

# Regret-Optimal Online Caching for Adversarial and Stochastic Arrivals<sup>\*</sup>

Fathima Zarin Faizal<sup>1</sup>[0000–0002–5663–8308], Priya Singh<sup>1</sup>[0000–0002–1658–1116],  
Nikhil Karamchandani<sup>1</sup>[0000–0002–7233–0717], and Sharayu  
Moharir<sup>1</sup>[0000–0001–9393–9276]

Indian Institute of Technology Bombay, Mumbai, Maharashtra, India - 400076  
[iitb.ac.in](http://www.iitb.ac.in)

**Abstract.** We consider the online caching problem for a cache of limited size. In a time-slotted system, a user requests one file from a large catalog in each slot. If the requested file is cached, the policy receives a unit reward and zero rewards otherwise. We show that a Follow the Perturbed Leader (FTPL)-based anytime caching policy is simultaneously regret-optimal for both adversarial and i.i.d. stochastic arrivals. Further, in the setting where there is a cost associated with switching the cached contents, we propose a variant of FTPL that is order-optimal with respect to time for both adversarial and stochastic arrivals and has a significantly better performance compared to FTPL with respect to the switching cost for stochastic arrivals. We also show that these results can be generalized to the setting where there are constraints on the frequency with which cache contents can be changed. Finally, we validate the results obtained on various synthetic as well as real-world traces.

**Keywords:** Online caching · algorithms · regret bounds

## 1 Introduction

The caching problem has been studied since the 1960s, initially motivated by memory management in computers [21]. More recently, there has been renewed interest motivated by Content Delivery Networks [3] used for applications such as Video-on-Demand services. Such applications rely on low latency to provide a good customer experience. The framework of this problem involves a library of  $L$  files and a cache located near the end-users that is capable of storing at most  $C$  files at any given time, the algorithmic challenge being to determine the most popular files to be stored in the cache.

Two types of arrival patterns have been considered in the existing literature and in our work. The first is known as the Independent Reference Model, where request arrivals are generated by an i.i.d. stochastic process and the distribution of the request process is unknown to the policy. The second arrival model is

---

<sup>\*</sup> This work is supported by a SERB grant on Leveraging Edge Resources for Service Hosting.

the adversarial arrival model where we make no structural assumptions on the arrival process. Here, the arrival pattern is generated by an oblivious adversary who knows which caching policy is being used but is not aware of the sample path of decisions made by the policy. In both models, as we are focused on the online caching problem, requests are revealed causally and therefore caching decisions have to be made based on past arrival patterns without any explicit knowledge of future arrivals.

Various metrics have been used to characterize the performance of caching policies, including hit-rate and competitive ratio. Regret is a popular metric for online learning algorithms [9] and is defined as the difference between the reward incurred by the optimal stationary policy and the policy under consideration. Our broad goal is to determine if there exists caching policies that have order-optimal regret with respect to time for i.i.d. stochastic and adversarial arrivals and therefore robust to the nature of the arrival process. Policies that perform well in the adversarial setting are primarily focused on not performing horribly on any arrival sequence. This often leads to sub-optimal performance for specific arrival sequences. Similarly, policies designed for specific arrival processes or under some structural assumptions on the arrival process have poor performance in the adversarial setting as they have very poor performance for specific arrival processes which affects the worst-case performance of the policy. For instance, policies designed for the independent reference model would not be ideal when requests are not stationary.

Prediction with expert advice [9,16] and Online Convex Optimization [23] are well-known settings in online learning for which optimal algorithms have been found. Though the caching problem is equivalent to the prediction with expert advice setting with  $\binom{L}{C}$  experts,  $\binom{L}{C}$  is typically a very large number resulting in standard algorithms being computationally inefficient. Least Frequently Used (LFU), Least Recently Used (LRU) and First-in-First-Out (FIFO) are popular caching policies that have been shown to achieve optimal competitive ratio [4]. There are also results on the closed form stationary hit probabilities of these algorithms under the Independent Reference Model [6,22].

Under stochastic arrivals, LFU achieves order-optimal regret [8] but under adversarial arrivals, LFU, LRU and FIFO have been shown to have suboptimal regret [20]. A sublinear regret upper bound was proved for a gradient-based coded caching policy (OGA) under adversarial arrivals [20] while the first uncoded caching policy to be shown to achieve sublinear regret is the Follow The Perturbed Leader (FTPL) policy [7]. Proposed in [11], FTPL has also been shown to achieve order-optimal regret under adversarial arrivals by proving a lower bound on the regret using a balls-into-bins argument in [7]. An FTPL-based policy has also been shown to be regret-optimal for bipartite caching networks [19].

These policies do not consider the overhead of fetching files into the cache from the library each time the cache updates, called the *switching cost* [18]. An  $\tilde{O}(C\sqrt{T})$  upper bound on the regret including the switching cost was shown for a variant of the Multiplicative-Weight policy (MW) under adversarial arrivals

which is also more computationally efficient compared to the original MW algorithm naively applied to the caching problem. An upper bound of  $\tilde{O}(\sqrt{CT})$  was shown for an FTPL-based policy which is also simpler to implement [18] and improves upon the earlier bound by a factor of  $\Theta(\sqrt{C})$ .

### 1.1 Our contributions

We consider the following two settings: unrestricted switching, where the objective is to minimize the regret including the switching cost, and restricted switching, where the cache is allowed to update at certain fixed points only and the objective is to minimize regret. In Section 4, we consider the unrestricted switching setting and show that FTPL with an adaptive learning rate achieves order-optimal regret under stochastic arrivals even after including the switching cost, while FTPL with a constant learning rate cannot have order-optimal regret for both stochastic and adversarial arrivals. We also propose the Wait then FTPL (W-FTPL) policy that improves the bound on the switching cost from  $\mathcal{O}(D)$  to  $\mathcal{O}(\log D)$ , where  $D$  is the per-file switching cost. In Section 5, we consider the restricted switching setting and prove a lower bound on the regret of any policy and an upper bound on the regret of FTPL. We show that FTPL achieves order-optimal regret under stochastic file requests and in a special case of this setting under adversarial file requests. Finally, in Section 6, we present the results of numerical experiments on synthetic as well as real-world traces that validate the results obtained. Due to a lack of space, the proofs of the theorems stated in this paper can be found in [1].

We thus show that FTPL with an adaptive learning rate applied to the online caching problem has order-optimal regret under stochastic and adversarial arrivals in the unrestricted switching and in a special case of the restricted switching setting.

## 2 Problem formulation

We consider the classical content caching problem where a user requests files from a library that is stored in a back-end server. There is a cache that is capable of serving user requests at a lower cost but has a storage size that is typically considerably smaller than the library size. Time is slotted and in each time slot, the user requests at most one file. The sequence of events in a time slot is as follows. The cache may first update its contents, after which it receives a request for a file from the user. If the requested file is available in the cache, the cache is said to have a *hit* and the request is fulfilled locally by the cache, and otherwise a *miss*, in which case the file request is fulfilled by the back-end server.

*Cache configuration.* We consider a cache of size  $C$  that stores files from a library  $\mathcal{L}$  of size  $L$ . Usually, the cache size is much smaller than the library size, i.e.,  $C \ll L$ . The file requested by the user at time  $t$  is denoted by  $x_t$  and is represented also in the form of the one-hot encoded vector  $\mathbf{x}_t \in \{0, 1\}^L$ . For  $\tau \geq 2$ , we denote by  $\mathbf{X}_\tau = \sum_{t=1}^{\tau-1} \mathbf{x}_t$  the  $L$ -length vector storing the cumulative sum

of requests for each file till time slot  $\tau - 1$ .  $\mathbf{X}_1$  is initialized to be the zero vector. Let  $C(t)$  denote the set of files cached in round  $t$  and let  $\mathbf{y}_t \in \{0, 1\}^L$  be a binary vector denoting the state of the cache at time  $t$ , such that  $\mathbf{y}_t = (y_t^1, y_t^2, \dots, y_t^L)$  with  $y_t^i = 1$  for  $i \in C(t)$  and 0 otherwise.

*File requests.* We consider two types of file requests: adversarial and stochastic. The file requests are said to be *adversarial* if no assumptions are made regarding the statistical properties of the file requests. We assume that the adversary is oblivious, i.e., the entire file request sequence is fixed before the first request is sent. The file requests are said to be *stochastic* if in each slot, the request is generated independently according to a popularity distribution  $\boldsymbol{\mu} = (\mu_1, \dots, \mu_L)$ , where  $\mathbb{P}(x_t = i) = \mu_i$  and  $\sum_i \mu_i = 1$ . Without loss of generality, we assume that  $\mu_1 \geq \dots \geq \mu_L$ . As is the case in most real-world applications, the popularity distribution is assumed to be unknown to the caching policy beforehand.

*Caching Policy.* At the beginning of any given time slot  $t \geq 2$ , a caching policy  $\pi(\cdot)$  maps the history of observations it has seen so far (denoted by  $h(t)$ ) to a valid cache configuration  $C(t)$ , i.e.,  $C(t) = \pi(h(t))$ . In the first time slot, we assume that the cache stores  $C$  files randomly chosen from the library and that this does not incur any switch cost. We define  $T$  to be the time horizon of interest. When the file requests are adversarial, the optimal stationary policy is defined to be the caching policy that stores the top  $C$  files in hindsight, i.e., stores the  $C$  files that received the maximum number of requests till time  $T$ . When the file requests are stochastic, the optimal stationary policy is defined to be the caching policy that stores the files with the top  $C$  popularities in the cache, i.e.,  $C(t) = \mathcal{C} \forall t$ , where  $\mathcal{C} = \{1, \dots, C\}$ .

*Reward and Switch Cost.* At each time step, the policy obtains a reward of 1 unit when the requested file is available in the cache, i.e., a hit occurs, and a reward of 0 units otherwise. A caching policy that fetches a large number of files into the cache each time the cache updates is not ideal as fetching files into the cache causes latency and consumes bandwidth. Thus, we also consider the switch cost, i.e., the cost of fetching files from the back-end server into the cache. We assume that fetching a file into the cache from the server incurs a cost of  $D$  units.

*Performance metric.* Policies are evaluated on the basis of the regret that they incur. Informally, the regret of a policy till time  $T$  is the difference between the net utility of the optimal stationary policy and the net utility the policy under consideration. The net utility of a policy till time  $T$  is the difference between the net reward accrued till  $T$  and the overall switch cost incurred till then. Stationary policies do not incur any switch cost and hence their net utility is determined by the overall number of hits they have till time  $T$ . As discussed in the next section, in some cases, we omit the switch cost.

*Problem settings.* We consider the following two variations of the classical content caching problem:

1. **Setting 1:** Unrestricted switching with switching cost.

In this setting, the system incurs a cost of  $D$  units every time a file is fetched from the back-end server to be stored in the cache. When following a policy

$\pi$  on a request sequence  $\{x_t\}_{t=1}^T$ , the regret till time  $T$  for  $\{x_t\}_{t=1}^T$  including the switching cost when the file requests are adversarial is defined as:

$$R_A^\pi(\{x_t\}_{t=1}^T, T, D) = \sup_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} \rangle - \sum_{t=1}^T \mathbb{E} \langle \mathbf{y}_t, \mathbf{x}_t \rangle + \frac{D}{2} \sum_{t=1}^{T-1} \mathbb{E} \|\mathbf{y}_{t+1} - \mathbf{y}_t\|_1, \quad (1)$$

where the expectation is with respect to any randomness introduced by the policy. The regret of a policy  $\pi$  till time  $T$  is defined as the worst-case regret over all possible request sequences, i.e.,

$$R_A^\pi(T, D) = \sup_{\{x_t\}_{t=1}^T} R_A^\pi(\{x_t\}_{t=1}^T, T, D).$$

When the file requests are stochastic, the regret including the switching cost after  $T$  time steps is defined as:

$$R_S^\pi(T, D) = \mathbb{E} \left[ \sum_{t=1}^T \mathbb{1}\{x(t) \in \mathcal{C}\} - \mathbb{1}\{x(t) \in C(t)\} \right] + \frac{D}{2} \sum_{t=1}^{T-1} \mathbb{E} \|\mathbf{y}_{t+1} - \mathbf{y}_t\|_1, \quad (2)$$

where  $\mathcal{C}$  denotes the set of files having the top  $C$  popularities. In the above expression, the expectation is taken with respect to the randomness in the file requests as well as any randomness introduced by the policy.

## 2. **Setting 2:** Restricted switching without switching cost.

Here, the cache is allowed to change its contents only at  $s + 1$  fixed time slots for some  $1 \leq s \leq T$ . To be precise, the cache is allowed to change its contents only at the beginning of the following time slots:  $1, r_1 + 1, r_1 + r_2 + 1, \dots, \sum_{i=1}^s r_i + 1$ , where  $1 \leq r_i \leq T, 1 \leq i \leq s$  denotes the  $i^{\text{th}}$  inter-switching period such that  $\sum_{i=1}^s r_i = T$ . Thus, within the time horizon  $T$ , the cache is allowed to update only at  $s$  fixed time slots. Note that the setting where the cache is allowed to change its contents only after every  $1 \leq r \leq T$  requests, i.e., at time slots  $1, r + 1, \dots, T + 1$  and  $s = \frac{T}{r}$  is a special case of this setting. For simplicity, we restrict our attention to the case where there is no switch cost, i.e.,  $D = 0$ . When following a policy  $\pi$ , the regret after  $T$  time steps when the file requests are adversarial is:

$$R_A^\pi(T) = \sup_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} \rangle - \sum_{t=1}^T \mathbb{E} \langle \mathbf{y}_t, \mathbf{x}_t \rangle, \quad (3)$$

where the expectation is with respect to the randomness introduced by the policy. When the file requests are stochastic, the regret after  $T$  time steps is:

$$R_S^\pi(T) = \mathbb{E} \left[ \sum_{t=1}^T \mathbb{1}\{x(t) \in \mathcal{C}\} - \mathbb{1}\{x(t) \in C(t)\} \right], \quad (4)$$

where  $\mathcal{C}$  denotes the set of files having the top  $C$  popularities. In the above expression, the expectation is taken with respect to the randomness introduced by the policy and the file requests.

To distinguish between results including switching cost and those without switching cost, we use the notation  $R_{(\cdot)}^\pi(T, D)$  for results involving the switch cost and  $R_{(\cdot)}^\pi(T)$  for results without the switch cost.

The overall goal of this work is to characterize the optimal regret in the two settings mentioned above, for both adversarial and stochastic file requests. This entails proving scheme-agnostic lower bounds on the regret as well as designing policies whose regret is of the same order as these lower bounds. As we will see, these results will also highlight the impact of switching cost and intermittent switching on the optimal achievable regret.

### 3 Policies

In this section, we introduce and formalize policies whose optimality (or suboptimality) will be discussed in later sections.

#### 3.1 Least Frequently Used (LFU)

The LFU algorithm (formally defined in Algorithm 1) keeps track of the number of times each file has been requested so far. At each time step  $t$ , the files with the  $C$  highest number of requests are cached. This policy is deterministic and thus performs poorly when faced with certain adversarial request sequences [20]. For the simplified case of  $L = 2, C = 1$ , one example is a round-robin request sequence of the form  $1, 2, 1, 2, \dots$  which would result in LFU obtaining essentially zero reward while the optimal stationary policy obtains a reward of  $T/2$ . For stochastic requests, it has been shown to achieve  $\mathcal{O}(1)$  regret when switching is allowed at all time slots and when the algorithm incurs no switch cost [8].

---

#### Algorithm 1 LFU algorithm

---

```

1: procedure LFU( $T$ )
2:    $\mathbf{c}_t \leftarrow \mathbf{0}$ 
3:   while  $t \leq T$  do
4:      $C_t \leftarrow \arg \max_C (c_t(1), \dots, c_t(L))$ 
5:     Receive file request  $x_t$ 
6:      $c_t(x_t) \leftarrow c_t(x_t) + 1$ 
7:   end while
8: end procedure

```

---

#### 3.2 Follow The Perturbed Leader (FTPL)

The FTPL algorithm (formally defined in Algorithm 2)) is a variation of the LFU algorithm and also keeps track of the number of times each file has been

requested so far, but adds an independent Gaussian perturbation with mean 0 and standard deviation  $\eta_t$  (referred to as the learning rate) to the counts of each file in each time slot. At each time step  $t$ , the files with the  $C$  highest perturbed counts are cached. Special cases of this policy are known to achieve order-optimal regret under adversarial requests (with and without switch cost) [7,18]. We will henceforth refer to the FTPL algorithm with the learning rate  $\eta_t$  by FTPL( $\eta_t$ ).

---

**Algorithm 2** FTPL algorithm
 

---

```

1: procedure FTPL( $T, \{\eta_t\}_{t=1}^T$ )
2:    $\mathbf{c}_t \leftarrow \mathbf{0}$ 
3:   Sample  $\boldsymbol{\gamma} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{L \times L})$ 
4:   while  $t \leq T$  do
5:      $C_t \leftarrow \arg \max_C (c_t(1) + \eta_t \gamma(1), \dots, c_t(L) + \eta_t \gamma(L))$ 
6:     Receive file request  $x_t$ 
7:      $c_t(x_t) \leftarrow c_t(x_t) + 1$ 
8:   end while
9: end procedure

```

---

### 3.3 Wait then FTPL (W-FTPL)

The algorithm that we propose, Wait then FTPL (formally defined in Algorithm 3), is a variant of the FTPL algorithm where the policy remains idle for an initial deterministic waiting period and then follows the normal FTPL algorithm. The motivation for this algorithm is to avoid the higher switch cost incurred initially by the FTPL algorithm under stochastic file requests until the policy has seen enough requests to have a good enough estimate of the underlying popularity distribution, while ensuring order-optimal regret in the adversarial setting. We will henceforth refer to the W-FTPL algorithm with the learning rate  $\eta_t$  by W-FTPL( $\eta_t$ ).

---

**Algorithm 3** W-FTPL algorithm
 

---

```

1: procedure W-FTPL( $T, \{\eta_t\}_{t=1}^T, D, t'$ )
2:    $\mathbf{c}_t \leftarrow \mathbf{0}$ 
3:   Sample  $\boldsymbol{\gamma} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{L \times L})$ 
4:   while  $t \leq T$  do
5:     if  $t > t'$  then
6:        $C_t \leftarrow \arg \max_C (c_t(1) + \eta_t \gamma(1), \dots, c_t(L) + \eta_t \gamma(L))$ 
7:     end if
8:     Receive file request  $x_t$ 
9:      $c_t(x_t) \leftarrow c_t(x_t) + 1$ 
10:  end while
11: end procedure

```

---

## 4 Setting 1: Unrestricted switching with switching cost

In this section, we consider the setting where there is no limitation on the switching frequency of the cache and the objective is to minimize the regret including the switching cost, i.e., minimize regret as well as the number of fetches into the cache. We consider both stochastic and adversarial file request sequences and show that  $\text{FTPL}(\alpha\sqrt{t})$  and  $\text{W-FTPL}(\alpha\sqrt{t})$  are order-optimal under both types of file requests. While the  $\text{FTPL}(\eta)$  algorithm is order-optimal under adversarial requests for a particular value of  $\eta$  [7], we prove that the same does not hold true for stochastic file requests.

### 4.1 Adversarial requests

In this section, we discuss the performance of the policies introduced in Section 3 under adversarial file requests. The key results of this section has been summarized in the following theorem:

**Theorem 1.** *Under adversarial requests, we have*

(a) [7, Theorem 2] *For any policy  $\pi$  and  $L \geq 2C$ ,*

$$R_A^\pi(T, D = 0) \geq \sqrt{\frac{CT}{2\pi}} - \Theta\left(\frac{1}{\sqrt{T}}\right).$$

(b) [20, Proposition 1] *The regret of the LFU policy can be characterized as:*

$$R_A^{\text{LFU}}(T, 0) = \Omega(T).$$

(c) [18, Theorem 4.1] *The regret of  $\text{FTPL}(\alpha\sqrt{t})$  is upper bounded as:*

$$R_A^{\text{FTPL}(\alpha\sqrt{t})}(T, D) \leq c_1\sqrt{T} + c_2 \ln T + c_3,$$

where  $c_1 = \mathcal{O}(\sqrt{\ln(Le/C)})$ , and  $c_2, c_3$  are small constants depending on  $L, C, D$  and  $\alpha$ .

(d) *The regret of  $\text{W-FTPL}(\alpha\sqrt{t})$  is upper bounded as:*

$$R_A^{\text{W-FTPL}(\alpha\sqrt{t})}(T, D) \leq \mathcal{O}(\sqrt{T}).$$

Part (a) has been proved in [7] and provides a lower bound on the regret of any policy under adversarial requests.

As argued before, LFU performs poorly under adversarial requests. This is also seen for many popular classical caching algorithms like LRU and FIFO (refer [20]).

Part (c) has been proved in [18] and provides an  $\mathcal{O}(\sqrt{T})$  upper bound on the regret including the switching cost of the  $\text{FTPL}(\alpha\sqrt{t})$  policy under adversarial requests, thus showing that this algorithm is order-optimal under adversarial requests.  $\text{FTPL}(\alpha\sqrt{T})$  has also been shown to be order-optimal under adversarial requests [7,18].

Part (d) provides an upper bound on the regret including the switching cost of  $\text{W-FTPL}(\alpha\sqrt{t})$  under adversarial requests. This result shows that  $\text{W-FTPL}(\alpha\sqrt{t})$  is order-optimal under adversarial requests. The proof of this result can be found in Appendix A.

## 4.2 Stochastic requests

To find a policy that achieves order-optimal regret under stochastic and adversarial arrivals, we were motivated by [7,18] where the regret for FTPL with the learning rates  $\alpha\sqrt{t}$  and  $\alpha\sqrt{T}$  ( $\alpha$  being some positive constant) under adversarial arrivals was characterized. In this section, we discuss the performance of these policies under stochastic file requests. The key results of this section has been summarized in the following theorem:

**Theorem 2.** *The file requests are stochastic i.i.d. with the popularity distribution  $\mu$ .*

(a) [8, Theorem 1] *When  $D = 0$ , the regret of the LFU policy can be upper bounded as:*

$$R_S^{LFU}(T, 0) < \min\left(\frac{16}{\Delta_{\min}^2}, \frac{4C(L-C)}{\Delta_{\min}}\right), \quad (5)$$

where  $\Delta_{\min} = \mu_C - \mu_{C+1}$ .

(b) *For  $L = 2, C = 1$  and  $D = 0$ , the regret of FTPL( $\eta$ ) can be lower bounded as:*

$$R_S^{FTPL(\eta)}(T, 0) \geq \frac{\eta e^{-(\frac{1+\eta}{\eta})^2}}{4}.$$

(c) *The regret of FTPL( $\alpha\sqrt{t}$ ) is upper bounded as:*

$$R_S^{FTPL(\sqrt{t})}(T, D) \leq (1 + DC)t_0 + \left(1 + \frac{D}{\Delta_{\min}}\right) \left(\frac{8}{\Delta_{\min}} + \frac{32\alpha^2}{\Delta_{\min}}\right),$$

where  $t_0 = \max\left\{\frac{8}{\Delta_{\min}^2} \log(L^3), \frac{32\alpha^2}{\Delta_{\min}^2} \log(L^3)\right\}$ .

(d) *The regret of W-FTPL( $\alpha\sqrt{t}$ ) is upper bounded as:*

$$R_S^{W-FTPL(\sqrt{t})}(T, D) \leq t' + \frac{16}{\Delta_{\min}} + \frac{64\alpha^2}{\Delta_{\min}} + 2L^3 D \left( e^{-u(\log D)^{1+\beta}} \frac{\Delta_{\min}^2}{8} \frac{8}{\Delta_{\min}^2} \right. \\ \left. + e^{-u(\log D)^{1+\beta}} \frac{\Delta_{\min}^2}{32\alpha^2} \frac{32\alpha^2}{\Delta_{\min}^2} \right),$$

where  $t' = \max\left\{\frac{8}{\Delta_{\min}^2} \log\left(\frac{L^3}{2}\right), \frac{32\alpha^2}{\Delta_{\min}^2} \log\left(\frac{L^3}{2}\right), u(\log D)^{1+\beta}\right\}$ .

Part (a) of the above theorem has been proved in [8] and shows that the regret of LFU is  $\mathcal{O}(1)$  when the file requests are stochastic. Thus, the regret of any policy that has order optimal regret under stochastic file requests should be  $\mathcal{O}(1)$ .

Part (b) gives a lower bound on the regret of the FTPL( $\eta$ ) algorithm under stochastic file requests. Note that for all  $\eta \geq 1$ , we have

$$R_S^{FTPL(\eta)}(T) \geq \frac{\eta e^{-(\frac{1+\eta}{\eta})^2}}{4} \geq \frac{\eta}{4e^4}.$$

This shows that the regret is  $\Omega(\eta)$  for the FTPL( $\eta$ ) algorithm when the file requests are stochastic. Using  $\eta = \alpha\sqrt{T}$ , where  $\alpha$  is a positive constant, from [7] which gave an  $\mathcal{O}(\sqrt{T})$  upper bound for FTPL( $\eta$ ), we get that with the constant learning rate  $\eta = \alpha\sqrt{T}$ , the regret of FTPL( $\eta$ ) is  $\Theta(\sqrt{T})$ . Thus, while FTPL( $\eta$ ) is order-optimal under adversarial requests (see Section 4.1), it cannot simultaneously be order-optimal under adversarial and stochastic file requests. The proof of this result can be found in Appendix B.

Part (c) shows that FTPL( $\alpha\sqrt{t}$ ),  $\alpha > 0$  is order-optimal when the file requests are stochastic. While FTPL( $\eta$ ) with  $\eta = \mathcal{O}(\sqrt{T})$  achieves  $\Omega(\sqrt{T})$  regret, this result shows an  $\mathcal{O}(1)$  upper bound on the regret of FTPL( $\alpha\sqrt{t}$ ) including the switching cost, thus showing that this algorithm is order-optimal. Note that the upper bound grows linearly with the per-file switch cost  $D$ . The proof of this result can be found in Appendix C

Recall that the Wait then FTPL( $\alpha\sqrt{t}$ ) (W-FTPL( $\alpha\sqrt{t}$ )) algorithm is a variant of the FTPL( $\alpha\sqrt{t}$ ) algorithm. The algorithm remains idle till time  $t' = u(\log D)^{1+\beta}$ ,  $u > 0$ , and then normal FTPL( $\alpha\sqrt{t}$ ) is followed. Part (d) proves an  $\mathcal{O}(1)$  upper bound on the regret including the switching cost of this algorithm with respect to the horizon  $T$  under stochastic file requests, thus showing that this algorithm is order-optimal. The main improvement over the FTPL( $\alpha\sqrt{t}$ ) algorithm is the  $\mathcal{O}((\log D)^{1+\beta})$  upper bound on the regret including the switching cost for a large enough value of  $D$  under stochastic file requests, as compared to the upper bound  $\mathcal{O}(D)$  for FTPL( $\alpha\sqrt{t}$ ). The key idea behind remaining idle for an initial period that depends logarithmically on  $D$  is to avoid the higher switch cost incurred at the beginning by the FTPL( $\alpha\sqrt{t}$ ) algorithm (refer to Section 6). The proof of this result can be found in Appendix D.

## 5 Restricted switching

In this section, we consider the setting where the cache is allowed to update its contents only at  $s + 1$  fixed number of time slots, where  $s \in \mathbb{Z}$ ,  $s \leq T$ . The first point is at time slot 1, the second at time slot  $r_1 + 1$ , the third point is at time slot  $r_1 + r_2 + 1$ , and so on till the  $s + 1^{\text{th}}$  point, which is at time slot  $\sum_{i=1}^s r_i + 1$ , where  $r_i \in \mathbb{Z}$ ,  $r_i \leq T$ ,  $1 \leq i \leq s$  and  $\sum_{i=1}^s r_i = T$ . Note that the cache is allowed to update its contents only  $s$  times till the time horizon  $T$ . Refer to Figure 1 for an illustration of this setting. As a special case of this setting, we also consider the homogenous case where the cache is allowed to update only after every  $r \in \mathbb{Z}$  requests, i.e.,  $s = \frac{T}{r}$ . We study the regret performance of FTPL and also provide lower bounds on the regret incurred by any online scheme. In the homogenous case, we also show that FTPL( $\sqrt{rt}$ ) achieves order-optimal regret.

### 5.1 Stochastic requests

**Theorem 3.** *The file requests are stochastic i.i.d. with the popularity distribution  $\mu$ . When cache updates are restricted to  $s + 1$  fixed points defined by the inter-switching periods  $\{r_i\}_{i=1}^s$  as outlined above,*




Fig. 1: The time slots where the cache is allowed to update its contents have been marked in yellow.

- (a) When  $L = 2, C = 1$ , for any online caching policy  $\pi$ , there exists a popularity distribution such that the popularities of the two files are greater than  $1 > a > 0$  and the difference in the popularities is  $\Delta$ , such that

$$R_S^\pi(T) \geq \frac{r_1 \Delta}{2} + \sum_{i=2}^s r_i \frac{\Delta}{4} \exp\left(-t_i \frac{\Delta^2}{a^2}\right).$$

- (b) The regret of  $FTPL(\alpha\sqrt{t})$  is upper bounded as:

$$R_S^{FTPL(\alpha\sqrt{t})}(T) \leq r_1 + 2 \sum_{j=1}^C \sum_{k=C+1}^L \sum_{i=2}^s r_i \Delta_{j,k} \left( e^{-t_i \Delta_{j,k}^2 / 8} + e^{-t_i \Delta_{j,k}^2 / 32\alpha^2} \right).$$

In part (a), we prove a fundamental lower bound on the regret of any policy  $\pi$  under stochastic file requests when cache updates are restricted to  $s + 1$  fixed time slots. The proof of this result can be found in Appendix E. In part (b), we prove an upper bound on the regret of the  $FTPL(\alpha\sqrt{t})$  policy when cache updates are restricted to  $s + 1$  fixed time slots. The proof of this result can be found in Appendix F. We thus have that the  $FTPL(\alpha\sqrt{t})$  policy has order-optimal regret in this setting under stochastic file requests. Next, we consider the special case where all  $r_i$  are equal to  $r$ , i.e.,  $s = T/r$ .

**Theorem 4.** *The file requests are stochastic i.i.d. with the popularity distribution  $\mu$ . When the cache is allowed to update only after every  $r$  requests,*

$$R_S^{FTPL(\alpha\sqrt{t})} \leq 1 + t'_0 + 2 \left( \frac{8}{\Delta_{\min}^2} + \frac{32\alpha^2}{\Delta_{\min}^2} \right).$$

where  $t'_0 = \max \left\{ r, \frac{8}{\Delta_{\min}^2} \log(L^2), \frac{32\alpha^2}{\Delta_{\min}^2} \log(L^2) \right\}$ .

In Theorem 4, we prove an  $\mathcal{O}(\max\{r, \log L\})$  upper bound on the regret of the  $FTPL(\alpha\sqrt{t})$  algorithm under stochastic file requests. While the order-optimality of this policy with respect to  $r$  follows from Theorem 3, we also note that the bound proved here improves upon the worst-case  $\mathcal{O}(L^2)$  dependency in the upper bound proved in part (b) of Theorem 3 to  $\mathcal{O}(\log L)$ . The proof of this result can be found in Appendix G.

## 5.2 Adversarial requests

**Theorem 5.** *Files are requested by an oblivious adversary. When cache updates are restricted to  $s + 1$  fixed points defined by  $\{r_i\}_{i=1}^s$  as outlined above,*

(a) *For any online caching policy  $\pi$  and  $L \geq 2C$ ,*

$$R_A^\pi(T) \geq \frac{1}{2} \left( 0.15 \sqrt{C \sum_{i=1}^s r_i^2} \left( 1 - \frac{(C-1) (\sum_{i=1}^s r_i^4)}{2 (\sum_{i=1}^s r_i^2)^2} \right) - 0.6 C \max_{1 \leq i \leq s} r_i \right).$$

(b) *The regret of  $FTPL(\alpha\sqrt{t})$  is upper bounded as:*

$$R_A^{FTPL(\alpha\sqrt{t})}(T) \leq \mathcal{O}(\alpha\sqrt{T}) + \sqrt{\frac{2}{\pi}} \sum_{i=1}^s \frac{r_i^2}{\alpha \sqrt{\sum_{j=0}^{i-1} r_j}}.$$

The proof of this theorem can be found in Appendix H and Appendix I respectively. In part (a), we prove a lower bound on the regret of any policy  $\pi$  under adversarial file requests when cache updates are restricted to  $s$  fixed time slots. Note that a necessary condition for this bound to be meaningful is

$$4 \max_{1 \leq i \leq s} r_i \leq \sqrt{\sum_{i=1}^s r_i^2}. \quad (6)$$

When this bound is meaningful, we have that  $R_A^\pi(T) = \Omega\left(\sqrt{C \sum_{i=1}^s r_i^2}\right)$  for any online caching policy  $\pi$ . When all the  $r_i$ 's are equal, this condition translates to  $r \leq T/16$ . This condition does not hold if any of the  $r_i$ 's is too large. For instance, when  $s = 3$  and  $r_1 = T/2, r_2 = r_3 = T/4$ , this condition does not hold. When  $\max_{1 \leq i \leq s} r_i$  and  $\min_{1 \leq i \leq s} r_i$  are known, a sufficient condition for (6) to hold is:

$$\frac{\max_{1 \leq i \leq s} r_i}{\min_{1 \leq i \leq s} r_i} \leq \frac{\sqrt{s}}{4}.$$

We now discuss the special case where all  $r_i$  are equal to  $r$ , i.e.,  $s = T/r$ . It follows from part (b) of Theorem 5 that  $FTPL(\sqrt{rt})$  achieves a regret of  $\mathcal{O}(\sqrt{rT})$ . The following theorem provides a matching lower bound for this setting that proves that  $FTPL(\sqrt{rt})$  is order-optimal, the proof of which can be found in Appendix J.

**Theorem 6.** *Files are requested by an oblivious adversary. When the cache is allowed to update only after every  $r$  requests, for any online caching policy  $\pi$  and  $L \geq 2C$ ,*

$$R_A^\pi(T) \geq \begin{cases} \sqrt{\frac{CrT}{2\pi}} - \Theta\left(\frac{r\sqrt{r}}{\sqrt{T}}\right), & \text{when } r = o(T), \\ \Omega(T), & \text{when } r = \Omega(T). \end{cases}$$

## 6 Numerical experiments

In this section, we present the results of numerical simulations for the various policies discussed in Section 3.

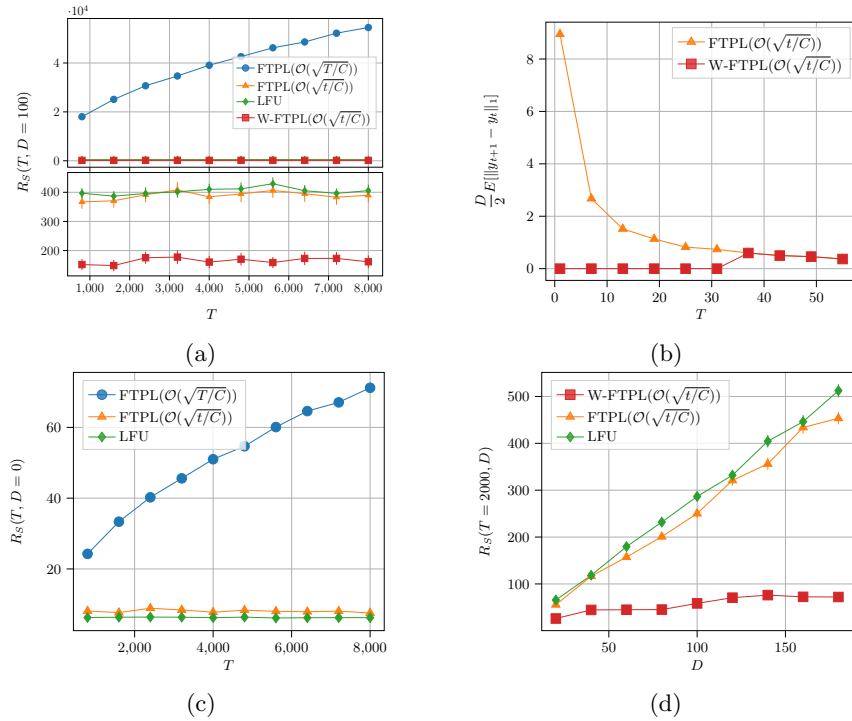


Fig. 2: The plots compare (a) the regret including the switching cost for  $D = 100$  as a function of  $T$ , (b) the switching cost in each time slot for  $D = 30, u = 5, \beta = 0.6$  as a function of  $T$ , (c) the regret as a function of  $T$  with  $D = 0$ , and (d) the regret including the switching cost for  $T = 2000$  as a function of  $D$ , of various caching policies under stochastic file requests.  $L = 10, C = 4$  for each of the plots and the popularity distribution is a dyadic distribution. Parts (a) and (c) show that the regret incurred by FTPL( $\mathcal{O}(\sqrt{T})$ ) is increasing with  $T$  while FTPL( $\mathcal{O}(\sqrt{t})$ ), W-FTPL( $\mathcal{O}(\sqrt{t})$ ) and LFU have essentially constant regret. Part (b) shows that FTPL( $\mathcal{O}(\sqrt{t})$ ) makes more switches at the beginning, thus motivating the W-FTPL( $\mathcal{O}(\sqrt{t})$ ) algorithm. Part (d) shows that while the regret including the switching cost of LFU and FTPL( $\mathcal{O}(\sqrt{t})$ ) increase linearly in  $D$ , it increases sublinearly for W-FTPL( $\mathcal{O}(\sqrt{t})$ ).

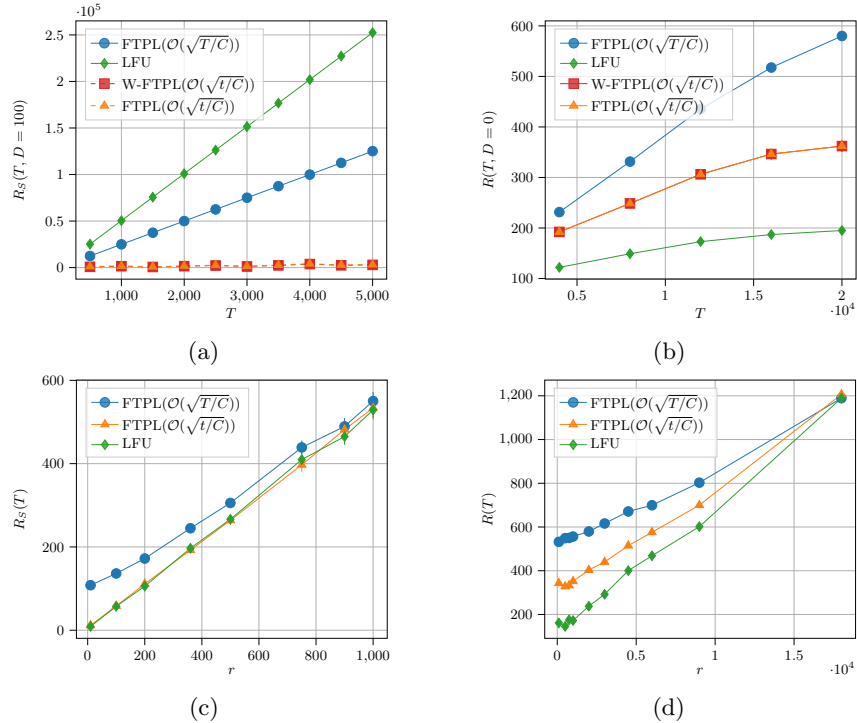


Fig. 3: The plots compare (a) the regret including the switching cost for  $D = 100$  as a function of  $T$  on a round robin request sequence, (b) the regret without the switching cost, i.e.,  $D = 0$  as a function of  $T$  on the MovieLens dataset, (c) the regret as a function of the constant switching frequency  $r$  under stochastic file requests from a dyadic distribution, and (d) the regret as a function of the switching frequency  $r$  on the MovieLens dataset, of various caching policies. We used  $L = 2, C = 1$  for Part (a) and  $L = 10, C = 4$  for Part (c). In Part (a), note that the regret including the switching cost scales linearly with  $T$  for LFU, while W-FTPL( $\mathcal{O}(\sqrt{t})$ ) and FTPL( $\mathcal{O}(\sqrt{t})$ ) show better performance. In Part (c), the regret scales linearly with  $r$  for all the three algorithms as expected. The regret scales sublinearly with  $T$  in Part (b) and linearly with  $r$  in Part (d).

### 6.1 Setting 1 with stochastic file requests

*Setup.* We use  $L = 10, C = 4$  throughout this section. The popularity distribution used is a dyadic distribution, i.e., for  $1 \leq i \leq L - 1$ ,  $\mu(i) = \frac{1}{2^i}$  and  $\mu(L) = \frac{1}{2^{L-1}}$ .

*Results.* Figure 2a shows that the regret of FTPL including the switching cost increases with  $T$ , while it is essentially constant for FTPL( $\mathcal{O}(\sqrt{t})$ ), W-FTPL( $\mathcal{O}(\sqrt{t})$ ) and LFU. One can also observe that W-FTPL( $\mathcal{O}(\sqrt{t})$ ) performs the best among all the four algorithms. In Figure 2c, we plot only the regret as a function of  $T$ . Note that the same trend is observed here as well. Here, we omit plotting

W-FTPL( $\mathcal{O}(\sqrt{t})$ ) as its regret would be the same as that of FTPL( $\mathcal{O}(\sqrt{t})$ ) in this case. Figure 2b shows that FTPL( $\mathcal{O}(\sqrt{t})$ ) indeed makes more switches at the beginning, which is the motivation for the W-FTPL( $\mathcal{O}(\sqrt{t})$ ) policy where no switches are made for an initial period. Figure 2d shows that the regret of LFU and FTPL( $\mathcal{O}(\sqrt{t})$ ) grows linearly with  $D$ , while that of W-FTPL( $\mathcal{O}(\sqrt{t})$ ) grows sublinearly with  $D$ .

## 6.2 Setting 1 with adversarial file requests

*Setup.* We consider a synthetic adversarial request sequence in Figure 3a and a real-world trace in Figure 3b. The synthetic adversarial request sequence used is  $1, 2, 1, 2, \dots$  for  $L = 2, C = 1$ , i.e., a round-robin request sequence. The real-world trace used is the first 20,000 rows of the MovieLens 1M dataset [2,12] which contains ratings for 2569 movies with timestamps that we model as requests to a CDN server of library size 2569 and a cache size of 25.

*Results.* Figure 3a shows that under the round-robin request sequence, the regret including the switching cost of LFU scales linearly with  $T$ , while that of W-FTPL( $\mathcal{O}(\sqrt{t})$ ) and FTPL( $\mathcal{O}(\sqrt{t})$ ) scales sublinearly with  $T$ . Figure 3b shows that on the MovieLens dataset, the regret scales sublinearly with  $T$  for all the four algorithms.

## 6.3 Setting 2

*Setup.* We consider stochastic file requests drawn from a dyadic distribution for  $L = 10, C = 4$  in Figure 3c and file requests from the MovieLens dataset in Figure 3d. We also used  $T = 18000$  and chose  $r$  to be factors of  $T$ . There were 2518 unique movies in the first 18000 rows of the MovieLens dataset and we set the cache size to be 25.

*Results.* Figure 3c shows that when file requests are drawn from a dyadic popularity distribution, the regret of all three policies vary linearly with  $r$ . Figure 3d shows that on the MovieLens dataset, the regret scales linearly with  $r$  for all the three policies.

## 7 Conclusion

We have shown that FTPL( $\mathcal{O}(\sqrt{t})$ ) achieves order-optimal regret even after including the switching cost under stochastic requests. Combining with prior results on the performance of FTPL, it is simultaneously order-optimal for both stochastic and adversarial requests. We also showed that FTPL( $\eta$ ) cannot possibly achieve order-optimal regret simultaneously under both stochastic and adversarial requests, while variants of this policy can individually be order-optimal under each type of request. We proposed the W-FTPL( $\mathcal{O}(\sqrt{t})$ ) policy as a way of preventing the high switching cost incurred by FTPL at the beginning under stochastic file requests. We also considered the restricted switching setting, where the cache is allowed to update its contents only at specific pre-determined

time slots and obtained a lower bound on the regret incurred by any policy. We proved an upper bound on the regret of  $\text{FTPL}(\mathcal{O}(\sqrt{t}))$  and showed that it is order-optimal under stochastic file requests and in the homogenous restricted switching case under adversarial file requests.

This work motivates several directions for future work: (1) Bringing the upper and lower bounds closer in the general restricted switching setting would help in proving whether  $\text{FTPL}(\mathcal{O}(\sqrt{t}))$  is order-optimal or not under adversarial file requests in this case too. (2) For the restricted switching setting, our results consider only the regret. Adding the switching cost too here would make the bounds complete.

## References

1. [https://github.com/fathimazarin/Valuetools2022\\_caching/blob/main/main.pdf](https://github.com/fathimazarin/Valuetools2022_caching/blob/main/main.pdf)
2. MovieLens 1M dataset. <https://grouplens.org/datasets/movielens/>, [accessed 8-August-2022]
3. Aggarwal, C., Wolf, J.L., Yu, P.S.: Caching on the world wide web. In: 125 Journal of Distributed and Parallel Systems (IJDPSS) Vol.2, No.6. pp. 94–107 (2000). <https://doi.org/10.1109/69.755618>
4. Albers, S.: Competitive online algorithms. BRICS (1996)
5. Alon, N., Spencer, J.: Appendix B: Paul Erdős. John Wiley and Sons, Ltd (2008). <https://doi.org/https://doi.org/10.1002/9780470277331.app2>
6. Aven, O.I., Coffman, E.G., Kogan, Y.A.: Stochastic Analysis of Computer Storage. Kluwer Academic Publishers, USA (1987)
7. Bhattacharjee, R., Banerjee, S., Sinha, A.: Fundamental limits of online network-caching. CoRR **abs/2003.14085** (2020). <https://doi.org/10.48550/arXiv.2003.14085>
8. Bura, A., Rengarajan, D., Kalathil, D., Shakkottai, S., Chamberland, J.F.: Learning to cache and caching to learn: Regret analysis of caching algorithms. IEEE/ACM Transactions on Networking **30**(1), 18–31 (2022). <https://doi.org/10.48550/arXiv.2004.00472>
9. Cesa-Bianchi, N., Lugosi, G.: Prediction, learning, and games. Cambridge university press (2006). <https://doi.org/10.1017/CBO9780511546921>
10. Cohen, A., Hazan, T.: Following the perturbed leader for online structured learning. In: Bach, F., Blei, D. (eds.) Proceedings of the 32nd International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 37, pp. 1034–1042. PMLR, Lille, France (07–09 Jul 2015)
11. Hannan, J.: Approximation to bayes risk in repeated play. Contributions to the Theory of Games **3**(2), 97–139 (1957). <https://doi.org/10.1515/9781400882151>
12. Harper, F.M., Konstan, J.A.: The movielens datasets: History and context. ACM Trans. Interact. Intell. Syst. **5**(4) (dec 2015). <https://doi.org/10.1145/2827872>
13. Hoeffding, W.: Probability inequalities for sums of bounded random variables. Journal of the American Statistical Association **58**(301), 13–30 (1963). <https://doi.org/10.1080/01621459.1963.10500830>
14. Kaas, R., Buhrman, J.M.: Mean, median and mode in binomial distributions. Statistica Neerlandica **34**(1), 13–18 (1980). <https://doi.org/10.1111/j.1467-9574.1980.tb00681.x>
15. Lattimore, T., Szepesvári, C.: Bandit Algorithms. Cambridge University Press (2020). <https://doi.org/10.1017/9781108571401>

16. Littlestone, N., Warmuth, M.K.: The weighted majority algorithm. *Information and computation* **108**(2), 212–261 (1994). <https://doi.org/10.1006/inco.1994.1009>
17. Mourtada, J., Gaïffas, S.: On the optimality of the hedge algorithm in the stochastic regime. *Journal of Machine Learning Research* **20**, 1–28 (2019). <https://doi.org/10.48550/arXiv.1809.01382>
18. Mukhopadhyay, S., Sinha, A.: Online caching with optimal switching regret. *CoRR* **abs/2101.07043** (2021). <https://doi.org/10.1109/ISIT45174.2021.9517925>
19. Paria, D., Sinha, A.: Leadcache: Regret-optimal caching in networks. In: *Thirty-Fifth Conference on Neural Information Processing Systems*. vol. 34, pp. 4435–4447 (2021). <https://doi.org/10.48550/arXiv.2009.08228>
20. Paschos, G.S., Destounis, A., Vigneri, L., Iosifidis, G.: Learning to cache with no regrets. In: *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*. p. 235–243. IEEE Press (2019). <https://doi.org/10.1109/INFOCOM.2019.8737446>
21. Silberschatz, A., Galvin, P., Gagne, G.: *Operating System Principles*. John Wiley & Sons (2006)
22. Starobinski, D., Tse, D.: Probabilistic methods for web caching. *Perform. Eval.* **46**(2–3), 125–137 (oct 2001). [https://doi.org/10.1016/S0166-5316\(01\)00045-1](https://doi.org/10.1016/S0166-5316(01)00045-1)
23. Zinkevich, M.: Online convex programming and generalized infinitesimal gradient ascent. In: *Proceedings of the 20th international conference on machine learning (icml-03)*. pp. 928–936 (2003)

## A Proof of part (d) of Theorem 1

In this section, we prove an upper bound on the regret of  $\text{W-FTPL}(\alpha\sqrt{t})$  under adversarial requests. The proof mostly follows that of Theorem 4.1 of [18] and is given here for completeness. We bound the regret till time  $t'$  by  $t'$  and bound the regret incurred from time  $t' + 1$  in a manner similar to [18]. Thus,

$$\begin{aligned} R_A^{\text{W-FTPL}(\alpha\sqrt{t})}(T) &= \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} \rangle - \sum_{t=1}^T \mathbb{E}_\gamma [\langle \mathbf{y}_t, \mathbf{x}_t \rangle] \\ &\leq t' + \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} - \mathbf{X}_{t'+1} \rangle - \sum_{t=t'+1}^T \mathbb{E}_\gamma [\langle \mathbf{y}_t, \mathbf{x}_t \rangle]. \end{aligned} \quad (7)$$

We define the potential function  $\Phi_t : \mathbb{R}^L \rightarrow \mathbb{R}$  for all time instants  $t$  in the following way:

$$\Phi_t(\mathbf{x}) = \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{x} + \eta_t \boldsymbol{\gamma} \rangle \right],$$

where  $\mathcal{Y}$  is the set of possible cache configurations, i.e., the set  $\{\mathbf{y} \in \{0, 1\}^L : \|\mathbf{y}\|_1 \leq C\}$ . As shown in the proof of Proposition 4.1 in [18], we have

$$\mathbb{E}_\gamma [\langle \mathbf{y}_t, \mathbf{x}_t \rangle] = \Phi_t(\mathbf{X}_{t+1}) - \Phi_t(\mathbf{X}_t) - \frac{1}{2} \langle \mathbf{x}_t, \nabla^2 \Phi_t(\widetilde{\mathbf{X}}_t) \mathbf{x}_t \rangle$$

where  $\widetilde{\mathbf{X}}_t = \mathbf{X}_t + \theta_t \mathbf{x}_t$ , for some  $\theta_t \in [0, 1]$ . Adding this from time  $t' + 1$  till  $T$  gives us:

$$\begin{aligned} &\sum_{t=t'+1}^T \mathbb{E}_\gamma [\langle \mathbf{y}_t, \mathbf{x}_t \rangle] \\ &= \sum_{t=t'+1}^T [\Phi_t(\mathbf{X}_{t+1}) - \Phi_t(\mathbf{X}_t)] - \frac{1}{2} \sum_{t=t'+1}^T \langle \mathbf{x}_t, \nabla^2 \Phi_t(\widetilde{\mathbf{X}}_t) \mathbf{x}_t \rangle \\ &= \Phi_T(\mathbf{X}_{T+1}) - \Phi_{t'+1}(\mathbf{X}_{t'+1}) + \sum_{t=t'+2}^T [\Phi_{t-1}(\mathbf{X}_t) - \Phi_t(\mathbf{X}_t)] \\ &\quad - \frac{1}{2} \sum_{t=t'+1}^T \langle \mathbf{x}_t, \nabla^2 \Phi_t(\widetilde{\mathbf{X}}_t) \mathbf{x}_t \rangle. \end{aligned}$$

Substituting this in (7) gives us:

$$\begin{aligned} R_A^{\text{W-FTPL}(\alpha\sqrt{t})}(T) &\leq t' + \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} - \mathbf{X}_{t'+1} \rangle - \Phi_T(\mathbf{X}_{T+1}) + \Phi_{t'+1}(\mathbf{X}_{t'+1}) \\ &\quad + \sum_{t=t'+1}^{T-1} [\Phi_{t+1}(\mathbf{X}_{t+1}) - \Phi_t(\mathbf{X}_{t+1})] + \frac{1}{2} \sum_{t=t'+1}^T \langle \mathbf{x}_t, \nabla^2 \Phi_t(\widetilde{\mathbf{X}}_t) \mathbf{x}_t \rangle. \end{aligned} \quad (8)$$

Using Jensen's inequality, we have that

$$\begin{aligned}\Phi_T(\mathbf{X}_{T+1}) &= \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} + \eta_T \gamma \rangle \right] \\ &\geq \max_{\mathbf{y} \in \mathcal{Y}} \mathbb{E}_\gamma [\langle \mathbf{y}, \mathbf{X}_{T+1} + \eta_T \gamma \rangle] \\ &= \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} \rangle.\end{aligned}$$

Substituting the above in (8) gives us:

$$\begin{aligned}R_A^{\text{W-FTPL}(\alpha\sqrt{t})}(T) &\leq t' + \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} - \mathbf{X}_{t'+1} \rangle - \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} \rangle + \Phi_{t'+1}(\mathbf{X}_{t'+1}) \\ &\quad - \sum_{t=t'+2}^T [\Phi_{t-1}(\mathbf{X}_t) - \Phi_t(\mathbf{X}_t)] + \frac{1}{2} \sum_{t=t'+1}^T \langle \mathbf{x}_t, \nabla^2 \Phi_t(\tilde{\mathbf{X}}_t) \mathbf{x}_t \rangle \\ &\leq t' + \Phi_{t'+1}(\mathbf{X}_{t'+1}) \\ &\quad + \sum_{t=t'+1}^{T-1} [\Phi_{t+1}(\mathbf{X}_{t+1}) - \Phi_t(\mathbf{X}_{t+1})] + \frac{1}{2} \sum_{t=t'+1}^T \langle \mathbf{x}_t, \nabla^2 \Phi_t(\tilde{\mathbf{X}}_t) \mathbf{x}_t \rangle,\end{aligned}\tag{9}$$

as  $\max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} - \mathbf{X}_{t'+1} \rangle \leq \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} \rangle$ . We also have:

$$\begin{aligned}\Phi_{t'+1}(\mathbf{X}_{t'+1}) &= \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{t'+1} + \eta_{t'+1} \gamma \rangle \right] \\ &\leq \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{t'+1} \rangle \right] + \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \eta_{t'+1} \gamma \rangle \right] \\ &\leq t' + \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \eta_{t'+1} \gamma \rangle \right],\end{aligned}$$

as  $\langle \mathbf{y}, \mathbf{X}_{t'+1} \rangle \leq t'$  for  $\mathbf{y} \in \mathcal{Y}$ . Substituting this back in (9), we get

$$\begin{aligned}R_A^{\text{W-FTPL}(\alpha\sqrt{t})}(T) &\leq 2t' + \eta_{t'+1} \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \gamma \rangle \right] \\ &\quad + \sum_{t=t'+1}^{T-1} [\Phi_{t+1}(\mathbf{X}_{t+1}) - \Phi_t(\mathbf{X}_{t+1})] + \frac{1}{2} \sum_{t=t'+1}^T \langle \mathbf{x}_t, \nabla^2 \Phi_t(\tilde{\mathbf{X}}_t) \mathbf{x}_t \rangle.\end{aligned}$$

As shown in [18],

$$\Phi_{t+1}(\mathbf{X}_{t+1}) - \Phi_t(\mathbf{X}_{t+1}) \leq |\eta_{t+1} - \eta_t| \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \gamma \rangle \right].$$

Also, [10] proves that:

$$\mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \gamma \rangle \right] \leq C \sqrt{2 \log \left( \frac{N}{C} \right)}.$$

It has also been proved in [7] that

$$\sum_{t=1}^T \langle \mathbf{x}_t, \nabla^2 \Phi_t(\tilde{\mathbf{X}}_t) \mathbf{x}_t \rangle \leq \sqrt{\frac{2}{\pi}} \sum_{t=1}^T \frac{1}{\eta_t}.$$

Combining all these results, we get

$$R_A^{\text{W-FTPL}(\alpha\sqrt{t})}(T) \leq 2t' + \eta_{t'+1} C \sqrt{2 \log\left(\frac{N}{C}\right)} + \eta_T C \sqrt{2 \ln(Ne/C)} + \sqrt{\frac{2}{\pi}} \sum_{t=1}^T \frac{1}{\eta_t}.$$

As  $\text{W-FTPL}(\alpha\sqrt{t})$  does not incur any switch cost till  $t'$  and then incurs the same switch cost as  $\text{FTPL}(\alpha\sqrt{t})$ , the total switch cost till time  $T$  incurred by  $\text{W-FTPL}(\alpha\sqrt{t})$  can be bounded from above by the switch cost of  $\text{FTPL}(\alpha\sqrt{t})$  proved in Proposition 4.2 of [18], which has been reproduced below:

$$\begin{aligned} \sum_{t=2}^T \mathbb{E} [\|\mathbf{y}_{t+1} - \mathbf{y}_t\|_1] &\leq \frac{3\sqrt{2}}{\alpha\sqrt{\pi}} (\sqrt{T} - 1) + (N-1) \frac{2 + \sqrt{2e \ln(2N)}}{\sqrt{e}} \ln T \\ &\quad + \frac{3(N-1)(2 + \sqrt{2e \ln(2N)})}{\sqrt{2\pi e} \alpha} (1 - T^{-1/2}). \end{aligned}$$

Thus, we get an  $\mathcal{O}(\sqrt{T})$  upper bound for  $R_A^{\text{W-FTPL}(\alpha\sqrt{t})}(T, D)$ .

## B Proof of part (b) of Theorem 2

In this section, we prove a lower bound on the regret of FTPL with a constant learning rate under stochastic file requests. For any file  $k$ , let  $\hat{\mu}_k(t) = \frac{c_k(t) + \eta_t \gamma(k)}{t}$  and let  $\alpha_k(t)$  denote the empirical average number of requests received by file  $k$ ,

where  $c_k(t)$  denotes the number of requests received by file  $k$  at time  $t$ .

$$\begin{aligned}
\mathbb{E}[R(T)] &= \mathbb{E} \left[ \sum_{t=1}^T \mathbb{1}\{x(t) \in \mathcal{C}\} - \mathbb{1}\{x(t) \in C(t)\} \right] \\
&\stackrel{(a)}{=} \mathbb{E} \left[ \sum_{t=1}^T \left( \sum_{j \in \mathcal{C}} \mu_j - \sum_{k \in C(t)} \mu_k \right) \right] \\
&\geq \mathbb{E} \left[ \sum_{t=1}^T \sum_{k \in \mathcal{C}(t) \setminus \mathcal{C}} (\mu_{\mathcal{C}} - \mu_k) \right] \\
&= \mathbb{E} \left[ \sum_{t=1}^T \sum_{k \in \mathcal{L} \setminus \mathcal{C}} (\mu_{\mathcal{C}} - \mu_k) \mathbb{1}\{k \in C(t)\} \right] \\
&= \sum_{k \in \mathcal{L} \setminus \mathcal{C}} (\mu_{\mathcal{C}} - \mu_k) \mathbb{E} \left[ \sum_{t=1}^T \mathbb{1}\{k \in C(t)\} \right] \\
&= \sum_{k \in \mathcal{L} \setminus \mathcal{C}} (\mu_{\mathcal{C}} - \mu_k) \sum_{t=1}^T \mathbb{P}(k \in C(t)) \\
&\geq (\mu_{\mathcal{C}} - \mu_{C+1}) \sum_{t=1}^T \sum_{k \in \mathcal{L} \setminus \mathcal{C}} \mathbb{P}(k \in C(t)) \\
&\geq (\mu_{\mathcal{C}} - \mu_{C+1}) \sum_{t=1}^T \mathbb{P} \left( \bigcup_{k \in \mathcal{L} \setminus \mathcal{C}} k \in C(t) \right), \tag{10}
\end{aligned}$$

where (a) follows from the tower property of conditional expectation (i.e., condition on  $C(t)$  inside the outer expectation). The event  $\left\{ \bigcup_{k \in \mathcal{L} \setminus \mathcal{C}} k \in C(t) \right\}$  corresponds to the event where file 2 belongs to  $C(t)$ , which is equivalent to saying that file 2 has a perturbed count greater than that of file 1 at time  $t$ . Thus,

$$\begin{aligned}
\mathbb{P} \left( \bigcup_{k \in \mathcal{L} \setminus \mathcal{C}} k \in C(t) \right) &= \mathbb{P}(\hat{\mu}_2(t) > \hat{\mu}_1(t)) \\
&= \mathbb{P} \left( (\alpha_{C+1}(t) - \alpha_{\mathcal{C}}(t)) - (\mu_{C+1} - \mu_{\mathcal{C}}) \right. \\
&\quad \left. + \frac{\eta}{t} (\gamma_{C+1}(t) - \gamma_{\mathcal{C}}(t)) > \Delta_{\min} \right).
\end{aligned}$$

Thus, we get

$$\begin{aligned} \mathbb{P}\left(\bigcup_{k \in \mathcal{L} \setminus \mathcal{C}} k \in C(t)\right) &\geq \mathbb{P}\left(\alpha_2(t) - \alpha_1(t) - (\mu_2 - \mu_1) > -\frac{2}{t}\right) \\ &\quad \times \mathbb{P}\left(\frac{\eta}{t}(\gamma_2(t) - \gamma_1(t)) > \Delta_{\min} + \frac{2}{t}\right), \end{aligned} \quad (11)$$

as the perturbation is independent of the file requests seen so far. We also have,

$$\begin{aligned} \mathbb{P}\left(\frac{\eta}{t}(\gamma_2(t) - \gamma_1(t)) > \Delta_{\min} + \frac{2}{t}\right) &= \mathbb{P}\left(\mathcal{N}(0, 2\eta^2) > t\Delta_{\min} + 2\right) \\ &\geq \frac{1}{4}e^{-(t\Delta_{\min}+2)^2/4\eta^2}. \end{aligned} \quad (12)$$

Also,

$$1 - \alpha_2(t) = \alpha_1(t). \quad (13)$$

Using (12) and (13) in (11),

$$\mathbb{P}\left(\bigcup_{k \in \mathcal{L} \setminus \mathcal{C}} k \in C(t)\right) \geq \frac{1}{4}\mathbb{P}\left(\alpha_2(t) > \frac{1 - \Delta_{\min}}{2} - \frac{1}{t}\right) e^{-(t\Delta_{\min}+2)^2/4\eta^2}, \quad (14)$$

where the last step follows from  $1 - \mu_1 = \mu_2$ . The paper [14] shows that any median  $m$  of a Binomial( $n, p$ ) distribution lies in the interval  $[\lfloor np \rfloor, \lceil np \rceil]$ . The first term in (14) is the probability of a Binomial( $t, \mu_2$ ) random variable exceeding  $t\mu_2 - 1 \leq \lfloor t\mu_2 \rfloor$ . Thus, we get that

$$\mathbb{P}\left(\bigcup_{k \in \mathcal{L} \setminus \mathcal{C}} k \in C(t)\right) \geq \frac{1}{8}e^{-(t\Delta_{\min}+2)^2/4\eta^2}.$$

Adding this over all  $T$  gives us:

$$\begin{aligned} \mathbb{E}[R(T)] &\geq \frac{\Delta_{\min}}{8} \sum_{t=1}^{\frac{2\eta}{\Delta_{\min}}} e^{-(t\Delta_{\min}+2)^2/4\eta^2} \\ &\geq \frac{\eta e^{-(\frac{1+\eta}{\eta})^2}}{4}. \end{aligned}$$

## C Proof of part (c) of Theorem 2

In this section, we prove an upper bound on the regret (including switching cost) of FTPL( $\alpha\sqrt{t}$ ) under stochastic arrivals. We assume  $L \geq 3$ . To prevent the

regret from scaling as a polynomial in  $L$ , we use the following idea from [17]: The algorithm's regret is upper bounded by  $t_0$  in the first  $t_0$  rounds and in the other  $T - t_0$  rounds, the regret is bounded using standard concentration inequalities. The switching cost is upper bounded by  $DCt_0$  for the first  $t_0$  requests and in the other  $T - t_0$  rounds, the switching cost is bounded using standard concentration inequalities. Recall that  $t_0 = \max \left\{ \frac{8}{\Delta_{\min}^2} \log(L^3), \frac{32\alpha^2}{\Delta_{\min}^2} \log(L^3) \right\}$ .

For any file  $k$ , at time  $t$ , let  $\alpha_k(t)$  denote the empirical average number of requests received by file  $k$ . We also define, for any file  $k$  after  $t$  files have been requested from the cache and for  $2 \leq i \leq s$ ,

$$\hat{\mu}_k(t) := \frac{c_k(t) + \eta_{t+1}\gamma(k)}{t}, \quad (15)$$

where  $c_k(t)$  denotes the number of requests received by file  $k$  at time  $t$ . We first consider the switching cost incurred by this policy. Any file can be fetched into the cache as the learning rate varies with time and as the number of switches is twice the number of fetches, the following equation from [18] holds:

$$\mathbb{E} [\|\mathbf{y}_{t+1} - \mathbf{y}_t\|_1] = 2 \sum_{f=1}^L \mathbb{P}(\text{The file index } f \text{ is fetched at time } t+1).$$

The probability of a file  $f$  being fetched at time  $t+1$  is evaluated by taking the following cases:

1.  $f \in \mathcal{C}$ :

Here, if file  $f$  is fetched at time  $t+1$ , then  $f \notin C(t)$  which implies that at time  $t$ ,  $\exists f' \in \mathcal{L} \setminus \mathcal{C}$  such that  $f' \in C(t)$ . This event can be upper bounded using a union bound over files in  $\mathcal{L} \setminus \mathcal{C}$  for  $f'$ .

2.  $f \notin \mathcal{C}$ :

Here, if file  $f$  is fetched at time  $t+1$ , then  $f \in C(t+1)$  which implies that at time  $t+1$ ,  $\exists f' \in \mathcal{C}$  such that  $f' \notin C(t+1)$ . This event can be upper bounded using a union bound over files in  $\mathcal{C}$  for  $f'$ .

Therefore,

$$\begin{aligned} \sum_{f=1}^L \mathbb{P}(\text{The file index } f \text{ is fetched at time } t+1) &\leq \sum_{j=1}^C \sum_{k=C+1}^L \mathbb{P}(\hat{\mu}_k(t-1) > \hat{\mu}_j(t-1)) \\ &\quad + \sum_{k=C+1}^L \sum_{j=1}^C \mathbb{P}(\hat{\mu}_k(t) > \hat{\mu}_j(t)). \end{aligned}$$

Thus the expected number of switches till  $T$  can be bounded as:

$$\begin{aligned}
\frac{D}{2} \sum_{t=2}^T \mathbb{E} [\|\mathbf{y}_{t+1} - \mathbf{y}_t\|_1] &\leq DCt_0 + \frac{D}{2} \sum_{t=t_0+1}^{T-1} \sum_{j=1}^C \sum_{k=C+1}^L \mathbb{P}(\hat{\mu}_k(t-1) > \hat{\mu}_j(t-1)) \\
&+ \frac{D}{2} \sum_{t=t_0+1}^{T-1} \sum_{j=1}^C \sum_{k=C+1}^L \mathbb{P}(\hat{\mu}_k(t) > \hat{\mu}_j(t)) \\
&\leq DCt_0 + D \sum_{t=t_0+1}^{T-1} \sum_{j=1}^C \sum_{k=C+1}^L \mathbb{P}(\hat{\mu}_k(t) > \hat{\mu}_j(t)) \\
&\leq DCt_0 + D \sum_{t=t_0+1}^{T-1} \sum_{j=1}^C \sum_{k=C+1}^L \mathbb{P}(\hat{\mu}_j(t) - \mu_j \leq -\Delta_{j,k}/2) \\
&+ \mathbb{P}(\hat{\mu}_k(t) - \mu_k > \Delta_{j,k}/2) \\
&\leq DCt_0 + D \sum_{t=t_0+1}^T \sum_{j=1}^C \sum_{k=C+1}^L \mathbb{P}(\alpha_j(t) - \mu_j \leq -\Delta_{j,k}/4) \\
&+ \mathbb{P}(\alpha_k(t) - \mu_k > \Delta_{j,k}/4) \\
&+ \mathbb{P}(\eta_t \gamma(j) \leq -t\Delta_{j,k}/4) + \mathbb{P}(\eta_t \gamma(k) > t\Delta_{j,k}/4),
\end{aligned}$$

where the last two steps follow from a union bounding argument. Now, using the Hoeffding inequality [13],

$$\begin{aligned}
\mathbb{P}(\alpha_j(t) - \mu_j \leq -\Delta_{j,k}/4) &\leq e^{-t\Delta_{j,k}^2/8}, \\
\mathbb{P}(\alpha_k(t) - \mu_k > \Delta_{j,k}/4) &\leq e^{-t^2\Delta_{j,k}^2/8}.
\end{aligned}$$

We also have:

$$\begin{aligned}
\mathbb{P}(\eta_t \gamma(j) \leq -t\Delta_{j,k}/4) &= \mathbb{P}(\eta_t \gamma(k) > t\Delta_{j,k}/4) \\
&\leq e^{-t^2\Delta_{j,k}^2/32\eta_t^2} = e^{-t\Delta_{j,k}^2/32\alpha^2}.
\end{aligned}$$

Thus, we get the following bound on the switching cost:

$$\sum_{t=1}^{T-1} \mathbb{E} [\|\mathbf{y}_{t+1} - \mathbf{y}_t\|_1] \leq DCt_0 + 2D \sum_{t=t_0+1}^{T-1} \sum_{j=1}^C \sum_{k=C+1}^L e^{-t\Delta_{j,k}^2/8} + e^{-t\Delta_{j,k}^2/32\alpha^2}. \quad (16)$$

The regret for the first  $t_0$  rounds is bounded by  $t_0$ . For  $t > t_0$ , we use ideas from [8] and [17] to upper bound the regret.

$$\begin{aligned}
 R_S^{\text{FTPL}(\alpha\sqrt{t})}(T) &= \mathbb{E} \left[ \sum_{t=1}^T \mathbb{1}\{x(t) \in \mathcal{C}\} - \mathbb{1}\{x(t) \in C(t)\} \right] \\
 &= \mathbb{E} \left[ \sum_{t=1}^T \left( \sum_{j \in \mathcal{C}} \mu_j - \sum_{k \in C(t)} \mu_k \right) \right] \\
 &\leq t_0 + \mathbb{E} \left[ \sum_{t=t_0+1}^T \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \mathbb{1}\{j \notin C(t), k \in C(t)\} \right] \\
 &\leq t_0 + \mathbb{E} \left[ \sum_{t=t_0+1}^T \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \mathbb{1}(\hat{\mu}_k(t) > \hat{\mu}_j(t)) \right] \\
 &\leq t_0 + \mathbb{E} \left[ \sum_{t=t_0+1}^T \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \left( \mathbb{1}\{\hat{\mu}_j(t) - \mu_j \leq -\Delta_{j,k}/2\} \right. \right. \\
 &\quad \left. \left. + \mathbb{1}\{\hat{\mu}_k(t) - \mu_k > \Delta_{j,k}/2\} \right) \right] \\
 &\leq t_0 + \mathbb{E} \left[ \sum_{t=t_0+1}^T \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \left( \mathbb{1}\{\alpha_j(t) - \mu_j \leq -\Delta_{j,k}/4\} \right. \right. \\
 &\quad \left. \left. + \mathbb{1}\{\eta_t \gamma(j) \leq -t\Delta_{j,k}/4\} \right) \right] \\
 &\quad + \mathbb{E} \left[ \sum_{t=t_0+1}^T \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \left( \mathbb{1}\{\alpha_k(t) - \mu_k > \Delta_{j,k}/4\} \right. \right. \\
 &\quad \left. \left. + \mathbb{1}\{\eta_t \gamma(k) > t\Delta_{j,k}/4\} \right) \right]. \tag{17}
 \end{aligned}$$

By taking expectation inside the summation, we obtain the following upper bound for (17):

$$R_S^{\text{FTPL}(\alpha\sqrt{t})}(T) \leq t_0 + \sum_{j=1}^C \sum_{k=C+1}^L \sum_{t=t_0+1}^T 2\Delta_{j,k} \left( e^{-t\Delta_{j,k}^2/8} + e^{-t^2\Delta_{j,k}^2/32\eta^2} \right) \tag{18}$$

Thus, combining (16) and (18) , we have:

$$\begin{aligned}
R_S^{\text{FTPL}(\alpha\sqrt{t})}(T, D) &\leq (1 + DC)t_0 \\
&\quad + 2D \sum_{t=t_0+1}^T \sum_{j=1}^C \sum_{k=C+1}^L \left( e^{-t\Delta_{j,k}^2/8} \right. \\
&\quad \left. + e^{-t\Delta_{j,k}^2/32\alpha^2} \right) \\
&\quad + 2 \sum_{t=t_0+1}^T \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \left( e^{-t\Delta_{j,k}^2/8} \right. \\
&\quad \left. + e^{-t\Delta_{j,k}^2/32\alpha^2} \right). \tag{19}
\end{aligned}$$

Each of these terms are bounded separately in the following way:

$$\begin{aligned}
\sum_{t=t_0+1}^T \sum_{j=1}^C \sum_{k=C+1}^L e^{-t\Delta_{j,k}^2/8} &\leq \sum_{t=t_0+1}^T C(L-C)e^{-t\Delta_{\min}^2/8} \\
&\leq C(L-C)e^{-t_0\Delta_{\min}^2/8} \sum_{t=t_0+1}^T e^{-(t-t_0)\Delta_{\min}^2/8} \\
&\leq \sum_{t=1}^{T-t_0} e^{-t\Delta_{\min}^2/8} \leq \frac{8}{\Delta_{\min}^2}.
\end{aligned}$$

Similarly,

$$\begin{aligned}
\sum_{t=t_0+1}^T \sum_{j=1}^C \sum_{k=C+1}^L e^{-t\Delta_{j,k}^2/32\alpha^2} &\leq \sum_{t=t_0+1}^T C(L-C)e^{-t\Delta_{\min}^2/32\alpha^2} \\
&\leq L^2 e^{-t_0\Delta_{\min}^2/32\alpha^2} \sum_{t=t_0+1}^T e^{-(t-t_0)\Delta_{\min}^2/32\alpha^2} \\
&\leq \sum_{t=1}^{T-t_0} e^{-t\Delta_{\min}^2/32\alpha^2} \leq \frac{32\alpha^2}{\Delta_{\min}^2}.
\end{aligned}$$

The function  $f(u) = ue^{-u^2/2}$  is decreasing on  $[1, +\infty)$ . Since  $\Delta_{j,k} \geq \Delta_{\min}$ ,  $j \in \mathcal{C}, k \notin \mathcal{C}$ , we get

$$\begin{aligned}
\Delta_{j,k} e^{-t\Delta_{j,k}^2/8} &= \frac{2}{\sqrt{t}} f\left(\frac{\sqrt{t}\Delta_{j,k}}{2}\right) \\
&\leq \frac{2}{\sqrt{t}} f\left(\frac{\sqrt{t}\Delta_{\min}}{2}\right) = \Delta_{\min} e^{-t\Delta_{\min}^2/8}.
\end{aligned}$$

Note that for  $t \geq t_0 + 1$ ,  $t > \frac{4}{\Delta_{\min}^2}$  holds as  $t_0 \geq \frac{8}{\Delta_{\min}^2}$ . Using this, we have that:

$$\sum_{t=t_0+1}^T \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} e^{-t\Delta_{j,k}^2/8} \leq \sum_{t=t_0+1}^T C(L-C)\Delta_{\min} e^{-t\Delta_{\min}^2/8} \leq \frac{8}{\Delta_{\min}}.$$

Similarly,

$$\sum_{t=t_0+1}^T \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} e^{-t\Delta_{j,k}^2/32\alpha^2} \leq \frac{32\alpha^2}{\Delta_{\min}}.$$

Substituting these bounds in (18), we have the following upper bound on the regret of FTPL( $\alpha\sqrt{t}$ ):

$$R_S^{\text{FTPL}(\alpha\sqrt{t})}(T, D) \leq (1 + DC) t_0 + 2 \left( 1 + \frac{D}{\Delta_{\min}} \right) \left( \frac{8}{\Delta_{\min}} + \frac{32\alpha^2}{\Delta_{\min}} \right).$$

## D Proof of part (d) of Theorem 2

In this section, we prove an upper bound on the regret (including the switching cost) of the W-FTPL algorithm under stochastic file requests. Using (16),

$$\begin{aligned} \frac{D}{2} \sum_{t=t'+1}^T \mathbb{E} [\|\mathbf{y}_{t+1} - \mathbf{y}_t\|_1] &\leq 2D \sum_{t=t'+1}^T \sum_{j=1}^C \sum_{k=C+1}^L \left( e^{-t\Delta_{j,k}^2/8} + e^{-t\Delta_{j,k}^2/32\alpha^2} \right) \\ &\leq 2DC(L-C) \sum_{t=t'+1}^T \left( e^{-t\Delta_{\min}^2/8} + e^{-t\Delta_{\min}^2/32\alpha^2} \right) \\ &\leq 2DC(L-C) \sum_{t=t'+1}^T \left( e^{-t'\Delta_{\min}^2/8} e^{-(t-t')\Delta_{\min}^2/8} \right. \\ &\quad \left. + e^{-t'\Delta_{\min}^2/32\alpha^2} e^{-(t-t')\Delta_{\min}^2/32\alpha^2} \right) \\ &\leq 2DC(L-C) \left( e^{-t'\frac{\Delta_{\min}^2}{8}} \frac{8}{\Delta_{\min}^2} + e^{-t'\frac{\Delta_{\min}^2}{32\alpha^2}} \frac{32\alpha^2}{\Delta_{\min}^2} \right) \\ &\leq 2DC(L-C) \left( e^{-u(\log D)^{1+\beta}} \Delta_{\min}^2/8 \frac{8}{\Delta_{\min}^2} \right. \\ &\quad \left. + e^{-u(\log D)^{1+\beta}} \Delta_{\min}^2/32\alpha^2 \frac{32\alpha^2}{\Delta_{\min}^2} \right). \end{aligned} \quad (20)$$

Till  $t'$ , the regret is bounded by  $t'$ . Thus, the regret can be bounded in the following way using (18) :

$$\begin{aligned} \mathbb{E} \left[ \sum_{t=1}^T \mathbb{1}\{x(t) \in \mathcal{C}\} - \mathbb{1}\{x(t) \in C(t)\} \right] &\leq t' + 2 \sum_{t=t'+1}^T \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \left( e^{-t\Delta_{j,k}^2/8} \right. \\ &\quad \left. + e^{-t\Delta_{j,k}^2/32\alpha^2} \right) \\ &\leq t' + 2 \left( \frac{8}{\Delta_{\min}} + \frac{32\alpha^2}{\Delta_{\min}} \right). \end{aligned} \quad (21)$$

Combining (20) and (21) gives the following upper bound on the regret of W-FTPL( $\alpha\sqrt{t}$ ):

$$\begin{aligned} R_S^{\text{W-FTPL}(\alpha\sqrt{t})}(T, D) &\leq t' + \frac{16}{\Delta_{\min}} + \frac{64\alpha^2}{\Delta_{\min}} + 2LDC(L-C) \left( e^{-u(\log D)^{1+\beta} \Delta_{\min}^2/8} \frac{8}{\Delta_{\min}^2} \right. \\ &\quad \left. + e^{-u(\log D)^{1+\beta} \Delta_{\min}^2/32\alpha^2} \frac{32\alpha^2}{\Delta_{\min}^2} \right). \end{aligned}$$

## E Proof of part (a) of Theorem 3

In this section, we prove a lower bound on the regret of any policy  $\pi$  under stochastic file requests when cache updates are restricted to  $s+1$  fixed points. To provide a lower bound on the regret, we first prove a lower bound on the

regret for the initial set of  $r_1$  requests.

$$\begin{aligned}
 R_S^\pi(r_1) &\geq \mathbb{E} \left[ \sum_{t=1}^{r_1} \mathbb{1}\{x(t) \in \mathcal{C}\} - \mathbb{1}\{x(t) \in C(t)\} \right] \\
 &\stackrel{(a)}{=} \mathbb{E} \left[ \sum_{t=1}^{r_1} \left( \sum_{j \in \mathcal{C}} \mu_j - \sum_{k \in C(t)} \mu_k \right) \right] \\
 &\geq \mathbb{E} \left[ \sum_{t=1}^{r_1} \sum_{k \in \mathcal{L} \setminus \mathcal{C}} (\mu_C - \mu_k) \right] \\
 &= \mathbb{E} \left[ \sum_{t=1}^{r_1} \sum_{k \in \mathcal{L} \setminus \mathcal{C}} (\mu_C - \mu_k) \mathbb{1}\{k \in C(t)\} \right] \\
 &= \sum_{k \in \mathcal{L} \setminus \mathcal{C}} (\mu_C - \mu_k) \mathbb{E} \left[ \sum_{t=1}^{r_1} \mathbb{1}\{k \in C(t)\} \right] \\
 &= \sum_{k \in \mathcal{L} \setminus \mathcal{C}} (\mu_C - \mu_k) \sum_{t=1}^{r_1} \mathbb{P}(k \in C(t)) \\
 &\geq (\mu_C - \mu_{C+1}) \sum_{t=1}^{r_1} \sum_{k \in \mathcal{L} \setminus \mathcal{C}} \mathbb{P}(k \in C(t)), \tag{22}
 \end{aligned}$$

where (a) follows from the tower property of conditional expectation (i.e., condition on  $C(t)$  inside the outer expectation). Note that  $C(t)$  remains constant during the first  $r_1$  requests and consists of  $C$  files chosen randomly from the library of  $L$  files. Thus, for any  $k \in \mathcal{L} \setminus \mathcal{C}$  and  $1 \leq t \leq r_1$ ,

$$\mathbb{P}(k \in C(t)) = \frac{\binom{N-1}{C-1}}{\binom{N}{C}} = \frac{C}{N}.$$

Using this in (22), we get that

$$R_S^\pi(r_1) \geq r_1(\mu_C - \mu_{C+1})(L - C)C/N.$$

For proving a lower bound on the regret for larger values of  $T$ , we consider  $L = 2, C = 1$ . Given a popularity distribution  $\mu = (\mu_1, \mu_2)$  such that  $\mu_1 > \mu_2$  and a policy  $\pi$ ,

$$\begin{aligned}
 R_{S,\mu}^\pi(T) &= \mathbb{E}_{\pi\mu} \left[ \sum_{t=1}^T \mathbb{1}\{x(t) \in \mathcal{C}\} - \mathbb{1}\{x(t) \in C(t)\} \right] \\
 &= \sum_{t=1}^T (\mu_1 - \mu_2) \mathbb{E}_{\pi\mu} [\mathbb{1}\{C(t) = \{2\}\}] \\
 &= (\mu_1 - \mu_2) \sum_{t=1}^T \mathbb{P}_{\pi\mu}(C(t) = \{2\}).
 \end{aligned}$$

Now, consider the popularity distribution  $\mu' = (\mu_2, \mu_1)$ . Following the same steps for this popularity distribution, we have that

$$R_{S,\mu'}^\pi(T) = (\mu_1 - \mu_2) \sum_{t=1}^T \mathbb{P}_{\pi\mu'}(C(t) = \{1\}).$$

Defining  $\Delta = \mu_1 - \mu_2$ , we have that

$$R_{S,\mu'}^\pi(T) + R_{S,\mu}^\pi(T) = \Delta \sum_{t=1}^T \mathbb{P}_{\pi\mu}(C(t) = \{2\}) + \Delta \sum_{t=1}^T \mathbb{P}_{\pi\mu'}(C(t) = \{1\}).$$

Using the fact that the cache configuration remains constant except at the places where the cache is allowed to change its contents, we have that

$$\begin{aligned} R_{S,\mu'}^\pi(T) + R_{S,\mu}^\pi(T) &\geq R_{S,\mu'}^\pi(r_1) + R_{S,\mu}^\pi(r_1) \\ &+ \Delta \sum_{i=2}^s r_i \left( \mathbb{P}_{\pi\mu}(C(t_i+1) = \{2\}) + \mathbb{P}_{\pi\mu'}(C(t_i+1) = \{1\}) \right), \end{aligned} \quad (23)$$

where  $t_i$  has been defined in (25).

**Lemma 1** ([15], **Theorem 14.2 (Bretagnolle-Huber inequality)**). *Let  $P$  and  $Q$  be probability distributions on the same measurable space  $(\Omega, \mathcal{F})$  and let  $A \in \mathcal{F}$  be an arbitrary event. Then,*

$$P(A) + Q(A^c) \geq \frac{1}{2} \exp(-D(P, Q)).$$

Consider a fixed  $i$  such that  $2 \leq i \leq s$ . Setting  $A$  to be the event  $\{C(t_i+1) = \{2\}\}$  in Lemma 1, we have that

$$\mathbb{P}_{\pi\mu}(C(t_i+1) = \{2\}) + \mathbb{P}_{\pi\mu'}(C(t_i+1) = \{1\}) \geq \frac{1}{2} \exp(-D(\mathbb{P}_{\pi\mu}, \mathbb{P}_{\pi\mu'})).$$

**Lemma 2.** *Let  $\mathbb{C}$  be the set of valid caching configurations. If  $\mathbb{P}_{\pi\mu}$  and  $\mathbb{P}_{\pi\mu'}$  are probability measures on the set  $\mathcal{G}_i := \{[2]^{t_i} \times \mathbb{C}^{t_i+1}\}$ , then*

$$D(\mathbb{P}_{\pi\mu}, \mathbb{P}_{\pi\mu'}) = t_i D(\mathbb{P}_\mu, \mathbb{P}_{\mu'}),$$

where  $\mathbb{P}_\mu, \mathbb{P}_{\mu'}$  are the corresponding marginal distributions.

Using the above lemma, we have that

$$\begin{aligned} \mathbb{P}_{\pi\mu}(C(t_i+1) = \{2\}) + \mathbb{P}_{\pi\mu'}(C(t_i+1) = \{1\}) &\geq \frac{1}{2} \exp(-t_i D(\mathbb{P}_\mu, \mathbb{P}_{\mu'})) \\ &\geq \frac{1}{2} \exp\left(-t_i \frac{\Delta^2}{\mu_1 \mu_2}\right), \end{aligned}$$

where the last step follows from an upper bound on the KL Divergence between two Bernoulli distributions. Substituting this back in (23), we have that

$$R_{S,\mu'}^\pi(T) + R_{S,\mu}^\pi(T) \geq r_1\Delta + \sum_{i=2}^s r_i \frac{\Delta}{2} \exp\left(-t_i \frac{\Delta^2}{\mu_1\mu_2}\right).$$

Thus,

$$\max\{R_{S,\mu'}^\pi(T), R_{S,\mu}^\pi(T)\} \geq \frac{r_1\Delta}{2} + \sum_{i=2}^s r_i \frac{\Delta}{4} \exp\left(-t_i \frac{\Delta^2}{\mu_1\mu_2}\right),$$

which gives us the result.

*Proof (Proof of Lemma 2).* Let  $p_\mu(x) = \mathbb{P}_\mu(x(t) = x)$  be the density function associated with  $\mathbb{P}_\mu$  and let  $p_{\pi\mu}$  be the density function associated with  $\mathbb{P}_{\pi\mu}$ . We denote by  $\pi(\mathcal{C}_1|h)$  to be the probability of the caching policy choosing cache configuration  $\mathcal{C}_1$  given the history  $h$ . Then,

$$\begin{aligned} p_{\pi\mu}(C(1), x_1, C(2), x_2, \dots, C(t_i + 1)) \\ = \pi(C(1))p_\mu(x_1)\pi(C(2)|C(1), x_1) \cdots \pi(C(t_i + 1)|C(1), x_1, \dots, x_{t_i}). \end{aligned}$$

Thus,

$$\begin{aligned} D(\mathbb{P}_{\pi\mu}, \mathbb{P}_{\pi\mu'}) &= \sum_{g \in \mathcal{G}_i} p_{\pi\mu}(g) \log\left(\frac{p_{\pi\mu}(g)}{p_{\pi\mu'}(g)}\right) \\ &= \sum_{X^i=(x(1), \dots, x(t_i))} p_\mu(X^i) \log\left(\frac{p_\mu(X^i)}{p_{\mu'}(X^i)}\right) \\ &= t_i D(\mathbb{P}_\mu, \mathbb{P}_{\mu'}), \end{aligned}$$

where the second last step follows from the fact that the log likelihood does not depend upon the caching policy and the last step follows from the independence of the file requests.

## F Proof of part (b) of Theorem 3

In this section, we prove an upper bound on the regret of FTPL( $\alpha\sqrt{t}$ ) when cache updates are restricted to  $s + 1$  fixed time slots under adversarial requests. For any file  $k$ , at time  $t$ , let  $\alpha_k(t)$  denote the empirical average number of requests received by file  $k$ . We also define, for any file  $k$  after  $t$  files have been requested from the cache and for  $2 \leq i \leq s$ ,

$$\hat{\mu}_k(t) := \frac{c_k(t) + \eta_t \gamma(k)}{t}, \quad (24)$$

$$t_i := \sum_{j=1}^{i-1} r_j, \quad (25)$$

$$\eta^i := \eta_{t_i},$$

where  $c_k(t)$  denotes the number of requests received by file  $k$  at time  $t$ . Thus,

$$\begin{aligned}
R_S^{FTPL(\alpha\sqrt{t})}(T) &= \mathbb{E} \left[ \sum_{t=1}^T \mathbb{1}\{x(t) \in \mathcal{C}\} - \mathbb{1}\{x(t) \in C(t)\} \right] \\
&\stackrel{(a)}{\leq} r_1 + \mathbb{E} \left[ \sum_{t=r_1+1}^T \mathbb{1}\{x(t) \in \mathcal{C}\} - \mathbb{1}\{x(t) \in C(t)\} \right] \\
&= r_1 + \mathbb{E} \left[ \sum_{i=2}^s \sum_{t=t_i+1}^{t_{i+1}} \mathbb{1}\{x(t) \in \mathcal{C}\} - \mathbb{1}\{x(t) \in C(t)\} \right] \\
&= r_1 + \sum_{i=2}^s \sum_{t=t_i+1}^{t_{i+1}} \mathbb{E} [\mathbb{1}\{x(t) \in \mathcal{C}\} - \mathbb{1}\{x(t) \in C(t)\}] \\
&\stackrel{(b)}{=} r_1 + \sum_{i=2}^s r_i \mathbb{E} [\mathbb{1}\{x(t_i+1) \in \mathcal{C}\} - \mathbb{1}\{x(t_i+1) \in C(t_i+1)\}],
\end{aligned}$$

where (a) follows from bounding the regret for the first  $r_1$  time slots by  $r_1$  and (b) follows from the fact that  $C(t)$  remains the same from  $t_i+1$  to  $t_{i+1}$  for  $1 \leq i \leq s$ .

$$\begin{aligned}
R_S^{FTPL(\alpha\sqrt{t})}(T) &\stackrel{(c)}{\leq} r_1 + \sum_{i=2}^s r_i \mathbb{E} \left[ \left[ \left( \sum_{j \in \mathcal{C}} \mu_j - \sum_{k \in C(t_i+1)} \mu_k \right) \right] \right] \\
&\leq r_1 + \sum_{i=2}^s r_i \mathbb{E} \left[ \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \mathbb{1}\{j \notin C(t_i+1), k \in C(t_i+1)\} \right] \\
&\leq r_1 + \sum_{i=2}^s r_i \mathbb{E} \left[ \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \mathbb{1}\{\hat{\mu}_k(t_i) > \hat{\mu}_j(t_i)\} \right] \\
&\stackrel{(d)}{\leq} r_1 + \sum_{i=2}^s r_i \mathbb{E} \left[ \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \left( \mathbb{1}\{\hat{\mu}_j(t_i) - \mu_j \leq -\Delta_{j,k}/2\} \right. \right. \\
&\quad \left. \left. + \mathbb{1}\{\hat{\mu}_k(t_i) - \mu_k > \Delta_{j,k}/2\} \right) \right],
\end{aligned}$$

where (c) follows from the tower property of conditional expectation (i.e., condition on  $C(t)$  inside the outer expectation) and (d) follows from adding  $\Delta_{j,k}$  on both sides and using the fact that at least one of  $(\hat{\mu}_k(t_i) - \mu_k)$  and  $(\mu_j - \hat{\mu}_j(t_i))$

must be greater than  $\Delta_{j,k}/2$ . Using (24),

$$\begin{aligned}
 R_S^{FTPL(\alpha\sqrt{t})}(T) &\leq r_1 + \sum_{i=2}^s r_i \mathbb{E} \left[ \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \left( \mathbb{1} \{ \alpha_j(t_i) - \mu_j \leq -\Delta_{j,k}/4 \} \right. \right. \\
 &\quad \left. \left. + \mathbb{1} \{ \eta_{t_i} \gamma(j) \leq -t_i \Delta_{j,k}/4 \} \right) \right] \\
 &\quad + \sum_{i=2}^s r_i \mathbb{E} \left[ \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \left( \mathbb{1} \{ \alpha_k(t_i) - \mu_k > \Delta_{j,k}/4 \} \right. \right. \\
 &\quad \left. \left. + \mathbb{1} \{ \eta_{t_i} \gamma(k) > t_i \Delta_{j,k}/4 \} \right) \right], \tag{26}
 \end{aligned}$$

which follows from a union bounding argument similar to (d). Using Hoeffding's inequality ([13]), we obtain

$$\begin{aligned}
 \mathbb{P}(\alpha_j(t) - \mu_j \leq -\Delta_{j,k}/4) &\leq e^{-t\Delta_{j,k}^2/8}, \\
 \mathbb{P}(\alpha_k(t) - \mu_k > \Delta_{j,k}/4) &\leq e^{-t^2\Delta_{j,k}^2/8}.
 \end{aligned}$$

Also,

$$\mathbb{P}\left(\frac{\eta_t}{t}\gamma(j) \leq -\Delta_{j,k}/4\right) = \mathbb{P}\left(\frac{\eta_t}{t}\gamma(k) > \Delta_{j,k}/4\right) \leq e^{-t^2\Delta_{j,k}^2/32\eta_t^2}.$$

By taking the expectation inside the summation in (26), we obtain the following upper bound for the regret:

$$\begin{aligned}
 R_S^{FTPL(\alpha\sqrt{t})}(T) &\leq r_1 + 2 \sum_{j=1}^C \sum_{k=C+1}^L \sum_{i=2}^s r_i \Delta_{j,k} \left( e^{-t_i\Delta_{j,k}^2/8} + e^{-t_i\Delta_{j,k}^2/32\alpha^2} \right) \\
 &= r_1 + 2 \sum_{j=1}^C \sum_{k=C+1}^L \sum_{i=2}^s r_i \Delta_{j,k} \left( e^{-(\sum_{j=1}^{i-1} r_j)\Delta_{j,k}^2/8} + e^{-(\sum_{j=1}^{i-1} r_j)\Delta_{j,k}^2/32\alpha^2} \right) \\
 &\leq r_1 + 2 \max_{2 \leq a \leq s} r_a \sum_{j=1}^C \sum_{k=C+1}^L \sum_{i=2}^s \Delta_{j,k} \left( e^{-\left((i-1) \min_{1 \leq a \leq s} r_a\right)\Delta_{j,k}^2/8} + e^{-\left((i-1) \min_{1 \leq a \leq s} r_a\right)\Delta_{j,k}^2/32\alpha^2} \right) \\
 &\leq r_1 + 2 \frac{\max_{2 \leq a \leq s} r_a}{\min_{1 \leq a \leq s} r_a} \sum_{j=1}^C \sum_{k=C+1}^L \frac{8}{\Delta_{j,k}} + \frac{32\alpha^2}{\Delta_{j,k}} \\
 &\leq r_1 + 2C(L-C) \frac{\max_{2 \leq a \leq s} r_a}{\min_{1 \leq a \leq s} r_a} \left( \frac{8}{\Delta_{\min}} + \frac{32\alpha^2}{\Delta_{\min}} \right). \tag{27}
 \end{aligned}$$

## G Proof of Theorem 4

In this section, we prove an upper bound on the regret of FTPL( $\alpha\sqrt{t}$ ) when cache updates are restricted to periodic time slots. Recall that  $t'_0 = \max\left\{r, \frac{8}{\Delta_{\min}^2} \log(L^2), \frac{32\alpha^2}{\Delta_{\min}^2} \log(L^2)\right\}$  and we assume that  $L \geq 3$ . We bound the regret for the first  $t'_0$  rounds by  $t'_0$ . This technique of bounding the regret of an initial period by its worst-case regret and then using normal methods to bound the regret for  $t > t'_0$  is based on [17]. Thus, we get the following expression for the regret:

$$\begin{aligned} R_S^{\text{FTPL}(\alpha\sqrt{t})} &\leq [t'_0] + r\mathbb{E} \left[ \sum_{t=\lceil \frac{t'_0}{r} \rceil + 1}^{T/r} \mathbb{1}\{x(rt) \in \mathcal{C}\} - \mathbb{1}\{x(rt) \in C(rt)\} \right] \\ &\leq 1 + t'_0 + 2r \sum_{t=\lceil \frac{t'_0}{r} \rceil + 1}^{T/r} \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \left( e^{-rt\Delta_{j,k}^2/8} + e^{-rt\Delta_{j,k}^2/32\alpha^2} \right), \end{aligned}$$

which follows from (27). The function  $f(u) = ue^{-u^2/2}$  is decreasing on  $[1, +\infty)$ . Since  $\Delta_{j,k} \geq \Delta_{\min}, j \in \mathcal{C}, k \notin \mathcal{C}$ , we get

$$\begin{aligned} \Delta_{j,k} e^{-rt\Delta_{j,k}^2/8} &= \frac{2}{\sqrt{rt}} f\left(\frac{\sqrt{rt}\Delta_{j,k}}{2}\right) \\ &\leq \frac{2}{\sqrt{rt}} f\left(\frac{\sqrt{rt}\Delta_{\min}}{2}\right) = \Delta_{\min} e^{-t\Delta_{\min}^2/8}. \end{aligned}$$

Note that for  $t \geq \lceil \frac{t'_0}{r} \rceil + 1$ ,  $t > \frac{4}{r\Delta_{\min}^2}$  holds as  $t'_0 \geq \frac{8}{\Delta_{\min}^2}$ . Thus,

$$\begin{aligned} 2r \sum_{t=\lceil \frac{t'_0}{r} \rceil + 1}^{T/r} \sum_{j=1}^C \sum_{k=C+1}^L \Delta_{j,k} \left( e^{-rt\Delta_{j,k}^2/8} + e^{-rt\Delta_{j,k}^2/32\alpha^2} \right) \\ \leq 2rL^2 \Delta_{\min} \sum_{t=\lceil \frac{t'_0}{r} \rceil + 1}^{T/r} e^{-rt\Delta_{\min}^2/8} + e^{-rt\Delta_{\min}^2/32\alpha^2}. \end{aligned}$$

Now,

$$\begin{aligned}
L^2 \sum_{t=\lceil \frac{t'_0}{r} \rceil + 1}^{T/r} e^{-rt\Delta_{\min}^2/8} + e^{-rt\Delta_{\min}^2/32\alpha^2} &\leq L^2 \sum_{t=\lceil \frac{t'_0}{r} \rceil + 1}^{T/r} \left\{ e^{-r\left(\lceil \frac{t'_0}{r} \rceil + t - \lceil \frac{t'_0}{r} \rceil\right)\Delta_{\min}^2/8} \right. \\
&\quad \left. + e^{-r\left(\lceil \frac{t'_0}{r} \rceil + t - \lceil \frac{t'_0}{r} \rceil\right)\Delta_{\min}^2/32\alpha^2} \right\} \\
&\leq \sum_{t=1}^{T/r - \lceil \frac{t'_0}{r} \rceil} e^{-\frac{rt\Delta_{\min}^2}{8}} + e^{-\frac{rt\Delta_{\min}^2}{32\alpha^2}} \\
&\leq \frac{1}{r} \left( \frac{8}{\Delta_{\min}^2} + \frac{32\alpha^2}{\Delta_{\min}^2} \right).
\end{aligned}$$

This gives the following upper bound on the regret:

$$R_S^{\text{FTPL}(\alpha\sqrt{t})} \leq 1 + t'_0 + 2 \left( \frac{8}{\Delta_{\min}^2} + \frac{32\alpha^2}{\Delta_{\min}^2} \right).$$

## H Proof of part (a) of Theorem 5

In this section, we prove a lower bound on the regret of any policy  $\pi$  when cache updates are restricted under adversarial file requests. The proof uses a technique called the probabilistic method from [7] which was pioneered by Erdős [5]. Consider a request sequence  $\{x_t\}_{t=1}^T$  with a joint probability distribution defined on it. Then, we have the following lower bound on the regret incurred by any policy  $\pi$ :

$$R_A^\pi(T) \geq \mathbb{E}_{\{x_t\}_{t=1}^T} R_A^\pi(\{x_t\}_{t=1}^T, T). \quad (28)$$

Thus, the proof proceeds by defining a request sequence first and then using (28) to lower bound the regret. Throughout this proof, if it is not explicitly mentioned with respect to what the expectation is being taken, it can be assumed that the expectation is being taken with respect to  $\{x_t\}_{t=1}^T$ .

### H.1 Lower Bound when $C = 1$

We consider an adversarial request sequence where files are requested from the first 2 files uniformly, and the same file is requested  $r_i$  times from  $t_i = \sum_{j=0}^{i-1} r_j$  to  $t_i + r_i - 1$ , where  $r_0 = 1$  and  $1 \leq i \leq s$ . To be precise, for each  $1 \leq i \leq s$ , a file is drawn uniformly at random from the first two files at  $t = \sum_{j=0}^{i-1} r_j$  and is repeatedly requested till  $t = t_i + r_i - 1$ , i.e., a total of  $r_i$  times. Throughout this proof, “phase  $i$ ” refers to the time period from  $t = \sum_{j=0}^{i-1} r_j$  to  $t = t_i + r_i - 1$  (both points included). Let  $W_i$  be a Bernoulli random variable indicating whether file

1 was requested in the  $i^{\text{th}}$  phase or not. Note that  $\{W_i\}_{i=1}^s$  are i.i.d. Thus, the reward obtained by the optimal offline policy corresponds to:

$$\begin{aligned} \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} \rangle &= \max_{\mathbf{y} \in \mathcal{Y}} \left\langle \mathbf{y}, \sum_{t=1}^T \mathbf{x}_t \right\rangle \\ &= \max_{\mathbf{y} \in \mathcal{Y}} \left\langle \mathbf{y}, \sum_{i=1}^s \sum_{t=t_i}^{t_i+r_i-1} \mathbf{x}_t \right\rangle \\ &= \max_{\mathbf{y} \in \mathcal{Y}} \left\langle \mathbf{y}, \sum_{i=1}^s r_i \mathbf{x}_{t_i} \right\rangle, \end{aligned}$$

as the specific request sequence chosen here is constant in each phase. As  $C = 1$ ,  $\mathbf{y}$  is a one-hot encoded vector and thus, we have that

$$\max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} \rangle = \max \left\{ \sum_{i=1}^s r_i W_i, T - \sum_{i=1}^s r_i W_i \right\} = \frac{T}{2} + \left| \frac{T}{2} - \sum_{i=1}^s r_i W_i \right|.$$

The reward of any algorithm can be upper bounded in the following way:

$$\begin{aligned} \sum_{t=1}^T \langle \mathbf{y}_t, \mathbf{x}_t \rangle &= \sum_{i=1}^s \sum_{t=t_i}^{t_i+r_i-1} \langle \mathbf{y}_t, \mathbf{x}_t \rangle \\ &= \sum_{i=1}^s r_i \langle \mathbf{y}_{t_i}, \mathbf{x}_{t_i} \rangle, \end{aligned}$$

as the specific request sequence chosen here and the cache configuration is constant in each phase. Thus, the expected reward of any algorithm can be upper bounded in the following way, where the expectation is taken with respect to the request sequence only:

$$\begin{aligned} \mathbb{E}_{\{x_t\}_{t=1}^T} \left[ \sum_{t=1}^T \langle \mathbf{y}_t, \mathbf{x}_t \rangle \right] &= \mathbb{E}_{\{x_t\}_{t=1}^T} \left[ \sum_{i=1}^s r_i \langle \mathbf{y}_{t_i}, \mathbf{x}_{t_i} \rangle \right] \\ &= \sum_{i=1}^s r_i \mathbb{E}_{\{x_t\}_{t=1}^T} [\langle \mathbf{y}_{t_i}, \mathbf{x}_{t_i} \rangle]. \end{aligned}$$

In each time slot, the cache update happens before the file request arrives and hence  $\mathbf{y}_t$  is independent of  $\mathbf{x}_t$ ,  $1 \leq t \leq T$ . Thus, we have

$$\begin{aligned}
 \mathbb{E}_{\{x_t\}_{t=1}^T} \left[ \sum_{t=1}^T \langle \mathbf{y}_t, \mathbf{x}_t \rangle \right] &= \sum_{i=1}^s r_i \mathbb{E}_{\{x_t\}_{t=1}^T} [\langle \mathbf{y}_{t_i}, \mathbf{x}_{t_i} \rangle] \\
 &= \sum_{i=1}^s r_i \left\langle \mathbf{y}_{t_i}, \mathbb{E}_{\{x_t\}_{t=1}^T} [\mathbf{x}_{t_i}] \right\rangle \\
 &= \frac{1}{2C} \sum_{i=1}^s r_i (\mathbf{y}_{t_i}(1) + \mathbf{y}_{t_i}(2)) \\
 &\leq \frac{1}{2} \sum_{i=1}^s r_i = \frac{T}{2}.
 \end{aligned} \tag{29}$$

Thus,

$$R_A^\pi(T) \geq \mathbb{E}_{\{x_t\}_{t=1}^T} \left[ \left| \frac{T}{2} - \sum_{i=1}^s r_i W_i \right| \right].$$

Now, we lower bound  $\mathbb{E}[|M|]$ , where  $M = \frac{T}{2} - \sum_{i=1}^s r_i W_i = \sum_{i=1}^s r_i (W_i - \frac{1}{2})$ , as  $\sum_{i=1}^s r_i = T$ . We denote  $m_i = r_i (W_i - \frac{1}{2})$ . Thus,  $M = \sum_{i=1}^s m_i$ , where  $\mathbb{E}[m_i] = 0$  and  $\sigma_i^2 = \mathbb{E}[m_i^2] = \frac{r_i^2}{4}$ . Note that  $\sigma_M = \sqrt{\sum_{i=1}^s \sigma_i^2} = \frac{1}{2} \sqrt{\sum_{i=1}^s r_i^2}$ . Now, using the Markov inequality,

$$\begin{aligned}
 \mathbb{E}[|M|] &\geq \sigma_M \mathbb{P}(|M| \geq \sigma_M) \geq \sigma_M \mathbb{P}\left(\frac{M}{\sigma_M} \geq 1\right) \\
 &\geq \sigma_M \mathbb{P}\left(\frac{\sum_{i=1}^s m_i}{\sigma_M} \geq 1\right).
 \end{aligned}$$

Using the Berry-Esseen theorem, we have that

$$\left| \Pr\left(\frac{\sum_{i=1}^s m_i}{\sigma_M} \leq 1\right) - \Phi(1) \right| \leq \frac{C_0}{\sigma_M} \max_{1 \leq i \leq s} \frac{\rho_i}{\sigma_i^2},$$

where  $\rho_i = \mathbb{E}[|m_i^3|] = \frac{r_i^3}{8}$ ,  $C_0$  is a constant and  $\Phi(\cdot)$  is the CDF of the standard Gaussian random variable. Thus, we get

$$\begin{aligned}
 \Pr\left(\frac{\sum_{i=1}^s m_i}{\sigma_M} \geq 1\right) &\geq 1 - \Phi(1) - C_0 \frac{\max_{1 \leq i \leq s} r_i}{2\sigma_M} \\
 &\geq 0.15 - C_0 \frac{\max_{1 \leq i \leq s} r_i}{\sqrt{\sum_{i=1}^s r_i^2}}.
 \end{aligned}$$

We thus have:

$$\begin{aligned} R_A^\pi(T) &\geq \mathbb{E}_{\{x_t\}_{t=1}^T} \left[ \max \left\{ \sum_{i=1}^s r_i W_i, T - \sum_{i=1}^s r_i W_i \right\} \right] \\ &\geq \frac{T}{2} + \frac{1}{2} \sqrt{\sum_{i=1}^s r_i^2} \left( 0.15 - C_0 \frac{\max_{1 \leq i \leq s} r_i}{\sqrt{\sum_{i=1}^s r_i^2}} \right). \end{aligned}$$

## H.2 Lower bound for general $L, C$

In this section, we extend the result in the previous section for a general  $L, C$  value. We consider an analogous adversarial request sequence where, files are requested from the first  $2C$  files uniformly, and the same file is requested  $r_i$  times from  $t_i = \sum_{j=0}^{i-1} r_j$  to  $t_i + r_i - 1$ , where  $r_0 = 1$  and  $1 \leq i \leq s$ . To be precise, for each  $1 \leq i \leq s$ , a file is drawn uniformly at random from the first two files at  $t = \sum_{j=0}^{i-1} r_j$  and is repeatedly requested till  $t = t_i + r_i - 1$ , i.e., a total of  $r_i$  times. We denote by  $W_i$  the file requested in the  $i^{\text{th}}$  phase.

We use the balls-into-bins technique from the proof of Lemma 1 of [7]. A bin is associated with each file from  $1, \dots, 2C$  where a request for that file is equivalent to a ball being thrown into that bin. The bins are numbered as  $1, 2, \dots, 2C$  and every two consecutive bins  $\{(2i-1, 2i)\}, 1 \leq i \leq C$  are combined to form  $C$  Super bins. Denote by  $Z_i, i = 1, 2, \dots, C$  the number of balls in the  $i^{\text{th}}$  super bin. Conditioned on  $Z_i$ , the number of balls in the bins  $2i-1$  and  $2i$  are jointly distributed as  $(V, Z_i - V)$ , where  $V$  is a binomial random variable with parameter  $(Z_i, \frac{1}{2})$ . Let  $H_i$  denote the number of balls in the bin containing the maximum number of balls among bins  $2i-1$  and  $2i$ . Then, as shown in the previous section, when  $\forall 1 \leq i \leq C, Z_i > 0$ :

$$\mathbb{E}(H_i | Z_i) \geq \frac{Z_i}{2} + \frac{1}{2} \sqrt{\sum_{j=1}^s r_j^2 \mathbb{I}_{W_j \in \{(2i-1, 2i)\}}} } \left( 0.15 - C_0 \frac{\max_{1 \leq j \leq s, W_j \in \{(2i-1, 2i)\}} r_j}{\sqrt{\sum_{j=1}^s r_j^2 \mathbb{I}_{W_j \in \{(2i-1, 2i)\}}}} \right),$$

The expected number of requests obtained by the best offline policy can be lower bounded as:

$$\begin{aligned} \max_{y \in \mathcal{Y}} \langle y, \mathbf{X}_{T+1} \rangle &\geq \mathbb{E} \left[ \sum_{i=1}^C H_i \right] \\ &= \sum_{i=1}^C \mathbb{E}[H_i] \\ &\stackrel{(a)}{=} CH_1 \\ &\stackrel{(b)}{=} C \mathbb{E}[\mathbb{E}[H_1 \mathbb{I}(Z_1 > 0) | Z_1]], \end{aligned}$$

where (a) follows from the request sequence being symmetric across the  $2C$  files and (b) follows from the tower property of conditional expectation as shown in

[7]. Thus,

$$\begin{aligned} \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} \rangle &\geq \frac{C\mathbb{E}[Z_1]}{2} + C\mathbb{E} \left[ \frac{1}{2} \sqrt{\sum_{i=1}^s r_i^2 \mathbb{I}_{W_i \in \{(1,2)\}}} \left( 0.15 - C_0 \frac{\max_{1 \leq i \leq s, W_i \in \{(1,2)\}} r_i}{\sqrt{\sum_{i=1}^s r_i^2 \mathbb{I}_{W_i \in \{(1,2)\}}}} \right) \right] \\ &\geq \frac{T}{2} + \frac{C}{2} \mathbb{E} \left[ 0.15 \sqrt{\sum_{i=1}^s r_i^2 \mathbb{I}_{W_i \in \{(1,2)\}}} - C_0 \max_{1 \leq i \leq s} r_i \right], \end{aligned} \quad (30)$$

as  $\mathbb{E}[Z_1] = T/C$ . Using (20) of [7],

$$\begin{aligned} \mathbb{E} \left[ \sqrt{\sum_{i=1}^s r_i^2 \mathbb{I}_{W_i \in \{(1,2)\}}} \right] &\geq \sqrt{\mathbb{E} \left( \sum_{i=1}^s r_i^2 \mathbb{I}_{W_i \in \{(1,2)\}} \right)} \left( 1 - \frac{\text{Var} \left( \sum_{i=1}^s r_i^2 \mathbb{I}_{W_i \in \{(1,2)\}} \right)}{2 \left( \mathbb{E} \left( \sum_{i=1}^s r_i^2 \mathbb{I}_{W_i \in \{(1,2)\}} \right) \right)^2} \right) \\ &= \sqrt{\frac{\sum_{i=1}^s r_i^2}{C}} \left( 1 - \frac{(C-1) \left( \sum_{i=1}^s r_i^4 \right)}{2 \left( \sum_{i=1}^s r_i^2 \right)^2} \right), \end{aligned}$$

as  $\text{Var}(\mathbb{I}_{W_i \in \{(1,2)\}}) = \frac{1}{C} \left( 1 - \frac{1}{C} \right)$  and  $\mathbb{E}(\mathbb{I}_{W_i \in \{(1,2)\}}) = \frac{1}{C}$ . Similar to (29),

$$\begin{aligned} \mathbb{E}_{\{\mathbf{x}_t\}_{t=1}^T} \left[ \sum_{t=1}^T \langle \mathbf{y}_t, \mathbf{x}_t \rangle \right] &= \frac{1}{2C} \sum_{i=1}^s r_i \left( \mathbf{y}_{t_i}(1) + \dots + \mathbf{y}_{t_i}(2C) \right) \\ &\leq \frac{1}{2} \sum_{i=1}^s r_i = \frac{T}{2}. \end{aligned}$$

Combining the above result and 30 gives us the following lower bound on the regret:

$$R_A^\pi(T) \geq \frac{1}{2} \left( 0.15 \sqrt{C \sum_{i=1}^s r_i^2} \left( 1 - \frac{(C-1) \left( \sum_{i=1}^s r_i^4 \right)}{2 \left( \sum_{i=1}^s r_i^2 \right)^2} \right) - 0.6 C \max_{1 \leq i \leq s} r_i \right).$$

## I Proof of part (b) of Theorem 5

In this section, we prove an upper bound on the regret of FTPL( $\alpha\sqrt{t}$ ) when cache updates are restricted to after every  $r$  slots only under adversarial file requests. The proof is based on the proof of Proposition 4.2 of [18]. Define the following potential function at each time slot  $t$ :

$$\Phi_t(\mathbf{x}) = \mathbb{E}_{\boldsymbol{\gamma} \sim \mathcal{N}(0, I)} \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{x} + \eta \boldsymbol{\gamma} \rangle \right]$$

The regret incurred in this setting can be expressed in the following way:

$$\begin{aligned}
R_A^{\text{FTPL}(\alpha\sqrt{t})}(T) &= \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} \rangle - \sum_{t=1}^T \mathbb{E}_\gamma [\langle \mathbf{y}_t, \mathbf{x}_t \rangle] \\
&= \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} \rangle - \sum_{i=1}^s \mathbb{E}_\gamma \left[ \left\langle \mathbf{y}_{\sum_{j=0}^{i-1} r_j}, \sum_{t=\sum_{k=0}^{i-1} r_k}^{\sum_{k=1}^i r_k} \mathbf{x}_t \right\rangle \right],
\end{aligned} \tag{31}$$

where we define  $r_0 = 1$ . This follows from the fact that the cache configuration can change only at the pre-defined  $s$  fixed time slots. Thus, we essentially have a time horizon of  $s$ , but each time slot  $i$ ,  $1 \leq i \leq s$  contains  $r_i$  requests instead. For brevity, we define

$$\mathbf{y}^i := \mathbf{y}_{\sum_{j=0}^{i-1} r_j}, \quad \mathbf{x}^i := \sum_{t=\sum_{k=0}^{i-1} r_k}^{\sum_{k=1}^i r_k} \mathbf{x}_t, \quad \eta^i := \eta_{\sum_{j=0}^{i-1} r_j}, \quad \mathbf{X}^i := \sum_{j=1}^i \mathbf{x}^j, \quad 1 \leq i \leq s.$$

We define the potential function  $\Phi_i : \mathbb{R}^L \rightarrow \mathbb{R}$  for  $1 \leq i \leq s$  in the following way:

$$\Phi_i(\mathbf{x}) = \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{x} + \eta^i \boldsymbol{\gamma} \rangle \right],$$

where  $\mathcal{Y}$  is the set of possible cache configurations, i.e., the set  $\{y \in \{0, 1\}^L : \|y\|_1 \leq C\}$ . As shown in the proof of Proposition 4.1 in [18],

$$\begin{aligned}
\nabla \Phi_i(\mathbf{X}^i) &= \nabla \mathbb{E}_\gamma [\langle \mathbf{y}^i, \mathbf{x} + \eta^i \boldsymbol{\gamma} \rangle] \Big|_{\mathbf{x}=\mathbf{X}^i} = \mathbb{E}_\gamma [\mathbf{y}^i]. \\
\mathbb{E}_\gamma [\langle \mathbf{y}^i, \mathbf{x}^i \rangle] &= \langle \nabla \Phi_i(\mathbf{X}^i), \mathbf{X}^i - \mathbf{X}^{i-1} \rangle \\
&= \Phi_i(\mathbf{X}^i) - \Phi_i(\mathbf{X}^{i-1}) - \frac{1}{2} \langle \mathbf{x}^i, \nabla^2 \Phi_i(\widetilde{\mathbf{X}}^i) \mathbf{x}^i \rangle,
\end{aligned}$$

where  $\widetilde{\mathbf{X}}^i = \mathbf{X}^{i-1} + \theta^i \mathbf{x}^i$ , for some  $\theta^i \in [0, 1]$  using Taylor's theorem. Thus, we have

$$\begin{aligned}
&\sum_{i=1}^s \mathbb{E}_\gamma [\langle \mathbf{y}^i, \mathbf{x}^i \rangle] \\
&= \Phi_s(\mathbf{X}^s) - \Phi_1(\mathbf{X}^0) + \sum_{i=2}^s [\Phi_{i-1}(\mathbf{X}^{i-1}) - \Phi_i(\mathbf{X}^{i-1})] \\
&\quad - \frac{1}{2} \sum_{i=1}^s \langle \mathbf{x}^i, \nabla^2 \Phi_i(\widetilde{\mathbf{X}}^i) \mathbf{x}^i \rangle.
\end{aligned}$$

Using Jensen's inequality, we have that

$$\begin{aligned}\Phi_s(\mathbf{X}_s) &= \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}^s + \eta^s \gamma \rangle \right] \\ &\geq \max_{\mathbf{y} \in \mathcal{Y}} \mathbb{E}_\gamma [\langle \mathbf{y}, \mathbf{X}^s + \eta^s \gamma \rangle] \\ &= \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}^s \rangle = \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}_{T+1} \rangle.\end{aligned}$$

Substituting the above results in (31), we get

$$\begin{aligned}R_A^{\text{FTPL}(\alpha\sqrt{t})}(T) &\leq \Phi_1(\mathbf{X}^0) + \sum_{i=2}^s [\Phi_i(\mathbf{X}^{i-1}) - \Phi_{i-1}(\mathbf{X}^{i-1})] \\ &\quad + \frac{1}{2} \sum_{i=1}^s \langle \mathbf{x}^i, \nabla^2 \Phi_i(\tilde{\mathbf{X}}^i) \mathbf{x}^i \rangle.\end{aligned}$$

Since  $\mathbf{x}^i$  contains  $r_i$  file requests, the quadratic form above may be upper bounded in the following way:

$$\begin{aligned}\langle \mathbf{x}^i, \nabla^2 \Phi_i(\tilde{\mathbf{X}}^i) \mathbf{x}^i \rangle &\leq r_i \max_{k,j,\mathbf{x}} (|\nabla^2 \Phi_i(\mathbf{x})|)_{kj} \langle \mathbf{x}^i, \mathbf{1} \rangle \\ &= r_i^2 \max_{k,j,\mathbf{x}} (|\nabla^2 \Phi_i(\mathbf{x})|)_{kj},\end{aligned}$$

where  $\mathbf{1}$  is the all ones vector. Moreover, from [7],

$$(|\nabla^2 \Phi_i(\mathbf{x})|)_{kj} \leq \frac{1}{\eta^i} \sqrt{\frac{2}{\pi}}.$$

Thus,

$$\frac{1}{2} \sum_{i=1}^s \langle \mathbf{x}^i, \nabla^2 \Phi_i(\tilde{\mathbf{X}}^i) \mathbf{x}^i \rangle \leq \sqrt{\frac{2}{\pi}} \sum_{i=1}^s \frac{r_i^2}{\eta^i}.$$

It can also be shown that:

$$\begin{aligned}\Phi_i(\mathbf{X}^{i-1}) - \Phi_{i-1}(\mathbf{X}^{i-1}) &= \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}^{i-1} + \eta^i \gamma \rangle \right] - \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \mathbf{X}^{i-1} + \eta^{i-1} \gamma \rangle \right] \\ &\leq \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, |\eta^i - \eta^{i-1}| \gamma \rangle \right] \\ &= |\eta^i - \eta^{i-1}| \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \gamma \rangle \right].\end{aligned}$$

We also have from [10] that:

$$\Phi_1(\mathbf{X}^0) = \eta^1 \mathbb{E}_\gamma \left[ \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \gamma \rangle \right] \leq \eta^1 C \sqrt{2 \log \left( \frac{N}{C} \right)}.$$

Thus, combining the above bounds, we get

$$\begin{aligned} R_A^{\text{FTPL}(\alpha\sqrt{\bar{t}})}(T) &\leq \eta^1 C \sqrt{2 \log \left( \frac{N}{C} \right)} + \eta^s C \sqrt{2 \log \left( \frac{N}{C} \right)} + \sqrt{\frac{2}{\pi}} \sum_{i=1}^s \frac{r_i^2}{\eta^i} \\ &\leq \alpha C \sqrt{2 \log \left( \frac{N}{C} \right)} + \alpha C \sqrt{T} \sqrt{2 \log \left( \frac{N}{C} \right)} + \sqrt{\frac{2}{\pi}} \sum_{i=1}^s \frac{r_i^2}{\alpha \sqrt{\sum_{j=0}^{i-1} r_j}}. \end{aligned}$$

## J Proof of part (a) of Theorem 6

In this section, we prove a lower bound on the regret of any policy  $\pi$  in the restricted switching case where the cache is allowed to update its contents only after every  $r$  requests. The proof uses (28) to bound the regret. We consider the adversarial request sequence where files are requested from the top  $2C$  files, and the same file is requested  $r$  times in each of the  $T/r$  ‘phases’ (using terminology defined in Appendix H). To be precise, at the beginning of each phase  $i$  where  $1 \leq i \leq T/r$ , i.e., at  $t = 1 + (i-1)r$ , a file is drawn uniformly at random from the first  $2C$  files. This file is repeatedly requested  $r$  times till  $t = ir$ . Throughout this proof, if it is not explicitly mentioned with respect to what the expectation is being taken, it can be assumed that the expectation is being taken with respect to  $\{x_t\}_{t=1}^T$ . The reward obtained by the optimal static configuration in hindsight can be bounded in the following way:

$$\begin{aligned} \mathbb{E} \left( \max_{y \in \mathcal{Y}} \left\langle y, \sum_{t=1}^T x_t \right\rangle \right) &= r \mathbb{E} \left( \max_{y \in \mathcal{Y}} \left\langle y, \sum_{t=1}^{T/r} x_t \right\rangle \right) \\ &\geq r \left( \frac{T}{2r} + \sqrt{\frac{CT}{2r\pi}} - \frac{\sqrt{r}(\sqrt{2}+1)C^{3/2}}{2\sqrt{2\pi T}} - \sqrt{\frac{2}{\pi}} \frac{rC^2}{T} \right) \\ &= \frac{T}{2} + \sqrt{\frac{CrT}{2\pi}} - \Theta \left( \frac{r\sqrt{r}}{\sqrt{T}} \right), \end{aligned}$$

where in the second last step, we used Theorem 2 of [7]. For  $1 \leq i \leq s$ , let

$$\mathbf{y}^i := \sum_{t=1+(i-1)r}^{ir} \mathbf{y}_t,$$

which is the sum of the caching configuration vectors in phase  $i$ . Note that the cache configuration at each time slot is independent of the file request that

arrives in that slot. Now, to upper bound the reward obtained by any policy  $\pi$ ,

$$\begin{aligned} \mathbb{E} \left( \sum_{t=1}^T \langle \mathbf{y}_t, \mathbf{x}_t \rangle \right) &= \sum_{i=1}^{T/r} \frac{1}{2C} \sum_{k=1}^{2C} y^i(k) \\ &\leq \sum_{t=1}^{T/r} \frac{1}{2C} \sum_{k=1}^L y^i(k) \\ &\leq \frac{T}{r} \cdot \frac{1}{2C} \cdot rC \\ &= \frac{T}{2}, \end{aligned}$$

The second last step follows from the fact that the sum of the elements of  $\mathbf{y}^i$  is exactly  $rC$  as the cache configuration remains constant in each phase.

$$\begin{aligned} \therefore R_A^\pi(T) &\geq \mathbb{E}_{\{\mathbf{x}_t\}_{t=1}^T} \left( \mathbf{y}^* \cdot \sum_{t=1}^T \mathbf{x}_t - \sum_{t=1}^T \mathbf{y}_t \cdot \mathbf{x}_t \right) \\ &\geq \sqrt{\frac{CrT}{2\pi}} - \Theta \left( \frac{r\sqrt{r}}{\sqrt{T}} \right). \end{aligned}$$

For this bound to be meaningful, the second term needs to be order-wise smaller than the first term:

$$\frac{r\sqrt{r}}{\sqrt{T}} < \sqrt{rT} \implies r < T$$

Hence this is an  $\mathcal{O}(\sqrt{rT})$  lower bound for  $r < o(T)$ .

For  $r = \Omega(T)$ , the first set of  $r$  requests always gives a regret that is  $\Omega(r)$  as in time slot 1, a random  $C$  files out of  $L$  are stored, and hence we get an overall regret of  $\Omega(T)$ .