Learning Unknown Service Rates in Queues: A Multi-Armed Bandit Approach

Subhashini Krishnasamy[1], Rajat Sen[1], Ramesh Johari[2], and Sanjay Shakkottai[1]

[1]The University of Texas at Austin
[2]Stanford University

Abstract

Consider a queueing system consisting of multiple servers. Jobs arrive over time and enter a queue for service; the goal is to minimize the size of this queue. At each opportunity for service, at most one server can be chosen, and at most one job can be served. Service is successful with a probability (the *service probability*) that is *a priori unknown* for each server. An algorithm that knows the service probabilities (the "genie") can always choose the server of highest service probability. We study algorithms that learn the unknown service probabilities. Our goal is to minimize *queue-regret*: the (expected) difference between the queue-lengths obtained by the algorithm, and those obtained by the "genie."

Since queue-regret cannot be larger than classical regret, results for the standard multi-armed bandit problem give algorithms for which queue-regret increases no more than logarithmically in time. Our paper shows surprisingly more complex behavior. In particular, as long as the bandit algorithm's queues have relatively long regenerative cycles, queue-regret is similar to cumulative regret, and scales (essentially) logarithmically. However, we show that this "early stage" of the queueing bandit eventually gives way to a "late stage", where the optimal queue-regret scaling is $O(1/t)$. We demonstrate an algorithm that (order-wise) achieves this asymptotic queue-regret in the late stage. Our results are developed in a more general model that allows for multiple job classes as well.

1 Introduction

Stochastic multi-armed bandits (MAB) have a rich history in sequential decision making (Gittins (1979), Mahajan and Teneketzis (2008), Bubeck and Cesa-Bianchi (2012)). In its simplest form, at each discrete time step the decision maker must choose a single arm from a collection of K arms. A random binary reward (i.e., a Bernoulli random variable taking value 0 or 1) is accrued each time an arm is pulled; if a reward is received, we refer to the outcome as a "success." The probabilities of success vary across arms, and are unknown a priori. The MAB problem is to determine which arm to choose at each time in order to minimize the *cumulative expected regret*: the cumulative loss of expected reward when compared to a genie that has knowledge of the arm success probabilities.

In this paper[1], we study a variant of this problem motivated by *queueing* applications, where success probabilities may be unknown.

To fix ideas, consider the problem of scheduling in a discrete-time queueing system with a single queue and multiple servers. Jobs arrive to the system and are stored in the queue until they are successfully served. At any time, only one server can be active and can serve at most one job in the queue. The active server is either successful (in which case the job departs the system) or

arXiv:1604.06377v4 [cs.SY] 21 Nov 2019

unsuccessful (in which case the job remains in the queue); success probabilities vary across servers. We consider this problem in the case where the success probabilities (a.k.a. service probabilies) are not known *a priori*.

We view this problem as a variant of the basic MAB model, where each arm is now a *server* that can serve a waiting job. From this perspective, the stochastic reward described above is equivalent to *service*, which takes binary values (1 or 0) depending on whether the job was successfully served or not. At any time the queue is empty, the only possible reward is 0. We observe that the basic MAB model is not suitable for queueing applications as it fails to capture an essential feature of service in many settings: in a queueing system, *jobs wait until they complete service.*

Such systems are *stateful*: when the chosen arm results in zero reward, the job being served remains in the queue. On the other hand, when there are no jobs waiting in the queue, there can be no accrual of reward. This makes it essential in this model to track the number of jobs waiting in the queue to be served. The *queue length*, which is the difference between the cumulative number of arrivals and departures, is the most common measure of the quality of the service strategy being employed.

Our paper is motivated by combining two well-established lines of research: first, the ubiquity of queueing as a means to model service systems; and second, the prevalence of MAB models as a way to study the balance between learning and performance optimization in a wide range of settings. Our paper develops an MAB model to study the role of unknown service rates on performance in scheduling systems. This problem clearly has the explore-exploit tradeoff inherent in the basic MAB problem: since the service probabilities across different servers are unknown, there is a tradeoff between learning (*exploring*) the different servers and scheduling (*exploiting*) the most promising server from past observations. For brevity, we refer to this problem as the *queueing bandit*. Since the queue length is simply the difference between the cumulative number of arrivals and departures, the natural notion of regret here is to compare the expected queue length under a bandit algorithm with the corresponding one under a genie policy that always chooses the arm with the highest expected reward.

Our results are actually developed for a switch network, i.e., a multi-class parallel-server system; this is a more general model than the single queue system described above. Specifically, we allow the possibility that there are U distinct classes of jobs with a queue for each class. Each of the job queues may be served by any of the K ($K \geq U$) servers, but the service probabilities are heterogeneous across queue-server pairs (also referred to as *links*). At each time, at most one queue can be assigned to each server, and at most one server can be assigned to each queue; thus in each time slot, scheduling amounts to choosing a matching in the complete bipartite graph between queues and servers. Arrivals to the queue and possible service offered by the links follow a product Bernoulli distribution, and are i.i.d. across time slots. A job remains in the queue if not successfully served; service on the same job can be attempted by different servers at different times. Statistical parameters (service probabilities) corresponding to the service distributions are considered unknown.

In the single queue setting, i.e., when $U = 1$, we assume that there is at least one server that has a service probability higher than the arrival probability. This ensures that the optimal "genie" policy—i.e., the policy with full *a priori* knowledge of all the service probabilities—is stable. When $U > 1$, we assume that for the true service probabilities, there exists a single unique matching (the *optimal matching*) which is strictly better than all other matchings for all queues. For $U > 1$, this assumption ensures that the optimal "genie" policy is one that selects the optimal matching in every time slot.

Let $\boldsymbol{Q}(t)$ be the queue length vector at time t under a given bandit algorithm, and let $\boldsymbol{Q}^*(t)$ be the corresponding queue length vector under the "genie" policy that always chooses the optimal

matching. We define the *queue-regret* vector $\boldsymbol{\Psi}(t)$ as the difference in expected queue lengths for the two policies:

$$\boldsymbol{\Psi}(t) := \mathbb{E}\left[\boldsymbol{Q}(t) - \boldsymbol{Q}^*(t)\right]. \tag{1}$$

Note that the difference in queue-lengths is equal to the difference in the total number of departures. Interpreting a departure as reward 1, the queue-regret $\boldsymbol{\Psi}(t)$ is the expected difference of the accumulated rewards. This is similar to the traditional MAB regret but with a slight difference: here, it is possible to accrue reward only if the queue is non-empty. As an example, in a single queue setting, suppose the number of waiting jobs is infinite; in that case, the queue-regret is identical to the regret in the basic MAB problem. Our goal is to develop bandit algorithms that lead to small queue-regret for each $u \in [U]$ at a finite time t. To develop some intuition, we compare this with the standard stochastic MAB problem; we focus on the setting of a single queue ($U = 1$) for ease of exposition, but the same insights hold in the general case of a switch network with a unique optimal matching when we consider each component of the queue-regret individually. For the standard MAB problem, well-known algorithms such as UCB, KL-UCB, and Thompson sampling achieve a cumulative regret of $O((K-1)\log t)$ at time t (Auer et al. (2002), Garivier and Cappé (2011), Agrawal and Goyal (2012)), and this result is essentially tight (Lai and Robbins (1985)). In the queueing bandit, it can be shown that the queue-regret cannot be any higher than the traditional regret (where a reward is accrued at each time whether a job is present or not). This leads to an upper bound of $O((K-1)\log t)$ for the queue-regret.

However, this upper bound does not tell the whole story for the queueing bandit: we show that there are two "stages" to the queueing bandit. If the arrival probability of a queue is higher than the service probabilities of all but the corresponding best server, the bandit algorithm is unable to even stabilize the queue in the *early* stage – i.e., on average, the queue length increases over time and is continuously backlogged; therefore the queue-regret grows similarly to the traditional regret. Once the algorithm is able to stabilize the queue — i.e., in the *late* stage — a dramatic shift occurs in the behavior of the queue-regret. A stochastically stable queue goes through **regenerative cycles** – a random cyclical behavior where queues build-up over time, then empty, and the cycle repeats. The associated recurring "zero-queue-length" epochs means that sample-path queue-regret essentially "resets" at (stochastically) regular intervals; i.e., the sample-path queue-regret becomes non-positive at these time instants. Thus the queue-regret should fall over time, as the algorithm learns.

1.1 Contributions

Our main results provide lower bounds on each component of the queue-regret vector in both the early and late stages, as well as algorithms that essentially match these lower bounds. Here, we discuss these results in the context of a single queue, but prove them in the more general switch network context described above.

Below we describe our main contributions.

The late stage: We first consider what happens to the queue-regret as $t \to \infty$. As noted above, a reasonable intuition for this regime comes from considering a standard bandit algorithm, but where the sample-path queue-regret "resets" at time points of regeneration. This is inexact since the optimal queueing system and bandit queueing system may not regenerate at the same time point; nevertheless, the intuition proves correct. In this case, the queue-regret is approximately a (discrete) *derivative* of the cumulative regret. Since the optimal cumulative regret scales like $\log t$, asymptotically the optimal queue-regret should scale like $1/t$. Indeed,

3

we show that the queue-regret for α-consistent policies (please see Definitions 1 and 2 in Section 4 for the definition of an α-consistent policy) is at least C/t infinitely often, where C is a constant independent of t. Further, we introduce scheduling algorithms called Q-UCB and Q-ThS, which are variants of the UCB1 and Thompson sampling algorithms tailored to the queueing bandit. We show an asymptotic regret upper bound of $O\left(\text{poly}(\log t)/t\right)$ for both these algorithms, thus matching the lower bound up to poly-logarithmic factors in t. The key feature of Q-UCB and Q-ThS is that they use *forced exploration*: they exploit the fact that the queue regenerates regularly to explore more systematically and aggressively.

The early stage: The preceding discussion might suggest that an algorithm that explores aggressively would dominate any algorithm that balances exploration and exploitation. However, an overly aggressive exploration policy will preclude the queueing system from ever stabilizing, which is *necessary* to induce the regenerative cycles that lead the system to the late stage. To even enter the late stage, therefore, we need an algorithm that exploits enough to actually stabilize the queue (i.e., we must choose good arms sufficiently often so that the service probability exceeds the arrival probability).

We refer to the early stage of the system, as noted above, as the period before the algorithm has learned to stabilize the queues. For a *heavily loaded system, where the arrival probability approaches the service probability of the optimal server,* we show a lower bound of $\Omega(\log t/\log\log t)$ on the queue-regret in the early stage. Thus up to a $1/\log\log t$ factor, the early stage regret behaves similarly to the cumulative regret (which scales like $\log t$). The heavily loaded regime is a natural asymptotic regime in which to study queueing systems, and has been extensively employed in the literature; see, e.g., Whitt (1974), Kushner (2013) for surveys.

Our results constitute the first insight into the behavior of regret in this queueing setting; as emphasized, it is quite different from minimization of cumulative regret in the standard MAB problem. The preceding discussion highlights why minimization of queue-regret presents a subtle learning problem. On one hand, if the queue has been stabilized, the presence of regenerative cycles allows us to establish that queue-regret must eventually decay to zero at rate $1/t$ under an optimal algorithm (the late stage). On the other hand, to actually have regenerative cycles in the first place, a learning algorithm needs to exploit enough to actually stabilize the queue (the early stage). Our analysis not only characterizes regret in both regimes, but also essentially exactly characterizes the transition point between the two regimes. In this way the queueing bandit is a remarkable new example of the tradeoff between exploration and exploitation.

2   Related Work

MAB algorithms: MAB models have been widely used in the past as a paradigm for various sequential decision making problems in industrial manufacturing, communication networks, clinical trials, online advertising and webpage optimization, and other domains requiring resource allocation and scheduling; see, e.g., Gittins (1979), Mahajan and Teneketzis (2008), Bubeck and Cesa-Bianchi (2012).

The classical MAB problem is based on the stochastic MAB model where the rewards of the arms are i.i.d. across time. The goal is to minimize the expected cumulative loss of reward relative to a "genie" policy that always chooses best arm; this objective is called *regret*. There is a vast body of literature which aims at getting the best possible *finite time* lower bounds for

regret and designing algorithms that match the lower bound (see Bubeck and Cesa-Bianchi (2012) for a survey). Lai and Robbins (1985) prove a lower bound of $O(\log t)$ with the scaling constant depending on the mean rewards of the best and the second best arm.

A variant of the MAB problem that is related to our problem is the Markovian bandit in which the state of the each arm evolves in a Markovian fashion, and the reward drawn at each time is a function of the state of the selected arm. Two kinds of model for the underlying Markov decision process (MDP) have been studied. The *rested* bandit model, introduced by Gittins (1979), assumes that the state of an arm is frozen at a time step unless it is pulled. Whittle (1988) introduced the *restless* bandit model, where the state of every arm can change at each step, regardless of which arm is pulled.

There are two lines of research which explore the Markovian MAB problem. The traditional approach has been to assume that the statistical parameters associated with the Markov chain of the arms are perfectly known to the decision maker. The aim is to optimize the *infinite horizon* discounted or average reward of the corresponding MDP. Papers that study these problems typically propose *index policies* fashioned after Gittin's index (Gittins (1979)) and Whittle's index (Whittle (1988)) that are computationally more efficient than solving Bellman equations; see Mahajan and Teneketzis (2008), Gittins et al. (2011) for a broad survey. These achieve approximately optimal infinite horizon cost.

A more recent line of research investigates the *reinforcement learning* problem, where the statistical parameters of the MDP are not known a priori. This assumption adds to the optimization problem the challenge of learning the transition structure, and thus presents an explore-exploit dilemma as in the stochastic MAB problem. For the restless bandit model with finite state space and action space, Jaksch et al. (2010), Ortner et al. (2014) prove finite time regret bounds of $O(\sqrt{t})$ where the scaling constant depends on the size of the state space and action space and the structure of the MDP. They also show that the regret bound scales as $O(\log t)$ with a scaling constant that depends on the gap between the average costs of the best and second best policy.

Although the queueing bandit problem studied in this paper has an underlying Markovian structure, the challenge is essentially different from either rested or restless bandits. Unlike the rested or restless bandits where the state of each arm evolves according to an independent Markov chain, the Markovian structure in this problem is captured by the state of the system (here queue-lengths) which is a complex function of the external arrival process, the decision rule and the rewards of the selected arm in the previous time slots. Nevertheless, it presents the same kind of exploration-exploitation dilemma as the above MAB problems when the statistical parameters are not known a priori.

Bandits for queues: There is a growing body of literature on the application of bandit models to queueing and scheduling problems – see Niño-Mora (2007), Mahajan and Teneketzis (2008), Gittins et al. (2011), Larrañaga et al. (2016). Specifically, Cox and Smith (1961), Buyukkoc et al. (1985), Van Mieghem (1995), Niño-Mora (2006), Jacko (2010), Lott and Teneketzis (2000), Ayesta et al. (2017) use bandit models to develop algorithms for various types of scheduling problems in queueing networks. The primary difference between these models and ours is that they assume a priori knowledge of the statistical parameters, while our focus is on learning these parameters. In these papers, the goal is to solve a stochastic scheduling MDP, and the problem is embedded in the bandit framework as a solution approach. Their aim is to optimize infinite-horizon costs (i.e., statistically steady-state behavior, where the focus typically is on conditions for optimality of index policies), whereas we focus on the analysis

of finite time regret. Further, the models do not typically consider user-dependent server statistics.

Several other problems in queueing systems such as routing and admission control have been studied in the bandit framework (Niño-Mora (2012), Avrachenkov et al. (2013)), again assuming known statistics. Here too, the focus is on designing index policies for these models and showing approximate optimality in steady state.

Switch scheduling: Scheduling in switch networks has received a great deal of attention in the last two decades; see, e.g., Srikant and Ying (2014) for a survey. Notably, many of the scheduling algorithms (e.g., queue-length-based backpressure scheduling algorithms) can yield near-optimal performance in the absence of information about arrival rates; but these algorithms still require information about server availability and capacity. By contrast, our work focuses on learning about all unknown aspects of the environment, as needed to deliver small overall queue lengths. Finally, the problem of identifying the right matchings in a bipartite graph has been formulated as a special case of the combinatorial/linear bandit problem (Gai et al. (2012), Cesa-Bianchi and Lugosi (2012), Combes et al. (2015b), Degenne and Perchet (2016)), but with a generic reward structure and not in the context of queue scheduling as in our case.

3   Problem Setting

In this section, we first introduce a baseline model of a discrete-time network with a single queue and K servers. We discuss our main results in the context of this model, but ultimately formally state and prove our main theorems in the context of a somewhat more general setting with multiple queues and multiple servers. By focusing on the single queue model in our presentation, we are able to better elucidate the main characterizations of regret, including the distinctions between the early and late stages.

3.1   A Single Queue Model

Formally, suppose that arrivals to the queue and service offered by the servers are according to a product Bernoulli distribution, i.i.d. across time slots, with arrival probability λ and service probabilities given by the vector $\boldsymbol{\mu} = [\mu_k]_{k \in [K]}$. Let the highest service probability among all the servers be denoted by μ^* and the lowest by μ_{min}. We assume that the system can be stabilized if the best server is known a priori, i.e., $\lambda < \mu^*$.

The scheduling decision at any time t is based on past observations corresponding to the services obtained for the scheduled servers until time $t - 1$. Statistical parameters corresponding to the service distributions are considered unknown.

The queue evolution for the single queue system can be described as follows. Let $\kappa(t)$ be the server scheduled at time t. Also, let $R_k(t)$ be the service offered to the queue by server k and $S(t)$ denote the service offered by server $\kappa(t)$ at time t. If $A(t)$ denotes the (binary) arrival at time t, then the queue-length at time t is given by:

$$Q(t) = (Q(t - 1) + A(t) - S(t))^+ .$$

We analyze the performance of a scheduling algorithm with respect to queue-regret as a function of time and key system parameters, particularly:

(a) the load on the system $\epsilon := \mu^* - \lambda$, and

(b) the difference between the service probabilities of the best and the next best servers $\Delta := \mu^* - \max_{k \neq k^*} \mu_k$.

3.2 The General Model: A Multi-Queue Switch Network

As mentioned, we prove the results in Sections 4 and 5 for a more general problem that deals with a discrete-time stochastic switch network described as follows. The multi-queue switch network consists of U queues and K servers, where $U \leq K$. The queues and servers are indexed by $u = 1, \ldots, U$ and $k = 1, \ldots, K$ respectively. Arrivals to queues and service offered by the links are according to product Bernoulli distribution and i.i.d. across time slots. The arrival probabilities are given by the vector $\boldsymbol{\lambda} = (\lambda_u)_{u \in [U]}$ and the service probabilities by the matrix $\boldsymbol{\mu} = [\mu_{uk}]_{u \in [U], k \in [K]}$.

We require the following notational definitions in our technical development:

$$\mu_u^* := \max_{k \in [K]} \mu_{uk}, \quad u \in [U];$$

$$k_u^* := \arg\max_{k \in [K]} \mu_{uk}, \quad u \in [U];$$

$$\epsilon_u := \mu_u^* - \lambda_u, \quad u \in [U];$$

$$\Delta_{uk} := \mu_u^* - \mu_{uk}, \quad u \in [U], k \in [K];$$

$$\Delta := \min_{u \in [U], k \notin k_u^*} \Delta_{uk};$$

$$\mu_{min} := \min_{u \in [U], k \in [K]} \mu_{uk};$$

$$\mu^* := \max_{u \in [U], k \in [K]} \mu_{uk};$$

$$\lambda_{min} := \min_{u \in [U]} \lambda_u.$$

In any time slot, each server can serve at most one queue and each queue can be served by at most one server. The task is to schedule, in every time slot, a matching in the complete bipartite graph between queues and servers. Let $\kappa_u(t)$ denote the server that is assigned to queue u at time t. Therefore, the vector $\boldsymbol{\kappa}(t) = (\kappa_u(t)_{u \in [U]})$ gives the matching scheduled at time t. Other important notation for the multi-queue setting can be found in Table 1.

Notation: Boldface letters are used to denote vectors or matrices and the corresponding non-bold letters to denote their individual components. Also, the notation

(i) $\mathbb{1}\{\cdot\}$ is used to denote the indicator function, and

(ii) for any $k \in \mathbb{N}$, $\log^k(\cdot)$ is used to denote $(\log(\cdot))^k$.

Unique Optimal Matching

We focus on a simple special case of the above switch scheduling system. In particular, we assume for every queue, there is a unique optimal server with the maximum service probability for that queue. Further, we assume that the optimal queue-server pairs form a matching in the complete bipartite graph between queues and servers, that we call the *optimal matching*; and that this optimal matching stabilizes every queue.

Formally, we make the following assumption on the switch scheduling system.

Assumption 1 (Unique Optimal Matching). *There is a unique optimal matching, i.e.:*

1. *There is a unique optimal server for each queue: k_u^* is a singleton, i.e., $\Delta_{uk} > 0$ for $k \neq k_u^*$, for all u,*

2. *The optimal queue-server pairs for a matching: For any $u' \neq u$, $k_u^* \neq k_{u'}^*$.*

The assumption of a unique optimal matching essentially means that the queues and servers are solving a pure coordination problem. It is most applicable in settings where there is strong horizontal differentiation across jobs and servers. For example, in a crowdsourcing system where the "servers" are worker types completing different types of jobs, the unique optimal matching assumption implies there is a unique worker type best suited to each type of job.

We evaluate the performance of scheduling policies against the policy that schedules the optimal matching in every time slot. Let $\mathbf{Q}(t)$ be the queue-length vector at time t under our specified algorithm, and let $\mathbf{Q}^*(t)$ be the corresponding vector under the optimal policy. We define *regret* as the difference in mean queue-lengths for the two policies. That is, the regret (vector) is given by: $\mathbf{\Psi}(t) := \mathbb{E}[\mathbf{Q}(t) - \mathbf{Q}^*(t)]$.

Throughout, when we evaluate queue-regret, we do so under the assumption that the system is "stable" under the optimal policy, i.e., the policy that chooses the optimal matching at each time step. With Bernoulli distributions for the arrival and service processes and under Assumption 1, most commonly used notions of stability Neely (2010), including the existence of a steady-state distribution, are equivalent to the following condition:

Assumption 2 (Stability). *$\epsilon_u > 0$ for all $u \in [U]$.*

In fact, the above condition guarantees more than just the existence of a steady-state distribution for the queue-length process $\{\mathbf{Q}^*(t)\}$; in fact, the condition ensures that the process is geometrically ergodic.

We also assume that the queueing system starts in the steady state distribution of the system induced by the optimal policy.

Assumption 3 (Initial State). *The queueing system starts with an initial state $\mathbf{Q}(0)$ distributed according to the stationary distribution of $\mathbf{Q}^*(t)$, which we denote $\pi_{(\boldsymbol{\lambda}, \boldsymbol{\mu}^*)}$.*

This assumption is made largely to ease the technical exposition. Throughout the paper, we discuss how the assumption may be weakened for each of our results.

4 The Late Stage

As a preview of the theoretical results, Figure 1 shows the evolution of queue-regret with time in a single queue system with five servers under a scheduling policy inspired by UCB. (Further discussion of the scheduling algorithm used can be found in Section 4.2.) It is observed that the regret goes through a phase transition. In the initial stage, when the algorithm has not estimated the service probabilities well enough to stabilize the queue, the regret grows poly-logarithmically as in the classical MAB setting. After a critical point when the algorithm has learned the system parameters well enough to stabilize the queue, the queue-length goes through regenerative cycles as the queue becomes empty. Thus at the beginning of every regenerative cycle, there is no accumulation of past errors and the sample-path queue-regret is at most zero. As the algorithm estimates the parameters better with time, the length of the regenerative cycles decreases and the queue-regret decays to zero.

An interesting question to ask is: how does the queue-regret scale as $t \to \infty$? For the basic MAB problem, it is well-known that regret, which is the cumulative sum of the rate loss in each

Table 1: General Notation.

Symbol	Description
λ_u	Probability of arrival to queue u
λ_{min}	Minimum arrival probability across all queues
$A_u(t)$	Arrival at time t to queue u
μ_{uk}	Service probability of server k for queue u
$R_{uk}(t)$	Service offered by server k to queue u at time t
k_u^*	Best server for queue u
μ_u^*	Best service probability for queue u
μ^*	Maximum service probability across all links
μ_{min}	Minimum service probability across all links
Δ	Minimum (among all queues) difference between the best and second best servers
$\kappa_u(t)$	server assigned to queue u at time t
$S_u(t)$	Potential service provided by server assigned to queue u at time t
$Q_u(t)$	queue-length of queue u at time t
$Q_u^*(t)$	queue-length of queue u at time t for the optimal strategy
$\Psi_u(t)$	Regret for queue u at time t
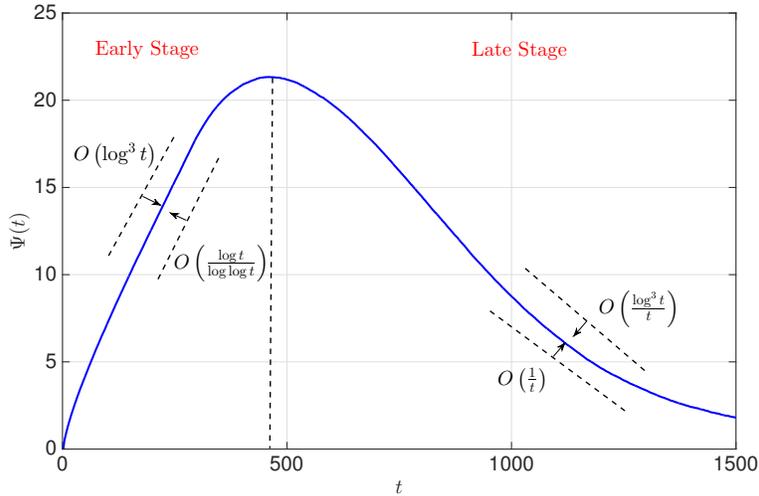
Figure 1: Variation of queue-regret $\Psi(t)$ for a particular user under Q-UCB in a 1×5 system with $\epsilon = 0.15$ and $\Delta = 0.17$

time-slot, scales as $O(\log t)$. However, queue-regret is the mean difference between the cumulative sum of departures. We show in the following lemma that, for any scheduling policy, queue-regret is upper bounded by the cumulative sum of rate loss, i.e., regret in the basic MAB problem. Let $\boldsymbol{S}(t)$ and $\boldsymbol{S}^*(t)$ denote the service offered by the servers scheduled at time t by the proposed algorithm and the genie policy respectively. Note that, for any queue $u \in [U]$, $\mathbb{E}\left[S_u^*(t) - S_u(t)\right]$ is the rate loss at time t. Then, the queue-regret $\boldsymbol{\Psi}(t)$ has the following upper bound:

Lemma 1. $\boldsymbol{\Psi}(t) = \mathbb{E}\left[\boldsymbol{Q}(t) - \boldsymbol{Q}^*(t)\right] \leq \sum_{l=1}^{t} \mathbb{E}\left[\boldsymbol{S}^*(l) - \boldsymbol{S}(l)\right].$

Please see Appendix 8.1 for the proof for this lemma. This result shows that, in the single queue setting, using classical bandit algorithms like UCB1 and Thompson Sampling, one can achieve a queue-regret of at most $O(\log t)$. However, our intuition suggests that the accumulation of errors is only over regenerative cycles. Now observe that the derivative of cumulative regret (which roughly corresponds to regret per time-slot) is $O(1/t)$, and that the regenerative cycle-lengths for the optimal policy are $\Theta(1)$. Thus, it is reasonable to believe that the queue-regret at time t is $O(1/t)$ times a constant factor that increases with the length of the regenerative cycle. To push this intuition through to a formal proof requires high probability results for bandits over finite intervals of time corresponding to queue busy periods, where the intervals are random variables which are themselves coupled to the bandit strategy (the regenerative cycle-lengths are coupled to the bandit scheduling decisions). Analyses of traditional algorithms like UCB1 and Thompson Sampling guarantee high probability results only after sufficient number of sub-optimal arm pulls (Auer et al. (2002)). Audibert et al. (2009) give upper tail bounds for the number of sub-optimal arm pulls for the UCB1 and UCB-V algorithms but these bounds are polynomial in the factor of deviation from the mean and not polynomial in time.

To the best of our knowledge, there is a lack of high probability upper and lower bounds in the multi-armed bandit literature for arbitrary time intervals. The lack of such results motivates us to use alternate proof strategies for the lower bound on queue-regret and the achievability results. For the lower bound, we use coupling arguments to derive lower bounds on queue regret growth over one time-step for any α-consistent policy. This allows us to show that no α-consistent policy can achieve a better scaling than $O(1/t)$. For the achievability result, we use a forced exploration variant of UCB/Thompson Sampling (a combination of ϵ-greedy and UCB/Thompson Sampling algorithms) that gives us good upper bounds on the expected number of sub-optimal schedules over finite time intervals. When combined with queueing arguments, this leads to an upper bound on queue regret.

As before, in each section, we first discuss all our main results in the context of a model with a single queue. In particular, Section 4.1 discusses our approach to the lower bound while Section 4.2 discusses our upper bound. We then extend the results in the single-queue setting to the more general model described in Section 3.2.

4.1 An Asymptotic Lower Bound

We first establish a lower bound in a single-queue system; the arguments leading to this lower bound also form the backbone of our lower bound for the switch scheduling system.

4.1.1 The Single-Queue System.

We establish an asymptotic lower bound on regret for the class of α-consistent policies. For the single-queue setting, we define this class of policies as in the traditional stochastic MAB problem (Lai and Robbins (1985), Salomon et al. (2013), Combes et al. (2015a)) as follows:

Definition 1. *A scheduling policy for a single queue network is said to be α-consistent (for some $\alpha \in (0,1)$) if for any problem instance $(\lambda, \boldsymbol{\mu})$, there exists a constant $C(\lambda, \boldsymbol{\mu})$ such that*

$$\mathbb{E}\left[\sum_{s=1}^{t} \mathbb{1}\{\kappa(s) = k\}\right] \leq C(\lambda, \boldsymbol{\mu}) t^{\alpha}$$

for all $k \neq k^*$.

This means that any policy under this class schedules a sub-optimal server not more than $O(t^{\alpha})$ (sub-linear in t) number of times. Without a restriction such as that imposed in the preceding definition, "trivial" policies that, e.g., schedule the same server every time step would be allowed. Such policies would have zero expected regret if the chosen server happened to be optimal, and otherwise would have linear regret. The α-consistency requirement rules out such policies, while ensuring the set under consideration is reasonable. Proposition 2 below gives an asymptotic lower bound on the queue regret for an arbitrary α-consistent policy.

Proposition 2. *For any problem instance $(\lambda, \boldsymbol{\mu})$ and any α-consistent policy, the regret $\Psi(t)$ satisfies*

$$\Psi(t) \geq \left(\frac{\lambda}{4} D(\boldsymbol{\mu})(1-\alpha)(K-1)\right)\frac{1}{t}$$

for infinitely many t, where

$$D(\boldsymbol{\mu}) := \frac{\Delta}{\mathrm{KL}\left(\mu_{min}, \frac{\mu^*+1}{2}\right)}. \tag{2}$$

Proof Outline for Proposition 2. The proof of the lower bound consists of three main steps. First, in Lemma 21, we show that the regret at any time-slot is lower bounded by the probability of a sub-optimal schedule in that time-slot (up to a constant factor that is dependent on the problem instance). The key idea in this lemma is to show, for any scheduling algorithm, the equivalence of two systems with the same marginal service distributions. This is achieved through a carefully constructed coupling argument that maps the original system with independent service across links to another system with service process that is dependent across links but with the same marginal distribution.

As a second step, the lower bound on the regret in terms of the probability of a sub-optimal schedule enables us to obtain a lower bound on the cumulative queue-regret in terms of the number of sub-optimal schedules. We then use a lower bound on the number of sub-optimal schedules for α-consistent policies (Lemma 19 and Corollary 20) to obtain a lower bound on the cumulative regret. In the final step, we use the lower bound on the cumulative queue-regret to obtain an *infinitely often* lower bound on the queue-regret. $\qquad\square$

We now generalize this result to the switch network (Theorem 3). We omit the proof details of Proposition 2 and give a detailed proof for Theorem 3 in Appendix 9.

4.1.2 The Multi-Queue Network.

Before stating the generalization of Proposition 2 for the switch network, we first extend the definition of the class of α-consistent policies for a multi-queue network with a unique optimal matching.

Definition 2. *A scheduling policy for a multi-queue network with a unique optimal matching is said to be α-consistent (for some $\alpha \in (0, 1)$) if for any problem instance $(\boldsymbol{\lambda}, \boldsymbol{\mu})$, there exists a constant $C(\boldsymbol{\lambda}, \boldsymbol{\mu})$ such that*

$$\mathbb{E}\left[\sum_{s=1}^{t} \mathbb{1}\{\kappa_u(s) = k\}\right] \leq C(\boldsymbol{\lambda}, \boldsymbol{\mu}) t^{\alpha}$$

for all $u \in [U]$ and $k \neq k_u^*$.

Theorem 3 below gives an asymptotic lower bound on the average queue-regret and per-queue regret for an arbitrary α-consistent policy in the multi-queue setting; it has the same scaling in t as in the single-queue setting, but now explicitly depends on the parameters of the switch problem.

Theorem 3. *For any problem instance $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ with a unique optimal matching, and any α-consistent policy, the regret $\boldsymbol{\Psi}(t)$ satisfies*

(a)

$$\frac{1}{U} \sum_{u \in [U]} \Psi_u(t) \geq \left(\frac{\lambda_{min}}{8} D(\boldsymbol{\mu})(1 - \alpha)(K - 1)\right) \frac{1}{t},$$

(b) and for any $u \in [U]$,

$$\Psi_u(t) \geq \left(\frac{\lambda_{min}}{8} D(\boldsymbol{\mu})(1 - \alpha) \max\{U - 1, 2(K - U)\}\right) \frac{1}{t}$$

for infinitely many t, where D is given by (2).

This result does not require Assumption 3 and holds irrespective of the distribution of the initial queue-length. The full proof of this theorem is given in Section 9.

4.2 Achieving the Asymptotic Bound

The lower bounds in Proposition 2 and Theorem 3 imply that no α-consistent policy can achieve a queue-regret better than $O(1/t)$. We next ask if straight-forward generalizations of standard bandit algorithms like UCB and Thompson sampling can achieve a scaling of $O(1/t)$, thus matching the lower bound in Theorem 3. To prove that these algorithms achieve this scaling, it is essential to show high probability bounds on scheduling errors over regenerative cycles in the late stage. A systematic way to show this would be to prove that the algorithm has a good estimate of all the service probabilities in the late stage leading to the correct scheduling decision. But for standard bandit algorithms, a lack of concentration results on the number of times each link is scheduled makes it difficult to prove a high probability bound on scheduling errors over a finite time-interval in the late stage.

Algorithms with Forced Exploration.

To get around this difficulty, we propose slightly modified versions of the standard bandit algorithms (UCB1 and Thompson Sampling) generalized to the multi-dimensional queueing bandit. These algorithms, which we call Q-UCB and Q-ThS (corresponding to UCB1 and Thompson sampling, respectively) have an explicit forced exploration component similar to ϵ-greedy algorithms. This

forced exploration may not be necessary to ensure good performance in practice (see Section 6 for a detailed discussion) but it helps us in theoretically proving that the algorithm has a sufficiently good estimate of all the service probabilities (including sub-optimal ones) in the late stage.

We first briefly describe the algorithms in the single queue setting. The algorithms are precisely defined in a more general setting with multiple queues in Section 4.2.1. In particular, notation and details of the algorithms are given for the general case in Table 2 and Algorithms 1 and 2 respectively.

In the single queue setting, at time-slot t ($t \geq 1$), both Q-UCB and Q-ThS *explore* with probability $\min\{1, 3K \log^2 t / t\}$, otherwise they *exploit*. To explore, the algorithm chooses a server uniformly at random from the K servers. The chosen exploration rate ensures that we are able to obtain concentration results for the number of times any link is sampled.[2] If it exploits, Q-UCB (resp., Q-ThS) chooses the best server using the UCB1 (resp., Thompson sampling) method. In other words, Q-UCB chooses the server that has the highest upper confidence bound for the corresponding service probability, while Q-ThS chooses the server according to random samples drawn from the posterior distribution of the service probabilities. Note that both UCB1 and Thompson sampling explore as well; the innovation in our algorithms is the *forced* exploration that we have added, to provide high confidence estimation of service probabilities (please see Lemma 11 in Appendix 8 for details).

Asymptotic Guarantees for Q-UCB and Q-ThS.

For a given problem instance $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ (and therefore fixed $\boldsymbol{\epsilon}$), we can show that the regret under Q-UCB and Q-ThS scale as $O\left(\mathrm{poly}(\log t)/t\right)$. We state our asymptotic upper bound in the case of a single queue in Proposition 4; this is a special case of Corollary 6 in Section 4.2.1, which holds for a more general setting with multiple queues.

Proposition 4. *Let $w(t) = t^{(1-1/\beta)}$ for some fixed $\beta > 1$. Then for both Q-UCB and Q-ThS,*

$$\Psi(t) = O\left(K \frac{\log^3 t}{\epsilon^2 t}\right)$$

for all t such that $\frac{w(t)}{\log t} \geq \frac{2}{\epsilon}$, $\frac{t}{w(t)} \geq \max\left\{\frac{24K}{\epsilon}, 15K^2 \log t\right\}$, $t \geq \exp\left(\frac{4}{\Delta^2(1-1/\beta)^3}\right)$ and $\frac{t}{\log t} \geq \frac{198}{\epsilon^2}$.

Proof Outline for Proposition 4. We outline how the proof proceeds here for a single queue, to lay bare the central arguments. The full proof for Theorem 5, which is a generalization of Proposition 4 for the setting of multiple queues, is provided in Section 8.

As mentioned earlier, the central idea in the proof is that the sample-path queue-regret is at most zero at the beginning of regenerative cycles, i.e., instants at which the queue becomes empty. The proof consists of two main parts – one which gives a high probability result on the number of sub-optimal schedules in the exploit phase in the late stage, and the other which shows that at any time, the beginning of the current regenerative cycle is not very far in time.

The former part is proved in Lemma 11, where we make use of the forced exploration component of Q-UCB and Q-ThS to show that all the links, including the sub-optimal ones, are sampled a sufficiently large number of times to give a good estimate of the service probabilities. This in turn ensures that the algorithm schedules the correct links in the exploit phase in the late stages with high probability.

For the latter part, we prove a high probability bound on the last time instant when the queue was zero (which is the beginning of the current regenerative cycle) in Lemma 17. Here, we make use of a recursive argument to obtain a tight bound. More specifically, we first use a coarse

high probability upper bound on the queue-length (Lemma 13) to get a first cut bound on the beginning of the regenerative cycle (Lemma 14). This bound on the regenerative cycle-length is then recursively used to obtain tighter bounds on the queue-length, and in turn, the start of the current regenerative cycle (Lemmas 16 and 17 respectively).

The proof proceeds by combining the two parts above to show that the main contribution to the queue-regret comes from the forced exploration component in the current regenerative cycle, which gives the stated result. □

This result, in combination with Proposition 2, shows that queue-regret for Q-UCB and Q-ThS in the long-term is within a poly$(\log t)$ factor of the optimal queue-regret for the α-consistent class. We will next describe how this result can be extended to the multi-queue switch network. We skip the details of the proof for Proposition 4 and give a detailed proof for its generalized version (Corollary 6) in Appendix 8.

4.2.1   The Multi-Queue Network.

We now describe the scheduling algorithm in the multi-queue setting. Let $\mathcal{M} \subset [K]^U$ be the set of all matchings. We represent a matching by a U-length vector in which the u^{th} element gives the server that is assigned to queue u. A vector $\boldsymbol{\kappa} \in \mathcal{M}$ if and only if for any $u' \neq u$, $\kappa_{u'} \neq \kappa_u$. Let $\mathcal{X} \subset \mathcal{M}$ be a subset of K perfect matchings such that their union covers the set of all edges in the complete bipartite graph (it is easy to show that such a decomposition is possible). Also, let $T_{uk}(t)$ be the number of times server k is assigned to queue u in the first t time-slots and $\hat{\boldsymbol{\mu}}(t)$ be the empirical mean of the service offered by the links at time t from past observations (until $t-1$).

At time-slot t, Q-UCB and Q-ThS both decide to *explore* with probability $\min\{1, 3K \log^2 t/t\}$, otherwise they *exploit*. To explore, the algorithms choose a matching uniformly at random from the set $\mathcal{X}$. If it exploits, Q-UCB (resp., Q-ThS) first assigns a best server for every queue using the UCB1 (resp., Thompson Sampling) method. In other words, for every queue, Q-UCB chooses the server that has the highest upper confidence bound for the corresponding service probability (breaking ties arbitrarily), while Q-ThS chooses the server according to random samples drawn from the posterior distribution of the service probabilities. Q-UCB (resp., Q-ThS) then schedules the projection of this U-length server vector $\hat{\boldsymbol{k}}(t)$ onto the space of all matchings $\mathcal{M}$. Notation and details of the algorithms are given in Table 2 and Algorithms 1 and 2.

Note that the scheduled matching $\boldsymbol{\kappa}(t)$ is the projection of $\hat{\boldsymbol{k}}(t)$ onto the space of all matchings $\mathcal{M}$ with Hamming distance as metric, i.e., a matching that matches the maximum number of queues with their corresponding "best" server. One way to compute this projection is to take the union of two matchings selected as follows:

1. The first is a maximal matching in the sub-graph induced by the best server configuration;

2. The second is a maximal matching in the sub-graph obtained by removing the matching chosen in the first step from the complete bipartite graph.

Let $\tau_1 = 5.8 \times 10^3$, and let τ_2 be a constant such that

$$t \exp\left(-\frac{1}{2} \log^2 t\right) + t^2 \exp\left(-\frac{1}{4} (2 \log t)^{4/3}\right) + t^2 \exp\left(-\frac{1}{4} \log^2 t\right) \leq \frac{1}{6t^3} \qquad (3)$$

for all $t \geq \tau_2$. Such a τ_2 exists since each term on the left-hand side of (3) is $o(1/t^3)$.

The following theorem gives an upper bound on the regret for each individual queue for both Q-UCB (Algorithm 1) and Q-ThS (Algorithm 2) in the multi-queue setting.

14

Table 2: Notation for Algorithms 1, 2

Symbol	Description
$\mathsf{E}(t)$	Indicates if the algorithm schedules a matching through *Explore*
$\mathsf{E}_{uk}(t)$	Indicates if Server k is assigned to Queue u at time t through *Explore*
$\mathsf{I}_{uk}(t)$	Indicates if Server k is assigned to Queue u at time t through *Exploit*
$T_{uk}(t)$	Number of time slots Server k is assigned to Queue u in time $[1, t-1]$
$\hat{\boldsymbol{\mu}}(t)$	Empirical mean of service at time t from past observations (until $t-1$)
$\boldsymbol{\kappa}(t)$	Matching scheduled in time-slot t

Algorithm 1 Q-UCB

At time $t \geq 1$,
Let $\mathsf{E}(t)$ be an independent Bernoulli sample of mean $\min\{1, 3K\frac{\log^2 t}{t}\}$.
if $\mathsf{E}(t) = 1$ **then**

    Explore:
    Schedule a matching from $\mathcal{X}$ uniformly at random.

else

    Exploit:
    Compute for all $u \in [U]$

$$\hat{k}_u(t) := \arg\max_{k \in [K]} \hat{\mu}_{uk}(t) + \sqrt{\frac{\log^2 t}{2T_{uk}(t-1)}}.$$

    Schedule a matching $\boldsymbol{\kappa}(t)$ such that

$$\boldsymbol{\kappa}(t) \in \arg\min_{\boldsymbol{\kappa} \in \mathcal{M}} \sum_{u \in [U]} \mathbb{1}\left\{\kappa_u \neq \hat{k}_u(t)\right\},$$

end if

15

Algorithm 2 Q-ThS

At time $t \geq 1$,

Let $\mathsf{E}(t)$ be an independent Bernoulli sample of mean $\min\{1, 3K\frac{\log^2 t}{t}\}$.

if $\mathsf{E}(t) = 1$ **then**

    Explore:

    Schedule a matching from $\mathcal{X}$ uniformly at random.

else

    Exploit:

    For each $k \in [K], u \in [U]$, pick a sample $\hat{\theta}_{uk}(t)$ of distribution,

$$\hat{\theta}_{uk}(t) \sim \text{Beta}\left(\hat{\mu}_{uk}(t)T_{uk}(t-1) + 1, (1 - \hat{\mu}_{uk}(t))\,T_{uk}(t-1) + 1\right).$$

    Compute for all $u \in [U]$

$$\hat{k}_u(t) := \arg\max_{k \in [K]} \hat{\theta}_{uk}(t)$$

    Schedule a matching $\boldsymbol{\kappa}(t)$ such that

$$\boldsymbol{\kappa}(t) \in \arg\min_{\boldsymbol{\kappa} \in \mathcal{M}} \sum_{u \in [U]} \mathbb{1}\left\{\kappa_u \neq \hat{k}_u(t)\right\},$$

end if

Theorem 5. *Consider any problem instance $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ which has a unique optimal matching. For any $u \in [U]$, let $w(t) = t^{(1-1/\beta)}$ for some fixed $\beta > 1$, $v'_u(t) = \frac{6K}{\epsilon_u}w(t)$ and $v_u(t) = \frac{24}{\epsilon_u^2}\log t + \frac{60K}{\epsilon_u}\frac{v'_u(t)\log^2 t}{t}$. Then, for Algorithm 1 (resp., Algorithm 2), the regret for queue u, $\Psi_u(t)$, satisfies*

$$\Psi_u(t) \leq 6K\frac{v_u(t)\log^2 t}{t} + \frac{24.004 + UK}{6t^2}$$

for all $t \geq \tau_1$ (resp., $t \geq \tau_2$) such that $t \geq \exp\left(\frac{4}{\Delta^2(1-1/\beta)^3}\right)$, $\frac{w(t)}{\log t} \geq \frac{2}{\epsilon_u}$ and $v_u(t) + v'_u(t) \leq t/2$.

Remark 1. *In our analysis, the constants τ_1 and τ_2 are determined by the upper bounds obtained in Lemmas 10 and 11. For the upper bound on the number of explore time-slots in Lemma 10, we show that is sufficient to have $t \geq 5.8 \times 10^3$ for both Q-UCB and Q-ThS. For the upper bound on the number of sub-optimal schedules in the exploit phase, our analysis in Lemma 11 shows that it is sufficient to take $\tau_1 = 5.8 \times 10^3$ for Q-UCB. For Q-ThS, we use the 'Beta-Binomial trick' from Kaufmann et al. (2012), Agrawal and Goyal (2012) to get an upper bound which is qualitatively similar to that for Q-UCB (please see the bounds in (13) and (16) for comparison). However, the constants in the exponent for Q-ThS are smaller than that of Q-UCB by a factor of 4 making τ_2 a much larger constant than τ_1. While the upper bounds in this paper are suggestive of the evolution of regret with time, the obtained constants do not accurately characterize the empirical performance of the proposed algorithms. In particular, the inferior constants for Q-ThS as compared to Q-UCB could just be an artefact of our analysis. As seen in Figure 3, simulations show that Q-ThS performs better than Q-UCB both in the early and late stages.*

Remark 2. *For any queue u, we state the regret bounds and the corresponding time-intervals in which these bounds hold as a function of ϵ_u, the gap between the arrival probability and the best*

service probability for that queue. Therefore, the time ranges for which the bounds hold may vary for different queues depending on ϵ.

Remark 3. *Although we assume Bernoulli distributions for arrival and service in our model, the result in Theorem 5 holds for general, non-Bernoulli distributions with bounded support if*

(i) *there is a unique matching that gives the best service rate for all users (similar to Assumption 1 in Section 3), and*

(ii) *the genie policy that defines the regret $\boldsymbol{\Psi}(t)$ is the one that always schedules the best matching.*

Remark 4. *The result in Theorem 5 can be proved without making any assumption on the distribution of $\mathbf{Q}^*(0)$ as in Assumption 3. Specifically, this assumption is used in Lemma 13 to show a first cut upper bound on the queue-length. Using the fact that the process $\{\mathbf{Q}^*(t)\}$ is geometrically ergodic, the same lemma can be extended to the case of a generic $\mathbf{Q}^*(0)$ but the convergence time would now be a function of $\mathbf{Q}^*(0)$.*

A slightly weaker version of Theorem 5 is given in Corollary 6. This corollary is useful to understand the dependence of the upper bound on the load ϵ and the number of servers K. (Proposition 4 is a special case of this corollary.)

Corollary 6. *For any $\beta > 1$,*

$$\Psi_u(t) \leq \frac{289K \log^3 t}{\epsilon_u^2 t}$$

for all $t \geq \tau_0$ such that

$$\log t \geq \max\left\{ \frac{4}{\Delta^2(1 - 1/\beta)^3}, \sqrt{2}\left(\log \frac{2}{\epsilon_u}\right)^{1.5}, \beta\frac{24K}{\epsilon_u}, \beta\left(\log\log t + \log(15K^2)\right), \frac{\beta}{\beta - 1}\log\left(\frac{13.2}{K^2\epsilon_u^2}\right) \right\}.$$

Full proofs of Theorem 5 and Corollary 6 are given in Section 8.

5 The Early Stage in the Heavily Loaded Regime

Proposition 4, Theorem 5 show that systematic exploration in Q-UCB and Q-ThS ensures an $O\left(\text{poly}(\log t)/t\right)$ queue-regret in the late stage. The penalty for aggressive exploration is likely to be more apparent in the initial stages when the queues have not yet stabilized and there are few regenerative cycles. As a result, the queueing system has a behavior similar to the traditional MAB system in the early stage. Thus, it is reasonable to expect that algorithms that achieve optimal performance for the traditional MAB problem also perform well in the early stages in the queueing system.

In order to study the performance of α-consistent policies in the early stage, we again focus on a single queue, and consider the *heavily loaded* system, where the arrival probability λ is close to the optimal service probability μ^*. Specifically, we characterize the behavior of queue-regret as the difference between the two probabilities, $\epsilon = \mu^* - \lambda \to 0$. As in the discussion of the late stage, we present all our results first within this single queue model to highlight the main insights, but these results are in fact special cases of results that are proven in the context of a general switch network setting.

Analyzing regret in the early stage in the heavily loaded regime has the effect that the optimal server is the only one that stabilizes the queue. As a result, in the heavily loaded regime, effective learning and scheduling of the optimal server play a crucial role in determining the transition point

from the early stage to the late stage. For this reason the heavily loaded regime reveals the behavior of regret in the early stage.

Proposition 7 gives a lower bound on the regret in the heavily loaded regime, roughly in the time interval $\left(K^{1/1-\alpha}, O\left(K/\epsilon\right)\right)$ for any α-consistent policy. Recall the definition of $D(\boldsymbol{\mu})$ given by Equation (2).

Proposition 7. *Given any single queue system $(\lambda, \boldsymbol{\mu})$, and for any α-consistent policy and $\gamma > \frac{1}{1-\alpha}$, there exist constants τ and C_1 (independent of $(\lambda, \boldsymbol{\mu})$) such that, for $\eta := \frac{(K-1)D(\boldsymbol{\mu})}{2\max\{C_1 K^\gamma, \tau\}}$, if $\epsilon < \eta$, then the regret $\Psi(t)$ satisfies*

$$\Psi(t) \geq \frac{D(\boldsymbol{\mu})}{2}(K-1)\frac{\log t}{\log\log t}$$

for $t \in \left[\max\{C_1 K^\gamma, \tau\}, (K-1)\frac{D(\boldsymbol{\mu})}{2\epsilon}\right]$.

Proof Outline for Proposition 7. The crucial idea in the proof is to show a lower bound on the queue-regret in terms of the number of sub-optimal schedules (Lemma 22). As in Theorem 2, we then use a lower bound on the number of sub-optimal schedules for α-consistent policies (given by Corollary 20) to obtain a lower bound on the queue-regret. □

Proposition 7 is a special case of Theorem 8 (given below) in the multi-queue setting. It gives lower bounds on the average queue regret and individual queue regret in the early stage. Assumption 3 is useful in this result to ensure that the initial queue-length is not far-off from the optimal (see Lemma 22). This assumption can be weakened to the following requirement: "the expected initial queue-length is bounded above by the expected stationary queue-length of $\mathbf{Q}^*(t)$". The proof of this theorem is given in Section 9. Recall that $\epsilon_u = \mu_u^* - \lambda_u$ for $u \in [U]$.

Theorem 8. *Given any problem instance $(\boldsymbol{\lambda}, \boldsymbol{\mu})$, and for any α-consistent policy and $\gamma > \frac{1}{1-\alpha}$, there exist constants τ and C_1 (independent of $(\lambda, \boldsymbol{\mu})$) such that, for $\eta := \frac{(K-1)D(\boldsymbol{\mu})}{2\max\{C_1 K^\gamma, \tau\}}$, the following holds:*

(a) if $\bar{\epsilon} := \frac{1}{U}\sum_{u\in[U]} \epsilon_u < \frac{\eta}{2}$, then the average regret satisfies

$$\frac{1}{U}\sum_{u\in[U]}\Psi_u(t) \geq \frac{D(\boldsymbol{\mu})}{4}(K-1)\frac{\log t}{\log\log t},$$

for $t \in \left[\max\{C_1 K^\gamma, \tau\}, (K-1)\frac{D(\boldsymbol{\mu})}{4\bar{\epsilon}}\right]$, and

(b) for any $u \in [U]$, if $\epsilon_u < \eta$, then the regret for queue u satisfies

$$\Psi_u(t) \geq \frac{D(\boldsymbol{\mu})}{4}\max\left\{U-1, 2(K-U)\right\}\frac{\log t}{\log\log t}$$

for $t \in \left[\max\{C_1 K^\gamma, \tau\}, (K-1)\frac{D(\boldsymbol{\mu})}{2\epsilon_u}\right]$.

Recall that the lower bound in Proposition 7 and Theorem 8 aim to characterize the behavior of regret in the heavily loaded regime, i.e., as $\epsilon := \mu^* - \lambda \to 0$. While the left end point of the interval in which the bound holds depends on the policy, the right end point of the interval is policy independent. Further, for a fixed policy, the left endpoint does not change as $\epsilon \to 0$ (i.e., it is independent of ϵ), while the right endpoint grows without bound as $\epsilon \to 0$. Therefore, for

any α-consistent policy, there is an ϵ small enough such that, for a non-empty time-interval (until $(K-1)\frac{D(\boldsymbol{\mu})}{2\epsilon}$), the regret grows at least as $\Omega\left(K\frac{\log t}{\log\log t}\right)$.[3] The main insight from this result is that it takes at least $\Omega\left(\frac{K}{\epsilon}\right)$ time for the queue-regret with any policy to transition from the early stage to the late stage.

Theorem 8 shows that, for any α-consistent policy, it takes at least $\Omega\left(K/\epsilon\right)$ time for the queue-regret to transition from the early stage to the late stage.

In this region, regret is growing with time, and the scaling $\Omega(\log t/\log\log t)$ reflects the fact that in this regime queue-regret is dominated by the fact that cumulative regret grows like $\Omega(\log t)$. A reasonable question then arises: after time $\Omega\left(K/\epsilon\right)$, should we expect the regret to transition into the late stage regime analyzed in the preceding section?

We answer this question by studying when Q-UCB and Q-ThS achieve their late-stage regret scaling of $O\left(\text{poly}(\log t)/\epsilon^2 t\right)$. Formally, we have Corollary 9, which helps in understanding the scaling of the upper bound in Theorem 5 with respect to the parameters K and ϵ.

Corollary 9. *For any problem instance* $(\boldsymbol{\lambda}, \boldsymbol{\mu})$*, any queue* $u \in [U]$*, any* $\beta < 2$*,* $\gamma > \frac{\beta}{\beta-1}$*,* $\delta > \beta$*, there exists a constant* C_2 *independent of* K *and* ϵ_u *(but depending on* Δ*,* β*,* γ *and* δ*) such that the regret for queue* u *satisfies*

$$\Psi_u(t) \leq \frac{289 K \log^3 t}{\epsilon_u^2 t}$$

$\forall t \geq C_2 \max\left\{\left(\frac{1}{\epsilon_u}\right)^\gamma, \left(\frac{K}{\epsilon_u}\right)^\beta, K^{2\delta}\right\}.$

As an illustrative example, we consider the case where $K = 1/\epsilon_u$ to compare the scaling of the upper bound with the lower bound in Theorem 8. In this case, for any $\delta > 1.5$, there is a constant $C(\Delta, \delta)$ such that the time taken to achieve $O\left(K\text{poly}(\log t)/\epsilon^2 t\right)$ regret scaling is $C(\Delta, \delta)\left(K/\epsilon\right)^\delta$. On the other hand, under any α-consistent policy, the regret scales as $\Omega\left(K\log t/\log\log t\right)$ for $t < K/\epsilon$.

We conclude by noting that Q-UCB and Q-ThS do not yield optimal regret performance in the early stage in general. In particular, recall that at any time t, the forced exploration component in Q-UCB and Q-ThS is invoked with probability $3K\log^2 t/t$. As a result, we see that, in the early stage, queue-regret under Q-UCB and Q-ThS could be a $\log^2 t$-factor worse than the $\Omega\left(\log t/\log\log t\right)$ lower bound shown in Theorem 8 for the α-consistent class. This intuition can be formalized: it is straightforward to show an upper bound of $2K\log^3 t$ for any $t > \max\{C_3, U\}$, where C_3 is a constant that depends on Δ but is independent of K and ϵ; we omit the details.

6 Simulation Results

In this section, we present simulation results to empirically evaluate the performance of the algorithms presented in this paper. These results corroborate our theoretical analysis in Sections 4 and 5. In particular a phase transition from unstable to stable behavior can be observed in all our simulations, as predicted by our analysis. Below, we demonstrate the performance of Q-ThS (Algorithm 2) under variations of system parameters like the traffic (ϵ), the gap between the optimal and the suboptimal servers (Δ), and the size of the system (U and K). We also compare the performance of our algorithms with the traditional bandit algorithms – UCB1 (Auer et al. (2002)) and Thompson Sampling (Thompson (1933)), which do not incorporate forced exploration.
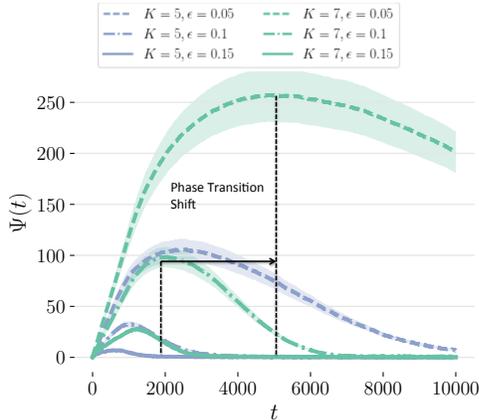
(a) Queue-Regret under Q-ThS for systems with 1 queue and 5, 7 servers with $\epsilon \in \{0.05, 0.1, 0.15\}$
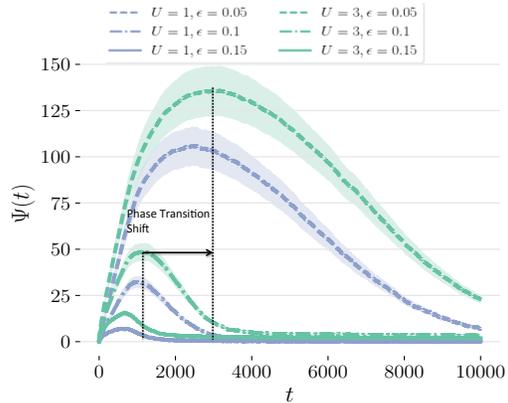
(b) Queue-Regret under Q-ThS for systems with 1, 3 queues and 5 servers with $\epsilon \in \{0.05, 0.1, 0.15\}$

Figure 2: Variation of Queue-regret $\Psi(t)$ with K, U and ϵ under Q-Ths. The phase-transition point shifts towards the right as ϵ decreases. The efficiency of learning decreases with increase in the size of the system. We plot the median statistics of the average queue regret over 1000 simulations and the shaded regions indicate the area between the first and third quartiles.

Variation with U, K and ϵ.

Figure 2a shows the evolution of queue-regret in single queue systems with 5 and 7 servers for different values of load ϵ. It can be observed that the regret decays faster in the smaller system as expected from the theoretical upper and lower bounds, which are proportional to the number of servers K. In Figure 2b, we show the evolution of queue regret in systems with 5 servers and $1, 3$ queues for different load values. For the system with 3 queues, we take the same load for all the queues ($\epsilon_u = \epsilon$ for $u = 1, 2, 3$) and plot the quantity $\Psi(t) = \max_u \Psi_u(t)$. We observe that the queue regret decays at a slower rate in the system with 3 queues as predicted by the theoretical analysis in Theorems 3 and 8.

It is also evident from Figure 2 that the regret of the queueing system grows with decreasing ϵ. We can observe that the time at which the phase transition occurs shifts towards the right with decreasing ϵ which is predicted by Corollaries 6 and 9.

Comparison with Implicit Exploration Algorithms.

In Figure 3, we benchmark the performance of our algorithms against traditional bandit algorithms UCB1 and Thompson sampling. It can be observed that, in the early stage, the traditional algorithms perform better than the proposed ones. This can be explained by the additional forced exploration required by Q-UCB and Q-ThS. In the late stage, we observe that Q-UCB gives slightly better performance than UCB1, however Thompson sampling has better performance in both the stages. In this numerical example, the results suggest that Thompson sampling is already achieving enough exploration to deliver adequate estimation of service probabilities to stabilize the queues. We introduced forced exploration as a vehicle to theoretically understand performance limits of queue regret, which brings some loss of performance in the early stage. Our simulation results thus pose an interesting open direction: to try to quantify the queue regret performance of classical bandit algorithms that do not explicitly include forced exploration.
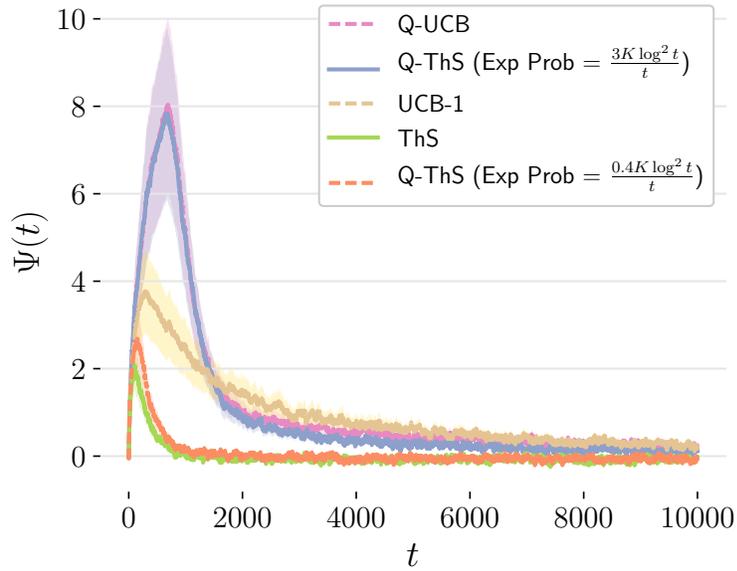
Figure 3: Comparison of queue-regret performance of Q-ThS, Q-UCB, UCB1, and Thompson sampling in a 5 server system with $\epsilon_u = 0.15$ and $\Delta = 0.17$. Two variants of Q-ThS are presented, with different exploration probabilities; note that $3K \log^2 t/t$ is the exploration probability suggested by theoretical analysis (which is necessarily conservative). Tuning the constant significantly improves performance of Q-ThS bringing it closer to Thompson sampling. We plot the median statistics of the average queue regret over 3000 simulations and the shaded regions in the same color indicate the area between the first and third quartiles.

7 Discussion and Conclusion

This paper provides the first regret analysis of the queueing bandit problem, including a characterization of regret in both early and late stages; and algorithms (Q-UCB, Q-ThS) that are asymptotically optimal (to within poly-logarithmic factors). Here we also highlight several additional substantial open directions for future work.

First, note that in the late stage we have a lower bound that is essentially of order $1/t$ (cf. Proposition 2), while the upper bound for our algorithm is of order $\log^3 t/t$ (cf. Proposition 4 and Theorem 5). It remains an open question whether one can close the gap between lower and upper bounds; in particular, this would require development of an algorithm that is able to achieve concentration of the distribution of busy period length with less exploration than the algorithms we have developed.

Second, and related to the previous point, is there a single algorithm that gives optimal performance in *both* early and late stages, as well as the optimal switching time between early and late stages? The price paid for forced exploration by Q-UCB/Q-ThS is an inflation of regret in the early stage. An important open question is to find a single, adaptive algorithm that gives good performance over all time. As we note in the Section 6, classic (no forced exploration) Thompson sampling is an intriguing candidate from this perspective.

Third, the most significant technical hurdle in finding a single optimal algorithm is the difficulty of establishing concentration results for the number of suboptimal arm pulls within a regenerative cycle whose length is dependent on the bandit strategy. Such concentration results would be needed in two different limits: first, as the start time of the regenerative cycle approaches infinity (for the asymptotic analysis of late stage regret); and second, as the load of the system increases (for the analysis of early stage regret in the heavily loaded regime). Any progress on the open directions described above would likely require substantial progress on these technical questions as well.

Finally, we conclude by noting that analysis of this problem in an adversarial/non-stochastic setting presents substantial challenges distinct from our analysis. At its core, the challenge is that even the formulation of an appropriate adversarial/non-stochastic setting is nontrivial. Suppose in particular that service rates are adversarial/non-stochastic, but the arrivals are a stationary stochastic process. Immediately within such a context, we are confronted with difficulty in defining both an appropriate notion of stability, as well as an appropriate notion of a genie policy. In the standard bandit problem in this setting, the goal of the scheduler is to match a genie which chooses the arm with the maximum cumulative (long term) reward; but this may not be appropriate in the queue setting. More precisely, two different regimes become apparent. In particular, if the long run average service rate of every single server is insufficient to serve the long-run average arrival rate, then the queue blows up, and optimal queue-regret (against the best arm in hindsight) is identical to the usual multi-armed bandit as $O(\sqrt{t})$. On the other hand, if all the servers give sufficient service to empty the queue regularly, then decisions made in the "recent" past will dominate queue-regret, and therefore we will want to compare against a genie that adaptively chooses the best server locally in time depending on the service capacity realized during the current busy period.

8 Proof of Regret Upper Bound

8.1 Proof of Lemma 1

Alternate Coupled Service Process.

We first construct an alternate service process such that, under any scheduling policy, the queue evolution for this system has the same distribution as that for the original system. We will use this

construction of the service process to prove *(i)* the upper bound in Lemma 1 and *(ii)* the lower bound in Theorem 3. The construction allows coupling of service offered to each queue across different servers. Specifically, the service process satisfies the following condition: *for each queue, the service offered by different servers at any time-slot could possibly be dependent on each other but has the same marginal distribution as that in the original system and is independent of the service offered to other queues.*

More precisely, the service process is constructed as follows: for each $u \in [U]$, let $\{U_u(t)\}_{t \geq 1}$ be i.i.d. random variables distributed uniformly in $(0, 1)$. The service process for queue u and server k be given by $R_{uk}(t) = \mathbb{1}\{U_u(t) \leq \mu_{uk}\}$ $\forall t$. Since $\mathbb{E}[R_{uk}(t)] = \mu_{uk}$, the marginals of the service offered by each of the servers is the same as the original system, and also, service offered to different queues are independent of each other.

Note that the evolution of the queues is a function of the process $(\mathbf{Z}(l))_{l \geq 1} := (\mathbf{A}(l), \boldsymbol{\kappa}(l), \mathbf{S}(l))_{l \geq 1}$. To prove that the process $\mathbf{Z}$ has the same distribution for both service processes – independent Bernoulli and the coupled process described above – we use induction on the size of the finite-dimensional distribution of $\mathbf{Z}$. In other words, we show that the distribution of the vector $(\mathbf{Z}(l))_{l=1}^{t}$ is the same for the two systems for all t by induction on t.

Suppose that the hypothesis is true for $t - 1$. Now consider the conditional distribution of $\mathbf{Z}(t)$ given $(\mathbf{Z}(l))_{l=1}^{t-1}$. Given $(\mathbf{Z}(l))_{l=1}^{t-1}$, the distribution of $(\mathbf{A}(t), \boldsymbol{\kappa}(t))$ is identical for the two systems for any scheduling policy since the two systems have the same arrival process. Also, given $\big((\mathbf{Z}(l))_{l=1}^{t-1}, \mathbf{A}(t), \boldsymbol{\kappa}(t)\big)$, the distribution of $\mathbf{S}(t)$ depends only on the marginal distribution of the scheduled servers given by $\boldsymbol{\kappa}(t)$ which is again the same for the two systems. Therefore, $(\mathbf{Z}(l))_{l=1}^{t}$ has the same distribution in the two systems. Since the statement is true for $t = 1$, it is true for all t.

Consequently, w.r.t. the process $\mathbf{Z}$, the original system with independent Bernoulli process is stochastically equivalent to the alternate system with the coupled service process. Notably, the queue-regret in the two systems are equal, and so is the traditional regret (given by the cumulative rate loss).

Proof of Lemma 1. We will prove the upper bound using the coupled service process described above. Consider any queue $u \in [U]$. Since $\mu_u^* > \mu_{uk}$ $\forall k \neq k_u^*$, we have $R_{uk_u^*}(t) \geq R_{uk}(t)$ $\forall k \neq k_u^*$, $\forall t$. This implies that $S_u^*(t) \geq S_u(t)$ and $Q_u^*(t) \leq Q_u(t)$ $\forall t$. Using this and the fact that the difference in queue-lengths is equal to the difference in the total number of departures, we have

$$
\begin{aligned}
Q_u(t) - Q_u^*(t) &= \sum_{l=1}^{t} S_u^*(l)\mathbb{1}\{Q_u^*(l-1) > 0\} - S_u(l)\mathbb{1}\{Q_u(l-1) > 0\} \\
&\leq \sum_{l=1}^{t} (S_u^*(l) - S_u(l))\,\mathbb{1}\{Q_u(l-1) > 0\} \\
&\leq \sum_{l=1}^{t} S_u^*(l) - S_u(l).
\end{aligned}
$$

This gives us the required result, i.e.,

$$
\mathbb{E}[\boldsymbol{Q}(t) - \boldsymbol{Q}^*(t)] \leq \sum_{l=1}^{t} \mathbb{E}[\boldsymbol{S}^*(l) - \boldsymbol{S}(l)].
$$

$\square$

8.2 Proof of Theorem 5

As mentioned, we prove all our results for the generalized setting of a switch network with a unique optimal matching. We state and prove a few intermediate lemmas that are useful in proving Theorem 5. Proof of Lemma 11 is given in separate parts for Q-UCB and Q-ThS. All other proofs in Appendix 8, although given for the algorithm Q-UCB, follow in an identical manner for Q-ThS. Important notation used in the two algorithms are summarized in Table 2.

As shown in Algorithm 1, $\mathsf{E}(t)$ indicates whether Q-UCB chooses to explore at time t. We now obtain a bound on the expected number of time-slots Q-UCB chooses to explore in an arbitrary time interval $(t_1, t_2]$. Since at any time t, Q-UCB decides to explore with probability $\min\{1, 3K\frac{\log^2 t}{t}\}$, we have

$$\mathbb{E}\left[\sum_{l=t_1+1}^{t_2} \mathsf{E}(l)\right] \leq 3K \sum_{l=t_1+1}^{t_2} \frac{\log^2 l}{l} \leq 3K \int_{t_1}^{t_2} \frac{\log^2 l}{l} \, \mathrm{d}l = K\left(\log^3 t_2 - \log^3 t_1\right). \qquad (4)$$

The following lemma gives a probabilistic upper bound on the same quantity. Recall that $w(t) = t^{(1-1/\beta)}$ as defined in Theorem 5. We will use this definition of $w(t)$ for all the lemmas in this section.

Lemma 10. (a) For any t and $t_1 < t_2$,

$$\mathbb{P}\left[\sum_{l=t_1+1}^{t_2} \mathsf{E}(l) \geq 5 \max\left(\log t, K\left(\log^3 t_2 - \log^3 t_1\right)\right)\right] \leq \frac{1}{t^4}.$$

(b) For $t \geq 5.8 \times 10^3$,

$$\mathbb{P}\left[\sum_{l=1}^{t} \mathsf{E}(l) > Kw(t)\right] \leq \frac{1}{t^{2K}}.$$

To prove the result, we will use the following Chernoff bound: for a sum of independent Bernoulli random variables Y with mean $\mathbb{E}Y$ and for any $\delta > 0$,

$$\mathbb{P}\left[Y \geq (1+\delta)\mathbb{E}Y\right] \leq \left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^{\mathbb{E}Y}. \qquad (5)$$

Proof of Lemma 10(a). If $\mathbb{E}Y \geq \log t$, the above bound for $\delta = 4$ gives

$$\mathbb{P}\left[Y \geq 5\mathbb{E}Y\right] \leq \frac{1}{t^4}.$$

Note that $\{\mathsf{E}(l)\}_{l=t_1+1}^{t_2}$ are independent Bernoulli random variables and let $X = \sum_{l=t_1}^{t_2} \mathsf{E}(l)$. Now consider the probability $\mathbb{P}\left[X \geq 5 \max\left(\log t, \mathbb{E}X\right)\right]$. If $\mathbb{E}X \geq \log t$, then the result is true from the above Chernoff bound. If $\mathbb{E}X < \log t$, then it is possible to construct a random variable Y which is a sum of independent Bernoulli random variables, has mean $\log t$ and stochastically dominates X, in which case we can again use the Chernoff bound on Y. Therefore,

$$\mathbb{P}\left[X \geq 5\log t\right] \leq \mathbb{P}\left[Y \geq 5\log t\right] \leq \frac{1}{t^4}.$$

Using inequality (4), we have the required result, i.e.,

$$\mathbb{P}\left[\sum_{l=t_1+1}^{t_2} \mathsf{E}(l) \geq 5 \max\left(\log t, K\left(\log^3 t_2 - \log^3 t_1\right)\right)\right] \leq \mathbb{P}\left[X \geq 5 \max\left(\log t, \mathbb{E}X\right)\right] \leq 1/t^4.$$

$\square$

Proof of Lemma 10(b). Using the following alternate form of the Chernoff bound (5)

$$\mathbb{P}\left[Y > y\right] \le \frac{1}{e\mathbb{E}Y}\left(\frac{e\mathbb{E}Y}{y}\right)^y,$$

we get,

$$\mathbb{P}\left[\sum_{l=1}^{t}\mathsf{E}(l) > Kw(t)\right] \le \mathbb{P}\left[\sum_{l=1}^{t}\mathsf{E}(l) > K\exp\left((2\log t)^{2/3}\right)\right]$$

$$\le \frac{1}{e^{K\log^3 t}}\left(\frac{e\log^3 t}{\exp\left((2\log t)^{2/3}\right)}\right)^{K\exp\left((2\log t)^{2/3}\right)}$$

$$\le \frac{1}{t^{2K}}$$

for all $t \ge 5.8 \times 10^3$. To justify the last inequality, we first verify that the function

$$e^{(2x)^{2/3}}\left((2x)^{2/3} - 1 - 3\log x\right) + x^3 - 2x$$

is positive for all $x \ge \log(5800)$. The result then follows by taking $x = \log t$. $\square$

The next lemma shows that, with high probability, Q-UCB (or Q-ThS) does not schedule a sub-optimal matching when it exploits in the late stage.

Lemma 11. *Define the event* $\mathcal{E}_1 := \sum_{l=w(t)+1}^{t}\sum_{u\in[U]}\sum_{k\neq k_u^*}\mathsf{I}_{uk}(l) = 0$. *Then, for Q-UCB (Q-ThS),*

$$\mathbb{P}\left[\mathcal{E}_1^c\right] = \mathbb{P}\left[\bigcup_{u\in[U]}\sum_{l=w(t)+1}^{t}\sum_{k\neq k_u^*}\mathsf{I}_{uk}(l) > 0\right] \le \frac{UK}{6t^3},$$

for all $t \ge \tau_1$ $(t \ge \tau_2)$, $t \ge \exp\left(\frac{4}{\Delta^2(1-1/\beta)^3}\right)$.

Proof. Let $X_{uk}(l), u = 1, 2, \ldots, K$, $k = 1, 2, \ldots, K$, $l = 1, 2, 3, \ldots$ be independent random variables denoting the service offered in the l^{th} assignment of server k to queue u and let $B_{uk}(s, t) = \frac{1}{s}\sum_{l=1}^{s}X_{uk}(l) + \sqrt{\frac{\log^2 t}{2s}}$. Consider the events

$$T_{uk}(l) \ge \frac{1}{2}\log^3(l-1) \quad \forall k \in [K], u \in [U], w(t)+1 \le l \le t, \tag{6}$$

$$B_{uk_u^*}(s, l) > \mu_u^* \quad \forall s, l \text{ s.t. } \frac{1}{2}\log^3(w(t)) \le s \le t-1, w(t)+1 \le l \le t, \ \forall u \in [U], \tag{7}$$

and

$$B_{uk}(s, l) \le \mu_u^* \quad \forall s, l \text{ s.t. } \frac{1}{2}\log^3(w(t)) \le s \le t-1, w(t)+1 \le l \le t, \forall k \neq k_u^*, \ \forall u \in [U]. \tag{8}$$

It can be seen that, given the above events, Q-UCB schedules the optimal matching in all time-slots in $(w(t), t]$ in which it decides to exploit, i.e., $\sum_{l=w(t)+1}^{t}\sum_{k\neq k_u^*}\mathsf{I}_{uk}(l) = 0$ for all $u \in [U]$. We now show that the events above occur with high probability.

25

Note that, since the matchings in $\mathcal{X}$ cover all the links in the system, $T_{uk}(l+1) \leq \frac{1}{2} \log^3(l)$ for some u, k implies that $\sum_{s=1}^{l} \mathbb{1}\{\boldsymbol{\kappa}(s) = \boldsymbol{\kappa}\} \leq \frac{1}{2} \log^3(l)$ for some $\boldsymbol{\kappa} \in \mathcal{E}$. Since $\sum_{s=1}^{l} \mathbb{1}\{\boldsymbol{\kappa}(s) = \boldsymbol{\kappa}\}$ is a sum of i.i.d. Bernoulli random variables with sum mean at least $\log^3(l)$, we use Chernoff bound to prove that event (6) occurs with high probability. Note that $t \geq \exp\left(\frac{4}{\Delta^2(1-1/\beta)^3}\right)$ implies that $\log^3(w(t)) = (1 - 1/\beta)^3 \log^3 t \geq \left(\frac{2\log t}{\Delta}\right)^2$. Therefore,

$$\log(w(t)) \geq \left(\frac{2\log t}{\Delta}\right)^{2/3}. \tag{9}$$

This gives us

$$
\begin{aligned}
\mathbb{P}\left[(6) \text{ is false}\right] &\leq \sum_{\boldsymbol{\kappa} \in \mathcal{E}} \sum_{l=w(t)}^{t-1} \mathbb{P}\left[\sum_{s=1}^{l} \mathbb{1}\{\boldsymbol{\kappa}(s) = \boldsymbol{\kappa}\} \leq \frac{1}{2}\log^3(l)\right] \\
&\leq Kt \exp\left(-\frac{1}{8}\log^3(w(t))\right) \\
&\leq Kt \exp\left(-\frac{1}{8}\left(\frac{2\log t}{\Delta}\right)^2\right) \\
&\leq Kt \exp\left(-\frac{1}{2}\log^2 t\right). \tag{10}
\end{aligned}
$$

Similarly, probability of events (7) and (8) can be bounded as follows –

$$
\begin{aligned}
\mathbb{P}\left[(7) \text{ is false}\right] &\leq \sum_{u \in [U]} \sum_{l=w(t)+1}^{t} \sum_{s=\frac{1}{2}\log^3(w(t))}^{t-1} \mathbb{P}\left[B_{uk_u^*}(s, l) \leq \mu_u^*\right] \\
&\leq U \sum_{l=w(t)+1}^{t} \sum_{s=\frac{1}{2}\log^3(w(t))}^{t-1} \exp\left(-\log^2(l)\right) \\
&\leq Ut^2 \exp\left(-\log^2(w(t))\right) \\
&\leq Ut^2 \exp\left(-\left(\frac{2\log t}{\Delta}\right)^{4/3}\right) \\
&\leq Ut^2 \exp\left(-(2\log t)^{4/3}\right). \tag{11}
\end{aligned}
$$

$$\mathbb{P}\left[(8) \text{ is false}\right] \leq \sum_{u \in [U], k \neq k_u^*} \sum_{l=w(t)+1}^{t} \sum_{s=\frac{1}{2}\log^3(w(t))}^{t-1} \mathbb{P}\left[B_{uk}(s,l) > \mu_u^*\right]$$

$$\leq \sum_{u \in [U], k \neq k_u^*} \sum_{l=w(t)+1}^{t} \sum_{s=\frac{1}{2}\log^3(w(t))}^{t-1} \mathbb{P}\left[B_{uk}(s,l) > \Delta + \mu_{uk}\right]$$

$$\leq UK \sum_{l=w(t)+1}^{t} \sum_{s=\frac{1}{2}\log^3(w(t))}^{t-1} \exp\left(-2s\left(\Delta - \sqrt{\frac{\log^2 l}{2s}}\right)^2\right)$$

$$\leq UKt^2 \exp\left(-\log^3(w(t))\left(\Delta - \sqrt{\frac{\log^2 t}{\log^3(w(t))}}\right)^2\right)$$

$$\leq UKt^2 \exp\left(-\log^2 t\right). \tag{12}$$

Combining the inequalities (10), (11) and (12) we have

$$\mathbb{P}\left[\mathcal{E}_1^c\right] \leq \mathbb{P}\left[(6) \text{ is false}\right] + \mathbb{P}\left[(7) \text{ is false}\right] + \mathbb{P}\left[(8) \text{ is false}\right]$$

$$\leq Kt \exp\left(-\frac{1}{2}\log^2 t\right) + Ut^2 \exp\left(-(2\log t)^{4/3}\right) + UKt^2 \exp\left(-\log^2 t\right) \tag{13}$$

$$\leq UKt\left(\exp\left(-\frac{1}{2}\log^2 t\right) + \frac{t}{2}\exp\left(-(2\log t)^{4/3}\right) + t\exp\left(-\log^2 t\right)\right)$$

$$\leq \frac{UK}{6t^3}$$

$\forall\, t \geq 5800$. This proves the result for Q-UCB.

The proof of this result for Q-ThS follows in a similar fashion. For Q-ThS, events (7) and (8) are substituted with the following events

$$\theta_{uk_u^*}(s) > \mu_u^* - \frac{\log(s-1)}{\sqrt{2T_{uk_u^*}(s)}}, \quad \forall s, \text{ s.t. } w(t)+1 \leq s \leq t, u \in [U] \tag{14}$$

and

$$\theta_{uk}(s) \leq \mu_u^* - \frac{\log(s-1)}{\sqrt{2T_{uk_u^*}(s)}}, \quad \forall s, k \text{ s.t. } w(t)+1 \leq s \leq t, k \neq k_u^*, u \in [U]. \tag{15}$$

It is then sufficient to prove that the above two events occur with high probability. Given events (6), (14), (15), Q-ThS schedules the optimal matching in all time-slots in $(w(t), t]$ in which it decides to exploit, i.e., $\sum_{l=w(t)+1}^{t} \sum_{k \neq k_u^*} I_{uk}(l) = 0$ for all $u \in [U]$.

Let $\Sigma_{u,k,l} = \sum_{r=1}^{l} X_{uk}(r)$ and $S_{uk}(l) = \hat{\mu}_{uk}(l)T_{uk}(l) = \Sigma_{u,k,T_{uk}(l)}$ for all $u \in [U], k \in [K], l \in \mathbb{N}$. We use the 'Beta-Binomial trick' (used in Kaufmann et al. (2012), Agrawal and Goyal (2012)), which gives a relation between the c.d.fs of Beta and Binomial distributions to prove the high probability results. Let $F_{a,b}^{\text{Beta}}$ and $F_{n,p}^{\text{B}}$ denote the c.d.f of Beta(a,b) distribution and the c.d.f. of Binomial(n,p) distribution respectively. Then

$$F_{a,b}^{\text{Beta}}(y) = 1 - F_{a+b-1,y}^{\text{B}}(a-1).$$

For each $s, l \in \mathbb{N}$, let $\{Z_{s,l}(r)\}_{r>0}$ be a sequence of i.i.d. Bernoulli random variables with mean $\mu_u^* - \frac{\log(s-1)}{\sqrt{2l}}$. Now, to bound probability of event (14),

$$\mathbb{P}\left[(14) \text{ is false} \bigcap (6) \text{ is true}\right] \leq \sum_{u \in [U]} \sum_{s=w(t)+1}^{t} \mathbb{P}\left[\theta_{uk_u^*}(s) \leq \mu_u^* - \frac{\log(s-1)}{\sqrt{2T_{uk_u^*}(s)}} \bigcap (6) \text{ is true}\right].$$

Now, for any $u \in [U]$, $w(t)+1 \leq s \leq t$, we have

$$\mathbb{P}\left[\theta_{uk_u^*}(s) \leq \mu_u^* - \frac{\log(s-1)}{\sqrt{2T_{uk_u^*}(s)}} \bigcap (6) \text{ is true}\right]$$

$$= \sum_{l=\frac{1}{2}\log^3(s-1)}^{s} \mathbb{E}\left[\mathbb{1}\left\{T_{uk_u^*}(s) = l\right\} F_{S_{uk_u^*}(s)+1,T_{uk_u^*}(s)-S_{uk_u^*}(s)+1}^{\text{Beta}}\left(\mu_u^* - \frac{\log(s-1)}{\sqrt{2T_{uk_u^*}(s)}}\right)\right]$$

$$= \sum_{l=\frac{1}{2}\log^3(s-1)}^{s} \mathbb{E}\left[1 - F_{l+1,\mu_u^*-\frac{\log(s-1)}{\sqrt{2l}}}^{\text{B}}\left(\Sigma_{u,k_u^*,l}\right)\right]$$

$$\leq \sum_{l=\frac{1}{2}\log^3(s-1)}^{s} \mathbb{P}\left[\Sigma_{u,k_u^*,l} \leq \sum_{r=1}^{l+1} Z_{s,l}(r)\right]$$

$$\leq \sum_{l=\frac{1}{2}\log^3(s-1)}^{s} \exp\left(-\frac{\log^2(s-1)}{4}\right)$$

$$\leq t \exp\left(-\frac{1}{4}\log^2(w(t))\right)$$

$$\leq t \exp\left(-\frac{1}{4}(2\log t)^{4/3}\right).$$

The last inequality follows by using (10) to bound the first term and Chernoff-Hoeffding inequality to bound the second term.

Similarly, the probability of event (15) can be bounded as follows.

$$\mathbb{P}\left[(15) \text{ is false} \bigcap (6) \text{ is true}\right] \leq \sum_{u \in [U], k \neq k_u^*} \sum_{s=w(t)+1}^{t} \mathbb{P}\left[\theta_{uk}(s) > \mu_u^* - \frac{\log(s-1)}{\sqrt{2T_{uk_u^*}(s)}} \bigcap (6) \text{ is true}\right]$$

28

Now, for any $u \in [U], k \neq k_u^*, w(t) + 1 \leq s \leq t$, we have

$$\mathbb{P}\left[\theta_{uk}(s) > \mu_u^* - \frac{\log(s-1)}{\sqrt{2T_{uk_u^*}(s)}} \bigcap (6) \text{ is true}\right]$$

$$\leq \sum_{l=\frac{1}{2}\log^3(s-1)}^{s} \mathbb{P}\left[T_{uk_u^*}(s) = l\right] \mathbb{E}\left[1 - F_{\Sigma_{u,k,l}+1,l-\Sigma_{u,k,l}+1}^{\text{Beta}}\left(\mu_u^* - \frac{\log(s-1)}{\sqrt{2l}}\right) \Bigg| T_{uk_u^*}(s) = l\right]$$

$$= \sum_{l=\frac{1}{2}\log^3(s-1)}^{s} \mathbb{P}\left[T_{uk_u^*}(s) = l\right] \mathbb{E}\left[F_{l+1,\mu_u^*-\frac{\log(s-1)}{\sqrt{2l}}}^{\text{B}}\left(\Sigma_{u,k,l}\right) \Bigg| T_{uk_u^*}(s) = l\right]$$

$$\leq \sum_{l=\frac{1}{2}\log^3(s-1)}^{s} \mathbb{P}\left[\sum_{r=1}^{l+1} Z_{s,l}(r) \leq \Sigma_{u,k,l}\right]$$

$$\leq \sum_{l=\frac{1}{2}\log^3(s-1)}^{s} \exp\left(-\frac{l}{2}\left(\Delta - \sqrt{\frac{\log^2(s-1)}{2l}}\right)^2\right)$$

$$\leq t \exp\left(-\frac{1}{4}\log^3(w(t))\left(\Delta - \sqrt{\frac{\log^2 t}{\log^3(w(t))}}\right)^2\right)$$

$$\leq t \exp\left(-\frac{1}{4}\log^2 t\right).$$

Combining the above inequalities and (10), we have

$$\mathbb{P}\left[\mathcal{E}_1^c\right] \leq \mathbb{P}\left[(6) \text{ is false}\right] + \mathbb{P}\left[(14) \text{ is false} \bigcap (6) \text{ is true}\right] + \mathbb{P}\left[(15) \text{ is false} \bigcap (6) \text{ is true}\right]$$

$$\leq Kt \exp\left(-\frac{1}{2}\log^2 t\right) + Ut^2 \exp\left(-\frac{1}{4}(2\log t)^{4/3}\right) + UKt^2 \exp\left(-\frac{1}{4}\log^2 t\right). \qquad (16)$$

Note from (3) that the last expression is $\leq \frac{UK}{6t^3}$ for $t \geq \tau_2$, which proves the result for Q-ThS. $\qquad \square$

For any time t, let

$$B_u(t) := \min\{s \geq 0 : Q_u(t-s) = 0\}$$

denote the time elapsed since the beginning of the current regenerative cycle for queue u. Alternately, at any time t, $t - B_u(t)$ is the last time instant at which queue u was zero.

The following lemma gives an upper bound on the sample-path queue-regret in terms of the number of sub-optimal schedules in the current regenerative cycle.

Lemma 12. *For any $t \geq 1$,*

$$Q_u(t) - Q_u^*(t) \leq \sum_{l=t-B_u(t)+1}^{t} \left(\mathsf{E}(l) + \sum_{k \neq k_u^*} \mathsf{I}_{uk}(l)\right).$$

Proof. If $B_u(t) = 0$, i.e., if $Q_u(t) = 0$, then the result is trivially true.

Consider the case where $B_u(t) > 0$. Since $Q_u(l) > 0$ for all $t - B_u(t) + 1 \leq l \leq t$, we have

$$Q_u(l) = Q_u(l-1) + A_u(l) - S_u(l) \quad \forall t - B_u(t) + 1 \leq l \leq t.$$

29

This implies that

$$Q_u(t) = \sum_{l=t-B_u(t)+1}^{t} A_u(l) - S_u(l).$$

Moreover,

$$Q_u^*(t) = \max_{1 \leq s \leq t} \left(Q_u^*(0) + \sum_{l=s}^{t} A_u(l) - S_u^*(l) \right)^+ \geq \sum_{l=t-B_u(t)+1}^{t} A_u(l) - S_u^*(l).$$

Combining the above two expressions, we have

$$
\begin{aligned}
Q_u(t) - Q_u^*(t) &\leq \sum_{l=t-B_u(t)+1}^{t} S_u^*(l) - S_u(l) \\
&= \sum_{l=t-B_u(t)+1}^{t} \sum_{k \in [K]} \left(R_{uk_u^*}(l) - R_{uk}(l) \right) \left(\mathsf{E}_{uk}(l) + \mathsf{I}_{uk}(l) \right) \\
&\leq \sum_{l=t-B_u(t)+1}^{t} \sum_{k \neq k_u^*} \left(\mathsf{E}_{uk}(l) + \mathsf{I}_{uk}(l) \right) \\
&\leq \sum_{l=t-B_u(t)+1}^{t} \left(\mathsf{E}(l) + \sum_{k \neq k_u^*} \mathsf{I}_{uk}(l) \right),
\end{aligned}
$$

where the second inequality follows from the assumption that the service provided by each of the links is bounded by 1, and the last inequality from the fact that $\sum_{k \in [K]} \mathsf{E}_{uk}(l) = \mathsf{E}(l) \ \forall l, \forall u \in [U]$. □

In the next lemma, we derive a coarse high probability upper bound on the queue-length. This bound on the queue-length is used later to obtain a first cut bound on the length of the regenerative cycle in Lemma 14.

Lemma 13. *Define the event* $\mathcal{E}_2 := \left\{ \sum_{l=1}^{t} \mathsf{E}(l) \leq K w(t) \right\}$. *Then for any* $l \in [1, t]$,

$$\mathbb{P}\left[\{ Q_u(l) > 2K w(t) \} \cap \mathcal{E}_1 \cap \mathcal{E}_2 \right] \leq \frac{1}{t^3}$$

$\forall t \ s.t. \ \frac{w(t)}{\log t} \geq \frac{2}{\epsilon_u}$.

Proof. We show this result for $l = t$ but the same argument holds for any $l \in [1, t]$. From Lemma 12,

$$Q_u(t) - Q_u^*(t) \leq \sum_{l=t-B_u(t)+1}^{t} \left(\mathsf{E}(l) + \sum_{k \neq k_u^*} \mathsf{I}_{uk}(l) \right) \leq \sum_{l=1}^{t} \left(\mathsf{E}(l) + \sum_{k \neq k_u^*} \mathsf{I}_{uk}(l) \right).$$

Since $Q_u^*(t)$ is distributed according to $\pi_{(\lambda_u, \mu_u^*)}$,

$$\mathbb{P}\left[Q_u^*(t) > w(t) \right] = \frac{\lambda_u}{\mu_u^*} \left(\frac{\lambda_u (1 - \mu_u^*)}{(1 - \lambda_u) \mu_u^*} \right)^{w(t)} \leq \exp\left(w(t) \log\left(\frac{\lambda_u (1 - \mu_u^*)}{(1 - \lambda_u) \mu_u^*} \right) \right) \leq \frac{1}{t^3}$$

if $\frac{w(t)}{\log t} \geq \frac{2}{\epsilon_u}$. The last inequality follows from the following bound –

$$\log\left(\frac{(1 - \lambda_u)\,\mu_u^*}{\lambda_u\,(1 - \mu_u^*)}\right) = \log\left(1 + \frac{\epsilon_u}{\lambda_u\,(1 - \mu_u^*)}\right)$$
$$\geq \log\left(1 + 4\epsilon_u\right) \quad \text{since } (\lambda_u\,(1 - \mu_u^*) < 1/4)$$
$$\geq \frac{3}{2}\epsilon_u.$$

Now, note that, given $\mathcal{E}_1$,

$$\sum_{l=1}^{t} \sum_{k \neq k_u^*} \mathsf{I}_{uk}(l) \leq (K - 1)w(t) + \sum_{l=w(t)+1}^{t} \sum_{k \neq k_u^*} \mathsf{I}_{uk}(l) = (K - 1)w(t).$$

Using the inequalities above, we have

$$\mathbb{P}\left[\{Q_u(l) > 2Kw(t)\} \cap \mathcal{E}_1 \cap \mathcal{E}_2\right] \leq \mathbb{P}\left[Q_u^*(t) > w(t)\right] \leq \frac{1}{t^3}.$$

$\square$

Lemma 14. *Let* $v_u'(t) = \frac{6K}{\epsilon_u}w(t)$ *and let* v_u *be an arbitrary function. Then,*

$$\mathbb{P}\left[\{B_u\,(t - v_u(t)) > v_u'(t)\} \cap \mathcal{E}_1 \cap \mathcal{E}_2\right] \leq \frac{2}{t^3}$$

$\forall t$ *s.t.* $\frac{w(t)}{\log t} \geq \frac{2}{\epsilon_u}$ *and* $v_u(t) + v_u'(t) \leq t/2$.

Proof. Let $r(t) := t - v_u(t)$. Consider the events

$$Q_u(r(t) - v_u'(t)) \leq 2Kw(t), \tag{17}$$

$$\sum_{l=r(t)-v_u'(t)+1}^{r(t)} A_u(l) - R_{uk_u^*}(l) \leq -\frac{\epsilon_u}{2}v_u'(t), \tag{18}$$

$$\sum_{l=r(t)-v_u'(t)+1}^{r(t)} \mathsf{E}(l) + \sum_{k \neq k_u^*} \mathsf{I}_{uk}(l) \leq Kw(t). \tag{19}$$

By the definition of $v_u'(t)$,

$$2Kw(t) - \frac{\epsilon_u}{2}v_u'(t) \leq -Kw(t).$$

Given Events (17)-(19), the above inequality implies that

$$Q_u(r(t) - v_u'(t)) + \sum_{l=r(t)-v_u'(t)+1}^{r(t)} A_u(l) \leq \sum_{l=r(t)-v_u'(t)+1}^{r(t)} R_{uk_u^*}(l) - \left(\mathsf{E}(l) + \sum_{k \neq k_u^*} \mathsf{I}_{uk}(l)\right)$$

$$\leq \sum_{l=r(t)-v_u'(t)+1}^{r(t)} S_u(l),$$

which further implies that $Q_u(l) = 0$ for some $l \in [r(t) - v_u'(t) + 1, r(t)]$. This gives us that $B_u(r(t)) \leq v_u'(t)$.

31

Since $v_u(t) + v_u'(t) \leq t/2$, we have $r(t) - v_u'(t) \geq t/2 > w(t)$. Thus, the event $\mathcal{E}_1 \cap \mathcal{E}_2$ implies event (19).

Now, consider the event (18) and note that $A_u(l) - R_{uk_u^*}(l)$ are i.i.d. random variables with mean $-\epsilon_u$ and bounded between -1 and 1. Using Chernoff bound for sum of bounded i.i.d. random variables, we have

$$\mathbb{P}\left[\sum_{l=r(t)-v_u'(t)+1}^{r(t)} A_u(l) - R_{uk_u^*}(l) > -\frac{\epsilon_u}{2}v_u'(t)\right] \leq \exp\left(-\frac{\epsilon_u^2}{8}v_u'(t)\right) \leq \frac{1}{t^3}$$

since $v_u'(t) \geq \frac{6K}{\epsilon_u}w(t) \geq \frac{24}{\epsilon_u^2}\log t$.

Further, by Lemma 13, $\forall t$ s.t. $\frac{w(t)}{\log t} \geq \frac{2}{\epsilon_u}$,

$$\mathbb{P}\left[\{Q_u(r(t) - v_u'(t)) > 2Kw(t)\} \cap \mathcal{E}_1 \cap \mathcal{E}_2\right] \leq \frac{1}{t^3}.$$

Therefore, $\forall t$ s.t. $\frac{w(t)}{\log t} \geq \frac{2}{\epsilon_u}$ and $v_u(t) + v_u'(t) \leq t/2$, we get

$$\mathbb{P}\left[\{B_u\left(t - v_u(t)\right) > v_u'(t)\} \cap \mathcal{E}_1 \cap \mathcal{E}_2\right]$$
$$\leq \mathbb{P}\left[\sum_{l=r(t)-v_u'(t)+1}^{r(t)} A_u(l) - R_{uk_u^*}(l) > -\frac{\epsilon_u}{2}v_u'(t)\right] + \mathbb{P}\left[\{Q_u(r(t) - v_u'(t)) > 2Kw(t)\} \cap \mathcal{E}_1 \cap \mathcal{E}_2\right]$$
$$\leq \frac{2}{t^3}.$$

$\square$

Using the preceding upper bound on the regenerative cycle-length, we derive tighter bounds on the queue-length and the regenerative cycle-length in Lemmas 16 and 17 respectively. The following lemma is a useful intermediate result.

Lemma 15. *For any $u \in [U]$ and t_2 s.t. $1 \leq t_2 \leq t$,*

$$\mathbb{P}\left[\max_{1 \leq s \leq t_2}\left\{\sum_{l=t_2-s+1}^{t_2} A_u(l) - R_{uk_u^*}(l)\right\} \geq \frac{2\log t}{\epsilon_u}\right] \leq \frac{1}{t^3}.$$

Proof. Let $X_s = \sum_{l=t_2-s+1}^{t_2} A_u(l) - R_{uk_u^*}(l)$. Since X_s is the sum of s i.i.d. random variables with mean ϵ_u and is bounded within $[-1, 1]$, Hoeffding's inequality gives

$$\mathbb{P}\left[X_s \geq \frac{2\log t}{\epsilon_u}\right] = \mathbb{P}\left[X_s - \mathbb{E}X_s \geq \epsilon_u s + \frac{2\log t}{\epsilon_u}\right]$$
$$\leq \exp\left(-\frac{2\left(\epsilon_u s + \frac{2\log t}{\epsilon_u}\right)^2}{4s}\right)$$
$$\leq \exp\left(-4\log t\right),$$

where the last inequality follows from the fact that $(a + b)^2 > 4ab$ for any $a, b \geq 0$. Using union bound over all $1 \leq s \leq t_2$ gives the required result. $\square$

Lemma 16. *Let $v'_u(t) = \frac{6K}{\epsilon_u} w(t)$ and v_u be an arbitrary function. Then,*

$$\mathbb{P}\left[\left\{Q_u(t - v_u(t)) > \left(\frac{2}{\epsilon_u} + 5\right) \log t + 30K \frac{v'_u(t) \log^2 t}{t}\right\} \cap \mathcal{E}_1 \cap \mathcal{E}_2\right] \leq \frac{3}{t^3} + \frac{1}{t^4}$$

$\forall t$ *s.t.* $\frac{w(t)}{\log t} \geq \frac{2}{\epsilon_u}$ *and* $v_u(t) + v'_u(t) \leq t/2$.

Proof. Let $r(t) = t - v_u(t)$. Now, consider the events

$$B_u(r(t)) \leq v'_u(t), \tag{20}$$

$$\sum_{l=r(t)-s+1}^{r(t)} A_u(l) - R_{uk_u^*}(l) \leq \frac{2 \log t}{\epsilon_u} \quad \forall 1 \leq s \leq v'_u(t), \tag{21}$$

$$\sum_{l=r(t)-v'_u(t)+1}^{r(t)} \mathsf{E}(l) + \sum_{k \neq k_u^*} \mathsf{I}_{uk}(l) \leq 5 \log t + 5K \left(\log^3 (r(t)) - \log^3 (r(t) - v'_u(t))\right). \tag{22}$$

Given the above events, we have

$$
\begin{aligned}
Q_u(r(t)) &= \sum_{l=r(t)-B_u(r(t))+1}^{r(t)} A_u(l) - S(l) \\
&\leq \sum_{l=r(t)-B_u(r(t))+1}^{r(t)} A_u(l) - R_{uk_u^*}(l) + \mathsf{E}(l) + \sum_{k \neq k^*} \mathsf{I}_{uk}(l) \\
&\leq \left(\frac{2}{\epsilon_u} + 5\right) \log t + 5K \left(\log^3 (r(t)) - \log^3 (r(t) - v'_u(t))\right) \\
&\leq \left(\frac{2}{\epsilon_u} + 5\right) \log t + 15K \frac{v'_u(t) \log^2 t}{(r(t) - v'_u(t))} \\
&\leq \left(\frac{2}{\epsilon_u} + 5\right) \log t + 30K \frac{v'_u(t) \log^2 t}{t},
\end{aligned}
$$

where the last inequality is true if $v_u(t) + v'_u(t) \leq t/2$.

Using Lemmas 10(a), 14, 15, $\forall t$ s.t. $\frac{w(t)}{\log t} \geq \frac{2}{\epsilon_u}$ and $v_u(t) + v'_u(t) \leq t/2$, we get

$$\mathbb{P}\left[\left\{Q_u(t - v_u(t)) > \left(\frac{2}{\epsilon_u} + 5\right) \log t + 30K \frac{v'_u(t) \log^2 t}{t}\right\} \cap \mathcal{E}_1 \cap \mathcal{E}_2\right]$$

$$\leq \mathbb{P}\left[\left\{B_u(t - v_u(t)) > v'_u(t)\right\} \cap \mathcal{E}_1 \cap \mathcal{E}_2\right] + \mathbb{P}\left[\max_{1 \leq s \leq t_2} \left\{\sum_{l=t_2-s+1}^{t_2} A_u(l) - R_{uk_u^*}(l)\right\} \geq \frac{2 \log t}{\epsilon_u}\right]$$

$$+ \mathbb{P}\left[\sum_{l=r(t)-v'_u(t)+1}^{r(t)} \mathsf{E}(l) > 5 \log t + 5K \left(\log^3 (r(t)) - \log^3 (r(t) - v'_u(t))\right)\right]$$

$$\leq \frac{3}{t^3} + \frac{1}{t^4}.$$

$\square$

Lemma 17. *Let $v'_u(t) = \frac{6K}{\epsilon_u}w(t)$ and $v_u(t) = \frac{24\log t}{\epsilon_u^2} + \frac{60K}{\epsilon_u}\frac{v'_u(t)\log^2 t}{t}$. Then,*

$$\mathbb{P}\left[\{B_u(t) > v_u(t)\} \cap \mathcal{E}_1 \cap \mathcal{E}_2\right] \leq \frac{4}{t^3} + \frac{2}{t^4}$$

$\forall t$ *s.t.* $\frac{w(t)}{\log t} \geq \frac{2}{\epsilon_u}$ *and* $v_u(t) + v'_u(t) \leq t/2$.

Proof. Let $r(t) = t - v_u(t)$. As in Lemma 14, consider the events

$$Q_u(r(t)) \leq \left(\frac{2}{\epsilon_u} + 5\right)\log t + 30K\frac{v'_u(t)\log^2 t}{t}, \tag{23}$$

$$\sum_{l=r(t)+1}^{t} A_u(l) - R_{uk_u^*}(l) \leq -\frac{\epsilon_u}{2}v_u(t), \tag{24}$$

$$\sum_{l=r(t)+1}^{t} \mathsf{E}(l) + \sum_{k\neq k_u^*} \mathsf{I}_{uk}(l) \leq 5\log t + 5K\left(\log^3 t - \log^3(r(t))\right). \tag{25}$$

The definition of $v_u(t)$ and events (23)-(25) imply that

$$Q_u(r(t)) + \sum_{l=r(t)+1}^{t} A_u(l) \leq \sum_{l=r(t)+1}^{t} R_{uk_u^*}(l) - \sum_{l=r(t)+1}^{t} \mathsf{E}(l) + \sum_{k\neq k_u^*} \mathsf{I}_{uk}(l)$$

$$\leq \sum_{l=r(t)+1}^{t} S_u(l),$$

which further implies that $Q(l) = 0$ for some $l \in [r(t)+1, t]$ and therefore $B_u(t) \leq v_u(t)$. We can again show that each of the events (23)-(25) occurs with high probability. Particularly, we can bound the probability of event (24) in the same way as event (21) in Lemma 14 to show that it occurs with probability at least $1-1/t^3$. Combining this with Lemmas 10(a) and 16, $\forall t$ s.t. $\frac{w(t)}{\log t} \geq \frac{2}{\epsilon_u}$ and $v_u(t) + v'_u(t) \leq t/2$, we get

$$\mathbb{P}\left[\{B_u(t) > v_u(t)\} \cap \mathcal{E}_1 \cap \mathcal{E}_2\right]$$

$$\leq \mathbb{P}\left[\left\{Q_u(r(t)) > \left(\frac{2}{\epsilon_u} + 5\right)\log t + 30K\frac{v'_u(t)\log^2 t}{t}\right\} \cap \mathcal{E}_1 \cap \mathcal{E}_2\right]$$

$$+ \mathbb{P}\left[\sum_{l=r(t)+1}^{t} A_u(l) - R_{uk_u^*}(l) > -\frac{\epsilon_u}{2}v_u(t)\right]$$

$$+ \mathbb{P}\left[\sum_{l=r(t)+1}^{t} \mathsf{E}(l) > 5\log t + 5K\left(\log^3 t - \log^3(r(t))\right)\right]$$

$$\leq \frac{4}{t^3} + \frac{2}{t^4}.$$

$\square$

Proof of Theorem 5. The proof is based on two main ideas: one is that the regenerative cycle length is not very large, and the other is that the algorithm has correctly identified the optimal matching in late stages. We combine Lemmas 11 and 17 to bound the regret at any time t s.t. $t \geq 5.8 \times 10^3$, $\frac{w(t)}{\log t} \geq \frac{2}{\epsilon_u}$ and $v_u(t) + v'_u(t) \leq t/2$:

$$\Psi_u(t) = \mathbb{E}\left[Q_u(t) - Q_u^*(t)\right]$$

$$\leq \mathbb{E}\left[(Q_u(t) - Q_u^*(t))\mathbb{1}\left\{B_u(t) \leq v_u(t)\right\} + (Q_u(t) - Q_u^*(t))\mathbb{1}\left\{B_u(t) > v_u(t)\right\}\bigg|\mathcal{E}_1\right]\mathbb{P}\left[\mathcal{E}_1\right]$$

$$+ \mathbb{E}\left[Q_u(t) - Q_u^*(t)\bigg|\mathcal{E}_1^c\right]\mathbb{P}\left[\mathcal{E}_1^c\right].$$

Now, the first term on the r.h.s. of the above inequality can be bounded as follows:

$$\mathbb{E}\left[(Q_u(t) - Q_u^*(t))\mathbb{1}\left\{B_u(t) \leq v_u(t)\right\}\bigg|\mathcal{E}_1\right]\mathbb{P}\left[\mathcal{E}_1\right] \leq \mathbb{E}\left[\sum_{l=t-v_u(t)+1}^{t}\mathsf{E}(l)\right]$$

$$\leq K\left(\log^3(t) - \log^3(t - v_u(t))\right)$$

$$\leq 3K\log^2 t \log\left(1 + \frac{v_u(t)}{t - v_u(t)}\right)$$

$$\leq 6K\frac{v_u(t)\log^2 t}{t}.$$

To bound the rest of the terms in the inequality, we have, for all $t \geq 5.8 \times 10^3$, ((for Q-UCB,))

$$\leq \mathbb{E}\left[(Q_u(t) - Q_u^*(t))\mathbb{1}\left\{B_u(t) > v_u(t)\right\}\bigg|\mathcal{E}_1\right]\mathbb{P}\left[\mathcal{E}_1\right] + \mathbb{E}\left[Q_u(t) - Q_u^*(t)\bigg|\mathcal{E}_1^c\right]\mathbb{P}\left[\mathcal{E}_1^c\right]$$

$$\leq t\left(\mathbb{P}\left[\{B_u(t) > v_u(t)\} \cap \mathcal{E}_1 \cap \mathcal{E}_2\right] + \mathbb{P}\left[\mathcal{E}_1^c\right] + \mathbb{P}\left[\mathcal{E}_2^c\right]\right)$$

$$\leq t\left(\frac{4}{t^3} + \frac{2}{t^4} + \frac{UK}{6t^3} + \frac{1}{t^{2K}}\right) \tag{26}$$

$$\leq \frac{24.004 + UK}{6t^2}.$$

Here, inequality (26) is obtained using Lemmas 10(b), 11 and 17.

The result then follows by combining the above two bounds. □

Proof of Corollary 6. If

$$\log t \geq \max\left\{\frac{4}{\Delta^2(1 - 1/\beta)^3}, \sqrt{2}\left(\log\frac{2}{\epsilon_u}\right)^{1.5}, \beta\frac{24K}{\epsilon_u}, \beta\left(\log\log t + \log(15K^2)\right), \frac{\beta}{\beta - 1}\log\left(\frac{13.2}{K^2\epsilon_u^2}\right)\right\},$$

then we can obtain the following.

(i) Using the fact that $\log t \geq \sqrt{2}\left(\log\log t\right)^{3/2} \ \forall t \geq 3$ and that $\log t \geq \sqrt{2}\left(\log\frac{2}{\epsilon_u}\right)^{3/2}$, we have

$$2\log t \geq \sqrt{2}\left(\log\log t\right)^{3/2} + \sqrt{2}\left(\log\frac{2}{\epsilon_u}\right)^{3/2} \geq \left(\log\log t + \log\frac{2}{\epsilon_u}\right)^{3/2},$$

35

where the last inequality follows from the fact that $\sqrt{2}(a^{3/2} + b^{3/2}) \geq (a+b)^{3/2}$. Combining this with (9), we get

$$\log(w(t)) \geq \left(\frac{2\log t}{\Delta}\right)^{2/3} \geq \log\log t + \log\frac{2}{\epsilon_u},$$

which gives us that $\frac{w(t)}{\log t} \geq \frac{2}{\epsilon_u}$.

(ii) $\log t \geq \max \beta \left\{\frac{24K}{\epsilon_u}, \log\log t + \log(15K^2)\right\}$ gives us that $\frac{t}{w(t)} = t^{\frac{1}{\beta}} \geq 15K^2 \log t$ and $\frac{t}{w(t)} \geq \frac{24K}{\epsilon_u}$.

(iii) Moreover, if $\log t \geq \frac{\beta}{\beta-1}\log\left(\frac{13.2}{K^2\epsilon_u^2}\right)$, then

$$w(t) = t^{\left(1-\frac{1}{\beta}\right)} \geq \frac{13.2}{K^2\epsilon_u^2},$$

which when combined with $t \geq 15K^2 w(t)\log t$ yields

$$\frac{t}{\log t} \geq 15K^2 w(t) \geq \frac{198}{\epsilon_u^2}.$$

Therefore, the lower bound on t in Corollory 6 yields the following inequalities $\frac{w(t)}{\log t} \geq \frac{2}{\epsilon_u}$, $\frac{t}{w(t)} \geq \max\left\{\frac{24K}{\epsilon_u}, 15K^2\log t\right\}$, and $\frac{t}{\log t} \geq \frac{198}{\epsilon_u^2}$. We now show that these conditions are sufficient for the result in Theorem 5.

(i) $\frac{t}{w(t)} \geq \frac{24K}{\epsilon_u}$ implies that $v_u'(t) \leq \frac{t}{4}$,

(ii) $\frac{t}{w(t)} \geq 15K^2\log t$ implies that $\frac{24}{\epsilon_u^2}\log t \geq \frac{60K}{\epsilon_u}\frac{v_u'(t)\log^2 t}{t}$, and therefore $v_u(t) \leq \frac{48}{\epsilon_u^2}\log t$

(iii) $\frac{t}{\log t} \geq \frac{198}{\epsilon_u^2}$ implies that $v_u(t) \leq \frac{t}{4}$.

These inequalities when applied to Theorem 5 gives the required result, i.e.,

$$\begin{aligned}
\Psi_u(t) &\leq 6K\frac{v_u(t)\log^2 t}{t} + \frac{24.004 + UK}{6t^2}\\
&\leq \frac{288K\log^3 t}{\epsilon_u^2 t} + \frac{24.004 + UK}{6t^2}\\
&\leq \frac{289K\log^3 t}{\epsilon_u^2 t}.
\end{aligned}$$

$\square$

Proof of Corollary 9. Let

$$B = \min\left\{1 - \frac{1}{\beta} - \frac{1}{\gamma}, \frac{1}{\beta} - \frac{1}{\delta}, 1 - \frac{2}{\gamma}\right\},$$

and let C_2 be such that

$$C_2 \geq \max\left\{\tau_2, \exp\left(\frac{4}{\Delta^2(1-1/\beta)^3}\right), 15^{\max\{\gamma,\delta\}}\right\},$$

and
$$t^B \geq \log t \quad \forall t \geq C_2.$$

Then $\forall t \geq C_2 \max \left\{ \left(\frac{1}{\epsilon_u} \right)^\gamma, \left(\frac{K}{\epsilon_u} \right)^\beta, K^{2\delta} \right\}$, we have the following:

(i) Since $B \leq 1 - \frac{1}{\beta} - \frac{1}{\gamma}$ and $C_2 > 2^\gamma$,

$$\frac{w(t)}{\log t} = \frac{t^{1 - \frac{1}{\beta}}}{\log t} \geq \frac{t^{\frac{1}{\gamma} + B}}{\log t} \geq t^{\frac{1}{\gamma}} \geq \frac{2}{\epsilon_u}.$$

(ii) Since $B \leq \frac{1}{\beta} - \frac{1}{\delta}$ and $C_2 > 15^\delta$,

$$\frac{t}{w(t)} = t^{\frac{1}{\beta}} \geq t^{\frac{1}{\delta} + B} \geq 15 K^2 \log t.$$

Similarly, we can also show that $\frac{t}{w(t)} \geq \frac{24K}{\epsilon_u}$.

(iii) Since $B \leq 1 - \frac{2}{\gamma}$ and $C_2 \geq 15^\gamma$,

$$\frac{t}{\log t} \geq t^{1-B} \geq t^{\frac{2}{\gamma}} \geq \frac{C_2^{\frac{2}{\gamma}}}{\epsilon_u^2} > \frac{198}{\epsilon_u^2}.$$

As shown in the proof of Corollary 6, the above conditions are sufficient for the upper bound. $\qquad \square$

9 Proof of Lower Bound for α-Consistent Policies

In order to prove Theorems 3 and 8, we use techniques from existing work in the MAB literature along with some new lower bounding ideas specific to queueing systems. Specifically, we use lower bounds for α-consistent policies on the expected number of times a sub-optimal server is scheduled. This lower bound, shown (in Lemma 19) specifically for the problem of scheduling a unique optimal matching, is similar in style to the traditional bandit lower bound by Lai and Robbins (1985) but holds in the non-asymptotic setting. Also, as opposed to the traditional change of measure technique used in Lai and Robbins (1985), the proof technique is similar to those used in more recent papers like Bubeck et al. (2013), Perchet et al. (2015), Combes et al. (2015a) and uses results from hypothesis testing (Lemma 18).

Lemma 18 (Tsybakov (2008))**.** *Consider two probability measures P and Q, both absolutely continuous with respect to a given measure. Then for any event $\mathcal{A}$ we have:*

$$P(\mathcal{A}) + Q(\mathcal{A}^c) \geq \frac{1}{2} \exp\{- \min(\mathrm{KL}(P\|Q), \mathrm{KL}(Q\|P))\}.$$

Lemma 19. *For any problem instance $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ and any α-consistent policy, there exist constants τ and C s.t. for any $u \in [U]$, $k \neq k_u^*$ and $t > \tau$,*

$$\mathbb{E}\left[T_{uk}(t+1) \right] + \sum_{u' \neq u} \mathbb{1}\left\{ k_{u'}^* = k \right\} \mathbb{E}\left[T_{u'k_u^*}(t+1) \right] \geq \frac{1}{\mathrm{KL}\left(\mu_{min}, \frac{\mu^*+1}{2} \right)} \left((1-\alpha)\log t - \log(4KC) \right).$$

Proof. Without loss of generality, let the optimal servers for the U queues be denoted by the first U indices. In other words, a server $k > U$ is not an optimal server for any queue, i.e., for any $u' \in [U]$, $K \geq k > U$, $\mathbb{1}\left\{k_{u'}^* = k\right\} = 0$. Also, let $\beta = \frac{\mu^*+1}{2}$.

We will first consider the case $k \leq U$. For a fixed user u and server $k \leq U$, let u' be the queue that has k as the best server, i.e., $k_{u'}^* = k$. Now consider the two problem instances $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ and $(\boldsymbol{\lambda}, \boldsymbol{\hat{\mu}})$, where $\boldsymbol{\hat{\mu}}$ is the same as $\boldsymbol{\mu}$ except for the two entries corresponding to indices $(u, k), (u', k_u^*)$ replaced by β. Therefore, for the problem instance $(\boldsymbol{\lambda}, \boldsymbol{\hat{\mu}})$, the best servers are swapped for queues u and u' and remain the same for all the other queues. Let $\mathbb{P}_{\boldsymbol{\mu}}^t$ and $\mathbb{P}_{\boldsymbol{\hat{\mu}}}^t$ be the distributions corresponding to the arrivals, chosen servers and service obtained in the first t plays for the two instances under a fixed α-consistent policy. Recall that $T_{uk}(t+1) = \sum_{s=1}^t \mathbb{1}\{\kappa_u(s) = k\} \; \forall u \in [U], k \in [K]$. Define the event $\mathcal{A} = \{T_{uk}(t+1) > t/2\}$. By the definition of α-consistency there exists a fixed integer τ and a fixed constant C such that for all $t > \tau$ we have,

$$\mathbb{E}_{\boldsymbol{\mu}}^t \left[\sum_{s=1}^t \mathbb{1}\{\kappa_u(s) = k\} \right] \leq Ct^\alpha$$

$$\mathbb{E}_{\boldsymbol{\hat{\mu}}}^t \left[\sum_{s=1}^t \mathbb{1}\{\kappa_u(s) = k'\} \right] \leq Ct^\alpha , \; \forall k' \neq k.$$

A simple application of Markov's inequality yields

$$\mathbb{P}_{\boldsymbol{\mu}}^t(\mathcal{A}) \leq \frac{2C}{t^{1-\alpha}}$$

$$\mathbb{P}_{\boldsymbol{\hat{\mu}}}^t(\mathcal{A}^c) \leq \frac{2C(K-1)}{t^{1-\alpha}}.$$

We can now use Lemma 18 to conclude that

$$\mathrm{KL}(\mathbb{P}_{\boldsymbol{\mu}}^t || \mathbb{P}_{\boldsymbol{\hat{\mu}}}^t) \geq (1-\alpha)\log t - \log(4KC). \tag{27}$$

It is, therefore, sufficient to show that

$$\mathrm{KL}\left(\mathbb{P}_{\boldsymbol{\mu}}^t || \mathbb{P}_{\boldsymbol{\hat{\mu}}}^t\right) = \mathrm{KL}\left(\mu_{uk}, \beta\right) \mathbb{E}_{\boldsymbol{\mu}}^t[T_{uk}(t+1)] + \mathrm{KL}\left(\mu_{u'k_u^*}, \beta\right) \mathbb{E}_{\boldsymbol{\mu}}^t[T_{u'k_u^*}(t+1)].$$

For the sake of brevity we write the scheduling sequence in the first t time-slots $\{\boldsymbol{\kappa}(1), \boldsymbol{\kappa}(2), ..., \boldsymbol{\kappa}(t)\}$ as $\boldsymbol{\kappa}^{(t)}$, and similarly we define $\mathbf{A}^{(t)}$ as the number of arrivals to the queue and $\mathbf{S}^{(t)}$ as the service offered by the scheduled servers in the first t time-slots. Let $\mathbf{Z}^{(t)} = (\boldsymbol{\kappa}^{(t)}, \mathbf{A}^{(t)}, \mathbf{S}^{(t)})$. The KL-divergence term can now be written as

$$\mathrm{KL}(\mathbb{P}_{\boldsymbol{\mu}}^t || \mathbb{P}_{\boldsymbol{\hat{\mu}}}^t) = \mathrm{KL}(\mathbb{P}_{\boldsymbol{\mu}}^t(\mathbf{Z}^{(t)}) || \mathbb{P}_{\boldsymbol{\hat{\mu}}}^t(\mathbf{Z}^{(t)})).$$

We can apply the chain rule of divergence to conclude that

$$\begin{aligned}
\mathrm{KL}(\mathbb{P}_{\boldsymbol{\mu}}^t(\mathbf{Z}^{(t)}) || \mathbb{P}_{\boldsymbol{\hat{\mu}}}^t(\mathbf{Z}^{(t)})) = \; & \mathrm{KL}(\mathbb{P}_{\boldsymbol{\mu}}^t(\mathbf{Z}^{(t-1)}) || \mathbb{P}_{\boldsymbol{\hat{\mu}}}^t(\mathbf{Z}^{(t-1)})) \\
& + \mathrm{KL}(\mathbb{P}_{\boldsymbol{\mu}}^t(\boldsymbol{\kappa}(t) \mid \mathbf{Z}^{(t-1)}) || \mathbb{P}_{\boldsymbol{\hat{\mu}}}^t(\boldsymbol{\kappa}(t) \mid \mathbf{Z}^{(t-1)})) \\
& + \mathbb{E}_{\boldsymbol{\mu}}^t \left[\mathbb{1}\{\kappa_u(t) = k\}\mathrm{KL}\left(\mu_{uk}, \beta\right) + \mathbb{1}\{\kappa_{u'}(t) = k_u^*\}\mathrm{KL}\left(\mu_{u'k_u^*}, \beta\right) \right].
\end{aligned}$$

We can apply this iteratively to obtain

$$\text{KL}(\mathbb{P}^t_{\boldsymbol{\mu}}||\mathbb{P}^t_{\hat{\boldsymbol{\mu}}}) = \sum_{s=1}^{t} \mathbb{E}^t_{\boldsymbol{\mu}}\left[\mathbb{1}\{\kappa_u(s) = k\}\text{KL}\left(\mu_{uk}, \beta\right)\right]$$

$$+ \sum_{s=1}^{t} \mathbb{E}^t_{\boldsymbol{\mu}}\left[\mathbb{1}\{\kappa_{u'}(s) = k^*_u\}\text{KL}\left(\mu_{u'k^*_u}, \beta\right)\right]$$

$$+ \sum_{l=1}^{t} \text{KL}(\mathbb{P}^t_{\boldsymbol{\mu}}(\boldsymbol{\kappa}(l) \mid \mathbf{Z}^{(l-1)})||\mathbb{P}^t_{\hat{\boldsymbol{\mu}}}(\boldsymbol{\kappa}(l) \mid \mathbf{Z}^{(l-1)})) \qquad (28)$$

Note that the second summation in (28) is zero, as over a sample path the policy pulls the same servers irrespective of the parameters. Therefore, we obtain

$$\text{KL}(\mathbb{P}^t_{\boldsymbol{\mu}}||\mathbb{P}^t_{\hat{\boldsymbol{\mu}}}) = \text{KL}\left(\mu_{uk}, \beta\right)\mathbb{E}^t_{\boldsymbol{\mu}}[T_{uk}(t+1)] + \text{KL}\left(\mu_{u'k^*_u}, \beta\right)\mathbb{E}^t_{\boldsymbol{\mu}}[T_{u'k^*_u}(t+1)],$$

which can be substituted in (27) to obtain the required result for $K \leq U$.

Now, consider the case $k > U$, where $\sum_{u \in U} \mathbb{1}\{k^*_u = k\} = 0$. We again compare the two problem instances $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ and $(\boldsymbol{\lambda}, \hat{\boldsymbol{\mu}})$, where $\hat{\boldsymbol{\mu}}$ is the same as $\boldsymbol{\mu}$ except for the entry corresponding to index (u, k) replaced by β. Therefore, for the problem instance $(\boldsymbol{\lambda}, \hat{\boldsymbol{\mu}})$, the best server for user u is server k while the best servers for all other queues remain the same. We can again use the same technique as before to obtain

$$\text{KL}(\mathbb{P}^t_{\boldsymbol{\mu}}||\mathbb{P}^t_{\hat{\boldsymbol{\mu}}}) = \text{KL}\left(\mu_{uk}, \beta\right)\mathbb{E}^t_{\boldsymbol{\mu}}[T_{uk}(t+1)],$$

which, along with (27), gives the required result for $K > U$. $\qquad \square$

As a corollary of the above result, we now derive lower bound on the total expected number of sub-optimal schedules summed across all queues. In addition, we also show, for each individual queue, a lower bound for those servers which are sub-optimal for all the queues. As in the proof of Lemma 19, we assume without loss of generality that the first U indices denote the optimal servers for the U queues.

Corollary 20. *For any problem instance $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ and any α-consistent policy, there exist constants τ and C s.t. for any $t > \tau$,*

(a)

$$2\Delta \sum_{u \in [U]} \sum_{k \neq k^*_u} \mathbb{E}\left[T_{uk}(t+1)\right] \geq U(K-1)D(\boldsymbol{\mu})\left((1-\alpha)\log t - \log(4KC)\right),$$

(b) for any $u \in [U]$,

$$2\Delta \sum_{k \neq k^*_u} \mathbb{E}\left[T_{uk}(t+1)\right] \geq (U-1)D(\boldsymbol{\mu})\left((1-\alpha)\log t - \log(4KC)\right),$$

(c) and for any $u \in [U]$,

$$\Delta \sum_{k > U} \mathbb{E}\left[T_{uk}(t+1)\right] \geq (K-U)D(\boldsymbol{\mu})\left((1-\alpha)\log t - \log(4KC)\right),$$

where $D(\boldsymbol{\mu})$ is given by (2).

Proof. To prove part (a), we observe that a unique optimal server for each queue in the system implies that

$$\sum_{u \in [U]} \sum_{k \neq k_u^*} \mathbb{E}\left[T_{uk}(t+1)\right] \geq \sum_{u \in [U]} \sum_{u' \neq u} \mathbb{E}\left[T_{uk_{u'}^*}(t+1)\right]$$

$$= \sum_{u \in [U]} \sum_{k \neq k_u^*} \sum_{u' \neq u} \mathbb{1}\left\{k_{u'}^* = k\right\} \mathbb{E}\left[T_{u'k_u^*}(t+1)\right].$$

Now, from Lemma 19, there exist constants C and τ such that for $t > \tau$,

$$2 \sum_{u \in [U]} \sum_{k \neq k_u^*} \mathbb{E}\left[T_{uk}(t+1)\right] \geq \sum_{u \in [U]} \sum_{k \neq k_u^*} \left(\mathbb{E}\left[T_{uk}(t+1)\right] + \sum_{u' \neq u} \mathbb{1}\left\{k_{u'}^* = k\right\} \mathbb{E}\left[T_{u'k_u^*}(t+1)\right]\right)$$

$$\geq \frac{U(K-1)}{\mathrm{KL}\left(\mu_{min}, \frac{\mu^*+1}{2}\right)} \left((1-\alpha)\log t - \log(4KC)\right).$$

Using the definition of $D(\boldsymbol{\mu})$ in the above inequality gives part (a) of the corollary.

To prove part (b), we can assume without loss of generality that a perfect matching is scheduled in every time-slot. Using this, and the fact that any server is assigned to at most one queue in every time-slot, for any $u \in [U]$, we have

$$T_{uk_u^*}(t+1) + \sum_{k \neq k_u^*} T_{uk}(t+1) = t \geq T_{uk_u^*}(t+1) + \sum_{u' \neq u} T_{u'k_u^*}(t+1),$$

which gives us

$$\sum_{k \neq k_u^*} T_{uk}(t+1) \geq \max\left\{\sum_{u' \neq u} T_{uk_{u'}^*}(t+1), \sum_{u' \neq u} T_{u'k_u^*}(t+1)\right\}. \tag{29}$$

From Lemma 19 we have, for any $u' \neq u$ and for $t > \tau$,

$$\mathbb{E}\left[T_{uk_{u'}^*}(t+1)\right] + \mathbb{E}\left[T_{u'k_u^*}(t+1)\right] \geq \frac{1}{\mathrm{KL}\left(\mu_{min}, \frac{\mu^*+1}{2}\right)} \left((1-\alpha)\log t - \log(4KC)\right),$$

which gives

$$\sum_{u' \neq u} \mathbb{E}\left[T_{uk_{u'}^*}(t+1)\right] + \mathbb{E}\left[T_{u'k_u^*}(t+1)\right] \geq \frac{U-1}{\mathrm{KL}\left(\mu_{min}, \frac{\mu^*+1}{2}\right)} \left((1-\alpha)\log t - \log(4KC)\right).$$

Combining the above with (29), we have for $t > \tau$

$$\sum_{k \neq k_u^*} \mathbb{E}\left[T_{uk}(t+1)\right] \geq \max\left\{\sum_{u' \neq u} \mathbb{E}\left[T_{uk_{u'}^*}(t+1)\right], \sum_{u' \neq u} \mathbb{E}\left[T_{u'k_u^*}(t+1)\right]\right\}$$

$$\geq \frac{U-1}{2\mathrm{KL}\left(\mu_{min}, \frac{\mu^*+1}{2}\right)} \left((1-\alpha)\log t - \log(4KC)\right).$$

To prove part (c), we use the fact that $\mathbb{1}\left\{k_{u'}^* = k\right\} = 0$ for any $u' \in [U]$, $K \geq k > U$. Therefore, for $t > \tau$, we have

$$
\begin{aligned}
\sum_{k>U} \mathbb{E}\left[T_{uk}(t+1)\right] &= \sum_{k>U} \left(\mathbb{E}\left[T_{uk}(t+1)\right] + \sum_{u' \neq u} \mathbb{1}\left\{k_{u'}^* = k\right\} \mathbb{E}\left[T_{u'k_u^*}(t+1)\right] \right) \\
&\geq \frac{K-U}{\mathrm{KL}\left(\mu_{min}, \frac{\mu^*+1}{2}\right)} \left((1-\alpha)\log t - \log(4KC)\right),
\end{aligned}
$$

which gives the required result. $\qquad\square$

9.1 Late Stage: Proof of Theorem 3

The following lemma, which gives a lower bound on the queue-regret in terms of probability of sub-optimal schedule in a single time-slot, is the key result used in the proof of Theorem 3. The proof for this lemma is based on the idea that the growth in regret in a single-time slot can be lower bounded in terms of the probability of sub-optimal schedule in that time-slot.

Lemma 21. *For any problem instance characterized by* $(\boldsymbol{\lambda}, \boldsymbol{\mu})$*, and for any scheduling policy, and user* $u \in [U]$*,*

$$
\Psi_u(t) \geq \lambda_u \sum_{k \neq k_u^*} \Delta_{uk} \mathbb{P}\left[\mathbb{1}\{\kappa_u(t) = k\} = 1\right].
$$

Proof. We lower bound the queue-regret for queue u for the alternate coupled service process described in Section 8.1. As seen in the proof of Lemma 1, since $\mu_u^* > \mu_{uk}$ $\forall k \neq k_u^*$, for the alternate system, we have $R_{uk_u^*}(t) \geq R_{uk}(t)$ $\forall k \neq k_u^*$, $\forall t$. This implies that $Q_u^*(t) \leq Q_u(t)$ $\forall t$. Now, for any given t, using the fact that $Q_u^*(t-1) \leq Q_u(t-1)$, it is easy to see that

$$
Q_u(t) - Q_u^*(t) \geq \mathbb{1}\left\{A_u(t) = 1\right\} \left(R_{k_u^*}(t) - \sum_{k=1}^{K} \mathbb{1}\{\kappa_u(t) = k\} R_{uk}(t) \right).
$$

Therefore,

$$
\begin{aligned}
\mathbb{E}\left[Q_u(t) - Q_u^*(t)\right] &\geq \mathbb{E}\left[\mathbb{1}\{A_u(t) = 1\} \left(R_{k_u^*}(t) - \sum_{k=1}^{K} \mathbb{1}\{\kappa_u(t) = k\} R_{uk}(t) \right)\right] \\
&= \lambda_u \sum_{k \neq k_u^*} \mathbb{P}\left[\mathbb{1}\{\kappa_u(t) = k\} = 1\right] \mathbb{P}\left[\mu_{uk} < U(t) \leq \mu_u^*\right] \\
&= \lambda_u \sum_{k \neq k_u^*} \Delta_{uk} \mathbb{P}\left[\mathbb{1}\{\kappa_u(t) = k\} = 1\right].
\end{aligned}
$$

$\qquad\square$

We now use Lemma 21 in conjunction with the lower bound for the expected number of sub-optimal schedules for an α-consistent policy (Corollary 20) to prove Theorem 3.

Proof of Theorem 3. From Lemma 21 we have,

$$
\begin{aligned}
\Psi_u(t) &\geq \lambda_u \sum_{k \neq k_u^*} \Delta_{uk} \mathbb{P}\left[\mathbb{1}\{\kappa_u(t) = k\} = 1\right] \\
&\geq \lambda_{min}\Delta \sum_{k \neq k_u^*} \mathbb{P}\left[\mathbb{1}\{\kappa_u(t) = k\} = 1\right].
\end{aligned}
\tag{30}
$$

Therefore,

$$\sum_{s=1}^{t} \sum_{u \in [U]} \Psi_u(s) \geq \lambda_{min} \Delta \sum_{u \in [U]} \sum_{k \neq k_u^*} \mathbb{E}\left[T_{uk}(t+1)\right].$$

We now claim that

$$\sum_{u \in [U]} \Psi_u(t) \geq \frac{U(K-1)}{8t} \lambda_{min} D(\boldsymbol{\mu})(1-\alpha) \tag{31}$$

for infinitely many t. This follows from part (a) of Corollary 20 and the following fact:

Fact 1. *For any bounded sequence $\{a_n\}$, if there exist constants $C > 0$ and $N_0 \in \mathbb{N}$ such that $\sum_{m=1}^{n} a_m \geq C \log n \ \forall n \geq N_0$, then $a_n \geq \frac{C}{2n}$ infinitely often.*

Proof. Proof by contradiction: Let B be an upper bound on the sequence. Suppose $\sum_{m=1}^{n} a_m \geq C \log n \ \forall n \geq N_0$ and there exists $N_1 \in \mathbb{N}$ such that $a_n \leq \frac{C}{2n}$ for all $n > N_1$. Wlog let N_1 satisfy the inequality $\exp\left(\frac{2BN_1}{C} + 1\right) \geq N_1$. Then for any $n > \exp\left(\frac{2BN_1}{C} + 1\right)$, we have

$$\sum_{m=1}^{n} a_m \leq BN_1 + \sum_{m=N_1+1}^{n} \frac{C}{2n} \leq BN_1 + \frac{C}{2}(\log n + 1) < C \log n,$$

which is a contradiction to our hypothesis that $\sum_{m=1}^{n} a_m \geq C \log n \ \forall n \geq N_0$. This proves the claim. $\qquad \square$

Similarly, for any $u \in U$, it follows from parts (b) and (c) of Corollary 20 that

$$\Psi_u(t) \geq \frac{\max\{U-1, 2(K-U)\}}{8t} \lambda_{min} D(\boldsymbol{\mu})(1-\alpha) \tag{32}$$

for infinitely many t. $\qquad \square$

9.2 Early Stage: Proof of Theorem 8

In order to prove Theorem 8, we first derive, in the following lemma, a lower bound on the queue-regret in terms of the expected number of sub-optimal schedules.

Lemma 22. *For any system with parameters $(\boldsymbol{\lambda}, \boldsymbol{\mu})$, any policy, and any user $u \in [U]$, the regret is lower bounded by*

$$\Psi_u(t) \geq \sum_{k \neq k_u^*} \Delta_{uk} \mathbb{E}\left[T_{uk}(t+1)\right] - \epsilon_u t.$$

Proof. Since $Q_u(0) \sim \pi_{\lambda_u, \mu_u^*}$, we have,

$$
\begin{aligned}
\Psi_u(t) &= \mathbb{E}\left[Q_u(t) - Q_u^*(t)\right] \\
&= \mathbb{E}\left[Q_u(t) - Q_u(0)\right] \\
&\geq \mathbb{E}\left[\sum_{l=1}^{t} A_u(l) - S_u(l)\right] \\
&= \lambda_u t - \sum_{k=1}^{K} \mathbb{E}\left[T_{uk}(t+1)\right] \mu_{uk} \\
&= \lambda_u t - \left(t - \sum_{k \neq k_u^*} \mathbb{E}\left[T_{uk}(t+1)\right]\right) \mu *_u - \sum_{k \neq k_u^*} \mathbb{E}\left[T_{uk}(t+1)\right] \mu_{uk} \\
&= \sum_{k \neq k_u^*} \Delta_{uk} \mathbb{E}\left[T_{uk}(t+1)\right] - \epsilon_u t.
\end{aligned}
$$

$\square$

We now use this lower bound along with the lower bound on the expected number of sub-optimal schedules for α-consistent policies (Corollary 20).

Proof of Theorem 8. To prove part (a) of the theorem, we use Lemma 22 and part (a) of corollary 20 as follows: For any $\gamma > \frac{1}{1-\alpha}$, there exist constants C_1 and τ such that for all $t \in [\max\{C_1 K^\gamma, \tau\}, (K-1)\frac{D(\boldsymbol{\mu})}{4\bar{\epsilon}}]$,

$$
\begin{aligned}
\frac{1}{U} \sum_{u \in [U]} \Psi_u(t) &\geq \frac{\Delta}{U} \sum_{u \in [U]} \left(\sum_{k \neq k_u^*} \mathbb{E}\left[T_{uk}(t+1)\right] - \epsilon_u t\right) \\
&\geq (K-1)\frac{D(\boldsymbol{\mu})}{2}\left((1-\alpha)\log t - \log(KC_1)\right) - \bar{\epsilon} t \\
&\geq (K-1)\frac{D(\boldsymbol{\mu})}{2}\frac{\log t}{\log \log t} - \bar{\epsilon} t \\
&\geq (K-1)\frac{D(\boldsymbol{\mu})}{4}\frac{\log t}{\log \log t},
\end{aligned}
$$

where the last two inequalities follow since $t \geq C_1 K^\gamma$ and $t \leq (K-1)\frac{D(\boldsymbol{\mu})}{4\bar{\epsilon}}$.

Part (b) of the theorem can be similarly shown using parts (b) and (c) of Corollary 20. $\square$

Endnotes

1. An earlier version of this work appeared in the Proceedings of the Thirtieth Annual Conference on Neural Information Processing Systems (NIPS), 2016 (Krishnasamy et al. (2016)).

2. The exploration rate could scale like $\log t/t$ if we knew Δ in advance; however, without this knowledge, additional exploration is needed.

3. Note that the interval in Proposition 7 is non-empty for ϵ satisfying the condition of the proposition; the same is analogously true for ϵ_u satisfying the conditions in Theorem 8.

Acknowledgments

References

Agrawal S, Goyal N (2012) Analysis of Thompson sampling for the multi-armed bandit problem. *Proceedings of the 25th Annual Conference on Learning Theory (COLT)*.

Audibert JY, Munos R, Szepesvári C (2009) Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science* 410(19):1876–1902.

Auer P, Cesa-Bianchi N, Fischer P (2002) Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47(2-3):235–256.

Avrachenkov K, Ayesta U, Doncel J, Jacko P (2013) Congestion control of TCP flows in internet routers by means of index policy. *Computer Networks* 57(17):3463–3478.

Ayesta U, Jacko P, Novak V (2017) Scheduling of multi-class multi-server queueing systems with abandonments. *J. Scheduling* 20(2):129–145.

Bubeck S, Cesa-Bianchi N (2012) Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Machine Learning* 5(1):1–122.

Bubeck S, Perchet V, Rigollet P (2013) Bounded regret in stochastic multi-armed bandits. *arXiv preprint arXiv:1302.1611* .

Buyukkoc C, Varaiya P, Walrand J (1985) The $c\mu$ rule revisited. *Advances in applied probability* 17(1):237–238.

Cesa-Bianchi N, Lugosi G (2012) Combinatorial bandits. *Journal of Computer and System Sciences* 78(5):1404–1422.

Combes R, Jiang C, Srikant R (2015a) Bandits with budgets: Regret lower bounds and optimal algorithms. *Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, 245–257 (ACM).

Combes R, Shahi MSTM, Proutiere A, et al. (2015b) Combinatorial bandits revisited. *Advances in Neural Information Processing Systems*, 2116–2124.

Cox D, Smith W (1961) Queues. *Wiley* .

Degenne R, Perchet V (2016) Combinatorial semi-bandit with known covariance. *Advances in Neural Information Processing Systems*, 2972–2980.

Gai Y, Krishnamachari B, Jain R (2012) Combinatorial network optimization with unknown variables: Multi-armed bandits with linear rewards and individual observations. *IEEE/ACM Transactions on Networking (TON)* 20(5):1466–1478.

Garivier A, Cappé O (2011) The KL-UCB algorithm for bounded stochastic bandits and beyond. *arXiv preprint arXiv:1102.2490* .

Gittins J, Glazebrook K, Weber R (2011) *Multi-Armed Bandit Allocation Indices* (Wiley-Blackwell), ISBN 9780470670026.

Gittins JC (1979) Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)* 148–177.

Jacko P (2010) Restless bandits approach to the job scheduling problem and its extensions. *Modern trends in controlled stochastic processes: theory and applications* 248–267.

Jaksch T, Ortner R, Auer P (2010) Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research* 11(Apr):1563–1600.

Kaufmann E, Korda N, Munos R (2012) Thompson sampling: An asymptotically optimal finite-time analysis. *Algorithmic Learning Theory*, 199–213 (Springer).

Krishnasamy S, Sen R, Johari R, Shakkottai S (2016) Regret of queueing bandits. *Proceedings of the Thirtieth Annual Conference on Neural Information Processing Systems (NIPS)* (Barcelona Spain).

Kushner H (2013) *Heavy traffic analysis of controlled queueing and communication networks*, volume 47 (Springer Science & Business Media).

Lai TL, Robbins H (1985) Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics* 6(1):4–22.

Larrañaga M, Ayesta U, Verloop IM (2016) Dynamic control of birth-and-death restless bandits: Application to resource-allocation problems. *IEEE/ACM Trans. Netw.* 24(6):3812–3825.

Lott C, Teneketzis D (2000) On the optimality of an index rule in multichannel allocation for single-hop mobile networks with multiple service classes. *Probability in the Engineering and Informational Sciences* 14:259–297.

Mahajan A, Teneketzis D (2008) Multi-armed bandit problems. *Foundations and Applications of Sensor Management*, 121–151 (Springer).

Neely MJ (2010) Stability and capacity regions or discrete time queueing networks. *arXiv preprint arXiv:1003.3396* .

Niño-Mora J (2006) Marginal productivity index policies for scheduling a multiclass delay-/loss-sensitive queue. *Queueing Systems* 54(4):281–312.

Niño-Mora J (2007) Dynamic priority allocation via restless bandit marginal productivity indices. *Top* 15(2):161–198.

Niño-Mora J (2012) Admission and routing of soft real-time jobs to multiclusters: Design and comparison of index policies. *Computers & Operations Research* 39(12):3431–3444.

Ortner R, Ryabko D, Auer P, Munos R (2014) Regret bounds for restless markov bandits. *Theoretical Computer Science* 558:62–76.

Perchet V, Rigollet P, Chassang S, Snowberg E (2015) Batched bandit problems. *arXiv preprint arXiv:1505.00369* .

Salomon A, Audiber JY, El Alaoui I (2013) Lower bounds and selectivity of weak-consistent policies in stochastic multi-armed bandit problem. *The Journal of Machine Learning Research* 14(1):187–207.

Srikant R, Ying L (2014) *Communication Networks: An Optimization, Control and Stochastic Networks Perspective* (Cambridge University Press).

Thompson WR (1933) On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 285–294.

Tsybakov AB (2008) *Introduction to nonparametric estimation* (Springer Science & Business Media).

Van Mieghem JA (1995) Dynamic scheduling with convex delay costs: The generalized $c\mu$ rule. *The Annals of Applied Probability* 809–833.

Whitt W (1974) Heavy traffic limit theorems for queues: a survey. *Mathematical Methods in Queueing Theory*, 307–350 (Springer).

Whittle P (1988) Restless bandits: Activity allocation in a changing world. *Journal of applied probability* 25(A):287–298.