

Learning to Schedule in Non-Stationary Wireless Networks With Unknown Statistics

Quang Minh Nguyen

nmquang@mit.edu

Massachusetts Institute of Technology
USA

Eytan Modiano

modiano@mit.edu

Massachusetts Institute of Technology
USA

ABSTRACT

The emergence of large-scale wireless networks with partially-observable and time-varying dynamics has imposed new challenges on the design of optimal control policies. This paper studies efficient scheduling algorithms for wireless networks subject to generalized interference constraint, where mean arrival and mean service rates are unknown and non-stationary. This model exemplifies realistic edge devices' characteristics of wireless communication in modern networks. We propose a novel algorithm termed MW-UCB for generalized wireless network scheduling, which is based on the Max-Weight policy and leverages the Sliding-Window Upper-Confidence Bound to learn the channels' statistics under non-stationarity. MW-UCB is provably throughput-optimal under mild assumptions on the variability of mean service rates. Specifically, as long as the total variation in mean service rates over any time period grows sub-linearly in time, we show that MW-UCB can achieve the stability region arbitrarily close to the stability region of the class of policies with full knowledge of the channel statistics. Extensive simulations validate our theoretical results and demonstrate the favorable performance of MW-UCB.

CCS CONCEPTS

• **Networks** → **Network performance modeling; Network performance analysis; Mobile ad hoc networks; Network control algorithms.**

KEYWORDS

Optimal Control, Scheduling, Wireless Network, Machine Learning, Partial Observability, Non-stationarity

ACM Reference Format:

Quang Minh Nguyen and Eytan Modiano. 2023. Learning to Schedule in Non-Stationary Wireless Networks With Unknown Statistics. In *The Twenty-fourth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (MobiHoc '23)*, October 23–26, 2023, Washington, DC, USA. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3565287.3610258>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

MobiHoc '23, October 23–26, 2023, Washington, DC, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9926-5/23/10.

<https://doi.org/10.1145/3565287.3610258>

1 INTRODUCTION

Wireless networks are increasingly large-scale and complex in response to the surge in edge-based Internet of Things (IoT) architecture [7], mobile communication [13] and wireless paradigm [23]. One fundamental challenge in the transition to large-scale networks is that minor inefficiencies can accumulate and severely limit performance [17]. Consequently, the advance of modern infrastructures toward massive scale has led to the re-design of operational management for various network tasks, such as traffic engineering [32], load-balancing [28], utility maximization [11], and link scheduling [2, 26]. In this work, we focus on designing scheduling algorithms that are theoretically efficient and meet the stringent requirements of emerging large-scale wireless networks.

Efficient scheduling of transmissions is essential for wireless devices to share the common spectrum while achieving high throughput. Despite its established throughput-optimality for a variety of classical stochastic network models, the celebrated Max-Weight scheduling policy [22, 27] requires the full knowledge of the channel statistics, which are often unknown a priori [26, 31] and thus hinder its direct adoption. First, due to the delay incurred by the accumulation of global network state information in emerging large-scale systems and multi-path fading, the instantaneous service capacities of wireless links and the packet arrivals to nodes are usually unavailable at the time of making scheduling decisions and can only be observed from channel feedback. We refer to this peculiar characteristic of large-scale networked systems as *partial observability*. Second, the mobility of edge devices [5, 24] and unreliable nature of wireless communication [25] impose *non-stationary dynamics*, whereby both the mean packet arrivals and mean service rates may vary over time, and are *unknown* in advance to the network operator. When the channel is not instantaneously observable, it is well-known that an optimal policy is to leverage the mean service rates in making Max-Weight scheduling decisions [22, 27]; however, in our setting, those statistics are *unknown*, *non-stationary* and must be learned. In this paper, we aim to develop throughput-optimal scheduling algorithms under the requirements of *partial observability*, *non-stationary dynamics* and *unknown statistics*.

A main challenge in the design of non-stationary network control algorithm under partially-observed and unknown statistics is that the analytical characterization of the capacity region for stationary network setting [22] no longer holds under non-stationarity due to the potential non-existence of steady state or well-defined long-term averages [3]. Previous works either consider simplified models [21, 26], or only achieves a constrained stability region for bipartite queueing system [31]. In particular, under *partial observability* and *unknown statistics*, [26] designed a throughput-optimal joint learning and scheduling policy for stationary network control. While establishing the effectiveness of the Max-Weight policy even

for non-stationary network control, [21] assumes the availability of instantaneous nodes' packet arrivals and links' service capacities to the controller for making decision. Closest to our work is [31], which proposes a stabilizing algorithm for bipartite queueing system that supports arrival rates within a stability region constrained by window-based (non-stationary) dynamics.

In this paper, we propose a new notion of stability for non-stationary network control, and a novel joint learning and scheduling algorithm that achieves a stability region arbitrarily close to the true stability region. Our contributions can be summarized as follows:

- We present a new class of approximate stability regions that is parameterized by a quantity capturing the closeness to the true stability region. Based on this notion of approximate stability region, we propose a new notion of throughput-optimality for non-stationary network control and, as a special case, prove its equivalence to the conventional notion of stability in the simplified setting of stationary network.
- We propose Max-Weight scheduling augmented by Sliding-Window Upper-Confidence Bound, hence termed MW-UCB, as a novel algorithm for non-stationary network control, subject to generalized wireless interference constraints, with partial observability and unknown statistics. Under mild assumptions on the system learnability, we establish the throughput-optimality of MW-UCB and its strong stability within the window-based region previously considered in the literature [31].
- We empirically validate our theoretical results and demonstrate that MW-UCB achieves the same stability region as that of the idealized Max-Weight policy with full knowledge of network statistics.

The rest of the paper is organized as follows. We present our system model and problem formulation in Section 2. In Section 3, we present our new notion of throughput-optimality for non-stationary networks. In Section 4, we propose the throughput-optimal MW-UCB algorithm and establish its stability results. We conduct numerical simulations to empirically validate the throughput-optimality of MW-UCB and demonstrate its favorable performance in Section 5, and conclude the paper in Section 6.

2 PRELIMINARIES AND PROBLEM FORMULATION

2.1 Network Model

A wireless network with arbitrary topology is represented by a directed graph $\mathcal{G}(V, E)$, where $V = \{1, 2, \dots, n\}$ is the set of nodes and $E = \{(i, j) : i, j \in V\}$ is the set of directed point-to-point links. Time is slotted. For simplicity of technical exposition, we consider single-hop traffic¹. For any $e = (i, j) \in E$, we denote by $a_e(t)$ the number of packets arriving at node i at time slot t to be transmitted to neighbouring node j . We consider $\{a_e(t)\}_{t \geq 0}$ to be independent with potentially time-varying means $\lambda_e(t) = \mathbb{E}[a_e(t)]$, and are bounded by a finite number, i.e. $a_e(t) \leq A_{\max}$ for all $e \in E$ and t .

¹The results of the paper naturally generalize to multi-hop setting by incorporating the Back-pressure mechanism [27].




Figure 1: Sequence of events in one time slot.

We assume a general wireless interference model. Denote by $\mathcal{M}$ the set of all admissible link activations and, at time slot t , by $\mathbf{x}(t) = \{x_e(t)\}_{e \in E} \in \mathcal{M}$ the scheduling decision of whether to activate link $e \in E$:

$$x_e(t) = \begin{cases} 1, & \text{if } e \text{ is activated at time } t \\ 0, & \text{if } e \text{ is not activated at time } t \end{cases}$$

We impose no structural restriction on the set $\mathcal{M}$, thereby capturing a wide range of practical wireless models including primary interference [19], k-hop interference [18], and protocol interference [9]. Let $\theta_e(t)$ be the service capacity of link e at time slot t , which is bounded by a finite number, i.e. $\theta_e(t) \leq \mu_{\max}$. For any link $e \in E$, we assume that $\{\theta_e(t)\}_{t \geq 0}$ are independent and that the mean service rate $\mu_e(t) = \mathbb{E}[\theta_e(t)]$ may vary over time. Additionally, we require the mean service rate to be lower bounded by a strictly positive constant, i.e. $\mu_e(t) \geq \mu_{\min} > 0$; this assumption is also often imposed by the literature on optimal control of queueing systems with time-varying statistics [14, 31]. The effective service rate of link e at time slot t is then given by:

$$b_e(t) = x_e(t)\theta_e(t), \quad (1)$$

which characterizes the achievable data rate of the link.

Let $Q_e(t)$ be the physical queue of backlogged packets at link $e \in E$ that are waiting to be transmitted at the end of time slot t . Since any link e receives $a_e(t)$ packet arrivals and can serve at most $b_e(t)$ packets during a time slot, the queueing dynamics evolves as:

$$Q_e(t+1) = (Q_e(t) + a_e(t) - b_e(t))^+, \quad \forall e \in E, \quad (2)$$

where $[x]^+ = \max\{x, 0\}$.

In order to capture the realistic characteristics of modern wireless network, we incorporate the following requirements in our model:

- *Partial Observability:* For every link $e \in E$, both the instantaneous packet arrivals $a_e(t)$ and link's service capacity $\theta_e(t)$ are not available at the start of the time slot t and thus cannot be used for making the scheduling decisions. At the end of time slot t , however, the nodes can accumulate statistics of the past time slot to obtain the packet arrivals $a_e(t)$'s and the service capacities of the activated links, i.e. those $\theta_e(t)$'s such that $x_e(t) = 1$. For unactivated links e where $x_e(t) = 0$, though the information of $\theta_e(t)$ is not revealed, the effective service rate is $b_e(t) = 0$. Thus, given the knowledge of $a_e(t)$ and $b_e(t)$, the queueing dynamics (2) for the next time slot $t+1$ can always be evaluated at the end of time slot t . The sequence of events within time slot t is depicted in Figure 1.

Table 1: Asymptotic Relationships

$f(\mathbf{x}) = O(g(\mathbf{x}))$	there exists constants M and C such that $ f(\mathbf{x}) \leq C g(\mathbf{x}) $ for all $\mathbf{x}$ with $\ \mathbf{x}\ _\infty > M$.
$f(\mathbf{x}) = \Omega(g(\mathbf{x}))$	there exists constants M and C such that $ f(\mathbf{x}) \geq C g(\mathbf{x}) $ for all $\mathbf{x}$ with $\ \mathbf{x}\ _\infty > M$.
$f(\mathbf{x}) = \Theta(g(\mathbf{x}))$	$f(\mathbf{x}) = O(g(\mathbf{x}))$ and $f(\mathbf{x}) = \Omega(g(\mathbf{x}))$.
$f(\mathbf{x}) = o(g(\mathbf{x}))$	for every $\varepsilon > 0$, there exists constant M such that $ f(\mathbf{x}) \leq \varepsilon g(\mathbf{x}) $ for all $\mathbf{x}$ with $\ \mathbf{x}\ _\infty > M$. In this case, we alternatively say $f(\mathbf{x})$ is sub-linear in $g(\mathbf{x})$.

- *Non-Stationary Dynamics:* We assume that both the mean packet arrivals $\lambda(t) = (\lambda_e(t))_{e \in E}$ and service rates $\mu(t) = (\mu_e(t))_{e \in E}$ vary over time, i.e. non-stationary.
- *Unknown Statistics:* All the statistics $\{\lambda(t)\}_{t \geq 0}$ and $\{\mu(t)\}_{t \geq 0}$ are unknown to the scheduler for making control decisions.

2.2 Asymptotic Relationships and Notations

Let $\mathbf{Q}(t) = (Q_e(t))_{e \in E}$, $\mathbf{a}(t) = (a_e(t))_{e \in E}$, $\boldsymbol{\theta}(t) = (\theta_e(t))_{e \in E}$ and $\mathbf{b}(t) = (b_e(t))_{e \in E}$ be respectively the vector of queue lengths, packet arrivals, service capacities and effective service rates. For any two real numbers x and y , we let $x \vee y = \max\{x, y\}$ and $x \wedge y = \min\{x, y\}$. For any vector $\mathbf{x} = (x_i)$ of real numbers and $p \in [1, \infty)$, we denote $\|\mathbf{x}\|_p$ as its ℓ_p -norm. For the two cases of $p \in \{1, \infty\}$ used in this paper, we have $\|\mathbf{x}\|_1 = \sum_i |x_i|$ and $\|\mathbf{x}\|_\infty = \max_i \{|x_i|\}$. For two positive multivariate functions $f(\mathbf{x})$ and $g(\mathbf{x})$, their asymptotic relationships [12] are given in Table 1.

2.3 Policy Space and Problem Statement

For any variable affected by the control of the scheduling decisions, we add the superscript π to acknowledge that it is under the action of the policy π . An *admissible* policy π at every time slot t generates a scheduling decision $\mathbf{x}^\pi(t) \in \mathcal{M}$ using only the knowledge of the past packet arrivals $\mathbf{a}(0), \mathbf{a}(1), \dots, \mathbf{a}(t-1)$, the past effective service rates $\mathbf{b}(0), \mathbf{b}(1), \dots, \mathbf{b}(t-1)$, and the past decisions $\mathbf{x}(0), \mathbf{x}(1), \dots, \mathbf{x}(t-1)$ up to time $t-1$. Additionally, we consider *idealized* policies, the definition of which is similar to that of an admissible policy except that at time slot t , it also has the full knowledge of the network statistics $\lambda(t)$ and $\mu(t)$ and can use them in making the scheduling decisions. The set of all admissible policies and the set of all idealized policies are respectively denoted by Π and Π_s . Under the simplified model whereby the network dynamics are stationary, [26] designed a joint learning and scheduling algorithm in Π that supports the same stability region, i.e. the set of arrival rates under which the system is stabilizable, as that of idealized policies in Π_s . Nevertheless, generalization to the case of non-stationary network dynamics is non-trivial due to the analytical intractability of the capacity region. Moreover, the only previous work [31] that attempts to learn non-stationary network dynamics could achieve only a reduced stability region that is constrained by the window-based dynamics, thereby being sub-optimal.

In this paper, we aim to develop a control scheme for the class of policies in Π that maximizes the stability region of the network under our considered setting.

3 NOTION OF STABILITY FOR NON-STATIONARY NETWORK WITH UNKNOWN STATISTICS

One main challenge of non-stationary network control is that the analytical characterization of the capacity region for the case of stationary network may no longer hold under the non-stationarity. In this Section, we propose a new notion of throughput-optimality that is more suitable to the non-stationary setting. For the simplified case of stationary network, we further establish the equivalence between our new notion and the conventional notion of throughput-optimality.

3.1 Assumption on Non-Stationary Dynamics

For any $t_1 < t_2$, we denote the total variation of the mean service rate by:

$$\gamma(t_1, t_2) = \sum_{t=t_1+1}^{t_2} \|\mu(t) - \mu(t-1)\|_\infty, \quad (3)$$

and stipulate the following mild assumption on the non-stationarity of the mean service rates.

Assumption 1. For any $t_1 < t_2$, the total variation is upper-bounded by $\gamma(t_1, t_2) = O(|t_2 - t_1|^\alpha)$ for some $\alpha \in [0, 1)$.

Our assumption only requires the total variation of mean service rates over any time period to grow sub-linearly in time, thereby ensuring that the network dynamics do not vary too aggressively. Similar assumptions have been extensively used the literature of learning in non-stationary environments [6, 10, 30].

3.2 Performance Metrics

Before characterizing the stability regions of interest, we first define the Q_T measure that captures the growth of queue size in expectation, and present the equivalent definition of mean rate stability [22] under the Q_T measure.

Definition 1 (Q_T Measure). The total expected queue length at time T under a control of policy π is quantified by $Q_T^\pi = \mathbb{E}[\sum_{e \in E} Q_e^\pi(T)]$.

Definition 2 (Mean Rate Stability). A network is mean rate stable under a policy π if:

$$\lim_{T \rightarrow \infty} \frac{\mathbb{E}[\sum_{e \in E} Q_e^\pi(T)]}{T} = 0,$$

or equivalently $Q_T^\pi = o(T)$.

The notion of stability region of a policy π describes the set of arrival rate vectors such that mean rate stability could be achieved under π . The stability region Λ is the region that can be achieved by the class Π of admissible policies, as formally defined below.

Definition 3 (Stability Region). The stability region of the class of admissible policies is defined as:

$$\Lambda = \{\{\lambda(t)\}_{t \geq 0} : \exists \pi \in \Pi \text{ such that } Q_T^\pi = o(T)\}.$$

Similarly, we define the idealized stability region Λ_s of the class Π_s of idealized policies.

Definition 4 (Idealized Stability Region). *The stability region of the class of idealized policies is defined as:*

$$\Lambda_s = \{ \{ \lambda(t) \}_{t \geq 0} : \exists \pi \in \Pi_s \text{ such that } Q_T^\pi = o(T) \}.$$

Note that we always have $\Lambda \subseteq \Lambda_s$, since the idealized policies have the full network statistics as opposed to the admissible policies. For the simplified model of stationary dynamics, a special case of our setting where Assumption 1 trivially holds with $\alpha = 0$, [26] designed an admissible policy stabilizing any arrival rate vector in $\text{int}(\Lambda_s)^2$, i.e. the interior of the idealized stability region. The algorithmic development and analysis of [26] heavily rely on the fact that under the stationary dynamics whereby $\lambda = \lambda(t)$ and $\mu = \mu(t)$ for all $t \geq 0$, the idealized stability region Λ_s can be further characterized by the existence of a policy $\pi^* \in \Pi_s$ such that:

$$\lambda_e = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} a_e(t) \leq \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} b_e^{\pi^*}(t), \quad \forall e \in E. \quad (4)$$

However, under non-stationarity, the above limits may not even exist, thereby hindering the adoption of stability region's characterization as in the case of stationary dynamics. Such analytical intractability of the stability region is central to the problem of optimal control for non-stationary network [3, 21].

3.3 Notion of Throughput-Optimality for Non-Stationary Network

In this Section, we propose a novel notion of throughput-optimality for non-stationary networks. For the simplified setting of stationary network, the conventional notion of stability defines a policy π to be throughput-optimal if it can stabilize the system for any arrival rate $\lambda \in \text{int}(\Lambda_s)$, i.e. in the interior of the idealized stability region. This is equivalent to $\lambda \in (1 - \varepsilon)\Lambda_s$ for some $\varepsilon > 0$, which is then usually incorporated with the analytical characterization of Λ_s for establishing the stability of MaxWeight-type algorithms. However, in the context of non-stationary networks, such an approach may not be directly applicable due to the analytical intractability of the idealized stability region Λ_s . To this end, we first present our new definition of approximate stability region, which is central to our throughput-optimality notion and algorithmic development.

Definition 5 (Approximate Stability Region). *Given any $\beta \in [0, 1]$, we define the approximate region $\Lambda_s(\beta)$ as:*

$$\Lambda_s(\beta) = \{ \{ \lambda(t) \}_{t \geq 0} : \exists \pi \in \Pi_s \text{ such that } Q_T^\pi = O(T^\beta) \}$$

We now derive the key properties of $\Lambda_s(\beta)$ and its relation to the idealized stability region Λ_s in the following Lemma, whose proof is deferred to Appendix A.

Lemma 1. *The set $\Lambda_s(\beta)$ is expanding for increasing β , i.e. if $0 \leq \beta_1 \leq \beta_2 \leq 1$, then $\Lambda_s(\beta_1) \subseteq \Lambda_s(\beta_2)$. Moreover, for any $\beta \in [0, 1]$, we have $\Lambda_s(\beta) \subseteq \Lambda_s \subseteq \Lambda_s(1)$.*

Lemma 1 suggests that the region $\Lambda_s(\beta)$ grows arbitrarily close to Λ_s as β approaches 1. Moreover, leveraging this notion of approximate stability region, the next Theorem establishes the characterization of the true stability region Λ .

²[26] considers maximal matching for scheduling under matching constraint, and thus achieves the interior of $\frac{1}{2}\Lambda_s$ as the stability region. However, by replacing maximal matching with maximum matching, we can achieve the full stability region $\text{int}(\Lambda_s)$.

Theorem 1. *Under Assumption 1, we have the following characterization of the stability region:*

$$\Lambda_s(\beta) \subseteq \Lambda \subseteq \Lambda_s,$$

for any $\beta \in [0, 1)$.

PROOF. As the idealized policies have the full network statistics as opposed to the admissible policies, the idealized stability region Λ_s trivially subsumes the stability region Λ . The proof of $\Lambda_s(\beta) \subseteq \Lambda$ is based on our development of the admissible policy MW-UCB in Section 4 that, given any $\beta \in [0, 1)$, achieves mean rate stability for any set of arrival rates $\{ \lambda(t) \}_{t \geq 0} \in \Lambda_s(\beta)$ (Theorem 3). $\square$

Motivated by Theorem 1, we propose the following notion of throughput-optimality for non-stationary network control.

Definition 6 (Throughput-Optimality). *A policy $\pi \in \Pi$ is throughput optimal if given any $\beta \in [0, 1)$, the network under π is mean rate stable for any $\{ \lambda(t) \}_{t \geq 0} \in \Lambda_s(\beta)$.*

Under the above definition, we aim to develop an admissible policy in Π that is throughput-optimal for our considered setting of non-stationary and partially-observable network.

3.4 Connection to Traditional Notion of Throughput-Optimality for Stationary Networks

We further demonstrate the equivalence of our throughput-optimality notion to the usual notion in the case of stationary network whereby $\lambda = \lambda(t)$ and $\mu = \mu(t)$ for all $t \geq 0$. As discussed in Section 3.2, the idealized stability region Λ_s can be characterized by:

$$\bar{\Lambda}_s = \{ \lambda : \exists \pi^* \in \Pi_s \text{ such that } \pi^* \text{ satisfies (4)} \},$$

and an admissible policy is throughput-optimal (in the usual notion) if it can stabilize any arrival rate $\lambda \in \bar{\Lambda}_s$ ³. The next Theorem illustrates that our new notion of throughput-optimality implies the usual notion of throughput-optimality in the stationary network control problem.

Theorem 2. *Under the stationary network setting, if a policy $\pi \in \Pi$ is throughput-optimal according to Definition 6, then the network under π is mean rate stable for any $\lambda \in \bar{\Lambda}_s$.*

PROOF SKETCH. Given $\lambda \in \bar{\Lambda}_s$, we can show via Lyapunov drift analysis that the Max-Weight (MW) policy with full knowledge of the statistics achieves $Q_T^{\text{MW}} = O(T^{\frac{1}{2}})$. Consequently, this implies that $\lambda \in \Lambda_s(\frac{1}{2})$. Since by our Definition 6, a throughput-optimal policy π would support the stability region $\Lambda_s(\frac{1}{2})$, it thus guarantees mean rate stability for any $\lambda \in \bar{\Lambda}_s$. The full proof is given in Appendix B.

4 SCHEDULING WITH NON-STATIONARY AND UNKNOWN CHANNEL STATISTICS

In this Section, we present MW-UCB as a provably throughput-optimal policy for non-stationary network control. We provide the preliminaries of Upper-Confidence Bound (UCB) for learning

³Here, we use $\bar{\Lambda}_s$ to emphasize that this is a special case of Λ_s where such characterization based on (4) only holds for the stationary network setting.

uncertain channel statistics in Section 4.1 and the algorithmic development of MW-UCB in Section 4.2. The throughput-optimality and stability results of MW-UCB then follow in Section 4.3.

4.1 Upper-Confidence Bound (UCB) for Learning Uncertain Channel Statistics

Central to our problem is the learning of not only the unknown links' service rates but also the scheduling decisions that can maximize the overall network's throughput. We start by considering a simplistic problem setting in which the network dynamics are stationary and the objective is to attain the maximum possible total service capacity (in expectation) of the network, and show that a simple Upper-Confidence Bound (UCB) algorithm is close-to-optimal in this scenario. However, while having the potential for being the solution for network control under uncertain channel statistics, the UCB algorithm in its original form lacks the adaptivity to deal with non-stationary dynamics and is a pure learning scheme in nature, which is not designed to deal with sophisticated control tasks as in our original problem.

4.1.1 A Simplistic Problem Setting and Application of UCB. At any time slot t , the scheduling decision $\mathbf{x}(t)$ yields in expectation the service of $\mu_e(t)x_e(t)$ for link $e \in E$ and thus the total service of:

$$\sum_{e \in E} \mu_e(t)x_e(t). \quad (5)$$

Now, we turn into a simplified objective of maximizing the total service (5) over the time horizon T of the network and further assume stationary dynamics of the links' service rates, i.e. $\mu_e(t) = \mu_e, \forall e \in E, t \geq 0$. Under this setting, an idealized policy with full knowledge of the statistics $\boldsymbol{\mu} = (\mu_e)_{e \in E}$ would make the scheduling decision at every time slot t that maximizes (5), i.e.

$$\mathbf{x}(t) = \mathbf{x}^* = \operatorname{argmax}_{\mathbf{x} \in \mathcal{M}} \left\{ \sum_{e \in E} \mu_e x_e \right\}. \quad (6)$$

However, such statistics $\boldsymbol{\mu} = (\mu_e)_{e \in E}$ are unknown in practice and thus must be learned, under our requirement of partial observability, via samples of service capacities of links having been activated. This gives rise to the exploration/exploitation tradeoff, where the controller must simultaneously learn the channel statistics $\boldsymbol{\mu}$ and utilize the existing information of observed service capacities to schedule transmissions achieving high total throughput. In particular, the problem of solving (6) over the time horizon T can be characterized as combinatorial multi-armed bandit (CMAB) with linear reward in stationary environment, which can be addressed by the class of UCB algorithms [4, 15, 20]. We hereby consider the UCB algorithm in [20], which proceeds as follows. At any time slot t and for every edge $e \in E$, the UCB algorithm keeps track of $T_t(e)$ and $\hat{\theta}_e(t)$, which respectively correspond to the number of times link e has been activated and observed up to time t , and the empirical mean of all the observations of the service capacities, i.e. those $\theta_e(s)$ such that $x_e(s) = 1$ for $s \in [0, t)$. The UCB weights are computed according to:

$$U_e(t) = \hat{\theta}_e(t) + \sqrt{\frac{3 \log(t)}{2T_t(e)}}, \forall e \in E,$$

which are then used for constructing the scheduling decision as:

$$\mathbf{x}(t) = \operatorname{argmax}_{\mathbf{x} \in \mathcal{M}} \left\{ \sum_{e \in E} U_e(t)x_e \right\}. \quad (7)$$

The total difference in achievable total expected service capacity between the maximizing policy with the full knowledge of the statistics $\boldsymbol{\mu}$ that makes the decision $\mathbf{x}^*$ as in (6) and the UCB algorithm that makes the decision $\mathbf{x}(t)$ as in (7) is captured by the regret:

$$\mathcal{R}_0(T) = T \sum_{e \in E} \mu_e x_e^* - \sum_{t=0}^{T-1} \sum_{e \in E} \mu_e x_e(t). \quad (8)$$

This type of metrics is also used by [26] for characterizing the performance and exploration/exploitation tradeoff of their joint learning and scheduling algorithm. From [20, Theorem 5], we have $\mathcal{R}_0(T) = O(\log(T))$, which guarantees only logarithmic growth in total error if the UCB algorithm is applied. Moreover, this regret bound is asymptotically tight [20, Proposition 1].

4.1.2 Limitations of The Conventional UCB Algorithm. While attaining competitive performance in the simplistic problem setting, the conventional UCB algorithm lacks the generality to readily be extended to deal with our problem of interest. First, vanilla UCB is known to be inappropriate for handling non-stationary dynamics [16]. Second, the formulation (6) that permits the adoption of UCB as a direct solution does not take into account the control of the system under arbitrary arrival rates: for example, if a policy aims to attain the maximum possible total service and hence always makes the scheduling decision $\mathbf{x}^*$ as in (6), it would inevitably overload certain unactivated links, i.e. $e \in E$ with $x_e^* = 0$, to which there are packet arrivals over time. On the other hand, the Max-Weight policy [22, 27] that incorporates the queue lengths into making scheduling decision can adapt to the dynamics of arbitrary arrival rates. Consequently, the solution for scheduling in non-stationary wireless networks with partial observability and unknown statistics requires the interplay between learning and network control. In the next Section, we present our main algorithm that combines the Max-Weight policy with UCB to address these aforementioned challenges through its joint learning and scheduling scheme.

4.2 The MW-UCB Algorithm

We proceed to develop our scheduling algorithm, termed MW-UCB, based on a frame-based variant [26] of the Max-Weight policy [22, 27] and the augmentation of the sliding-window UCB [10] in the weight construction for adaptively learning the channels' statistics under non-stationarity. The full MW-UCB policy is depicted in Algorithm 1 with the convention that $0/0 = 0$.

For the class of idealized policies in Π_s , the Max-Weight policy [22] that at time slot t weights each edge e by $Q_e(t)\mu_e(t)$ and consequently schedule the link activation vector according to:

$$\mathbf{x}(t) = \operatorname{argmax}_{\mathbf{x} \in \mathcal{M}} \left\{ \sum_{e \in E} Q_e(t)\mu_e(t)x_e \right\} \quad (9)$$

is known to be throughput-optimal for the case of stationary network [22, 26] and obtains competitive performance on adversarial network control [21]. Nevertheless, under our considered model, the vector $\boldsymbol{\mu}(t)$ of mean service rates is unknown a priori, thereby hindering any direct adoption of the Max-Weight policy. The algorithmic design for joint network control and learning of the

Algorithm 1: Max-Weight with UCB (MW-UCB)

Input: graph $\mathcal{G}(V, E)$, restart period τ , window size $d \leq \tau$

```

1 for  $t = 1, \dots, T$  do
2   if  $t = \tau_j \in \mathcal{T} = \{\tau_0, \tau_1, \dots, \tau_K\}$  then
3     Initialize  $\phi_e(\tau_j) = 0, N_e(\tau_j) = 0, \hat{\mu}_e(t) = 0, \forall e \in E$ 
4     Reset the weights:  $w_e(\tau_j) = \frac{Q_e(\tau_j)}{\|\mathbf{Q}(\tau_j)\|_\infty}, \forall e \in E$ 
5   end
6   if  $t \in (\tau_j, \tau_{j+1})$  then
7     for  $e \in E$  do
8        $\phi_e(t) = \phi_e(t-1) - \mathbb{1}_{\{t \geq \tau_j + d\}} b_e(t-d) + b_e(t-1)$ 
9        $N_e(t) = N_e(t-1) - \mathbb{1}_{\{t \geq \tau_j + d\}} x_e(t-d) + x_e(t-1)$ 
10       $\hat{\mu}_e(t) = \frac{\phi_e(t)}{N_e(t)}$   $\triangleright$  empirical estimate of  $\mu_e(t)$ 
11       $w_e(t) = w_e(\tau_j)$   $\triangleright$  fixing queue lengths
12    end
13  end
14   $\rho_e(t) = \sqrt{\frac{3 \log(\tau)}{2N_e(t)}}$  (or  $\infty$  if  $N_e(t) = 0$ ),  $\forall e \in E$ 
15   $\bar{W}_e(t) = \min\{w_e(t)\hat{\mu}_e(t) + \rho_e(t), 1\}, \forall e \in E$   $\triangleright$ 
    computing UCB weights
16  [Scheduling] Activate the link activation vector:
     $\mathbf{x}(t) = \operatorname{argmax}_{\mathbf{x} \in \mathcal{M}} \{\sum_{e \in E} \bar{W}_e(t)x_e\}$ 
17  Observe  $\mathbf{a}(t)$  and  $\mathbf{b}(t)$ , and update the queues:
     $Q_e(t+1) = (Q_e(t) + a_e(t) - b_e(t))^+, \forall e \in E,$ 
18 end

```

weights $\{Q_e(t)\mu_e(t)\}_{e \in E}$ faces two challenges. First, $Q_e(t)\mu_e(t)$ is time-varying due to the dynamics of the queue length $Q_e(t)$ and the non-stationarity of $\mu_e(t)$. Second, the evolution of the weight $Q_e(t)\mu_e(t)$ is coupled with the scheduling decision $\mathbf{x}(t)$ due to its interdependence with effective service rate via (1) and thus the queueing dynamics (2). To address these challenges, we periodically freeze the queue length information in the weight instantiation, which helps to alleviate a source of non-stationarity and decouples the weight evolution from the scheduling decision. Specifically, our method partitions the time horizon T into frames of size τ , where the j^{th} frame begins at time slot $\tau_j = j\tau$, called restart point. We allow the last frame to have size potentially less than τ and let $\mathcal{T} = \{\tau_0, \tau_1, \dots, \tau_K\}$ be the set of all restart points, i.e. K is the largest number such that $\tau_K < T$. Then for any $t \in [\tau_j, \tau_{j+1})$, we use the normalized queue backlogs at the restart point τ_j as the unified weights (Line 4 and Line 11 of Algorithm 1):

$$w_e(\tau_j) = \frac{Q_e(\tau_j)}{\|\mathbf{Q}(\tau_j)\|_\infty}, \forall e \in E, \quad (10)$$

and aim to solve the following "relaxed" problem of (9) with simplified time-varying weight structure:

$$\mathbf{x}(t) = \operatorname{argmax}_{\mathbf{x} \in \mathcal{M}} \left\{ \sum_{e \in E} Q_e(\tau_j) \mu_e(t) x_e \right\} \quad (11)$$

$$\stackrel{(10)}{=} \operatorname{argmax}_{\mathbf{x} \in \mathcal{M}} \left\{ \sum_{e \in E} w_e(\tau_j) \mu_e(t) x_e \right\}. \quad (12)$$

Note that the objective in (11) is the "approximation" of the objective in (9) with error growing linearly in τ . Moreover, the problem

of solving (12) over τ time slots from τ_j to $\tau_{j+1} - 1$ can be characterized as stochastic combinatorial multi-armed bandit (SCMAB) problem in non-stationary environment, whereby the mean reward of each arm, i.e. $w_e(\tau_j)\mu_e(t)$, varies over time and is independent of the action, i.e. scheduling decision. To this end, we adopt the combinatorial UCB with sliding window (CUCB-SW) algorithm [10] for dealing with SCMAB under non-stationarity. Specifically, CUCB-SW is restarted at the beginning of each frame with the newly updated queue weights for the joint learning of the mean service rates $\mu(t)$'s and control of the system. Given the sliding window of size d as a hyper-parameter to be chosen later, the algorithm computes the estimate $\hat{\mu}_e(t)$ of the true mean service rate $\mu_e(t)$ as the local empirical average of the observed service capacities in the last d time slots. Formally, for $t \in [\tau_j, \tau_{j+1})$, i.e. within the j^{th} frame, and any $e \in E$, the following quantities:

$$\phi_e(t) = \sum_{s=\tau_j \vee (t-d)}^{t-1} \mathbb{1}_{\{x_e(s)=1\}} \theta_e(s) = \sum_{s=\tau_j \vee (t-d)}^{t-1} b_e(s), \quad (13)$$

$$N_e(t) = \sum_{s=\tau_j \vee (t-d)}^{t-1} \mathbb{1}_{\{x_e(s)=1\}} = \sum_{s=\tau_j \vee (t-d)}^{t-1} x_e(s), \quad (14)$$

respectively denote the total *observed* service capacities of link e and the number of times it had been *activated* up to time $t-1$ over the last d time slots. Line 8 and Line 9 of Algorithm 1 equivalently rewrite (13) and (14) in recursive forms for actual iterative updates in the algorithm. Then the local empirical average can be computed accordingly via $\hat{\mu}_e(t) = \phi_e(t)/N_e(t)$ as in Line 10. Finally, the UCB weights are computed by (Line 15):

$$\bar{W}_e(t) = \min \left\{ w_e(\tau_j) \hat{\mu}_e(t) + \sqrt{\frac{3 \log(\tau)}{2N_e(t)}}, 1 \right\}, \forall e \in E,$$

to be used for constructing of the scheduling decision (Line 16) as:

$$\mathbf{x}(t) = \operatorname{argmax}_{\mathbf{x} \in \mathcal{M}} \left\{ \sum_{e \in E} \bar{W}_e(t) x_e \right\}. \quad (15)$$

In order to capture the loss due to learning when CUCB-SW is applied to solve (12) within the j^{th} frame, we consider the regret:

$$\begin{aligned} \mathcal{R}(\tau_j) &= \sum_{t=\tau_j}^{\tau_{j+1}-1} \max_{\mathbf{x} \in \mathcal{M}} \sum_{e \in E} w_e(\tau_j) \mu_e(t) x_e \\ &\quad - \sum_{t=\tau_j}^{\tau_{j+1}-1} \sum_{e \in E} w_e(\tau_j) \mu_e(t) x_e(t), \end{aligned} \quad (16)$$

which characterizes the gap in objective between the maximizing policy with the full knowledge of the statistics $\mu(t)$ for $t \in [\tau_j, \tau_{j+1})$ that solves (12) and the considered policy that solves (15). The following Lemma 2, whose proof is given in Appendix D.3, provides an upper bound for $\mathcal{R}(\tau_j)$.

Lemma 2. Under MW-UCB, the regret $\mathcal{R}(\tau_j)$ can be bounded by:

$$\begin{aligned} \mathcal{R}(\tau_j) &\leq |E| \left(\frac{\tau}{d} + 1 \right) \left(2\sqrt{6\log(\tau)} + 48\sqrt{d}\log(\tau) \right) \\ &\quad + 4|E|d \cdot \gamma(\tau_j, \tau_{j+1}) + |E| \frac{\tau}{\sqrt{d}} + \frac{\pi^2}{3} |E|^2 \mu_{\max} \\ &\quad + \frac{\pi^2}{6} |E|^2 \mu_{\max} \log(2d^{1/2}). \end{aligned}$$

Under Assumption 1, which gives $\gamma(\tau_j, \tau_{j+1}) = O(\tau^\alpha)$, and by setting $d = \Theta(\tau^{\frac{2}{3}(1-\alpha)})$, we further have $\mathcal{R}(\tau_j) = O(\log(\tau)\tau^{\frac{1}{3}(\alpha+2)})$.

The guarantee in Lemma 2 demonstrates that the average loss due to learning over τ time slots of the j^{th} frame vanishes as τ grows in the sense that $\frac{\mathcal{R}(\tau_j)}{\tau} \rightarrow 0$ as $\tau \rightarrow \infty$ under our mild assumption on the learnability of the system. This is crucial for establishing the throughput-optimality of MW-UCB in the next Section.

4.3 Throughput-Optimality and Stability Results

In this Section, we prove the throughput-optimality of MW-UCB and, as a byproduct, its strong stability in a region constrained by window-based dynamics. The key components of the proof leverage the regret bound for learning non-stationarity (Lemma 2) in the analysis of the frame-based Lyapunov drift and non-trivially generalize the shedding technique from the adversarial network control literature [21]. In particular, to address the analytical intractability of the capacity region, we shed the traffic of the original system to obtain a new imaginary system whose traffic is within the window-based region, which is formally described in Definition 7. While the shedding process incurs additional term in the queue bound as a tradeoff, it makes the imaginary network dynamics tractable, from which stability of MW-UCB can be derived.

Definition 7 (Window-Based Region). The window-based region $C(W, \varepsilon)$ of the class of idealized policies is defined as:

$$C(W, \varepsilon) = \{ \{ \lambda(t) \}_{t \geq 0} : \exists \pi \in \Pi_s \text{ such that for } q = 0, W, 2W, \dots \\ \sum_{t=q}^{q+W-1} \lambda_e(t) \leq (1-\varepsilon) \sum_{t=q}^{q+W-1} \mathbb{E}[b_e^\pi(t)], \forall e \in E \}.$$

Specifically, the sequence of arrival rates $\{ \lambda(t) \}_{t \geq 0}$ satisfies the window-based region, parameterized by window size W and a shrinkage term ε , if there exists an idealized policy such that the total mean arrivals are less than a fraction $1 - \varepsilon$ of the total mean services over a window of W slots starting at every starting point $t = 0, W, 2W, \dots$. Next, we proceed to establish the main Theorem on the throughput-optimality of MW-UCB.

Theorem 3. Under Assumption 1, MW-UCB is throughput-optimal.

PROOF. Given any $\beta \in [0, 1)$, we show that MW-UCB can achieve mean rate stability for any $\{ \lambda(t) \}_{t \geq 0} \in \Lambda_s(\beta)$. We now consider an *imaginary* system that is obtained by imitating the same link service process $\{ \theta(t) \}_{t \geq 0}$ as the original system's and shedding a certain amount of traffic from the original system's arrivals $\{ a(t) \}_{t \geq 0}$ to obtain a new sequence of arrivals $\{ \tilde{a}(t) \}_{t \geq 0}$ with

$\tilde{\lambda}_e(t) = \mathbb{E}[\tilde{a}_e(t)], \forall e \in E, t \geq 0$. Denote the amount of shed traffic within the time horizon by:

$$X_T = \sum_{t=0}^{T-1} \sum_{e \in E} a_e(t) - \sum_{t=0}^{T-1} \sum_{e \in E} \tilde{a}_e(t). \quad (17)$$

For some $\varepsilon \in (0, 1)$ to be determined later, we consider the shedding scheme as in the following Lemma 3, whose proof is deferred to Appendix C.

Lemma 3. Given $\{ \lambda(t) \}_{t \geq 0} \in \Lambda_s(\beta)$ and any $\varepsilon \in (0, 1)$, there exists a shedding procedure such that $\{ \tilde{\lambda}(t) \}_{t \geq 0} \in C(\tau, \varepsilon)$ and:

$$\mathbb{E}[X_T] = O(T^{\beta+1}\tau^{-1} + \varepsilon T + \tau). \quad (18)$$

Intuitively, Lemma 3 suggests that, for analysis, despite the potential analytical intractability of the approximate region $\Lambda_s(\beta)$, we can shed traffic to obtain an imaginary system that is constrained in $C(\tau, \varepsilon)$ and thus more tractable with the tradeoff as characterized by (18). When MW-UCB is applied to the original system, it produces the sequence of decisions $\{ \mathbf{x}^{\text{MW-UCB}}(t) \}_{t \geq 0}$. Let $\tilde{Q}(t) = \{ \tilde{Q}_e(t) \}_{e \in E}$ be the virtual queue length vector at time slot t if such sequence of decisions $\{ \mathbf{x}^{\text{MW-UCB}}(t) \}_{t \geq 0}$ is applied to the imaginary system. Then the following Lemma 4 upper-bounds the Q_T measure of MW-UCB in the original system. The proof is given in Appendix D.1.

Lemma 4. We have the following bound:

$$Q_T^{\text{MW-UCB}} \leq \mathbb{E}[X_T] + \mathbb{E} \left[\sum_{e \in E} \tilde{Q}_e(T) \right]. \quad (19)$$

In particular, though the shedding process incurs the term $\mathbb{E}[X_T]$, which can be bounded as in Lemma 3, in the queue bound of the original system, we are now left with bounding the virtual queues $\tilde{Q}(t)$ which evolve over the imaginary system that is more tractable. Next, we further provide guarantee for the total expected queue length of the imaginary system, i.e. the second term of (19), as follows.

Lemma 5. Under Assumption 1 and given $\{ \tilde{\lambda}(t) \}_{t \geq 0} \in C(\tau, \varepsilon)$, there exists some universal constant $c_1 > 0$ such that if $c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)} \leq \varepsilon < 1$, MW-UCB with $d = \Theta(\tau^{\frac{2}{3}(1-\alpha)})$ being applied to the imaginary system satisfies:

$$\mathbb{E} \left[\sum_{e \in E} \tilde{Q}_e(T) \right] = O(T^{\frac{1}{2}} \tau^{\frac{1}{2}} + T^{1+\frac{\beta}{2}} \tau^{-\frac{1}{2}} + \varepsilon^{\frac{1}{2}} T). \quad (20)$$

The proof of Lemma 5 can be found in Appendix D.2. Notice that under our assumption $\alpha < 1$ (implying $\lim_{\tau \rightarrow \infty} \log(\tau) \tau^{\frac{1}{3}(\alpha-1)} = 0$), without loss of generality (WLOG), we can consider τ large enough so that $c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)} < 1$. The requirement that the considered "shrinkage" ε must be bounded away from 0 by such quantity reflects the loss due to learning. Note that the whole shedding process only serves for analytical purposes, i.e. we can shed the original system into the new imaginary system in the sense of Lemma 3 for any arbitrary $\varepsilon \in (0, 1)$. Finally, by setting $\varepsilon = c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)}$, we plug (18) and (20) into (19) to obtain the following bound for the Q_T measure of MW-UCB with $d = \Theta(\tau^{\frac{2}{3}(1-\alpha)})$:

$$Q_T^{\text{MW-UCB}} = O(T^{\beta+1} \tau^{-1} + T \sqrt{\log(\tau) \tau^{\frac{1}{3}(\alpha-1)}} + T^{1+\frac{\beta}{2}} \tau^{-\frac{1}{2}} + T^{\frac{1}{2}} \tau^{\frac{1}{2}}).$$

Therefore, MW-UCB with $\tau = T^{\frac{1}{2}(\beta+1)}$ and $d = \Theta(\tau^{\frac{2}{3}(1-\alpha)})$ can achieve:

$$\begin{aligned} Q_T^{\text{MW-UCB}} &= O(T^{\frac{1}{2}(\beta+1)} + \sqrt{\log(T)} T^{1-\frac{1}{12}(\beta+1)(1-\alpha)} + T^{\frac{1}{4}(\beta+3)}) \\ &= o(T), \end{aligned} \quad (21)$$

where the last line holds because $\alpha, \beta \in [0, 1]$. Since (21) asserts the mean rate stability of MW-UCB for $\{\lambda(t)\}_{t \geq 0} \in \Lambda_s(\beta)$ given any $\beta \in [0, 1]$, we conclude that MW-UCB is throughput-optimal. $\square$

Additionally, we derive the strong stability of MW-UCB for the window-based region in the following Corollary, whose proof is given in Appendix E.

Corollary 1. *Under Assumption 1, MW-UCB with a fixed window size τ and $d = \Theta(\tau^{\frac{2}{3}(1-\alpha)})$ achieves strong stability for any $\{\lambda(t)\}_{t \geq 0} \in C(\tau, \epsilon)$ with any $\epsilon > c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)}$ for some universal constant $c_1 > 0$, i.e.*

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \sum_{e \in E} \mathbb{E}[Q_e(t)] < \infty. \quad (22)$$

5 NUMERICAL SIMULATION

In this Section, we empirically evaluate the performance of MW-UCB and validate its throughput-optimality. We compare our proposed algorithm with two baseline algorithms:

- The idealized Max-Weight (MW) policy [22] with full knowledge of network statistics, which schedules the link activation vector at time slot t according to (9). This represents the class of idealized policies Π_s and serves as an *unrealistic* baseline.
- The MW with restart UCB [26], which can be thought of as a special case of MW-UCB for $d = \tau$ and represents the class of admissible policies Π . While originally proposed for stationary network control with partial observability and unknown statistics, this algorithm was empirically verified as a heuristic for non-stationary settings in [26], and is the only algorithm in the literature that is directly applicable to our model.




Figure 2: The 3×3 grid topology; an example of link schedule subject to node-exclusive interference, which forms a matching, is shown in solid lines.

For both MW-UCB and MW with restart UCB [26], we set the restart period to $\tau = T^{2/3}$. Sliding window size of MW-UCB is set to $d = 2\lceil \tau^{\frac{2}{3}(1-\alpha)} \rceil + 150$. We perform extensive testing on the 3×3 grid network with node-exclusive wireless interference constraints [8], as depicted in Figure 2. To model the non-stationary service rates, for any time slot t and link $e \in E$, we consider $\mu_e(t)$ evolving over time according to the Markov chain in Figure 3,

whereby $\mu_e(t)$ would change its state (between 0.25 and 0.75) with probability δ_t which itself may vary over time. Then, given $\mu_e(t)$, the instantaneous service capacity $\theta_e(t) \sim \text{Rayleigh}(\sqrt{\frac{2}{\pi}} \mu_e(t))$ is sampled from the Rayleigh distribution with the scale parameter that ensures $\mathbb{E}[\theta_e(t)] = \mu_e(t)$. We consider two settings of δ_t , which governs the dynamics of the non-stationary service rates:

- (1) Time-invariant $\delta_t = \frac{0.5}{T^{1/2}}$: This corresponds to the uniformly changing dynamics and was considered in the literature [26] for simulations.
- (2) Time-varying $\delta_t = \frac{0.5}{(t+1)^{1/2}}$: This non-stationary aperiodic setting captures more abruptly changing environments.




Figure 3: Markov chain for the time-varying mean service rate $\mu_e(t)$, $\forall e \in E$. At time slot t and for any edge e , given $\mu_e(t) \in \{0.25, 0.75\}$, the instantaneous service capacity $\theta_e(t)$ follows Rayleigh distribution with mean $\mu_e(t)$.

Moreover, both of the above settings satisfy Assumption 1 in the sense that $\mathbb{E}[\gamma(t_1, t_2)] = O(|t_2 - t_1|^{1/2})$ for any $0 \leq t_1 < t_2 \leq T$ ⁴ (see [1, Appendix F] for the proof). We thus use $\alpha = 1/2$ in our simulations.




5.1 Throughput-Optimality and Stability

We first consider fixed arrival rates where, at time slot t , every link receives Poisson arrivals with the same packet generation rate $\lambda = \lambda_e(t)$, $\forall e \in E$. To demonstrate the stability properties of the algorithms, we investigate the evolution of the total queue backlog, i.e. $\sum_{e \in E} Q_e(t)$ at time t , for $\lambda = 0.11$ and 0.12, which respectively represent the regimes of moderately loaded and highly loaded network. We run the simulations for $T = 10^6$ time slots and report the results for both settings of time-invariant δ_t and time-varying δ_t in Figure 4 and Figure 5, respectively.

Throughput-Optimality of MW-UCB: The results from Figure 4 and Figure 5 demonstrate that MW-UCB preserves the stability property of the idealized MW policy and thus supports the same stability region as achieved by the class of idealized policies with full statistics. In particular, the total queue backlogs of both algorithms remain stable for the moderate-load regime $\lambda = 0.11$, and start to explode in the high-load regime $\lambda = 0.12$.

Performance Evaluation of MW-UCB: In all experiments from Figure 4 and Figure 5, MW-UCB consistently outperforms MW with restart UCB. Whenever the arrival rate λ is inside the stability region (Figure 4a and Figure 5a), MW-UCB can learn the channels' statistics under non-stationarity and consequently stabilize the system. Additionally, for MW-UCB and idealized MW, we gradually increase $\lambda = 0.03 \rightarrow 0.22$ and report in Figure 6 the value of $\log(\sum_{e \in E} Q_e(T)/T)$ at $T = 1.5 \cdot 10^6$ to empirically measure the closeness of Q_T/T to 0 as well as its growth outside the stability region. The result suggests that MW-UCB preserves the pattern of idealized MW.

⁴While this is $\mathbb{E}[\gamma(t_1, t_2)]$, we can strictly enforce Assumption 1, i.e. without expectation, by deterministically simulating a feasible trajectory of $\mu(t)$'s evolution.

Figure 4: Total queue length over time for time-invariant $\delta_t = \frac{0.5}{T/2}$.Figure 5: Total queue length over time for time-varying $\delta_t = \frac{0.5}{(t+1)^{1/2}}$.Figure 6: Measuring $\log(Q_T/T)$ at $T = 1.5 \cdot 10^6$ for $\lambda = 0.03 \rightarrow 0.22$.

5.2 Time-Varying Arrival Rates

Next, we provide additional simulations for time-varying arrival rates $\lambda(t)$. We let all links in any time slot t to receive Poisson arrivals with the same packet generation rate $\bar{\lambda}(t) = \lambda_e(t)$, $\forall e \in E$. Given the mean service rates $\mu(t)$, an upper-bound on the maximum arrival rate supported by the network is given by:

$$\bar{\lambda}(t) \leq \min_{v \in V} \frac{1}{\sum_{e \in \mathcal{A}(v)} \mu_e(t)^{-1}}, \quad (23)$$

where $\mathcal{A}(v)$ is the set of links adjacent to node v . In our simulations, we set $\bar{\lambda}(t)$ to be exactly the right-hand side of (23) to model highly loaded network. We plot the total queue length over time, i.e. $\sum_{e \in E} Q_e(t)$ at time t , for both settings of time-invariant δ_t and time-varying δ_t respectively in Figure 4c and Figure 5c. The results demonstrate that MW-UCB can well adapt to the time-varying arrival rates to achieve stability, and consistently improves over MW with restart UCB.

6 CONCLUSION

In this paper, we present MW-UCB as a novel joint learning and scheduling algorithm for non-stationary wireless network control under partial observability and non-stationary dynamics. Our algorithmic development is based on the Max-Weight policy for network control and sliding-window UCB for learning uncertain and time-varying channel statistics. We propose a new notion of stability for non-stationary networks and prove that the MW-UCB algorithm achieves a stability region that is arbitrarily close to the true stability region. Extensive simulations on both uniformly changing and abruptly changing dynamics confirm the throughput-optimality and the favorable performance of the algorithm. We believe that our analytical framework can be extended to study stability properties of algorithms for non-stationary network control under stringent requirements of emerging large-scale wireless networks.

ACKNOWLEDGMENT

This work was supported by ONR grant N00014-20-1-2119.

REFERENCES

- [1] 2023. Full version of the paper. https://www.dropbox.com/s/j3say7gix4cpkxz/mobihoc_non_stationary_full_paper.pdf?dl=0
- [2] Arjun Anand, Gustavo De Veciana, and Sanjay Shakkottai. 2018. Joint Scheduling of URLLC and eMBB Traffic in 5G Wireless Networks. In *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*. 1970–1978. <https://doi.org/10.1109/INFOCOM.2018.8486430>
- [3] M. Andrews and L. Zhang. 2004. Scheduling over nonstationary wireless channels with finite rate sets. In *IEEE INFOCOM 2004*, Vol. 3. 1694–1704 vol.3. <https://doi.org/10.1109/INFCOM.2004.1354581>
- [4] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning* 47 (05 2002), 235–256. <https://doi.org/10.1023/A:1013689704352>
- [5] Cosmin Avasalcui, Christos Tsigkanos, and Schahram Dustdar. 2021. Adaptive Management of Volatile Edge Systems at Runtime With Satisfiability. *ACM Trans. Internet Technol.* 22, 1, Article 26 (sep 2021), 21 pages. <https://doi.org/10.1145/3470658>
- [6] Omar Besbes, Yonatan Gur, and Assaf Zeevi. 2014. Stochastic Multi-Armed-Bandit Problem with Non-stationary Rewards. In *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger (Eds.), Vol. 27. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2014/file/903ce9225fca3e988c2af215d4e544d3-Paper.pdf>
- [7] Matthew Bradbury, Arshad Jhumka, and Tim Watson. 2021. Trust Trackers for Computation Offloading in Edge-Based IoT Networks. In *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*. 1–10. <https://doi.org/10.1109/INFOCOM42981.2021.9488844>
- [8] Loc X. Bui, Sujay Sanghavi, and R. Srikant. 2009. Distributed Link Scheduling With Constant Overhead. *IEEE/ACM Transactions on Networking* 17, 5 (2009), 1467–1480. <https://doi.org/10.1109/TNET.2009.2013621>
- [9] Xin Che, Xiaohui Liu, Xi Ju, and Hongwei Zhang. 2010. Adaptive Instantiation of the Protocol Interference Model in Mission-Critical Wireless Networks. In *2010 7th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*. 1–9. <https://doi.org/10.1109/SECON.2010.5508292>
- [10] Wei Chen, Liwei Wang, Haoyu Zhao, and Kai Zheng. 2020. Combinatorial Semi-Bandit in the Non-Stationary Environment. In *Conference on Uncertainty in Artificial Intelligence*.
- [11] Guozhen Cheng, Hongchang Chen, Zhiming Wang, and Shuqiao Chen. 2015. DHA: Distributed decisions on the switch migration toward a scalable SDN control plane. In *2015 IFIP Networking Conference (IFIP Networking)*. 1–9. <https://doi.org/10.1109/IFIPNetworking.2015.7145319>
- [12] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2009. *Introduction to Algorithms, Third Edition* (3rd ed.). The MIT Press.
- [13] Salvatore Costanzo, Laura Galluccio, Giacomo Morabito, and Sergio Palazzo. 2012. Software Defined Wireless Networks: Unbridling SDNs. In *2012 European Workshop on Software Defined Networking*. 1–6. <https://doi.org/10.1109/EWSN.2012.12>
- [14] Daniel Freund, Thodoris Lykouris, and Wentao Weng. 2022. Efficient decentralized multi-agent learning in asymmetric queueing systems. In *Proceedings of Thirty Fifth Conference on Learning Theory (Proceedings of Machine Learning Research, Vol. 178)*, Po-Ling Loh and Maxim Raginsky (Eds.). PMLR, 4080–4084. <https://proceedings.mlr.press/v178/freund22a.html>
- [15] Yi Gai, Bhaskar Krishnamachari, and Rahul Jain. 2012. Combinatorial Network Optimization With Unknown Variables: Multi-Armed Bandits With Linear Rewards and Individual Observations. *IEEE/ACM Transactions on Networking* 20, 5 (2012), 1466–1478. <https://doi.org/10.1109/TNET.2011.2181864>
- [16] Cédric Hartland, Nicolas Baskiotis, Sylvain Gelly, Michèle Sebag, and Olivier Teytaud. 2007. Change Point Detection and Meta-Bandits for Online Learning in Dynamic Environments. In *Cap 2007 : 9^e Conférence francophone sur l'apprentissage automatique*. Grenoble, France, 237–250. <https://hal.inria.fr/inria-00164033>
- [17] Predrag R. Jelenkovic, Petar Momcilovic, and Mark S. Squillante. 2007. Scalability of Wireless Networks. *IEEE/ACM Transactions on Networking* 15, 2 (2007), 295–308. <https://doi.org/10.1109/TNET.2007.892846>
- [18] C. Joo, X. Lin, and N. B. Shroff. 2008. Understanding the Capacity Region of the Greedy Maximal Scheduling Algorithm in Multi-Hop Wireless Networks. In *IEEE INFOCOM 2008 - The 27th Conference on Computer Communications*. 1103–1111. <https://doi.org/10.1109/INFOCOM.2008.165>
- [19] Changhee Joo, Xiaojun Lin, and Ness B. Shroff. 2009. Greedy Maximal Matching: Performance Limits for Arbitrary Network Graphs Under the Node-Exclusive Interference Model. *IEEE Trans. Automat. Control* 54, 12 (2009), 2734–2744. <https://doi.org/10.1109/TAC.2009.2031719>
- [20] Branislav Kveton, Zheng Wen, Azin Ashkan, and Csaba Szepesvari. 2015. Tight Regret Bounds for Stochastic Combinatorial Semi-Bandits. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 38)*, Guy Lebanon and S. V. N. Vishwanathan (Eds.). PMLR, San Diego, California, USA, 535–543. <https://proceedings.mlr.press/v38/kveton15.html>
- [21] Qingkai Liang and Eytan Modiano. 2018. Minimizing Queue Length Regret Under Adversarial Network Models. *Proc. ACM Meas. Anal. Comput. Syst.* 2, 1, Article 11 (apr 2018), 32 pages. <https://doi.org/10.1145/3179414>
- [22] Michael J. Neely. 2010. *Stochastic Network Optimization with Application to Communication and Queueing Systems*. Morgan and Claypool Publishers.
- [23] Quang Minh Nguyen, M. Shahir Rahman, Xinzhe Fu, Sastry Kompella, Joseph Macker, and Eytan H. Modiano. 2022. An Optimal Network Control Framework for Wireless SDN: From Theory to Implementation. In *MILCOM 2022 - 2022 IEEE Military Communications Conference (MILCOM)*. 102–109. <https://doi.org/10.1109/MILCOM55135.2022.10017713>
- [24] Aniq Ur Rahman, Gourab Ghatak, and Antonio De Domenico. 2020. An Online Algorithm for Computation Offloading in Non-Stationary Environments. *IEEE Communications Letters* 24, 10 (2020), 2167–2171. <https://doi.org/10.1109/LCOMM.2020.3004523>
- [25] Ermin Sakic and Wolfgang Kellerer. 2020. Decoupling of Distributed Consensus, Failure Detection and Agreement in SDN Control Plane. In *2020 IFIP Networking Conference (Networking)*. 467–475.
- [26] Thomas Stahlbuhk, Brooke Shrader, and Eytan Modiano. 2019. Learning algorithms for scheduling in wireless networks with unknown channel statistics. *Ad Hoc Networks* 85 (2019), 131–144. <https://doi.org/10.1016/j.adhoc.2018.10.006>
- [27] L. Tassiulas and A. Ephremides. 1992. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automat. Control* 37, 12 (1992), 1936–1948. <https://doi.org/10.1109/9.182479>
- [28] Mark van der Boor, Sem Borst, and Johan van Leeuwen. 2017. Load balancing in large-scale systems with multiple dispatchers. In *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*. 1–9. <https://doi.org/10.1109/INFOCOM.2017.8057012>
- [29] Qinshi Wang and Wei Chen. 2017. Improving Regret Bounds for Combinatorial Semi-Bandits with Probabilistically Triggered Arms and Its Applications. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (Long Beach, California, USA) (NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 1161–1171.
- [30] Chen-Yu Wei, Yi-Te Hong, and Chi-Jen Lu. 2016. Tracking the Best Expert in Non-Stationary Stochastic Environments. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (Barcelona, Spain) (NIPS'16)*. Curran Associates Inc., Red Hook, NY, USA, 3979–3987.
- [31] Zixian Yang, R. Srikant, and Lei Ying. 2022. MaxWeight With Discounted UCB: A Provably Stable Scheduling Policy for Nonstationary Multi-Server Systems With Unknown Statistics. <https://doi.org/10.48550/ARXIV.2209.01126>
- [32] Minghao Ye, Junjie Zhang, Zehua Guo, and H. Jonathan Chao. 2021. Federated Traffic Engineering with Supervised Learning in Multi-region Networks. In *2021 IEEE 29th International Conference on Network Protocols (ICNP)*. 1–12. <https://doi.org/10.1109/ICNP52444.2021.9651918>

APPENDIX

A PROOF OF LEMMA 1

First, we prove that $\Lambda_s(\beta_1) \subseteq \Lambda_s(\beta_2)$ for any $0 \leq \beta_1 \leq \beta_2 \leq 1$. Take any $\{\lambda(t)\}_{t \geq 0} \in \Lambda_s(\beta_1)$. Then there exists some $\pi \in \Pi_s$ such that $Q_T^\pi = O(T^{\beta_1})$. Since $\beta_1 \leq \beta_2$, this implies $Q_T^\pi = O(T^{\beta_2})$ and thus $\{\lambda(t)\}_{t \geq 0} \in \Lambda_s(\beta_2)$. Therefore, we have $\Lambda_s(\beta_1) \subseteq \Lambda_s(\beta_2)$.

Second, we prove that $\Lambda_s(\beta) \subseteq \Lambda_s$ for any $\beta \in [0, 1)$. Take any $\{\lambda(t)\}_{t \geq 0} \in \Lambda_s(\beta)$. Then there exists some $\pi \in \Pi_s$ such that $Q_T^\pi = O(T^\beta)$. Since $\beta < 1$, this implies $Q_T^\pi = o(T)$ and thus $\{\lambda(t)\}_{t \geq 0} \in \Lambda_s$. Therefore, we have $\Lambda_s(\beta) \subseteq \Lambda_s$.

Finally, we prove that $\Lambda_s \subseteq \Lambda_s(1)$. Take any $\{\lambda(t)\}_{t \geq 0} \in \Lambda_s$. Then there exists some $\pi \in \Pi_s$ such that $Q_T^\pi = o(T)$, which implies $Q_T^\pi = O(T)$ and thus $\{\lambda(t)\}_{t \geq 0} \in \Lambda_s(1)$. Therefore, we have $\Lambda_s \subseteq \Lambda_s(1)$.

B PROOF OF THEOREM 2

Since $\lambda \in \bar{\Lambda}_s$, by definition there exists some $\pi^* \in \Pi_s$ that satisfies:

$$\lambda_e = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} a_e(t) \leq \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} b_e^{\pi^*}(t), \forall e \in E. \quad (24)$$

Let $\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^{|\mathcal{M}|}$ be all the admissible link activations in $\mathcal{M}$. For any time slot $t > 0$, we consider the empirical counter and distribution $\forall i \in [1, |\mathcal{M}|]$:

$$n_i(t) = \sum_{s=0}^{t-1} \mathbb{1}_{\{\mathbf{x}^{\pi^*}(s) = \mathbf{x}^i\}}, \quad (25)$$

$$p_i(t) = \frac{n_i(t)}{t}, \quad (26)$$

where $n_i(t)$ and $p_i(t)$ respectively represent the number of times and the time fraction that the link activation vector $\mathbf{x}^i$ has been chosen by the policy π^* by time t . Additionally, we consider $\forall i \in [1, |\mathcal{M}|]$:

$$S_{i,e}(t) = \begin{cases} 0, & \text{if } x_e^i = 0 \\ \frac{\sum_{s=0}^{t-1} \theta_e(s) \mathbb{1}_{\{\mathbf{x}^{\pi^*}(s) = \mathbf{x}^i\}}}{n_i(t)}, & \text{otherwise} \end{cases}. \quad (27)$$

Let $\mathbf{p}(t) = (p_i(t))_{i \in [1, |\mathcal{M}|]}$ be the vector of all such empirical distributions. Also note that $\sum_{i=1}^{|\mathcal{M}|} p_i(t) = 1$ for any $t > 0$. Since $\mathbf{p}(t) \in [0, 1]^{|\mathcal{M}|}$, by the Bolzano–Weierstrass Theorem, there exists a convergent subsequence $\{\mathbf{p}(t_k)\}_{k=1,2,\dots}$ where $\lim_{k \rightarrow \infty} t_k = \infty$. Define the limit of this convergent subsequence by:

$$\mathbf{p}^* = \lim_{k \rightarrow \infty} \mathbf{p}(t_k). \quad (28)$$

Now, we note that for any $i \in [1, |\mathcal{M}|]$ such that $p_i^* > 0$ (which also implies $\lim_{k \rightarrow \infty} n_i(t_k) = \infty$), by SLLN, we have:

$$\lim_{k \rightarrow \infty} S_{i,e}(t_k) = \mathbb{1}_{\{x_e^i=1\}} \mu_e, \forall e \in E. \quad (29)$$

With the convention that $0/0 = 0$, we obtain from (24) that $\forall e \in E$:

$$\begin{aligned} \lambda_e &\leq \lim_{k \rightarrow \infty} \frac{1}{t_k} \sum_{s=0}^{t_k-1} b_e^{\pi^*}(s) \\ &= \lim_{k \rightarrow \infty} \sum_{i=1}^{|\mathcal{M}|} p_i(t_k) \cdot S_{i,e}(t_k) \\ &\stackrel{(28)+(29)}{=} \mu_e \sum_{i=1}^{|\mathcal{M}|} p_i^* \cdot \mathbb{1}_{\{x_e^i=1\}} \end{aligned} \quad (30)$$

Consider a stationary randomized policy π_r that at any time t activates the link schedule $\mathbf{x}^{\pi_r}(t) = \mathbf{x}^i$ with probability p_i^* . Next, we consider the Max-Weight (MW) policy with the full knowledge of the statistics (i.e. MW is in Π_s) that at any time slot t , schedules the link according to:

$$\mathbf{x}^{\text{MW}}(t) = \operatorname{argmax}_{\mathbf{x} \in \mathcal{M}} \left\{ \sum_{e \in E} Q_e^{\text{MW}}(t) \mu_e x_e \right\}, \quad (31)$$

and proceeds to show that $Q_T^{\text{MW}} = O(T^{\frac{1}{2}})$. For brevity, we use π_M to denote the MW policy. Under the MW policy, we consider the quadratic Lyapunov function of the queue lengths $Q^{\pi_M}(t)$ as:

$$\Phi^{\pi_M}(Q^{\pi_M}(t)) = Q^{\pi_M}(t)^T Q^{\pi_M}(t) = \sum_{e \in E} Q_e^{\pi_M}(t)^2. \quad (32)$$

We consider the 1-step Lyapunov drift conditioned on the queue length as follows:

$$\Delta^{\pi_M}(t) = \mathbb{E}[\Phi^{\pi_M}(Q^{\pi_M}(t+1)) - \Phi^{\pi_M}(Q^{\pi_M}(t)) | Q^{\pi_M}(t)]. \quad (33)$$

From the queue process (2), we first obtain that $\forall e \in E, t \geq 0$:

$$\begin{aligned} Q_e^{\pi_M}(t+1)^2 &\leq (Q_e^{\pi_M}(t) + a_e(t) - b_e^{\pi_M}(t))^2 \\ &= Q_e^{\pi_M}(t)^2 + (a_e(t) - b_e^{\pi_M}(t))^2 \\ &\quad + 2Q_e^{\pi_M}(t)(a_e(t) - b_e^{\pi_M}(t)) \\ &\leq Q_e^{\pi_M}(t)^2 + (A_{\max} + \mu_{\max})^2 \\ &\quad + 2Q_e^{\pi_M}(t)(a_e(t) - b_e^{\pi_M}(t)), \end{aligned}$$

where in the last line we use $0 \leq a_e(t) \leq A_{\max}$, $0 \leq b_e^{\pi_M}(t) \leq \theta_e(t) \leq \mu_{\max}$. Summing the above over all $e \in E$ and taking the expectation conditioned on $Q^{\pi_M}(t)$, we obtain that:

$$\begin{aligned} \Delta^{\pi_M}(t) &\leq |E|(A_{\max} + \mu_{\max})^2 + 2 \sum_{e \in E} Q_e^{\pi_M}(t) \lambda_e(t) \\ &\quad - 2 \sum_{e \in E} Q_e^{\pi_M}(t) \mu_e \mathbb{E}[x_e^{\pi_M}(t) | Q^{\pi_M}(t)] \\ &\stackrel{(31)}{\leq} |E|(A_{\max} + \mu_{\max})^2 + 2 \sum_{e \in E} Q_e^{\pi_M}(t) \lambda_e(t) \\ &\quad - 2 \sum_{e \in E} Q_e^{\pi_M}(t) \mu_e \mathbb{E}[x_e^{\pi_r}(t) | Q^{\pi_M}(t)] \\ &= |E|(A_{\max} + \mu_{\max})^2 + 2 \sum_{e \in E} Q_e^{\pi_M}(t) \lambda_e(t) \\ &\quad - 2 \sum_{e \in E} Q_e^{\pi_M}(t) \mu_e \sum_{i=1}^{|\mathcal{M}|} p_i^* \cdot \mathbb{1}_{\{x_e^i=1\}} \\ &\stackrel{(30)}{\leq} |E|(A_{\max} + \mu_{\max})^2. \end{aligned}$$

Taking expectation of both sides with respect to $\mathbf{Q}^{\pi_M}(t)$, we have:

$$\mathbb{E}[\Phi^{\pi_M}(\mathbf{Q}^{\pi_M}(t+1))] - \mathbb{E}[\Phi^{\pi_M}(\mathbf{Q}^{\pi_M}(t))] \leq |E|(A_{\max} + \mu_{\max})^2.$$

Telescoping over $t = 0 \rightarrow T-1$ and noting that $\Phi^{\pi_M}(\mathbf{Q}^{\pi_M}(0)) = 0$, we get:

$$\mathbb{E}[\Phi^{\pi_M}(\mathbf{Q}^{\pi_M}(T))] \leq T|E|(A_{\max} + \mu_{\max})^2. \quad (34)$$

Following the same argument as the proof of Lemma 10, we can similarly obtain that:

$$\begin{aligned} Q_T^{\text{MW}} &= \mathbb{E}\left[\sum_{e \in E} Q_e^{\pi_M}(T)\right] \leq \sqrt{|E| \cdot \mathbb{E}[\Phi^{\pi_M}(\mathbf{Q}^{\pi_M}(T))]} \\ &\stackrel{(34)}{\leq} T^{\frac{1}{2}} |E|(A_{\max} + \mu_{\max}), \end{aligned}$$

which implies $Q_T^{\text{MW}} = O(T^{\frac{1}{2}})$. Since $\pi_M \in \Pi_S$, we thus have $\lambda \in \Lambda_S(\frac{1}{2})$. Since by our Definition 6, a throughput-optimal policy π , given $\beta \in [0, 1]$, achieves the stability region $\Lambda_S(\beta)$, the network is mean rate stable under π for any $\lambda \in \bar{\Lambda}_S$.

C PROOF OF LEMMA 3

Since $\{\lambda(t)\}_{t \geq 0} \in \Lambda_S(\beta)$, by definition there exists some policy $\pi^* \in \Pi_S$ such that:

$$Q_T^{\pi^*} = O(T^\beta).$$

For any τ_j , we consider $\forall e \in E$:

$$r_e(\tau_j) = \begin{cases} 0, & \text{if } \sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) = 0 \\ \left(\frac{(1-\varepsilon)(\sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) - \mathbb{E}[Q_e^{\pi^*}(\tau_{j+1})])}{\sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t)} \right)^+, & \text{otherwise} \end{cases} \quad (35)$$

and sheds the traffics such that for any $t \in [\tau_j, \tau_{j+1})$,

$$\tilde{a}_e(t) = r_e(\tau_j) a_e(t), \forall e \in E. \quad (36)$$

We now proceed to show that this shedding scheme guarantees $\{\tilde{\lambda}(t)\}_{t \geq 0} \in C(\tau, \varepsilon)$ and (18).

From the queue dynamics (2), we have $\forall e \in E$:

$$\begin{aligned} Q_e^{\pi^*}(\tau_{j+1}) &= (Q_e^{\pi^*}(\tau_{j+1}-1) + a_e(\tau_{j+1}-1) - b_e^{\pi^*}(\tau_{j+1}-1))^+ \\ &\geq Q_e^{\pi^*}(\tau_{j+1}-1) + a_e(\tau_{j+1}-1) - b_e^{\pi^*}(\tau_{j+1}-1). \end{aligned}$$

By repeating the above argument, we get $\forall e \in E$:

$$\begin{aligned} Q_e^{\pi^*}(\tau_{j+1}) &\geq Q_e^{\pi^*}(\tau_j) + \sum_{t=\tau_j}^{\tau_{j+1}-1} a_e(t) - \sum_{t=\tau_j}^{\tau_{j+1}-1} b_e^{\pi^*}(t) \\ &\geq \sum_{t=\tau_j}^{\tau_{j+1}-1} a_e(t) - \sum_{t=\tau_j}^{\tau_{j+1}-1} b_e^{\pi^*}(t). \end{aligned}$$

Taking expectation of the above, we obtain that:

$$\begin{aligned} \sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) - \mathbb{E}[Q_e^{\pi^*}(\tau_{j+1})] &\leq \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[b_e^{\pi^*}(t)], \quad \forall e \in E. \\ \therefore \left(\sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) - \mathbb{E}[Q_e^{\pi^*}(\tau_{j+1})] \right)^+ &\leq \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[b_e^{\pi^*}(t)], \quad \forall e \in E. \end{aligned} \quad (37)$$

Taking expectation of (36) and summing over $t = \tau_j \rightarrow \tau_{j+1}-1$, we obtain in view of (35) and (37) that $\forall e \in E$:

$$\sum_{t=\tau_j}^{\tau_{j+1}-1} \tilde{\lambda}_e(t) \stackrel{(36)}{=} r_e(\tau_j) \left(\sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) \right) \stackrel{(35)+(37)}{\leq} (1-\varepsilon) \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[b_e^{\pi^*}(t)],$$

which verifies $\{\tilde{\lambda}(t)\}_{t \geq 0} \in C(\tau, \varepsilon)$.

To prove (18), we note that from (17), $\mathbb{E}[X_T] = \sum_{t=0}^{T-1} \sum_{e \in E} \lambda_e(t) - \sum_{t=0}^{T-1} \sum_{e \in E} \tilde{\lambda}_e(t)$, and first proceeds to upper-bound $\sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) - \sum_{t=\tau_j}^{\tau_{j+1}-1} \tilde{\lambda}_e(t)$ by considering two cases as follows.

Case 1: $\sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) \leq \mathbb{E}[Q_e^{\pi^*}(\tau_{j+1})]$. From (35), we have $r_e(\tau_j) = 0$ and thus (from (36)) $\tilde{\lambda}_e(t) = \mathbb{E}[\tilde{a}_e(t)] = 0$. Then we obtain that:

$$\sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) - \sum_{t=\tau_j}^{\tau_{j+1}-1} \tilde{\lambda}_e(t) = \sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) \leq \mathbb{E}[Q_e^{\pi^*}(\tau_{j+1})].$$

Case 2: $\sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) > \mathbb{E}[Q_e^{\pi^*}(\tau_{j+1})] \geq 0$. We have:

$$\begin{aligned} \sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) - \sum_{t=\tau_j}^{\tau_{j+1}-1} \tilde{\lambda}_e(t) &\stackrel{(36)}{=} (1-r_e(\tau_j)) \left(\sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) \right) \\ &\stackrel{(35)}{=} \varepsilon \left(\sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) \right) + (1-\varepsilon) \mathbb{E}[Q_e^{\pi^*}(\tau_{j+1})] \\ &\leq \varepsilon \tau A_{\max} + \mathbb{E}[Q_e^{\pi^*}(\tau_{j+1})] \end{aligned}$$

Thus, in any case, we have:

$$\sum_{t=\tau_j}^{\tau_{j+1}-1} \lambda_e(t) - \sum_{t=\tau_j}^{\tau_{j+1}-1} \tilde{\lambda}_e(t) \leq \varepsilon \tau A_{\max} + \mathbb{E}[Q_e^{\pi^*}(\tau_{j+1})].$$

Summing the above over all $e \in E$ and $j = 0 \rightarrow K-1$, we obtain:

$$\begin{aligned} \mathbb{E}[X_{\tau_K}] &= \sum_{t=0}^{\tau_K-1} \sum_{e \in E} \lambda_e(t) - \sum_{t=0}^{\tau_K-1} \sum_{e \in E} \tilde{\lambda}_e(t) \\ &\leq \varepsilon \tau K A_{\max} + \sum_{j=0}^{K-1} \mathbb{E}\left[\sum_{e \in E} Q_e^{\pi^*}(\tau_{j+1})\right] \\ &= \varepsilon \tau K A_{\max} + \sum_{j=1}^K Q_{\tau_j}^{\pi^*} \\ &= O(\varepsilon \tau K + \sum_{j=1}^K \tau_j^\beta) \\ &= O(\varepsilon \tau K + \tau_K^\beta + \sum_{j=1}^{K-1} \frac{1}{\tau} \int_{\tau_j}^{\tau_{j+1}} x^\beta dx) \\ &= O(\varepsilon \tau K + \tau_K^\beta + \frac{1}{\tau} \int_{\tau_1}^{\tau_K} x^\beta dx) \\ &= O(\varepsilon \tau K + \tau_K^\beta + \frac{1}{\tau} \int_{\tau_1}^{\tau_K} x^\beta dx) \\ &= O(\varepsilon \tau K + \tau_K^\beta + \frac{\tau_K^{\beta+1}}{\tau}) \\ &= O(T^{\beta+1} \tau^{-1} + \varepsilon T). \end{aligned} \quad (38)$$

Finally, we have:

$$\begin{aligned}\mathbb{E}[X_T] &= \mathbb{E}[X_{\tau_K}] + \sum_{t=\tau_K}^{T-1} \sum_{e \in E} \lambda_e(t) - \sum_{t=\tau_K}^{T-1} \sum_{e \in E} \tilde{\lambda}_e(t) \\ &\leq \mathbb{E}[X_{\tau_K}] + \tau|E|A_{max} \\ &\stackrel{(38)}{=} O(T^{\beta+1}\tau^{-1} + \varepsilon T + \tau).\end{aligned}$$

D GUARANTEES OF THE IMAGINARY SYSTEM'S QUEUE PROCESS $\{\tilde{Q}(t)\}_{t \geq 0}$

Recall from Section 4.3 that:

- When MW-UCB is applied to the original system, it produces the sequence of decisions $\{\mathbf{x}^{\text{MW-UCB}}(t)\}_{t \geq 0}$ and thus effective service rate $\{\mathbf{b}^{\text{MW-UCB}}(t)\}_{t \geq 0}$. The queueing dynamics of the original system evolves as $\{\mathbf{Q}^{\text{MW-UCB}}(t)\}_{t \geq 0}$ via:

$$Q_e^{\text{MW-UCB}}(t+1) = (Q_e^{\text{MW-UCB}}(t) + a_e(t) - b_e^{\text{MW-UCB}}(t))^+, \forall e \in E.$$

- The "imaginary" queue lengths $\tilde{Q}(t)$ evolve as the sequence of decisions $\{\mathbf{x}^{\text{MW-UCB}}(t)\}_{t \geq 0}$ is applied to the imaginary system, i.e.

$$\tilde{Q}_e(t+1) = (\tilde{Q}_e(t) + \tilde{a}_e(t) - b_e^{\text{MW-UCB}}(t))^+, \forall e \in E.$$

In the proofs of this Appendix D, we refer to $\{\mathbf{x}^{\text{MW-UCB}}(t)\}_{t \geq 0}$, $\{\mathbf{b}^{\text{MW-UCB}}(t)\}_{t \geq 0}$ and $\{\mathbf{Q}^{\text{MW-UCB}}(t)\}_{t \geq 0}$ respectively as $\{\mathbf{x}(t)\}_{t \geq 0}$, $\{\mathbf{b}(t)\}_{t \geq 0}$ and $\{\mathbf{Q}(t)\}_{t \geq 0}$ for brevity. Consequently, the queueing dynamics of the original system and imaginary system can be respectively expressed as:

$$Q_e(t+1) = (Q_e(t) + a_e(t) - b_e(t))^+, \forall e \in E \quad (39)$$

$$\tilde{Q}_e(t+1) = (\tilde{Q}_e(t) + \tilde{a}_e(t) - b_e(t))^+, \forall e \in E. \quad (40)$$

D.1 Proof of Lemma 4

From Lemma 6, we have $\forall e \in E$:

$$Q_e(T) \leq \tilde{Q}_e(T) + \sum_{q=0}^{T-1} (a_e(q) - \tilde{a}_e(q))$$

Summing up the above over all $e \in E$ and taking expectation, we conclude that:

$$\begin{aligned}\mathbb{E}[\sum_{e \in E} Q_e(T)] &\leq \mathbb{E}[\sum_{q=0}^{T-1} (a_e(q) - \tilde{a}_e(q))] + \mathbb{E}[\sum_{e \in E} \tilde{Q}_e(T)] \\ \therefore Q_T^{\text{MW-UCB}} &\leq \mathbb{E}[X_T] + \mathbb{E}[\sum_{e \in E} \tilde{Q}_e(T)],\end{aligned}$$

where the last line follows from the definitions of Q_T and X_T .

D.2 Proof of Lemma 5

Since $\{\tilde{\lambda}(t)\}_{t \geq 0} \in C(\tau, \varepsilon)$, by definition there exists some $\pi_0 \in \Pi_s$ such that for any $\tau_j \in \{\tau_0, \tau_1, \dots, \tau_{K-1}\}$:

$$\sum_{t=\tau_j}^{\tau_{j+1}-1} \tilde{\lambda}_e(t) \leq (1-\varepsilon) \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[b_e^{\pi_0}(t)], \forall e \in E. \quad (41)$$

We consider the quadratic Lyapunov function of the queue lengths $\tilde{Q}(t)$ of the imaginary system under MW-UCB as:

$$\Phi(\tilde{Q}(t)) = \tilde{Q}(t)^T \tilde{Q}(t) = \sum_{e \in E} \tilde{Q}_e(t)^2. \quad (42)$$

We consider the τ -step Lyapunov drift of $\Phi(\cdot)$ conditioned on the queue lengths of both the original system and the imaginary system as follows:

$$\Delta(\tau_j) = \mathbb{E}[\Phi(\tilde{Q}(\tau_{j+1})) - \Phi(\tilde{Q}(\tau_j)) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)], \quad (43)$$

where we recall that $\tau_j = j\tau$. From Lemma 7 (in Appendix D.3), the drift can be upper-bounded by:

$$\begin{aligned}\Delta(\tau_j) &\leq B_1 \tau(\tau+1) + 2 \sum_{t=\tau_j}^{\tau_{j+1}-1} \tilde{Q}(\tau_j)^T \tilde{\lambda}(t) \\ &\quad - 2 \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\tilde{Q}(\tau_j)^T \mathbf{b}(t) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)],\end{aligned} \quad (44)$$

where $B_1 = \frac{1}{2}|E|(A_{max} + \mu_{max})^2$. Now, we consider the normalized queue lengths of both the original and imaginary systems as:

$$\tilde{w}_e(\tau_j) = \frac{\tilde{Q}_e(\tau_j)}{\|\tilde{Q}(\tau_j)\|_\infty}, \forall e \in E, \quad (45)$$

$$w_e(\tau_j) = \frac{Q_e(\tau_j)}{\|\mathbf{Q}(\tau_j)\|_\infty}, \forall e \in E, \quad (46)$$

with the convention that $0/0 = 0$. Note that (46) is the same weight updating rule as Line 4 of MW-UCB (Algorithm 1). Then we consider the following two regrets which respectively use (45) and (46) in their weight instantiations:

$$\begin{aligned}\bar{\mathcal{R}}(\tau_j) &= \sum_{t=\tau_j}^{\tau_{j+1}-1} \max_{\mathbf{x}' \in \mathcal{M}} \mathbb{E}[\sum_{e \in E} \tilde{w}_e(\tau_j) x'_e \theta_e(t) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)] \\ &\quad - \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\sum_{e \in E} \tilde{w}_e(\tau_j) x_e(t) \theta_e(t) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)] \\ &= \frac{1}{\|\tilde{Q}(\tau_j)\|_\infty} \sum_{t=\tau_j}^{\tau_{j+1}-1} \max_{\mathbf{x}' \in \mathcal{M}} \mathbb{E}[\sum_{e \in E} \tilde{Q}_e(\tau_j) x'_e \theta_e(t) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)] \\ &\quad - \frac{1}{\|\tilde{Q}(\tau_j)\|_\infty} \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\tilde{Q}(\tau_j)^T \mathbf{b}(t) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)],\end{aligned} \quad (48)$$

and (recalling from (16)),

$$\begin{aligned}\mathcal{R}(\tau_j) &= \sum_{t=\tau_j}^{\tau_{j+1}-1} \max_{\mathbf{x}' \in \mathcal{M}} \mathbb{E}[\sum_{e \in E} w_e(\tau_j) x'_e \theta_e(t) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)] \\ &\quad - \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\sum_{e \in E} w_e(\tau_j) x_e(t) \theta_e(t) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)] \\ &= \frac{1}{\|\mathbf{Q}(\tau_j)\|_\infty} \sum_{t=\tau_j}^{\tau_{j+1}-1} \max_{\mathbf{x}' \in \mathcal{M}} \mathbb{E}[\sum_{e \in E} Q_e(\tau_j) x'_e \theta_e(t) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)] \\ &\quad - \frac{1}{\|\mathbf{Q}(\tau_j)\|_\infty} \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\mathbf{Q}(\tau_j)^T \mathbf{b}(t) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)].\end{aligned} \quad (49)$$

$$\quad (50)$$

From Lemma 8 (in Appendix D.3), we further relate (44) to the regret $\bar{\mathcal{R}}(\tau_j)$ as follows:

$$\begin{aligned}
& \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\tilde{\mathbf{Q}}(\tau_j)^T \mathbf{b}(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)] \\
& \geq (1-\varepsilon) \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\tilde{\mathbf{Q}}(\tau_j)^T \mathbf{b}^{\pi_0}(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)] \\
& \quad + \varepsilon \tau \|\tilde{\mathbf{Q}}(\tau_j)\|_{\infty} \mu_{\min} - \|\tilde{\mathbf{Q}}(\tau_j)\|_{\infty} \bar{\mathcal{R}}(\tau_j) \\
& = \sum_{e \in E} \tilde{Q}_e(\tau_j) \left[(1-\varepsilon) \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[b_e^{\pi_0}(t)] \right] \\
& \quad + \varepsilon \tau \|\tilde{\mathbf{Q}}(\tau_j)\|_{\infty} \mu_{\min} - \|\tilde{\mathbf{Q}}(\tau_j)\|_{\infty} \bar{\mathcal{R}}(\tau_j) \\
& \stackrel{(41)}{\geq} \sum_{e \in E} \tilde{Q}_e(\tau_j) \left[\sum_{t=\tau_j}^{\tau_{j+1}-1} \tilde{\lambda}_e(t) \right] \\
& \quad + \varepsilon \tau \|\tilde{\mathbf{Q}}(\tau_j)\|_{\infty} \mu_{\min} - \|\tilde{\mathbf{Q}}(\tau_j)\|_{\infty} \bar{\mathcal{R}}(\tau_j) \tag{51}
\end{aligned}$$

where for (51) we recall that $\pi_0 \in \Pi_s$ is the idealized policy that satisfies (41). Combining (52) and (44), we have:

$$\Delta(\tau_j) \leq B_1 \tau(\tau+1) - 2\varepsilon \tau \|\tilde{\mathbf{Q}}(\tau_j)\|_{\infty} \mu_{\min} + 2\|\tilde{\mathbf{Q}}(\tau_j)\|_{\infty} \bar{\mathcal{R}}(\tau_j). \tag{53}$$

Now, we have from Lemma 9 (in Appendix D.3) that:

$$\|\tilde{\mathbf{Q}}(\tau_j)\|_{\infty} \bar{\mathcal{R}}(\tau_j) \leq \|\mathbf{Q}(\tau_j)\|_{\infty} \mathcal{R}(\tau_j) + \tau \mu_{\max} \mathbb{E}[X_{\tau_j}]. \tag{54}$$

and from Lemma 6 (in Appendix D.3) that:

$$\mathbb{E}[\|\tilde{\mathbf{Q}}(\tau_j)\|_{\infty} | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)] \geq \mathbb{E}[\|\mathbf{Q}(\tau_j)\|_{\infty} | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)] - \mathbb{E}[X_{\tau_j}] \tag{55}$$

Plugging (54) and (55) into (53), we obtain that:

$$\begin{aligned}
\Delta(\tau_j) & \leq B_1 \tau(\tau+1) + 2(\mu_{\max} + \varepsilon \mu_{\min}) \tau \mathbb{E}[X_{\tau_j}] \\
& \quad - 2\varepsilon \tau \|\mathbf{Q}(\tau_j)\|_{\infty} \mu_{\min} + 2\|\mathbf{Q}(\tau_j)\|_{\infty} \mathcal{R}(\tau_j). \tag{56}
\end{aligned}$$

Recall from Section 4.2 that MW-UCB, during every time frame $[\tau_j, \tau_{j+1})$, fixes the queue length in the original system to $\mathbf{Q}(\tau_j)$ (and thus the normalized weights $\{w_e(\tau_j)\}_{e \in E}$), and adopts the CUCB-SW algorithm for scheduling while learning the non-stationary mean service rate. Thus, the regret $\mathcal{R}(\tau_j)$ as in (49) serves to capture the learning efficiency of CUCB-SW with theoretical guarantee in [10]. To this end, the regret bound under our choice of parameter $d = \tau^{\frac{2}{3}(1-\alpha)}$ is given by Lemma 2 (in Appendix D.3) as follows:

$$\mathcal{R}(\tau_j) \leq c_0 \log(\tau) \tau^{\frac{1}{3}(\alpha+2)}, \tag{57}$$

for some universal constant $c_0 > 0$ that can be explicitly determined. Substituting (57) into (56) and defining $c_1 = \frac{c_0}{\mu_{\min}}$, we have:

$$\begin{aligned}
\Delta(\tau_j) & \leq B_1 \tau(\tau+1) + 2(\mu_{\max} + \varepsilon \mu_{\min}) \tau \mathbb{E}[X_{\tau_j}] \\
& \quad - 2\tau \|\mathbf{Q}(\tau_j)\|_{\infty} \mu_{\min} (\varepsilon - c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)}). \tag{58}
\end{aligned}$$

Now, if $\varepsilon \geq c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)}$, we obtain from (58) that:

$$\Delta(\tau_j) \leq B_1 \tau(\tau+1) + 2(\mu_{\max} + \varepsilon \mu_{\min}) \tau \mathbb{E}[X_{\tau_j}].$$

Recall from (43) that $\Delta(\tau_j) = \mathbb{E}[\Phi(\tilde{\mathbf{Q}}(\tau_{j+1})) - \Phi(\tilde{\mathbf{Q}}(\tau_j)) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)]$. Taking expectation on both sides of the above with respect to $\tilde{\mathbf{Q}}(\tau_j)$ and $\mathbf{Q}(\tau_j)$, we get:

$$\begin{aligned}
& \mathbb{E}[\Phi(\tilde{\mathbf{Q}}(\tau_{j+1}))] - \mathbb{E}[\Phi(\tilde{\mathbf{Q}}(\tau_j))] \\
& \leq B_1 \tau(\tau+1) + 2(\mu_{\max} + \varepsilon \mu_{\min}) \tau \mathbb{E}[X_{\tau_j}].
\end{aligned}$$

Summing up the above for $j = 0, 1, \dots, K-1$ and noting that $\Phi(\tilde{\mathbf{Q}}(\tau_0)) = \Phi(\tilde{\mathbf{Q}}(0)) = 0$, we obtain that:

$$\begin{aligned}
\mathbb{E}[\Phi(\tilde{\mathbf{Q}}(\tau_K))] & \leq B_1 \tau(\tau+1) K + 2(\mu_{\max} + \varepsilon \mu_{\min}) \tau \sum_{j=0}^{K-1} \mathbb{E}[X_{\tau_j}] \\
& \stackrel{(18)}{=} O(\tau^2 K + \tau \sum_{j=0}^{K-1} (\tau_j^{\beta+1} \tau^{-1} + \varepsilon \tau_j)) \\
& = O(\tau^2 K + \tau^{\beta+1} \sum_{j=0}^{K-1} j^{\beta+1} + \varepsilon \tau^2 \sum_{j=0}^{K-1} j) \\
& = O(\tau^2 K + \tau^{\beta+1} \int_0^{K-1} x^{\beta+1} dx + \varepsilon \tau^2 \sum_{j=0}^{K-1} j) \\
& = O(\tau^2 K + \tau^{\beta+1} K^{\beta+2} + \varepsilon \tau^2 K^2) \\
& = O(\tau T + T^{\beta+2} \tau^{-1} + \varepsilon T^2), \tag{59}
\end{aligned}$$

where for last line, we recall that K is the largest number such that $\tau_K = \tau K < T$. From Lemma 10 (in Appendix D.3), we have:

$$\begin{aligned}
\mathbb{E}[\sum_{e \in E} \tilde{Q}_e(\tau_K)] & \leq \sqrt{|E| \cdot \mathbb{E}[\Phi(\tilde{\mathbf{Q}}(\tau_K))]} \\
& \stackrel{(59)}{=} O(T^{\frac{1}{2}} \tau^{\frac{1}{2}} + T^{1+\frac{\beta}{2}} \tau^{-\frac{1}{2}} + \varepsilon^{\frac{1}{2}} T). \tag{60}
\end{aligned}$$

Noting that $0 < T - \tau_K \leq \tau$ and using Lemma 6 (in Appendix D.3), we have:

$$\begin{aligned}
\mathbb{E}[\sum_{e \in E} \tilde{Q}_e(T)] & = \mathbb{E}[\sum_{e \in E} \tilde{Q}_e(\tau_K)] + \mathbb{E}[\sum_{e \in E} (\tilde{Q}_e(T) - \tilde{Q}_e(\tau_K))] \\
& \leq \mathbb{E}[\sum_{e \in E} \tilde{Q}_e(\tau_K)] + (T - \tau_K)(A_{\max} + \mu_{\max}) \\
& \stackrel{(60)}{=} O(T^{\frac{1}{2}} \tau^{\frac{1}{2}} + T^{1+\frac{\beta}{2}} \tau^{-\frac{1}{2}} + \varepsilon^{\frac{1}{2}} T + \tau) \\
& = O(T^{\frac{1}{2}} \tau^{\frac{1}{2}} + T^{1+\frac{\beta}{2}} \tau^{-\frac{1}{2}} + \varepsilon^{\frac{1}{2}} T),
\end{aligned}$$

which concludes the proof of the Lemma.

D.3 Supplementary Lemmas

Lemma 6. We have the following bounds $\forall e \in E$:

$$|Q_e(t_1) - Q_e(t_2)| \leq |t_1 - t_2| (A_{\max} + \mu_{\max}) \tag{61}$$

$$|\tilde{Q}_e(t_1) - \tilde{Q}_e(t_2)| \leq |t_1 - t_2| (A_{\max} + \mu_{\max}) \tag{62}$$

$$\tilde{Q}_e(t) \leq Q_e(t) \leq \tilde{Q}_e(t) + \sum_{q=0}^{t-1} (a_e(q) - \tilde{a}_e(q)) \tag{63}$$

$$\|\tilde{\mathbf{Q}}(t)\|_{\infty} \leq \|\mathbf{Q}(t)\|_{\infty} \leq \|\tilde{\mathbf{Q}}(t)\|_{\infty} + X_t \tag{64}$$

PROOF. (61) trivially holds for $t_1 = t_2$. If $t_1 \neq t_2$, WLOG, we assume that $t_1 > t_2$. From the queue dynamics (39), we have:

$$\begin{aligned} Q_e(t_1) &\geq Q_e(t_1 - 1) + a_e(t_1 - 1) - b_e(t_1 - 1) \\ &\geq Q_e(t_1 - 1) - \mu_{\max}, \\ Q_e(t_1) &\leq Q_e(t_1 - 1) + a_e(t_1 - 1) \leq Q_e(t_1 - 1) + A_{\max}, \end{aligned}$$

where we use $0 \leq a_e(t) \leq A_{\max}$ and $0 \leq b_e(t) \leq \theta_e(t) \leq \mu_{\max}$. Iterating the above, we obtain that:

$$\begin{aligned} Q_e(t_1) &\geq Q_e(t_2) - (t_1 - t_2)\mu_{\max}, \\ Q_e(t_1) &\leq Q_e(t_2) + (t_1 - t_2)A_{\max}. \end{aligned}$$

Combining the two above, we have (61). Similarly, we obtain (62).

Next we proceed to prove (63) by induction on t .

Base case $t = 0$: Now, (63) trivially holds since $\tilde{Q}_e(0) = Q_e(0) = 0, \forall e \in E$.

Inductive step $t \rightarrow t+1$: First, we note that the number of packet arrivals of the imaginary system are shed from and thus upper-bounded by the number of packet arrivals of the original system, i.e. $\tilde{a}_e(q) \leq a_e(q), \forall e \in E, q = 0, 1, 2, \dots$. From the inductive hypothesis $\tilde{Q}_e(t) \leq Q_e(t) \leq \tilde{Q}_e(t) + \sum_{q=0}^{t-1} (a_e(q) - \tilde{a}_e(q))$ and the queue dynamics (39) and (40), we have $\forall e \in E$:

$$\begin{aligned} Q_e(t+1) &= (Q_e(t) + a_e(t) - b_e(t))^+ \\ &\geq (\tilde{Q}_e(t) + a_e(t) - b_e(t))^+ \\ &\stackrel{(40)}{=} \tilde{Q}_e(t+1), \end{aligned}$$

and,

$$\begin{aligned} Q_e(t+1) &= (Q_e(t) + a_e(t) - b_e(t))^+ \\ &\leq (\tilde{Q}_e(t) + \sum_{q=0}^{t-1} (a_e(q) - \tilde{a}_e(q)) + a_e(t) - b_e(t))^+ \\ &\leq (\tilde{Q}_e(t) + \tilde{a}_e(t) - b_e(t) + \sum_{q=0}^t (a_e(q) - \tilde{a}_e(q)))^+ \\ &\leq (\tilde{Q}_e(t) + \tilde{a}_e(t) - b_e(t))^+ + \sum_{q=0}^t (a_e(q) - \tilde{a}_e(q)) \\ &\stackrel{(40)}{=} \tilde{Q}_e(t+1) + \sum_{q=0}^t (a_e(q) - \tilde{a}_e(q)). \end{aligned}$$

Thus, (63) also holds for $t+1$.

To prove (64), we first obtain from (63) that:

$$\tilde{Q}_e(t) \leq Q_e(t) \leq \tilde{Q}_e(t) + X_t, \forall e \in E.$$

Next, we consider $e_1^* = \arg \max_{e \in E} \tilde{Q}_e(t)$ and $e_2^* = \arg \max_{e \in E} Q_e(t)$ and get that:

$$\begin{aligned} \|\tilde{Q}(t)\|_{\infty} &= \tilde{Q}_{e_1^*}(t) \leq Q_{e_1^*}(t) \leq \|\mathbf{Q}(t)\|_{\infty}, \\ \|\mathbf{Q}(t)\|_{\infty} &= Q_{e_2^*}(t) \leq \tilde{Q}_{e_2^*}(t) + X_t \leq \|\tilde{Q}(t)\|_{\infty} + X_t, \end{aligned}$$

which concludes the proof of the Lemma. $\square$

Lemma 7. We have the following bound:

$$\begin{aligned} \Delta(\tau_j) &\leq B_1\tau(\tau+1) + 2 \sum_{t=\tau_j}^{\tau_{j+1}-1} \tilde{Q}(\tau_j)^T \tilde{\lambda}(t) \\ &\quad - 2 \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\tilde{Q}(\tau_j)^T \mathbf{b}(t) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)], \end{aligned} \quad (65)$$

where $B_1 = \frac{1}{2}|E|(A_{\max} + \mu_{\max})^2$.

PROOF. From the queue process (40) of the imaginary system, we first obtain that $\forall e \in E, t \in [\tau_j, \tau_{j+1}]$:

$$\begin{aligned} \tilde{Q}_e(t+1)^2 &\leq (\tilde{Q}_e(t) + \tilde{a}_e(t) - b_e(t))^2 \\ &= \tilde{Q}_e(t)^2 + (\tilde{a}_e(t) - b_e(t))^2 + 2\tilde{Q}_e(t)(\tilde{a}_e(t) - b_e(t)) \\ &= \tilde{Q}_e(t)^2 + (\tilde{a}_e(t) - b_e(t))^2 + 2\tilde{Q}_e(\tau_j)(\tilde{a}_e(t) - b_e(t)) \\ &\quad + 2(\tilde{Q}_e(t) - \tilde{Q}_e(\tau_j))(\tilde{a}_e(t) - b_e(t)) \\ &\leq \tilde{Q}_e(t)^2 + (\tilde{a}_e(t) - b_e(t))^2 + 2\tilde{Q}_e(\tau_j)(\tilde{a}_e(t) - b_e(t)) \\ &\quad + 2|\tilde{Q}_e(t) - \tilde{Q}_e(\tau_j)||\tilde{a}_e(t) - b_e(t)| \\ &\leq \tilde{Q}_e(t)^2 + (A_{\max} + \mu_{\max})^2(1 + t - \tau_j) \\ &\quad + 2\tilde{Q}_e(\tau_j)(\tilde{a}_e(t) - b_e(t)), \end{aligned}$$

where in the last line we use $0 \leq \tilde{a}_e(t) \leq a_e(t) \leq A_{\max}$, $0 \leq b_e(t) \leq \theta_e(t) \leq \mu_{\max}$ and Lemma 6 which gives $|\tilde{Q}_e(t) - \tilde{Q}_e(\tau_j)| \leq (t - \tau_j)(A_{\max} + \mu_{\max})$. Telescoping the above for $t = \tau_j \rightarrow \tau_{j+1} - 1$ and summing over all $e \in E$, we obtain that:

$$\begin{aligned} \sum_{e \in E} \tilde{Q}_e(\tau_{j+1})^2 &\leq \sum_{e \in E} \tilde{Q}_e(\tau_j)^2 + \frac{1}{2}|E|(A_{\max} + \mu_{\max})^2\tau(\tau+1) \\ &\quad + 2 \sum_{e \in E} \tilde{Q}_e(\tau_j)(\tilde{a}_e(t) - b_e(t)). \end{aligned}$$

Taking the expectation conditioned on $\tilde{Q}(\tau_j)$ and $\mathbf{Q}(\tau_j)$ of the above and noting that $\tilde{\lambda}_e(t) = \mathbb{E}[\tilde{a}_e(t)] = \mathbb{E}[\tilde{a}_e(t) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)]$, i.e. the packet arrivals are independent of the queue lengths, we conclude the proof of (65). $\square$

Lemma 8. We have the following bound:

$$\begin{aligned} &\sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\tilde{Q}(\tau_j)^T \mathbf{b}(t) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)] \\ &\geq (1 - \varepsilon) \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\tilde{Q}(\tau_j)^T \mathbf{b}^{\tau_0}(t) | \tilde{Q}(\tau_j), \mathbf{Q}(\tau_j)] \\ &\quad + \varepsilon\tau\|\tilde{Q}(\tau_j)\|_{\infty}\mu_{\min} - \|\tilde{Q}(\tau_j)\|_{\infty}\bar{\mathcal{R}}(\tau_j) \end{aligned} \quad (66)$$

PROOF. First, letting $e^* = \arg \max_{e \in E} \tilde{Q}_e(\tau_j)$, we note that:

$$\begin{aligned} & \sum_{t=\tau_j}^{\tau_{j+1}-1} \max_{\mathbf{x}' \in \mathcal{M}} \mathbb{E} \left[\sum_{e \in E} \tilde{Q}_e(\tau_j) x'_e \theta_e(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j) \right] \\ & \geq \sum_{t=\tau_j}^{\tau_{j+1}-1} \|\tilde{\mathbf{Q}}(\tau_j)\|_\infty \cdot 1 \cdot \mathbb{E}[\theta_{e^*}(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)] \\ & \geq \sum_{t=\tau_j}^{\tau_{j+1}-1} \|\tilde{\mathbf{Q}}(\tau_j)\|_\infty \mu_{\min} = \tau \|\tilde{\mathbf{Q}}(\tau_j)\|_\infty \mu_{\min}, \end{aligned} \quad (67)$$

where for the first inequality, we compare the maximizing solution with the feasible activation link vector that activates only the link e^* . Also, we have:

$$\begin{aligned} & \sum_{t=\tau_j}^{\tau_{j+1}-1} \max_{\mathbf{x}' \in \mathcal{M}} \mathbb{E} \left[\sum_{e \in E} \tilde{Q}_e(\tau_j) x'_e \theta_e(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j) \right] \\ & \geq \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\tilde{\mathbf{Q}}(\tau_j)^T \mathbf{b}^{\pi_0}(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)], \end{aligned} \quad (68)$$

where we compare the maximizing policy with the policy π_0 with also the full knowledge of every link e 's weight, i.e.

$$\mathbb{E}[\tilde{Q}_e(\tau_j) \theta_e(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)] = \tilde{Q}_e(\tau_j) \mu_e(t).$$

Now, from (48), we have:

$$\begin{aligned} & \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\tilde{\mathbf{Q}}(\tau_j)^T \mathbf{b}(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)] \\ & = \sum_{t=\tau_j}^{\tau_{j+1}-1} \max_{\mathbf{x}' \in \mathcal{M}} \mathbb{E} \left[\sum_{e \in E} \tilde{Q}_e(\tau_j) x'_e \theta_e(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j) \right] - \|\tilde{\mathbf{Q}}(\tau_j)\|_\infty \bar{\mathcal{R}}(\tau_j) \\ & \stackrel{(67)+(68)}{\geq} (1-\epsilon) \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\tilde{\mathbf{Q}}(\tau_j)^T \mathbf{b}^{\pi_0}(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)] \\ & \quad + \epsilon \tau \|\tilde{\mathbf{Q}}(\tau_j)\|_\infty \mu_{\min} - \|\tilde{\mathbf{Q}}(\tau_j)\|_\infty \bar{\mathcal{R}}(\tau_j), \end{aligned}$$

which concludes the proof of the Lemma. $\square$

Lemma 9. *We have the following bound:*

$$\|\tilde{\mathbf{Q}}(\tau_j)\|_\infty \bar{\mathcal{R}}(\tau_j) \leq \|\mathbf{Q}(\tau_j)\|_\infty \mathcal{R}(\tau_j) + \tau \mu_{\max} \mathbb{E}[X_{\tau_j}] \quad (69)$$

PROOF. From (48), we have:

$$\begin{aligned} \|\tilde{\mathbf{Q}}(\tau_j)\|_\infty \bar{\mathcal{R}}(\tau_j) & = \sum_{t=\tau_j}^{\tau_{j+1}-1} \max_{\mathbf{x}' \in \mathcal{M}} \mathbb{E} \left[\sum_{e \in E} \tilde{Q}_e(\tau_j) x'_e \theta_e(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j) \right] \\ & \quad - \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\tilde{\mathbf{Q}}(\tau_j)^T \mathbf{b}(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)]. \end{aligned} \quad (70)$$

From Lemma 6, we have $\forall e \in E$:

$$\begin{aligned} & \mathbb{E}[Q_e(\tau_j) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)] - \sum_{t=0}^{\tau_j-1} (\lambda_e(t) - \tilde{\lambda}_e(t)) \\ & \leq \mathbb{E}[\tilde{Q}_e(\tau_j) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)] \leq \mathbb{E}[Q_e(\tau_j) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)]. \end{aligned} \quad (71)$$

Plugging (71) into (70), we have:

$$\begin{aligned} \|\tilde{\mathbf{Q}}(\tau_j)\|_\infty \bar{\mathcal{R}}(\tau_j) & = \sum_{t=\tau_j}^{\tau_{j+1}-1} \max_{\mathbf{x}' \in \mathcal{M}} \mathbb{E} \left[\sum_{e \in E} \tilde{Q}_e(\tau_j) x'_e \theta_e(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j) \right] \\ & \quad - \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\tilde{\mathbf{Q}}(\tau_j)^T \mathbf{b}(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)] \\ & \leq \sum_{t=\tau_j}^{\tau_{j+1}-1} \max_{\mathbf{x}' \in \mathcal{M}} \mathbb{E} \left[\sum_{e \in E} Q_e(\tau_j) x'_e \theta_e(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j) \right] \\ & \quad - \sum_{t=\tau_j}^{\tau_{j+1}-1} \mathbb{E}[\mathbf{Q}(\tau_j)^T \mathbf{b}(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)] \\ & \quad + \sum_{t=\tau_j}^{\tau_{j+1}-1} \sum_{t=0}^{\tau_j-1} \sum_{e \in E} (\lambda_e(t) - \tilde{\lambda}_e(t)) \mathbb{E}[b_e(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)] \\ & \stackrel{(50)}{\leq} \|\mathbf{Q}(\tau_j)\|_\infty \mathcal{R}(\tau_j) + \tau \mu_{\max} \mathbb{E}[X_{\tau_j}], \end{aligned}$$

which concludes the proof of the Lemma. $\square$

Lemma 2. [Restated] Under MW-UCB, the regret $\mathcal{R}(\tau_j)$ can be bounded by:

$$\begin{aligned} \mathcal{R}(\tau_j) & \leq |E| \left(\frac{\tau}{d} + 1 \right) \left(2\sqrt{6 \log(\tau)} + 48\sqrt{d} \log(\tau) \right) \\ & \quad + 4|E|d \cdot \gamma(\tau_j, \tau_{j+1}) + |E| \frac{\tau}{\sqrt{d}} + \frac{\pi^2}{3} |E|^2 \mu_{\max} \\ & \quad + \frac{\pi^2}{6} |E|^2 \mu_{\max} \log(2d^{1/2}). \end{aligned}$$

Under Assumption 1 and by setting $d = \Theta(\tau^{\frac{2}{3}(1-\alpha)})$, we further have $\mathcal{R}(\tau_j) = O(\log(\tau) \tau^{\frac{1}{3}(\alpha+2)})$.

PROOF. When we fix the queue lengths $\mathbf{Q}(\tau_j)$ throughout the frame $[\tau_j, \tau_{j+1})$, we aim to find a scheduling policy that solves

$$\begin{aligned} & \max_{\mathbf{x}' \in \mathcal{M}} \mathbb{E} \left[\sum_{e \in E} w_e(\tau_j) x'_e \theta_e(t) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j) \right] \\ & = \max_{\mathbf{x}' \in \mathcal{M}} \sum_{e \in E} x'_e \mu_e(t) \mathbb{E}[w_e(\tau_j) | \tilde{\mathbf{Q}}(\tau_j), \mathbf{Q}(\tau_j)], \end{aligned}$$

over the τ time slots from τ_j to τ_{j+1} despite not knowing $\boldsymbol{\mu}(t)$ and thus the true rewards at the time of making decisions. This problem, whereby the mean reward of each arm, i.e. $\mu_e(t) w_e(\tau_j)$, varies over time, can be characterized as stochastic combinatorial multi-armed bandit (SCMAB) problem in non-stationary environment and solved via the CUCB-SW algorithm [10]. To derive the bound for $\mathcal{R}(\tau_j)$, we first verify the conditions required by [10] and adapt the notations therein to our case. In particular, our model corresponds to SCMAB without probabilistically triggered arms, each of which is associated with a link $e \in E$. At any time slot $t \in [\tau_j, \tau_{j+1})$, an action is a link activation vector $\mathbf{x} \in \mathcal{M}$. The expected reward of arm $e \in E$ at time $t \in [\tau_j, \tau_{j+1})$, if it's activated, is denoted by $W_e(t) = w_e(\tau_j) \mu_e(t)$. Let $\mathbf{W}(t) = (W_e(t))_{e \in E}$ be the vector of the arms' expected rewards at time t . The total variation of the mean reward statistics inside

the frame $[\tau_j, \tau_{j+1})$ is thus depicted by:

$$V(\tau_j) = \sum_{t=\tau_j+1}^{\tau_{j+1}} \|\mathbf{W}(t) - \mathbf{W}(t-1)\|_\infty \quad (72)$$

$$\begin{aligned} &= \sum_{t=\tau_j+1}^{\tau_{j+1}} \|(w_e(\tau_j)(\mu_e(t) - \mu_e(t-1)))_{e \in E}\|_\infty \\ &\leq \sum_{t=\tau_j+1}^{\tau_{j+1}} \|\mu(t) - \mu(t-1)\|_\infty \quad (73) \\ &= \gamma(\tau_j, \tau_{j+1}) = O(\tau^\alpha), \quad (74) \end{aligned}$$

where (73) holds since $w_e(\tau_j) \leq 1$, and (74) is by Assumption 1. For convenience, we denote the total expected reward under the arms' expected rewards $\mathbf{W}$ and the action $\mathbf{x} \in \mathcal{M}$ as:

$$r(\mathbf{W}, \mathbf{x}) = \sum_{e \in E} W_e x_e. \quad (75)$$

In view of the requirements imposed by [10], given two vectors of expected rewards $\mathbf{W}$ and $\mathbf{W}'$ and any action $\mathbf{x}$, we can verify that our model satisfies both the ℓ_1 TPM bounded smoothness assumption of with constant $B = 1$, i.e.

$$|r(\mathbf{W}, \mathbf{x}) - r(\mathbf{W}', \mathbf{x})| \leq \sum_{e: x_e=1} |W_e - W'_e|,$$

and the monotonicity assumption, i.e. if $\mathbf{W} \leq \mathbf{W}'$ (entry-wise), we have:

$$r(\mathbf{W}, \mathbf{x}) \leq r(\mathbf{W}', \mathbf{x}).$$

Furthermore, for each action $\mathbf{x} \in \mathcal{M}$, the optimality gap with respect to the reward $\mathbf{W}$ is defined as $\Delta_{\mathbf{x}}^{\mathbf{W}} = \max_{\mathbf{x}' \in \mathcal{M}} r(\mathbf{W}, \mathbf{x}') - r(\mathbf{W}, \mathbf{x})$. Then for each arm $e \in E$ and $t \in [\tau_j, \tau_{j+1})$, we define:

$$\begin{aligned} \Delta_{\min}^{e,t} &= \min_{\mathbf{x} \in \mathcal{M}: \Delta_{\mathbf{x}}^{\mathbf{W}(t)} > 0} \Delta_{\mathbf{x}}^{\mathbf{W}(t)}, \\ \Delta_{\max}^{e,t} &= \max_{\mathbf{x} \in \mathcal{M}: \Delta_{\mathbf{x}}^{\mathbf{W}(t)} > 0} \Delta_{\mathbf{x}}^{\mathbf{W}(t)}. \end{aligned}$$

We define $\Delta_{\min}^{e,t} = \infty$ and $\Delta_{\max}^{e,t} = 0$ if they are not properly defined by the above definitions. Then, $\Delta_{\min} = \inf_{e \in E, t \in [\tau_j, \tau_{j+1})} \Delta_{\min}^{e,t}$ and $\Delta_{\max} = \sup_{e \in E, t \in [\tau_j, \tau_{j+1})} \Delta_{\max}^{e,t}$ are respectively the minimum and maximum gap. For our problem instance, noting that $W_e(t) = w_e(\tau_j)\mu_e(t) \leq \mu_{\max}$, we have the following bounds on the optimality gaps:

$$\begin{aligned} \Delta_{\mathbf{x}}^{\mathbf{W}(t)} &\leq \max_{\mathbf{x}' \in \mathcal{M}} r(\mathbf{W}(t), \mathbf{x}') \leq \sum_{e \in E} W_e(t) \leq |E|\mu_{\max} \\ \therefore \Delta_{\max} &\leq |E|\mu_{\max}. \quad (76) \end{aligned}$$

Following [10], given a set of positive parameters $\{M_e\}_{e \in E}$ and for any action $\mathbf{x} \in \mathcal{M}$, we define $M_{\mathbf{x}} = \max_{e: x_e=1} M_e$ with the convention that $M_{\mathbf{x}} = 0$ if $x_e = 0, \forall e \in E$. From the proof of [10, Theorem 4], we have the following regret bound given the frame size τ , sliding-window size d and any arbitrary set of positive

parameters $\{M_e\}_{e \in E}$:

$$\begin{aligned} \mathcal{R}(\tau_j) &\leq \sum_{e \in E} \left(\frac{\tau}{d} + 1 \right) \left(2\sqrt{6 \log(\tau)} + \frac{48|E| \log(\tau)}{M_e} \right) \\ &\quad + 4|E|d \cdot V(\tau_j) + \sum_{t=\tau_j}^{\tau_{j+1}-1} M_{\mathbf{x}(t)} + \frac{\pi^2}{3} |E| \Delta_{\max} \\ &\quad + \frac{\pi^2}{6} \Delta_{\max} \sum_{e \in E} j_{\max}^e, \quad (77) \end{aligned}$$

where $j_{\max}^e = \max \{ \lceil \log(\frac{2|E|}{M_e}) \rceil, 0 \}$ (also see [29]). Finally, by plugging (74) and (76) into (77) and setting $M_e = M = \frac{|E|}{\sqrt{d}}, \forall e \in E$, which also implies that $M_{\mathbf{x}(t)} = M, \forall t \in [\tau_j, \tau_{j+1})$, we conclude the required statements of the Lemma. $\square$

Lemma 10. *We have the following bound for any time slot T :*

$$\mathbb{E} \left[\sum_{e \in E} \tilde{Q}_e(T) \right] \leq \sqrt{|E| \cdot \mathbb{E}[\Phi(\tilde{\mathbf{Q}}(T))]} \quad (78)$$

PROOF. By Cauchy-Schwarz inequality, we first have:

$$\sum_{e \in E} \tilde{Q}_e(T) \leq \sqrt{|E| \left(\sum_{e \in E} \tilde{Q}_e(T)^2 \right)} = \sqrt{|E| \Phi(\tilde{\mathbf{Q}}(T))}.$$

Taking expectation of the above and by Jensen's inequality, we obtain that:

$$\mathbb{E} \left[\sum_{e \in E} \tilde{Q}_e(T) \right] \leq \mathbb{E} \left[\sqrt{|E| \Phi(\tilde{\mathbf{Q}}(T))} \right] \leq \sqrt{|E| \cdot \mathbb{E}[\Phi(\tilde{\mathbf{Q}}(T))]},$$

which concludes the proof of the Lemma. $\square$

E PROOF OF COROLLARY 1

The proof follows as a side result of a special case of the proof of Lemma 5 (Appendix D.2). In particular, since now $\{\lambda(t)\}_{t \geq 0} \in \Lambda_s(\beta)$, we can consider the shedding scheme that sheds no traffic, i.e. $\tilde{a}_e(t) = a_e(t)$ for any $e \in E$ and $t \geq 0$ and thus the total amount of shed traffic from (17) is:

$$X_T = \sum_{t=0}^{T-1} \sum_{e \in E} a_e(t) - \sum_{t=0}^{T-1} \sum_{e \in E} \tilde{a}_e(t) = 0. \quad (79)$$

Then the original and imaginary systems are now the same where $\{\mathbf{Q}(t)\}_{t \geq 0} \equiv \{\tilde{\mathbf{Q}}(t)\}_{t \geq 0}$ and the Lyapunov drift (43) can be equivalently written as:

$$\Delta(\tau_j) = \mathbb{E}[\Phi(\mathbf{Q}(\tau_{j+1})) - \Phi(\mathbf{Q}(\tau_j))] \tilde{\mathbf{Q}}(\tau_j, \mathbf{Q}(\tau_j)). \quad (80)$$

Now from (58) in the proof of Lemma 5

$$\begin{aligned} \Delta(\tau_j) &\leq B_1 \tau(\tau+1) + 2(\mu_{\max} + \varepsilon \mu_{\min}) \tau \mathbb{E}[X_{\tau_j}] \\ &\quad - 2\tau \|\mathbf{Q}(\tau_j)\|_\infty \mu_{\min} (\varepsilon - c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)}) \\ &\stackrel{(79)}{=} B_1 \tau(\tau+1) - 2\tau \|\mathbf{Q}(\tau_j)\|_\infty \mu_{\min} (\varepsilon - c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)}) \\ &\leq B_1 \tau(\tau+1) - 2\tau \left(\sum_{e \in E} Q_e(\tau_j) \right) |E|^{-1} \mu_{\min} (\varepsilon - c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)}), \quad (81) \end{aligned}$$

where for the last line we use $\sum_{e \in E} Q_e(\tau_j) \leq \sum_{e \in E} \|Q(\tau_j)\|_\infty = |E| \|Q(\tau_j)\|_\infty$. Taking expectation on both sides of (81) with respect to $\tilde{Q}(\tau_j)$ and $Q(\tau_j)$ in view of (80), we get:

$$\begin{aligned} & \mathbb{E}[\Phi(Q(\tau_{j+1}))] - \mathbb{E}[\Phi(Q(\tau_j))] \\ & \leq B_1 \tau(\tau+1) - 2\tau \mathbb{E}\left[\sum_{e \in E} Q_e(\tau_j)\right] |E|^{-1} \mu_{\min}(\varepsilon - c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)}). \end{aligned} \quad (82)$$

Now, if $\varepsilon > c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)}$, by summing (82) from $j = 0 \rightarrow K-1$ and noting that $\Phi(Q(\tau_K)) \geq 0$ and $\Phi(Q(0)) = 0$, we obtain that:

$$\sum_{j=0}^{K-1} \sum_{e \in E} \mathbb{E}[Q_e(\tau_j)] \leq \frac{B_1(\tau+1)K|E|}{2\mu_{\min}(\varepsilon - c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)})}. \quad (83)$$

By Lemma 6 and noting that $\tau_K = \tau K < T$ and $T - \tau_K \leq \tau$, we have:

$$\begin{aligned} \sum_{t=0}^{T-1} \sum_{e \in E} \mathbb{E}[Q_e(t)] &= \sum_{j=0}^{K-1} \sum_{t=\tau_j}^{\tau_{j+1}-1} \sum_{e \in E} \mathbb{E}[Q_e(t)] \\ &\quad + \sum_{t=\tau_K}^{T-1} \sum_{e \in E} \mathbb{E}[Q_e(t)] \\ &\leq \sum_{j=0}^{K-1} \sum_{t=\tau_j}^{\tau_{j+1}-1} \left[\sum_{e \in E} \mathbb{E}[Q_e(\tau_j)] \right. \\ &\quad \left. + (t - \tau_j)|E|(A_{\max} + \mu_{\max}) \right] \\ &\quad + \sum_{t=\tau_K}^{T-1} \left[\sum_{e \in E} \mathbb{E}[Q_e(\tau_{K-1})] \right. \\ &\quad \left. + (t - \tau_{K-1})|E|(A_{\max} + \mu_{\max}) \right] \\ &\leq 2\tau \sum_{j=0}^{K-1} \sum_{e \in E} \mathbb{E}[Q_e(\tau_j)] \\ &\quad + \frac{|E|(A_{\max} + \mu_{\max})}{2} (K\tau(\tau-1) + 2\tau^2 - \tau) \\ &\stackrel{(83)}{\leq} \frac{B_1(\tau+1)K|E|}{\mu_{\min}(\varepsilon - c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)})} \\ &\quad + \frac{|E|(A_{\max} + \mu_{\max})}{2} (K\tau(\tau-1) + 2\tau^2 - \tau) \\ &\leq \frac{2B_1 T |E|}{\mu_{\min}(\varepsilon - c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)})} \\ &\quad + \frac{|E|(A_{\max} + \mu_{\max})}{2} (T(\tau-1) + 2\tau^2 - \tau) \\ \therefore \frac{\sum_{t=0}^{T-1} \sum_{e \in E} \mathbb{E}[Q_e(t)]}{T} &\leq \frac{2B_1 |E|}{\mu_{\min}(\varepsilon - c_1 \log(\tau) \tau^{\frac{1}{3}(\alpha-1)})} \\ &\quad + \frac{|E|(A_{\max} + \mu_{\max})}{2} \left(\tau - 1 + \frac{2\tau^2 - \tau}{T} \right). \end{aligned}$$

Taking lim sup of the above, we conclude that required statement of the Corollary, i.e. for fixed τ ,

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \sum_{e \in E} \mathbb{E}[Q_e(t)] < \infty.$$

F SUPPLEMENTARIES FOR SIMULATIONS

For the experimental setup in Section 5, we justify that $\mathbb{E}[\gamma(t_1, t_2)] = O(|t_2 - t_1|^{1/2})$ for any $0 \leq t_1 < t_2 \leq T$ for both settings of $\delta_t = \frac{0.5}{T^{1/2}}$ and $\delta_t = \frac{0.5}{(t+1)^{1/2}}$. Recall from (3) that $\gamma(t_1, t_2) = \sum_{t=t_1+1}^{t_2} \|\mu(t) - \mu(t-1)\|_\infty$. Under our experimental setting, since $\mu_e(t), \mu_e(t-1) \in \{0.25, 0.75\}, \forall e \in E, t > 0$, we obtain that:

$$\begin{aligned} \mathbb{E}[\|\mu(t) - \mu(t-1)\|_\infty] &= 0.5 \cdot P(\mu_e(t) \text{ changes its state for some } e) \\ &= 0.5 \cdot [1 - (1 - \delta_t)^{|E|}] \\ &\leq 0.5|E|\delta_t, \end{aligned}$$

where the last line follows Bernoulli's inequality. Summing up the above for $t = t_1 + 1 \rightarrow t_2$, we have:

$$\mathbb{E}[\gamma(t_1, t_2)] \leq 0.5|E| \sum_{t=t_1+1}^{t_2} \delta_t. \quad (84)$$

If $\delta_t = \frac{0.5}{T^{1/2}}$, we obtain from (84) that:

$$\mathbb{E}[\gamma(t_1, t_2)] \leq 0.25|E| \frac{t_2 - t_1}{T^{1/2}} \leq 0.25|E| \sqrt{t_2 - t_1} = O(|t_2 - t_1|^{1/2}).$$

If $\delta_t = \frac{0.5}{(t+1)^{1/2}}$, we obtain from (84) that:

$$\begin{aligned} \mathbb{E}[\gamma(t_1, t_2)] &\leq 0.25|E| \sum_{t=t_1+1}^{t_2} \frac{1}{(t+1)^{1/2}} \\ &\leq 0.25|E| \int_{t=t_1+2}^{t_2+1} \frac{1}{x^{1/2}} dx \\ &= 0.5|E|(\sqrt{t_2+1} - \sqrt{t_1+2}) \\ &\leq 0.5|E|\sqrt{t_2 - t_1} = O(|t_2 - t_1|^{1/2}). \end{aligned}$$