_f^{(b)} = 1$) or not ($X_f^{(b)} = 0$). Then, the vector $\mathbf{X}_f = (X_f^{(b)})_{b \in [B]}$ describes the allocation of f across the caches, and the matrix $\mathbf{X} = (X_f^{(b)})_{b \in [B], f \in [F]}$ describes the allocation of the entire catalog. Given UE u and allocation $\mathbf{X}_f$, we denote by $J_u(\mathbf{X}_f)$ the set of neighboring BSs of u that are caching f , that is $J_u(\mathbf{X}_f) = \{b \in I_u : X_f^{(b)} = 1\}$ is actually caching f .

When a UE requires a file, BSs can use CoMP techniques to jointly transmit the file. The *wireless channel access delay* [18], [19] for UE u to download f from k BSs is:

$$d_{u,f}^{\text{WC}}(k) \triangleq \frac{s_f}{w \cdot \log_2(1 + \bar{h} \cdot \min(k, |I_u|))}, \quad (1)$$

where $w \in \mathbb{R}^+$ is the channel bandwidth and the denominator is the aggregate capacity. The min operator captures the fact

that u can download from at most $|I_u|$ neighboring BSs. We consider $d_{u,f}^{\text{WC}}(0) = +\infty$.

The *backhaul-access delay* for any BS to fetch file f from the back-end servers through the backhaul network is:

$$d_f^{\text{BH}} \triangleq r + \frac{s_f}{c_{\text{BH}}}, \quad (2)$$

where c_{BH} is the backhaul network capacity and r is a constant that represents any sort of latency for accessing the back-end servers (e.g., the round-trip time in the backhaul network), henceforth generically referred to as *backhaul latency*.

When UE u wants to retrieve file f , it broadcasts an inquiry message, which is received by u 's neighboring BSs in I_u . Then, BSs in I_u estimate the download delay in two cases:

- The set of BSs caching f , i.e., $J_u(\mathbf{X}_f)$, directly transmit f to u with delay $d_{u,f}^{\text{WC}}(|J_u(\mathbf{X}_f)|)$.
- A random BS $b' \in I_u$, not caching f , fetches f from the backhaul. Then, BSs in set $J_u(\mathbf{X}_f) \cup \{b'\}$ transmit f to u with delay $d_f^{\text{BH}} + d_{u,f}^{\text{WC}}(|J_u(\mathbf{X}_f)| + 1)$.

Once the delays are estimated, the BSs proceed to transmit f to u according to the case resulting in the smallest delay.

Therefore, we define the total *end-to-end delay* experienced by UE u to download file f under allocation $\mathbf{X}_f$ as:

$$d_{u,f}(|J_u(\mathbf{X}_f)|) \triangleq \min(d_{u,f}^{\text{WC}}(|J_u(\mathbf{X}_f)|), d_f^{\text{BH}} + d_{u,f}^{\text{WC}}(|J_u(\mathbf{X}_f)| + 1)). \quad (3)$$

Note that (3) also captures the delay when *misses* at all caches occur ($J_u(\mathbf{X}_f) = \emptyset$): one neighboring BS b' will fetch the file from the backhaul and transmit it to u .

III. OPTIMIZING STATIC CACHE ALLOCATIONS

In this section, we consider the static cache allocation problem whose goal is to minimize the average end-to-end delay, assuming that probabilities $\{\lambda_f, \forall f \in [F]\}$ and UEs positions are known. Assuming these quantities are relatively stable over time, the operator could use historical data to estimate them, find the optimal allocation and then prefetch the contents to caches during low-traffic periods of the day, as considered in related works [5], [20].

In particular, if we assume that all UEs are equally likely to generate a request, the delay minimization problem can be formulated as follows:

Problem 1 (Average Delay Minimization Problem).

$$\underset{\mathbf{X}}{\text{minimize}} \quad \bar{d}(\mathbf{X}) \triangleq \sum_{f \in [F]} \lambda_f \frac{1}{U} \sum_{u \in [U]} d_{u,f}(|J_u(\mathbf{X}_f)|) \quad (4)$$

$$\text{subject to} \quad \sum_{f \in [F]} s_f \cdot X_f^{(b)} \leq C^{(b)}, \forall b \in [B], \quad (5)$$

$$X_f^{(b)} \in \{0, 1\}, \forall b \in [B], \forall f \in [F]. \quad (6)$$

The objective (4) is the average experienced delay for a request over all files and UEs and $d_{u,f}(\cdot)$ is given by (3). The set of constraints (5) guarantees that any feasible solution meets each BS's cache capacity. Problem 1 is NP-Hard because it is a generalization of the single-cache problem with capacity

(“knapsack”) constraints, which is NP-Hard [21]. For a general network setting with multiple caches, the problem is NP-hard even in the homogeneous size case [22].

A. Approximate Solution

The following greedy algorithm was introduced in [9, Algorithm 1] to solve the general submodular multiple knapsack problem (SMKP) with a $(1 - 1/e)$ approximation guarantee. Starting from empty caches ($X_f^{(b)} = 0, \forall b, f$), the algorithm iteratively finds the placement (b^*, f^*) that maximizes the ratio between the delay gain and the file size given the current cache allocation $\mathbf{X}$ and adds a copy of f^* to b^* , i.e., it sets $X_{f^*}^{(b^*)} = 1$. Whenever the placement (b^*, f^*) makes b^* 's occupancy reach or exceed its caching capacity, b^* is considered “full” and disregarded in the upcoming iterations. The algorithm stops when all BSs are “full.” From now on, we will refer to it as *Infeasible Greedy Algorithm* (IGA), since the resulting allocation is likely to violate constraints (5).

Symmetric instances of Problem 1, where BSs cover equivalent groups of UEs, can be directly mapped to a general instance of SMKP. Therefore, IGA can also be used as an approximate solution in these cases. The objective function (4) was studied in [22] as a set function and the authors proved that it is monotone and submodular for the case where SNRs are homogeneous.¹ Although IGA's solution is likely infeasible and the approximation guarantee is only valid for symmetric setups, it can still be used as a heuristic to approximate the minimum achievable average delay in general instances of Problem 1. We use this approximation as a baseline for the techniques introduced in Section IV. We show in detail in our technical report [23] how the general solution proposed in [9] can be adapted to Problem 1.

IV. ONLINE CACHING POLICIES

In this section, we assume that there is no centralized intelligence controlling the caching decisions. Instead, BSs manage their cache content on-the-fly, as new requests arrive. Consider that BSs' caches are implemented as ordered queues. Files in the cache are ordered from the most recently used one (at the *front*) to the least recently used one (at the *rear*). The cache can perform three operations: (i) *insert* a new file to the front, (ii) *evict* files from the rear, and (iii) *move-to-the-front* a file already present.

At every request (u, f) from UE u for file f , after u 's neighboring BSs serve the file, the respective caches react to that request by performing some of the operations described above. In what follows, we define a variant of q LRU [12] caching policy whose operation depends on the quantity:

$$\frac{\Delta d_{u,f}(k)}{s_f} = \frac{d_{u,f}(k-1) - d_{u,f}(k)}{s_f}, \quad (7)$$

which is the delay reduction UE u experiences thanks to the k -th copy of file f divided by file f 's size (i.e., its delay gain per byte occupied in the cache).

¹Note that, as the guarantee holds for the maximization of a non-decreasing submodular function, we need to transform the minimization in Problem 1 in an equivalent maximization problem.

We call our policy q LRU-HS as it is inspired by q LRU and takes explicitly into account files with heterogeneous sizes. We describe q LRU-HS operation as follows:

Upon a request (u, f) , given the current allocation $\mathbf{X}_f$:

- All neighboring BSs caching f ($\forall b \in J_u(\mathbf{X}_f)$) independently move f from its current position in the queue to the front with probability:

$$p_{u,f}(|J_u(\mathbf{X}_f)|) \triangleq \beta \cdot \frac{\Delta d_{u,f}(|J_u(\mathbf{X}_f)|)}{s_f}, \quad (8)$$

where constant β ensures that $p_{u,f}(\cdot) \in (0, 1]$, e.g.,

$$\beta = \min_{u', f', k' > 0} \left\{ \frac{s_{f'}}{\Delta d_{u', f'}(k')} \right\}. \quad (9)$$

- For the remaining BSs ($\forall b \in I_u \setminus J_u(\mathbf{X}_f)$): (i) If there is enough cache space, f is directly inserted at the front; (ii) otherwise, with probability q , they evict from the rear enough files to make room for f and insert it to the front. We refer to the file at the rear of the queue as f_{rear} .

We formalize q LRU-HS caching policy in Algorithm 1 from the perspective of each BS b .

Algorithm 1: q LRU-HS Caching Policy (for BS b)

Input: $w, c^{\text{BH}}, r, s_f, \mathbf{X}_f, I_u$, and $\bar{h}$.

```

1 if  $X_f^{(b)} = 1$  then
2   with probability  $p_{u,f}(|J_u(\mathbf{X}_f)|)$  in (8) do
3     Move-to-the-front  $f$ 
4   end
5 else
6   if  $C^{(b)} - \sum_{f' \in [F]} s_{f'} \cdot X_{f'}^{(b)} \geq s_f$  then
7     Insert  $f$  to the front;  $X_f^{(b)} \leftarrow 1$ 
8   else
9     with probability  $q$  do
10      while  $C^{(b)} - \sum_{f' \in [F]} s_{f'} \cdot X_{f'}^{(b)} < s_f$  do
11        Evict file  $f_{\text{rear}}$  from the rear;  $X_{f_{\text{rear}}}^{(b)} \leftarrow 0$ 
12      end
13      Insert  $f$  to the front;  $X_f^{(b)} \leftarrow 1$ 
14    end
15  end
16 end

```

Remark 1. We note that probability $p_{u,f}(\cdot)$ only depends on (i) sets $J_u(\mathbf{X}_f)$ and I_u , (ii) the file size s_f , and (iii) the average SNR $\bar{h}$. In cellular networks, UEs can measure the SNR of neighboring BSs [24] and piggyback this information in an uplink transmission from u to its neighboring BSs, with negligible overhead.²

When each cache deploys q LRU-HS, the whole network's cache allocation probabilistically changes with time as new requests arrive. Under the Characteristic Time Approximation

²The setting (9) would require each BS to be aware of the sizes of all files in the catalog. To avoid this issue in practice, every BS can estimate β on-the-fly, based on previous requests.

(CTA) [25], [26], and the Exponentialization Approximation (EA) [11], we can represent such process, as a set of coupled Continuous-Time Markov Chains (MCs). When q tends to 0, these MCs admit stationary distributions, which, in turn, correspond to the optimal solution of the continuous relaxation of Problem 1. Therefore, although BSs run the q LRU-HS policy individually, they implicitly coordinate to achieve the optimal cache allocation. This result is formalized as follows:

Proposition 1. *Under IRM, CTA, and EA, a network of q LRU-HS caches asymptotically achieves an optimal caching configuration, when $q \rightarrow 0$, even if files have different sizes.*

Proposition 1 is based on [27, Prop. IV.1], which states the optimality of another policy for the case where all files have the same size. Due to space limitations, we present the detailed proof of Proposition 1 in our technical report [23]. Here, we provide an intuitive explanation of why optimality holds.

Intuition: We observe that, as q converges to 0, the cache exhibits two different dynamics: The *insertion* of new files tends to happen more and more rarely (q converges to 0), while the frequency of *moves-to-the-front* for files already in the cache is unchanged ($p_{u,f}(\cdot)$ does not depend on q). A file f at cache b is moved to the front with a probability proportional to the placement's cost-benefit $\Delta d_{u,f}(|J_u(\mathbf{X}_f)|)/s_f$, i.e., (i) proportional to how much the file contributes to reducing the delay of that specific request and (ii) inversely proportional to how much cache space it takes. The expected number of moves-to-the-front file f experiences depends on (i) how often it is requested (λ_f) and (ii) how likely it is to be moved to the front upon a request ($p_{u,f}(\cdot)$). By the law of large numbers, the random number of moves-to-the-front will be close to its expected value and the least valuable file in the cache likely occupies the last position. We can then think that, when a new file is inserted in the cache, it will replace files that contribute the least to the decrease of the expected cost. q LRU-HS progressively replaces the least useful files from the cache, until it reaches a global minimum.

V. NUMERICAL RESULTS

In this section, we first study q LRU-HS convergence to the optimal cache allocation when q tends to 0 and then we evaluate its performance in different scenarios by comparing it against other policies from related literature, including:

- q LRU- Δd [14], it aims to minimize the average delay in a small-cell CoMP-aware setup, but considers that all files have the same size.
- *greedy-dual-size* [16], it aims to maximize the hit ratio in a single-cache setup, considering sizes are heterogeneous. We consider that *all* BSs run an instance of greedy-dual-size and react independently to each request in their cell. We refer to such operation as GDSIZE-ALL in analogy to MULTI-LRU-ALL in [13].
- IGA greedy algorithm [9], as discussed in Section III, its average delay reduction is guaranteed to be $(1 - 1/e)$ far from the optimal. Thus, we use it as a baseline for the other policies.

In our experiments, we consider the *Berlin topology*: a cellular network consisting of $B = 10$ BSs located according to the positions of T-mobile BSs in Berlin extracted from [28]. We call *network density*, ρ , the average number of BSs covering a UE and we assume that UEs are homogeneously distributed within the BSs' coverage area. In this network, all BSs have the same cache capacity, i.e., $C^{(b)} = C, \forall b \in [B]$, and can store up to $C = 50$ GB. Unless otherwise specified, we consider that the backhaul network is able to transmit data at $c^{\text{BH}} = 100$ Mbps with backhaul latency $r = 10$ ms. The wireless channel bandwidth is $w = 5$ MHz and all connected pairs BS-UE have average SNR of $\bar{h} = 10$ dB. All these values are consistent with related literature [6], [27].

In our simulations, we consider that, at every request, a file is chosen from a catalog of $F = 10^4$ files with probability determined by a Zipf law with exponent $\alpha = 0.8$. As suggested by [7], real file sizes may be represented by a truncated exponential distribution. We randomly generate the file sizes according to an exponential distribution within the interval $[s_{\min}, s_{\min} + \Delta s]$. Unless otherwise specified, we consider $s_{\min} = 1$ GB and $\Delta s = 9$ GB. We split the simulation into warm-up and measurement phases, each having 10^7 requests.

A. Convergence Analysis

According to Proposition 1, as q tends to 0, q LRU-HS converges to an optimal allocation. In our first experiments, our goal is to observe this convergence in practice. We consider the Berlin topology with density of $\rho = 5.9$ BSs/UE.

In Fig. 1, we show the average delay (left) and the hit ratio (right) versus the parameter q . As a reference, we include the result of IGA for the same setup, which is independent of parameter q . We emphasize that, although IGA may be unfeasible, its delay saving is not farther than $(1 - 1/e)$ from the optimal. As we observe in Fig. 1 (left), q LRU-HS gets closer to IGA as q decreases, suggesting its convergence to the optimal allocation. In addition to q LRU-HS results, we also plot the results for q LRU- Δd , that is also guaranteed to converge to the minimum delay as q vanishes, but only when files have all the same size [22]. However, q LRU- Δd converges to a value of average delay larger than q LRU-HS's one. This is due to fact that q LRU- Δd , while trying to minimize the delay, tends to store large files, that indeed incur large transmission delay, ignoring that they also occupy a large amount of space in the cache. In particular, given two files f_1 and f_2 with $\lambda_{f_1} > \lambda_{f_2}$ and $s_{f_2} \gg s_{f_1}$, q LRU- Δd would prefer f_2 , while our caching policy q LRU-HS correctly bias its choices in favor of f_1 that leads to a larger benefit for byte occupied in the cache. From Fig. 1 (right), we see that, for this particular scenario, better average delay is associated with a better hit ratio, which is not always necessarily the case.

In Fig. 2, we show the average delay (left) and the hit ratio (right) versus the number of requests in the simulation. For this plot, we simulate q LRU-HS and q LRU- Δd for $q = 10^{-3}$ and $q = 10^{-4}$, and we indicate the results of IGA as reference. As we observe in Fig. 2 (left), the average delay achieved by each policy decreases over time, and reaches its minimum value after about 10^6 requests (10^5 requests per BS).

Fig. 1. Average delay $\bar{d}$ (left) and hit ratio (right) versus q . $C = 50.0$ GB, $\rho = 5.9$ BSs/UE, $r = 10.0$ ms, $s_{\min} = 1.0$ GB, and $\Delta s = 9.0$ GB.

Fig. 2. Average delay $\bar{d}$ (left) and hit ratio (right) versus number of requests. Results of qLRU-HS and qLRU- Δd are shown for $q = 10^{-3}$ and $q = 10^{-4}$.

B. Performance Evaluation

Now, we compare the performance of qLRU-HS with other caching solutions in different scenarios. From now on, we consider $q = 10^{-3}$ for qLRU-HS and qLRU- Δd .

In Fig. 3, we show the performance for different values of caching capacity, ranging from $C = 10$ GB to $C = 100$ TB. We present the average delay (left) and the hit ratio (right) versus the cache capacity size C . qLRU-HS provides a more efficient management of the cache, outperforming all other policies and presenting results close to the IGA ones. The difference of performance across policies is maximal for smaller values of C , in particular, for $C = 10$ GB. qLRU-HS achieves a delay about 20% smaller than GDSIZE-ALL. As expected, when the capacity increases, all policies perform better because they can store more files and also differences reduce until all policies perform equally when the cache is so large to be able to store the whole catalog.

Fig. 3. Average delay $\bar{d}$ (left) and hit ratio (right) versus cache capacity C . $\rho = 5.9$ BSs/UE, $s_{\min} = 1$ GB, $q = 10^{-3}$, and $\Delta s = 9$ GB.

Fig. 4. Average delay $\bar{d}$ (left) and hit ratio (right) versus network density ρ . $C = 30$ GB, $s_{\min} = 1$ GB, $\Delta s = 9$ GB, and $q = 10^{-3}$.

In Fig. 4, we fix the cache capacity to $C = 30$ GB and observe the policies' performances for different levels of density, from $\rho = 1.4$ BSs/UE to $\rho = 9.1$ BSs/UE. We control the network density by simply increasing the BSs' transmission range, although we keep constant the SNR to $h = 10$ dB. In this scenario, qLRU-HS again outperforms all other policies and has results close to the IGA ones.

We observe in Fig. 4 (left) that all policies experience a delay reduction as ρ increases. The reason is that the aggregate cache available to each UE gets larger with ρ , then more files are found in the neighboring caches. Because of the larger aggregate cache, also the difference between qLRU-HS and qLRU- Δd becomes slightly smaller as ρ increases (similarly to what observed in Fig. 3). On the contrary the performance gap with GDSIZE-ALL increases: the fact that all BSs in I_u react to a request from u leads to poor coordination.

Fig. 5 shows the average delay $\bar{d}_P$ achieved by policy P normalized by the average delay $\bar{d}_{IGA}$ achieved by IGA. Results are presented for different size variability (captured by the parameter Δs), on the left, and backhaul latency r , on the right. For these experiments, we fix the network density to $\rho = 5.9$ BSs/UE. We chose to show the results in a normalized fashion due to the large excursion of $\bar{d}_P$ values when both Δs and r change.

In Fig. 5 (left) we evaluate $\bar{d}_P/\bar{d}_{IGA}$ for fixed $s_{\min} = 1$ GB and change the Δs from $\Delta s = 0$ (homogeneous file sizes) to $\Delta s = 49$ GB. We first observe that qLRU-HS and qLRU- Δd both have results close to IGA in the homogeneous size case. The more heterogeneous is the catalog, in terms of size, the more noisy is the convergence process, as the insertion of a single large file can lead to the eviction of many other files and significantly change the quality of the current allocation. This fact explains why the relative performance of all dynamic policies worsens when size variability increases. Despite the increasing trend shared by all policies, we observe that qLRU-HS is always the closest to IGA. Interestingly, although GDSIZE-ALL has the worst performance, it is less sensitive to the variability of file sizes.

Finally, one interesting aspect in our model is how the backhaul latency constant affects the policies operation and results. In Fig. 5 (right), we show $\bar{d}_P/\bar{d}_{IGA}$ when the backhaul latency increases from $r = 30$ ms to $r = 1$ s. In this case, we fixed $s_{\min} = 1$ MB and the size variability to $\Delta s = 9.0$ MB.

Fig. 5. $\bar{d}_P/\bar{d}_{IGA}$ by size variability Δs (left) and backhaul latency r (right). $C = 30.0$ GB, $q = 10^{-3}$, and $\rho = 5.9$ BSs/UE.

In this experiment, we also observe dynamic policies perform worse in comparison to IGA as the backhaul latency r increases. When r becomes larger, the optimal caching strategy changes from a scenario where it is convenient to store more copies of the same files across the BSs' caches (to create CoMP opportunities) to a scenario where file diversity across caches is preferred because it minimizes cache misses that cause the largest delay. This means that, for large enough values of backhaul latency, $qLRU-\Delta d$ and $qLRU-HS$ take an equivalent strategy, to diversify files throughout the network of caches. However, $qLRU-\Delta d$ still erroneously prefer to store large files. This leads to $qLRU-\Delta d$ storing on average less files, which decreases the hit probability and, in turn, worsens $qLRU-\Delta d$'s performance. On the contrary, GDSIZE-ALL correctly prefer the smallest files, but, as all caches react at the same time, BSs tend to have similar cache content. This replication of files throughout the BSs is suboptimal for high latency, which explains GDSIZE-ALL's worse performance.

VI. CONCLUSION

In this paper we proposed an online caching policy for average delay minimization in a small cell architecture with CoMP joint transmissions and heterogeneous file sizes. We formulated the static optimization problem for which an infeasible greedy algorithm provides approximation guarantees in the homogeneous SNR regime. Then, we introduced a novel online caching policy able to converge to the optimal caching allocation that minimizes the delay under IRM. In our experiments, we observed $qLRU-HS$'s convergence and evaluated its performance under different request processes and SNR regimes. We conclude that $qLRU-HS$ achieves considerable performance gains with negligible additional deployment complexity.

REFERENCES

- [1] CISCO, "Cisco annual internet report (2018–2023)," CISCO, Tech. Rep., March 2020.
- [2] N. Bhushan *et al.*, "Network densification: the dominant theme for wireless evolution into 5g," *IEEE Communications Magazine*, vol. 52, no. 2, pp. 82–89, Feb. 2014.
- [3] D. Lee, H. Seo, B. Clerckx, E. Hardouin, D. Mazzarese, S. Nagata, and K. Sayana, "Coordinated multipoint transmission and reception in LTE-advanced: deployment scenarios and operational challenges," *IEEE Communications Magazine*, vol. 50, no. 2, pp. 148–155, Feb. 2012.
- [4] N. Golrezaei, K. Shanmugam, A. G. Dimakis, A. F. Molisch, and G. Caire, "Femtocaching: Wireless video content delivery through distributed caching helpers," in *2012 Proceedings IEEE INFOCOM*, March 2012, pp. 1107–1115.
- [5] K. Shanmugam, N. Golrezaei, A. G. Dimakis, A. F. Molisch, and G. Caire, "Femtocaching: Wireless content delivery through distributed caching helpers," *IEEE Transactions on Information Theory*, vol. 59, no. 12, pp. 8402–8413, Dec 2013.
- [6] A. Tuholukova, G. Neglia, and T. Spyropoulos, "Optimal cache allocation for Femto helpers with joint transmission capabilities," in *IEEE ICC 2017, 21-25 May 2017, Paris, France*, Paris, France, 05 2017.
- [7] T. Mihretu Ayenew, D. Xenakis, N. Passas, and L. Merakos, "A novel content placement strategy for heterogeneous cellular networks with small cells," *IEEE Networking Letters*, vol. 2, no. 1, pp. 10–13, 2020.
- [8] T. M. Ayenew, D. Xenakis, N. Passas, and L. Merakos, "Cooperative content caching in mec-enabled heterogeneous cellular networks," *IEEE Access*, vol. 9, pp. 98 883–98 903, 2021.
- [9] X. Sun, J. Zhang, and Z. Zhang, "Deterministic algorithms for the submodular multiple knapsack problem," 2020, arXiv:2003.11450.
- [10] S. Traverso *et al.*, "Temporal Locality in Today's Content Caching: Why It Matters and How to Model It," *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 5, pp. 5–12, Nov. 2013.
- [11] E. Leonardi and G. Neglia, "Implicit coordination of caches in small cell networks under unknown popularity profiles," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 6, pp. 1276–1285, June 2018.
- [12] M. Garetto, E. Leonardi, and V. Martina, "A unified approach to the performance analysis of caching systems," *ACM Trans. Model. Perform. Eval. Comput. Syst.*, vol. 1, no. 3, pp. 12:1–12:28, May 2016.
- [13] A. Giovanidis and A. Avranas, "Spatial multi-lru caching for wireless networks with coverage overlaps," *SIGMETRICS Perform. Eval. Rev.*, vol. 44, no. 1, pp. 403–405, Jun. 2016.
- [14] G. Ricardo, G. Neglia, and T. Spyropoulos, "Caching policies for delay minimization in small cell networks with joint transmissions," in *IEEE ICC 2020, Dublin, Ireland, 06 2020*.
- [15] G. Neglia, D. Carra, M. Feng, V. Janardhan, P. Michiardi, and D. Tsigkari, "Access-time-aware cache algorithms," *ACM Trans. Model. Perform. Eval. Comput. Syst.*, vol. 2, no. 4, pp. 21:1–21:29, Nov. 2017.
- [16] Shudong Jin and A. Bestavros, "Popularity-aware greedy dual-size web proxy caching algorithms," in *Proceedings 20th IEEE International Conference on Distributed Computing Systems*, 2000, pp. 254–261.
- [17] D. S. Berger, R. K. Sitaraman, and M. Harchol-Balter, "Adaptsize: Orchestrating the hot object memory cache in a content delivery network," in *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI) 17*, 2017, pp. 483–498.
- [18] D. Tse and P. Viswanath, *Fundamentals of wireless communication*. Cambridge university press, 2005.
- [19] W. C. Ao and K. Psounis, "Distributed caching and small cell cooperation for fast content delivery," in *MobiHoc. ACM*, 2015, pp. 127–136.
- [20] K. Poularakis, G. Iosifidis, and L. Tassioulas, "Approximation algorithms for mobile data caching in small cell networks," *IEEE Transactions on Communications*, vol. 62, no. 10, pp. 3665–3677, Oct 2014.
- [21] M. Chrobak, G. J. Woeginger, K. Makino, and H. Xu, "Caching is hard—even in the fault model," *Algorithmica*, vol. 63, no. 4, pp. 781–794, 2012.
- [22] G. I. Ricardo, A. Tuholukova, G. Neglia, and T. Spyropoulos, "Caching policies for delay minimization in small cell networks with coordinated multi-point joint transmissions," *IEEE/ACM Transactions on Networking*, to appear.
- [23] G. I. Ricardo, G. Neglia, and T. Spyropoulos, "Caching heterogeneous size content in small cell networks with comp joint transmissions," Tech. Rep., 2021. [Online]. Available: <https://guilhermeir.github.io/papers/qlruhs.pdf>
- [24] S. Sesia, I. Toufik, and M. Baker, *LTE - The UMTS Long Term Evolution: From Theory to Practice*. Wiley, 2011.
- [25] H. Che, Y. Tung, and Z. Wang, "Hierarchical Web caching systems: modeling, design and experimental results," *Selected Areas in Communications, IEEE Journal on*, vol. 20, no. 7, pp. 1305–1314, Sep 2002.
- [26] R. Fagin, "Asymptotic miss ratios over independent references," *Journal of Computer and System Sciences*, vol. 14, no. 2, pp. 222 – 250, 1977.
- [27] G. Neglia, E. Leonardi, G. I. Ricardo, and T. Spyropoulos, "A swiss army knife for online caching in small cell networks," *IEEE/ACM Transactions on Networking*, pp. 1–12, 2021.
- [28] "Openmobilenetwork." [Online]. Available: openmobilenetwork.org/