

---

# LeadCache: Regret-Optimal Caching in Networks

---

**Debjit Paria\***

Department of Computer Science  
Chennai Mathematical Institute  
Chennai 603103, India  
debjit.paria1999@gmail.com

**Abhishek Sinha**

Department of Electrical Engineering  
Indian Institute of Technology Madras  
Chennai 600036, India  
abhishek.sinha@ee.iitm.ac.in

## Abstract

We consider an online prediction problem in the context of network caching. Assume that multiple users are connected to several caches via a bipartite network. At any time slot, each user may request an arbitrary file chosen from a large catalog. A user’s request at a slot is met if the requested file is cached in at least one of the caches connected to the user. Our objective is to predict, prefetch, and optimally distribute the files on the caches at each slot to maximize the total number of cache hits. The problem is non-trivial due to the non-convex and non-smooth nature of the objective function. In this paper, we propose LeadCache - an efficient online caching policy based on the Follow-the-Perturbed-Leader paradigm. We show that LeadCache is regret-optimal up to a factor of  $\tilde{O}(n^{3/8})$ , where  $n$  is the number of users. We design two efficient implementations of the LeadCache policy, one based on Pipage rounding and the other based on Madow’s sampling, each of which makes precisely one call to an LP-solver per iteration. Furthermore, with a Strong-Law-type assumption, we show that the total number of file fetches under LeadCache remains almost surely finite over an infinite horizon. Finally, we derive an approximately tight regret lower bound using results from graph coloring. We conclude that the learning-based LeadCache policy decisively outperforms the state-of-the-art caching policies both theoretically and empirically.

## 1 Introduction

We consider an online structured learning problem, called **Bipartite Caching**, that lies at the core of many large-scale internet services, including Content Distribution Networks (CDN) and Cloud Computing. Formally, a set  $\mathcal{I}$  of  $n$  users is connected to a set  $\mathcal{J}$  of  $m$  caches via a bipartite network  $G(\mathcal{I} \cup \mathcal{J}, E)$ . Each cache is connected to at most  $d$  users, and each user is connected to at most  $\Delta$  caches (see Figure 1 (b)). There is a catalog consisting of  $N$  unique files, and each of the  $m$  caches can host at most  $C$  files at a time (in practice,  $C \ll N$ ). The system evolves in discrete time slots. Each of the  $n$  users may request any file from the catalog at each time slot. The file requests could be dictated by an adversary. Given the storage capacity constraints, an online caching policy decides the files to be cached on different caches at each slot before the requests for that slot arrive. The objective is to maximize the total number of *hits* by the unknown incoming requests by coordinating the caching decisions among multiple caches in an online fashion. The **Bipartite Caching** problem is a strict generalization of the online  **$k$ -sets** problem that predicts a set of  $k$  items at each round so that the predicted set includes the item chosen by the adversary [Koolen et al., 2010, Cohen and Hazan, 2015]. However, unlike the  **$k$ -sets** problem, which predicts a single subset at a time, in this problem, we are interested in sequentially predicting multiple subsets, each corresponding to one of the caches. The interaction among the caches through the non-linear reward function makes this problem challenging.

---

\*Work done at the Indian Institute of Technology Madras as a part of the first author’s Master’s thesis.




Figure 1: Reduction of the Network Caching problem (a) to the Bipartite Caching problem (b). In this schematic, we assumed that a cache, located within two hops, is reachable to a user.

The Bipartite Caching problem is a simplified abstraction of the more general Network Caching problem central to the commercial CDNs, such as Akamai [Nygren et al., 2010], Amazon Web Services (AWS), and Microsoft Azure [Paschos et al., 2020]. In the Network Caching problem, one is given an arbitrary graph  $\mathcal{G}(V, E)$ , a set of users  $\mathcal{I} \subseteq V$ , and a set of caches  $\mathcal{J} \subseteq V$ . A user can retrieve a file from a cache only if the cache hosts the requested file. If the  $i^{\text{th}}$  user retrieves the requested file from the  $j^{\text{th}}$  cache, the user receives a reward of  $r_{ij} \geq 0$  for that slot. If the requested file is not hosted in any of the caches reachable to the user, the user receives zero rewards for that slot. The goal of a network caching policy is to dynamically place files on the caches so that cumulative reward obtained by all users is maximized. The Network Caching problem reduces to the Bipartite Caching problem when the rewards are restricted to the set  $\{0, 1\}$ . It will be clear from the sequel that the algorithms presented in this paper can be extended to the general Network Caching problem as well.

### 1.1 Problem Formulation

Denote the file requested by the  $i^{\text{th}}$  user by the one-hot encoded  $N$ -dimensional vector  $\mathbf{x}_t^i$ . In other words,  $x_{tf}^i = 1$  if the  $i^{\text{th}}$  user requests file  $f \in [N]$  at time slot  $t$ , or  $x_{tf}^i = 0$  otherwise. Since a user may request at most one file per time slot, we have:  $\sum_{f=1}^N x_{tf}^i \leq 1, \forall i \in \mathcal{I}, \forall t$ . An online caching policy prefetches files on the caches at every time slot based on past requests. Unlike classical caching policies, such as LRU, LFU, FIFO, Marker, that fetch a file immediately upon a cache-miss, we do not enforce this constraint in the problem statement. The set of files placed on the  $j^{\text{th}}$  cache at time  $t$  is represented by the  $N$ -dimensional incidence vector  $\mathbf{y}_t^j \in \{0, 1\}^N$ . In other words,  $y_{tf}^j = 1$  if the  $j^{\text{th}}$  cache hosts file  $f \in [N]$  at time  $t$ , or  $y_{tf}^j = 0$  otherwise. Due to cache capacity constraints, the following inequality must be satisfied at each time slot  $t$ :  $\sum_{f=1}^N y_{tf}^j \leq C, \forall j \in \mathcal{J}$ .

The set of all admissible caching configurations, denoted by  $\mathcal{Y} \subseteq \{0, 1\}^{Nm}$ , is dictated by the cache capacity constraints. In principle, the caching policy is allowed to replace all elements of the caches at every slot, incurring a potentially huge downloading cost over an interval. However, in Section 4, we show that the total number of files fetched to the caches under the proposed LeadCache policy remains almost surely finite under very mild assumptions on the file request process.

The  $i^{\text{th}}$  user receives a *cache hit* at time slot  $t$  if and only if *any* of the caches connected to the  $i^{\text{th}}$  user hosts the file requested by the user at slot  $t$ . In the case of a cache hit, the user obtains a unit reward. On the other hand, in the case of a *cache miss*, the user receives zero rewards for that slot. Hence, for a given aggregate request vector from all users  $\mathbf{x}_t = (\mathbf{x}_t^i, i \in \mathcal{I})$  and the aggregate cache configuration vector of all caches  $\mathbf{y}_t = (\mathbf{y}_t^j, j \in \mathcal{J})$ , the total reward  $q(\mathbf{x}_t, \mathbf{y}_t)$  obtained by the users at time  $t$  may be expressed as follows:

$$q(\mathbf{x}_t, \mathbf{y}_t) \equiv \sum_{i \in \mathcal{I}} \mathbf{x}_t^i \cdot \min \left\{ \mathbf{1}_{N \times 1}, \left( \sum_{j \in \partial^+(i)} \mathbf{y}_t^j \right) \right\}, \quad (1)$$

where  $\mathbf{a} \cdot \mathbf{b}$  denotes the inner-product of the vectors  $\mathbf{a}$  and  $\mathbf{b}$ ,  $\mathbf{1}_{N \times 1}$  denotes the  $N$ -dimensional all-one column vector, the set  $\partial^+(i)$  denotes the set of all caches connected to the  $i^{\text{th}}$  user, and the “min” operator is applied component wise. The total reward  $Q(T)$  accrued in a time-horizon of length  $T$

is obtained by summing the slot-wise rewards, *i.e.*,  $Q(T) = \sum_{t=1}^T q(\mathbf{x}_t, \mathbf{y}_t)$ . Following the standard practice in the online learning literature, we measure the performance of any online policy  $\pi$  using the notion of (static) *regret*  $R^\pi(T)$ , defined as the maximum difference in the cumulative rewards obtained by the optimal fixed caching-configuration in hindsight and that of the online policy  $\pi$ , *i.e.*,

$$R^\pi(T) \stackrel{(\text{def.})}{=} \sup_{\{\mathbf{x}_t\}_{t=1}^T} \left( \sum_{t=1}^T q(\mathbf{x}_t, \mathbf{y}^*) - \sum_{t=1}^T q(\mathbf{x}_t, \mathbf{y}_t^\pi) \right), \quad (2)$$

where  $\mathbf{y}^*$  is the best static cache-configuration in *hindsight* for the file request sequence  $\{\mathbf{x}_t\}_{t=1}^T$ , *i.e.*,  $\mathbf{y}^* = \arg \max_{\mathbf{y} \in \mathcal{Y}} \sum_{t=1}^T q(\mathbf{x}_t, \mathbf{y})$ . We assume that the file request sequence is generated by an *oblivious adversary*, *i.e.*, the entire request sequence  $\{\mathbf{x}_t\}_{t \geq 1}$  is fixed a priori. Note that the problem is *non-convex*, as we seek binary cache allocations. With an eye towards efficient implementation, later we will also consider the problem of designing efficient policies that guarantee a sub-linear  $\alpha$ -regret for a suitable value of  $\alpha < 1$  [Garber, 2021, Kakade et al., 2009, Fujita et al., 2013].

## 2 Background and Related Work

Online Linear Optimization (OLO) is a canonical online learning problem that can be formulated as a repeated game played between a learner (also known as the forecaster) and an adversary [Cesa-Bianchi and Lugosi, 2006]. In this model, at every time slot  $t$ , the policy selects an action  $\mathbf{y}_t$  from a feasible set  $\mathcal{Y} \subseteq \mathbb{R}^d$ . After that, the adversary reveals a reward vector  $\mathbf{x}_t$  from a set  $\mathcal{X} \subseteq \mathbb{R}^d$ . The adversary is assumed to be *oblivious*, *i.e.*, the sequence of reward vectors is fixed before the game begins. With the above choices, the policy receives a scalar reward  $q(\mathbf{x}_t, \mathbf{y}_t) := \langle \mathbf{x}_t, \mathbf{y}_t \rangle$  at slot  $t$ . A classic objective in this setting is to design a policy with a small regret. Follow the Perturbed Leader (FTPL), is a well-known online policy for the OLO problem [Hannan, 1957]. At time slot  $t$ , the FTPL policy adds a random noise vector  $\gamma_t$  to the cumulative reward vector  $\mathbf{X}_t = \sum_{\tau=1}^{t-1} \mathbf{x}_\tau$ , and then selects the best action against this perturbed reward, *i.e.*,  $\mathbf{y}_t := \arg \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{X}_t + \gamma_t, \mathbf{y} \rangle$ . See Abernethy et al. [2016] for a unifying treatment of the FTPL policies through the lens of stochastic smoothing.

A large number of papers on caching assume some stochastic model for the file request sequence, *e.g.*, Independent Reference Model (IRM) and Shot Noise Model (SNM) [Traverso et al., 2013]. Classic page replacement algorithms, such as MIN, LRU, LFU, and FIFO, are designed to minimize the *competitive ratio* with adversarial requests [Borodin and El-Yaniv, 2005, Van Roy, 2007, Lee et al., 1999, Dan and Towsley, 1990]. These algorithms, being *non-prefetching* in nature, replace a page on demand upon a cache-miss. However, since the competitive ratio metric is multiplicative in nature, there can be a large gap between the hit ratio of a competitively optimal policy and the optimal offline policy. To design better algorithms, the caching problem has recently been investigated through the lens of regret minimization with *prefetching* policies that *learn* from the past request sequence [Vitter and Krishnan, 1996, Krishnan and Vitter, 1998]. Daniely and Mansour [2019] considered the problem of minimizing the regret plus the switching cost for a single cache. The authors proposed a variant of the celebrated exponential weight algorithm [Littlestone and Warmuth, 1994, Freund and Schapire, 1997] that ensures the minimum competitive ratio and a small but sub-optimal regret. The Bipartite Caching model was first proposed in a pioneering paper by Shanmugam et al. [2013], where they considered a stochastic version of the problem with known file popularities. Paschos et al. [2019] proposed an Online Gradient Ascent (OGA)-based Bipartite Caching policy that allows caching a fraction of the Maximum Distance Separable (MDS)-coded files. Closely related to this paper is the recent work by Bhattacharjee et al. [2020], where the authors designed a regret-optimal single-cache policy and a Bipartite Caching policy for fountain-coded files. However, the fundamental problem of designing a regret-optimal uncoded caching policy for the Bipartite Caching problem was left as an open problem.

**Why standard approaches fail:** A straightforward way to formulate the Bipartite Caching problem is to pose it as an instance of the classic Prediction with Expert Advice problem [Cesa-Bianchi and Lugosi, 2006], where each of the possible  $\binom{N}{C}^m$  cache configurations is treated as experts. However, this approach is computationally infeasible due to the massive number of resulting experts. Furthermore, the expected FTPL policy for convex losses in Hazan [2019] and its sampled version in Hazan and Minasyan [2020] are both computationally intensive. In view of the above challenges, we now present our main technical contributions.




Figure 2: Reduction pipeline illustrating the translation of the Bipartite Caching problem with a non-linear reward function to an online learning problem with a linear reward function.

### 3 Main Results

**1. FTPL for a non-linear non-convex problem:** We propose LeadCache, a network caching policy based on the Follow the Perturbed Leader paradigm. The *non-linearity* of the reward function and the non-convexity of the feasible set (due to the integrality of cache allocations) pose a significant challenge in using the generic FTPL framework [Abernethy et al., 2016]. To circumvent this difficulty, we switch to a *virtual action* domain  $\mathcal{Z}$  where the reward function is linear. We use an anytime version of the FTPL policy for designing a virtual policy for the linearized virtual learning problem. Finally, we translate the virtual policy back to the original action domain  $\mathcal{Y}$  with the help of a mapping  $\psi : \mathcal{Z} \rightarrow \mathcal{Y}$ , obtained by solving a combinatorial optimization problem (see Figure 2 for the overall reduction pipeline).

**2. New Rounding Techniques and  $\alpha$ -regret:** The mapping  $\psi$ , which translates the virtual actions to physical caching actions in the above scheme, turns out to be an NP-hard Integer Linear Program. As our second contribution, we design a linear-time Pipage rounding technique for this problem [Ageev and Sviridenko, 2004]. Incidentally, our rounding process substantially improves upon a previous rounding scheme proposed by Shanmugam et al. [2013] in the context of caching with i.i.d. requests. Next, we propose a linear-time randomized rounding scheme that yields an efficient online policy with a provable sub-linear  $\alpha \equiv 1 - 1/e$  regret. The proposed randomized rounding algorithm exploits a classical sampling technique used in statistical surveys.

**3. Bounding the Switching Cost:** As our third contribution, we show that if the file requests are generated by a stochastic process satisfying a mild Strong Law-type property, then the caching configuration under the LeadCache policy converges to the corresponding optimal configuration *almost surely in finite time*. As new file fetches to the caches from the remote server consume bandwidth, this result implies that the proposed policy offers the best of both worlds - (1) a sub-linear regret for adversarial requests, and (2) finite downloads for “stochastically regular” requests.

**4. New Regret Lower Bound:** As our final contribution, we derive minimax regret lower bound that is tight up to a factor of  $\tilde{O}(n^{3/8})$ . Our lower bound sharpens a result in Bhattacharjee et al. [2020]. The proof of the lower bound critically utilizes graph coloring theory and the probabilistic Balls-into-Bins framework.

### 4 The LeadCache Policy

In this section, we propose LeadCache - an efficient network caching policy that guarantees near-optimal regret. Since the reward function (1) is non-linear, we *linearize* the problem by switching to a *virtual* action domain  $\mathcal{Z}$ , as detailed below.

**The Virtual Caching Problem:** First, we consider an associated Online Linear Optimization (OLO) problem, called Virtual Caching, as defined next. In this problem, at each slot  $t$ , a virtual action  $\mathbf{z}_t \equiv (\mathbf{z}_t^i, i \in \mathcal{I})$  is taken in response to the file requests received so far. The  $i^{\text{th}}$  component of the virtual action, denoted by  $\mathbf{z}_t^i \in \{0, 1\}^N$ , roughly indicates the availability of the files in the caches connected to the  $i^{\text{th}}$  user. The set of all admissible virtual actions, denoted by  $\mathcal{Z} \subseteq \{0, 1\}^{N \times n}$ , is defined below in Eqn. (4). The reward  $r(\mathbf{x}_t, \mathbf{z}_t)$  accrued by the virtual action  $\mathbf{z}_t$  for the file request vector  $\mathbf{x}_t$  at the  $t^{\text{th}}$  slot is given by their inner product, *i.e.*,

$$r(\mathbf{x}_t, \mathbf{z}_t) := \langle \mathbf{x}_t, \mathbf{z}_t \rangle = \sum_{i \in \mathcal{I}} \mathbf{x}_t^i \cdot \mathbf{z}_t^i. \quad (3)$$

**Virtual Actions:** The set  $\mathcal{Z}$  of all admissible virtual actions is defined as the set of all binary vectors  $\mathbf{z} \in \{0, 1\}^{N \times n}$  such that the following component wise inequalities hold for some admissible physical cache configuration vector  $\mathbf{y} \in \mathcal{Y}$ :

$$\mathbf{z}^i \leq \min \left\{ \mathbf{1}_{N \times 1}, \left( \sum_{j \in \partial^+(i)} \mathbf{y}^j \right) \right\}, \quad 1 \leq i \leq n. \quad (4)$$

More explicitly, the set  $\mathcal{Z}$  can be characterized as the set of all binary vectors  $\mathbf{z} \in \{0, 1\}^{N \times n}$  satisfying the following constraints for some feasible  $\mathbf{y} \in \mathcal{Y}$ :

$$z_f^i \leq \sum_{j \in \partial^+(i)} y_f^j, \quad \forall i \in \mathcal{I}, f \in [N] \quad (5)$$

$$\sum_{f=1}^N y_f^j \leq C, \quad \forall j \in \mathcal{J}, \quad (6)$$

$$y_f^j, z_f^i \in \{0, 1\}, \quad \forall i \in \mathcal{I}, \forall j \in \mathcal{J}, f \in [N]. \quad (7)$$

Let  $\psi : \mathcal{Z} \rightarrow \mathcal{Y}$  be a mapping that maps any admissible virtual action  $\mathbf{z} \in \mathcal{Z}$  to a corresponding physical caching action  $\mathbf{y}$  satisfying the condition (4). Hence, the binary variable  $z_{t,f}^i = 1$  *only if* the file  $f$  is hosted in one of the caches connected to the  $i^{\text{th}}$  user at time  $t$  in the physical configuration  $\mathbf{y}_t = \psi(\mathbf{z}_t)$ . The mapping  $\psi$  may be used to translate any virtual caching policy  $\pi^{\text{virtual}} = \{\mathbf{z}_t\}_{t \geq 1}$ , to a physical caching policy  $\pi^{\text{phy}} \equiv \psi(\pi^{\text{virtual}}) = \{\mathbf{y}_t\}_{t \geq 1}$  through the correspondence  $\mathbf{y}_t = \psi(\mathbf{z}_t), \forall t \geq 1$ . The following lemma relates the regrets incurred by these two online policies:

**Lemma 1.** *For any virtual caching policy  $\pi^{\text{virtual}}$ , define a physical caching policy  $\pi^{\text{phy}} = \psi(\pi^{\text{virtual}})$  as above. Then the regret of the policy  $\pi^{\text{phy}}$  is bounded above by that of the policy  $\pi^{\text{virtual}}$ , i.e.,*

$$R_T^{\pi^{\text{phy}}} \leq R_T^{\pi^{\text{virtual}}}, \quad \forall T \geq 1.$$

Please refer to Section 10.1 in the supplementary material for the proof of Lemma 1. Lemma 1 implies that any low-regret virtual caching policy may be used to design a low-regret physical caching policy using the non-linear mapping  $\psi(\cdot)$ . The key advantage of the virtual caching problem is that it is a standard OLO problem. Hence, in our proposed LeadCache policy, we use an anytime version of the FTPL policy for solving virtual caching problem. The overall LeadCache policy is described below in Algorithm 1:

---

**Algorithm 1** The LeadCache Policy

---

- 1:  $\mathbf{X}(0) \leftarrow \mathbf{0}$
  - 2: Sample  $\boldsymbol{\gamma} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \mathbf{1}_{Nn \times 1})$
  - 3: **for**  $t = 1$  to  $T$  **do**
  - 4:    $\mathbf{X}(t) \leftarrow \mathbf{X}(t-1) + \mathbf{x}_t$
  - 5:    $\eta_t \leftarrow \frac{n^{3/4}}{(2d(\log \frac{N}{C} + 1))^{1/4}} \sqrt{\frac{t}{Cm}}$
  - 6:    $\boldsymbol{\Theta}(t) \leftarrow \mathbf{X}(t) + \eta_t \boldsymbol{\gamma}$
  - 7:    $\mathbf{z}_t \leftarrow \max_{\mathbf{z} \in \mathcal{Z}} (\boldsymbol{\Theta}(t), \mathbf{z})$ .
  - 8:    $\mathbf{y}_t \leftarrow \psi(\mathbf{z}_t)$ .
  - 9: **end for**
- 

In Algorithm 1, the flattened  $Nn \times 1$  dimensional vector  $\mathbf{X}(t)$  denotes the cumulative count of the file requests (for each (user, file) tuple), and the vector  $\boldsymbol{\Theta}(t) = (\boldsymbol{\theta}^i(t), 1 \leq i \leq n)$  denotes the perturbed cumulative file request counts obtained upon adding a scaled i.i.d. Gaussian noise vector to  $\mathbf{X}(t)$ . It is also possible to sample a fresh Gaussian vector  $\boldsymbol{\gamma}_t$  in step 6 at every time slot, leading to a high probability regret bound [Devroye et al., 2015]. The following Theorem gives an upper bound on the regret achieved by LeadCache:

**Theorem 1.** *The expected regret of the LeadCache policy is upper bounded as:*

$$\mathbb{E}(R_T^{\text{LeadCache}}) \leq \kappa n^{3/4} d^{1/4} \sqrt{mCT},$$

where  $\kappa = O(\text{poly-log}(N/C))$ , and the expectation is taken with respect to the random noise added by the policy.

Note that in contrast with the generic regret bound of Suggala and Netrapalli [2020] for non-convex problems, our regret-bound has only logarithmic dependence on the ambient dimension  $N$ .

**Proof outline:** Our analysis of the LeadCache policy uses the elegant stochastic smoothing framework developed in [Abernethy et al. \[2016, 2014\]](#), [Cohen and Hazan \[2015\]](#), [Lee \[2018\]](#). See Section 10.2 of the supplementary material for the proof of Theorem 1.

#### 4.1 Fast approximate implementation

The computationally intensive procedures in Algorithm 1 are (I) solving the optimization problem in step 7 to determine the virtual caching actions and (II) translating the virtual actions back to the physical caching actions in step (8). Since the perturbed vector  $\Theta(t)$  is obtained by adding white Gaussian noise to the cumulative request vector  $\mathbf{X}(t)$ , some of its components could be negative. For maximizing the objective (7), it is clear that if some coefficient  $\theta_f^i(t)$  is negative for some  $(i, f)$  tuple, it is feasible and optimal to set the virtual action variable  $z_f^i$  to zero. Hence, steps (7) and (8) of Algorithm 1 may be combined as:

$$\mathbf{y}(t) \leftarrow \arg \max_{\mathbf{y} \in \mathcal{Y}} \underbrace{\sum_{i \in \mathcal{I}, f \in [N]} (\theta_f^i(t))^+ \left( \min \left( 1, \sum_{j \in \partial^+(i)} y_f^j \right) \right)}_{L(\mathbf{y})}, \quad (8)$$

where  $x^+ \equiv \max(0, x)$ . Incidentally, we find that problem (8) is mathematically identical to the uncoded *Femtocaching* problem with known file request probabilities studied by [Shanmugam et al. \[2013, Section III\]](#). In the same paper, the authors proved the problem (8) to be **NP-Hard**. The authors also proposed a complex iterative rounding method for the LP relaxation of the problem, where, in each iteration, one needs to compute certain matchings. We now propose two simple *linear-time* rounding techniques that enjoy the same approximation guarantee.

**LP Relaxation:** We now introduce a new set of variables  $z_f^i := \min(1, \sum_{j \in \partial^+(i)} y_f^j)$ ,  $\forall i, f$ , and relax the integrality constraints to arrive at the following LP:

$$\max \sum_{i, f} (\theta_f^i(t))^+ z_f^i, \quad (9)$$

Subject to,

$$z_f^i \leq \sum_{j \in \partial^+(i)} y_f^j, \quad \forall i \in \mathcal{I}, f \in [N] \quad (10)$$

$$\sum_{f=1}^N y_f^j \leq C, \quad \forall j \in \mathcal{J}, \quad (11)$$

$$0 \leq y_f^j \leq 1, \quad \forall j \in \mathcal{J}, f \in [N]; 0 \leq z_f^i \leq 1, \quad \forall i \in \mathcal{I}, f \in [N]. \quad (12)$$

Denote the objective function for the problem (8) by  $L(\mathbf{y})$  and its optimal value (over  $\mathcal{Z}$ ) by  $\text{OPT}$ . Let  $\mathbf{y}^*$  be an optimal solution to the relaxed LP (9) and  $\mathcal{Z}_{\text{rel}}$  be the corresponding relaxed feasible set. Since LP (9) is a relaxation to (8), it naturally holds that  $L(\mathbf{y}^*) \geq \text{OPT}$ . To round the resulting cache allocation vector  $\mathbf{y}^*$  to an integral one, we consider the following two rounding schemes - (1) Deterministic Pipage rounding and (2) Randomized sampling-based rounding.

**1. Pipage Rounding:** The general framework of *Pipage rounding* was introduced by [Ageev and Sviridenko \[2004\]](#). Our rounding technique, given in Algorithm 2, is markedly simpler compared to Algorithm 1 of [Shanmugam et al. \[2013\]](#). While we round two fractional allocations of a single cache at a time, the rounding procedure of [Shanmugam et al. \[2013\]](#) jointly rounds several allocations in multiple caches at the same time by computing matchings in a bipartite graph.

**Design:** The key to our deterministic rounding procedure is to consider the following surrogate objective function  $\phi(\mathbf{y})$  instead of the original objective  $L(\mathbf{y})$  as given in Eqn. (8):

$$\phi(\mathbf{y}) \equiv \sum_{i, f} (\theta_f^i(t))^+ \left( 1 - \prod_{j \in \partial^+(i)} (1 - y_f^j) \right). \quad (15)$$

Following a standard algebraic argument [[Ageev and Sviridenko, 2004](#), Eqn. (16)], we have:

$$L(\mathbf{y}) \stackrel{(a)}{\geq} \phi(\mathbf{y}) \geq \left( 1 - \left( 1 - \frac{1}{\Delta} \right)^\Delta \right) L(\mathbf{y}), \quad \forall \mathbf{y} \in [0, 1]^N, \quad (16)$$

---

**Algorithm 2** Cache-wise Deterministic Pipage rounding
 

---

- 1:  $\mathbf{y} \leftarrow$  Solution of the LP (9).
- 2: **while**  $\mathbf{y}$  is not integral **do**
- 3:   Select a cache  $j$  with two fractional variables  $y_{f_1}^j$  and  $y_{f_2}^j$ .
- 4:   Set  $\epsilon_1 \leftarrow \min(y_{f_1}^j, 1 - y_{f_2}^j)$ ,  $\epsilon_2 \leftarrow \min(1 - y_{f_1}^j, y_{f_2}^j)$ .
- 5:   Define two new feasible cache-allocation vectors  $\alpha, \beta$  as follows:

$$\alpha_{f_1}^j \leftarrow y_{f_1}^j - \epsilon_1, \alpha_{f_2}^j \leftarrow y_{f_2}^j + \epsilon_1, \quad \text{and} \quad \alpha_f^k \leftarrow y_f^k, \text{ otherwise,} \quad (13)$$

$$\beta_{f_1}^j \leftarrow y_{f_1}^j + \epsilon_2, \beta_{f_2}^j \leftarrow y_{f_2}^j - \epsilon_2, \quad \text{and} \quad \beta_f^k \leftarrow y_f^k, \text{ otherwise.} \quad (14)$$

- 6:   Set  $\mathbf{y} \leftarrow \arg \max_{\mathbf{x} \in \{\alpha, \beta\}} \phi(\mathbf{x})$ .
  - 7: **end while**
  - 8: **return**  $\mathbf{y}$ .
- 

where  $\Delta \equiv \max_{i \in \mathcal{I}} |\partial^+(i)|$ . Note that inequality (a) holds with equality for all binary vectors  $\mathbf{y} \in \{0, 1\}^{mN}$ . Our Pipage rounding procedure, given in Algorithm 2, begins with an optimal solution of the LP (9). Then it iteratively perturbs two fractional variables (if any) in a single cache in such a way that the value of the surrogate objective function  $\phi(\mathbf{y})$  never decreases while at least one of the two fractional variables is rounded to an integer. Step (4) ensures that the feasibility is maintained at every step of the roundings. Upon termination (which occurs within  $O(mN)$  steps), the rounding procedure yields a feasible integral allocation vector  $\hat{\mathbf{y}}$  with an objective value  $L(\hat{\mathbf{y}})$ , which is within a factor of  $1 - (1 - \frac{1}{\Delta})^\Delta$  of the optimum objective. The following theorem formalizes this claim.

**Theorem 2.** *Algorithm 2 is an  $\alpha = 1 - (1 - \frac{1}{\Delta})^\Delta$  approximation algorithm for the problem 8.*

See Section 11 of the supplementary material for the proof of Theorem 2. Note that the Pipage rounding procedure, although effective in practice, is not known to have a formal regret bound.

**2. An efficient policy for achieving a sub-linear  $\alpha$ -regret:** Since the offline problem (8) is **NP-Hard**, a natural follow-up problem is to design a policy with a sub-linear  $\alpha$ -regret. Recall that  $\alpha$ -regret is defined similarly as the usual static regret where the reward accrued by the offline oracle policy (*i.e.*, the first term in Eqn. (2)) is discounted by a factor of  $\alpha$  [Kalai and Vempala, 2005]. Note that directly using the Pipage-rounded solution from Algorithm 2 does not necessarily yield an online policy with a sub-linear  $\alpha$ -regret [Kakade et al., 2009, Garber, 2021, Hazan et al., 2018, Fujita et al., 2013]. In the following, we give an efficient offline-to-online reduction that makes only a *single* query to a linear-time randomized rounding procedure per iteration. In our reduction, the following notion of an  $\alpha$  point-wise approximation algorithm [Kalai and Vempala, 2005] plays a pivotal role.

**Definition 1** ( $\alpha$  point-wise approximation). *For a feasible set  $\mathcal{Z}$  in the non-negative orthant and a non-negative input vector  $\mathbf{x}$ , consider the Integer Linear Program  $\max_{\mathbf{z} \in \mathcal{Z}} \mathbf{z} \cdot \mathbf{x}$ . Let  $\mathcal{Z}_{\text{rel}} \supseteq \mathcal{Z}$  be a relaxation of the feasible set  $\mathcal{Z}$  and let  $\mathbf{z} \in \arg \max_{\mathbf{z} \in \mathcal{Z}_{\text{rel}}} \mathbf{z} \cdot \mathbf{x}$  be an optimal solution of the relaxed ILP. If for some  $\alpha > 0$ , a (randomized) rounding algorithm  $A$  returns a feasible solution  $\hat{\mathbf{z}} \in \mathcal{Z}$  such that  $\mathbb{E} \hat{z}_i \geq \alpha z_i, \forall i$  and for any input  $\mathbf{x}$ , we call the algorithm  $A$  an  $\alpha$  point-wise approximation.*

It immediately follows that for any  $\alpha$  point-wise approximation algorithm for the problem 9, if  $\mathbf{z}_t \in \mathcal{Z}_{\text{rel}}$  be the relaxed virtual action at time  $t$ , we have  $\sum_{t=1}^T \mathbb{E}[\hat{\mathbf{z}}_t] \cdot \mathbf{x}_t \geq \alpha \sum_{t=1}^T \mathbf{z}_t \cdot \mathbf{x}_t$ , where the inequality follows from the point-wise approximation property. Thus, for any  $\mathbf{z}^* \in \mathcal{Z}$ , the  $\alpha$ -regret of the virtual policy may be upper bounded as

$$\alpha \sum_{t=1}^T \mathbf{x}_t \cdot \mathbf{z}^* - \sum_{t=1}^T \mathbb{E}[\hat{\mathbf{z}}_t] \cdot \mathbf{x}_t \leq \alpha \left( \sum_{t=1}^T \mathbf{x}_t \cdot \mathbf{z}^* - \sum_{t=1}^T \mathbf{x}_t \cdot \mathbf{z}_t \right) \stackrel{(b)}{\leq} \alpha \mathbb{E}(\tilde{R}_T^{\text{LeadCache}}), \quad (17)$$

where  $\tilde{R}_T^{\text{LeadCache}}$  is an upper bound to the regret of the LeadCache policy with the relaxed actions given by the solution of the LP (9). We bound the quantity  $\mathbb{E}(\tilde{R}_T^{\text{LeadCache}})$  in the following Proposition.

**Proposition 1.** *For an appropriate learning rate sequence  $\{\eta_t\}_{t \geq 1}$ , the expected regret of the LeadCache policy with the relaxed action set  $\mathcal{Z}_{\text{rel}}$  can be bounded as follows:*

$$\mathbb{E}(\tilde{R}_T^{\text{LeadCache}}) \leq \kappa_1 n^{3/4} \sqrt{dmCT},$$

where  $\kappa_1$  is poly-logarithmic in  $N$  and  $n$ .

See Section 13 of the supplementary materials for a proof sketch of the above result. In the following, we design an  $\alpha$  point-wise approximate randomized rounding scheme.

**Randomized Rounding via Madow’s sampling:** The key ingredient to our  $\alpha$  point-wise approximation oracle is Madow’s systematic sampling scheme taken from the statistical sampling literature [Madow et al., 1949]. For a set of *feasible* inclusion probabilities  $\mathbf{p}$  on a set of items  $[N]$ , Madow’s scheme outputs a subset  $S$  of size  $C$  such that the  $i^{\text{th}}$  element is included in the subset  $S$  with probability  $p_i$ ,  $1 \leq i \leq N$ . For this sampling scheme to work, it is necessary and sufficient that the inclusion probability vector  $\mathbf{p}$  satisfies the following feasibility constraint:

$$\sum_{i=1}^N p_i = C, \text{ and } 0 \leq p_i \leq 1, \forall i \in [N]. \quad (18)$$

The pseudocode for Madow’s sampling is given in Section 12 in the supplement. Our proposed  $\alpha$  point-wise approximate rounding scheme independently samples  $C$  files in each cache in accordance with the inclusion probability given by the fractional allocation vector  $\mathbf{y}$  obtained from the solution of the relaxed LP (9). From the constraints of the LP, it immediately follows that the inclusion vector  $\mathbf{y}^j$  satisfies the feasibility constraint (18); hence the above process is sound. The overall rounding scheme is summarized in Algorithm 3. To show that the resulting rounding scheme satisfies the  $\alpha$  point-wise approximation property, note that for all  $i \in \mathcal{I}$  and  $f \in [N]$ :

$$\mathbb{P}(\hat{z}_f^i = 1) = \mathbb{P}\left(\bigvee_{j \in \partial^+(i)} \hat{y}_f^j = 1\right) \stackrel{(a)}{=} 1 - \prod_{j \in \partial^+(i)} (1 - y_f^j) \stackrel{(b)}{\geq} 1 - e^{-\sum_{j \in \partial^+(i)} y_f^j} \stackrel{(c)}{\geq} 1 - e^{-z_f^i} \stackrel{(d)}{\geq} \left(1 - \frac{1}{e}\right) z_f^i,$$

where the equality (a) follows from the fact that rounding in each caches are done independently of each other, the inequality (b) follows from the standard inequality  $\exp(x) \geq 1 + x$ ,  $\forall x \in \mathbb{R}$ , the inequality (c) follows from the feasibility constraint (10) of the LP, and finally the inequality (d) follows from the concavity of the function  $1 - \exp(-x)$  and the fact that  $0 \leq z_f^i \leq 1$ . Since  $\hat{z}_f^i$  is binary, it immediately follows that the randomized rounding scheme in Algorithm 3 is an  $\alpha$  point-wise approximation with  $\alpha = 1 - 1/e$ . We formally state the result in the following Theorem.

**Theorem 3.** *The LeadCache policy, in conjunction with the randomized rounding scheme with Madow’s sampling (Algorithm 3), achieves an  $\alpha = 1 - e^{-1}$ -regret bounded by  $\tilde{O}(n^{3/4} \sqrt{dmCT})$ .*

---

#### Algorithm 3 Randomized Rounding with Madow’s Sampling Scheme

---

**Input:** Fractional cache allocation  $(\mathbf{y}, \mathbf{z})$ .

**Output:** A rounded allocation  $(\hat{\mathbf{y}}, \hat{\mathbf{z}})$

- 1: Let  $\mathbf{y}$  be the output of the LP (9).
  - 2: **for** each cache  $j \in \mathcal{J}$  **do**
  - 3:   Independently sample  $S_j$  a set of  $C$  files with the probability vector  $\mathbf{y}^j$  using Madow’s sampling scheme 4.
  - 4:    $\hat{y}_f^j \leftarrow 1(f \in S_j)$ .
  - 5: **end for**
  - 6:  $\hat{z}_f^i \leftarrow \bigvee_{j \in \partial^+(i)} \hat{y}_f^j, \forall i \in \mathcal{I}$ .
  - 7: **Return**  $(\hat{\mathbf{y}}, \hat{\mathbf{z}})$
- 

## 5 Bounding the Number of Fetches

Fetching files from the remote server to the local caches consumes bandwidth and increases network congestion. Under non-prefetching policies, such as LRU, FIFO, and LFU, a file is fetched if and only if there is a cache miss. Hence, for these policies, it is enough to bound the cache miss rates in order to control the download rate. However, since the LeadCache policy decouples the fetching process from the cache misses, in addition to a small regret, we need to ensure that the number of file fetches remains small as well. We now prove the surprising result that if the file request process satisfies a mild regularity property, the file fetches *stop almost surely after a finite time*. Note that our result is of a different flavor from the long line of work that minimizes the switching regret under adversarial inputs but yield a much weaker bound on the number of fetches [Mukhopadhyay and Sinha, 2021, Devroye et al., 2015, Kalai and Vempala, 2005, Geulen et al., 2010].



**A. Stochastic Regularity Assumption:** Let  $\{\mathbf{X}(t)\}_{t \geq 1}$  be the cumulative request-arrival process. We assume that, there exists a set of non-negative numbers  $\{p_f^i\}_{i \in \mathcal{I}, f \in [N]}$  such that for any  $\epsilon > 0$  :

$$\sum_{t=1}^{\infty} \sum_{i \in \mathcal{I}, f \in [N]} \mathbb{P}\left(\left|\frac{\mathbf{X}_f^i(t)}{t} - p_f^i\right| \geq \epsilon\right) < \infty. \quad (19)$$

Using the first Borel-Cantelli Lemma, the regularity assumption **A** implies that the process  $\{\mathbf{X}(t)\}_{t \geq 1}$  satisfies the strong-law:  $\mathbf{X}_f^i(t)/t \rightarrow p_f^i$ , a.s.,  $\forall i \in \mathcal{I}, f \in [N]$ . However, the converse may not be true. Nevertheless, the assumption **A** is quite mild and holds, *e.g.*, when the file request sequence is generated by a renewal process having an inter-arrival distribution with a finite fourth moment (See Section 15 of the supplementary material for the proof). Define a fetch event  $F(t)$  to take place at time slot  $t$  if the cache configuration at time slot  $t$  is different from that of at time slot  $t - 1$ , *i.e.*,  $F(t) := \{\mathbf{y}(t) \neq \mathbf{y}(t - 1)\}$ . The following is our main result in this section:

**Theorem 4.** *Under the stochastic regularity assumption **A**, the file fetches to the caches stop after a finite time with probability 1 under the LeadCache policy.*

Please refer to Section 14 of the supplementary material for the proof of Theorem 4. This Theorem implicitly assumes that the optimization problems in steps 7 and 8 are solved exactly at every time.

## 6 A Minimax Lower Bound

In this section, we establish a minimax lower bound to the regret for the Bipartite Caching problem. Recall that our reward function (1) is non-linear. As such, the standard lower bounds, such as Theorem 5.1 and 5.12 of Orabona [2019], Theorem 5 of Abernethy et al. [2008] are insufficient for our purpose. The first regret lower bound for the Bipartite Caching problem was given by Bhattacharjee et al. [2020]. We now strengthen this bound by utilizing results from graph coloring.

**Theorem 5** (Regret Lower Bound). *For a catalog of size  $N \geq \max(2 \frac{d^2 C m}{n}, 2mC)$  the regret of any online caching policy  $\pi$  is lower bounded as:  $R_T^\pi \geq \max(\sqrt{\frac{mnCT}{2\pi}}, d\sqrt{\frac{mCT}{2\pi}}) - \Theta(\frac{1}{\sqrt{T}})$ .*

**Proof outline:** We use the standard probabilistic method where the worst-case regret is lower bounded by the average regret over an ensemble of problems. To lower-bound the cumulative reward accrued by the optimal offline policy, we construct an offline static caching configuration, where each user sees different files on each of the connected caches (*local exclusivity*). This construction effectively linearizes the reward function that can be analyzed using the probabilistic *Balls-into-Bins* framework and graph coloring theory. See Section 16 of the supplementary material for the proof.

**Approximation guarantee for LeadCache:** Theorem 1 and Theorem 5, taken together, imply that the LeadCache policy achieves the optimal regret within a factor of  $\tilde{O}(\min((nd)^{1/4}, (n/d)^{3/4}))$ . Hence, irrespective of  $d$ , the LeadCache policy is regret-optimal up to a factor of  $\tilde{O}(n^{3/8})$ .

## 7 Experiments

In this section, we compare the performance of the LeadCache policy with standard caching policies.

**Baseline policies:** Under the LRU policy [Borodin and El-Yaniv, 2005], each cache considers the set of all requested files from its connected users independently of other caches. In the case of a *cache-miss*, the LRU policy fetches the requested file into the cache while evicting a file that was requested *least recently*. The LFU policy works similarly to the LRU policy with the exception that, in the case of a cache-miss, the policy evicts a file that was requested the *least number of times* among all files currently on the cache. Finally, for the purpose of benchmarking, we also experiment with Belady's offline MIN algorithm [Aho et al., 1971], which is optimal in the class of non-pre-fetching reactive policies for each *individual* cache when the entire file request sequence is known a priori. Note that Belady's algorithm is *not* optimal in the network caching setting as it does not consider the adjacency relations between the caches and the users. The heuristic multi-cache policy by Bhattacharjee et al. [2020] uses an FTPL strategy for file prefetching while approximating the reward function (1) by a linear function.

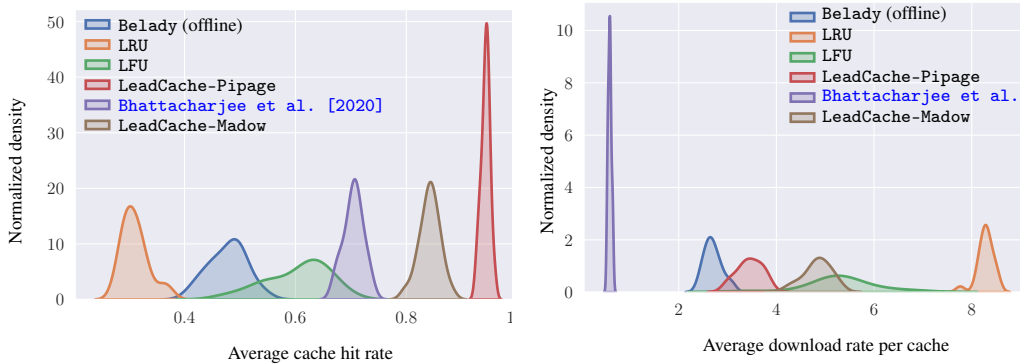


Figure 3: Empirical distributions of (a) Cache hit rates and (b) Fetch rates of different caching policies.

**Experimental Setup:** In our experiments, we use a publicly available anonymized production trace from a large CDN provider available under a BSD 2-Clause License [Berger et al., 2018, Berger, 2018]. The trace data consists of three fields, namely, request number, file-id, and file size. In our experiments, we construct a random Bipartite caching network with  $n = 30$  users and  $m = 10$  caches. Each cache is connected to  $d = 8$  randomly chosen users. Thus, every user is connected to  $\approx 2.67$  caches on average. The storage capacity  $C$  of each cache is taken to be 10% of the catalog size. We divide the trace consisting of the first  $\sim 375K$  requests into 20 consecutive sub-intervals. File requests are assigned to the users sequentially before running the experiments on each of the sub-intervals. The code for the experiments is available online [Paria and Sinha, 2021].

**Results and Discussion:** Figure 3 shows the performance of different caching policies in terms of the average cache hit rate per file and the average number of file fetches per cache. From the plots, it is clear that the LeadCache policy, with any of its Pipage and Madow rounding variants, outperforms the baseline policies in terms of the cache hit rate. Furthermore, we find that the deterministic Pipage rounding variant empirically outperforms the randomized Madow rounding variant (that has a guaranteed sublinear  $\alpha$ -regret) in terms of both hit rate and fetch rates. On the flip side, the heuristic policy of Bhattacharjee et al. [2020] incurs a fewer number of file fetches compared to the LeadCache policy. From the plots, it is clear that the LeadCache policy excels by effectively coordinating the caching decisions among different caches and quickly adapting to the dynamic file request pattern. Section 17 of the supplementary material gives additional plots for the popularity distribution and the temporal dynamics of the policies for a given file request sequence.

## 8 Conclusion and Future Work

In this paper, we proposed an efficient network caching policy called LeadCache. We showed that the policy is competitive with the state-of-the-art caching policies, both theoretically and empirically. We proved a lower bound for the achievable regret and established that the number of file-fetches incurred by our policy is finite under reasonable assumptions. Note that LeadCache optimizes the cumulative cache hits without considering fairness among the users. In particular, LeadCache could potentially ignore a small subset of users who request unpopular content. A future research direction could be to incorporate fairness into the LeadCache policy so that *each* user incurs a small regret [Destounis et al., 2017]. Finally, it will be interesting to design caching policies that enjoy strong guarantees for the regret and the competitive ratio simultaneously [Daniely and Mansour, 2019].

## 9 Acknowledgement

This work was partially supported by the grant IND-417880 from Qualcomm, USA, and a research grant from the Govt. of India under the IoE initiative. The computational work reported on this paper was performed on the AQUA Cluster at the High Performance Computing Environment of IIT Madras. The authors would also like to thank Krishnakumar from IIT Madras for his help with a few illustrations appearing on the paper.

## References

- Wouter M Koolen, Manfred K Warmuth, Jyrki Kivinen, et al. Hedging structured concepts. In *COLT*, pages 93–105. Citeseer, 2010.
- Alon Cohen and Tamir Hazan. Following the perturbed leader for online structured learning. In *International Conference on Machine Learning*, pages 1034–1042, 2015.
- Erik Nygren, Ramesh K. Sitaraman, and Jennifer Sun. The akamai network: A platform for high-performance internet applications. *SIGOPS Oper. Syst. Rev.*, 44(3):2–19, August 2010. ISSN 0163-5980. doi: 10.1145/1842733.1842736. URL <https://doi.org/10.1145/1842733.1842736>.
- Georgios Paschos, George Iosifidis, and Giuseppe Caire. Cache optimization models and algorithms. *Foundations and Trends® in Communications and Information Theory*, 16(3–4):156–345, 2020. ISSN 1567-2190. doi: 10.1561/0100000104. URL <http://dx.doi.org/10.1561/0100000104>.
- Dan Garber. Efficient online linear optimization with approximation algorithms. *Mathematics of Operations Research*, 46(1):204–220, 2021.
- Sham M Kakade, Adam Tauman Kalai, and Katrina Ligett. Playing games with approximation algorithms. *SIAM Journal on Computing*, 39(3):1088–1106, 2009.
- Takahiro Fujita, Kohei Hatano, and Eiji Takimoto. Combinatorial online prediction via metarounding. In *International Conference on Algorithmic Learning Theory*, pages 68–82. Springer, 2013.
- Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- James Hannan. Approximation to bayes risk in repeated play. *Contributions to the Theory of Games*, 3:97–139, 1957.
- Jacob Abernethy, Chansoo Lee, and Ambuj Tewari. Perturbation techniques in online learning and optimization. *Perturbations, Optimization, and Statistics*, page 233, 2016.
- Stefano Traverso, Mohamed Ahmed, Michele Garetto, Paolo Giaccone, Emilio Leonardi, and Saverio Niccolini. Temporal locality in today’s content caching: why it matters and how to model it. *ACM SIGCOMM Computer Communication Review*, 43(5):5–12, 2013.
- Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. cambridge university press, 2005.
- Benjamin Van Roy. A short proof of optimality for the min cache replacement algorithm. *Information processing letters*, 102(2-3):72–73, 2007.
- Donghee Lee, Jongmoo Choi, Jong-Hun Kim, Sam H Noh, Sang Lyul Min, Yookun Cho, and Chong-Sang Kim. On the existence of a spectrum of policies that subsumes the least recently used (lru) and least frequently used (lfu) policies. In *SIGMETRICS*, volume 99, pages 1–4. Citeseer, 1999.
- Asit Dan and Don Towsley. *An approximate analysis of the LRU and FIFO buffer replacement schemes*, volume 18. ACM, 1990.
- Jeffrey Scott Vitter and Padmanabhan Krishnan. Optimal prefetching via data compression. *Journal of the ACM (JACM)*, 43(5):771–793, 1996.
- P Krishnan and Jeffrey Scott Vitter. Optimal prediction for prefetching in the worst case. *SIAM Journal on Computing*, 27(6):1617–1636, 1998.
- Amit Daniely and Yishay Mansour. Competitive ratio vs regret minimization: achieving the best of both worlds. In *Algorithmic Learning Theory*, pages 333–368, 2019.
- Nick Littlestone and Manfred K Warmuth. The weighted majority algorithm. *Information and computation*, 108(2):212–261, 1994.

- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- Karthikeyan Shanmugam, Negin Golrezaei, Alexandros G Dimakis, Andreas F Molisch, and Giuseppe Caire. Femtocaching: Wireless content delivery through distributed caching helpers. *IEEE Transactions on Information Theory*, 59(12):8402–8413, 2013.
- Georgios S Paschos, Apostolos Destounis, Luigi Vigneri, and George Iosifidis. Learning to cache with no regrets. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*, pages 235–243. IEEE, 2019.
- Rajarshi Bhattacharjee, Subhankar Banerjee, and Abhishek Sinha. Fundamental limits on the regret of online network-caching. *Proc. ACM Meas. Anal. Comput. Syst.*, 4(2), June 2020. doi: 10.1145/3392143. URL <https://doi.org/10.1145/3392143>.
- Elad Hazan. Introduction to online convex optimization. *arXiv preprint arXiv:1909.05207*, 2019.
- Elad Hazan and Edgar Minasyan. Faster projection-free online learning. In *Conference on Learning Theory*, pages 1877–1893. PMLR, 2020.
- Alexander A Ageev and Maxim I Sviridenko. Pipe rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8(3): 307–328, 2004.
- Luc Devroye, Gábor Lugosi, and Gergely Neu. Random-walk perturbations for online combinatorial optimization. *IEEE Transactions on Information Theory*, 61(7):4099–4106, 2015.
- Arun Sai Suggala and Praneeth Netrapalli. Online non-convex learning: Following the perturbed leader is optimal. In *Algorithmic Learning Theory*, pages 845–861, 2020.
- Jacob Abernethy, Chansoo Lee, Abhinav Sinha, and Ambuj Tewari. Online linear optimization via smoothing. In *Conference on Learning Theory*, pages 807–823, 2014.
- Chansoo Lee. *Analysis of Perturbation Techniques in Online Learning*. PhD thesis, 2018.
- Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- Elad Hazan, Wei Hu, Yuanzhi Li, and Zhiyuan Li. Online improper learning with an approximation oracle. *arXiv preprint arXiv:1804.07837*, 2018.
- William G Madow et al. On the theory of systematic sampling, ii. *The Annals of Mathematical Statistics*, 20(3):333–354, 1949.
- Samrat Mukhopadhyay and Abhishek Sinha. Online caching with optimal switching regret. In *2021 IEEE International Symposium on Information Theory (ISIT)*, pages 1546–1551, 2021. doi: 10.1109/ISIT45174.2021.9517925.
- Sascha Geulen, Berthold Vöcking, and Melanie Winkler. Regret minimization for online buffering problems using the weighted majority algorithm. In *COLT*, pages 132–143. Citeseer, 2010.
- Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.
- Jacob Abernethy, Peter L Bartlett, Alexander Rakhlin, and Ambuj Tewari. Optimal strategies and minimax lower bounds for online convex games. 2008.
- Alfred V Aho, Peter J Denning, and Jeffrey D Ullman. Principles of optimal page replacement. *Journal of the ACM (JACM)*, 18(1):80–93, 1971.
- Daniel S Berger, Nathan Beckmann, and Mor Harchol-Balter. Practical bounds on optimal caching with variable object sizes. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 2(2):1–38, 2018.
- Daniel S Berger. CDN Traces. <https://github.com/dasebe/optimalwebcaching>, 2018.

- Debjit Paria and Abhishek Sinha. Source code for LeadCache: Regret-Optimal Caching in Networks. <https://github.com/AbhishekMITIITM/LeadCache-NeurIPS21>, 2021.
- Apostolos Destounis, Mari Kobayashi, Georgios Paschos, and Asma Ghorbel. Alpha fair coded caching. In *2017 15th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, pages 1–8. IEEE, 2017.
- Dimitri P Bertsekas. Stochastic optimization problems with nondifferentiable cost functionals. *Journal of Optimization Theory and Applications*, 12(2):218–231, 1973.
- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Pascal Massart. *Concentration inequalities and model selection*, volume 6. Springer, 2007.
- Sheldon M Ross. *Stochastic processes*, volume 2. Wiley New York, 1996.
- Reinhard Diestel. Graph theory 3rd ed. *Graduate texts in mathematics*, 173, 2005.
- Noga Alon and Joel H Spencer. *The probabilistic method*. John Wiley & Sons, 2004.
- Godfrey Harold Hardy, John Edensor Littlewood, George Pólya, György Pólya, DE Littlewood, et al. *Inequalities*. Cambridge university press, 1952.
- F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.

## 10 Supplementary Material for the paper *LeadCache: Regret-Optimal Caching in Networks* by Debjit Paria and Abhishek Sinha

### 10.1 Proof of Lemma 1

From equation (4), we have that for any file request vector  $\mathbf{x}$  and virtual caching configuration  $\mathbf{z}$  :

$$r(\mathbf{x}, \mathbf{z}) = \langle \mathbf{x}, \mathbf{z} \rangle \leq q(\mathbf{x}, \psi(\mathbf{z})).$$

Thus for any file request sequence  $\{\mathbf{x}_t\}_{t \geq 1}$ , we have:

$$\sum_{t=1}^T r(\mathbf{x}_t, \mathbf{z}_t) \leq \sum_{t=1}^T q(\mathbf{x}_t, \psi(\mathbf{z}_t)). \quad (20)$$

On the other hand, let  $\mathbf{y}_* \in \arg \max_{\mathbf{y} \in \mathcal{Y}} \sum_{t=1}^T q(\mathbf{x}_t, \mathbf{y})$  be an optimal static offline cache configuration vector corresponding to the file requests  $\{\mathbf{x}_t\}_{t=1}^T$ . Consider a candidate static virtual cache configuration vector  $\mathbf{z}_* \in \mathcal{Z}$  defined as:

$$\mathbf{z}_*^i \equiv \min \left\{ \mathbf{1}_{N \times 1}, \left( \sum_{j \in \partial^+(i)} \mathbf{y}_*^j \right) \right\}, \quad 1 \leq i \leq n.$$

We have

$$\max_{\mathbf{z} \in \mathcal{Z}} \sum_{t=1}^T r(\mathbf{x}_t, \mathbf{z}) \geq \left\langle \sum_{t=1}^T \mathbf{x}_t, \mathbf{z}_* \right\rangle = \max_{\mathbf{y} \in \mathcal{Y}} \sum_{t=1}^T q(\mathbf{x}_t, \mathbf{y}). \quad (21)$$

Combining Eqns. (20) and (21), we conclude that

$$\max_{\mathbf{y} \in \mathcal{Y}} \sum_{t=1}^T q(\mathbf{x}_t, \mathbf{y}) - \sum_{t=1}^T q(\mathbf{x}_t, \psi(\mathbf{z}_t)) \leq \max_{\mathbf{z} \in \mathcal{Z}} \sum_{t=1}^T r(\mathbf{x}_t, \mathbf{z}) - \sum_{t=1}^T r(\mathbf{x}_t, \mathbf{z}_t). \quad (22)$$

Taking supremum of both sides of inequality (22) over all possible file request sequences  $\{\mathbf{x}_t\}_{t \geq 1}$  yields the result.  $\blacksquare$

### 10.2 Proof of Theorem 1

Keeping Lemma 1 in view, to prove the desired regret upper bound for the LeadCache policy, it is enough to bound the regret for the virtual policy  $\pi^{\text{virtual}}$  only. Following Cohen and Hazan [2015] we derive a general expression for the regret upper bound applicable to any linear reward function under an anytime FTPL policy. This is accomplished in the following steps. First, we extend the argument of Cohen and Hazan [2015] to the anytime setting. Then, we specialize this bound to our problem setting.

Recall the notations used in the paper - the aggregate file-request sequence from all users is denoted by  $\{\mathbf{x}_t\}_{t \geq 1}$  and the virtual cache configuration sequence is denoted by  $\{\mathbf{z}_t\}_{t \geq 1}$ . Define the cumulative requests up to time  $t$  as:

$$\mathbf{X}_t = \sum_{\tau=1}^{t-1} \mathbf{x}_\tau.$$

Note that the LeadCache policy chooses the virtual cache configuration at time slot  $t$  by solving the following optimization problem at time slot  $t$ :

$$\mathbf{z}_t = \arg \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \mathbf{X}_t + \eta_t \boldsymbol{\gamma} \rangle, \quad (23)$$

where each of the  $Nn$  components of the random vector  $\boldsymbol{\gamma}$  is sampled independently from the standard Gaussian distribution, and  $\mathcal{Z}$  denotes the set of all feasible virtual cache configurations as defined earlier in the paper. Next, we define the following potential function:

$$\Phi_{\eta_t}(\mathbf{x}) \equiv \mathbb{E}_{\boldsymbol{\gamma}} \left[ \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \mathbf{x} + \eta_t \boldsymbol{\gamma} \rangle \right]. \quad (24)$$

Since the perturbation r.v.  $\boldsymbol{\gamma}$  is Gaussian, it follows that the potential function  $\Phi_{\eta_t}(\mathbf{x})$  is twice continuously differentiable [Abernethy et al., 2016, Lemma 1.5]. Furthermore, since the max function

is convex, we may interchange the expectation and gradient to obtain  $\nabla\Phi_{\eta_t}(\mathbf{X}_t) = \mathbb{E}(z_t)$  [Bertsekas, 1973, Proposition 2.2]. Thus we have:

$$\langle \nabla\Phi_{\eta_t}(\mathbf{X}_t), \mathbf{x}_t \rangle = \mathbb{E}\langle z_t, \mathbf{x}_t \rangle. \quad (25)$$

To upper bound the regret of the LeadCache policy, we expand  $\Phi_{\eta_t}(\mathbf{X}_{t+1})$  in second-order Taylor's series as follows:

$$\begin{aligned} & \Phi_{\eta_t}(\mathbf{X}_{t+1}) \\ &= \Phi_{\eta_t}(\mathbf{X}_t + \mathbf{x}_t) \\ &= \Phi_{\eta_t}(\mathbf{X}_t) + \langle \nabla\Phi_{\eta_t}(\mathbf{X}_t), \mathbf{x}_t \rangle + \frac{1}{2} \mathbf{x}_t^T \nabla^2\Phi_{\eta_t}(\tilde{\mathbf{X}}_t) \mathbf{x}_t, \end{aligned} \quad (26)$$

where  $\tilde{\mathbf{X}}_t$  lies on the line segment joining  $\mathbf{X}_t$  and  $\mathbf{X}_{t+1}$ . Plugging in the expression of the inner product from Eqn. (25) in expression (26), we obtain:

$$\mathbb{E}\langle z_t, \mathbf{x}_t \rangle = \Phi_{\eta_t}(\mathbf{X}_{t+1}) - \Phi_{\eta_t}(\mathbf{X}_t) - \frac{1}{2} \mathbf{x}_t^T \nabla^2\Phi_{\eta_t}(\tilde{\mathbf{X}}_t) \mathbf{x}_t. \quad (27)$$

Summing up Eqn. (27) from  $t = 1$  to  $T$ , the total expected reward accrued by the LeadCache policy may be computed to be:

$$\begin{aligned} & \mathbb{E}(Q^{\text{LeadCache}}(T)) \\ &= \sum_{t=1}^T \mathbb{E}\langle z_t, \mathbf{x}_t \rangle \\ &= \sum_{t=1}^T \left( \Phi_{\eta_t}(\mathbf{X}_{t+1}) - \Phi_{\eta_t}(\mathbf{X}_t) \right) - \frac{1}{2} \sum_{t=1}^T \mathbf{x}_t^T \nabla^2\Phi_{\eta_t}(\tilde{\mathbf{X}}_t) \mathbf{x}_t \\ &= \sum_{t=1}^T \left( \Phi_{\eta_t}(\mathbf{X}_{t+1}) - \Phi_{\eta_{t+1}}(\mathbf{X}_{t+1}) + \Phi_{\eta_{t+1}}(\mathbf{X}_{t+1}) - \Phi_{\eta_t}(\mathbf{X}_t) \right) - \frac{1}{2} \sum_{t=1}^T \mathbf{x}_t^T \nabla^2\Phi_{\eta_t}(\tilde{\mathbf{X}}_t) \mathbf{x}_t \\ &= \sum_{t=1}^T \left( \Phi_{\eta_t}(\mathbf{X}_{t+1}) - \Phi_{\eta_{t+1}}(\mathbf{X}_{t+1}) \right) + \Phi_{\eta_{T+1}}(\mathbf{X}_{T+1}) - \Phi_{\eta_1}(\mathbf{X}_1) - \frac{1}{2} \sum_{t=1}^T \mathbf{x}_t^T \nabla^2\Phi_{\eta_t}(\tilde{\mathbf{X}}_t) \mathbf{x}_t. \end{aligned}$$

Next, note that

$$\begin{aligned} \Phi_{\eta_{T+1}}(\mathbf{X}_{T+1}) &= \mathbb{E}_{\gamma} \left[ \max_{z \in \mathcal{Z}} \langle z, \mathbf{X}_{T+1} + \eta_{T+1} \gamma \rangle \right] \\ &\stackrel{\text{(Jensen's ineq.)}}{\geq} \max_{z \in \mathcal{Z}} \left[ \mathbb{E}_{\gamma} \langle z, \mathbf{X}_{T+1} + \eta_{T+1} \gamma \rangle \right] \\ &= \max_{z \in \mathcal{Z}} \langle z, \mathbf{X}_{T+1} \rangle \\ &= Q^*(T), \end{aligned}$$

where recall that  $Q^*(T)$  denotes the optimal cumulative reward up to time  $T$  obtained by the best static policy in hindsight. Hence, from the above, we can upper bound the expected regret (2) of the LeadCache policy as:

$$\begin{aligned} & \mathbb{E}(R_T^{\text{LeadCache}}) \\ &= Q^*(T) - \mathbb{E}(Q^{\text{LeadCache}}(T)) \\ &\leq \underbrace{\Phi_{\eta_1}(\mathbf{X}_1) + \sum_{t=1}^T \left( \Phi_{\eta_{t+1}}(\mathbf{X}_{t+1}) - \Phi_{\eta_t}(\mathbf{X}_{t+1}) \right)}_{(a)} + \frac{1}{2} \sum_{t=1}^T \mathbf{x}_t^T \nabla^2\Phi_{\eta_t}(\tilde{\mathbf{X}}_t) \mathbf{x}_t. \end{aligned} \quad (28)$$

**Bounding the term (a):** Next, to upper bound the expected regret, we control term (a) in inequality (28). From Eqns. (23) and (24), we can write:

$$\Phi_{\eta_{t+1}}(\mathbf{X}_{t+1}) = \mathbb{E}[\langle z_{t+1}, \mathbf{X}_{t+1} + \eta_{t+1} \gamma \rangle],$$

and

$$\Phi_{\eta_t}(\mathbf{X}_{t+1}) \geq \mathbb{E}[\langle z_{t+1}, \mathbf{X}_{t+1} + \eta_t \gamma \rangle].$$

Hence, each term in the summation (a) may be upper bounded as follows:

$$\begin{aligned}
\Phi_{\eta_{t+1}}(\mathbf{X}_{t+1}) - \Phi_{\eta_t}(\mathbf{X}_{t+1}) &\leq \mathbb{E}[\langle \mathbf{z}_{t+1}, \mathbf{X}_{t+1} + \eta_{t+1}\boldsymbol{\gamma} \rangle] - \mathbb{E}[\langle \mathbf{z}_{t+1}, \mathbf{X}_{t+1} + \eta_t\boldsymbol{\gamma} \rangle] \\
&= \mathbb{E}[\langle \mathbf{z}_{t+1}, (\eta_{t+1} - \eta_t)\boldsymbol{\gamma} \rangle] \\
&= (\eta_{t+1} - \eta_t)\mathbb{E}[\langle \mathbf{z}_{t+1}, \boldsymbol{\gamma} \rangle] \\
&\leq (\eta_{t+1} - \eta_t)\mathbb{E}[\max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \boldsymbol{\gamma} \rangle] \\
&= (\eta_{t+1} - \eta_t)\mathcal{G}(\mathcal{Z}),
\end{aligned}$$

where the quantity  $\mathcal{G}(\mathcal{Z})$  is known as the *Gaussian complexity* of the set  $\mathcal{Z}$  of virtual configurations [Wainwright \[2019\]](#). Since the Gaussian perturbation  $\boldsymbol{\gamma}$  has zero mean,  $\mathcal{G}(\mathcal{Z})$  is non-negative due to Jensen's inequality. Substituting the above upper bound back in Eqn. (28), we notice that the summation in part (a) telescopes, yielding the following bound for the expected regret:

$$\mathbb{E}(R_T^{\text{LeadCache}}) \leq \eta_{T+1} \underbrace{\mathcal{G}(\mathcal{Z})}_{(b)} + \frac{1}{2} \sum_{t=1}^T \underbrace{\mathbf{x}_t^T \nabla^2 \Phi_{\eta_t}(\tilde{\mathbf{X}}_t) \mathbf{x}_t}_{(c)}. \quad (29)$$

We now upper bound each of the terms (b) and (c) as defined in the above regret bound.

**Bounding term (b) in Eqn. (29):** In the following, we upper bound the Gaussian complexity of the set  $\mathcal{Z}$ :

$$\mathcal{G}(\mathcal{Z}) \equiv \mathbb{E}_{\boldsymbol{\gamma}}[\max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \boldsymbol{\gamma} \rangle].$$

From equation (4), we have for any feasible  $\mathbf{z} \in \mathcal{Z}$ :

$$\begin{aligned}
\sum_{i \in \mathcal{I}, f \in [N]} z_f^i &\leq \sum_{i \in \mathcal{I}, f \in [N]} \sum_{j \in \partial^+(i)} y_f^j \\
&= \sum_{j \in \mathcal{J}} \sum_{f \in [N]} \sum_{i \in \partial^-(j)} y_f^j \\
&\stackrel{(a)}{\leq} d \sum_{j \in \mathcal{J}} \sum_{f \in [N]} y_f^j \\
&\stackrel{(b)}{\leq} dmC.
\end{aligned} \quad (30)$$

where, in step (a), we have used our assumption that the right-degree of the bipartite graph  $\mathcal{G}$  is upper bounded by  $d$ , and in (b), we have used the fact that the capacity of each cache is bounded by  $C$ .

For any fixed  $\mathbf{z} \in \mathcal{Z}$ , the random inner-product  $\langle \mathbf{z}, \boldsymbol{\gamma} \rangle$  follows a normal distribution with mean zero and variance  $\sigma^2$  where

$$\sigma^2 \equiv \mathbb{E} \langle \mathbf{z}, \boldsymbol{\gamma} \rangle^2 \stackrel{(a)}{=} \sum_{i \in \mathcal{I}, f \in [N]} (z_f^i)^2 \stackrel{(b)}{=} \sum_{i \in \mathcal{I}, f \in [N]} z_f^i \stackrel{(c)}{\leq} dmC.$$

In the above, equality (a) follows from the fact that  $\boldsymbol{\gamma}$  is a standard normal r.v., equality (b) follows from the fact that the components  $z_f^i$ 's are binary-valued (hence,  $(z_f^i)^2 = z_f^i$ ), and equality (c) follows from the upper bound given in Eqn. (30).

Next, observe that since the feasible set  $\mathcal{Z}$  is downward closed, if  $\mathbf{z}_* \in \arg \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \boldsymbol{\gamma} \rangle$ , then  $\gamma_f^i < 0$  implies  $z_{*f}^i = 0$ . Hence, we can simplify the expression for the Gaussian complexity of the set  $\mathcal{Z}$  as

$$\mathcal{G}(\mathcal{Z}) \equiv \mathbb{E}_{\boldsymbol{\gamma}}[\max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \boldsymbol{\gamma} \rangle] = \mathbb{E}_{\boldsymbol{\gamma}}[\max_{\mathbf{z} \in \mathcal{Z}} \sum_{(i,f): \gamma_f^i > 0} z_f^i \gamma_f^i].$$

Since all coefficients  $\gamma_f^i$  in the above summation are positive, we conclude that there exists an optimal vector  $\mathbf{z}_* \in \mathcal{Z}$  such that the inequality in Eqn. (4) is met with an equality for other components of  $\mathbf{z}$ , i.e.,  $\forall (i, f) : \gamma_f^i > 0$ , we have

$$z_{*f}^i = \min(1, \sum_{j \in \partial^+(i)} y_{*f}^j), \quad (31)$$



for some  $\mathbf{y}_* \in \mathcal{Y}$ . Let  $\mathcal{Z}_*$  be the set of all feasible virtual caching vectors satisfying (31) for some feasible  $\mathbf{y}_* \in \mathcal{Y}$ . Since the optimal virtual caching vector  $\mathbf{z}_* \in \mathcal{Z}_*$  is completely determined by the corresponding physical caching vector  $\mathbf{y}_* \in \mathcal{Y}$ , we have that  $|\mathcal{Z}_*| \leq |\mathcal{Y}|$ . Furthermore, since any of the  $m$  caches can be loaded with any  $C$  files, we have the bound:

$$|\mathcal{Y}| \leq \binom{N}{C}^m \leq \left(\frac{Ne}{C}\right)^{mC}, \quad (32)$$

where the last inequality is a standard upper bound for binomial coefficients. Finally, using Massart [2007]'s lemma for Gaussian variables, we have

$$\begin{aligned} \mathcal{G}(\mathcal{Z}) &\equiv \mathbb{E}_\gamma \left[ \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \gamma \rangle \right] = \mathbb{E}_\gamma \left[ \max_{\mathbf{z} \in \mathcal{Z}_*} \sum_{(i,f): \gamma_f^i > 0} z_f^i \gamma_f^i \right] \leq \sqrt{dmC} \sqrt{2 \log |\mathcal{Z}_*|} \\ &\leq mC \sqrt{2d \left( \log \frac{N}{C} + 1 \right)}. \end{aligned} \quad (33)$$

**Bounding term (c) in Eqn. (29):** Let us denote the file requested by the  $i^{\text{th}}$  user at time  $t$  by  $f_i$ . Using Abernethy et al. [2016, Lemma 1.5], we have

$$\left( \nabla^2 \Phi_{\eta_t}(\tilde{\mathbf{X}}_t) \right)_{\mathbf{p}\mathbf{q}} = \frac{1}{\eta_t} \mathbb{E}[\hat{z}_{\mathbf{p}}^i \gamma_{\mathbf{q}}^j], \quad (34)$$

where  $\hat{z} \in \arg \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \tilde{\mathbf{X}}_t + \eta_t \gamma \rangle$ , and each of the indices  $\mathbf{p}, \mathbf{q}$  is a (user, file) tuple. Hence, using Eqn. (34), and noting that each user requests only one file at a time, we have:

$$\begin{aligned} \mathbf{x}_t^T \nabla^2 \Phi_{\eta_t}(\tilde{\mathbf{X}}_t) \mathbf{x}_t &= \frac{1}{\eta_t} \sum_{i,j \in \mathcal{I}} \mathbb{E}[\hat{z}_{f_i}^i \gamma_{f_j}^j] \\ &= \frac{1}{\eta_t} \mathbb{E} \left( \sum_{i \in \mathcal{I}} \hat{z}_{f_i}^i \right) \left( \sum_{j \in \mathcal{I}} \gamma_{f_j}^j \right) \\ &\stackrel{(a)}{\leq} \frac{1}{\eta_t} \sqrt{\mathbb{E} \left( \sum_{i \in \mathcal{I}} \hat{z}_{f_i}^i \right)^2 \mathbb{E} \left( \sum_{j \in \mathcal{I}} \gamma_{f_j}^j \right)^2} \\ &\stackrel{(b)}{\leq} \frac{1}{\eta_t} \sqrt{n^2 \times n} \\ &= \frac{1}{\eta_t} n^{3/2}, \end{aligned} \quad (35)$$

where the inequality (a) follows from the Cauchy-Schwartz inequality and the inequality (b) follows from the facts that  $\mathbf{z}$  are binary variables and that the components of the random vector  $\gamma$  are i.i.d. Finally, substituting the upper bounds from Eqns. (33) and (35) in the regret upper bound in Eqn. (29), we may upper bound the expected regret of the LeadCache policy as:

$$\begin{aligned} \mathbb{E}(R_T^{\text{LeadCache}}) &\leq \eta_{T+1} \mathcal{G}(\mathcal{Z}) + \frac{n^{3/2}}{2} \sum_{t=1}^T \frac{1}{\eta_t} \\ &\leq \eta_{T+1} mC \sqrt{2d \left( \log \frac{N}{C} + 1 \right)} + \frac{n^{3/2}}{2} \sum_{t=1}^T \frac{1}{\eta_t}, \end{aligned}$$

where the bound in the last inequality follows from Eqn. (33). Choosing the learning rates  $\eta_t = \beta \sqrt{t}$  with an appropriate constant  $\beta > 0$  yields the following regret upper bound:

$$\mathbb{E}(R_T^{\text{LeadCache}}) \leq kn^{3/4} d^{1/4} \sqrt{mCT},$$

for some  $k = O(\text{poly-log}(N/C))$ . ■

## 11 Proof of Theorem 2

Denote the objective function of Problem (8) by  $L(\mathbf{y}) \equiv \sum_{i,f} \theta_f^i \min(1, \sum_{j \in \partial^+(i)} y_f^j)$ , where, to simplify the notations, we have not explicitly shown the dependence of the  $\theta$  coefficients on the time

index  $t$ . Recall the definition of surrogate objective function  $\phi(\mathbf{y})$  given in Eqn. (15):

$$\phi(\mathbf{y}) = \sum_{i,f} (\theta_f^i)^+ \left(1 - \prod_{j \in \partial^+(i)} (1 - y_f^j)\right), \quad (36)$$

From Ageev and Sviridenko [2004, Eqn. (16)], we have the following algebraic inequality:

$$L(\mathbf{y}) \stackrel{(a)}{\geq} \phi(\mathbf{y}) \geq \left(1 - \left(1 - \frac{1}{\Delta}\right)^\Delta\right) L(\mathbf{y}), \quad (37)$$

where  $\Delta \equiv \max_{i \in \mathcal{I}} |\partial^+(i)|$ . Note that inequality (a) holds with equality for binary vectors  $\mathbf{y} \in \{0, 1\}^{mN}$ .

Let  $\mathbf{y}^*$  be a solution of the relaxed LP (9), and OPT be the optimal value of the problem (8). Obviously,  $L(\mathbf{y}^*) \geq \text{OPT}$ , which, combined with the estimate in Eqn. (37), yields:

$$\phi(\mathbf{y}^*) \geq \left(1 - \left(1 - \frac{1}{\Delta}\right)^\Delta\right) \text{OPT}. \quad (38)$$

Since  $\mathbf{y}^*$  is a solution to the relaxed LP, it may possibly contain fractional coordinates. In the following, we show that the Pipage rounding procedure, described in Algorithm 2, rounds at least one fractional variable of a cache at a round *without* decreasing the value of the surrogate objective function  $\phi(\cdot)$  (Steps 3-6).

For a given fractional allocation vector  $\mathbf{y}$ , and another vector  $\mathbf{v}_y$  of our choice depending on  $\mathbf{y}$ , define a one-dimensional function  $g_y(\cdot)$  as:

$$g_y(s) = \phi(\mathbf{y} + s\mathbf{v}_y). \quad (39)$$

The vector  $\mathbf{v}_y$  denotes the direction along which the fractional allocation vector  $\mathbf{y}$  is rounded in the current step. The Pipage rounding procedure, Algorithm 2, chooses the vector  $\mathbf{v}_y$  as follows: consider any cache  $j$  that has at least two fractional coordinates  $y_{f_1}^j$  and  $y_{f_2}^j$  in the current allocation  $\mathbf{y}$  (Step 3 of Algorithm 2)<sup>2</sup>. Take  $\mathbf{v}_y = e_{j,f_1} - e_{j,f_2}$ , where  $e_{j,l}$  denotes the standard unit vector with 1 in the coordinate corresponding to the  $l^{\text{th}}$  file of the  $j^{\text{th}}$  cache,  $l = f_1, f_2$ . We now claim that the function  $g_y(s) = \phi(\mathbf{y} + s\mathbf{v}_y)$  is linear in  $s$ . To see this, consider any one of the constituent terms of  $g_y(s)$  as given in Eqn. (36). Examining each term, we arrive at the following two cases:

1. If both  $f \neq f_1$  and  $f \neq f_2$  then that term is independent of  $s$ .
2. If either  $f = f_1$ , or  $f = f_2$ , the variables  $y_{f_1}^j$  or  $y_{f_2}^j$  may appear in each product term in (36) at most once. Since the product terms contain exactly one variable corresponding to each file, the variables  $y_{f_1}^j$  and  $y_{f_2}^j$  can not appear in the same product term together.

The above two cases imply that the function  $g_y(s)$  is linear in  $s$ . By increasing and decreasing the variable  $s$  to the maximum extent possible, so that the candidate allocation  $\mathbf{y} + s\mathbf{v}_y$  does not violate the constraint (12), we construct two new candidate allocation vectors  $\alpha = \mathbf{y} - \epsilon_1 \mathbf{v}_y$  and  $\beta = \mathbf{y} + \epsilon_2 \mathbf{v}_y$ , where the constants  $\epsilon_1$  and  $\epsilon_2$  are chosen in such a way that at least one of the fractional variables of  $\mathbf{y}$  becomes integral (Steps 4-5). It is easy to see that, by design, all cache capacity constraints in Eqn. (11) continue to hold in both of these two candidate allocations. In step 6, we choose the best of the candidate allocations  $\alpha$  and  $\beta$ , corresponding to the surrogate function  $\phi(\cdot)$ . Let  $\mathbf{y}^{\text{new}}$  denote the new candidate allocation vector. Since the maximum of a linear function over an interval is achieved on one of its two boundaries, we conclude that  $\phi(\mathbf{y}^{\text{new}}) \geq \phi(\mathbf{y})$ . As argued above, the rounded solution is feasible and has at least one less fractional coordinate. Hence, by repeated application of the above procedure, we finally arrive at a feasible integral allocation  $\hat{\mathbf{y}}$  such that:

$$L(\hat{\mathbf{y}}) = \phi(\hat{\mathbf{y}}) \geq \phi(\mathbf{y}^*) \geq \left(1 - \left(1 - \frac{1}{\Delta}\right)^\Delta\right) \text{OPT},$$

where the first equality follows from that fact that the functions  $\phi(\mathbf{y}) = L(\mathbf{y})$  on integral points. ■

<sup>2</sup>Since the cache capacities are integers, there cannot be a cache with only one fractional allocation variable.

## 12 Madow's Sampling Scheme

Madow's sampling scheme is a simple statistical procedure for randomly sampling a subset of items without replacement from a larger universe with a specified set of inclusion probabilities [Madow et al., 1949]. The pseudocode for Madow's sampling scheme is given in Algorithm 4. It samples  $C$  items without replacement from a universe with  $N$  items such that the item  $i$  is included in the sampled set with probability  $p_i$ ,  $1 \leq i \leq N$ . The inclusion probabilities satisfy the feasibility constraint given by Eqn. (18).

---

### Algorithm 4 Madow's Systematic Sampling Scheme without Replacement

---

**Input:** A universe  $[N]$  of size  $N$ , cardinality of the sampled set  $C$ , marginal inclusion probability vector  $\mathbf{p} = (p_1, p_2, \dots, p_N)$  satisfying the feasibility condition (18),

**Output:** A random set  $S$  with  $|S| = C$  such that,  $\mathbb{P}(i \in S) = p_i, \forall i \in [N]$

- 1: Define  $\Pi_0 = 0$ , and  $\Pi_i = \Pi_{i-1} + p_i, \forall 1 \leq i \leq N$ .
  - 2: Sample a uniformly distributed random variable  $U$  from the interval  $[0, 1]$ .
  - 3:  $S \leftarrow \emptyset$
  - 4: **for**  $i \leftarrow 0$  to  $C - 1$  **do**
  - 5:     Select element  $j$  if  $\Pi_{j-1} \leq U + i < \Pi_j$ .
  - 6:      $S \leftarrow S \cup \{j\}$ .
  - 7: **end for**
  - 8: **Return**  $S$
- 

**Correctness:** The correctness of Madow's sampling scheme is easy to verify. Due to the feasibility condition (18), Algorithm 4 selects exactly  $C$  elements. Furthermore, the element  $j$  is selected if the random variable  $U$  falls in the interval  $\cup_{i=1}^N [\Pi_{j-1} - i, \Pi_j - i]$ . Since  $U$  is uniformly distributed in  $[0, 1]$ , the probability that the element  $j$  is selected is equal to  $\Pi_j - \Pi_{j-1} = p_j, \forall j \in [N]$ .

## 13 Proof of Proposition 1

The proof of the regret bound with the relaxed action set  $\mathcal{Z}_{\text{rel}}$  follows the same line of arguments as the proof of Theorem 1 with integral cache allocations. In particular, we decompose the regret bound as in Eqn. (29), with the difference that we now replace the feasible set  $\mathcal{Z}$  in term (b) with the relaxed feasible set  $\mathcal{Z}_{\text{rel}}$ . Observe that, for bounding the term (c), we did not exploit the fact that the variables are integral. Hence, the bound (35) holds in the case of the relaxed feasible set as well. However, for bounding the Gaussian complexity in term (b) in the proof of Theorem 1, we explicitly used the fact that the cache allocations (and hence, the virtual actions) are integral (*viz.* the counting argument in Eqn. (32)). To get around this issue, we now give a different argument for bounding the Gaussian complexity in term (b) for the relaxed action set  $\mathcal{Z}_{\text{rel}}$ . Note that for any feasible  $(\mathbf{z}, \mathbf{y}) \in \mathcal{Z}_{\text{rel}}$ , we have

$$\|\mathbf{z}\|_1 = \sum_{i,f} z_f^i \leq \sum_i \sum_{j \in \partial^+(i),f} y_f^j \leq d \sum_j \sum_f y_f^j \leq mCd, \quad (40)$$

where we have used the fact that each cache is connected to at most  $d$  users and that each cache can hold  $C$  files at a time. Hence, we have

$$\mathcal{G}(\mathcal{Z}_{\text{rel}}) = \mathbb{E}_{\gamma} \left[ \max_{\mathbf{z} \in \mathcal{Z}_{\text{rel}}} \langle \mathbf{z}, \gamma \rangle \right] \stackrel{\text{(H\"older's ineq.)}}{\leq} \mathbb{E}_{\gamma} \left[ \max_{\mathbf{z} \in \mathcal{Z}_{\text{rel}}} \|\mathbf{z}\|_1 \|\gamma\|_{\infty} \right] \leq mCd \sqrt{4 \ln(Nn)}, \quad (41)$$

where, in the last inequality, we have used the  $\ell_1$ -norm bound (40) along with a standard upper bound on the expectation of the maximum of the absolute value of a set of i.i.d. standard Gaussian random variables [Wainwright, 2019]. Now proceeding similarly as in the proof of the regret bound for the action set  $\mathcal{Z}$ , we conclude that with an appropriate learning rate sequence, we have the following regret upper bound for the relaxed action set  $\mathcal{Z}_{\text{rel}}$ :

$$\mathbb{E}(\tilde{R}_T^{\text{LeadCache}}) \leq \kappa_1 n^{3/4} \sqrt{dmCT},$$

for some polylogarithmic factor  $\kappa_1$ . ■

## 14 Proof of Theorem 4

**Discussion:** To intuitively understand why the total number of fetches is expected to be small under the LeadCache policy, consider the simplest case of a single cache with a single user [Bhattacharjee et al., 2020]. At every slot, the LeadCache policy populates the cache with a set of  $C$  files having the highest perturbed cumulative count of requests  $\Theta(t)$ . For the sake of argument, assume that the learning rate  $\eta_t$  is time-invariant. Since at most one file is requested by the user per slot, only one component of  $\Theta(t)$  changes at a slot, and hence, the LeadCache policy fetches at most *one* new file at any time slot. Surprisingly, the following rigorous argument proves a far stronger result: the total number of fetches over an infinite time interval remains almost surely finite, even with a time-varying learning rate in the network caching setting.

**Proof:** Recall that, under the LeadCache policy, the optimal virtual caching configuration  $z_t$  for the  $t^{\text{th}}$  slot is obtained by solving the optimization problem  $\mathbf{P}$ :

$$\max_{z \in \mathcal{Z}} \sum_{i \in \mathcal{I}} \langle \theta^i(t), z^i \rangle, \quad (42)$$

where we assume that the ties (if any) are broken according to some *fixed* tie-breaking rule. As discussed before, the corresponding physical cache configuration  $\mathbf{y}_t$  may be obtained using the mapping  $\psi(\cdot)$ . Now consider a static virtual cache configuration  $\tilde{z}$  obtained by replacing the perturbed-count vectors  $\theta^i(t)$  in the objective function (42) with the vectors  $\mathbf{p}^i, \forall i \in \mathcal{I}$ , where  $\mathbf{p} = (\mathbf{p}^i, i \in \mathcal{I})$  is defined to be the vector of long-term file-request probabilities, given by Eqn. (19). In other words,

$$\tilde{z} \in \arg \max_{z \in \mathcal{Z}} \sum_{i \in \mathcal{I}} \langle \mathbf{p}^i, z^i \rangle. \quad (43)$$

Since the set of all possible virtual caching configurations  $\mathcal{Z}$  is finite, the objective value corresponding to any other non-optimal caching configuration must be some non-zero gap  $\delta > 0$  away from that of an optimal configuration. Let us denote the set of all *sub-optimal* virtual cache configuration vectors by  $\mathcal{B}$ . Hence, for any  $z \in \mathcal{B}$ , we must have:

$$\sum_{i \in \mathcal{I}} \langle \mathbf{p}^i, \tilde{z}^i \rangle \geq \sum_{i \in \mathcal{I}} \langle \mathbf{p}^i, z^i \rangle + \delta. \quad (44)$$

Let us define an “error” event  $\mathcal{E}(t)$  to be event such that the LeadCache policy yields a sub-optimal virtual cache configuration (and possibly, a sub-optimal physical cache configuration (4)) at time  $t$ . Let  $G$  be a zero-mean Gaussian random variable with variance  $2Nn$ . We now upper bound the probability of the error event  $\mathcal{E}(t)$  as below:

$$\begin{aligned} & \mathbb{P}(\mathcal{E}(t)) \\ & \stackrel{(a)}{\leq} \mathbb{P}\left(\sum_{i \in \mathcal{I}} \langle \theta^i(t), z^i(t) \rangle > \sum_{i \in \mathcal{I}} \langle \theta^i(t), \tilde{z}^i \rangle, z(t) \in \mathcal{B}\right) \\ & \stackrel{(b)}{\leq} \mathbb{P}\left(\eta_t G \geq \sum_{i \in \mathcal{I}} \langle \mathbf{X}^i(t), \tilde{z}^i \rangle - \sum_{i \in \mathcal{I}} \langle \mathbf{X}^i(t), z^i(t) \rangle, z(t) \in \mathcal{B}\right) \\ & \stackrel{(c)}{=} \mathbb{P}\left(\eta_t G \geq \sum_{i \in \mathcal{I}} \langle \mathbf{X}^i(t), \tilde{z}^i \rangle - \sum_{i \in \mathcal{I}} \langle \mathbf{X}^i(t), z^i(t) \rangle, \sum_{i \in \mathcal{I}} \langle \mathbf{X}^i(t), \tilde{z}^i \rangle - \sum_{i \in \mathcal{I}} \langle \mathbf{X}^i(t), z^i(t) \rangle > \frac{\delta t}{2}, z(t) \in \mathcal{B}\right) \\ & \quad + \mathbb{P}\left(\eta_t G \geq \sum_{i \in \mathcal{I}} \langle \mathbf{X}^i(t), \tilde{z}^i \rangle - \sum_{i \in \mathcal{I}} \langle \mathbf{X}^i(t), z^i(t) \rangle, \sum_{i \in \mathcal{I}} \langle \mathbf{X}^i(t), \tilde{z}^i \rangle - \sum_{i \in \mathcal{I}} \langle \mathbf{X}^i(t), z^i(t) \rangle \leq \frac{\delta t}{2}, z(t) \in \mathcal{B}\right) \\ & \stackrel{(d)}{\leq} \mathbb{P}\left(\eta_t G \geq \frac{\delta t}{2}\right) + \mathbb{P}\left(\frac{1}{t} \sum_{i \in \mathcal{I}} \langle \mathbf{X}^i(t), \tilde{z}^i - z^i(t) \rangle \leq \frac{\delta}{2}, z(t) \in \mathcal{B}\right) \\ & = \mathbb{P}\left(\eta_t G \geq \frac{\delta t}{2}\right) + \mathbb{P}\left(\sum_{i \in \mathcal{I}} \left\langle \frac{\mathbf{X}^i(t)}{t} - \mathbf{p}^i, \tilde{z}^i - z^i(t) \right\rangle + \sum_{i \in \mathcal{I}} \langle \mathbf{p}^i, \tilde{z}^i - z^i(t) \rangle \leq \frac{\delta}{2}, z(t) \in \mathcal{B}\right) \\ & \stackrel{(e)}{\leq} \mathbb{P}\left(\eta_t G \geq \frac{\delta t}{2}\right) + \mathbb{P}\left(\sum_{i \in \mathcal{I}} \left\langle \frac{\mathbf{X}^i(t)}{t} - \mathbf{p}^i, \tilde{z}^i - z^i(t) \right\rangle \leq -\frac{\delta}{2}\right) \\ & \stackrel{(f)}{\leq} \mathbb{P}\left(\eta_t G \geq \frac{\delta t}{2}\right) + \mathbb{P}\left(\sum_{i \in \mathcal{I}} \sum_{f \in [N]} \left| \frac{\mathbf{X}_f^i(t)}{t} - p_f^i \right| \geq \frac{\delta}{2}\right) \end{aligned}$$

$$\begin{aligned}
&\stackrel{(g)}{\leq} \mathbb{P}\left(\eta_t G \geq \frac{\delta t}{2}\right) + \mathbb{P}\left(\bigcup_{i,f} \left| \frac{\mathbf{X}_f^i(t)}{t} - p_f^i \right| \geq \frac{\delta}{2Nn}\right) \\
&\stackrel{(h)}{\leq} \mathbb{P}\left(\eta_t G \geq \frac{\delta t}{2}\right) + \sum_{i,f} \mathbb{P}\left(\left| \frac{\mathbf{X}_f^i(t)}{t} - p_f^i \right| \geq \frac{\delta}{2Nn}\right) \\
&\stackrel{(i)}{\leq} \exp(-ct) + Nn\alpha_\epsilon(t).
\end{aligned}$$

for some positive constants  $c$  and  $\epsilon$ , which depend on the problem parameters.

In the above chain of inequalities:

- (a) follows from the fact that on the error event  $\mathcal{E}(t)$ , the virtual cache configuration vector  $\mathbf{z}(t)$  must be in the sub-optimal set  $\mathcal{B}$  and, by definition, it must yield more objective value in (42) than the optimal virtual cache configuration vector  $\tilde{\mathbf{z}}$ ,
- (b) follows by writing  $\Theta(t) = \mathbf{X}(t) + \eta_t \gamma$ , and observing that the virtual configurations  $\mathbf{z}(t) \in \mathcal{B}$  and  $\tilde{\mathbf{z}}(t)$  may differ in at most  $Nn$  coordinates, and that the normal random variables are increasing (in the convex ordering sense) with their variances,
- (c) follows from the law of total probability,
- (d) follows from the monotonicity of the probability measures,
- (e) follows from Eqn. (44),
- (f) follows from the fact that for any two equal-length vectors  $\mathbf{a}, \mathbf{b}$ , triangle inequality yields:

$$\langle \mathbf{a}, \mathbf{b} \rangle \geq - \sum_k |\mathbf{a}_k| |\mathbf{b}_k|,$$

and that  $|\tilde{\mathbf{z}}_f^i - \mathbf{z}_f^i(t)| \leq 1, \forall i, f$ ,

(g) follows from the simple observation that at least one number in a set of some numbers must be at least as large as the average,

(h) follows from the union bound,

and finally, the inequality (i) follows from the concentration inequality for Gaussian variables and the concentration inequality for the request process  $\{\mathbf{X}(t)\}_{t \geq 1}$ , as given by Eqn. (19). Using the above bound and the assumptions on the request sequence, we have

$$\sum_{t \geq 1} \mathbb{P}(\mathcal{E}(t)) \leq \sum_{t \geq 1} \exp(-ct) + Nn \sum_{t \geq 1} \alpha_\epsilon(t) < \infty.$$

Hence, the first Borel-Cantelli Lemma implies that

$$\mathbb{P}(\mathcal{E}(t) \text{ i.o.}) = 0.$$

Hence, almost surely, the error events stop after a finite time. Thus, with a fixed tie-breaking rule, the new file fetches stop after a finite time w.p. 1.  $\blacksquare$

## 15 Renewal Processes satisfies the Regularity Condition (A)

Suppose that, for any  $i \in \mathcal{I}, f \in [N]$ , the cumulative request process  $\{\mathbf{X}_f^i(t)\}_{t \geq 1}$  constitutes a renewal process such that the renewal intervals have a common expectation  $1/p_f^i$  and a finite fourth moment. Let  $S_k$  be the time of the  $k^{\text{th}}$  renewal,  $k \geq 1$  [Ross, 1996]. In other words, the  $i^{\text{th}}$  user requests file  $f$  for the  $k^{\text{th}}$  time at time  $S_k, k \geq 1$ . Then we have

$$\begin{aligned}
\mathbb{P}\left(\frac{\mathbf{X}_f^i(t)}{t} - p_f^i \leq -\epsilon\right) &= \mathbb{P}\left(\mathbf{X}_f^i(t) \leq t(p_f^i - \epsilon)\right) \\
&\leq \mathbb{P}\left(S_{\lfloor t(p_f^i - \epsilon) \rfloor} \geq t\right) \\
&\leq \mathbb{P}\left((S_{\lfloor t(p_f^i - \epsilon) \rfloor} - \lfloor t(p_f^i - \epsilon) \rfloor)^4 \geq (t(1 - p_f^i + \epsilon))^4\right) \\
&\stackrel{(a)}{\leq} O\left(\frac{1}{t^2}\right),
\end{aligned}$$

where, in (a), we have used the Markov Inequality along with a standard upper bound on the fourth moment of a centered random walk. Using a similar line of arguments, we can show that

$$\mathbb{P}\left(\frac{\mathbf{X}_f^i(t)}{t} - p_f^i \geq \epsilon\right) \leq O\left(\frac{1}{t^2}\right).$$

Combining the above two bounds, we conclude that

$$\sum_{t \geq 1} \mathbb{P}\left(\left|\frac{\mathbf{X}_f^i(t)}{t} - p_f^i\right| \geq \epsilon\right) < \infty.$$

The above derivation verifies the regularity condition A for renewal request processes.  $\blacksquare$

## 16 Proof of Theorem 5

We establish a slightly stronger result by proving the announced lower bound for a regular bipartite network with uniform left-degree  $d_L$  and uniform right-degree  $d$ . Counting the total number of edges in two different ways, we have  $nd_L = md$ . Hence,  $d_L = \frac{md}{n}$ . For pedagogical reasons, we divide the entire proof into several logically connected parts.

**(A) Some Observations and Preliminary Lemmas:** To facilitate the analysis, we introduce the following surrogate linear reward function:

$$q_{\text{linear}}(\mathbf{x}, \mathbf{y}) \equiv \sum_{i \in \mathcal{I}} \mathbf{x}_t^i \cdot \left( \sum_{j \in \partial^+(i)} \mathbf{y}_t^j \right). \quad (45)$$

We begin our analysis with the following two observations:

1. **Upper Bound:** From the definition (1) of the rewards, we clearly have:

$$q(\mathbf{x}, \mathbf{y}) \leq q_{\text{linear}}(\mathbf{x}, \mathbf{y}), \quad \forall \mathbf{x}, \mathbf{y}. \quad (46)$$

2. **Local Exclusivity implies Linearity:** In the case when all caches connected to each user host *different* files, *i.e.*, the cached files are *locally exclusive* in the sense that they are not duplicated from each user's local point-of-view, *i.e.*,

$$\mathbf{y}^{j_1} \cdot \mathbf{y}^{j_2} = 0, \quad \forall j_1 \neq j_2 : j_1, j_2 \in \partial^+(i), \forall i \in \mathcal{I}, \quad (47)$$

the reward function (1) reduces to a linear one:

$$q(\mathbf{x}, \mathbf{y}) = q_{\text{linear}}(\mathbf{x}, \mathbf{y}), \quad \forall \mathbf{x}, \mathbf{y}. \quad (48)$$

The equation (48) follows from the fact that with the local exclusivity constraint, we have

$$\sum_{j \in \partial^+(i)} \mathbf{y}_t^j \leq \mathbf{1}, \quad \forall i \in \mathcal{I},$$

where the inequality holds component wise. Hence, the 'min' operator in the definition of the reward function (Eqn. (1)) is vacuous in this case. To make use of the linearity of the rewards as in Eqn. (48), the caches need to store items in such a way that the local exclusivity condition (47) holds. Towards this, we now define a special coloring of the nodes in the set  $\mathcal{J}$  for a given bipartite graph  $\mathcal{G}(\mathcal{I} \cup \mathcal{J}, E)$ .

**Definition 2** (Valid  $\chi$ -coloring of the caches). *Let  $\chi$  be a positive integer. A valid  $\chi$ -coloring of the caches of a bipartite network  $\mathcal{G}(\mathcal{I} \cup \mathcal{J}, E)$  is an assignment of colors from the set  $\{1, 2, \dots, \chi\}$  to the vertices in  $\mathcal{J}$  (*i.e.*, the caches) in such a way that all neighboring caches  $\partial^+(i)$  to every node  $i \in \mathcal{I}$  (*i.e.*, the users) are assigned different colors.*

Obviously, for a given bipartite graph  $\mathcal{G}$ , a valid  $\chi$ -coloring of the caches exists only if the number of possible colors  $\chi$  is large enough. The following lemma gives an upper bound to the value of  $\chi$  so that a valid  $\chi$ -coloring of the caches exists.

**Lemma 2.** Consider a bipartite network  $\mathcal{G}(\mathcal{I} \cup \mathcal{J}, E)$ , where each user  $i \in \mathcal{I}$  is connected to at most  $d_L$  caches, and each cache  $j \in \mathcal{J}$  is connected to at most  $d$  users. Then there exists a valid  $\chi$ -coloring of the caches where  $\chi \leq d_L d$ .

*Proof.* From the given bipartite network  $\mathcal{G}(V, E)$ , construct another graph  $H(V', E')$ , where the caches form the vertices of  $H$ , i.e.,  $V' \equiv \mathcal{J}$ . For any two vertices in  $u, v \in V'$ , there is an edge  $(u, v) \in E'$  if and only if a user  $i \in \mathcal{I}$  is connected to both the caches  $u$  and  $v$  in the bipartite network  $\mathcal{G}$ . Next, consider any cache  $j \in \mathcal{J}$ . By our assumption, it is connected to at most  $d$  users. On the other hand, each of the users is connected to at most  $d_L - 1$  caches other than  $j$ . Hence, the degree of any node  $j$  in the graph  $H$  is upper bounded as:

$$\Delta' \leq d(d_L - 1) \leq d_L d - 1.$$

Finally, using Brook's theorem [Diestel \[2005\]](#), we conclude that the vertices of the graph  $H$  may be colored using at most  $1 + \Delta' \leq d_L d$  different colors.  $\square$

**(B) Probabilistic Method for Regret Lower Bounds:** With the above results at our disposal, we now employ the well-known probabilistic method for proving the regret lower bound [[Alon and Spencer, 2004](#)]. The basic principle of the probabilistic method is quite simple. We compute a lower bound to the *expected regret* for any online network caching policy  $\pi$  for a chosen joint probability distribution  $\mathbf{p}(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T)$  over an *ensemble* of incoming file request sequence. Since the maximum of a set of numbers is at least as large as the expectation, the above quantity also gives a lower bound to the regret for the worst-case file request sequence. Clearly, the tightness of the resulting bound largely depends on our ability to identify a suitable input distribution  $\mathbf{p}(\cdot)$  that is amenable to analysis and, at the same time, yields a good bound. In the following, we show how this program can be elegantly carried out for the network caching problem.

Fix a valid  $\chi$ -coloring of the caches, and let  $k = \chi C$ . Consider a library consisting of  $N = 2k$  different files. We now choose a randomized file request sequence  $\{\mathbf{x}_t^i \equiv \alpha_t\}_{t=1}^T$  where each user  $i$  requests the same (random) file  $\alpha_t$  at slot  $t$  such that the common file request vector  $\alpha_t$  is sampled uniformly at random from the set of the *first*  $2k$  unit vectors  $\{\mathbf{e}_i \in \mathbb{R}^{2k}, 1 \leq i \leq 2k\}$  independently at each slot<sup>3</sup>. Formally, we choose:

$$\mathbf{p}(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T) := \prod_{t=1}^T \left( \frac{1}{2k} \mathbb{1}(\mathbf{x}_t^{i_1} = \mathbf{x}_t^{i_2}, \forall i_1, i_2 \in \mathcal{I}) \right).$$

**(C) Upper-bounding the Total Reward accrued by any Online Policy:** Making use of observation (46), the expected total reward  $G_T^\pi$  accrued by any online network caching policy  $\pi$  may be upper bounded as follows:

$$\begin{aligned} G_T^\pi &\leq \mathbb{E} \left( \sum_{t=1}^T \sum_{i \in \mathcal{I}} \mathbf{x}_t^i \cdot \sum_{j \in \partial^+(i)} \mathbf{y}_t^j \right) \\ &\stackrel{(a)}{=} \sum_{t=1}^T \sum_{(i,j) \in E} \mathbb{E}(\mathbf{x}_t^i \cdot \mathbf{y}_t^j) \\ &\stackrel{(b)}{=} \sum_{t=1}^T \sum_{(i,j) \in E} \mathbb{E}(\mathbf{x}_t^i) \cdot \mathbb{E}(\mathbf{y}_t^j) \\ &\stackrel{(c)}{=} \frac{1}{2k} \sum_{t=1}^T \mathbb{E} \left( \sum_{(i,j) \in E} \sum_{f \in [N]} \mathbf{y}_{tf}^j \right) \\ &\stackrel{(d)}{\leq} \frac{d}{2k} \sum_{t=1}^T \mathbb{E} \left( \sum_{j \in \mathcal{J}} \sum_{f \in [N]} \mathbf{y}_{tf}^j \right) \\ &\stackrel{(e)}{\leq} \frac{dmCT}{2k}, \end{aligned} \tag{49}$$

<sup>3</sup>Recall that, according to our one-hot encoding convention, a request for a file  $f$  by any user corresponds to the unit vector  $\mathbf{e}_f$ .

where the eqn. (a) follows from the linearity of expectation, eqn. (b) follows from the fact that the cache configuration vector  $\mathbf{y}_t$  is independent of the file request vector  $\mathbf{x}_t$  at the same slot, as the policy is online, the eqn. (c) follows from the fact that each of the  $N = 2k$  components of the vector  $\mathbb{E}(\mathbf{x}_t^i)$  is equal to  $\frac{1}{2k}$ , the inequality (d) follows from the fact that each cache is connected to at most  $d$  users, and finally, the inequality (e) follows from the cache capacity constraints.

**(D) Lower-bounding the Total Reward Accrued by the Static Oracle:** We now lower bound the total expected reward accrued by the optimal static offline policy (*i.e.*, the first term in the regret definition (2)). Note that, due to the presence of the ‘min’ operator in (1), obtaining an exactly optimal static cache configuration vector  $\mathbf{y}^*$  is non-trivial, as it requires solving an **NP-hard** optimization problem (8) (with the vector  $\theta(t)$  in the objective replaced by the cumulative file request vector  $\mathbf{X}(T) \equiv \sum_{t=1}^T \mathbf{x}_t$ ). However, since we only require a good lower bound to the total expected reward, a suitably constructed sub-optimal caching configuration will serve our purpose, provided that we can compute a lower bound to its expected reward. Towards this end, in the following, we construct a joint cache configuration vector  $\mathbf{y}_\perp$  that satisfies the local exclusivity constraint (47).

**D.1 Construction of a “good” cache configuration vector  $\mathbf{y}_\perp$ :** Let  $\mathcal{X}$  be a valid  $\chi$ -coloring of the caches. Let the color  $c$  appear in  $f_c$  different caches in the coloring  $\mathcal{X}$ . To simplify the notations, we relabel the colors in non-increasing order of their frequency of appearance in the coloring  $\mathcal{X}$ , *i.e.*,

$$f_1 \geq f_2 \geq \dots \geq f_\chi. \quad (50)$$

Let the vector  $\mathbf{v}$  be obtained by sorting the components of the vector  $\sum_{t=1}^T \alpha_t$  in non-increasing order. Partition the vector  $\mathbf{v}$  into  $\frac{2k}{C} = 2\chi$  segments by sequentially merging  $C$  contiguous coordinates of  $\mathbf{v}$  at a time. Let  $c_j$  denote the color of the cache  $j \in \mathcal{J}$  in the coloring  $\mathcal{X}$ . The cache configuration vector  $\mathbf{y}_\perp$  is constructed by loading each cache  $j \in \mathcal{J}$  with the set of all  $C$  files in the  $c_j^{\text{th}}$  segment of the vector  $\mathbf{v}$ . See Figure 4 for an illustration.

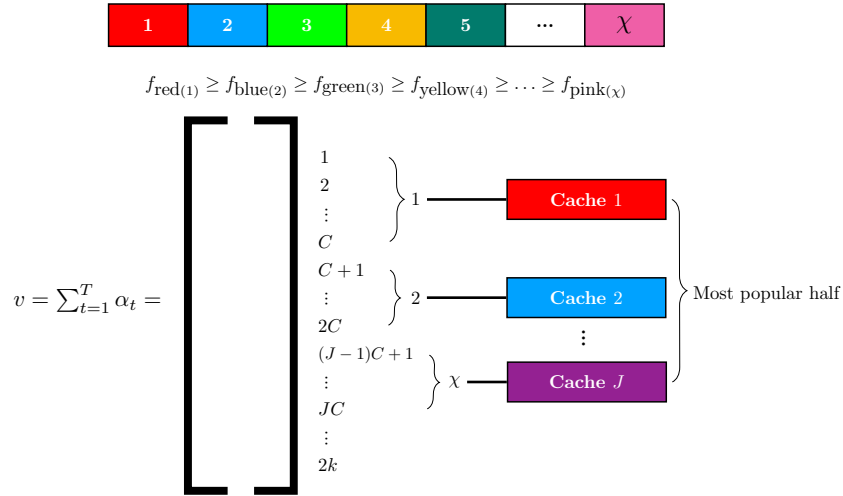


Figure 4: Construction of the caching configuration  $\mathbf{y}_\perp$ .

**D.2 Observation:** Since the vector  $\mathbf{v}$  has  $2\chi$  segments (each containing  $C$  different files) and the number of possible colors in the coloring  $\mathcal{X}$  is  $\chi$ , it follows that only the most popular half of the files get mapped to some caches under  $\mathbf{y}_\perp$ . Moreover, it can be easily verified that the cache configuration vector  $\mathbf{y}_\perp$  satisfies the local exclusivity property (47).

**D.3 Analysis:** Let  $S_v(m)$  denote the sum of the frequency counts in the  $m^{\text{th}}$  segment of the vector  $\mathbf{v}$ . In other words,  $S_v(m)$  gives the aggregate frequency of requests of the files in the  $m^{\text{th}}$  segment of the vector  $\mathbf{v}$ . By construction, we have

$$S_v(1) \geq S_v(2) \geq \dots \geq S_v(\chi). \quad (51)$$



Since under the distribution  $\mathbf{p}$ , all users request the same file at each time slot (*i.e.*,  $\mathbf{x}_t^i = \alpha_t, \forall i$ ), and since the linearity in rewards holds due to the local exclusivity property of the cache configuration  $\mathbf{y}_\perp$ , the reward obtained by the files in the  $j^{\text{th}}$  cache under the caching configuration  $\mathbf{y}_\perp$  is given by:

$$\begin{aligned} \mathbf{y}_\perp^j \cdot \left( \sum_{i \in \partial^-(j)} \sum_{t=1}^T \mathbf{x}_t^i \right) &= \mathbf{y}_\perp^j \cdot \left( \sum_{i \in \partial^-(j)} \sum_{t=1}^T \alpha_t \right) \\ &\stackrel{(a)}{=} d \mathbf{y}_\perp^j \cdot \left( \sum_{t=1}^T \alpha_t \right) \\ &\stackrel{(b)}{=} d S_{\mathbf{v}}(c_j), \end{aligned} \quad (52)$$

where the equation (a) follows from the fact that, in this converse proof, we are investigating a regular bipartite network where each cache is connected to exactly  $d$  users, and the equation (b) follows from the construction of the cache configuration vector  $\mathbf{y}_\perp$ . Hence, the expected aggregate reward accrued by the optimal stationary configuration  $\mathbf{y}^*$  may be lower-bounded by that of the configuration  $\mathbf{y}_\perp$  as follows:

$$\begin{aligned} G_T^{\pi^*} &\stackrel{(a)}{\geq} \mathbb{E} \left( \sum_{j \in \mathcal{J}} \mathbf{y}_\perp^j \cdot \left( \sum_{i \in \partial^-(j)} \sum_{t=1}^T \mathbf{x}_t^i \right) \right) \\ &\stackrel{(b)}{=} d \mathbb{E} \left( \sum_{j \in \mathcal{J}} S_{\mathbf{v}}(c_j) \right) \\ &\stackrel{(c)}{=} d \mathbb{E} \left( \sum_{c=1}^{\chi} f_c S_{\mathbf{v}}(c) \right) \\ &\stackrel{(d)}{\geq} \frac{d}{\chi} \left( \sum_{c=1}^{\chi} f_c \right) \mathbb{E} \left( \sum_{c=1}^{\chi} S_{\mathbf{v}}(c) \right), \end{aligned}$$

where

(a) follows from the local exclusivity property of the configuration  $\mathbf{y}_\perp$ ,

(b) follows from Eqn. (52),

(c) follows after noting that the color  $c$  appears on  $f_c$  different caches in the coloring  $\mathcal{X}$ ,

(d) follows from an algebraic inequality presented in Lemma 3 below, used in conjunction with the conditions (50) and (51).

Next, we lower bound the quantity  $\mathbb{E} \left( \sum_{c=1}^{\chi} S_{\mathbf{v}}(c) \right)$  appearing on the right hand side of the equation (53). Conceptually, identify the catalog with  $N = 2k$  "bins", and the random file requests  $\{\alpha_t\}_{t=1}^T$  as "balls" thrown uniformly into one of the "bins." With this correspondence in mind, a little thought reveals that the random variable  $\sum_{c=1}^{\chi} S_{\mathbf{v}}(c)$  is distributed identically as the total load in the most popular  $k = \chi C$  bins when  $T$  balls are thrown uniformly at random into  $2k$  bins. Continuing with the above chain of inequalities, we have

$$\begin{aligned} G_T^{\pi^*} &\stackrel{(e)}{\geq} \frac{dmC}{k} \mathbb{E}(\text{load in the most popular half of } 2k \text{ bins with } T \text{ balls thrown u.a.r.}) \\ &\stackrel{(f)}{\geq} \frac{dmC}{k} \left( \frac{T}{2} + \sqrt{\frac{kT}{2\pi}} \right) - \Theta\left(\frac{1}{\sqrt{T}}\right) \\ &= \frac{dmCT}{2k} + dmC \sqrt{\frac{T}{2\pi k}} - \Theta\left(\frac{1}{\sqrt{T}}\right), \end{aligned} \quad (53)$$

where, in the inequality (e), we have used the fact that  $\sum_{c=1}^{\chi} f_c = m$ , and the inequality (f) follows from Lemma 4 stated below. Hence, combining Eqns. (49) and (53), and noting that  $k = \chi C \leq d_L d C = \frac{mC d^2}{n}$  from Lemma 2, we conclude that for any caching policy  $\pi$ :

$$R_T^\pi \geq G_T^{\pi^*} - G_T^\pi \geq \sqrt{\frac{mnCT}{2\pi}} - \Theta\left(\frac{1}{\sqrt{T}}\right). \quad (54)$$

Moreover, making use of a *Globally* exclusive configuration (where all caches store different files), in Theorem 7 of their paper, Bhattacharjee et al. [2020] proved the following regret lower bound for any

online caching policy  $\pi$ :

$$R_T^\pi \geq d\sqrt{\frac{mCT}{2\pi}} - \Theta\left(\frac{1}{\sqrt{T}}\right). \quad (55)$$

Hence, combining the bounds from Eqns. (54) and (55), we conclude that

$$R_T^\pi \geq \max\left(\sqrt{\frac{mnCT}{2\pi}}, d\sqrt{\frac{mCT}{2\pi}}\right) - \Theta\left(\frac{1}{\sqrt{T}}\right).$$

■

**Lemma 3.** *Let  $f_1 \geq f_2 \geq \dots \geq f_n$  and  $s_1 \geq s_2 \geq \dots \geq s_n$  be two non-increasing sequences of  $n$  real numbers each. Then*

$$\sum_{i=1}^n f_i s_i \geq \frac{1}{n} \left(\sum_{i=1}^n f_i\right) \left(\sum_{i=1}^n s_i\right).$$

*Proof.* From the rearrangement inequality (Hardy et al. [1952]), we have for each  $0 \leq j \leq n-1$ :

$$\sum_{i=1}^n f_i s_i \geq \sum_{i=1}^n s_i f_{(i+j) \pmod{n} + 1}, \quad (56)$$

where the modulo operator is used to cyclically shift the indices. Summing over the inequalities (56) for all  $0 \leq j \leq n-1$ , we have

$$n \sum_{i=1}^n f_i s_i \geq \left(\sum_{i=1}^n f_i\right) \left(\sum_{i=1}^n s_i\right),$$

which yields the result. □

**Lemma 4** (Bhattacharjee et al. [2020]). *Suppose that  $T$  balls are thrown independently and uniformly at random into  $2C$  bins. Let the random variable  $M_C(T)$  denote the number of balls in the most populated  $C$  bins. Then*

$$\mathbb{E}(M_C(T)) \geq \frac{T}{2} + \sqrt{\frac{CT}{2\pi}} - \Theta\left(\frac{1}{\sqrt{T}}\right).$$

## 17 Additional Experimental Results

In this section, we compare the performance of the LeadCache policy (with Pipage rounding) with other standard caching policies on two datasets taken from two different application domains. We observe that the relative performance of the algorithms remains qualitatively the same across the datasets, with the LeadCache policy consistently maintaining the highest hit rate. In our experiments, we instantiated a randomly generated bipartite network with  $n = 30$  users and  $m = 10$  caches. Each cache is connected to  $d = 8$  randomly chosen users. The capacity of each cache is taken to be 10% of the library size. Our experiments are run on HPE Apollo XL170rGen10 Servers with Dual Intel Xeon Gold 6248 20-core and 192 GB RAM.

### 17.1 Experiments with the CMU dataset [Berger et al., 2018]

**Dataset description:** This dataset is obtained from the production trace of a commercial content distribution network. It consists of 500M requests for a total of 18M objects. The popularity distribution of the requests follows approximately a Zipf distribution with the parameter  $\alpha$  between 0.85 and 1. Since we are interested only in the order in which the requests arrive, we ignore the size

Table 1: Performance Evaluation with the CMU dataset [Berger et al., 2018]

Policies	Hit Rate	Fetch Rate
LeadCache (with Pipage rounding)	<b>0.864</b>	<b>1.754</b>
LRU	0.472	13.375
LFU	0.504	13.643
Belady (offline)	0.581	5.128




Figure 5: Plot showing the popularity distribution of the files in the CMU dataset

of the objects in our experiments. Due to the massive volume of the original dataset, we consider only the first  $\sim 375K$  requests in our experiments.

Figure 5 shows the sorted overall popularity distribution of the most popular files in the dataset. It is easy to see that the popularity distribution has a light tail - a small fraction of the files are extremely popular, while others are rarely requested. The *Recall distance* measures the number of file requests between two successive requests of the same file. Figure 8 shows a plot of the empirical Recall distance distribution for this dataset.

**Experimental Results:** Figure 6 compares the dynamics of the caching policies for a particular file request pattern. It shows that the proposed LeadCache policy maintains a high cache hit rate right from the beginning. In other words, the proposed policy quickly learns the file request pattern from all users and distributes the files near-optimally on different caches. This plot also shows that the fetch rate of the LeadCache policy remains small compared to the other three caching policies. Figure 7 gives a bivariate plot of the cache hits and downloads by the LeadCache policy. From the plots, it is clear that the LeadCache policy outperforms the benchmarks on this dataset in terms of

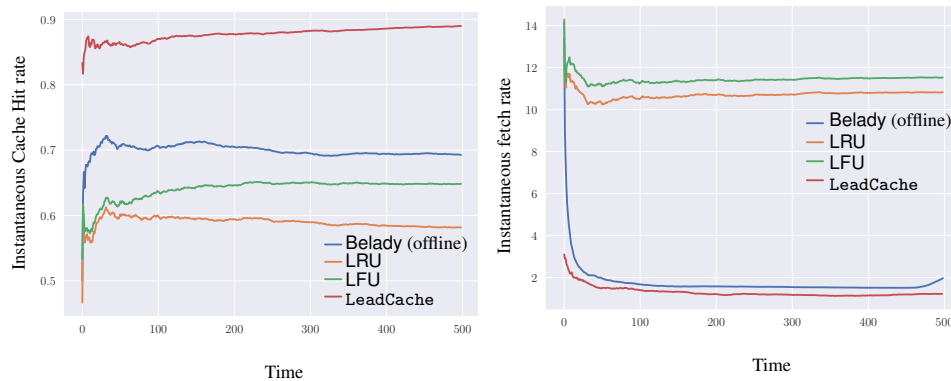


Figure 6: Temporal dynamics of instantaneous (a) cache hit rates and (b) fetch rates of different caching policies for a given file request sequence taken from the CMU dataset [Berger et al., 2018]




Figure 7: Bivariate plot of cache hit rates and the Fetch rates of different caching policies for the CMU dataset [Berger et al., 2018]

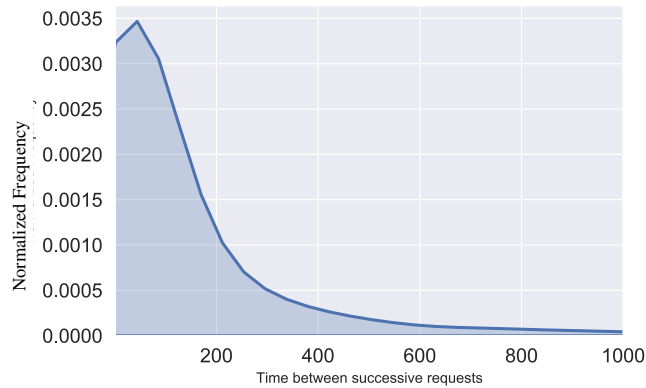


Figure 8: Distribution of time between two successive request of the same file on the CMU Dataset

both Hit rate and Fetch rate. The average values of the performance indices are summarized in Table 1.

## 17.2 Experiments with the MovieLens Dataset [Harper and Konstan, 2015]

**Dataset Description:** MovieLens<sup>4</sup> is a popular dataset containing  $\sim 20\text{M}$  ratings for  $N \sim 27278$  movies along with the timestamps of the ratings [Harper and Konstan, 2015]. The ratings were assigned by 138493 users over a period of approximately twenty years. Our working assumption is that a user rates a movie in the same sequence as she requests the movie file for download from the Content Distribution Network. Due to the sheer size of the dataset, in our experiments, we consider the first 1M ratings only. Figure 9 shows the empirical distribution of the number of times the movies have been rated (and hence, downloaded) by the users. Figure 10 shows the empirical distribution of time between two successive requests of the same file (*i.e.*, the *Recall distance*).

<sup>4</sup>This dataset is freely available from <https://www.kaggle.com/grouplens/movielens-20m-dataset>




Figure 9: Empirical Popularity distribution of the number of ratings for the MovieLens Dataset [Harper and Konstan, 2015]




Figure 10: Distribution of time between two successive request of the same file on the MovieLens Dataset

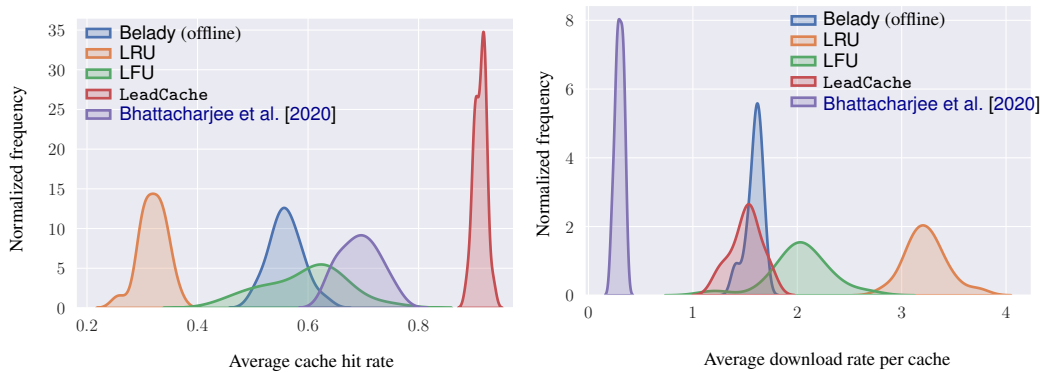


Figure 11: Empirical distributions of (a) Cache hit rates and (b) Fetch rates of different caching policies on the MovieLens Dataset.

Table 2: Performance Evaluation with the MovieLens dataset [Harper and Konstan, 2015]

Policies	Hit Rate	Fetch Rate
LeadCache (with Pipage rounding)	<b>0.991</b>	1.509
Heuristic [Bhattacharjee et al., 2020]	0.694	<b>0.297</b>
LRU	0.312	3.234
LFU	0.595	2.028
Belady (offline)	0.560	1.589

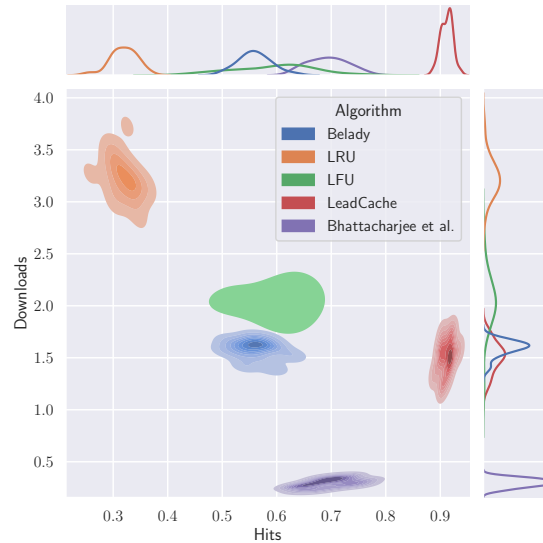


Figure 12: Bivariate plot of cache hit rates and the fetch rates of different caching policies for the MovieLens dataset [Harper and Konstan, 2015]

**Experimental Results** Figure 11 compares the performance of different policies in terms of the hit rates and fetch rates. The average values of the key performance indicators are shown in Table 2. From the plots and the table, we see that the LeadCache policy achieves the highest hit rate among all other policies, which is about 32% more than that of the Heuristic policy proposed by Bhattacharjee et al. [2020]. On the other hand, it incurs more file fetches compared to only the heuristic policy proposed by Bhattacharjee et al. [2020]. Figure 12 gives a joint plot of the hit rate and the fetch rate of different policies. It is clear from the plots that the LeadCache policy robustly learns the file request patterns and caches them on the caches near-optimally.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
  - (b) Did you describe the limitations of your work? [Yes] See Section 8.
  - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
  - (b) Did you include complete proofs of all theoretical results? [Yes] See the supplementary material.
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The code is available at <https://github.com/AbhishekMITIITM/LeadCache-NeurIPS21>.
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Please see Figure 3.
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 17 in the supplementary.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [Yes] Included the full reference to the dataset used [Berger et al., 2018, Berger, 2018].
  - (b) Did you mention the license of the assets? [Yes] In Section 7, we mentioned that the dataset is available under a BSD 2-Clause License.
  - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
  - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] Mentioned that we used an anonymized production trace. See Section 7.
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]