Online Resource Allocation with Matching Constraints*

John P. Dickerson
University of Maryland, College Park
College Park, MD, USA
john@cs.umd.edu

Karthik Abinav Sankararaman
University of Maryland, College Park
College Park, MD, USA
kabinav@cs.umd.edu

Kanthi Kiran Sarpatwar
IBM Research AI
Yorktown Heights, NY, USA
sarpatwa@us.ibm.com

Aravind Srinivasan
University of Maryland, College Park
College Park, MD, USA
srin@cs.umd.edu

Kun-Lung Wu
IBM Research AI
Yorktown Heights, NY, USA
klwu@us.ibm.com

Pan Xu
University of Maryland, College Park
College Park, MD, USA
panxu@cs.umd.edu

ABSTRACT

Matching markets with historical data are abundant in many applications, *e.g.,* matching candidates to jobs in hiring, workers to tasks in crowdsourcing markets, and jobs to servers in cloud services. In all these applications, a match consumes one or more *shared* and *limited* resources and the goal is to best utilize these to maximize a global objective. Additionally, one often has historical data and hence some statistics (usually first-order moments) of the arriving agents (*e.g.,* candidates, workers, and jobs) can be learnt. To model these scenarios, we propose a unifying framework, called *Multi-Budgeted Online Assignment with Known Adversarial Distributions*. In this model, we have a set of *offline* servers with different deadlines and a set of *online* job types. At each time, a job of type j arrives. Assigning this job to a server i yields a profit $w_{i,j}$ while consuming $\mathbf{a}_e \in [0,1]^K$ quantities of distinct resources. The goal is to design an (online) assignment policy that maximizes the total expected profit without violating the (hard) budget constraint. We propose and theoretically analyze two linear programming (LP) based algorithms which are almost optimal among all LP-based approaches. We also propose several heuristics adapted from our algorithms and compare them to other LP-agnostic algorithms using both synthetic as well as real-time cloud scheduling and public safety datasets. Experimental results show that our proposed algorithms are effective and *significantly* out-perform the baselines. Moreover, we show empirically the trade-off between *fairness* and *efficiency* of our algorithms which does well even on fairness metrics without explicitly optimizing for it.

CCS CONCEPTS

• **Theory of computation** → **Scheduling algorithms**; **Packing and covering problems**; **Stochastic approximation**; **Online algorithms**; • **Mathematics of computing** → *Matchings and factors*;

KEYWORDS

Online Scheduling, Online Matching, Randomized Algorithms, Fairness

ACM Reference Format:
John P. Dickerson, Karthik Abinav Sankararaman, Kanthi Kiran Sarpatwar, Aravind Srinivasan, Kun-Lung Wu, and Pan Xu. 2019. Online Resource Allocation with Matching Constraints. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 19 pages.

1 INTRODUCTION

Large-scale matching markets are abundant in many modern applications. A canonical example is the online advertising market, which is the main source of revenue for internet companies like Google. Online bipartite matching models and their variants provide mathematical insight into the design and analysis of these ubiquitous markets. In the basic version, we are given a bipartite graph $G = (U, V, E)$ where U and V represent sets of advertisers and keywords, respectively. There is an edge $e = (u, v)$ if and only if the advertisement of u is relevant to a keyword v. Keywords arrive one-by-one in an online manner and must be matched to a potential advertiser *immediately and irrevocably*. Matching a keyword v to an advertiser u gives a profit of $w_{u,v}$.

However, the above abstraction for online advertising can be used to model the more general *assignment* problem in various emerging applications, ranging from crowdsourcing marketplaces (*e.g.,* Amazon Mechanical Turk [50], matching online workers to offline tasks), ride-sharing platforms (*e.g.,* Uber, matching online requests to drivers [41]), to assignment of jobs to servers in cloud services [28]. There are several other applications of online matching models in advance admission scheduling and online recommendations (*e.g.,* matching online users to service providers or offline products [17, 43, 53]). These problems can be abstracted as a variant of online bipartite matching where (1) there are two sets of agents with at least one coming online; in any online step an immediate and irrevocable decision has to be made; (2) there is a set of offline

*Part of this work is done when Pan Xu was an intern at the IBM T. J. Watson Research Center during the summer of 2016. Aravind Srinivasan's research was supported in part by NSF Awards CNS-1010789, CCF-1422569 and CCF-1749864, and by research awards from Adobe, Inc. Karthik Sankararaman's research was supported in part by NSF Awards CNS-1010789 and CCF-1422569. John Dickerson's research was supported by NSF IIS RI CAREER Award #1846237. Pan Xu's research was supported by NSF Awards CNS-1010789, CCF-1422569 and NSF IIS RI CAREER Award #1846237. We thank Aditya Parameswaran and Xingjie Liu for useful discussions on datasets related to our problem.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

(limited) resources with each having a given total budget; every match consumes a subset of these resources.

Assadi et al. [9] considered the Online Task Assignment (OTA) problem arising in crowdsourcing marketplaces. They assumed a global budget on a single resource. Each match of an online worker to an offline task will require a payment to the worker. The goal is to design an online matching policy such that the expected number of tasks completed is maximized without violating the budget constraint. Ho and Vaughan [33] studied a capacitated OTA where every task has several copies and thus can be matched multiple times. Hence the number of copies of each task is an offline resource. Huang et al. [34], Ma et al. [41], Tong et al. [52] considered OTA emerging in the real-time spatial crowdsourcing platforms (*e.g.,* Grubhub in the online food-ordering business). In this context, we can assign multiple online orders to a single worker where each worker has two kinds of budgets: the number of orders they can handle in each trip and the total working hours and/or travel distance over all trips.

In this paper, we propose a unifying framework to handle the various budget constraints in the above applications. Additionally, we consider a realistic arrival assumption inspired from real datasets; these help us get much better provable performance. The following are the two distinctive features in the model.

Multi-Budgeted Constraints. We have a set of K resources with each resource having a known total budget. Each online match (or assignment) is associated with a vector-valued cost of dimension K with the k^{th} element denoting the amount of resource k the match consumes. We call a resource integral if and only if the value consumed by all possible matches is integral (*e.g.,* number of sub-jobs); otherwise we call it non-integral (*e.g.,* the total running time).

Known Adversarial Distributions. Common assumptions on the arrival sequence include Adversarial Order (AO) where the arrival sequence is fixed by an adversary (*e.g.,* [9]), Random Arrival Order (RAO) where the arrival sequence forms a random permutation over the set of online agents (unknown but fixed) (*e.g.,* [51, 57]) and Known Independent and Identical Distribution (KIID) where online agents present themselves, in every time-step, as a sample (with replacement) from a known and identical distribution (*e.g.,* [24, 48, 49]). In this paper, we consider a generalization of KIID, called Known Adversarial Distributions (KAD), where the arrival distributions are allowed to change over time [23]. We motivate KIID and its generalization KAD as follows. In practice, allocation algorithms are implemented in episodes. We have L episodes (where an episode could last a few hours as in cloud platforms to a day as in ride-sharing). Within each episode, algorithms use the information from the past episodes to "learn" the arrival patterns which are then used as an estimate for the current episode. This model has a two-fold challenge; first is to learn the patterns across episodes and the second is to have an efficient allocation mechanism within an episode (which is the focus of this paper).

We now give a few concrete examples and show how our model can be used to capture these with experiments on real world datasets in the experiments section.

Resource allocation in datacenters. In modern data centers, one of the challenges is to allocate various resources such as CPU, memory, clusters, to various heterogenous tasks with different requirements. The tasks usually fall into broad *categories* with very structured demands for various resources [37]. A number of recent works in the systems literature have empirically studied both efficient allocation as well fair division of resources [20, 27, 28, 36]. Our model can efficiently capture this multi-resource setting where we have multiple shared resources with high *sparsity* (*e.g.,* every computer is associated with its CPU, while the set of resources are CPU's for all computers). The tasks arrive online and should be allocated to a machine immediately and irrevocably. We look at a small time-span of one-two minutes where multiple tasks are run simultaneously on a machine with *hard* constraints on limited resources (*i.e.,* every machine has a finite amount of CPU and memory to simultaneously be utilized). The goal is to maximize the number of tasks performed and/or to drop as few requests as possible.

Public safety. Law enforcement in a given city or region is an important task for any government. The challenge is to allocate limited resources such as cops, vehicles, breath analyzers, etc. to various regions where potential violations can occur while reducing the response time and maximizing the efficiency (see, *e.g.,* [16, 29, 40, 47] for some work in this area). This general problem can be captured as an online resource allocation problem and naturally fits in our model. We have a set of offline vertices which corresponds to patrol team with multiple resources (*e.g.,* number of cops, type of vehicle, chase capabilities) stationed at various locations. When a potential violation occurs, it has to immediately be matched to a certain patrol team where the match fetches a *reward* proportional to the nature of the violation.

Hiring candidates to jobs. Consider the scenario where an HR division wants to hire new employees for a set of job positions (see, *e.g.,* [19, 55]). In practice they can hire at most one or two new employees per post (due to a total budget) and need to train each new employee to acquire the other skills needed. The training process demands several kinds of related resources such as experienced staff, time, money (paid to trainers), machines, to name a few. Suppose we have a finite budget allotted for each resource and every successful hire fetches a monetary reward to the HR. Then the problem facing the HR division can be precisely captured by our model: each match of a candidate j to a post i will consume various related resources due to the training process for j to acquire skills required by i while absent from j.

Our contributions. We make several contributions in this paper. First, we propose a general model that effectively captures the online resource allocation problems in various matching markets with historical data. Our model exploits the fact that historical data can be used to learn the arrival distributions of online agents at various times (*e.g.,* [19] in case of hiring). Second, we present algorithms that are provably correct and yield improved performance over models where distributions are assumed to be unknown. In particular, we first consider the simple case when only integral resources are involved with sparse resource consumption. We show

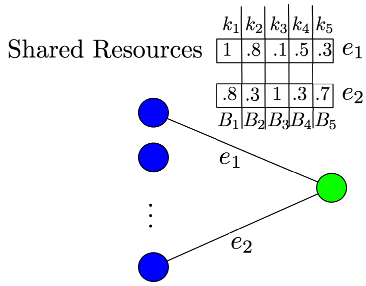

Figure 1: Five shared resources with budgets $B_1, B_2, \ldots, B_5$.

that in this case, our algorithms are near-optimal among all LP-based algorithms. Next we consider the general case when both integral and non-integral resources are involved. We show that to achieve a target competitive ratio, our model admits algorithms which have significantly improved lower-bound requirements over the budgets of non-integral resources, compared to previous ones when arrival distributions are unknown. In particular, our results completely eliminate the dependence on the ratio of largest to the smallest bid [9]. This is crucial since this ratio can typically be very large if the bids are non-uniform. Third, we consider a special case when each assignment consumes a single non-integral resource with no assumptions on its budget. We devise an algorithm with theoretical guarantees and also show hardness results. Finally, we give an empirical study of our algorithms and compare them with natural heuristics on real datasets to validate and complement our theoretical results. We also define two natural metrics for fairness for this setting and explore how efficiency maximizing algorithms perform on these fairness metrics.

2 PRELIMINARIES

We first formally define the model considered in this paper and then describe the required background for the technical sections of this paper. As a notation, denote $[k] := \{1, 2, \ldots, k\}$ for any positive integer k.

Multi-Budgeted Online Assignment (MBOA-KAD**).** Let $I = \{i \in [m]\}$ be the set of (*offline*) servers, $J = \{j \in [n]\}$ be the set of types of (*online*) jobs and T be the time horizon. Every server i has a (hard) time-out $d_i \in [T]$ after which it shuts down. Let $G = (I, J, E)$ be the bipartite graph with an edge $e = (i, j)$ iff job-type j can be run on server i. Let $N(j) = \{i : (i, j) \in E\}$ be the set of servers that can handle job-type j and $N(i) = \{j : (i, j) \in E\}$ be the set of job-types that can run on server i. Each edge $e = (i, j)$ has a weight w_e denoting the profit obtained by allocating server i to job-type j. Each assignment $e = (i, j)$ consumes one or more of a given set of K resources. The cost of an allocation e is given by a K-dimensional vector $\mathbf{a}_e \in [0, 1]^K$, where the k^{th} dimension $a_{e,k}$ represents the amount of resource k consumed by assignment e. Each resource k has a budget $B_k \in \mathbb{R}_+$ that must not be exceeded. For each e, let $S_e = \{k \in [K] : a_{e,k} > 0\}$, *i.e.,* the set of resources it consumes.

At any instant $t \in [T]$, a job of type j arrives with a probability p_{jt} such that $\sum_j p_{jt} \leq 1$ (thus, with probability $1 - \sum_j p_{jt}$, no job arrives at time t). Let $E_{jt} = \{e = (i, j), i \in N(j) : d_i \geq t\}$ denote the set of *available* assignments (*i.e.,* the corresponding servers should be active at time t) for the job-type j at time t.[1] For each $e \in E_{jt}$, we say e is *safe or valid* iff for each $k \in S_e$, resource k has a remaining budget larger or equal to $a_{e,k}$. When a job of type j arrives at t, we have to make an immediate and irrevocable decision: either reject it or choose a safe option $e \in E_{jt}$ and get a resultant profit w_e. Once a safe assignment e is scheduled, the budget of each resource $k \in S_e$ will be reduced by $a_{e,k}$. Our goal is to design an online assignment policy such that the expected profit is maximized.

Note that all algorithms presented in this paper are applicable to a more general setting where each successful match e yields a *random* profit W_e (independent from others). All the algorithms need to know is $w_e \doteq \mathbb{E}[W_e]$ for each e [2].

The performance of online algorithms is usually measured using the notion of competitive ratio (see [13]). For our problem, we define the competitive ratio as follows.

Definition 2.1 (Competitive Ratio). Let ALG denote a given online algorithm whose performance we want to measure. Consider an instance $\mathcal{I}$ of the problem. Let $\mathbb{E}[\text{ALG}(\mathcal{I})]$ denote the expected profit obtained by ALG for this instance (here expectation is over the randomness in the input as well as any randomness the algorithm uses). Similarly, let $\mathbb{E}[\text{OPT}(\mathcal{I})]$ denote the expected value of the optimal offline solution (*i.e.,* the expected value of the optimal solution on seeing the entire arrival sequence). The competitive ratio is defined as $\inf_{\mathcal{I}} \mathbb{E}[\text{ALG}(\mathcal{I})]/\mathbb{E}[\text{OPT}(\mathcal{I})]$.

For any maximization problem like the one studied here, we say ALG achieves a ratio at least $\alpha \in (0, 1)$ if for any instance of the problem the expected profit obtained by ALG is at least a fraction α of the offline optimal solution. Typically computing the value of $\mathbb{E}[\text{OPT}(\mathcal{I})]$ directly is hard. A common methodology to bypass this is to construct a linear program (called *benchmark* LP) whose optimal value is an upper bound on $\mathbb{E}[\text{OPT}(\mathcal{I})]$. Hence comparing $\mathbb{E}[\text{ALG}(\mathcal{I})]$ to the optimal value of this LP gives a lower bound on the competitive ratio. We will now describe the benchmark LP used in this paper.

Recall E_{jt} be the set of available assignments for a job of type j arriving at time-step t. For any t, let $E_t = \bigcup_j E_{jt}$ be the set of all possible assignments (a-priori before the execution of any online algorithm) at t. Further, for each t and $e \in E_t$, let $x_{e,t}$ be the probability that an assignment e is made at round t in an offline optimal algorithm. Then the benchmark LP we use is as follows.

$$\text{maximize} \sum_t \sum_j \sum_{e \in E_{j,t}} w_e x_{e,t} \tag{1}$$

$$\text{subject to} \sum_{e \in E_{jt}} x_{e,t} \leq p_{jt} \quad \forall j \in J, t \in [T] \tag{2}$$

$$\sum_t \sum_{e \in E_t} x_{e,t} a_{e,k} \leq B_k \quad \forall k \in [K] \tag{3}$$

$$0 \leq x_{e,t} \leq 1 \quad \forall e \in E, t \in [T] \tag{4}$$

This LP can be interpreted as follows. Constraint (2) — for any given job of type j and time t, the probability that we assign a server to j is at most the probability that j arrives at step t. Constraint (3) — for any (integral or non-integral) resource k, the expected consumption cannot be larger than its budget (B_k). The last constraint (4)

[1]In this paper, we assume w.l.o.g. that each server can be allocated for an arbitrary number of times before its shutdown. Any potential restriction on the number of allocations can easily be modeled by an additional budget constraint.

[2]An allocation does not imply a completion of a job, hence this uncertainty can be handled in our model.

is due to the fact that all $\{x_{e,t}\}$ are probability values and hence should lie in the interval $[0,1]$. The above analysis suggests that any offline optimal solution $\{x_{e,t}\}$ should be feasible for the above LP. Formally, we have Lemma 2.2 which claims that the optimal solution of this LP is an upper bound on the expected offline optimal value.

Lemma 2.2. *The optimal value to* LP-(1) *is a valid upper bound for the offline optimal solution.*

The benchmark LP-(1) has previously been used in [7] and the proof of Lemma 2.2 can be found there.

Integral and non-integral resources. For an integral resource k, we have that for any $e \in E$, $a_{e,k} \in \{0,1\}$ while for a non-integral resource k, we have that for any $e \in E$, $a_{e,k} \in [0,1]$. For any integral resource k, WLOG we assume that $B_k \in \mathbb{Z}_+$. Let $\mathcal{K}_1 = \{1,2,\cdots,K_1\}$ and $\mathcal{K}_2 = \{K_1+1,\cdots,K_1+K_2\}$ denote the set of integral and non-integral resources respectively. For any assignment e, we assume $|S_e \cap \mathcal{K}_1| \le \ell_1$ and $|S_e \cap \mathcal{K}_2| \le \ell_2$, where ℓ_1 and ℓ_2 are the integral and non-integral sparsity respectively.

Adaptive and non-adaptive algorithms. For an LP-based ALG, we say ALG is *non-adaptive* if for a given LP solution, the computation of strategy in each round t does not depend on the strategies in the previous rounds from $1,2,\ldots,t-1$. Otherwise, we call it "adaptive". Here we distinguish "adaptive" and "non-adaptive" to highlight the computations of strategies in the "online" phase.

3 OTHER RELATED WORK

We now describe some related works other than those already described. Our problem falls under the *online packing* family of problems. Some representative works on this include [4, 6, 14, 15, 22, 38]. The most relevant to this paper is that of Devanur et al. [22] who study it in the unknown i.i.d. setting. We have the following important distinctions. First, our work assumes the known i.i.d. setting. Second, Devanur et al. [22] study this problem in the large budget regime (*i.e.,* $B \ge \Omega(\frac{1}{\epsilon^2})$) and obtain $1-\epsilon$ approximation. This paper however considers regimes where the lower bound on the budget is only $\Omega(\frac{1}{\log \epsilon})$. We circumvent the necessity for a large lower-bound on the budget is by assuming sparsity in the packing program. Closely related to the online packing problems literature is another line of work on a family of problems called *bandits with knapsacks* [2, 3, 5, 10, 11, 46]. These works consider the learning variant of the online packing problems and obtain approximation ratios that are comparable to [22]. Many special cases of online packing problems and bandits with knapsacks have been studied across communities including dynamic pricing ([21] and references within), network routing and optimization ([10] and references within), network revenue management ([12] and references within).

4 ALGORITHMS

In this section, we describe our two main algorithms NADAP and ADAP.

Non-adapative algorithm (NADAP**).** Algorithm NADAP is a non-adaptive algorithm based on LP. Suppose $\{x^*_{e,t} | t \in [T], e \in E_t\}$ is an optimal solution to the LP-(1). The main idea behind NADAP (described in Algorithm 1) is as follows. Suppose a job of type j arrives at time t: sample a server i from E_{jt} with probability $\alpha x^*_{e,t}/p_{jt}$, where $\alpha \in (0,1]$ is a parameter optimized in the analysis. Make the assignment $e=(i,j)$ iff e is safe (*i.e.,* it will not violate any budget constraint at t).

ALGORITHM 1: The non-adaptive algorithm (NADAP)

For each time t, let the arriving job be denoted by j.
Let $\hat{E}_{jt} \subseteq E_{jt}$ be the set of *safe* assignments available for j.
If $\hat{E}_{jt} = \emptyset$, then reject j; else sample an assignment $e \in \hat{E}_{jt}$ with probability $\alpha x^*_{e,t}/p_{jt}$.

The last step of Algorithm 1 is well defined since we have $\sum_{e \in \hat{E}_{jt}} \alpha x^*_{e,t}/p_{jt} \le \sum_{e \in E_{jt}} x^*_{e,t}/p_{jt}$, which is at most 1.[3]

Adapative algorithm (ADAP**).** Algorithm ADAP is an adaptive algorithm which uses Monte-Carlo simulations. The main idea is as follows. Suppose we aim to develop an online algorithm achieving a competitive ratio of $\gamma \in [0,1]$. Consider an assignment $e=(i,j) \in E_t$ for a job j at time t. Let $\mathcal{S}_{e,t}$ be the event that e is safe (*i.e.,* we can choose this assignment without budget violation for all resources) conditioning on the arrival of e at t. By using Monte-Carlo simulation of the strategy up to t, we can get a sharp estimate of $\Pr[\mathcal{S}_{e,t}]$, say $\beta_{e,t}$, with polynomial number of samples. Therefore if e is safe at t, we choose it with probability $\frac{x_{e,t}}{p_{j,t}}\frac{\gamma}{\beta_{e,t}}$, which implies that e is chosen with probability $\gamma x_{e,t}$ unconditionally.

The simulation-based attenuation technique has been used previously in other stochastic optimization problems (*e.g.,* stochastic knapsack [42], stochastic matching [1]). We assume that the sharp estimate $\beta_{e,t}$ of $\Pr[\mathcal{S}_{e,t}]$ for all t and e is *exact*, since the sampling error can be accounted as a multiplicative factor of $(1-\epsilon)$ in the competitive ratio by a standard Chernoff bound argument. Formally our algorithm, denoted by ADAP, is described in Algorithm 2. The running time of this algorithm is polynomial in $1/\varepsilon$.

ALGORITHM 2: The adaptive algorithm (ADAP)

At time t, let j be the job that arrives.
Let $\hat{E}_{j,t} \subseteq E_{j,t}$ be the set of *safe* assignments available for j.
If $\hat{E}_{j,t} = \emptyset$, then reject j; else sample an assignment $e \in \hat{E}_{j,t}$ with probability $\frac{x^*_{e,t}}{p_{j,t}}\frac{\gamma}{\beta_{e,t}}$.

To ensure the above algorithm is mathematically well-defined with parameter γ, we need to show that $\beta_{e,t} \ge \gamma$ for every t and e.

Lemma 4.1 (Validity of ADAP). *By choosing $\gamma = 1/(\ell+1)$, we have $\beta_{e,t} \ge \gamma$ for all $t \in [T]$ and $e \in E_t$.*

Both the algorithms presented do not work in the regime when B is small. To overcome this, we propose a new algorithm with additional restrictions when $B_k = 1$. Consider MBOA-S which has the following setting: (1) all resources are non-integral and have a unit budget $B_k = 1$; (2) each assignment requires only one

[3] In other words, with probability $1 - \sum_{e \in \hat{E}_{jt}} \alpha x^*_{e,t}/p_{jt}$, we will do nothing and reject job j.

single resource ($\ell = 1$); (3) servers have no deadline ($d_i = T$); (4) the arrival distributions over all job-types are *identical* across all rounds, *i.e.,* for each job-type j, $p_{jt} = p_j$ for all $t \in [T]$. It can be shown that the performance of ALG_1 and ALG_2 can be arbitrarily bad (see supplementary).

This new setting requires a modified benchmark LP. For each $e \in E$, let x_e be the expected number of times the assignment e is made in the offline optimal over the T rounds. For each e, we use a_e to denote the cost of e for the unique resource it consumes. For a given threshold $\alpha = 1/2$, we say e is *big* if $a_e > \alpha$ and *small* otherwise. For each resource k, let BIG(k) (SM(k)) refers to the set of big assignments (small) participating on constraint (or resource) k. We use the following benchmark LP.

$$\max \textstyle\sum_{e \in E} w_e x_e \tag{5}$$
$$\text{s.t.} \textstyle\sum_{e \in E_j} x_e \le p_j * T \quad \forall j \in J, t \in [T] \tag{6}$$
$$\textstyle\sum_{e \in \text{BIG}(k)} x_e a_e + \sum_{e \in \text{SM}(k)} x_e a_e \le 1 \quad \forall k \in [K] \tag{7}$$
$$\textstyle\sum_{e \in \text{BIG}(k)} x_e \le 1 \quad \forall k \in [K] \tag{8}$$

Constraint (8) is valid since in any offline optimal, at most one big assignment from BIG(k) can be made, for each k. We design a new algorithm for this setting, whose main idea is as follows. We split the whole T rounds into two stages, where the first stage consists of the first $T * \beta$ rounds and the second stage consists of the remaining rounds. In the first stage, our algorithm only considers big assignments and drops any small assignment while in the second stage it considers only small ones while dropping the big ones. For each job-type j, let BIG(j) (SM(j)) be the set of big (small) assignment with respect to j. Suppose $\{x_e^*\}$ is an optimal solution to the new benchmark LP (5). We can then formally describe the algorithm as in Algorithm 3.

ALGORITHM 3: MBOA-S(β, γ_1, γ_2)

The first stage:

For each time $t \in [T * \beta]$, assume some job j arrives.

Sample a big assignment $e \in \text{BIG}(j)$ with probability $\frac{\gamma_1 x_e^*}{p_j * T}$.

If e is safe, then make it; otherwise reject it.

The second stage:

For each time $t \in \{T * \beta + 1, T * \beta + 2, \ldots, T\}$, assume some job j arrives.

Sample a small assignment $e \in \text{SM}(j)$ with probability $\frac{\gamma_2 x_e^*}{p_j * T}$.

If e is safe, then make it; otherwise reject it.

---

5 MAIN RESULTS AND TECHNIQUES

We now describe the main results and theoretical techniques used. Detailed proofs are deferred to the supplementary.

First, we present two algorithms based on LP-(1), NADAP and ADAP, which are non-adaptive and adaptive respectively. For the integral MBOA-KAD where all resources are integral, we have the following theorems.

Theorem 5.1 (Performance of NADAP for integral case). *For* MBOA-KAD *when all resources are integral with sparsity ℓ,* NADAP *with $\alpha = \frac{1}{\ell+1}$ achieves a competitive ratio of at least $\frac{1}{\ell+1}(1 - \frac{1}{\ell+1})^{\ell} \ge \frac{1}{e(\ell+1)}$ using* LP-(1) *as the benchmark. The analysis for this is tight.*

Theorem 5.2 (Performance of ADAP for integral case). *For* MBOA-KAD *when all resources are integral with sparsity ℓ,* ADAP *with $\gamma = \frac{1}{\ell+1}$ achieves a competitive ratio of at least $\frac{1-\epsilon}{\ell+1}$, for any given constant $\epsilon > 0$. Moreover, no adaptive algorithm can achieve a ratio better than $\frac{1}{(\ell-1+1/\ell)}$.*

From the above two theorems, we have that NADAP and ADAP are almost optimal among all algorithms that use LP-(1) as benchmark, for the integral MBOA-KAD. Additionally, Theorem 5.1 achieves the best possible ratio that NADAP can get based on LP-(1). For algorithms which do not use LP-(1) as benchmark, the hardness result is $O(\ln \ell/\ell)$ since the inapproxibility result of ℓ-uniform hypergraph matching [31] carries over to this setting.

We now show an example to show that the results proved in Theorem 5.1 is tight.

Example 5.3 (Tight Example for Integral Resources). Consider a star graph $G = (I, J, E)$ where $|I| = 1, |J| = \ell+1, E = \{e_j | j \in [\ell+1]\}$ with $T = \ell + 1$. Let $d_1 = T$, *i.e.,* no deadline constraints. For each $t \in [T], p_{jt} = 1$ iff $j = t$ and 0 otherwise. We use $\mathbf{a}_j$ and x_j^* to denote the terms $\mathbf{a}_{e_j}$ and $x_{e_j, t=j}^*$. Let $K = \ell$ with $B_k = 1$ for each $k \in [\ell]$ and $\mathbf{a}_j = \mathbf{e}_j$ for each $j \le \ell$, where $\mathbf{e}_j$ is the jth standard-basis unit vector of dimension K, and $\mathbf{a}_j = \mathbf{1}$ (of dimension K) for $j = \ell + 1$. Let the optimal solution to LP-(1) be $x_j^* = 1 - \epsilon$ for each $j \le \ell$ and $x_{\ell+1}^* = \epsilon$ for a proper weight vector. Now consider the assignment $e = e_{\ell+1}$ when $j = \ell + 1$ comes at $t = T$. Let us compute the probability $\Pr[\mathcal{S}_{e,T}]$ that e is safe at T in NADAP(α). Assignment e will be safe at $t = T$ iff none of $e_j, j \le \ell$ is made before. At each time $t < T$, NADAP(α) makes the assignment $e_{j=t}$ with probability $\frac{\alpha x_j^*}{p_j} = \alpha(1 - \epsilon)$. This implies that $\Pr[\mathcal{S}_{e,T}] = (1 - \alpha(1 - \epsilon))^{\ell}$, which matches the lower bound. □

Next, we consider a general case, where we have both integral and non-integral resources while making a mild assumption that the budget of any non-integral resource is large enough. Let B be the minimum budget for any non-integral resource. We then prove the following two theorems.

Theorem 5.4 (Performance of NADAP for the general case). *For* MBOA-KAD *with integral and non-integral sparsity ℓ_1 and ℓ_2,* NADAP *with $\alpha = \frac{1}{\ell_1+1}$ achieves a competitive ratio of $\frac{1}{\ell_1+1}\Big((1 - \frac{1}{\ell_1+1})^{\ell_1} - \varepsilon\Big)$, for any $\varepsilon > 0$, assuming $B \ge 2\ln(\frac{\ell_2}{\varepsilon})\Big(1 + \frac{3\ell_1+2}{\ell_1^2}\Big) + 2$.*

Theorem 5.5 (Performance of ADAP for the general case). *For* MBOA-KAD *with integral and non-integral sparsity ℓ_1 and ℓ_2,* ADAP *with $\gamma = \frac{1-\epsilon}{\ell_1+1}$ achieves a competitive ratio of $\frac{1-2\varepsilon}{\ell_1+1}$ for any given $\varepsilon > 0$, assuming $B \ge 3\ln(\frac{\ell_2}{\varepsilon})(1 + \frac{1}{\ell_1}) + 2$.*

Our results imply that with the knowledge about arrival distributions we can obtain significant improvements over the results for the adversarial model. Let us compare our results with those of [9]. The setting in [9] can be viewed as a special case of our model with $\ell_1 = \ell_2 = 1$. From Theorem 5.5, we obtain a $(\frac{1}{2} - \epsilon)$ competitive ratio assuming $B \ge 12\ln(1/\epsilon)$ while [9] obtain a ratio of $O(\frac{1}{R^{\epsilon}\ln R})$, assuming $B \ge \frac{R}{\epsilon}$ and $R \doteq \frac{\max b_{i,j}}{\min b_{i,j}}$ (*i.e.,* the ratio of the largest bid to the smallest bid over all possible assignments). *Note that our results completely removes the dependency on R and also significantly relax* *the lower bound assumption on B*. This is a theoretical evidence to advocate the use of historical data to learn arrival distributions.

Third, we consider the case of MBOA-KAD when both integral and non-integral resources are involved but no lower bound is known for the budgets of non-integral resources. To make the problem tractable, we make the following three assumptions: (1) each assignment consumes only a single resource ($\ell = 1$); (2) deadline constraints on all servers are removed; (3) the arrival distributions over all job types are *identical* across all rounds. We refer to MBOA-KAD under these three simplifying conditions as MBOA-S. Note that MBOA-S still generalizes the well-known online bipartite matching problem and several variants such as Adwords and Display Ads (*e.g.,* [25, 30, 35, 44]). It can be shown that the performance of the two previous algorithms, NADAP and ADAP, can be arbitrarily bad in this case. We propose a strengthened benchmark LP and obtain the following theorem.

THEOREM 5.6. *There exists an online algorithm which achieves an online ratio of $\frac{1}{4}$ for* MBOA-S. *Meanwhile, no online algorithm can achieve a ratio better than $\frac{e-1}{2e-1} \sim 0.387$.*

6 EXPERIMENTS

In this section we describe our experimental results. We use the Google cluster trace data [18, 32] which was used by Kash et al. [37]. This dataset contains traces of job allocation to servers within Google's datacenters. We process this dataset for our purposes, which we describe below. To further show the generality of our model, we also run additional experiments modeling an allocation problem in the public safety domain.

Experimental setup. Every machine is characterized by an id, the total CPU capacity and the total memory capacity. Each sample is from an interval of 2 minutes in the dataset and hence is a short enough time-span to consider *hard total budget* on the resources. A job type is characterized by the CPU and memory requirements. Hence there can be multiple jobs of the same type (which we use to construct our arrivals). Every machine consumes two resources and these resources are not shared by other machines (this application has a simpler notion of *shared* resources compared to the generality our model can handle). Therefore we have total of $2m$ resources with a sparsity $\ell = 2$. We assume that all machines are active throughout.

Dataset and preprocessing. Our experimental setup is inspired from the experiments of Kash et al. [37]. We use a random subset of the dataset by sampling m machines and n jobs randomly for $m = 10$ and $n = \{20, 100\}$. Our experiments are run for both the (m, n) pairs. We assign arrival rates randomly to each of the jobs and use it for all the experiments. For every (m, n) pair we generate the compatibility graph by choosing 5 machines at random, that a job j can be run on. All experiments are reported by running 100 independent trials and taking the sample average. Assignment weights are assigned by generating an independent random number between 0 and 1.

Algorithms. We compare our main algorithm NADAP with the following three baselines. These baselines have been used previously in the literature [23]. The baselines are as follows. Suppose job j arrives at time t. (1) SCALED: sample an assignment $e \in E_{jt}$ with probability $\frac{x^*_{e,t}}{\sum_{e \in E_{jt}} x^*_{e,t}}$ and assign e iff safe. (2) USamp: sample an assignment $e \in E_{jt}$ uniformly from E_{jt} and assign e iff safe. (3) Greedy: choose the assignment $e \in E_{jt}$, which has the largest weight w_e among all safe options in E_{jt}. Finally note that ADAP had similar performances as NADAP in our experiments and for clarity we omit those from the figures. Deviating from the conservative estimate predicted by theory for α in NADAP, throughout the experimental section we set $\alpha = 1$. We chose this by tuning for optimal performance on a small holdout of the dataset. The reason this value is different from what theory predicts is that, in theory we are optimizing for the worst case input, while, as will be evident from the results, these datasets do not represent the worst-case graphs.

Throughput experiments. We compare the *total weight* of all the assignments made by each of these algorithms. The first column in Figure 2 describes the details. It is clear that our main algorithm NADAP performs the best. Among the baselines Greedy is better than the other algorithms. Despite SCALED having the information of the optimal solution from the LP, it is not able to perform as good as Greedy. This shows the inherent power of *adaptive* algorithms.

Fairness experiments. We run further additional experiments to study the *fairness* of these algorithms. Fairness is a broad topic and our goal in this paper is to show that despite not explicitly optimizing for it, NADAP performs well compared to the baselines. We discuss two notions of fairness which are inspired from the *max-sum* and *max-min* fairness of Kash et al. [37]. Our setting differs from theirs and hence we cannot directly use their definitions. However we define two notions, namely *drop-sum* and *drop-max* fairness. For a give type j, let ds_j be defined as the expected difference between the number of times this type appears in an arrival sequence and the number of times an algorithm assigns it successfully. *Drop-sum* metric calculates the sum of ds_j for every type j while *drop-max* calculates the maximum over all types j of ds_j. Intuitively, higher the value of either of these metrics, the more *unfair* the algorithm is. In the second column of Figure 2 we show how the baselines and our algorithm performs on the *drop-sum* metric. Our algorithm once again comes on top with almost no difference among the other three baselines. On the *drop-max* metric in the third column of Figure 2, however, we see an interesting result where our algorithm is slightly worse than the three baselines. However the difference is not too significant, suggesting that alongside maximizing throughput, NADAP is also inherently *fair*.

Does graph sparsity matter? [4] We further study if either throughput or the two fairness metrics significantly change for the algorithms when the graph sparsity is varied. We use the setup of 10 machines, 20 jobs types and 1000 arrivals as the underlying parameters. Graph sparsity is varied by controlling the number of neighbors each job-arrival can be assigned to. Figure 3 shows the results for this experiments. As evident, there is some variation in the absolute numbers, but these are not significant enough in

[4]Term sparsity is overloaded. Here it refers to the number of edges, which is different from sparsity of resources used throughout the paper.

Figure 2: Allocation experiments on the Google cluster trace dataset. The upper row corresponds to 10 machines and 20 job types, while the lower row corresponds to 10 machines and 100 job types. First column is the results for Throughput experiments, second column is the results for the *drop-sum* fairness experiments and the third column is the results for the *drop-max* fairness experiments.

absolute terms. Additionally, these numbers do not change the relative ordering for the performance of the algorithms in either of the three metrics.

Discussion. The main experiments suggest that NADAP performs well in practice, alongside having good theoretical guarantees. In fact, the throughput experiments indicate that the performance guarantee is far better than the theoretical prediction of around 0.3. We also show that for a reasonable definition of fairness metrics, our algorithm performs as good or better than the baselines. This suggests that our model and algorithm is suitable for scenarios where we want to maximize both throughput and fairness properties. It is also an interesting future direction to explicitly account for fairness to further improve the performance of NADAP.

6.1 Additional Experiments for the Public Safety Application

We use a large-scale policing dataset [45] for the public safety application.

Dataset and assumptions. The policing dataset [45] contains records from the Texas state since 2010. Every record in this dataset contains the following: County, Latitude and Longitude, Time, Violation, Officer id, whether there was a search, stop outcome and driver details such as gender, age, race (which we do not require for our purposes). We use this information to create offline and online vertices as well as the graph as follows. For every county we create one offline vertex and pick one of the locations within this county as the station (we assume that dispatch for this county is from this station). For the online side, we create one online vertex for every pair of (location, offense).[5] For each vertex or type (offline and online), the set of neighbors is the set of 5 counties that are within a radius $r = 2$ (in terms of differences in latitude and longitude) of the location associated with this vertex. The dataset provides the time rounded to the nearest minute. We regard every minute as a time-step and set the time horizon $T = 24 * 60 = 1440$ (24-hour period). We sample 20 frequently appearing counties for the offline side, 120 frequently appearing (location, offense) pairs for the online types. To learn p_{jt}'s, we average the arrival frequencies in a randomly chosen 90-day period and use another randomly chosen 14-day period for testing. We consider the unweighted case and are only interested in maximizing the total number of unlawful activities handled.

We assume that there are 7-types of resources, namely, total travel distance, officers, patrol vehicles and equipments for speeding, search, citation and arrest. The first 3 resources are owned by each offline vertex (station) while the last four are shared across vertices (therefore a total of $3 * |I| + 4$ resources). The first resource captures the total working hours each station can continuously be engaged in and account for it as the total travel distance travelled. The last four represent the resources involved for detecting and issuing a speeding violation, for conducting a search on the vehicle, for issuing a citation and for making an arrest on the spot respectively. Each match $e = (i, j)$ (assignment of online offense type j to station i) will consume 7 resources; the consumption of the first three resources is jointly determined by the pair (i, j) (denoted by $\mathbf{a}_{e,1}$) and the consumption of the last four resources is determined by type j (denoted by $\mathbf{a}_{e,2}$). By averaging over all records from the dataset, we compute and normalize such that $\mathbf{a}_{e,1} \in [0, 1]^3$ and $\mathbf{a}_{e,2} \in [0, 1]^4$. To make $\mathbf{a}_{e,2}$ depend on station i, we uniformly sample a number λ_i from $[0, 1]$ and set the final cost $\mathbf{a}_e$ as $(\mathbf{a}_{e,1}, \lambda_i * \mathbf{a}_{e,2})$.

[5] Location is a latitude longitude pair rounded to the nearest integer.

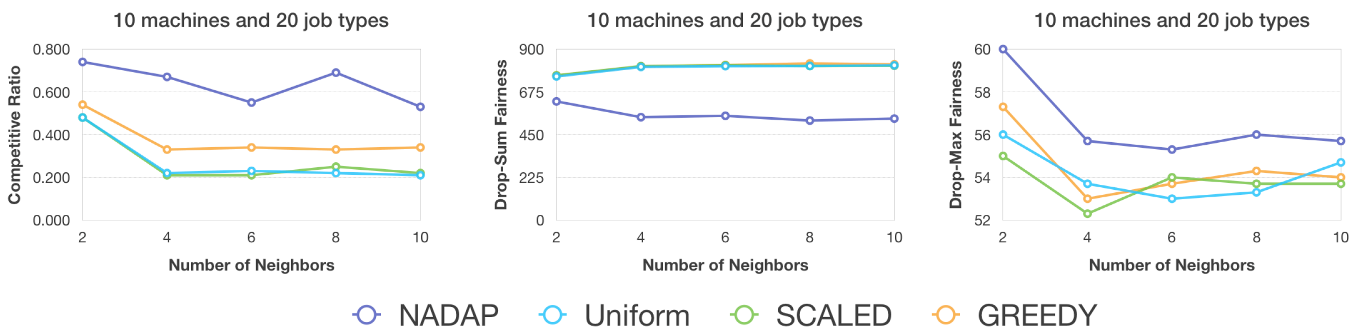

Figure 3: Varying the graph sparsity for instance with 10 machines, 20 job types and 100 arrivals of jobs. First, second and third column corresponds to throughput, drop-sum and drop-max metrics respectively.

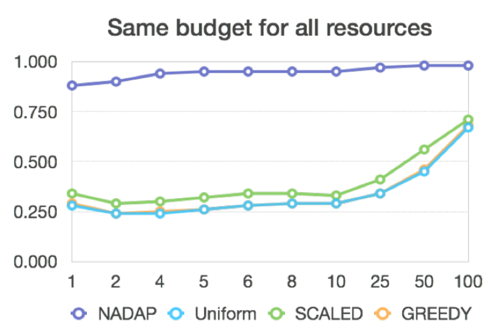

(a) Comparing algorithms when all resources have a uniform budget. x and y-axises represent the budget value and competitive ratio respectively.

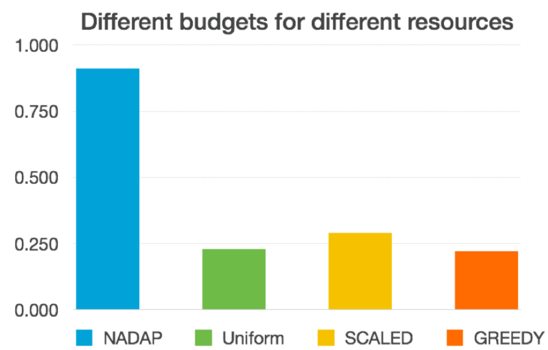

(b) Comparing algorithms when different resources have different budgets

Experimental parameters and baselines. First, we conduct experiments by choosing a unform budget for all resources. In particular we choose the budget value in the range $\{1, 2, 4, 5, 6, 8, 10, 25, 50{,}100\}$. We also run an experiment with different budgets for different resources. For every resource we choose a random integer between 1 and 100 as the budget. The value of these budgets are chosen to ensure that it accounts for both the small budget (*i.e.,* optimal solution is much smaller than T) and the large budget (*i.e.,* optimal is very close to T) case. For each given experiment and algorithm, we run 10 independent runs on each of the 14-day testing period and take the average. We use the following parameters for the algorithms. For NADAP, we set $\alpha = 1$ (which is higher than the theoretical value which was fine-tuned via standard cross-validation approach) and set $\gamma = 1$ for ADAP.

Results and discussion. Figures 4a and 4b describe the respective experimental results when the total budget takes a uiform value or different values among resources. From the results we can observe that our algorithms ADAP and NADAP significantly out-perform all our natural baselines. Among the baselines the LP-based baseline (SCALED) beats the LP-agnostic baselines (albeit by a small amount). Additionally, our algorithms almost approach the optimal on this dataset which is far better than the theoretical guarantee (around $1/8 = 0.125$). Hence experimentally we show that our algorithms are useful, beat the natural baselines comprehensively and achieve near optimal in practice.

7 CONCLUSIONS & FUTURE RESEARCH

In this paper, we studied the multi-budgeted allocation problem in the context of matching markets such as crowdsourcing, candidate hiring, etc. In this context, we proposed a novel model and provided efficient LP-based algorithms with improved competitive ratios. In particular, we showed two algorithms and analyzed their performance formally. In the theoretical analysis of these algorithms, we used novel ideas which can potentially be of independent interest in both the analysis of online matching and allocation algorithms. Finally, our algorithms were compared against several heuristic baselines experimentally on a real-world dataset to validate our theoretical results. We also explored properties of our algorithms in the context of tradeoffs in economic efficiency and fairness.

In addition to further exploration of the interplay between various fairness objectives and the traditional efficiency-maximizing objective used in many settings, we believe that the inclusion of incentives in our models would be of future interest. For example, Kash et al. [37] explore various design desiderata in their dynamic fair divison (of divisible tasks) model; similar qualitative steps could be taken in our model (which focuses on *in*divisible allocation). As another example, resource allocation in the *security games* setting [39], limited resources are deployed to prevent a strategic adversary from attacking targets. Recent work has explored security games under various forms of dynamism [8, 54, 56].

REFERENCES

[1] Marek Adamczyk, Fabrizio Grandoni, and Joydeep Mukherjee. 2015. Improved approximation algorithms for stochastic matching. In *ESA*.

[2] Shipra Agrawal and Nikhil Devanur. 2016. Linear contextual bandits with knapsacks. In *NIPS*.

[3] Shipra Agrawal and Nikhil R Devanur. 2014. Bandits with concave rewards and convex knapsacks. In *EC*.

[4] Shipra Agrawal and Nikhil R Devanur. 2014. Fast algorithms for online stochastic convex programming. In *SODA*.

[5] Shipra Agrawal, Nikhil R Devanur, and Lihong Li. 2016. An efficient algorithm for contextual bandits with knapsacks, and an extension to concave objectives. In *COLT*.

[6] Shipra Agrawal, Zizhuo Wang, and Yinyu Ye. 2014. A dynamic near-optimal algorithm for online linear programming. *Operations Research* 62, 4 (2014).

[7] Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. 2013. The online stochastic generalized assignment problem. In *APPROX-RANDOM*.

[8] Tansu Alpcan and Sonja Buchegger. 2011. Security games for vehicular networks. *IEEE Transactions on Mobile Computing* 10, 2 (2011), 280–290.

[9] Sepehr Assadi, Justin Hsu, and Shahin Jabbari. 2015. Online Assignment of Heterogeneous Tasks in Crowdsourcing Markets. In *AAAI-HComp*.

[10] Ashwinkumar Badanidiyuru, Robert Kleinberg, and Aleksandrs Slivkins. 2013. Bandits with knapsacks. In *FOCS*.

[11] Ashwinkumar Badanidiyuru, John Langford, and Aleksandrs Slivkins. 2014. Resourceful contextual bandits. In *COLT*.

[12] Omar Besbes and Assaf Zeevi. 2012. Blind network revenue management. *Operations research* 60, 6 (2012), 1537–1550.

[13] Niv Buchbinder, Kamal Jain, and Joseph Seffi Naor. 2007. Online primal-dual algorithms for maximizing ad-auctions revenue. In *ESA*.

[14] Niv Buchbinder and Joseph Naor. 2009. Online primal-dual algorithms for covering and packing. *Mathematics of Operations Research* 34, 2 (2009), 270–286.

[15] Niv Buchbinder, Joseph Seffi Naor, et al. 2009. The design of competitive online algorithms via a primal–dual approach. *Foundations and Trends® in Theoretical Computer Science* 3, 2–3 (2009), 93–263.

[16] Jan M. Chaiken and Peter Dormont. 1978. A Patrol Car Allocation Model: Background. *Management Science* 24, 12 (1978).

[17] Xi Chen, Will Ma, David Simchi-Levi, and Linwei Xin. 2016. Dynamic recommendation at checkout under inventory constraint. *http://dx.doi.org/10.2139/ssrn.2853093* (2016).

[18] Yanpei Chen, Archana Sulochana Ganapathi, Rean Griffith, and Randy H Katz. 2010. Analysis and lessons from a publicly available google cluster trace. (2010).

[19] V.E. Chenthamarakshan, N. Kambhatla, R.C. Kanjiranthinkal, A.K.R. Singh, and K. Visweswariah. 2012. Systems and methods for matching candidates with positions based on historical assignment data. (May 17 2012). https://www.google.com/patents/US20120123956 US Patent App. 12/944,868.

[20] Alan Demers, Srinivasan Keshav, and Scott Shenker. 1989. Analysis and simulation of a fair queueing algorithm. In *ACM SIGCOMM*.

[21] Arnoud V den Boer. 2014. Dynamic pricing with multiple products and partially specified demand distribution. *Mathematics of operations research* 39, 3 (2014), 863–888.

[22] Nikhil R Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A Wilkens. 2011. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *EC*.

[23] John P. Dickerson, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. 2018. Allocation Problems in Ride-Sharing Platforms: Online Matching with Offline Reusable Resources. In *AAAI*.

[24] John P. Dickerson, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. 2018. Assigning Tasks to Workers Based on Historical Data: Online Task Assignment with Two-sided Arrivals. In *AAMAS*. 318–326.

[25] Jon Feldman, Nitish Korula, Vahab Mirrokni, S Muthukrishnan, and Martin Pál. 2009. Online ad assignment with free disposal. In *WINE*.

[26] C. M. Fortuin, P. W. Kasteleyn, and J. Ginibre. 1971. 2 (1971).

[27] Ali Ghodsi, Vyas Sekar, Matei Zaharia, and Ion Stoica. 2012. Multi-resource fair queueing for packet processing. *ACM SIGCOMM* (2012).

[28] Ali Ghodsi, Matei Zaharia, Scott Shenker, and Ion Stoica. 2013. Choosy: Max-min fair sharing for datacenter jobs with constraints. In *ACM ECCS*.

[29] R. Guedes, V. Furtado, and T. Pequeno. 2014. Multiagent models for police resource allocation and dispatch. In *2014 IEEE Joint Intelligence and Security Informatics Conference*.

[30] Bernhard Haeupler, Vahab S Mirrokni, and Morteza Zadimoghaddam. 2011. Online stochastic weighted matching: Improved approximation algorithms. In *WINE*.

[31] Elad Hazan, Shmuel Safra, and Oded Schwartz. 2006. On the complexity of approximating k-set packing. *computational complexity* 15, 1 (2006).

[32] J. L. Hellerstein. 2010. Google Cloud Trace Dataset. (2010). https://github.com/google/cluster-data

[33] Chien-Ju Ho and Jennifer Wortman Vaughan. 2012. Online Task Assignment in Crowdsourcing Markets.. In *AAAI*.

[34] Yan Huang, Favyen Bastani, Ruoming Jin, and Xiaoyang Sean Wang. 2014. Large Scale Real-time Ridesharing with Service Guarantee on Road Networks. *VLDB Endow.* 7, 14 (2014).

[35] Patrick Jaillet and Xin Lu. 2013. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research* 39, 3 (2013).

[36] Carlee Joe-Wong, Soumya Sen, Tian Lan, and Mung Chiang. 2013. Multiresource allocation: Fairness-efficiency tradeoffs in a unifying framework. *IEEE/ACM TON* (2013).

[37] Ian Kash, Ariel D Procaccia, and Nisarg Shah. 2014. No agent left behind: Dynamic fair division of multiple resources. *JAIR* (2014).

[38] Thomas Kesselheim, Andreas Tönnis, Klaus Radke, and Berthold Vöcking. 2014. Primal beats dual on online packing LPs in the random-order model. In *STOC*.

[39] Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordóñez, and Milind Tambe. 2009. Computing optimal randomized resource allocations for massive security games. In *AAMAS*.

[40] Sang M. Lee, Lori Sharp Franz, and A. James Wynne. 1979. Optimizing State Patrol Manpower Allocation. *Journal of the Operational Research Society* 30, 10 (01 Oct 1979).

[41] S. Ma, Y. Zheng, and O. Wolfson. 2013. T-share: A large-scale dynamic taxi ridesharing service. In *ICDE*.

[42] Will Ma. 2014. Improvements and generalizations of stochastic knapsack and multi-armed bandit approximation algorithms. In *SODA*.

[43] Will Ma and David Simchi-Levi. 2017. Online resource allocation under arbitrary arrivals: Optimal algorithms and tight competitive ratios. *http://dx.doi.org/10.2139/ssrn.2989332* (2017).

[44] Vahideh H Manshadi, Shayan Oveis Gharan, and Amin Saberi. 2012. Online stochastic matching: Online actions based on offline statistics. *MOR* (2012).

[45] Emma Pierson, Camelia Simoiu, Jan Overgoor, Sam Corbett-Davies, Vignesh Ramachandran, Cheryl Phillips, and Sharad Goel. 2017. A large-scale analysis of racial disparities in police stops across the United States. *arXiv preprint arXiv:1706.05678* (2017).

[46] Karthik Abinav Sankararaman and Aleksandrs Slivkins. 2018. Combinatorial Semi-Bandits with Knapsacks. In *AIStats*.

[47] Robert P Shumate and Richard F Crowther. 1966. Quantitative methods for optimizing the allocation of police resources. *J. Crim. L. Criminology & Police Sci.* 57 (1966).

[48] Yaron Singer and Manas Mittal. 2013. Pricing mechanisms for crowdsourcing markets. In *WWW*.

[49] Adish Singla and Andreas Krause. 2013. Truthful incentives in crowdsourcing tasks using regret minimization mechanisms. In *WWW*.

[50] Aleksandrs Slivkins and Jennifer Wortman Vaughan. 2014. Online decision making in crowdsourcing markets: Theoretical challenges. *ACM SIGecom Exchanges* 12, 2 (2014).

[51] Ashwin Subramanian, G Sai Kanth, Sharayu Moharir, and Rahul Vaze. 2015. Online incentive mechanism design for smartphone crowd-sourcing. In *WiOPT*.

[52] Yongxin Tong, Libin Wang, Zimu Zhou, Bolin Ding, Lei Chen, Jieping Ye, and Ke Xu. 2017. Flexible Online Task Assignment in Real-time Spatial Data. *Proc. VLDB Endow.* (2017).

[53] Xinshang Wang, Van-Anh Truong, and David Bank. 2018. Online advance admission scheduling for services with customer preferences. *arXiv preprint arXiv:1805.10412* (2018).

[54] Rong Yang, Benjamin Ford, Milind Tambe, and Andrew Lemieux. 2014. Adaptive resource allocation for wildlife protection against illegal poachers. In *AAMAS*.

[55] Xing Yi, James Allan, and W Bruce Croft. 2007. Matching resumes and jobs based on relevance models. In *SIGIR*.

[56] Yue Yin, Haifeng Xu, Jiarui Gan, Bo An, and Albert Xin Jiang. 2015. Computing Optimal Mixed Strategies for Security Games with Dynamic Payoffs.. In *IJCAI*.

[57] Dong Zhao, Xiang-Yang Li, and Huadong Ma. 2014. How to crowdsource tasks truthfully without sacrificing utility: Online incentive mechanisms with budget constraint. In *INFOCOM*.

Supplementary Materials

A MISSING PROOFS AND THEORY FROM MAIN SECTION

A.1 Proof of Lemma 2.2

The main idea is to show that all the constraints in the above LP are valid for the offline optimal. For each given t and worker j, $\sum_{e \in E_{j,t}} x_{e,t}$ can be interpreted as the sum of the expected number of assignments related to j we could make in the offline optimal, which is surely no larger than the probability that j comes at t. This justifies constraints (6). Any offline algorithm should satisfy the budget constraints as well and by linearity of expectation, we see constraints (7) are valid.

A.2 Performance of NADAP for integral case

Before we prove the main theorem we will start with a weaker analysis which gives the main ideas.

Theorem A.1 (Weaker Analysis). *By choosing $\alpha = \frac{1}{2\ell}$, NADAP achieves an online competitive ratio of at least $\frac{1}{4\ell}$.*

Proof. WLOG assume that $t = T$. Consider an assignment $e \in E_T$. Recall that S_e is the set of resources consumed by e. For each $k \in S_e$, let $\mathcal{S}_k$ be the event that e consumes less than the remaining resource k at T (we call this as *safe* with respect to k, henceforth). We now lower bound the value $\Pr[\wedge_{k \in S_e} \mathcal{S}_k]$. Consider a given $k \in S_e$. Let U_k be the usage of resource k at the beginning of $t = T$ and $X_{e',t'}$ be the indicator random variable if the assignment $e' \in E_{t'}$ is chosen at $t' \in [T-1]$. We have $U_k = \sum_{t'<T} \sum_{e' \in E_{t'}} X_{e',t'} a_{e',k}$. By definition, e is safe with respect to the resource k iff $U_k \leq B_k - 1$. Observe that $\mathbb{E}[X_{e',t'}] \leq \alpha x^*_{e',t'}$. By Markov's inequality we have

$$\Pr[U_k \leq B_k - 1] = 1 - \Pr[U_k \geq B_k] \geq 1 - \alpha \tag{9}$$

Thus, we get

$$\Pr[\wedge_{k \in S_e} \mathcal{S}_k] = \Pr\left[\bigwedge_{k \in S_e} \left(U_k \leq B_k - 1\right)\right] \geq 1 - \ell\alpha \tag{10}$$

Therefore, for a given (e, t), we have that the assignment e is made with probability at least $\alpha x^*_{e,t}(1 - \ell\alpha)$. By setting $\alpha = \frac{1}{2\ell}$, we get that each assignment e is made with probability at least $x^*_{e,t}/(4\ell)$. Hence, from linearity of expectation, we get the ratio of $\frac{1}{4\ell}$. □

A tight analysis for NADAP with unit budget. We first consider a special case when $B_k = 1$ for all $k \in K$ and show a *tight* analysis for NADAP. To obtain this, we first show an example which motivates us for the tighter analysis.

Example A.2. Consider an unweighted star graph $G = (I, J, E)$ where $|I| = 1, |J| = 3, E = (e_1, e_2, e_3)$ with $T = 2$ and $d_1 = T$ (no deadline constraints). Suppose at $t = 1$, $j = 1, 2$ arrives with equal probability $1/2$ and at $t = 2$, $j = 3$ will arrive with probability 1. Let e_1, e_2, e_3 denote respectively the assignment we consider when $j = 1$ comes at $t = 1$, $j = 2$ comes at $t = 1$ and $j = 3$ comes at $t = 2$. Let $K = 2$ with $\mathbf{B} = (1, 1)$ and $\mathbf{a}_{e_1} = (1, 0)$, $\mathbf{a}_{e_2} = (0, 1)$ and $\mathbf{a}_{e_3} = (1, 1)$. Suppose an optimal solution to the LP-(1) is $x^*_{e_1} = x^*_{e_2} = 1/2$ and $x^*_{e_3} = 1/2$. Let us analyze the assignment e_3 when $j = 3$ comes at $t = 2$ by running NADAP(α). According to NADAP(α), at $t = 1$ we will choose e_j with probability α whenever $j = 1$ or $j = 2$ comes. *Notice that at $t = 2$, the first and the second resource are each safe with respective probabilities $1 - \alpha/2$ and both are safe with probability $1 - \alpha$.*

In the above example, we observe that: (a) the events that two different resources are safe are not correlated in the "right" direction to use the FKG inequality [26]. (b) Inequality (9) obtained via Markov's inequality is not tight. Now we use these observations to present a tight analysis for NADAP for the case of unit budget.

Theorem A.3 (Tight analysis of NADAP for unit budget). *By choosing $\alpha = \frac{1}{\ell+1}$, NADAP has an online competitive ratio of $\frac{1}{\ell+1}(1 - \frac{1}{\ell+1})^\ell$ with unit budget.*

Proof. As before, we consider the case that $t = T$ and an assignment $e \in E_T$. For each $t' < T$ and $k \in S_e$, let $E_{k,t'} = \{e' | e' \in E_{t'}, S_{e'} \ni k\}$ be the set of assignments that are available at t' and consume resource k. Let $B_{k,t'}$ be the (random) budget of k at the beginning of t'. Define $F_{k,t'} = (B_{k,t'+1} = 1 | B_{k,t'} = 1)$ and $F_{t'} = \wedge_{k \in S_e} F_{k,t'} = \left(\wedge_{k \in S_e} B_{k,t'+1} = 1 | \wedge_{k \in S_e} B_{k,t'} = 1\right)$.

We have that $\Pr[F_{k,t'}] = 1 - \sum_{e' \in E_{k,t'}} \alpha x^*_{e',t'}$, $\Pr[F_{t'}] \geq 1 - \sum_{e' \in \cup_{k \in S_e} E_{k,t'}} \alpha x^*_{e',t'}$. Therefore,

$$\Pr[\wedge_{k \in S_e} \mathcal{S}_k] = \prod_{t'<T} \Pr[F_{t'}] \tag{11}$$

$$\geq \prod_{t'<T} \left(1 - \sum_{e' \in \cup_{k \in S_e} E_{k,t'}} \alpha x^*_{e',t'}\right) \doteq \prod_{t'<T} (1 - H_{t'}) \tag{12}$$

Here are two useful observations. First, for each $t' < T$, $\cup_{k \in S_e} E_{k,t'} \subseteq E_{t'}$ where $E_{t'}$ is the set of all assignment available at t'. Due to Constraint (2), we have that $\sum_{e \in E_{t'}} x^*_{e,t'} \leq 1$, which implies that $H_{t'} \leq \alpha$. Second, from Constraint (3) we have that $\sum_{t'<T} H_{t'} \leq \alpha\ell$. These two observations lead to the fact that the rightmost expression of Inequality (12) has a minimum value of $(1 - \alpha)^\ell$. Therefore assignment e will be successfully made at t with probability at least $x^*_{e,t} * \alpha * (1 - \alpha)^\ell$. By choosing $\alpha = 1/(\ell + 1)$ and using linearity of expectation, we prove our claim. □

A tight analysis for NADAP **with general integral budget.** We now give a tight analysis for NADAP by extending the result in Theorem A.3 to the case of general integral budget.

Consider the following equivalent model for NADAP(α). Suppose we have K types of balls and for each type $k \in [K]$, the number of balls is $B_k \in \mathbb{Z}_+$. We have a set of choices $E = \{e | e \in E\}$ and each choice is associated with a binary vector $\mathbf{a}_e \in \{0,1\}^K$, with at most ℓ non-zero elements. When we make the choice e, we use one ball of type k whenever $a_{e,k} = 1$. For each time $t \in [T]$, a single choice e will arrive with probability $x^*_{e,t}$ such that $\sum_{e \in E} x^*_{e,t} \le 1$ for each t. At any time t, for any arrival we accept it (non-adaptively) with probability $\alpha \in (0,1)$. Consider a fixed choice e and $t = T$ and let $S_e \subseteq [K]$ be the set of types of balls e takes. For any $k \in S_e$, let $\mathcal{S}_k$ be the event that at $t = T$, we still have at least one ball of type k left. Hence, the adversary solves the following mathematical program: minimize $\Pr[\wedge_{k \in S_e} \mathcal{S}_k]$ subject to (1) $\sum_{t<T} x^*_{e',t} a_{e',k} \le B_k$ for each $k \in S_e$ and (2) $\sum_{e' \in E} x^*_{e',t} \le 1$ for each $t < T$.

The equivalence between this new model and NADAP(α) can be seen as follows: (1) each assignment corresponds to a choice; (2) for some assignment e with deadline t, we set $x^*_{e,t'} = 0$ for all $t' > t$. This makes the two models equivalent.

Let us now show how to solve the mathematical program for the adversary described above. Consider a given $k \in S_e$. Let $E_k = \{e \in E | a_{e,k} = 1\}$ be the set of choices e that participate in the resource constraint of k. Let $x^*_{k,t} = \sum_{e \in E_k} x^*_{e,t}$. Notice that $x^*_{k,t} \le 1$ and at each time t, one choice from E_k arrives with probability $x^*_{k,t}$. Let $A_{k,t}$ be the indicator random variable that one choice from E_k arrives at t and $A_k = \sum_{t \le T-1} A_{k,t}$, which denotes the random number of arrivals of choices from E_k over the first $T-1$ rounds. For an integral $a \ge 0$ and $b \ge 1$, let $p(a,\alpha,b) \doteq \Pr[Z \le b-1]$ where $Z \sim \mathrm{Bi}(a,\alpha)$ (binomial distribution) and we assume $p(a,\alpha,b) = 1$ for any $0 \le a \le b-1$. Consider a given configuration $\mathcal{A} = \{A_k = a_k \in \mathbb{Z}_+ | k \in S_e\}$. We can show the following three lemmas.

Lemma A.4. *(1)* $\Pr[\mathcal{S}_k | A_k = a_k] \ge p(a_k, \alpha, B_k)$ *for each given* $a_k \in \mathbb{Z}_+$*; (2)* $\Pr[\wedge_{k \in S_e} \mathcal{S}_k | \mathcal{A}] \ge \prod_{k \in S_e} p(a_k,\alpha,B_k)$ *for each given configuration* $\mathcal{A} = \{A_k = a_k \in \mathbb{Z}_+ | k \in S_e\}$.

Proof. Consider a given k and A_k. For A_k trials we take a ball per time-step independently with probability at most α. Thus we have at at least $B_k - 1$ balls with probability at least $p(A_k,\alpha,B_k)$. Since the events $\{(\mathcal{S}_k | A_k) | k \in S_e\}$ are positively correlated, applying FKG inequality [26], yields the second inequality. □

Lemma A.5. $\Pr[\wedge_{k \in S_e} \mathcal{S}_k] \ge \prod_{k \in S_e} \exp\left(\mathbb{E}\left[\ln(p(A_k,\alpha,B_k))\right]\right)$.

Proof. First notice that $\Pr[\wedge_{k \in S_e} \mathcal{S}_k] = \mathbb{E}_A\left[\Pr[\wedge_{k \in S_e} \mathcal{S}_k | A]\right]$ by conditioning on the event A. From Lemma A.4, we have that the latter should be at least $\mathbb{E}_A\left[\prod_{k \in S_e} p(A_k,\alpha,B_k)\right]$. Thus

$$\begin{aligned}
\Pr[\wedge_{k \in S_e} \mathcal{S}_k] &= \mathbb{E}_A\left[\Pr[\wedge_{k \in S_e} \mathcal{S}_k | A]\right] \ge \mathbb{E}_A\left[\prod_{k \in S_e} p(A_k,\alpha,B_k)\right] \\
&= \mathbb{E}_A\left[\exp\left(\sum_k \ln(p(A_k,\alpha,B_k))\right)\right] \\
&\ge \exp\left(\sum_k \mathbb{E}\left[\ln(p(A_k,\alpha,B_k))\right]\right) \\
&= \prod_{k \in S_e} \exp\left(\mathbb{E}\left[\ln(p(A_k,\alpha,B_k))\right]\right)
\end{aligned}$$
□

The inequality from the second line to the third line is due to Jensen's inequality. Recall that $A_k = \sum_{t \le T-1} A_{k,t}$ where $A_{k,t}$ is a Bernoulli random variable indicating if a choice $e \in E_k$ arrives at t. Notice that $\mathbb{E}[A_k] = \sum_{t \le T-1} x^*_{k,t} \le B_k$.

Lemma A.6. *For any* $\alpha \in [0, \frac{1}{2}]$ *and integer* $B_k \ge 1$, $\mathbb{E}_{A_k}[\ln(p(A_k,\alpha,B_k))] \ge \ln(1-\alpha)$.

We now show the proof of Lemma A.6. Let us reiterate the lemma in simple language. Suppose we are given an integer $B \ge 1$ and $A = \sum_t A_t$ where each A_t is a Bernoulli random variable with mean $0 \le x_t \le 1$. For a given integral A and α, let $p(A,\alpha,B) \doteq \Pr[Z \le B-1]$ where $Z \sim \mathrm{Bi}(A,\alpha)$ (binomial distribution) and we assume $p(A,\alpha,B) = 1$ for any $0 \le A \le B-1$. Lemma A.6 states the for any given $\alpha \in [0, \frac{1}{2}]$ and integer $B \ge 1$, we have $\mathbb{E}_A[\ln(p(A,\alpha,B))] \ge \ln(1-\alpha)$ if $\sum_t x_t \le B$.

The whole proof consists of the following two Lemmas.

Lemma A.7. $\mathbb{E}_A[\ln(p(A,\alpha,B))]$ *gets minimized when* $A \sim \mathrm{Pois}(B)$ *(Poisson distribution).*

Lemma A.8. *Consider a given* $\alpha \in [0, \frac{1}{2}]$ *and integer* $B \ge 1$ *and assume* $A \sim \mathrm{Pois}(B)$*. We have*

$$\mathbb{E}_A[\ln(p(A,\alpha,B))] \ge \ln(1-\alpha)$$

Proof of Lemma A.7

Proof. Here are two special cases irrespective of configuration of A: when $B = 1$, $\mathbb{E}[\ln(p(A,\alpha,B))] = \ln(1-\alpha)$ and when $\alpha = 0$, $\mathbb{E}[\ln(p(A,\alpha,B))] = 0$. Now we assume $B \ge 2$ and $\alpha > 0$.

Consider the WS where $A = \sum_i X_i$. Arbitrarily select one random variable out, say X with $\mathbb{E}[X] = x > 0$ and let $A' = A - X$. Consider the following perturbation to the WS: split X into two independent Bernoulli random variables X' and X'' and each has mean $x/2$; all the rest keep unchanged. Condition on $A' = k \in Z_+$.

Before perturbation, we see

$$\mathbb{E}[\ln(p(A, \alpha, B))|A' = k] = (1 - x)\ln(p(k, \alpha, B)) + x\ln(p(k + 1, \alpha, B))$$

After perturbation, we see

$$\mathbb{E}[\ln(p(A, \alpha, B))|A' = k] = (1 - \frac{x}{2})^2 \ln(p(k, \alpha, B)) + x(1 - \frac{x}{2})\ln(p(k + 1, \alpha, B)) + \frac{x^2}{4}\ln(p(k + 2, \alpha, B))$$

Let F and F' be the value of $\mathbb{E}[\ln(p(A, \alpha, B))|A' = k]$ before and after perturbation respectively. We have

$$F - F' = \frac{x^2}{4}\left(\ln(p(k + 1, \alpha, B))^2 - \ln(p(k, \alpha, B)) - \ln(p(k + 2, \alpha, B))\right)$$

We have that $F \geq F'$ and $F > F'$ when $k \geq B - 2$. Summarizing all above analysis we reach our claim. □

Proof of Lemma A.8

Proof. Note that

$$\mathbb{E}_A[\ln(p(A, \alpha, B))] = \sum_{A \geq B} \frac{e^{-B}B^A}{A!} \ln\left(\sum_{i=0}^{B-1} \alpha^i (1 - \alpha)^{A-i} \binom{A}{i}\right)$$

Let $F(A, \alpha, B) = \mathbb{E}_A[\ln(p(A, \alpha, B))] - \ln(1 - \alpha)$. Suppose we try to show that for $\alpha \in [0, 1/2]$, F is an increasing function of α for any given B.

$$\begin{aligned}
\frac{dF}{d\alpha} > 0 \quad &\Leftrightarrow \quad \sum_{A \geq B} \frac{e^{-B}B^A}{A!} \frac{(\Pr[Z \leq B - 1])'}{\Pr[Z \leq B - 1]} + \frac{1}{1 - \alpha} > 0 \\
&\Leftrightarrow \quad \sum_{A \geq B} \frac{e^{-B}B^A}{A!} \frac{(\Pr[Z \leq B - 1])'(1 - \alpha)}{\Pr[Z \leq B - 1]} > -1 \\
&\Leftrightarrow \quad \sum_{A \geq B} \frac{e^{-B}B^A}{A!} \frac{(-1)\Pr[Z = B - 1](A - B + 1)}{\Pr[Z \leq B - 1]} > -1 \\
&\Leftrightarrow \quad \sum_{A \geq B} \frac{e^{-B}B^A}{A!} \frac{\Pr[Z = B - 1](A - B + 1)}{\Pr[Z \leq B - 1]} < 1 \\
&\Leftrightarrow \quad \sum_{A \geq B'+1} \frac{e^{-(B'+1)}(B' + 1)^A}{A!} \frac{\Pr[Z = B'](A - B')}{\Pr[Z \leq B']} < 1 \text{ for any } B' \geq 0
\end{aligned}$$

WLOG assume $B \geq 2$ and $B' = B - 1 \geq 1$. (For the case $B = 1$, we can verify that our claim is trivially valid.) Notice that $\Pr[Z = B'] = \binom{A}{B'}\alpha^{B'}(1 - \alpha)^{A-B'}$ and $\Pr[Z = B' - 1] = \binom{A}{B'-1}\alpha^{B'-1}(1 - \alpha)^{A-B'+1}$. Assume $\alpha \leq 1/2$. Thus

$$\frac{\Pr[Z = B']}{\Pr[Z \leq B']} \leq \frac{1}{1 + (1/\alpha - 1)B'/(A - B' + 1)} \leq \frac{1}{1 + B'/(A - B' + 1)} = \frac{A - B' + 1}{A + 1}$$

Therefore by continuing the above analysis, we have

$$\begin{aligned}
&\Leftrightarrow \quad \sum_{A \geq B'+1} \frac{e^{-(B'+1)}(B' + 1)^A}{A!} \frac{\Pr[Z = B'](A - B')}{\Pr[Z \leq B']} < 1 \text{ for any } B' \geq 1 \\
&\Leftarrow \quad \sum_{A \geq B'+1} \frac{e^{-(B'+1)}(B' + 1)^A}{A!} \frac{(A - B')(A - B' + 1)}{A + 1} < 1 \\
&\Leftrightarrow \quad (B' + 1)^{B'} \sum_{i=0}^{\infty} \frac{e^{-(B'+1)}(B' + 1)^i}{(B + 1 + i)!} i(i + 1) < 1 \text{ (set } i = A - B')
\end{aligned}$$

After simplification, the left-hand-side of the last inequality is

$$(B' + 1)^{B'} \sum_{i=0}^{\infty} \frac{e^{-(B'+1)}(B' + 1)^i}{(B + 1 + i)!} i(i + 1) = 1 + \frac{(5 + 3B') - e^{1+B'}(1 + B')^{-(1+B')}\Gamma(3 + B', 1 + B')}{(1 + B')^{-(1+B')}e^{1+B'}\Gamma(3 + B')}$$

We can show that the numerator is negative for any $B' \geq 1$. Let $F(B') = (5 + 3B') - e^{1+B'}(1 + B')^{-(1+B')}\Gamma(3 + B', 1 + B')$. We find that the maximum value of $F(B')$ is -1.5 when $B' \geq 1$. □

We can now prove the competitive ratio for NADAP using the above lemmas.

Theorem A.9 (Tight Analysis for NADAP). *By choosing* $\alpha = \frac{1}{\ell+1}$, ALG_1 *has an online competitive ratio of* $\frac{1}{\ell+1}(1-\frac{1}{\ell+1})^{\ell}$ *with general integral budget.*

Proof. The proof is very similar to that of Theorem A.3. Consider a given assignment e and $t = T-1$ (WLOG). Notice that $\alpha = \frac{1}{\ell+1} \le \frac{1}{2}$. From Lemma A.5 and A.6, we see that $\Pr[\wedge_{k \in S_e} \mathcal{S}_k] \ge (1-\alpha)^{\ell}$. Thus by plugging in $\alpha = \frac{1}{\ell+1}$, we prove our claim. □

A.3 Performance of ADAP for integral case

Proof of Lemma 4.1 The proof is similar to that of Theorem A.1. Consider a given t and $e \in E_t$. Focus on a given $k \in S_e$ and let $U_{k,t}$ be the usage of resource k at the beginning of t. For each $t' < t$ and $e' \in E_{t'}$, let $X_{e',t'}$ be the indicator random variable that e' is chosen at t'. Notice that $U_{k,t} = \sum_{t'<t} X_{e',t'} a_{e',k}$.

Now we prove by induction on t. For the base case $t = 1$, we see $\beta_{e,t} = 1$ for all $e \in E_t$. Thus we claim is valid. Assume our claim works for all $t' < t$, which leads to the fact that for all $e' \in E_{t'}$ with $t' < t$, e' will be made at t' with probability *exactly* equal to $x^*_{e',t'}\gamma$. In other words, $\mathbb{E}[X_{e',t'}] = x^*_{e',t'}\gamma$. Consider the event that e is safe at t with respect to resource k. By Markov's inequality, we have

$$\Pr[U_{k,t} \le B_k - 1] = 1 - \Pr[U_{k,t} \ge B_k] \ge 1 - \gamma$$

Thus we have

$$\Pr[\mathcal{S}_{e,t}] = \Pr\left[\bigwedge_{k \in S_e} \left(U_{k,t} \le B_k - 1\right)\right] \ge 1 - \ell\gamma \ge \gamma$$

The last inequality is valid since $\gamma \le 1/(\ell+1)$.

A.4 Extension to combined integral and non-integral resources

In this section, we show that our algorithms NADAP and ADAP can be extended to the case with both integral and non-integral resources with an assumption of sufficiently large B. In particular, we analyze a lower bound on B obtain a $1-\epsilon$ competitive ratio. WLOG we assume $\mathcal{K}_1 \neq \emptyset$ and $\ell_1 \ge 1$.

Extension of NADAP. In this section, we analyze the performance of NADAP with parameter $\alpha = 1/(\ell_1+1) \le 1/2$ when non-integral resources are included in the model. Recall that in NADAP, each assignment e is made at t non-adaptively with probability at most $\alpha x^*_{e,t}$. Let $X_{e,t}, Y_{e,t}$ be random variables to indicate if assignment e is made at time t and if e is safe at t respectively. Let $Z_{e,t}$ to be a Bernoulli random variable with mean $\alpha x^*_{e,t}$ and *independent* from $Y_{e,t}$. The value of this random variable is determined as follows. For an assignment $e = (i,j)$, when worker j comes at time t if e is not safe, set $Z_{e,t} = 1$ with probability $\alpha x^*_{e,t}/p_{j,t}$ and 0 otherwise. We then have the following observations (1) $X_{e,t} = Y_{e,t} Z_{e,t} \le Z_{e,t}$ (2) For any two random variables $Z_{e,t}$ and $Z_{e',t'}$, they are independent if $t \neq t'$ and negatively correlated if $t = t'$. We now provide a proof sketch for Theorem 5.4.

Proof sketch for Theorem 5.4. Consider a time-step t and an assignment $e \in E_t$. Let $S_1 = S_e \cap \mathcal{K}_1$ and $S_2 = S_e \cap \mathcal{K}_2$. Let $\mathcal{S}_{k,t}$ be the event that e is safe with respect to resource k at t. From the analysis of Theorem 5.1, we have $\Pr[\wedge_{k \in S_1} \mathcal{S}_{k,t}] \ge (1-\alpha)^{\ell_1}$. Thus it remains to show $\Pr[\wedge_{k \in S_2} \mathcal{S}_{k,t}] \ge 1-\epsilon$ when B is an appropriate function of ϵ. Let $U_{k,t}$ be the consumption of resource k at the beginning of t, i.e., $U_{k,t} = \sum_{t'<t}\sum_{e' \in E_{t'}} X_{e',t'} a_{e',k}$. Observe that for each $k \in S_2$,

$$\begin{aligned} 1 - \Pr[\mathcal{S}_{k,t}] &\le \Pr[U_k \ge B_k - 1] \\ &= \Pr[\textstyle\sum_{t'<t}\sum_{e' \in E_{t'}} X_{e',t'} a_{e',t'} \ge B_k - 1] \\ &\le \Pr[\textstyle\sum_{t'<t}\sum_{e' \in E_{t'}} Z_{e',t'} a_{e',t'} \ge B_k - 1] \end{aligned}$$

Let $H_{k,t} = \sum_{t'<t}\sum_{e' \in E_{t'}} Z_{e',t'} a_{e',t'}$. We have that
(1) $\mathbb{E}[H_{k,t}] \le \alpha B_k$ and (2) $\{Z_{e',t'} | e' \in E_{t'}, t' < t\}$ are 1-correlated. Thus, we can essentially treat them as "independent" random variables and apply the Chernoff bound to $\Pr[Z_{k,t} \ge B_k - 1]$ to find the upper bound. Taking an union bound over all $k \in S_2$, we can get the appropriate lower bound on B as a function of ϵ, such that $\Pr[\wedge_{k \in S_2} \mathcal{S}_{k,t}] \ge 1-\epsilon$. □

Extension of ADAP. Suppose we aim for a competitive ratio of $\gamma = \frac{1-\epsilon}{\ell_1+1}$ for ALG_2 where the multiplicative loss ϵ is due to the addition of non-integral resources (we first ignore errors from Monte-Carlo simulations. Later we show how to account for them). This implies for each time t and assignment e, we need to ensure that assignment e is made at t with probability exactly equal to $\frac{1-\epsilon}{\ell_1+1}$. Following the analysis for the integral case, it would suffice to show at each time t, e is safe with probability $\beta_{e,t} \ge \gamma$. Consider a given assignment e and let $\mathcal{S}_{k,t}$ be the event that e is safe at t with respect to the resource k. Let $S_1 = S_e \cap \mathcal{K}_1$ and $S_2 = S_e \cap \mathcal{K}_2$. From the proof of Lemma 4.1, we have that all integral resources are safe at t with probability $\Pr[\wedge_{k \in S_1} \mathcal{S}_{k,t}] \ge 1 - \frac{(1-\epsilon)\ell_1}{\ell_1+1}$. Thus it remains to show that $\Pr[\wedge_{k \in S_2} \mathcal{S}_{k,t}] \ge 1-\epsilon$, which by union bound leads to the fact that $\beta_{e,t} = \Pr[\wedge_{k \in S_e} \mathcal{S}_{k,t}] \ge \gamma = \frac{1-\epsilon}{\ell_1+1}$.

In Section A.4 we show that if B is large enough, all non-integral resources are almost safe throughout T in NADAP by applying Chernoff bound and union bound. However, this does not work for the analysis of ADAP. In fact, the random variables $\{X_{e,t}\}$ themselves can be positively correlated.

To bypass this, we use a technique called *virtual algorithms*. Suppose we run ALG_2 with some parameter γ up to the time t such that for each e' and $t' < t$, $\Pr[X_{e',t'} = 1] = \gamma x^*_{e',t'}$.It suffices to show that $\beta_{k,t} \doteq \Pr[\mathcal{S}_{k,t}] \geq 1 - \epsilon/\ell_2$ for a given e and $k \in S_2$ and thus by union bound over all $k \in S_2$, we are done.

Consider the simple setting where only one non-integral resource k is involved. Suppose we run NADAP with parameter $\alpha = \frac{\gamma}{1-\delta}$ as a virtual algorithm up to time t and let $\beta'_{k,t}(\delta) = \Pr[\mathcal{S}'_{k,t}]$ be the probability that e is safe at time t with respect to resource k in this virtual algorithm. Here $\delta = o(1)$ when $B \to \infty$. We then have the following Lemma.

Lemma A.10. *For any δ with $\beta'_{k,t}(\delta) \geq 1 - \delta$, we have $\beta_{k,t} \geq 1 - \delta$.*

Proof. Consider a feasible δ with $\beta'_{k,t}(\delta) \geq 1 - \delta$. For each e' and $t' < t$, let $X'_{e',t'}$ indicate that assignment e' is made at t' in the virtual algorithm. We have
$\Pr[X'_{e',t'}] = \frac{\gamma x^*_{e',t'}}{1-\delta} \Pr[e' \text{ is safe at } t'] \geq \gamma x^*_{e',t'}$. Notice that in the algorithm ADAP with parameter γ, each assignment e' will be made with probability equal to $\gamma x^*_{e',t'}$. Therefore we claim that in ADAP, $\beta_{k,t} = \Pr[\mathcal{S}_{k,t}] \geq \Pr[\mathcal{S}'_{k,t}] = \beta'_{k,t} \geq 1 - \delta$. ◻

From Lemma A.10, it suffices to set $\delta = \epsilon/\ell_2$ and set B appropriately such that $\beta'_{k,t}(\delta) \geq 1 - \delta$. By re-applying the analysis from the previous section to the virtual algorithm (NADAP) we can get the required ratio for the general case.

A.5 General case without the large budget assumption

Before proving Theorem 5.6, it is natural to ask if the previous algorithms work well for this setting. Here we provide a negative answer.

Lemma A.11. *The competitive ratios (based on LP (1)) of* NADAP *and* ADAP *both can be arbitrarily small on Example (A.12) when ϵ is sufficiently small and M is sufficiently large.*

Example A.12. Consider a star graph $G = (I, J, E)$ where $|I| = 1, |J| = n + 1, E = \{e_j = (1, j) | j \in [J]\}$. Set $d_1 = T$, *i.e.,* , no deadline constraint. Let $K = \ell = 1$, *i.e.,* , we have only one resource. Let a_j be the cost associated with assignment e_j for each j. Set $a_j = \epsilon$ for $j \leq n$ and $a_j = 1$ when $j = n + 1$. The arrival distributions are independent and identical over all rounds, which are specified as follows: $p_{jt} = (1-\epsilon)/n$ for each $j \leq n$ and $t \in [T]$ and and $p_{jt} = \epsilon$ for $j = n + 1$ and $t \in [T]$. Let $T = 1/(2\epsilon - \epsilon^2)$ which is assumed to an integer WLOG. Let w_j be the weight associated with assignment e_j for each j. Set $w_j = 1$ for $j \leq n$ and $w_j = M/\epsilon$ for $j = n + 1$ where M is a fixed large constant.

We can verify that LP (1) has such an optimal solution: $x^*_{e_j,t} = p_j$ for all j and t. The resultant objective value is $\text{OPT}(LP\ (1)) = (1-\epsilon) * T + \epsilon * T * \frac{M}{\epsilon} \geq M * T$. ◻

Proof. Consider NADAP with parameter α. Note that the weight of e_{n+1} dominates over the sum of rest and thus it will suffice to analyze the assignment $e = e_{n+1}$. According to NADAP, e_{n+1} is safe at t iff none assignment is made before t, which occurs with probability $(1-\alpha)^{t-1}$. Therefore the asymptotical performance of e_{n+1} (the total weights obtained from assignment e_{n+1}) when $\epsilon \to 0, T \to \infty$ is as follows:

$$\lim_{T\to\infty} \sum_{t=1}^{T} w_{n+1} \alpha x^*_{e,t} (1-\alpha)^{t-1} = M$$

Thus the total weights of all assignments made in NADAP with parameter α will be at most $\alpha(1-\epsilon) * T + M$. Note that the optimal value of LP (1) is at least $M * T$, which implies that NADAP achieves a competitive ratio at most $1/M$ when $T \to \infty$. Since M is an arbitrarily fixed large constant, we claim that the performance of NADAP can be arbitrarily bad based on the benchmark (1).

Now we consider ADAP with parameter γ. In order to make ADAP work with γ, we need to ensure that the probability that each assignment e is safe at t $\Pr[\mathcal{S}_{e,t}] \geq \gamma$ for all e and t. Consider $e = e_{n+1}$ and notice that e is safe at t iff none assignment is made before t. Thus we see e is safe at t with probability equal to $(1-\gamma)^{t-1}$. To make ADAP work with γ, we need to set $\gamma \leq (1-\gamma)^{T-1}$. This implies the choice γ will approach to o when $T \to \infty$. ◻

We now restate the two parts of Theorem 5.6 as the following two lemmas.

Lemma A.13. *By choosing $\gamma_1 = 0.943$, $\gamma_2 = 0.739$ and $\beta = 0.323$,* $\text{ALG}_3(\beta, \gamma_1, \gamma_2)$ *achieves a ratio of $\frac{1}{4}$ based on* LP-(5).

Lemma A.14. *No adaptive algorithm can achieve a ratio that is at least $\frac{e-1}{2e-1}$ for* MBOA-S *using LP-(5) as the benchmark.*

Proof of Lemma A.13

Proof. For each e, let k_e be the unique resource it requires.

First consider a big assignment e. We see e is safe at t iff none assignment from $\mathrm{BIG}(k_e)$ is added so far. Then we have that each big edge e will be added with probability

$$\Pr[e \text{ is added}] = \frac{\gamma_1 x_e}{T} \sum_{t=1}^{T\beta} \left(1 - \frac{\sum_{e' \in \mathrm{BIG}(k_e)} x_{e'} \gamma_1}{T}\right)^{t-1} \geq x_e \left(1 - (1 - \frac{\gamma_1}{T})^{T\beta}\right) \geq x_e \left(1 - e^{-\gamma_1 \beta}\right) \tag{13}$$

Notice that the first inequality in the above expression is mainly due to Constraint (8), which is $\sum_{e' \in \mathrm{BIG}(k_e)} x_{e'} \leq 1$. Now consider a small assignment e. Observe that a sufficient condition for e is safe at t is that (1) none assignments from $\mathrm{BIG}(k_e)$ is added before (denoted by A_1); (2) the consumption of budget k_e from small assignments in $\mathrm{SM}(k_e)$ is no more than $\alpha = 1/2$ (denoted by A_2) and we just consider the worst case when $a_e = \frac{1}{2}$. Notice that

$$\Pr[A_1] = \left(1 - \frac{\sum_{e' \in \mathrm{BIG}(k_e)} x_{e'} \gamma_1}{T}\right)^{T\beta} \sim \exp(-\gamma_1 \beta \tau_b)$$

where $\tau_b = \sum_{e' \in \mathrm{BIG}(k_e)} x_{e'} \leq 1$, which denotes the sum of $x_{e'}$ for all big assignments in the constraint k_e.

Count the remaining $T(1-\beta)$ rounds from $t' = 1$ to $t' = T(1-\beta)$. Now we try to bound A_2. Notice that each single round the expected consumption of budget k_e from small assignments in k_e should be $\sum_{e' \in \mathrm{SM}(k_e)} \frac{a_{e'} x_{e'} \gamma_2}{T} \doteq \frac{\tau_s \gamma_2}{T}$, where $\tau_s = \sum_{e' \in \mathrm{SM}(k_e)} a_{e'} x_{e'}$, which is the expected cost for all small assignments in k_e. Therefore after Δt rounds, the total consumption is no larger than $1/2$ should be at least

$$\Pr[A_2] \geq 1 - \frac{2\Delta t \tau_s \gamma_2}{T}$$

Combining A_1 and A_2 together, we see the expected number of times that e will be added should be at least

$$\frac{x_e \gamma_2}{T} \exp(-\gamma_1 \beta \tau_b) \sum_{t=0}^{T(1-\beta)-1} \left(1 - \frac{2t\tau_s \gamma_2}{T}\right) = x_e \exp(-\gamma_1 \beta \tau_b) \gamma_2 \Big((1-\beta) - (1-\beta)^2 \gamma_2 \tau_s\Big) \tag{14}$$

Recall that τ_b is the sum of $x_{e'}$ values over all large e' in k_e while τ_s is the sum of mass $\sum_{e'} a_{e'} x_{e'}$ over all small e' in k_e. Thus we have $0 \leq \tau_b \leq 1$ and $\frac{1}{2} * \tau_b + \tau_s \leq 1$ from Constraint (8) and (7) respectively.

The ratio for a big and small assignment are lower bounded respectively in Inequality (13) and (14) (after removing the term x_e from each). Therefore to maximize the final ratio achieved by $\mathsf{ALG}_3(\beta, \gamma_1, \gamma_2)$, we essentially need to maximize the minimum of the two, which is summarized in the following program:

$$\max_{\gamma_1, \gamma_2, \beta} \min \left\{ 1 - e^{-\gamma_1 \beta}, \min \left[\exp(-\gamma_1 \beta \tau_b) \gamma_2 (1-\beta) \Big(1 - (1-\beta)\gamma_2 \tau_s\Big), s.t.\ 0 \leq \tau_b \leq 1, \frac{1}{2} * \tau_b + \tau_s \leq 1 \right] \right\} \tag{15}$$

From Lemma A.15, Program (15) has an optimal value of $\frac{1}{4}$, which arrives at $\gamma_1 = 0.943$, $\gamma_2 = 0.739$ and $\beta = 0.323$. Thus we are done. □

Lemma A.15. *The maxmin program* (15) *has an optimal value of* $\frac{1}{4}$.

Proof. For the inside minimization program, consider only two boundary points: $(\tau_b, \tau_s) = (0, 1)$ and $(\tau_b, \tau_s) = (1, 1/2)$. Replace the inside minimization program with minimum value over these two points. Consider the relaxed version of the maxmin program (15) as follows:

$$\begin{aligned} &\max t \\ &t \leq 1 - e^{-\gamma_1 \beta} \\ &t \leq \gamma_2 (1-\beta)(1 - \gamma_2(1-\beta)) \\ &t \leq e^{-\gamma_1 \beta} \gamma_2 (1-\beta) \Big(1 - \gamma_2 (1-\beta) \frac{1}{2}\Big) \\ &0 \leq \gamma_1, \gamma_2, \beta \leq 1 \end{aligned} \tag{16}$$

The above program has an optimal value $1/4$ when $\gamma_1^* = 0.94328$, $\gamma_2^* = 0.73855$ and $\beta^* = 0.323004$. We can prove the optimal value of Program (15) is upper bounded by that of Program (16). Fix the value γ_1^*, γ_2^* and β^*, we can verify that the inside minimization program has an optimal value of $1/4$ as well. This implies that the optimal value of Program (15) is least $1/4$. Summarizing all analysis, we reach our conclusion.

□

Proof of Lemma A.14

Example A.16. Consider such an instance of MBOA-S: $G = (I, J, E)$ with $|I| = 1, |J| = 2, E = \{e_j = (1, j) | j \in [J]\}$. Let $K = \ell = 1$, *i.e.*, , we have only one resource. Let (a_1, a_2) and (w_1, w_2) be the costs and weights respectively for the two assignments. Set $a_1 = \epsilon$ and $a_2 = 1$. The arrival distributions over each round are as follows: $p_1 = 1 - \alpha\epsilon$ and $p_2 = \alpha\epsilon$ where $\alpha > 0$ is a constant. Set $T = \frac{1}{\epsilon(1+\alpha) - \alpha\epsilon^2}$. We can verify that LP (5) has such an optimal solution: $x_1^* = T * (1 - \alpha\epsilon)$ and $x_2^* = T * (\alpha * \epsilon)$. Notice that we assume ϵ is sufficiently small and thus $x_2^* \leq 1$. The resultant objective value is $w_1 x_1^* + w_2 x_2^*$ for whatever w_1 and w_2. □

Lemma A.17. *Assume* $w_2\left(1-e^{-\alpha\epsilon T}\right) \geq w_1 T(1-\alpha\epsilon)$. *Then no online algorithm could have a performance better than* $w_2\left(1-e^{-\alpha\epsilon T}\right)$ *on the instance shown in Example A.16.*

Proof. Let us refer to the two assignments e_1 and e_2 as *small* and *big* respectively. Focus on an online algorithm ALG. Suppose we are at t and assume prior to t, we make no any assignment. ALG essentially need to answer such a question: if a small comes, we should reject it or accept. (of course we should accept the big if it comes.) We can show that ALG will always reject the small if it comes at any t.

Let $T' = T - t + 1$. Assume a small comes at t. If ALG accepts it, then it has to stick on accepting the small ones in all future rounds. Thus the expected weight should be $T' * (1-\alpha\epsilon) * w_1$. If ALG rejects it and reserve the budget for a big one, then the expected weight should be $\left(1-(1-\alpha*\epsilon)^{T'}\right) * w_2$. Notice that

$$\frac{\left(1-(1-\alpha*\epsilon)^{T'}\right)w_2}{T'(1-\alpha\epsilon)w_1} \geq \frac{\left(1-e^{-\alpha\epsilon T'}\right)w_2}{T'(1-\alpha\epsilon)w_1} \geq \frac{\left(1-e^{-\alpha\epsilon T}\right)w_2}{T(1-\alpha\epsilon)w_1} \geq 1$$

Thus we claim that ALG will get an expected weight of $w_2\left(1-e^{-\alpha\epsilon T}\right)$ by always rejecting a small and accept a big one once it comes. □

Now we start to prove Lemma A.14.

Proof. Focus on the instance shown in Example A.16. From the analysis in Lemma A.17, we see no online algorithm could achieve a ratio better than $\frac{w_2\left(1-e^{-\alpha\epsilon T}\right)}{w_1 x_1^* + w_2 x_2^*}$ if using LP (5) as the benchmark. Therefore we essentially need to solve the below minimization program:

$$\left\{\min_{w_1\geq 0,\, w_2\geq 0} \frac{w_2\left(1-e^{-\alpha\epsilon T}\right)}{w_1 x_1^* + w_2 x_2^*} \quad s.t.\; w_2\left(1-e^{-\alpha\epsilon T}\right) \geq w_1 T(1-\alpha\epsilon)\right\} \tag{17}$$

We can verify that when $w_2\left(1-e^{-\alpha\epsilon T}\right) = w_1 T(1-\alpha\epsilon)$, $\alpha \longrightarrow \infty$ and $\epsilon \longrightarrow 0$, the optimal objective value will approach to $\frac{e-1}{2e-1} \sim 0.387$. □

B MISSING DETAILS IN EXPERIMENTAL SECTION AND FURTHER EXPERIMENTAL RESULTS

B.1 Public Safety Application on Policing Dataset

Here is the list of offenses in the dataset.

(1) DUI
(2) Registration and plates related issues
(3) Paperwork related issues
(4) License related issues
(5) Not wearing seat belt
(6) Illegal Equipment in the vehicle
(7) Speeding
(8) Skipping Stop sign
(9) Skipping red light
(10) Unsafe Maneuvering of the vehicle

Exploratory analysis. We first describe some exploratory analysis on the dataset, which justify our modeling assumptions. The first *inexplicit* assumption is that the arrival probabilities for different online types are different (as opposed to the random arrival order). In Figure 5a we show the total number of arrivals (on a randomly chosen 10-day period) for 600 randomly chosen types from the dataset. It can be seen that there is a huge variance in the number of arrivals among different types. Hence this justifies different arrival probabilities for different types. The second assumption is *Known Adversarial Distributions* (KAD). In other words, the distributions change in different time-steps. To justify this, in Figure 5b we pick a random online type and plot the arrival rates for a randomly chosen 90-day period. It can be seen that there is a large variance in the arrival rates among different time-steps which again validates the KAD assumption.

B.2 Experiments on Synthetic Dataset

Throughout this sub-section we refer to ALG_1 as our algorithm NADAP and NAdap as the algorithm SCALED.

We now describe how we generate our synthetic dataset.

(1) For each j, recall that $N(j)$ is the set of tasks that interest j. We generate $N(j)$ by sampling each $i \in [m]$ independently with some probability, say 0.3.
(2) Let P_1 be the arrival probability matrix of size $n \times T$ such that $P_1(i,j) = p_{i,j}$. We first generate a random "seed" matrix P_0 of size $n \times T_1$ such that for each $t \in [T_1]$, the values in the t^{th} column of P_0 are uniformly distributed over $[0,1]$ *conditioning on the column sum is*

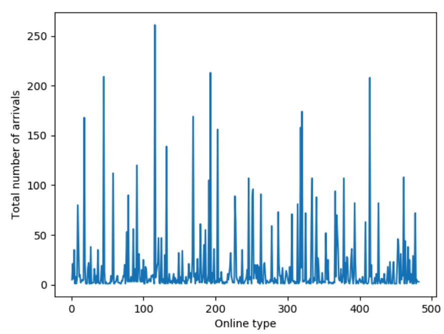

(a) Plot indicating the total number of arrivals in a given 24-hour period among 600 randomly chosen online-types

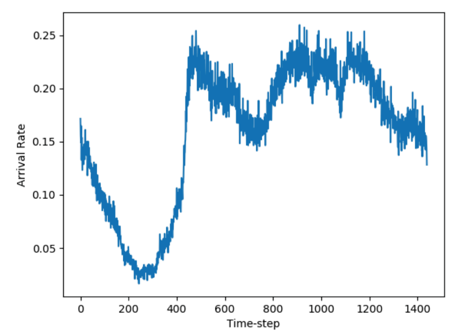

(b) Plot showing the arrival rates of a randomly chosen online-type over $24 * 60 = 1440$ time-steps

1, i.e., $\sum_t P_0(i,t) = 1$. We achieve this by running the file "randfixedsum.m" due to Roger Stafford [6]. Once we have a fixed P_0, we generate P_1 by sampling one column from P_0 uniformly for T times. Notice that if we generate P_1 in the direct way as P_0, then each j will have almost the same arrivals over T rounds since T assumes to be very large. In our case we set $T_1 = m \ll T$ and we hope we can create some potential bias of the arrivals over all $j \in [n]$ and that can pass to P_1.

(3) Let E be the set of assignments generated as shown in the first point. For each assignment $e \in E$, we independently choose a uniform value $w_e \in [0, 1]$.
(4) Recall that $\mathcal{K}_1$ and $\mathcal{K}_2$ are the set of integral and non-integral resources respectively. We generate a budget B_k by uniformly sampling an integer from $[UB] = \{1, 2, 3, \ldots, UB\}$ for each $k \in \mathcal{K}_1$ and from $[LB, 5 * LB]$ for each $k \in \mathcal{K}_2$ respectively. Here UB and LB are parameters specified in advance.
(5) Recall S_e is the set of resources requested by e. For each e, we first generate a random permutation π_1 over $\mathcal{K}_1$ and then set $S_e \cap \mathcal{K}_1$ as the first $\lceil \rho_0 * K_1 \rceil$ elements of π_1. Set $a_{e,k} = 1$ for each $k \in S_e \cap \mathcal{K}_1$. We then generate another random permutation π_2 over $\mathcal{K}_2$ and set $S_e \cap \mathcal{K}_2$ as the first $\lceil \rho_0 * K_2 \rceil$ elements of π_2. Sample a uniform value from $[0, 1]$ for $a_{e,k}$ for each $k \in S_e \cap \mathcal{K}_2$. Here $\rho_0 \in [0, 1]$ is a parameter given in advance.
(6) For each e, let d_e be the deadline of e. We sample a random integer from $[T/2, T]$ uniformly as d_e for each $e \in E$. In this experiment we consider a relative more flexible setting: allow assignments with respect to a single task to have potentially distinctive deadlines.

We start with the case when only integral resources are involved. In these experiments, we show that the performance of our main algorithm is significantly better than the theoretical worst case bounds (such bounds hold only for some extremely specialized cases such as the one shown in Example 5.3).

Our experimental setup is as follows. Theorem 5.1 shows that $\text{ALG}_1(\frac{1}{\ell+1})$ can achieve a ratio at least $\frac{1}{\ell+1}\frac{1}{e}$. Our experimental results suggest that it will be too conservative for the choice of $\alpha = \frac{1}{\ell+1}$.

For each set of parameters $\mathcal{P} = (m, n, K_1, K_2, T, UB, LB, \rho_0)$, we generate a set $\mathcal{I}(\mathcal{P})$ of 5 random instances as described before. For each instance $I \in \mathcal{I}(\mathcal{P})$, we run the above five algorithms each on I for 100 times and take the mean as the final performance. For each given

[6]https://www.mathworks.com/matlabcentral/fileexchange/9700-random-vectors-with-fixed-sum/content/randfixedsum.m

instance I, let $\text{OPT}(I)$ be the LP optimal value and $\text{ALG}(I)$ be the final performance on I. We define $\rho(\text{ALG}, I) = \text{ALG}(I)/\text{OPT}(I)$, which is the ratio of performance of ALG to the LP value on I. For each set of parameters $\mathcal{P} = (m, n, K_1, K_2, T, UB, LB, \rho_0)$, we generate 5 random instances as described before and set the mean ratio as $\rho(\text{ALG}, \mathcal{P})$ for each ALG. The results can be seen in Figures 6 and 7.

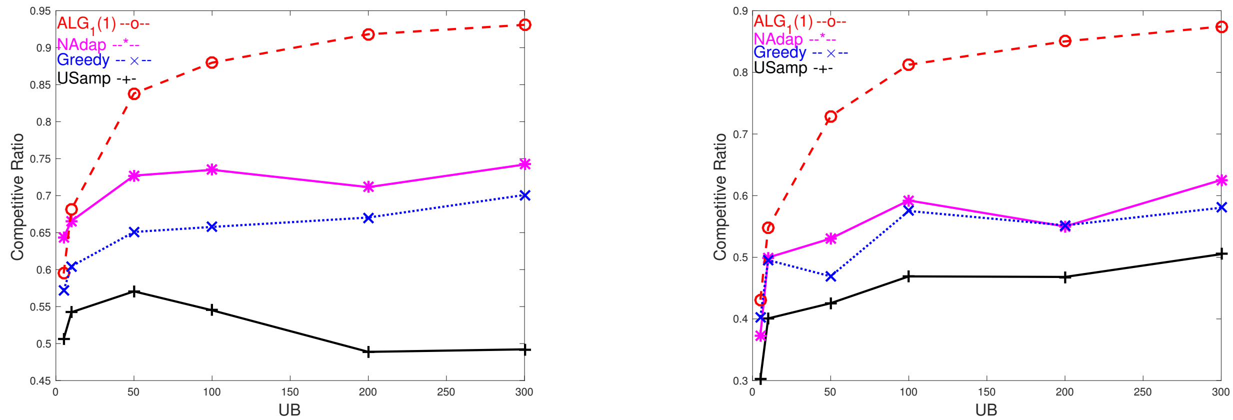

(a) $m = 10$, $n = 50$, $K_1 = 90$, $K_2 = 0$, $T = 3000$, $\rho_0 = 0.1$

(b) $m = 10$, $n = 50$, $K_1 = 90$, $K_2 = 0$, $T = 3000$, $\rho_0 = 0.5$

Figure 6: Performance of the four algorithms as UB increases where: $m = 10$, $n = 50$, $K_1 = 90$, $K_2 = 0$, $T = 3000$, $\rho_0 = 0.1$ (Left) and $\rho_0 = 0.5$ (Right). The best LP-based heuristic $\text{ALG}_1(1)$ **(red-colored) strictly beats the best non-LP-based strategy** Greedy **(blue-colored).**

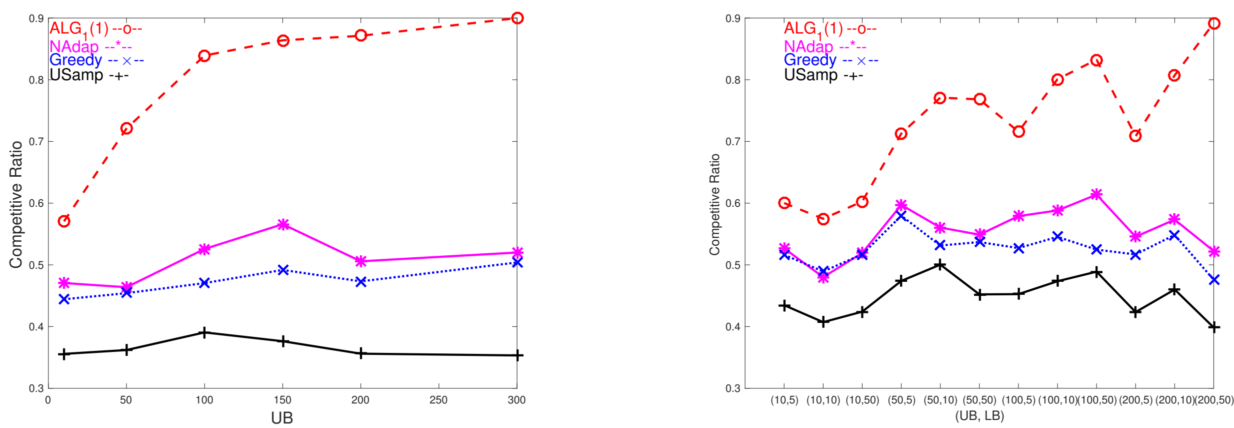

(a) $m = 20$, $n = 100$, $K_1 = 50$, $K_2 = 0$, $T = 3000$, $\rho_0 = 0.5$

(b) $m = 10$, $n = 50$, $K_1 = 50$, $K_2 = 40$, $T = 2000$, $\rho_0 = 0.5$

Figure 7: Performance of the four algorithms as UB increases when only integral resources are involved (Left) and as (UB, LB) increase when both integral and non-integral are involved (Right). The best LP-based heuristic $\text{ALG}_1(1)$ **(red-colored) strictly beats the best non-LP-based strategy** Greedy **(blue-colored).**

Observations from our synthetic experiments

(1) In nearly all instances, our LP-based heuristic $\text{ALG}_1(1)$ dominates the other three algorithms. Also, Greedy dominates the other non-LP-based strategy USamp universally.

(2) The competitive ratios of all the four algorithms seem insensitive to the number of resources requested by each assignment when only integral resources are involved. That can be seen from the results in Figure 6a and 6b.

(3) When only integral resources are involved, both of $\text{ALG}_1(1)$ and Greedy have overall increasing ratios as UB increases. The results in Figures 6 and 7a show that ratio of $\text{ALG}_1(1)$ increases much faster than that of Greedy as UB increases. In particularly, Figure 7a shows that the ratio gap between $\text{ALG}_1(1)$ and Greedy can increase to nearly 0.4 when $UB = 300$.

(4) When both of integral and non-integral resources are involved, $\text{ALG}_1(1)$ has an overall improving ratios when UB and LB increase. That constrasts sharply with the fact that Greedy has a nearly steady ratio ranging from $[0.47, 0.57]$. See Figure 7b for more details.