

Dynamic Resource Allocation in the Cloud

with Near-Optimal Efficiency

Sebastian Perez-Salazar* Ishai Menache† Mohit Singh‡ Alejandro Toriello§

August 15, 2019

Abstract

Cloud computing has motivated renewed interest in resource allocation problems with new consumption models. A common goal is to share a resource, such as CPU or I/O bandwidth, among distinct users with different demand patterns as well as different quality of service requirements. To ensure these service requirements, cloud offerings often come with a service level agreement (SLA) between the provider and the users. An SLA specifies the amount of a resource a user is entitled to utilize. In many cloud settings, providers would like to operate resources at high utilization while simultaneously respecting individual SLAs. There is typically a tradeoff between these two objectives; for example, utilization can be increased by shifting away resources from idle users to “scavenger” workload, but with the risk of the former then becoming active again. We study this fundamental tradeoff by formulating a resource allocation model that captures basic properties of cloud computing systems, including SLAs, highly limited feedback about the state of the system, and variable and unpredictable input sequences. Our main result is a simple and practical algorithm that achieves near-optimal performance on the above two objectives. First, we guarantee nearly optimal utilization of the resource even if compared to the omniscient offline dynamic optimum. Second, we simultaneously satisfy all individual SLAs up to a small error. The main algorithmic tool is a multiplicative weight update algorithm, and a duality argument to obtain its guarantees. Experiments on both synthetic and real production traces demonstrate the merits of our algorithm in practical settings.

Key words: Cloud Computing, Online Algorithms

*Georgia Institute of Technology, sperez@gatech.edu

†Microsoft Research, ishai@microsoft.com

‡Georgia Institute of Technology, mohitsinghr@gmail.com

§Georgia Institute of Technology, atoriello@isye.gatech.edu

1 Introduction

Cloud computing has motivated renewed interest in resource allocation, manifested in new consumption models (e.g., AWS spot pricing), as well as the design of resource-sharing platforms [1, 2]. These platforms need to support a heterogeneous set of users, also called tenants, that share the same physical computing resource, e.g., CPU, memory, I/O bandwidth. Providers such as Amazon, Microsoft and Google offer cloud services with the goal of benefiting from economies of scale. However, the inefficient use of resources – over-provisioning on the one hand or congestion on the other – could result in a low return on investment or in loss of customer goodwill, respectively. Hence, resource allocation algorithms are key for efficiently utilizing cloud resources.

To ensure quality of service, cloud offerings often come with a *service level agreement* (SLA) between the provider and the users. An SLA specifies the amount of resource the user is entitled to consume. Perhaps the most common example is renting a virtual machine (VM) that guarantees an explicit amount of CPU, memory, etc. Naturally, VMs that guarantee more resources are more expensive. In this context, a simple allocation policy is to assign each user the resources specified by their SLAs. However, such an allocation can be wasteful, as users may not need the resource at all times. In principle, a dynamic allocation of resources can increase the total efficiency of the system. However, allocating resources dynamically without carefully accounting for SLAs can lead to user dissatisfaction.

Recent scheduling proposals address these challenges through work-maximizing yet fair schedulers [3, 4]. However, such schedulers do not have explicit SLA guarantees. On the other hand, other works focus on enforcing SLAs [5–7], but do not explicitly optimize the use of extra resources.

Our goal in this work is to understand the fundamental tradeoff between high utilization of resources and SLA satisfaction of individual users. In particular, we design algorithms that guarantee *both* near optimal utilization as well as the satisfaction of individual SLAs, simultaneously. To that end, we formulate a basic model for online dynamic resource allocation. We focus on a single divisible resource, such as CPU or I/O bandwidth, that has to be shared among multiple users. Each user also has an SLA that specifies the fraction of the resource it expects to obtain. The actual demand of the user is in general time-varying, and may exceed the fraction specified in the SLA. As in

many real systems, the demand is not known in advance, but rather arrives in an online manner. Arriving demand is either processed or queued up, depending on the resource availability. In many real-world scenarios, it is difficult to measure the actual demand size (see, e.g., [8]). Accordingly, we assume that the system (and the underlying algorithm) receives only a simple *binary feedback* per user at any given time: whether the user queue is empty (the user’s work arriving so far has been completed), or not. This is a plausible assumption in many systems, because one can observe workload activity, yet anticipating how much of the resource a job will require is more difficult. Additionally, it also models settings where demands are not known in advance.

While online dynamic resource allocation problems have been studied in different contexts and communities (see Section 1.3 for an overview), our work aims to address the novel aspects arising in the cloud computing paradigm, particularly the presence of SLAs, the highly limited feedback about the state of the system, and a desired robustness over arbitrary input sequences. For the algorithm design itself, we pay close attention to practicality; our approach involves fairly simple computations that can be implemented with minimal overhead of space or time. Our algorithm achieves nearly optimal utilization of the resource, as well as approximately satisfying the SLA of each individual user. We see two main use-cases for the algorithm:

- In enterprise settings (“private cloud”), different applications or organizations share the same infrastructure. These often have SLAs, but providers would still like to maximize the ROI by maximizing utilization [9].
- In public clouds, users buy VMs, which are offered at different “sizes” (which is practically the SLA). In addition, the service providers offer “best-effort” alternatives, such as Azure Batch (MS) or Spot instances (AWS). In our model, these services can be modeled by giving an SLA of zero. Here, not only satisfying the SLA for the VMs is important but achieving high utilization is also important. Indeed, the provider is paid for the latter workloads only if work is done for these services. Our work can be viewed as a principled way to accommodate such services, and even give VMs better service than expected, an important consideration as public cloud offerings gradually become commoditized.

1.1 The Model

We consider the problem of having multiple tenants or users sharing a single resource, such as CPU, I/O or networking bandwidth. For simplicity, we assume that the total resource capacity is normalized to 1. We have N users sharing the resource, a finite but possibly unknown discrete time horizon indexed $t = 1, \dots, T$, and an underlying queuing system. For each user i , we are also given an expected share of resource $\beta(i) \geq 0$ satisfying $\sum_{i=1}^N \beta(i) \leq 1$. The input is an online sequence of workloads $L_1, \dots, L_T \in \mathbf{R}_+^N$, where $L_t(i) \geq 0$ corresponds to i 's workload arising at time t . The system maintains a queue $Q_t(i)$, denoting i 's remaining work at time t . In our model, the decision maker does *not* have direct access to the values of the queues or the workloads. This allows us to consider settings where the job sizes are not known in advance and minimal information is available about the underlying system, a regular occurrence in many cloud settings. At time t , the following happens:

1. **Feedback:** The decision maker observes which queues are non-empty (the set of users i with $Q_t(i) > 0$, the *active* users), and which are empty ($Q_t(i) = 0$, the *inactive* users).
2. **Decision:** The decision maker updates user resource allocations $h_t(i)$, satisfying $\sum_i h_t(i) \leq 1$.
3. **Update:** The load $L_t(i)$ for each i arrives and each user processes as much of the work from the queue plus the arriving workload as possible. The work completed by user i in step t is

$$w_t(i) := \min\{h_t(i), L_t(i) + Q_t(i)\}.$$

The queues at the end of the time step are updated accordingly,

$$Q_{t+1}(i) = \max\{0, L_t(i) + Q_t(i) - h_t(i)\}.$$

We assess the performance of any algorithm based on two measures.

1. **Work Maximization.** The algorithm should maximize the total work completed over all users, and thus utilize the resource as much as possible.

2. **SLA Satisfaction.** The algorithm should (approximately) satisfy the SLAs in the following manner. The work completed by user i up to any time $1 \leq t \leq T$ should be no less than the work completed for this user up to t if it were given a constant fraction $\beta(i)$ of the resource over the whole horizon.

Achieving either of the criteria on their own is straightforward. A greedy strategy that takes away resources from an idle user and gives them to any user whose queue is non-empty is approximately work-maximizing (see Appendix C for details). On the other hand, to satisfy the SLAs, we give each user a *static* assignment of $h_t(i) := \beta(i)$ for all t . Naturally, the two criteria compete with each other; the following examples illustrate why these simple algorithms do not satisfy both simultaneously.

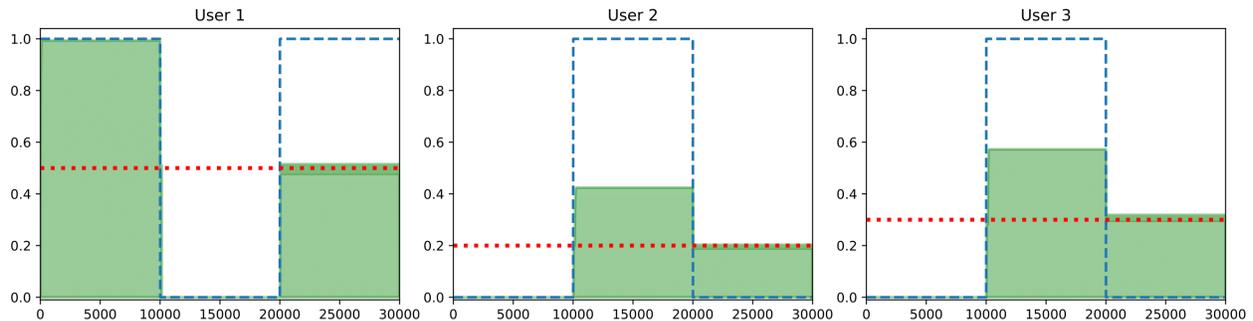


Figure 1: Example of loads and work for 3 users. The blue dashed lines show each user’s workload. The green areas represent the work done by the users. The red dashed lines depict SLAs.

Example 1.1 We have a shared system with three users and corresponding SLAs $\beta(1) = 0.5$, $\beta(2) = 0.2$ and $\beta(3) = 0.3$. Loads are defined by

$$L_t(1) = \begin{cases} 1 & t = 1, \dots, T/3, 2T/3 + 1, \dots, T \\ 0 & t = T/3 + 1, \dots, 2T/3 \end{cases}, \text{ and } L_t(2) = L_t(3) = 1 - L_t(1).$$

We assume that T is divisible by 3. In Figure 1 we show the three users’ loads in blue dashed lines and the corresponding SLAs in dotted red lines. The static solution given by the SLAs, i.e. $h_t(i) = \beta(i)$ for all t , ensures a total of $5T/6$ work done. However, the dynamic policy shown in the green line, given by

| | | | |
|----------|------------|-------------------|-----------------|
| t | $[0, T/3]$ | $[T/3 + 1, 2T/3]$ | $[2T/3 + 1, T]$ |
| $h_t(1)$ | 1 | 0 | 0.5 |
| $h_t(2)$ | 0 | 0.4 | 0.2 |
| $h_t(3)$ | 0 | 0.6 | 0.3 |

ensures T work is done (the green area). Moreover, it also ensures SLA satisfaction at all times. An alternative policy is

| | | | |
|----------|------------|-------------------|-----------------|
| t | $[0, T/3]$ | $[T/3 + 1, 2T/3]$ | $[2T/3 + 1, T]$ |
| $h_t(1)$ | 1 | 0 | 0 |
| $h_t(2)$ | 0 | 1 | 0 |
| $h_t(3)$ | 0 | 0 | 1 |

which is also work maximizing. However, it does not ensure SLA satisfaction. Indeed, this policy does not satisfy user 3's SLA at any time in $(T/3, 2T/3]$.

1.2 Our Results and Contributions

We design a simple and efficient online algorithm that achieves approximate work maximization as well as approximate SLA satisfaction even in the limited feedback model that we consider. For work maximization, we analyze the performance by comparing our algorithm to the *optimal offline dynamic allocation* that knows all the data up front. In contrast, our online algorithm receives limited feedback even in an online setting. Thus, our aim is to minimize the quantity

$$\text{work}_{\mathbf{h}_1^*, \dots, \mathbf{h}_T^*} - \text{work}_{\text{alg}} = \sum_{t=1}^T \sum_i w_t^*(i) - \sum_{t=1}^T \sum_i w_t(i),$$

where $\text{work}_{\mathbf{h}_1^*, \dots, \mathbf{h}_T^*}$ is the optimal offline work done by dynamic allocations $\mathbf{h}_1^*, \dots, \mathbf{h}_T^*$, $\mathbf{w}_t^* = (w_t^*(1), \dots, w_t^*(N))$ is the work done at time t by these allocations, and work_{alg} is the work done by the algorithm with allocations $\mathbf{h}_1, \dots, \mathbf{h}_T$ and work $\mathbf{w}_t = (w_t(1), \dots, w_t(N))$ at time t . The objective of the decision maker is to *minimize* this quantity by constructing a sequence of good allocations that approach the best allocations in hindsight. Note that our benchmark is *dynamic*, rather than the more common *static* offline optimum usually considered in regret minimization [10–12]. Similarly, for SLA satisfaction, our benchmark is the total work done for a user if they

were given $\beta(i)$ resources for each time $1 \leq t \leq T$. We give a bi-criteria online algorithm that achieves nearly the same performance as the benchmarks if the resources for the latter are slightly more constrained than that of the algorithm. Algorithm 1, which we formally describe in Section 2, follows a multiplicative weight approach. The idea is to boost the allocations of active users by a factor greater than 1, with more emphasis on users with current allocation below their SLA. This intuition translates into a simple update that ensures high utilization of the resource and SLA satisfaction; formally:

Theorem 1 *For any $0 < \varepsilon \leq \frac{1}{10}$, SLAs $\beta = (\beta(1), \dots, \beta(N))$ satisfying $\beta(i) \geq 2\frac{\varepsilon}{N}$, and online loads $L_1, \dots, L_T \in \mathbf{R}_{\geq 0}^N$, Algorithm 1 achieves the following guarantees:*

1. **Approximate Work Maximization.** *Let $\mathbf{h}_1^*, \dots, \mathbf{h}_T^* \in [0, 1]^N$ be the optimal offline sequence of allocations such that $\sum_i h_t^*(i) = 1$ for all $1 \leq t \leq T$. Then*

$$\text{work}_{alg} \geq (1 - \varepsilon) \text{work}_{\mathbf{h}_1^*, \dots, \mathbf{h}_T^*} - \mathcal{O}(N\varepsilon^{-2} \log(N/\varepsilon)).$$

2. **Approximate SLA Satisfaction.** *There exists $\tilde{s} = \tilde{s}(N, \varepsilon, \beta) = \mathcal{O}(N^2\varepsilon^{-3} \log(N/\varepsilon))$ such that for any user i and time t , if we take $\mathbf{h}'_1, \dots, \mathbf{h}'_T \in [0, 1]^N$ to be any sequence of allocations with $h'_t(i) \leq \beta(i)$, then*

$$\sum_{\tau=1}^t w_{\tau}(i) \geq (1 - 2\varepsilon) \sum_{\tau=1}^t w'_{\tau}(i) - \beta(i)\tilde{s},$$

where w'_t is the work performed by the allocations $\mathbf{h}'_1, \dots, \mathbf{h}'_T$.

The first part of Theorem 1 asserts that if T is known for the system, then we can achieve $\text{work}_{\text{offline}} - \text{work}_{alg} \leq \mathcal{O}(T^{2/3} \log T)$ with $\varepsilon = \Theta\left(\frac{1}{T^{1/3}}\right)$. Yet this choice of ε is suboptimal for SLA satisfaction. We can achieve an improved bound for SLA satisfaction (when $t = T$) by picking $\varepsilon = \Theta\left(\frac{1}{T^{1/4}}\right)$. If T is unknown, we can use a standard doubling trick, see for example [12]. Summarizing, we obtain the following result.

Corollary 1 *For $\varepsilon = \Theta\left(\frac{1}{T^{1/4}}\right)$, Algorithm 1 guarantees*

$$\text{work}_{\mathbf{h}_1^*, \dots, \mathbf{h}_T^*} - \text{work}_{alg} \leq \mathcal{O}(T^{3/4}),$$

where $\mathbf{h}_1^*, \dots, \mathbf{h}_T^*$ are the optimal offline dynamic allocations.

As a consequence of Corollary 1 we obtain that, as T grows, the *rate* of work done by our algorithm $\frac{\text{work}_{\text{alg}}}{T}$ approaches the rate of work done by the optimal dynamic solution $\frac{\text{work}_{\mathbf{h}_1^*, \dots, \mathbf{h}_T^*}}{T}$. We emphasize again that this is a stronger guarantee using the much powerful optimal dynamic solution as a benchmark rather than the typical static allocation used in regret analysis of online algorithms. Such a stronger guarantee can be obtained in our model because the incomplete work remains stored in the queues until we are able to finish it. This allows the algorithm to *catch up* with the incomplete work. In Theorem 2, we give a general lower bound showing our result is nearly optimal.

The second result in Theorem 1 states that the work done by any individual user will be comparable to the work done by their promised SLA. Rephrased, this says that the *delay* or *latency*, measured as queue length, will not be much larger than the delay when the user receives their SLA continuously. By using $\varepsilon = \Theta\left(\frac{1}{T^{1/4}}\right)$ we obtain the following result.

Corollary 2 *For a user i , $Q_T(i) \leq Q'_T(i) + \mathcal{O}(T^{1/2} \log T)$, where Q_t is the queue given by Algorithm 1 and Q'_t is the queue induced by any dynamic policy $\mathbf{h}'_1, \dots, \mathbf{h}'_T$ with $h'_t(i) \leq \beta(i)$.*

Corollary 1's guarantee is near-optimal in terms of work maximization, as the following result shows. The proof of this result appears in Appendix D.

Theorem 2 *For any online deterministic algorithm \mathcal{A} for our model, there is a sequence of online loads L_1, \dots, L_T such that $\text{work}_{\mathbf{h}_1^*, \dots, \mathbf{h}_T^*} - \text{work}_{\mathcal{A}} = \Omega\left(\sqrt{T}\right)$, where $\mathbf{h}_1^*, \dots, \mathbf{h}_T^*$ are the optimal offline dynamic allocations.*

Our algorithm follows a mirror descent approach [11, 13]. Unable to see the lengths of the queues, a (naive) approach is to pretend that active users have gigantic queues. From this approach we extract a simple update rule that multiplicatively boosts active users; however, active users who are under their SLA are boosted slightly more than other active users. If inactive users are assigned more than ε of the resource, where ε is an algorithm parameter, active users ramp up their allocation in few iterations. In the opposite case, at least $1 - \varepsilon$ of the resource is assigned to active users, and the slight boost to users below their SLA ensures a healthy re-balancing of the resource. We show that this efficient heuristic strategy is enough to achieve approximate work maximization and SLA

satisfaction. We remark that the mirror descent interpretation is only used to provide intuition for the algorithm, and our proofs follow a different path than the usual mirror descent analysis. Later, we detail a slight modification that enjoys the same theoretical guarantees as Algorithm 1 but in practice exhibits more desirable behavior (see Algorithm 2). Intuitively, among active users, the modified algorithm tries to keep allocations proportional to their SLAs. This behavior is appealing; for example, a user A with twice the SLA of another user B would expect in practice to perform at least twice as much work. Similarly, user B would expect to receive no less than half of user A 's allocation.

The analysis of the algorithm relies on a primal-dual fitting approach. For work maximization, we can write the offline dynamic optimal allocation as a solution to a linear program and then construct feasible dual solutions with objective value close to the algorithm's resource utilization. A crucial ingredient of the algorithm is the use of entropic projection on the *truncated* simplex, which ensures every user gets at least a fraction ε/N of the resource at all times. Intuitively, this means any user with a non-empty queue will recover their SLA requirement in a few steps.

We do an extensive analysis of our algorithm on synthetic data as well as real data obtained from CPU traces of a production service in Microsoft's cloud. We aim to quantify the performance of the algorithm on three objectives, (i) work maximization, (ii) SLA guarantee and (iii) queue behaviour. While our theoretical results give guarantees for these objectives, we show experimentally that the algorithm exceeds these guarantees. We benchmark the algorithm against natural online algorithms, such as a static allocation as given by the SLA guarantee, or the algorithm that aims to proportionally distribute the resource among active clients (according to their SLA). We also benchmark against offline algorithms that know all input data up front; our algorithm's performance on various measures is comparable to the offline algorithms.

This work is organized as follows. In Section 2, we present the preliminaries and the basic version of the multiplicative weight algorithm. Section 3 contains the proof of Theorem 1 in the bi-criteria form. We split the proof into two parts: work maximization in Section 3.2 and SLA satisfaction in Section 3.3. In subsection 3.4, we present the extension of our algorithm and its guarantees. Finally, in Section 4 we present numerical experiments that empirically validate our results, and conclude in Section 5.

1.3 Related Work

There has been significant work in resource allocation problems arising from cloud computing applications, both from a practical as well as a theoretical standpoint [1, 2, 5–9, 14, 15]. A focus of many of these works has been to understand the trade-offs between efficiency and ensuring guarantees to individual users. Also, a large literature has been dedicated to stochastic arrival with unknown but random demands [16–18]. More broadly, online resource allocation is a classical problem and has been studied for networking in adversarial as well as stochastic settings; we refer the reader to the books [19–21] on the topic.

Our work differs from the aforementioned lines of work by combining the three following elements already in the literature. First, we digress from the stochastic arrival model to a more adversarial scenario and we do a worst-case analysis. This makes our algorithm robust against unpredictable users’ demands. For instance, demands in the morning could be totally different from demands in the afternoon or the morning of the next day. We are able to provide a single strategy that adapts easily to any scenario. Second, and in addition to the previous point, we provide simple online algorithms that perform well even under *limited feedback*, a typical situation in cloud systems in which we can only read afterwards the utilization of a resource. Finally, we consider SLA satisfaction as a measure of user contentment and we seek to satisfy it up to a small error.

There is now an extensive literature devoted to the pricing of cloud computing services. In [22] the authors study a genetic model for generating a suitable pricing function in the cloud market. In [23], pricing is studied via a revenue management formulation to address resource provisioning decisions. See also [24, 25] for more pricing models. A closely related topic is fairness in resource allocation [3, 4]. Although we do not directly consider pricing, work maximization could be interpreted as a way to obtain extra revenue by allocating unused resources to active users.

More recently, there has been work considering over-commitment in the cloud [26–28], that is, selling resources beyond server capacity. One of the objectives of over-commitment is to reduce the number of servers opened in order to minimize energy consumption. In our work, we do not assume over-commitment, yet our algorithm can still be applied to this setting (see Algorithm 2). Besides work maximization, no other strong guarantees can be ensured under this regime.

A fundamental tool in our design is *mirror descent* algorithms [13]. These first-order iterative algorithms have been widely used in optimization [13], online optimization and machine learning [11, 12, 29] to generate update policies under limited feedback. Similarly, multiplicative weights algorithms have been widely studied in optimization [10, 30], online convex optimization [11], online competitive analysis [31] and learning theory [12, 32]. Our results bear some resemblance to regret analysis [11, 33], where the benchmark is the optimal offline static policy.

2 Algorithm

2.1 Preliminaries

For $N \geq 1$, we identify the set of *users* with the set $[N] = \{1, \dots, N\}$. For $0 < \varepsilon < 1$, we call an allocation $\mathbf{h} = (h(1), \dots, h(N)) \in [0, 1]^N$ a $(1 - \varepsilon)$ -*allocation* if $\sum_i h(i) \leq 1 - \varepsilon$. In what follows, we will assume that $N \geq 2$, that is, the system consists of at least two users.

For any t , we define the set of *active users* at that time as the set of users with non-empty queue, and denote this set by A_t . Observe that $h_t(i) = w_t(i)$ for all active users. Let $B_t = [N] \setminus A_t$ be the sets of users with empty queues at time t ; we call these users *inactive*. A_t and B_t correspond to the feedback given to the decision maker. Also, let $A_t^1 = \{i \in A_t : h_t(i) < \beta(i)\}$ be the set of active users with allocation below their SLA and $A_t^2 = A_t \setminus A_t^1$ be the set of active users receiving at least their SLAs.

We assume without loss of generality that the allocations set by the decision maker always add up to 1. We propose an algorithm that uses a multiplicative weight strategy to boost a subset of users by multiplying their allocation by factor greater than one. Because the allocations do not sum to one after applying the update, we then project them onto the simplex using the KL-divergence metric. Furthermore, to ensure no user gets an allocation arbitrarily close to zero, we in fact project onto the *truncated simplex*,

$$\Delta_\varepsilon = \{\mathbf{x} = (x(1), \dots, x(N)) : \|\mathbf{x}\|_1 = 1, x(i) \geq \varepsilon/N, \forall i\}.$$

To fix notation, let $\pi_{\Delta_\varepsilon}(\cdot)$ be the *projection function* onto Δ_ε using Kullback–Leibler divergence

(KL-divergence for short), i.e., $\pi_{\Delta_\varepsilon}(\mathbf{y}) := \operatorname{argmin}_{\mathbf{x} \in \Delta_\varepsilon} \sum_i x(i) \log(x(i)/y(i))$, where $\mathbf{y} = (y(1), \dots, y(N)) \in \mathbf{R}_{\geq 0}^N$. In Appendix A, we show how to efficiently compute this projection. The following proposition states some basic facts that are useful in our analysis. The proof appears in Appendix A.

Proposition 1 *Let $\mathbf{y} \in \mathbf{R}_+^N$, $\mathbf{x} = \pi_{\Delta_\varepsilon}(\mathbf{y})$, and $S = \{i : x(i) = \frac{\varepsilon}{N}\}$. Then:*

- (a) *If $y(1) \leq y(2) \leq \dots \leq y(N)$, then $S = \{1, \dots, k\}$ for some $k \geq 0$.*
- (b) *$x(i) = y(i)e^{\mu_i}C$, where $C = \left(\frac{1 - \frac{\varepsilon}{N}|S|}{\sum_{j \notin S} y(j)}\right)$, $\mu_i \geq 0$ for all i and $\mu_i = 0$ for $i \notin S$.*
- (c) *\mathbf{x} can be computed in $\mathcal{O}(N \log N)$ time.*

2.2 The Multiplicative Weight Algorithm

We propose an algorithm that follows a multiplicative weight strategy (see Algorithm 1). We describe here the basic approach given by the mirror descent algorithm. In Section 3.4 we present an extension of the algorithm that in practice shows a better relation between the allocations and the ratios between the SLAs.

Algorithm 1: Multiplicative Weight Update Algorithm

Input: Parameters $0 < \varepsilon \leq \frac{1}{10}, 0 < \eta < \frac{1}{3}$.

1 Initialization: \mathbf{h}_1 any allocation over Δ_ε and $\lambda = \frac{\varepsilon^2}{8N}$.

2 for $t = 1, \dots, T$ do

3 Set allocation \mathbf{h}_t .

4 Read active and inactive users A_t and B_t . $A_t^1 = \{i \in A_t : h_t(i) < \beta(i)\}$, $A_t^2 = A_t \setminus A_t^1$.

5 Set gain function $g_t(i) = \begin{cases} 1 + \lambda & i \in A_t^1 \\ 1 & i \in A_t^2 \\ 0 & i \in B_t \end{cases}$.

6 Update allocation:

$$\hat{h}_{t+1}(i) = h_t(i)e^{\eta g_t(i)}.$$

$$\mathbf{h}_{t+1} = \pi_{\Delta_\varepsilon}(\hat{h}_{t+1}).$$

Intuitively, the algorithm boosts active users at the expense of inactive ones, and boosts users slightly more if they are currently under their SLA. The algorithm update rule comes from a mirror descent approach applied to a Lagrangian relaxation of a work-maximization linear function at time t . More formally, under the assumption that active users have a huge queue, we aim to maximize the objective $\sum_{i \in A_t} w_t(i)$ subject to $w_t(i) \geq \beta(i)$ for $i \in A_t$. The update rule is obtained after applying a mirror descent with a KL-divergence distance generating function over the simplex to the Lagrangian relaxation of the previous problem (see [13, 34]). We restrict the projection to the truncated simplex so no user gets an allocation too close to 0. We use this update rule solely to guide the algorithm's decisions; however, the proof of work maximization and SLA satisfaction do not follow from the standard analysis of mirror descent.

3 Analysis

To give the analysis of the algorithm and prove Theorem 1, we prove the following stronger guarantees about Algorithm 1. We compare its performance to the optimal offline dynamic strategy that uses at most $1 - 4\varepsilon$ fraction of the resources at each time step.

Theorem 3 *Given loads L_1, \dots, L_T , for any $\varepsilon > 0$ and $\eta > 0$ such that $\varepsilon \leq 1/10$, Algorithm 1 guarantees*

$$\text{work}_{\mathbf{h}_1^*, \dots, \mathbf{h}_T^*} - \text{work}_{\text{alg}} \leq 8 \frac{N}{\varepsilon^2 \eta} \ln(N/\varepsilon),$$

where $\mathbf{h}_1^*, \dots, \mathbf{h}_T^*$ is the optimal offline sequence of $(1 - 4\varepsilon)$ -allocations.

The first guarantee of Theorem 1 regarding work maximization now follows simply from Theorem 3. Given any offline dynamic policy $\mathbf{h}_1, \dots, \mathbf{h}_T$ such that $\sum_i h_t(i) = 1$, we define $\bar{\mathbf{h}}_t := (1 - \varepsilon)\mathbf{h}_t$ which satisfies the assumption of Theorem 3. Now we have

$$\begin{aligned} \text{work}_{\text{alg}} &\geq \text{work}_{\bar{\mathbf{h}}_1, \dots, \bar{\mathbf{h}}_T} - 8 \frac{N}{(\varepsilon/4)^2 \eta} \ln(4N/\varepsilon) \\ &\geq (1 - \varepsilon) \cdot \text{work}_{\mathbf{h}_1, \dots, \mathbf{h}_T} - 2000 \frac{N}{\varepsilon^2 \eta} \ln(N/\varepsilon), \end{aligned}$$

where the first inequality follows from Theorem 3 and the second inequality follows since the work

$$\begin{array}{l|l}
(P_\varepsilon) \quad \max & \sum_{i=1}^N \sum_{t=1}^T w_t(i) \\
\text{s.t.} & \\
& \forall t, i \quad \sum_{s=1}^t w_s(i) \leq \sum_{s=1}^t L_s(i) \quad (1) \\
& \forall t \quad \sum_{i=1}^N w_t(i) \leq 1 - \varepsilon \quad (2) \\
& \forall t \quad \mathbf{w}_t \geq 0 \quad (3) \\
(D_\varepsilon) \quad \min & \sum_{i=1}^N \sum_{t=1}^T L_t(i) \gamma_t(i) + (1 - \varepsilon) \sum_{t=1}^T \beta_t \\
\text{s.t.} & \\
& \forall t, i \quad \gamma_t(i) + \beta_t \geq 1 \quad (4) \\
& \forall t \quad \gamma_t \geq \gamma_{t+1} \quad (5) \\
& \forall t \quad \beta_t, \gamma_t \geq 0 \quad (6)
\end{array}$$

Figure 2: The primal and dual LP formulation for the maximum work problem.

done by the dynamic policy $\mathbf{h}_1, \dots, \mathbf{h}_T$, say $\mathbf{w}_1, \dots, \mathbf{w}_T$, induces a feasible work $\bar{\mathbf{w}}_1, \dots, \bar{\mathbf{w}}_T$ under the dynamic policy $\bar{\mathbf{h}}_1, \dots, \bar{\mathbf{h}}_T$. However, the real work done by policy $\bar{\mathbf{h}}_1, \dots, \bar{\mathbf{h}}_T$ is at least the work done by $\bar{\mathbf{w}}_1, \dots, \bar{\mathbf{w}}_T$, which is $(1 - \varepsilon) \text{work}_{\mathbf{h}_1, \dots, \mathbf{h}_T}$.

Similarly, for SLA satisfaction, we prove a stronger bi-criteria result that implies the SLA guarantee in Theorem 1.

Theorem 4 *Let $0 < \varepsilon \leq 1/10$ and $0 < \eta \leq 1/3$. Take any SLAs $\beta(1), \dots, \beta(N)$ such that $\beta(i) \geq 2\frac{\varepsilon}{N}$ and let $\tilde{s} = 32\frac{N^2}{\varepsilon^3\eta} \ln(N/\varepsilon)$. Then, for any user i and time t , if we take $\mathbf{h}'_1, \dots, \mathbf{h}'_T$ to be allocations such that $h'_t(i) = (1 - 2\varepsilon)\beta(i)$, the work done by Algorithm 1 for user i satisfies*

$$\sum_{\tau=1}^{t+\tilde{s}} w_\tau(i) \geq \sum_{\tau=1}^t w'_\tau(i),$$

where \mathbf{w}'_t is the work done by the allocations $\mathbf{h}'_1, \dots, \mathbf{h}'_T$. Moreover, $\sum_{\tau=1}^t w_\tau(i) \geq \sum_{\tau=1}^t w'_\tau(i) - \beta(i)\tilde{s}$.

3.1 The Offline Formulation

Before presenting the proof of Theorem 3, we state the offline LP formulation of the maximum work problem for $(1 - \varepsilon)$ -allocations. We denote by $\mathbf{w}_t = (w_t(1), \dots, w_t(N))$ the work done for each user at time t . Given loads L_1, \dots, L_T , the offline formulation and its dual LP are given in Figure 2. As written, the dual LP includes a change of variable; see Appendix B for details. Constraints (1) state that the work done by the allocation for any user cannot be greater than the loads at all time. Constraints (2) limit the work performed at time t to at most $1 - \varepsilon$ of the resource. The LP

(D_ε) will be of special importance in the analysis. Using our algorithm, we will construct a dual feasible solution.

Observe that (P_ε) is feasible and bounded since the feasible region is a non-empty polytope. Let v_{P_ε} be the optimal value of (P_ε) . The following proposition gives a simple characterization of v_{P_ε} ; the proof appears in Appendix B. One way to prove this result is using the Max-Flow/Min-Cut theorem (see [35]) over a time-indexed network.

Proposition 2 $v_{P_\varepsilon} = \min_{0 \leq t \leq T} (\sum_{s=1}^t \sum_i L_s(i) + (1 - \varepsilon)(T - t))$.

3.2 Work Maximization

In this section we prove Theorem 3. Our first Lemma characterizes the implications of the update rule. The proof follows from a careful analysis of the dynamics using the KL-divergence and appears in Appendix B.

The first result of the lemma shows the behavior of active users' allocations when the system is under-utilized ($\leq 1 - \varepsilon$). In this case, all the active users receive a multiplicative boost in their allocation. The second result shows a more general behavior (see also Lemma 2). In this case, active users with allocation below their SLA do not decrease their allocations while the other active users might decrease their allocation, but in this case, the multiplicative penalization will be less severe.

Lemma 1 *Let $c = \frac{\varepsilon\eta}{4N}$. Then Algorithm 1 satisfies the following:*

1. *Suppose $\sum_{i \in A_t} h_t(i) \leq 1 - \varepsilon$. If $i \in A_t$, then $h_{t+1}(i) \geq h_t(i)(1 + c)$.*
2. *In general, Algorithm 1 satisfies $h_{t+1}(i) \geq h_t(i)$ for $i \in A_t^1$ and $h_{t+1}(i) \geq h_t(i)(1 - \varepsilon c)$ for $i \in A_t^2$.*

Proof of Theorem 3 Given loads $L_1, \dots, L_T \in \mathbf{R}_+^N$, consider the following $\{0, 1\}$ -matrix M of dimension $N \times T$ that encodes the information about the status of queues obtained while running Algorithm 1

$$M_{i,t} = \begin{cases} 0 & i\text{'s queue is empty at } t, Q_t(i) = 0, \\ 1 & i\text{'s queue is not empty at } t, Q_t(i) > 0. \end{cases}$$

Let $\tilde{s} = \frac{\ln(N/\varepsilon)}{\varepsilon c}$, where c is defined in Lemma 1. Now, pick s^* to be the maximum non-negative integer s (including 0) such that

$$\sum_{t=1}^s \sum_i L_t(i) \leq \sum_{t=1}^{s+\tilde{s}} \sum_i w_t(i) \quad (7)$$

Claim 3.1 *Consider any block of time $[r, r + \tilde{s}]$ where $r > s^*$; then there exists a user i such that $M_{i,r'} = 1$ for all $r' \in [r, r + \tilde{s}]$.*

Proof Suppose not. Then we claim that r satisfies condition (7). Consider any user i and let $r'_i \in [r, r + \tilde{s}]$ be such that $M_{i,r'_i} = 0$. Then work done by the user i up to time $[r + \tilde{s}]$ is at least

$$\sum_{t=1}^{r+\tilde{s}} w_t(i) \geq \sum_{t=1}^{r'_i} w_t(i) = \sum_{t=1}^{r'_i} L_t(i) \geq \sum_{t=1}^r L_t(i).$$

Now summing over all i , we get the desired contradiction. \square

We now prove the following claim that shows that the algorithm ensures that, on average, the total resource utilization after s^* is close to $1 - 4\varepsilon$. The proof of the Claim relies on Lemma 1 and appears in Appendix B.

Claim 3.2 *Let $B = [r, r + \tilde{s}]$ where $r > s^*$ be a consecutive block of \tilde{s} timesteps and let $B' = \{t \in B : \sum_{j \in A_t} h_t(j) \leq 1 - \varepsilon\}$ be the time steps in L with low utilization. Then $|B'| \leq 4\varepsilon\tilde{s}$ and therefore, $\sum_{t=s^*+1}^T \sum_i w_t(i) \geq (1 - 4\varepsilon)(T - s^*) - \tilde{s}$.*

Now, consider the following feasible dual solution of $(D_{4\varepsilon})$: $\gamma_t(i) = 1, \beta_t = 0$ for all users i and $t = 1, \dots, s^*$, and $\gamma_t(i) = 0, \beta_t = 1$ for all users i and $t = s^* + 1, \dots, T$. Observe that $\sum_{t=1}^T \beta_t = T - s^*$.

For optimal $(1 - 4\varepsilon)$ -allocations $\mathbf{h}_1^* \dots, \mathbf{h}_T^*$ we obtain

$$\begin{aligned}
\text{work}_{\mathbf{h}_1^*, \dots, \mathbf{h}_T^*} &\leq v_{\text{dual}}(\gamma_1, \dots, \gamma_T, \beta_1, \dots, \beta_T) && \text{(weak duality)} \\
&= \sum_{t=1}^{s^*} \sum_i L_t(i) + (1 - 4\varepsilon)(T - s^*) \\
&\leq \sum_{t=1}^{s^* + \tilde{s}} \sum_i w_t(i) + (1 - 4\varepsilon)(T - s^*) && (s^* \text{ choice}) \\
&\leq \sum_{t=1}^{s^*} \sum_i w_t(i) + \tilde{s} + \sum_{t=s^*+1}^T \sum_i w_t(i) + \tilde{s} && \text{(Claim 3.2)} \\
&= \text{work}_{\text{alg}} + 8 \frac{N}{\varepsilon^2 \eta} \ln(N/\varepsilon).
\end{aligned}$$

where we have used $\sum_{t=s^*+1}^{s^*+\tilde{s}} w_t(i) \leq \tilde{s}$ and the definition of \tilde{s} . □

3.3 SLA Satisfaction

In this section, we prove Theorem 4. Recall that $\lambda = \frac{\varepsilon^2}{8N}$ and $A_t^1 = \{i \in A_t : h_t(i) < \beta(i)\}$ is the set of active users receiving less than their SLAs and $A_t^2 = A_t \setminus A_t^1$ is the set of active user receiving at least their SLA. Analogous to Lemma 1, we have the following lemma, whose proof appears in Appendix B.

Lemma 2 *Assume $\varepsilon \leq \frac{1}{10}$, $\eta \leq \frac{1}{3}$ and $\beta(i) \geq 2 \frac{\varepsilon}{N}$ for all users. Then for any $i \in A_t^1$, Algorithm 1 guarantees $h_{t+1}(i) \geq h_t(i)(1 + c')$, where $c' = \frac{\varepsilon \eta \lambda}{2N}$.*

Proof of Theorem 4 Let $\tilde{s} = \left\lceil \frac{\ln(N/\varepsilon)}{\ln(1+c')} \right\rceil$, where c' is defined in Lemma 2. Now, we proceed by induction on t to prove that $\sum_{\tau=1}^{t+\tilde{s}} w_\tau(i) \geq \sum_{\tau=1}^t w'_\tau(i)$, where \mathbf{w}'_t is the work done by the allocations $\mathbf{h}'_1, \dots, \mathbf{h}'_T$. Clearly, the case $t = 0$ is direct.

Take $t \geq 1$ and suppose the result true for $t - 1$. If there exists $r \in [t, t + \tilde{s}]$ such that user i 's queue is empty, then

$$\sum_{\tau=1}^{t+\tilde{s}} w_\tau(i) \geq \sum_{\tau=1}^r w_\tau(i) = \sum_{\tau=1}^r L_\tau(i) \geq \sum_{\tau=1}^t w'_\tau(i).$$

Therefore, assume that for all $\tau \in [t, t + \tilde{s}]$ we have that user i 's queue is non-empty. By induction

hypothesis

$$\sum_{\tau=1}^{t-1+\tilde{s}} w_{\tau}(i) \geq \sum_{\tau=1}^{t-1} w'_{\tau}(i).$$

In order to complete the proof, we need to prove that $w_{t+\tilde{s}}(i) \geq w'_t(i)$. We proceed as follows. Suppose that for all $\tau \geq t$ we have $w_{\tau}(i) < (1 - \varepsilon)\beta(i)$. By Lemma 2, in each time $\tau \in [t, t + \tilde{s}]$, user's i allocation increases multiplicatively by a rate $(1 + c')$. Therefore,

$$w_{t+\tilde{s}}(i) \geq \frac{\varepsilon}{N}(1 + c')^{\tilde{s}} \geq 1 \geq \beta(i),$$

a contradiction. From the previous analysis we obtain the existence of $\tau^* \in [t, t + \tilde{s}]$ such that $w_{\tau^*}(i) \geq (1 - \varepsilon)\beta(i)$. By using Lemmas 1 and 2, we can show that the allocation $h_{\tau}(i)$ will never go below $(1 - \varepsilon c)(1 - \varepsilon)\beta(i)$ for all $\tau \geq \tau^*$. In particular $w_{t+\tilde{s}}(i) \geq (1 - \varepsilon c)(1 - \varepsilon)\beta(i) \geq (1 - 2\varepsilon)\beta(i) \geq w'_t(i)$. \square

3.4 Extension to Proportionality

In the previous subsections, we have introduced the first version of the multiplicative weights algorithm. We explained how we deduced our algorithm using mirror descent and proved its theoretical guarantees. Even though Algorithm 1 guarantees individual SLA satisfaction, this simple policy can lead to undesirable results that do not respect the ratio between allocations. If one user has an SLA twice the size of another, it would be reasonable for the former to expect allocations at least twice as big as the latter's if both are consistently busy. Likewise, the second user would expect allocations no less than half of the first user's. In other words, both users should expect shares that respect the ratio between their SLAs.

To illustrate this unsatisfactory behavior in Algorithm 1, we run it with three users having SLAs $\beta(1) = 0.5$, $\beta(2) = 0.3$ and $\beta(3) = 0.2$. We set $\eta \leq 1/3$ and $\varepsilon \leq 1/10$. For simplicity, the initial allocation will be uniform. In our example, user 1 is always idle. User 2 consistently demands 1 unit of resource. User 3 begins idle and remains so until user 2's allocation reaches $1 - \varepsilon$. This takes roughly $\frac{1}{\eta\varepsilon}$ time steps; call this time t_0 . Starting at time t_0 , user 3 demands unit loads every time step for the rest of the horizon. Initially, the allocations are uniformly $1/3$ for everyone. Between

time 1 and t_0 , user 2's allocation increases until it hits $1 - \varepsilon$, since they are the only active user. After t_0 , user 3 becomes active, and has an allocation below their SLA. Therefore, the algorithm redistributes allocation from user 2 to 3 until user 3's allocation hits 0.2. After this, allocations remain stable at approximately $h_t(1) = \frac{\varepsilon}{3}$, $h_t(2) = 0.8 - \frac{\varepsilon}{3}$ and $h_t(3) = 0.2$. User 2 receives about 4 times the allocation of user 3 if ε is small enough. However, a better allocation for user 2 and 3 is $\frac{\beta(2)}{\beta(2)+\beta(3)} = \frac{3}{5}$ and $\frac{\beta(3)}{\beta(2)+\beta(3)} = \frac{2}{5}$ respectively. These allocations reflect the ratio $\frac{\beta(2)}{\beta(3)}$ between active users.

Given this, we propose a slight modification of Algorithm 1, shown in Algorithm 2. As before, the plan is always to benefit active users. However, this time, we boost active users slightly more if they fall behind their ‘‘proportional SLA’’ among active users. Intuitively, if there are $n < N$ active users for a long period of time, the allocation of these active users should converge to their proportional share.

Algorithm 2: Extended Multiplicative Weight Update Algorithm

Input: Parameters $0 < \varepsilon \leq \frac{1}{10}, 0 < \eta < \frac{1}{3}$.

- 1 Initialization: \mathbf{h}_1 any allocation over Δ_ε and $\lambda = \frac{\varepsilon^2}{8N}$.
 - 2 **for** $t = 1, \dots, T$ **do**
 - 3 Set allocation \mathbf{h}_t .
 - 4 Read active and inactive users A_t and B_t . $A_t^1 = \left\{ i \in A_t : h_t(i) < (1 - \varepsilon) \frac{\beta(i)}{\sum_{j \in A_t} \beta(j)} \right\}$,
 $A_t^2 = A_t \setminus A_t^1$.
 - 5 Set gain function $g_t(i) = \begin{cases} 1 + \lambda & i \in A_t^1 \\ 1 & i \in A_t^2 \\ 0 & i \in B_t \end{cases}$.
 - 6 Update allocation: $\hat{h}_{t+1}(i) = h_t(i)e^{\eta g_t(i)}, \forall i$ and $\mathbf{h}_{t+1} = \pi_{\Delta_\varepsilon}(\hat{h}_{t+1})$
-

For technical reasons, the set A_t^1 , the active users with allocation below their proportional share among active users at time t , has to be defined as $\left\{ i \in A_t : h_t(i) < (1 - \varepsilon) \frac{\beta(i)}{\sum_{j \in A_t} \beta(j)} \right\}$. The reason behind this choice is to ensure that if $A_t^1 \neq \emptyset$ and the resource is nearly fully utilized, i.e., $\sum_{i \in A_t} h_t(i) \geq 1 - \varepsilon$, then there is a different active user $j \neq i$ from which we can move allocation to i . This is fundamental in the proof of Theorem 6 below.

In terms of work maximization and SLA satisfaction, Algorithm 2 provides exactly the same guarantees as Algorithm 1.

Theorem 5 *Given loads L_1, \dots, L_T , for any $\varepsilon > 0$ and $\eta > 0$ such that $\varepsilon \leq 1/10$, Algorithm 2 guarantees*

$$\text{work}_{\mathbf{h}_1^*, \dots, \mathbf{h}_T^*} - \text{work}_{\text{alg}} \leq 8 \frac{N}{\varepsilon^2 \eta} \ln(N/\varepsilon),$$

where $\mathbf{h}_1^*, \dots, \mathbf{h}_T^*$ is the optimal offline sequence of $(1 - 4\varepsilon)$ -allocations.

The proof of Theorem 5 is exactly the same as the proof of Theorem 3. To see this, observe that the proof of Theorem 3 only uses the fact that the allocations of *every* active users get a multiplicative boost whenever the usage is below $1 - \varepsilon$. The last statement is true since Lemma 1 also holds in this case.

Theorem 6 *Let $0 < \varepsilon \leq 1/10$ and $0 < \eta \leq 1/3$. Take any SLAs $\beta(1), \dots, \beta(N)$ such that $\beta(i) \geq 2\frac{\varepsilon}{N}$ and let $\tilde{s} = 32\frac{N^2}{\varepsilon^3 \eta} \ln(N/\varepsilon)$. Then, for any user i and time t , if we take $\mathbf{h}'_1, \dots, \mathbf{h}'_T$ to be any allocations such that $h'_t(i) = (1 - 2\varepsilon)\beta(i)$, the work done by Algorithm 1 for user i satisfies*

$$\sum_{\tau=1}^{t+\tilde{s}} w_{\tau}(i) \geq \sum_{\tau=1}^t w'_{\tau}(i),$$

where w'_t is the work done by the allocations $\mathbf{h}'_1, \dots, \mathbf{h}'_T$. Moreover $\sum_{\tau=1}^t w_{\tau}(i) \geq \sum_{\tau=1}^t w'_{\tau}(i) - \beta(i)\tilde{s}$.

The proof of this result is similar to the proof of Theorem 4. A subtle difference is that the analogue of Lemma 2 holds if we add the hypothesis $\sum_{i \in A_t} h_t(i) > 1 - \varepsilon$. We skip the proof for brevity.

Lemma 3 *Assume $\varepsilon \leq \frac{1}{10}$, $\eta \leq \frac{1}{3}$ and $\beta(i) \geq 2\frac{\varepsilon}{N}$ for all users. In Algorithm 2, if $\sum_{k \in A_t} h_t(k) > 1 - \varepsilon$ then for any $i \in A_t^1$ we have $h_{t+1}(i) \geq (1 + c')h_t(i)$, where $c' = \frac{\varepsilon \eta \lambda}{2N}$.*

Lemma 1 and 3 ensure that any active user gets a multiplicative boost of at least $(1 + c')$. Therefore, any user that is active \tilde{s} consecutive times will have an allocation of at least $(1 - 2\varepsilon)\beta(i)$. Then, by following the same inductive proof of Theorem 4, we obtain Theorem 6.

4 Experiments

In this section, we empirically test Algorithm 2 against a family of offline and online algorithms. We aim to measure the performance on both synthetic data as well as real CPU traces from a production service in Microsoft’s cloud. We quantify the performance on the following three criteria.

- **Work maximization.** We compare the overall work done by Algorithm 2 against various benchmark algorithms.
- **SLA guarantee.** We examine the extent to which our algorithm achieves the cumulative work of the static SLA policy for each user. We do so by measuring the cumulative work over plausible time windows.
- **Queue length.** We compare the 2-norm of the individual queues over time. We use this metric as a proxy for the system latency, which is not captured by our theoretical results.

We consider the following *online algorithms*, against which we benchmark our algorithm.

- **Static SLA Policy (Static).** Each user gets their SLA as a constant, static allocation. We call this algorithm Static.
- **Proportional Online (PO).** In each iteration, every active user will get their SLA normalized by the sum of SLAs of active users (just their SLA if there are no active users). This simple algorithm seems suitable for a practical implementation; however, its performance can be arbitrarily bad in terms of work maximization. The formal description appears in Algorithm 4 in Appendix E. We call this Algorithm PO.
- **Online Work Maximizing (OWM).** This algorithm divides users into three categories: A , B and I (active users with allocation, active users without allocation and inactive users). At each iteration, the resource is divided uniformly among users in A . If a user in A becomes inactive, they are moved to I . If a user in I becomes active, they are moved to B . When A becomes empty, we move all users from B to A . In Appendix C we prove that this method is work maximizing. However, this greedy strategy is not guaranteed to satisfy SLA constraints for general input loads. We call this algorithm OWM.

We also consider the following *offline algorithms*, against which we benchmark our algorithm.

- * **Optimal 1-allocations (PG)**. The optimal offline solution to the work maximization problem. This solution is computed using Algorithm 3, which we call Proportional Greedy (PG). This algorithm can be considered as the offline counterpart of Proportional Online.
- * **Optimal $(1 - \varepsilon)$ -allocation (restPG)**. Offline solution to the work maximization problem with resource restricted to $1 - \varepsilon$, where ε is the parameter of Algorithm 2.

4.1 Synthetic Experiment

4.1.1 Description of the Experiment

In this experiment, we consider a synthetically generated input sequence, which we use to examine how online algorithms adapt to different conditions. Specifically, our system consists of three users with SLAs of $\beta(1) = 0.2$, $\beta(2) = 0.3$ and $\beta(3) = 0.5$. We consider a time horizon of $T = 3,000,000$. The load input sequence is divided into six periods: $P_i = \left[\frac{(i-1)T}{6}, \frac{iT}{6} \right)$ for $i = 1, \dots, 6$. In each period, only two users demand new resources. During the first 3 periods, the random demand has a mean proportional to the users' SLA. In the following 3 periods, the random demand changes to a distribution with uniform mean among users demanding resources. Specifically:

- During P_1 , only users 2 and 3 demand the following loads. At the beginning of P_1 , i.e., $t = 1$, user 2 demands a large load of $L_1(2) \sim \frac{T}{6} \cdot \text{Gamma} \left(2000, \frac{1}{2000} \cdot \frac{\beta(2)}{\beta(2)+\beta(3)} \right)$ and $L_1(3) = 0$. During the rest of period P_1 , $L_t(2) = 0$ and $L_t(3) \sim \text{Gamma} \left(2000, \frac{1}{2000} \cdot \frac{\beta(3)}{\beta(2)+\beta(3)} \right)$. User 1 demands nothing during this entire period. Similar loads are set for period P_2 and P_3 .
- Similarly, during P_4 users 2 and 3 demand $L_t(i) \sim \text{Gamma} \left(2000, \frac{1}{2000} \cdot \frac{1}{2} \right)$ for $i = 2, 3$. During P_5 users 1 and 2 demand $L_t(i) \sim \text{Gamma} \left(2000, \frac{1}{2000} \cdot \frac{1}{2} \right)$ for $i = 1, 2$. During P_6 users 1 and 3 demand $L_t(i) \sim \text{Gamma} \left(2000, \frac{1}{2000} \cdot \frac{1}{2} \right)$ for $i = 1, 3$.

The expectation of a $\text{Gamma}(k, \theta)$ random variable is given by $k\theta$ and the variance is given by $k\theta^2$ (see e.g. [36]). For example, in period P_1 user 2's expected load is $\frac{T}{6} \frac{\beta(2)}{\beta(2)+\beta(3)}$, with variance $\frac{1}{2000} \left(\frac{\beta(2)}{\beta(2)+\beta(3)} \right)^2$. Similarly, user 3's expected total load is $\frac{T}{6} \frac{\beta(3)}{\beta(2)+\beta(3)}$. Thus, the expected overall

load is $T/6$, exactly the length of the period. A brief summary of Gamma distribution’s properties is given in Appendix F.

We instantiate Algorithm 2 with $\eta = \frac{1}{3}$, $\varepsilon = 0.02$ and $T = 3,000,000$.

4.1.2 Results

Work maximization. In Figure 3 we present the cumulative work difference between PG and Algorithm 2 (solid blue line with star), restPG and Algorithm 2 (red dashed line with triangle), PO and Algorithm 2 (solid magenta line), Static and Algorithm 2 (solid green line with small circle) and OWM and Algorithm 2 (solid cyan line with large circle). Intuitively, one positive unit of difference implies the corresponding algorithm is ahead of Algorithm 2 by one unit of time.

First, we consider the comparison to online algorithms Static, PO and OWM. Algorithm 2 outperforms Static significantly, by roughly 700,000 units of time. During the first half of the experiment PO shows good performance, but in the second half of the experiment (when the load distribution changes), Algorithm 2 outperforms PO. This shows that Algorithm 2 can adapt to changing input sequences that PO cannot adapt to. Finally, OWM surpasses Algorithm 2 during the whole experiment, with a performance similar to PG; this is an expected result since OWM is work maximizing.

With respect to offline algorithms, PG outperforms Algorithm 2 by roughly 10,000 time units. On the other hand, Algorithm 2 surpasses restPG by approximately 20,000 units. This shows that Algorithm 2 indeed performs better than the offline optimum with a slightly reduced amount of resources, as guaranteed by the theoretical results.

SLA satisfaction. Among online algorithms, Static and PO satisfy SLA restrictions by design, but as seen earlier they are not competitive in terms of work maximization. We focus on the comparison between OWM and Algorithm 2 as far as SLA satisfaction is concerned. Although OWM performs extremely well in work maximization, this comes at a significant price in SLA satisfaction. In Figure 4 we depict empirically this behavior by plotting the instantaneous work done by users 2 and 3 by OWM and Algorithm 2 during period P_1 . We empirically observe that Algorithm 2 approximately satisfies user 3’s SLA, but OWM does not allocate the user any resources. Such

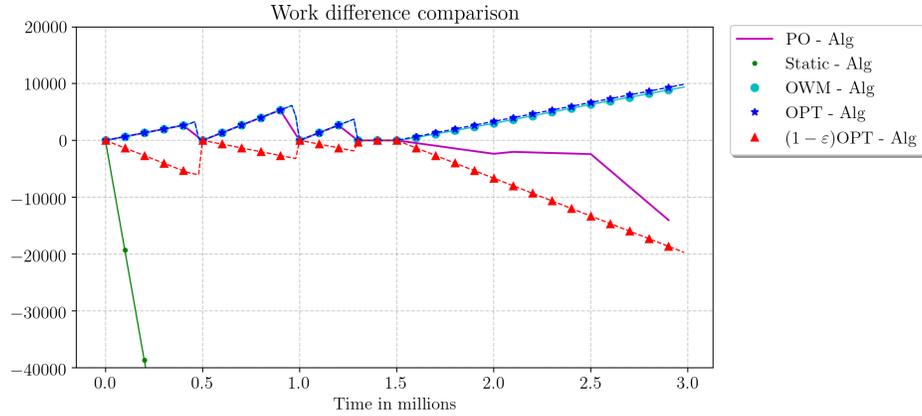


Figure 3: Difference between algorithms’ work. Alg is short for Algorithm 2. Observe that the differences “OPT - Alg” and “OWM - Alg” slightly overlap.

an extreme behavior arises because OWM is geared towards work maximization rather than SLA satisfaction.

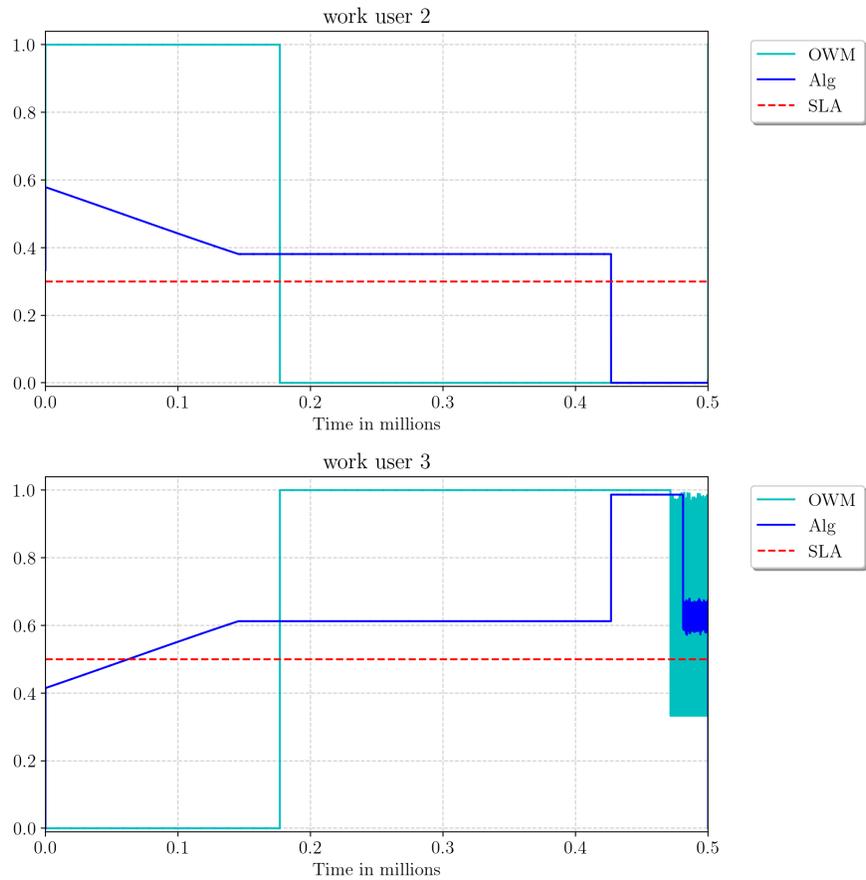


Figure 4: Instantaneous work for user 2 and 3 during period P_1 . User 3 does not receive any allocation in OWM until user 2 finishes all of their work.

Queue lengths. In Figure 5 we present the 2-norm of queues induced by Algorithm 2 (solid blue), Static (solid magenta with star), PO (solid cyan with large circle) and OWM (solid black with triangle). Once again, we can interpret one unit of norm as one unit of latency. Experimentally, we observe that Static shows the worst performance with a final 2-norm of 524,000 units. Algorithm 2 ends with a 2-norm of 10,000 units, PO with 26,970 units, and OWM with 381 units. As remarked above, even though OWM induces very small queues, this comes at the cost of not satisfying SLA requirements.

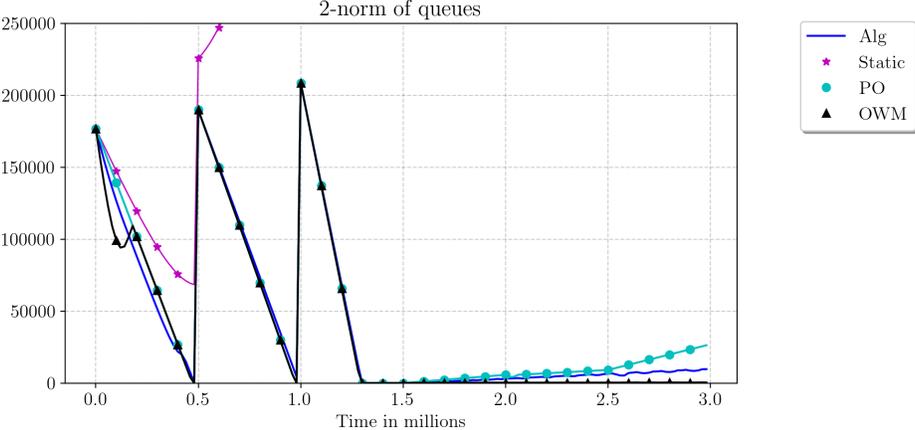


Figure 5: 2-norm of queues.

Summary. Among all online algorithms, Algorithm 2 is able to best balance work maximization and SLA satisfaction. In particular, Algorithm 2 is only slightly worse in terms of work maximization compared to OWM; this is expected since Algorithm 2 always reserves a small fraction of the resource for each user, regardless of activity. Furthermore, our results show that the actual total work done by Algorithm 2 is much better than the theoretical guarantee of Theorem 5, as it substantially outperforms $(1 - \epsilon)\text{OPT}$ (restPG). Finally, we observe small queues throughout the entire horizon, which is key for maintaining reasonable latencies.

4.2 Experiment with Real Data

We next describe the results of our computational study using real world data. For these experiments, we obtained CPU traces of a production service on Azure, Microsoft’s public cloud. The data consists of demand traces of six different users over a time window of approximately ten days.

To show Algorithm 2’s robustness with respect to (short-term) real data, we also consider the following measurement:

- **Instantaneous SLA.** We focus on a modified SLA satisfaction criterion because of the relatively short horizon (about 14,000 minutes); we assess Algorithm 2’s performance in the following way. For a user i and any time t , we compare the cumulative work done by user i in Algorithm 2 during a time window $[t, t + \tau)$ versus the cumulative work done by the same user i under a Static SLA Policy during the time window $[t, t + \tau)$. In order to have a meaningful comparison, at time t , we run the Static SLA Policy with the queues of Algorithm 2 at time t . The motivating question is, *what happens if at time t and the next τ time steps we run the static policy instead of Algorithm 2?* For this experiment, we used $\tau = 500$ minutes.

The data set consists of demand traces of six users of exactly 14,628 minutes (approximately ten days). Each user is assigned their normalized average workload as SLA. (Since the data is proprietary, we cannot disclose actual SLAs.) For the purpose of the experiment, we run Algorithm 2 with parameters $\varepsilon = 0.01$ and $\eta = \frac{1}{3}$.

4.2.1 Results

Work maximization. We depict in Figure 6 the following differences: cumulative work done until time t by optimal 1-allocations (PG) and Algorithm 2, $(1 - \varepsilon)$ -allocations (restPG) and Algorithm 2, PO and Algorithm 2, Static and Algorithm 2, and OWM and Algorithm 2. In a similar fashion to the previous experiment, one positive unit can be interpreted as Algorithm 2 being one unit (minute) of work behind, and one negative unit means Algorithm 2 is ahead by a minute.

We observe that Algorithm 2 outperforms all online benchmarks. Against Static, the final difference is 105 units, with a maximum difference of 167 units. Against PO the final difference is 25 units with a maximum difference of 37. Finally, against OWM, the final difference is 20 with a maximum difference of 21. Surprisingly, for this data set Algorithm 2 is able to surpass even OWM.

Regarding the offline algorithms, PG surpasses Algorithm 2 during the whole experiment as expected, with a final difference of 14 units. On the other hand, Algorithm 2 outperforms restPG by 26 units by the end of the experiment.

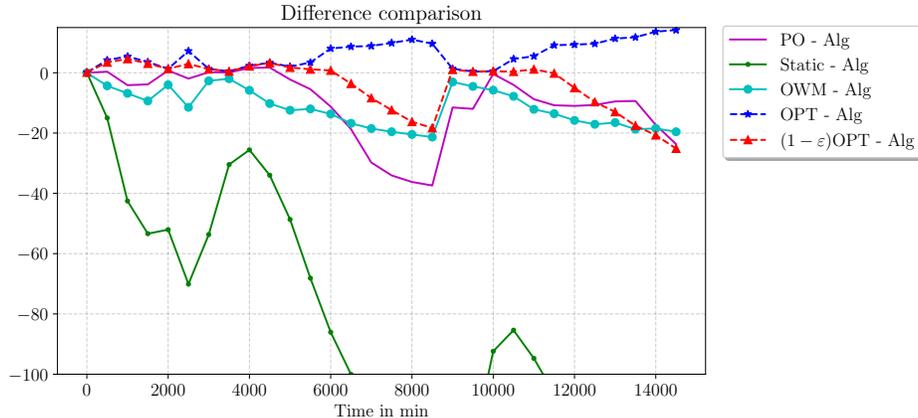


Figure 6: Difference of cumulative works.

Instantaneous SLA. In Table 1 we report statistics on the differences between the cumulative work done in time windows $[t, t + \tau)$ by Static and Algorithm 2 for each user. For each t and each user i , the exact formula is $r_i(t) = \sum_{r=t}^{t+\tau-1} w'_r(i) - \sum_{r=t}^{t+\tau-1} w_r(i)$, where $w_r(i)$ is the work done by user i under Algorithm 2 and $w'_r(i)$ is the work done by user i with Static (with queues at t given by Algorithm 2). The table shows the min, max, average and standard deviation of $\{r_i(t)\}_t$ when $\tau = 500$ minutes. A positive unit means Algorithm 2 is outperformed by Static during $[t, t + \tau)$ under the same initial conditions by one unit of time. In general, we observe that all users show negative empirical average difference. This result empirically suggests that Algorithm 2 ensures approximate SLA satisfaction, even for small time windows. For instance, user 3 occasionally has a high difference (46.4 units), mostly due to times t where Algorithm 2 allocates the user a small amount of resource but a huge load is incoming during the window $[t, t + \tau)$. The experiment tells us that averaging out these “bad” times ensures good performance under the SLA criterion. Furthermore, we tested values of $\tau = 60, 500$ and 1000 minutes; larger windows improve the results, with lower maximum and average values.

Queue lengths. In Figure 7 we present the 2-norms of queues given by the online benchmark algorithms and Algorithm 2. As usual, we can interpret one unit as the respective algorithm’s latency, that is, lateness with respect to the overall users’ demand. Compared against the online algorithms, we empirically observe the superiority of Algorithm 2, as it has the smallest latency most of the time. Algorithm 2 ends with a 2-norm of roughly 44 units, average length of 22 and a maximum length of 92. PO ends with a 2-norm of approximately 68, average length of 32 and

Table 1: Statistics for the difference between the cumulative works of our algorithm and static over time windows $[t, t + \tau)$. For each user i we present the min, max, average, and standard deviation over the sequence of differences $\{r_i(t)\}_t$.

| User | min | max | mean | std |
|--------|--------|------|-------|------|
| User 1 | -31.1 | 21.7 | -4.6 | 13.8 |
| User 2 | -122.3 | 46.9 | -31.2 | 49.2 |
| User 3 | -90.8 | 46.4 | -7.2 | 29.5 |
| User 4 | -49.7 | 18.5 | -0.8 | 12.9 |
| User 5 | -42.6 | 22.6 | -5.85 | 14.0 |
| User 6 | -21.9 | 14.9 | -0.6 | 6.4 |

a maximum length of 126. OWM ends with a 2-norm of 68, average of 32 and maximum of 113. Finally, Static shows the worst behavior, with a final 2-norm of 103, average of 91 and maximum of 231. For this data set, Algorithm 2 shows a remarkable performance, considering particularly that Algorithm 2 always reserves ε/N resource for each user.

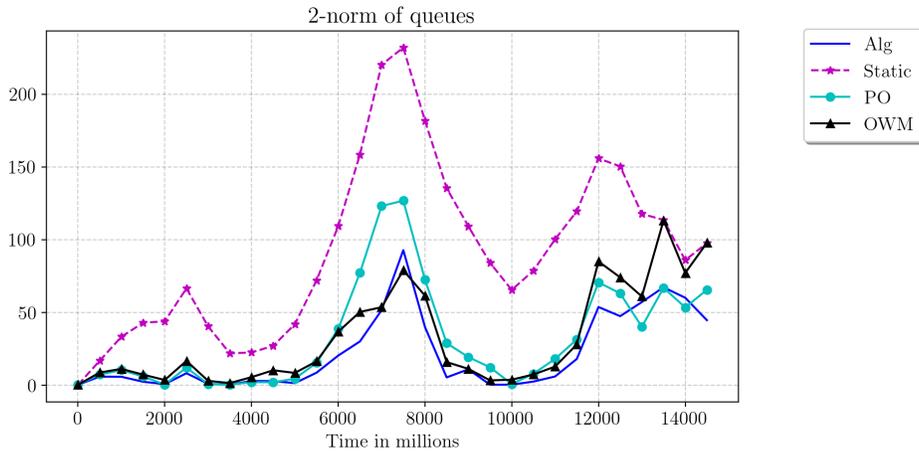


Figure 7: 2-norm of queues.

Summary. Algorithm 2 performs very well in terms of work maximization compared to all other online algorithms, and exceeds the theoretical guarantees. Furthermore, SLA requirements are typically satisfied, even when measured over relatively short time windows. Finally, as in the previous experiment, the algorithm maintains small queues compared to other online algorithms.

5 Conclusion

We have proposed a new online model for dynamic resource allocation of a single divisible resource in a shared system. Our framework captures basic properties of cloud systems, including SLAs, limited system feedback and unpredictable (even adversarial) input sequences. We designed an algorithm that is near-optimal in terms of both work maximization and SLA satisfaction (Theorems 1, 5 and 6). We derived a simple expression for the offline work maximization problem that allowed us to reinterpret the algorithm’s dynamics as an approximate solution of the optimal (offline) work maximization LP. Numerical experiments show that our algorithm is indeed able to achieve a desirable trade-off between work maximization and SLA satisfaction. In particular, comparisons with offline algorithms (PG and restPG) indicate that our algorithm is empirically work maximizing. Further, unlike other plausible online algorithms, our algorithm is able to quickly adapt to unexpected changes in demand and still approximately satisfy the underlying user SLAs. Our model and results may be extended in various directions of interest to the operations research and cloud computing communities. Some potential future directions include exploiting predictions about future workloads (as in [37]), and simultaneously allocating multiple resources.

References

- [1] V. K. Vavilapalli, A. C. Murthy, C. Douglas, S. Agarwal, M. Konar, R. Evans, T. Graves, J. Lowe, H. Shah, S. Seth, *et al.*, “Apache hadoop yarn: Yet another resource negotiator,” in *Proceedings of the 4th annual Symposium on Cloud Computing*, p. 5, ACM, 2013.
- [2] B. Hindman, A. Konwinski, M. Zaharia, A. Ghodsi, A. D. Joseph, R. H. Katz, S. Shenker, and I. Stoica, “Mesos: A platform for fine-grained resource sharing in the data center.,” in *NSDI*, vol. 11, pp. 22–22, 2011.
- [3] M. Zaharia, D. Borthakur, J. Sen Sarma, K. Elmeleegy, S. Shenker, and I. Stoica, “Delay scheduling: a simple technique for achieving locality and fairness in cluster scheduling,” in *Proceedings of the 5th European conference on Computer systems*, pp. 265–278, ACM, 2010.
- [4] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica, “Dominant

- resource fairness: Fair allocation of multiple resource types.,” in *Nsdi*, vol. 11, pp. 24–24, 2011.
- [5] C. Curino, D. E. Difallah, C. Douglas, S. Krishnan, R. Ramakrishnan, and S. Rao, “Reservation-based scheduling: If you’re late don’t blame us!,” in *Proceedings of the ACM Symposium on Cloud Computing*, pp. 1–14, ACM, 2014.
- [6] S. A. Jyothi, C. Curino, I. Menache, S. M. Narayanamurthy, A. Tumanov, J. Yaniv, R. Mavlyutov, I. Goiri, S. Krishnan, J. Kulkarni, *et al.*, “Morpheus: Towards automated slos for enterprise clusters.,” in *OSDI*, pp. 117–134, 2016.
- [7] R. Grandl, M. Chowdhury, A. Akella, and G. Ananthanarayanan, “Altruistic scheduling in multi-resource clusters.,” in *OSDI*, pp. 65–80, 2016.
- [8] V. Narasayya, S. Das, M. Syamala, B. Chandramouli, and S. Chaudhuri, “SQLVM: Performance isolation in multi-tenant relational database-as-a-service,” 2013.
- [9] J. Rasley, K. Karanasos, S. Kandula, R. Fonseca, M. Vojnovic, and S. Rao, “Efficient queue management for cluster scheduling,” in *Proceedings of the Eleventh European Conference on Computer Systems*, p. 36, ACM, 2016.
- [10] S. Arora, E. Hazan, and S. Kale, “The multiplicative weights update method: a meta-algorithm and applications.,” *Theory of Computing*, vol. 8, no. 1, pp. 121–164, 2012.
- [11] E. Hazan *et al.*, “Introduction to online convex optimization,” *Foundations and Trends® in Optimization*, vol. 2, no. 3-4, pp. 157–325, 2016.
- [12] S. Shalev-Shwartz *et al.*, “Online learning and online convex optimization,” *Foundations and Trends® in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2012.
- [13] A. Ben-Tal and A. Nemirovski, *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*, vol. 2. Siam, 2001.
- [14] V. Narasayya, I. Menache, M. Singh, F. Li, M. Syamala, and S. Chaudhuri, “Sharing buffer pool memory in multi-tenant relational database-as-a-service,” *Proceedings of the VLDB Endowment*, vol. 8, no. 7, pp. 726–737, 2015.

- [15] I. Menache and M. Singh, “Online caching with convex costs,” in *Proceedings of the 27th ACM symposium on Parallelism in Algorithms and Architectures*, pp. 46–54, ACM, 2015.
- [16] S. T. Maguluri, R. Srikant, and L. Ying, “Stochastic models of load balancing and scheduling in cloud computing clusters,” in *2012 Proceedings IEEE Infocom*, pp. 702–710, IEEE, 2012.
- [17] S. T. Maguluri, R. Srikant, and L. Ying, “Heavy traffic optimal resource allocation algorithms for cloud computing clusters,” *Performance Evaluation*, vol. 81, pp. 20–39, 2014.
- [18] S. T. Maguluri and R. Srikant, “Scheduling jobs with unknown duration in clouds,” *IEEE/ACM Transactions On Networking*, vol. 22, no. 6, pp. 1938–1951, 2014.
- [19] S. Albers, “Online algorithms: a survey,” *Mathematical Programming*, vol. 97, no. 1-2, pp. 3–26, 2003.
- [20] A. Borodin and R. El-Yaniv, *Online computation and competitive analysis*. cambridge university press, 2005.
- [21] R. Srikant and L. Ying, *Communication networks: an optimization, control, and stochastic networks perspective*. Cambridge University Press, 2013.
- [22] M. Macías and J. Guitart, “A genetic model for pricing in cloud computing markets,” in *Proceedings of the 2011 ACM Symposium on Applied Computing*, pp. 113–118, ACM, 2011.
- [23] A. Gera and C. H. Xia, “Learning curves and stochastic models for pricing and provisioning cloud computing services,” *Service Science*, vol. 3, no. 1, pp. 99–109, 2011.
- [24] M. Passacantando, D. Ardagna, and A. Savi, “Service provisioning problem in cloud and multi-cloud systems,” *INFORMS Journal on Computing*, vol. 28, no. 2, pp. 265–277, 2016.
- [25] B. Sharma, R. K. Thulasiram, P. Thulasiraman, S. K. Garg, and R. Buyya, “Pricing cloud compute commodities: A novel financial economic model,” in *Proceedings of the 2012 12th IEEE/ACM international symposium on cluster, cloud and grid computing (ccgrid 2012)*, pp. 451–457, IEEE Computer Society, 2012.
- [26] M. C. Cohen, P. W. Keller, V. Mirrokni, and M. Zadimoghaddam, “Overcommitment in cloud services: Bin packing with chance constraints,” *Management Science*, 2019.

- [27] A. Gordon, M. Hines, D. Da Silva, M. Ben-Yehuda, M. Silva, and G. Lizarraga, “Ginkgo: Automated, application-driven memory overcommitment for cloud computing,” *Proc. RESoLVE*, 2011.
- [28] M. Dabbagh, B. Hamdaoui, M. Guizani, and A. Rayes, “Efficient datacenter resource utilization through cloud resource overcommitment,” in *2015 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 330–335, IEEE, 2015.
- [29] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of machine learning*. MIT press, 2018.
- [30] S. A. Plotkin, D. B. Shmoys, and É. Tardos, “Fast approximation algorithms for fractional packing and covering problems,” *Mathematics of Operations Research*, vol. 20, no. 2, pp. 257–301, 1995.
- [31] N. Buchbinder, J. S. Naor, *et al.*, “The design of competitive online algorithms via a primal–dual approach,” *Foundations and Trends® in Theoretical Computer Science*, vol. 3, no. 2–3, pp. 93–263, 2009.
- [32] Y. Freund and R. E. Schapire, “A decision-theoretic generalization of on-line learning and an application to boosting,” *Journal of computer and system sciences*, vol. 55, no. 1, pp. 119–139, 1997.
- [33] S. Bubeck, “Introduction to online optimization,” *Lecture Notes*, pp. 1–86, 2011.
- [34] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [35] A. Schrijver, *Combinatorial optimization: polyhedra and efficiency*, vol. 24. Springer Science & Business Media, 2003.
- [36] W. Feller, “An introduction to probability theory and its applications,” 1957.
- [37] T. Lykouris and S. Vassilvitskii, “Competitive caching with machine learned advice,” *arXiv preprint arXiv:1802.05399*, 2018.

A Projecting on Δ_ε

Proof Proof of Proposition 1 Let $\mathbf{y} \in \mathbf{R}_+^N$. The projection of \mathbf{y} on Δ_ε corresponds to the solution of the convex problem

$$(Q) \quad \begin{aligned} \min \quad & \sum_i x(i) \ln \left(\frac{x(i)}{y(i)} \right) \\ & \sum_i x(i) = 1 \\ & x(i) \geq \varepsilon/N \end{aligned}$$

Its lagrangian (see [34]) is

$$\mathcal{L}(\mathbf{x}, \lambda, \mu) = \sum_i x(i) \ln \left(\frac{x(i)}{y(i)} \right) - \lambda \left(\sum_i x(i) - 1 \right) - \sum_i \mu_i \left(x(i) - \frac{\varepsilon}{N} \right).$$

Using the FO conditions:

$$\forall i : x(i) = y(i)e^{\mu_i + \lambda - 1} = y(i)e^{\mu_i} C.$$

and the SO conditions:

$$\mu_i \geq 0, \forall i, \quad \text{and} \quad x(i) > \frac{\varepsilon}{N} \implies \mu_i = 0.$$

Let $S = \{i : x(i) = \varepsilon/N\}$ and $T = [N] \setminus S$. Then, using $\sum_i x(i) = 1$ we obtain

$$e^{\lambda-1} = \frac{1 - \frac{\varepsilon}{N}|S|}{\sum_{i \in T} y(i)}.$$

This proves part (b). Now, suppose we have $y(1) \leq \dots \leq y(N)$. If $i, j \in T$, then $x(i) = y(i)e^{\lambda-1}$ and $x(j) = y(j)e^{\lambda-1}$ and then

$$x(i) \leq x(j) \iff y(i) \leq y(j).$$

That is, in T the variables preserve their ordering.

If $i \in S$ and $j \in T$, then $y(i)e^{\lambda-1+\mu_i} = x(i) = \frac{\varepsilon}{N} < x(j) = y(j)e^{\lambda-1}$, which implies $y(i) < y(j)$ using that $\mu_i \geq 0$. Now, let $k = \min\{i \in T\}$ which is a well-defined number using constraint

$\sum_{i=1}^N x(i) = 1$. We claim that for any $j \geq k$, $j \in T$, that is, T corresponds to the interval $[k, N]$. By contradiction, suppose that $j > k$ does not belong to T , then $y(j) < y(k)$ by previous calculus. However $y(j) \geq y(k)$ by the ordering of \mathbf{y} . A contradiction. With this, the algorithm to project is clear, we sort \mathbf{y} and then we test increasingly the possible set $S = \{1, \dots, k-1\}$ for $k = 1, \dots, N$ and select the best candidate. This proves (a).

We say that S is *feasible* if there is a feasible solution \mathbf{x} such that $S = \{i : x(i) = \varepsilon/N\}$. In the following paragraphs we prove that the first feasible solution found in this process is the right one.

Observe that once $S = \{1, \dots, k\}$ is feasible, then $S' = \{1, \dots, j\}$ remains feasible for all $j \geq k$. Indeed, if $S = \{1, \dots, k\}$ is feasible, then

$$1 = \frac{\varepsilon}{N}k + \sum_{i \in T} x(i).$$

Now, increasing S to $S' = \{1, \dots, k+1\}$ means that we pick $x(k+1) > \frac{\varepsilon}{N}$ and we decrease it to $\frac{\varepsilon}{N}$. Therefore, $x(k+2), \dots, x(N)$ must increase. Therefore, S' remains feasible. The proof for general case $j \geq k$ follows by induction.

Now, we claim that if $S = \{1, \dots, k\}$ is feasible, then $S' = \{1, \dots, k+1\}$ cannot have better optimal value. Indeed, the difference between the objective S' and S is

$$\begin{aligned} & \sum_{i=1}^{k+1} \frac{\varepsilon}{N} \ln \frac{\varepsilon}{Ny(i)} + \left(1 - \frac{\varepsilon}{N}(k+1)\right) \ln \frac{1 - \frac{\varepsilon}{N}(k+1)}{\sum_{i \geq k+2} y(i)} - \sum_{i=1}^k \frac{\varepsilon}{N} \ln \frac{\varepsilon}{Ny(i)} - \left(1 - \frac{\varepsilon}{N}k\right) \ln \frac{1 - \frac{\varepsilon}{N}k}{\sum_{i \geq k+1} y(i)} \\ &= \frac{\varepsilon}{N} \ln \frac{\varepsilon}{Ny(k+1)} + \left(1 - \frac{\varepsilon}{N}(k+1)\right) \ln \frac{1 - \frac{\varepsilon}{N}(k+1)}{\sum_{i \geq k+2} y(i)} - \left(1 - \frac{\varepsilon}{N}k\right) \ln \frac{1 - \frac{\varepsilon}{N}k}{\sum_{i \geq k+1} y(i)} \end{aligned}$$

The function $f(x) = x \ln x$ is convex for $x > 0$. Now, pick $x = \frac{\varepsilon}{Ny(k+1)}$, $y = \frac{1 - \frac{\varepsilon}{N}(k+1)}{\sum_{i \geq k+2} y(i)}$ and $\lambda = \frac{y(k+1)}{\sum_{i \geq k+1} y(i)}$. Then

$$\lambda x + (1 - \lambda)y = \frac{y(k+1)}{\sum_{i \geq k+1} y(i)} \left(\frac{\varepsilon}{Ny(k+1)} \right) + \frac{\sum_{i \geq k+2} y(i)}{\sum_{i \geq k+1} y(i)} \left(\frac{1 - \frac{\varepsilon}{N}(k+1)}{\sum_{i \geq k+2} y(i)} \right) = \frac{1 - \frac{\varepsilon}{N}k}{\sum_{i \geq k+1} y(i)}.$$

Then, using the convexity of f we obtain the result. This implies that the first feasible prefix S that we find is the optimal one. Therefore, by ordering \mathbf{y} in $\mathcal{O}(N \log N)$ time and then running binary search we can find S in $\mathcal{O}(N \log N)$ time. This finishes the proof of (c). \square

B Omitted Proofs

B.1 Proofs of Section 3.1

Here we present dual stated in the offline formulation of the maximum work problem. We have the LP

$$\begin{aligned}
 \max \quad & \sum_{i=1}^N \sum_{t=1}^T w_t(i) \\
 (P_\varepsilon) \quad & \sum_{s=1}^t w_s(i) \leq \sum_{s=1}^t L_s(i) \quad \forall t, i \quad (1) \\
 & \sum_{i=1}^N w_t(i) \leq 1 - \varepsilon \quad \forall t \quad (2) \\
 & \mathbf{w}_t \geq 0 \quad \forall t
 \end{aligned}$$

Using the variables $\alpha_t(i)$ for constraint (1) and β_t for constraint (2) we obtain the dual

$$\begin{aligned}
 \min \quad & \sum_{i=1}^N \sum_{t=1}^T \alpha_t(i) \sum_{s=1}^t L_s(i) + (1 - \varepsilon) \sum_{t=1}^T \beta_t \\
 (D_\varepsilon) \quad & \sum_{s=t}^T \alpha_s(i) + \beta_t \geq 1 \quad \forall t, i \quad (1') \\
 & \alpha, \beta \geq 0
 \end{aligned}$$

Using the change of variable $\gamma_t(i) = \sum_{s=t}^T \alpha_s(i)$ we obtain the stated dual

$$\begin{aligned}
 \min \quad & \sum_{i=1}^N \sum_{t=1}^T L_t(i) \gamma_t(i) + (1 - \varepsilon) \sum_{t=1}^T \beta_t \\
 (D_\varepsilon) \quad & \gamma_t(i) + \beta_t \geq 1 \quad \forall t, i \quad (1') \\
 & \gamma_t \geq \gamma_{t+1} \quad \forall t \quad (2') \\
 & \beta, \gamma \geq 0
 \end{aligned}$$

Proof Proof of Proposition 2

We prove each inequality separately. Let $0 \leq t^* \leq T$ be such that $\sum_{s=1}^{t^*} \sum_i L_s(i) + (1 - \varepsilon)(T - t^*) = \min_{0 \leq t \leq T} \sum_{s=1}^t \sum_i L_s(i) + (1 - \varepsilon)(T - t)$. Consider the dual solution (β, γ) such that $\gamma_t = 1, \beta_t = 0$ for $t = 1, \dots, t^*$ and $\gamma_t = 0, \beta_t = 1$ for $t = t^* + 1, \dots, T$. Then, by weak duality,

$$v_{P_\varepsilon} \leq v_{\text{dual}}(\beta, \gamma) = \sum_{s=1}^{t^*} \sum_i L_s(i) + (1 - \varepsilon)(T - t^*) = \min_{0 \leq t \leq T} \sum_{s=1}^t \sum_i L_s(i) + (1 - \varepsilon)(T - t).$$

Now, consider the greedy algorithm that, in each iteration, gives enough allocation to the users in order to complete their work starting with user 1, then user 2, and so on. We restrict the algorithms'

allocations to $(1 - \varepsilon)$ -allocations. We denote by $\text{work}_{\text{greedy}}$ the work done by this algorithm. As usual, we denote by \mathbf{w}_t the vector of work done at time t . Let t^* be the maximum non-negative t such that $\sum_i w_t(i) < 1 - \varepsilon$. Observe that $\sum_{s=1}^{t^*} \sum_i w_s(i) = \sum_{s=1}^{t^*} \sum_i L_s(i)$. Then

$$\min_{0 \leq t \leq T} \sum_{s=1}^t \sum_i L_s(i) + (1 - \varepsilon)(T - t) \leq \sum_{s=1}^{t^*} \sum_i L_s(i) + (1 - \varepsilon)(T - t^*) = \text{work}_{\text{greedy}} \leq v_{P_\varepsilon},$$

since v_{P_ε} is the optimal solution. □

Remark 1 *This max-min result shows that the greedy algorithm is optimal for solving (P_ε) and also shows how to compute the dual variables. Finally, solving (P_ε) can be done efficiently in $\mathcal{O}(NT)$ by running the greedy algorithm.*

B.2 Proofs of Section 3.2

In what follows, we denote by S_t the users with allocation $\frac{\varepsilon}{N}$ at time t .

Proof Proof of Lemma 1

1. First, for $i \in A_t$ we have

$$h_{t+1}(i) = \frac{h_t(i)e^{\eta g_t(i)} e^{\mu_i} (1 - \frac{\varepsilon}{N}|S_{t+1}|)}{\sum_{j \in \bar{S}_{t+1}} \hat{h}_{t+1}(j)} \geq h_t(i) \frac{(1 - \frac{\varepsilon}{N}|S_{t+1}|)}{e^{-\eta} \sum_{j \in \bar{S}_{t+1}} \hat{h}_{t+1}(j)}.$$

We divide the analysis into two cases: $B_t \cap \bar{S}_{t+1} \neq \emptyset$ and $B_t \subseteq S_{t+1}$.

For $B_t \cap \bar{S}_{t+1} \neq \emptyset$ we have

$$\begin{aligned}
e^{-\eta} \sum_{j \in \bar{S}_{t+1}} \widehat{h}_{t+1}(j) &\leq e^{\lambda\eta} \sum_{j \in \bar{S}_{t+1} \cap A_t} h_t(j) + e^{-\eta} \sum_{j \in \bar{S}_{t+1} \cap B_t} h_t(j) \\
&= e^{\lambda\eta} \sum_{j \in \bar{S}_{t+1}} h_t(j) - (e^{\lambda\eta} - e^{-\eta}) \sum_{j \in B_t \cap \bar{S}_{t+1}} h_t(j) \\
&\leq e^{\lambda\eta} \left(1 - \frac{\varepsilon}{N} |S_{t+1}|\right) - \frac{\varepsilon}{N} (e^{\lambda\eta} - e^{-\eta}) \quad (\text{since } h_t(i) \geq \frac{\varepsilon}{N}) \\
&\leq (1 + 2\lambda\eta) \left(1 - \frac{\varepsilon}{N} |S_{t+1}|\right) - \frac{\varepsilon}{N} \left(\lambda\eta + \eta - \frac{\eta^2}{2}\right) \\
&\quad (\text{using } 1 + \lambda\eta \leq e^{\lambda\eta} \leq 1 + 2\lambda\eta, e^{-\eta} \leq 1 - \eta + \frac{\eta^2}{2}) \\
&\leq 1 - \frac{\varepsilon}{N} |S_{t+1}| + 2\lambda\eta - \frac{\varepsilon}{N} \eta \left(1 - \frac{\eta}{2}\right) \\
&\leq 1 - \frac{\varepsilon}{N} |S_{t+1}| - \frac{\varepsilon\eta}{4N}. \quad (\eta \leq 1 \text{ and } 2\lambda \leq \frac{\varepsilon}{4N})
\end{aligned}$$

Therefore,

$$\frac{\left(1 - \frac{\varepsilon}{N} |S_{t+1}|\right)}{e^{-\eta} \sum_{j \in \bar{S}_{t+1}} \widehat{h}_{t+1}(j)} \geq \frac{1 - \frac{\varepsilon}{N} |S_{t+1}|}{1 - \frac{\varepsilon}{N} |S_{t+1}| - \frac{\varepsilon\eta}{4N}} \geq \frac{1}{1 - \frac{\varepsilon\eta}{4N}} \geq 1 + \frac{\varepsilon\eta}{4N},$$

using $\frac{1}{1-x} \geq 1+x$ when $x \in (0, 1)$. Hence $h_{t+1}(i) \geq h_t(i) \left(1 + \frac{\varepsilon\eta}{4N}\right)$.

Now, if $B_t \subseteq S_{t+1}$, then $\bar{S}_{t+1} \subseteq A_t$. We have

$$\begin{aligned}
\frac{e^{-\eta} \sum_{j \in \bar{S}_{t+1}} \widehat{h}_{t+1}(j)}{1 - \frac{\varepsilon}{N} |S_{t+1}|} &\leq \frac{(1 - \varepsilon)e^{\lambda\eta}}{1 - \frac{\varepsilon}{N} |S_{t+1}|} \\
&\leq \frac{(1 - \varepsilon)e^{\lambda\eta}}{1 - \varepsilon + \frac{\varepsilon}{N}} \\
&\leq \frac{(1 - \varepsilon)(1 + 2\lambda\eta)}{1 - \varepsilon + \frac{\varepsilon}{N}} \quad (e^{\lambda\eta} \leq 1 + 2\lambda\eta \text{ since } 2\lambda\eta < 1) \\
&\leq \frac{1 - \varepsilon + 3\lambda\eta}{1 - \varepsilon + \frac{\varepsilon}{N}} = 1 - \frac{\frac{\varepsilon}{N} + 3\lambda\eta}{1 - \varepsilon + \frac{\varepsilon}{N}}.
\end{aligned}$$

Since $\frac{\frac{\varepsilon}{N} + 3\lambda\eta}{1 - \varepsilon + \frac{\varepsilon}{N}} \geq \frac{\varepsilon}{N}$ we obtain

$$h_{t+1}(i) \geq h_t(i) \frac{1}{1 - \frac{\varepsilon}{N}} \geq h_t(i) \left(1 + \frac{\varepsilon}{N}\right) \geq h_t(i) \left(1 + \frac{\varepsilon\eta}{4N}\right).$$

2. The monotonicity of $h_t(i)$ with $i \in A_t^1$ is easy to see. Let us prove the second statement:

$$\begin{aligned}
\frac{e^{-\eta} \sum_{j \in \bar{S}_{t+1}} \widehat{h}_{t+1}(j)}{1 - \frac{\varepsilon}{N} |S_{t+1}|} &\leq \frac{e^{\lambda\eta} \sum_{j \in \bar{S}_{t+1}} h_t(j)}{1 - \frac{\varepsilon}{N} |S_{t+1}|} \\
&\leq \frac{e^{\lambda\eta} (1 - \frac{\varepsilon}{N} |S_{t+1}|)}{1 - \frac{\varepsilon}{N} |S_{t+1}|} \\
&\leq e^{\lambda\eta} \\
&\leq 1 + 2\lambda\eta = 1 + \varepsilon c,
\end{aligned}$$

since $\lambda = \frac{\varepsilon^2}{8N}$. Then, for $i \in A_t^2$,

$$h_{t+1}(i) \geq h_t(i) \frac{e^{\eta(1 - \frac{\varepsilon}{N} |S_{t+1}|)}}{\sum_{j \in \bar{S}_{t+1}} \widehat{h}_{t+1}(j)} \geq h_t(i) \frac{1}{1 + \varepsilon c} \geq h_t(i)(1 - \varepsilon c).$$

□

Proof Proof of Claim 3.2 Now, let $[s^* + 1, \dots, T]$ and let us divide this interval into blocks of length \tilde{s} with a possible last piece of length of length at most \tilde{s} . Let L be one of these blocks and let i be the user given by claim 3.1, that is, $M_{i,r} = 1$ for all $r \in L$. Consider $L' = \{t \in L : \sum_{j \in A_t} h_t(j) < 1 - \varepsilon\}$. By using part 1 of Lemma 1, user i increases her allocation multiplicatively in L' by a factor of $(1 + c)$. Observe that for $t \notin L'$, user's i allocation can increase or decrease depending on $h_t(i)$. However, by Lemma by part 2 of 1, we know that $h_t(i)$ will not decrease by a huge amount. Let $k' = |L'|$, then i increases her allocation for k' times and decreases it for at most $\tilde{s} - k'$ times. Therefore, k' maximum value is such that

$$\frac{\varepsilon}{N} (1 + c)^{k'} (1 - \varepsilon c)^{\tilde{s} - k'} = 1$$

and therefore,

$$\begin{aligned}
k' &\leq \frac{\ln(N/\varepsilon) + \tilde{s} \ln(1 - \varepsilon c)^{-1}}{\ln((1+c)/(1-\varepsilon c))} \\
&\leq \frac{1+c}{c(1+\varepsilon)} \ln(N/\varepsilon) + \frac{\varepsilon(1+c)}{(1+\varepsilon)(1-\varepsilon c)} \tilde{s} && (\ln \frac{1+c}{1-\varepsilon c} \geq \frac{c(1+\varepsilon)}{1+c}, \ln \frac{1}{1-\varepsilon c} \leq \frac{\varepsilon c}{1-\varepsilon c}) \\
&\leq \frac{\varepsilon(1+c)}{(1+\varepsilon)} \tilde{s} + \frac{\varepsilon(1+c)}{(1+\varepsilon)(1-\varepsilon c)} \tilde{s} && (\tilde{s} = \frac{\ln(N/\varepsilon)}{\varepsilon c}) \\
&= \varepsilon \frac{1+c}{1+\varepsilon} \left(1 + \frac{1}{1-\varepsilon c}\right) \tilde{s} \\
&\leq 3\varepsilon \tilde{s}. && (\text{for } N \geq 2 \text{ and } \varepsilon \leq \frac{1}{10})
\end{aligned}$$

Hence, L' is at most a fraction of \tilde{s} and with this

$$\sum_{t \in L} \sum_i w_t(i) \geq (1-\varepsilon)(\tilde{s} - k') \geq (1-\varepsilon)(1-3\varepsilon)\tilde{s} \geq (1-4\varepsilon)|L|.$$

Summing over all blocks we conclude the desired result. \square

B.3 Proof of Section 3.3

Proof Proof of Lemma 2 If $A_t^1 = \emptyset$, the result is vacuously true. Suppose that $A_t^1 \neq \emptyset$. First, we prove that under the assumption of Lemma 2, we have $\bar{A}_t^1 \cap \bar{S}_{t+1} \neq \emptyset$. For $j \in S_{t+1}$ we have

$$\begin{aligned}
\frac{\varepsilon}{N} &= h_{t+1}(j) \\
&= \frac{\hat{h}_{t+1}(j) e^{\mu_j} (1 - \frac{\varepsilon}{N} |S_{t+1}|)}{\sum_{k \in \bar{S}_{t+1}} \hat{h}_{t+1}(k)} \\
&\geq \frac{h_t(j) (1 - \frac{\varepsilon}{N} |S_{t+1}|)}{\sum_{k \in \bar{S}_{t+1}} h_t(k) e^{\eta(1+\lambda)}}, && (\text{since } \mu_j \geq 0)
\end{aligned}$$

This implies $h_t(j) \leq \frac{\varepsilon}{N} e^{\eta(1+\lambda)} \leq \frac{\varepsilon}{N} (1 + 2\eta(1+\lambda)) < 2\frac{\varepsilon}{N} \leq \beta(j)$. Since $\sum_j h_t(j) = 1$, $i \in A_t^1$ and $\sum_j \beta(j) = 1$ we must have that there is a user $j \neq i$ with allocation $h_t(j) \geq \beta(j)$. Clearly, $j \notin S_{t+1}$ and $j \notin A_t^1$. Therefore, $\bar{A}_t^1 \cap \bar{S}_{t+1} \neq \emptyset$.

Following the proof of Lemma 1, for $i \in A_t^1$, we have

$$\begin{aligned}
e^{-\eta(1+\lambda)} \sum_{j \in \bar{S}_{t+1}} \hat{h}_{t+1}(j) &\leq \sum_{j \in \bar{S}_{t+1} \cap A_t^1} h_t(j) + e^{-\eta\lambda} \sum_{j \in \bar{S}_{t+1} \cap \bar{A}_t^1} h_t(j) \\
&= \sum_{j \in \bar{S}_{t+1}} h_t(j) - (1 - e^{-\lambda\eta}) \sum_{j \in \bar{S}_{t+1} \cap \bar{A}_t^1} h_t(j) \\
&\leq 1 - \frac{\varepsilon}{N} |S_{t+1}| - \frac{\varepsilon}{N} (1 - e^{-\lambda\eta}) \quad (\text{since } \bar{S}_{t+1} \cap \bar{A}_t^1 \neq \emptyset) \\
&\leq 1 - \frac{\varepsilon}{N} |S_{t+1}| - \frac{\varepsilon\lambda\eta}{2N}. \quad (1 - e^{-x} \geq \frac{x}{2} \text{ for } x \in [0, 1])
\end{aligned}$$

Therefore,

$$h_{t+1}(i) \geq h_t(i) \frac{1 - \frac{\varepsilon}{N} |S_{t+1}|}{1 - \frac{\varepsilon}{N} |S_{t+1}| - \frac{\varepsilon\lambda\eta}{2N}} \geq h_t(i) \left(1 + \frac{\varepsilon\lambda\eta}{2N} \right).$$

□

C Greedy Online Algorithm

In this section, we prove that the following greedy allocation strategy is almost optimal in work maximization. The algorithm divides the users into 3 categories: A non-empty queue users with non-zero allocation, B non-empty queue users with zero allocation and I empty queue users with zero allocation. At time t , a user $i \in A$ is left in A if she still has non-empty queue, otherwise we will move her to I ; a user $i \in I$ will be moved to B if her queue is becomes non-empty, otherwise she will remain in I ; finally, if all users from A are moved to I , then we will move all B to A , otherwise we will left B untouched. In any case, we will distribute uniformly among the users that remain in A .

Users move from A to I , I to B and B to A . Let \mathbf{w}_t be the work done by the algorithm and let \mathbf{w}'_t be the optimal offline work.

Theorem 7 *For any loads $L_1, \dots, L_T \in \mathbf{R}_{\geq 0}^N$, and any $\varepsilon > 0$, this greedy Algorithm guarantees*

$$\sum_{t=1}^T \sum_i w_t(i) + 2 \frac{N^2}{\varepsilon} \geq \sum_{t=1}^T \sum_i w'_t(i)$$

where \mathbf{w}_t^l is the work done by the optimal offline sequence of $(1 - 2\varepsilon/N)$ -allocations.

Proof Let t^* be the maximum $t \geq 0$ such that $\sum_{s=1}^{t+N^2} \sum_i w_s(i) \geq \sum_{s=1}^t \sum_i L_s(i)$. By claim 3.1 we know that each interval $[r, r + N^2/\varepsilon)$, with $r > t^*$, has a user with no-empty queue. As in claim 3.2 we divide the interval $[t^* + 1, T]$ into blocks of length N^2/ε with a last block of length at most N^2/ε . Pick any of these blocks, say L , and let $L' = \{t \in L : \sum_i w_t(i) < 1\}$. It is easy to see that $|L'| \leq 2N$ and therefore, summing over all block, we have $\sum_{t \geq t^*+1} \sum_i w_t(i) + N^2/\varepsilon \geq (T - t^*)(1 - 2\varepsilon/N)$. The conclusion follows applying weak duality to $(P_{2\varepsilon/N})$. \square

Remark 2 Against the best 1-allocations we can optimize ε and obtain $\varepsilon = \sqrt{N^3/T}$. This greedy strategy will be $\mathcal{O}(\sqrt{NT})$ far from the optimal dynamic work. Observe that this matches the lower bound in Theorem 8.

D Lower Bound

Theorem 8 For any online deterministic algorithm \mathcal{A} setting at each time 1-allocations, with an underlying queuing system, and with the same limited feedback as Algorithm 1, there exists a sequence of online loads L_1, \dots, L_T such that $\text{work}_{\mathbf{h}_1^*, \dots, \mathbf{h}_T^*} - \text{work}_{\mathcal{A}} = \Omega(\sqrt{T})$, where $\mathbf{h}_1^*, \dots, \mathbf{h}_T^*$ are the optimal offline dynamic 1-allocations.

Proof We consider the case with $N = 2$ users, the general case reduces to $N = 2$ by loading jobs only to two users. Let \mathcal{A} be an online algorithm for allocating a divisible resource for 2 users and with underlying queuing system and limited feedback. Without loss of generality, we can assume that the allocations sets by \mathcal{A} always add up to 1 at every time step.

We will construct a sequence of loads $L_t = (L_t(1), L_t(2))$ that at every time will add up to 1. This will ensure that the overall work done by the optimal offline dynamic policy will be T . On the other hand, we will show that this sequence of loads will lead to large queue length for at least one of the users. The main ingredient is to use the fact the algorithm receives limited feedback about the state of the system, i.e., which users have empty queue. In particular, this implies that if there are two distinct set of load vectors L_t and L'_t for some interval $t \in [r, s]$ such that the queues

remain non-empty on both these sequences, then the resource allocation to the users in the two load sequences must be identical.

We will divide the time window $[1, T]$ into phases. Each phase will begin with a configuration of queues, say $Q = (Q(1), Q(2))$, where one of the queues is empty and the other one nonempty. We set $q = Q(1) + Q(2)$ and we denote by q_i the q at phase i . We define $q_0 = 0$. We will prove that at the end of each phase $i \geq 1$, $q_{i+1} \geq q_i + \frac{1}{4}$ with all q_{i+1} cumulated in one queue and the other queue empty.

Initially, the algorithm has a fixed deterministic allocation $\mathbf{h}_1 = (h_1(1), h_1(2))$. If $h_1(1) \leq h_1(2)$, then we load $L_1 = (0, 1)$. Otherwise, we load $L_1 = (1, 0)$. In any case, we have $q_1 \geq \frac{1}{2}$ and all q_1 in one queue.

Now, we will describe how the general phases work. For the sake of simplicity, we will describe the phase starting at time $t = 1$. We have queue configuration $Q = (Q(1), Q(2))$ with $q > 0$. By the initial phase, we can assume $q \geq 1/2$. Moreover, we can assume that only one of the queues is nonempty, this point will be clear after we describe how the phase works and it is clearly true for phase 1. Phase i with $q = q_i$ will last at most $2q + 2$ time steps.

Suppose that $Q(1) = 0$ and $Q(2) > 0$. If $h_1(1) = 1$, then we load $L_1 = (0, 1)$ and the phase ends with q increased by 1 and user 1's queue empty. Therefore, we can assume $h_1(1) < 1$. Our first load will be $L_1 = (h_1(1) + \varepsilon, h_2(2) - \varepsilon)$ with $0 < \varepsilon < \frac{1}{4}$ small enough and such that both queues are nonempty. The following loads will be $L_t = \mathbf{h}_t$, the allocation of \mathcal{A} at time t . Observe that the first load will ensure that both users see nonempty queues until the end of the phase. Moreover, user 1 always has exactly ε remaining in her queue.

- If there is a time $\tau^* \in [1, 2q + 1]$ such that $h_{\tau^*}(1) \geq 1/2$, then we change the load at time τ^* for $L'_{\tau^*} = (0, 1)$. This will increment q by at least $1/2 - \varepsilon \geq 1/4$ and the phase ends. Observe that $Q_{\tau^*+1}(1) = 0$.
- We can assume now that for the loads $L_t = \mathbf{h}_t$ we always have $h_t(1) < 1/2$ for all $t \in [2, 2q+1]$. We change the loads to $L'_t = (1, 0)$ until time τ^* in which user 2 empties her queue. Recall that the feedback of the algorithm is only the set of empty queues at every time step. Thus the behavior of \mathcal{A} under L_t and L'_t will be the same until time τ^* . Now, we change load L_{τ^*} by

$L'_{\tau^*} = (1 - h_{\tau^*}(2) + Q_{\tau^*}(2) - \varepsilon', h_{\tau^*}(2) - Q_{\tau^*}(2) + \varepsilon')$ with $0 < \varepsilon' < 1/4$ small enough. This will ensure that queue 2 will be exactly ε' . Now, in an extra step, we load $L_{\tau^*+1} = (1, 0)$. Again, we have q increased by at least $1/4$ and this ends the phase. Observe that $Q_{\tau^*+2}(2) = 0$.

The analysis is similar for $Q(1) > 0$ and $Q(2) = 0$. Observe that at the end of each phase, only one queue is nonempty and the other one is empty. In any case, we have the desired increment. With this, we can set the following recurrence, $q_0 = 0$, $q_{i-1} + \frac{1}{4} \leq q_i \quad \forall i \geq 1$. We deduce that $q_i \geq i/4$. Now, let m be the number of phases. By construction, each phase last at most $2q_i + 2$. Then $T \leq \sum_{i=1}^m (2q_i + 2) \leq 40q_m^2$, where we have used $4q_m \geq 4q_i \geq i \geq 1$. From here we deduce that $q_m \geq \sqrt{T/40}$.

Now, the work done by the algorithm and the unfulfilled work in the queues must add up the overall load. Then $q_m + \text{work}_{\mathcal{A}} = T = \text{work}_{\mathbf{h}_1^*, \dots, \mathbf{h}_T^*}$ from which we obtain the result. \square

E Additional Algorithms

Algorithm 3: Proportional Greedy

Input: Sequence of loads $(L_t)_{t=1}^T$ and SLAs $\beta(1), \dots, \beta(N)$.

```

1 for  $t = 1, \dots, T$  do
2    $w_t(i) \leftarrow 0, \forall i \in [N]$ .
3    $\text{Rem} \leftarrow 1$ .
4    $\text{Rem}(i) \leftarrow Q_{t-1}(i) + L_t(i), \forall i \in [N]$ .
5   repeat
6      $A \leftarrow \{i : \text{Rem}(i) > 0\}$ .
7      $\Sigma \leftarrow \sum_{i \in A_t} \beta(i)$ .
8      $i^* \leftarrow \text{argmin}_{k \in A} \text{Rem}(k)$ .
9     if  $\text{Rem}(i) < \frac{\beta(i)}{\Sigma} \text{Rem}$  then
10       $w_t(i^*) \leftarrow w_t(i^*) + \text{Rem}(i^*)$ .
11       $\text{Rem} \leftarrow \text{Rem} - \text{Rem}(i^*)$ .
12       $\text{Rem}(i^*) \leftarrow 0$ .
13    else
14      for  $k \in A$  do
15         $w_t(k) \leftarrow w_t(k) + \frac{\beta(k)}{\Sigma} \text{Rem}$ .
16         $\text{Rem}(k) \leftarrow \text{Rem}(k) - \frac{\beta(k)}{\Sigma} \text{Rem}$ .
17       $\text{Rem} \leftarrow 0$ .
18  until  $\text{Rem} = 0$  or  $\{i : \text{Rem}(i) > 0\} = \emptyset$ ;

```

Algorithm 4: Online Proportional

Input: Sequence of loads $(L_t)_{t=1}^T$ and SLAs $\beta(1), \dots, \beta(N)$.

- 1 Initial distribution $\mathbf{h}_1 = (\beta(1), \dots, \beta(N))$,
 - 2 **for** $t = 1, \dots, T$ **do**
 - 3 Set \mathbf{h}_t and obtain $A_t = \{i : Q_t(i) \neq 0\}$,
 - 4 Update $h_{t+1}(i) = \frac{\beta(i)}{\sum_{j \in A} \beta(j)}$
 - 5 If $A = \emptyset$, then $\mathbf{h}_{t+1} = \mathbf{h}_1$.
-

F Gamma Distribution

Recall that a Gamma distribution ([36]) is characterized by two parameters: the *shape* $k > 0$ and the *scale* $\theta > 0$. The PDF of a $\text{Gamma}(k, \theta)$ is given by $\frac{1}{\Gamma(k)\theta^k} x^{k-1} e^{-x/\theta}$ where $\Gamma(k) = \int_0^\infty u^{k-1} e^{-u} du$ is the standard Gamma function. See Figure 8 for PDFs of different Gamma distribution for various choices of k and θ .

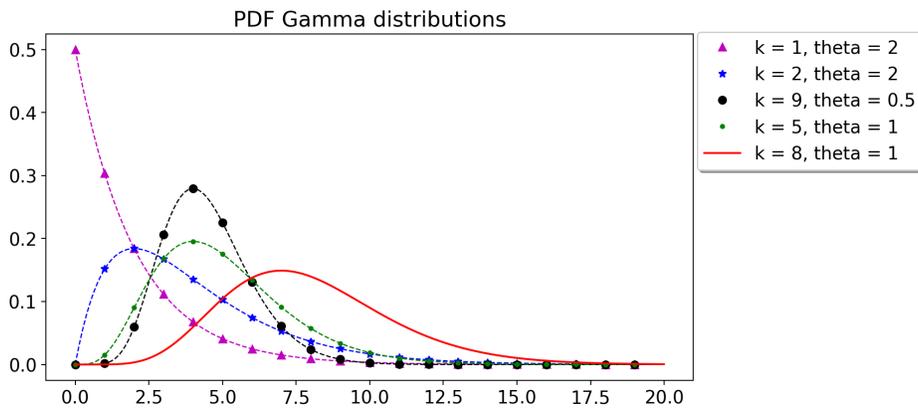


Figure 8: PDF of different Gamma distributions.

Proposition 3 *Let $X \sim \text{Gamma}(k, \theta)$. Then, $\mathbb{E}[X] = k\theta$ and $\text{Var}(X) = k\theta^2$.*