Throughput-Optimal Wireless Scheduling with Regulated Inter-Service Times

Ruogu Li, Bin Li, and Atilla Eryilmaz

***Abstract*—Motivated by the low-jitter requirements of streaming multi-media traffic, we focus on the development of scheduling strategies under fading conditions that not only *maximize throughput performance* but also *provide regular inter-service times* to users. Since the service regularity of the traffic is related to the higher-order statistics of the arrival process and the policy operation, it is highly challenging to characterize and analyze directly. We overcome this obstacle by introducing a new quantity, namely the *time-since-last-service*, which has a unique evolution different from a tradition queue. By combining it with the queue-length in the weight, we propose a novel maximum-weight type scheduling policy that is proven to be throughput-optimal and also provides provable service regularity guarantees. In particular, our algorithm can achieve a degree of service regularity within a constant factor of a fundamental lower bound we derive. This constant is independent of the higher-order statistics of the arrival process and can be as low as two. Our results, both analytical and numerical, exhibit significant service regularity improvements over the traditional throughput-optimal policies, which reveals the importance of incorporating the metric of time-since-last-service into the scheduling policy for providing regulated service.**

I. Introduction

During the past years, there has been increasing deployment of a variety of real-time applications over the wireless networks, especially streaming multi-media applications. Unlike its non-real-time counterpart, the real-time traffic often has various quality-of-service (QoS) requirements besides throughput. Such requirements usually include end-to-end delay constraints, packet delivery ratio requirements, and the regularity of the inter-service times. Unlike the traditional long-term mean throughput based requirements, these QoS requirements often have a complex dependence on the higher-order statistics of the arrival process as well as the system operation. Thus, the canonical optimization-based approaches that aim to optimize the throughput performance (e.g., [1], [2], [3], [4], [5]) do not apply.

Recently, valuable efforts have been exerted in the design of algorithms that improve various aspects of the QoS, especially on the delay performance of the algorithms. For example, some works focus on designing algorithms with low end-to-end delay performance, such as [6], [7], [8]. Constant delay bounds (e.g. [9]) and delivery ratio requirements for deadline-constrained traffic (e.g. [10], [11], [12], [13]) are some of the other QoS metrics considered in the literature.

However, to the authors' best knowledge, service regularity of the scheduling policies has not been theoretically studied. Yet, the service regularity is an important metric in serving multi-media streaming applications that are subject to jitter requirements. Such applications often have a constant playback data rate, while the incoming flow may experience different source of randomness as it traverses a network, such as the arrival process, retransmissions, and the channel variations. The traditional throughput-based schedulers can guarantee the rate requirements under such stochastic characteristics of the network. However, the service received by the user under these schemes may have large variations in the inter-service times, and hence disrupt the regularity of service they need.

With these motivations, in this work, we focus on the development of scheduling policies that not only maximize throughput performance but also provide regulated inter-service times to users with heterogeneous arrival processes. However, the inter-service time characteristics are difficult to analyze directly due to: its complex dependence on the high-order statistics of the arrival and service processes, and its non-Markovian evolution. To overcome this, we need to find new approaches to study the inter-service time behavior. Our contributions in this work can be summarized as follows:

- We propose a new quantity (cf. Section II), namely the *time-since-last-service,* that has a tight relationship with the service regularity performance, and hence enables novel design strategies. Yet, this new parameter has its unique evolution, drastically different from a queue, which introduces new challenges in its analysis.
- We develop a novel maximum-weight type scheduling policy that combines the time-since-last-service parameter and the queue-length in its weight measure (cf. Section III). Using a non-traditional stochastic stability argument, we then show that the proposed scheduling policy possesses the desirable throughput optimality property (cf. Section IV-A).
- We derive upper and lower bounds on the service regularity performance by utilizing a novel Lyapunov-drift-based argument, inspired by the approach in [14]. We further show that, by properly scaling the design parameter in our policy, we can guarantee a degree of service regularity within a constant factor of our fundamental lower bound (cf. Section IV-B). This constant is independent of the higher-order statistics of the arrival process and can be as low as two under symmetric arrival rates.
- We support our analytical results with extensive numerical investigations (cf. Section V), which show significant performance gains in the service regularity over the traditional queue-length-based policies. Furthermore, the numerical investigations indicate that the service regularity performance of

Ruogu Li, Bin Li and Atilla Eryilmaz, ({lir, lib, eryilmaz}@ece.osu.edu) are with the Department of Electrical and Computer Engineering at The Ohio State University, Columbus, Ohio 43210 USA.

our policy actually approaches the lower bounds as the weight of the time-since-last-service increases.

II. System Model

We consider a time-slotted system in which N different users compete for a common resource that is subject to time variations. Each user receives externally generated packets into its dedicated queue for service by the common server. The arrival process for user i is composed of a sequence of random variables $A_i[t]$ that are independently and identically distributed (i.i.d.) over time with a fixed *arrival rate* $\lambda_i \triangleq \mathbb{E}[A_i[t]]$. The arrival processes for different users are independently (not necessarily identically) distributed. We also assume that all arrival processes have a finite support[1] $\{0, 1, \ldots, A_{max}\}$, i.e., $P(A_i[t] \leq A_{max}) = 1$ for all user i and time slot t.

Channel Fading Model: We assume the N users share a common ON/OFF fading channel, whose state in time slot t is denoted by $C[t] \in \{0, 1\}$. We assume the channel condition is i.i.d. over time, and we use $p \triangleq P(C[t] = 1)$ to denote the probability that the channel is in its ON state. We use this particular type of channel model to simplify our analytical study of the problem. We would like to point out that this simplified fading model is not restrictive: this type of common fading channel for all users is typically used to model the channel for the secondary users in the context of cognitive radio network (see [15], [16] for examples), where all secondary users within the interfering range of a transmitting primary user must keep silent.

Queueing and Service of the Packets: The packets for user i waiting to be served are placed into the queue Q_i, and we denote its length at the beginning of time slot t by $Q_i[t]$. We assume that at any given time slot t, the node can only schedule one of the users, and we use $S_i[t] \in \{0, 1\}$ to denote the scheduling decision for user i at time slot t, where $S_i[t] = 1$ represents user i is scheduled in time slot t, and

$$\sum_{i=1}^{N} S_i[t] = 1, \quad \forall i, t.$$

When user i is scheduled in time slot t, and the channel is ON for that slot, then r_i of its packets will be served. We assume $r_i > 0$ is a fixed integer for each user i. When the channel is OFF, nothing can be served. Thus, the evolution of the length of Q_i is given by

$$\begin{aligned} Q_i[t+1] &= (Q_i[t] + A_i[t] - r_i S_i[t] C[t])^+ \\ &= Q_i[t] + A_i[t] - r_i S_i[t] C[t] + U_i[t], \ \forall i, \end{aligned} \tag{1}$$

where $U_i[t] \triangleq \max(0, r_i S_i[t] C[t] - A_i[t] - Q_i[t])$ denotes the unused service by Q_i in time slot t. For convenience, we use $\boldsymbol{Q}, \boldsymbol{A}, \boldsymbol{S}$ and $\boldsymbol{U}$ to denote the N-dimensional vectors of the queue-length, arrival, scheduling decision and unused service, respectively.

We use the following definition for the stability of the system:

[1]The assumption of finite support for $(A_i[t])_i$ is not critical for our results, and can be relaxed to arrival processes with bounded moment generating functions.

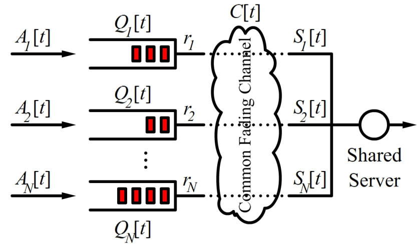

Fig. 1. The system model we use in the paper, where N users compete for a shared server with a common fading channel for all users.

Definition 1 (Stability of the System): We say the system is *stable* if the underlying Markov chain $\boldsymbol{Q}[t]$ is positive recurrent, and its corresponding stationary distribution, denoted by $\overline{\boldsymbol{Q}} \triangleq \{\overline{Q}_i\}_i$, has $\mathbb{E}[\overline{Q}_i] < \infty$ for each user i. $\diamond$

When the system is stable, we use $\overline{\boldsymbol{S}}$ and $\overline{\boldsymbol{U}}$ to denote the steady-state distribution of the scheduling vector and the unused service vector, respectively.

If under some scheduling policy, the system is stable with an arrival rate vector $\boldsymbol{\lambda} \triangleq \{\lambda_i\}_i$, then we say $\boldsymbol{\lambda}$ is *supportable* by the system. We call the set of all supportable rate vector $\boldsymbol{\lambda}$ the *capacity region* of the system. Under our system model, the capacity region, denoted by Λ, is the following set:

$$\Lambda \triangleq \left\{ \boldsymbol{\lambda} : \sum_{i=1}^{N} \frac{\lambda_i}{r_i} < p \right\}. \tag{2}$$

Thus, for any arrival rate vector within the capacity region, there exists a positive constant $\epsilon = p - \sum_{i=1}^{N} \lambda_i > 0$, which represents how close the arrival rate vector is to the boundary of the capacity region. The capacity region can be interpreted as the fraction of the busy period of the shared server cannot exceed the probability that the channel is ON.

If a policy can support all arrival rates that lies within the capacity region, then we call it *throughput-optimal*. We shall see examples of scheduling policies that possess this property in Section III.

Inter-Service Time Metric: In this work, we are interested in providing regular service to users, which relates to the statistics of the inter-service time. To that end, we use $I_i[m]$ to denote the time between the $(m-1)^{st}$ and the m^{th} service for user i. It is determined by the evolution of Q_i. Thus when the system is in steady-state, the inter-service time has a steady-state distribution, denoted as $\overline{I}_i$. We write $\overline{\boldsymbol{I}} = \{\overline{I}_i\}_i$ as the N-dimensional vector of $\overline{I}_i$.

We use the second moment of the inter-service time under steady-state distribution, i.e., $\mathbb{E}[\overline{I}_i^2]$, as a measure of the *regularity* of the service that user i receives. This metric conveniently captures information about both the average service rate that the user receives, and also the variance of its inter-service time. As such, our results in this work can be easily extended to characterize the variance of the inter-service time, which we also uses as an alternative metric for service regularity.

In this work, we are interested in the development of throughput-optimal policies that achieve low values of $\sum_{i=1}^{N} \mathbb{E}[\overline{I}_i^2]$ in steady-state, implying regular service. However, unlike queue-lengths with Markovian evolution, the dynamics of inter-service times do not lend themselves to commonly used Markovian analysis methods. To overcome this obstacle, we introduce the following related quantity, namely the *time-since-last-service*, which has a much more tractable form of evolution, and whose mean is tightly related to $\mathbb{E}[\overline{I}_i^2]$ (cf. Lemma 1).

Time-Since-Last-Service (TSLS): For each user i, we use a timer T_i to keep track of the time since it was lastly *served*, i.e., it was scheduled and the channel was ON. By letting

$$\tau_i[t] \triangleq \max_{\tau=\{1,\dots,t\}} \left\{ \begin{array}{l} S_i[\tau]C[\tau]=1, S_i[\tau+1]C[\tau+1]= \\ \cdots = S_i[t-1]C[t-1]=0 \end{array} \right\},$$

be the last time when user i was served before time slot t, we can write $T_i[t] = t - \tau_i[t] - 1$. By definition, each counter T_i increases by 1 in each time slot, and drops to 0 whenever user i is served. More precisely, the evolution of the counter T_i can be written as

$$\begin{aligned} T_i[t+1] &= \begin{cases} 0 & \text{, if } S_i[t]C[t]=1; \\ T_i[t]+1 & \text{, if } S_i[t]C[t]=0; \end{cases} \\ &= \mathbb{1}\{S_i[t]C[t]=0\}(T_i[t]+1), \end{aligned} \tag{3}$$

where $\mathbb{1}\{\cdot\}$ is the indicator function. We also concisely denote the N-dimensional vector $\{T_i\}_i$ by $\boldsymbol{T}$.

It can be seen from (3) that the evolution of $T_i[t]$ differs significantly from that of a traditional queue (also see Fig. 2). In particular, unlike the slowly evolving nature of queue-lengths, the $T_i[t]$ is incremented until user i receives service at which time it drops to zero. In our design, we will consider policies that not only use $\boldsymbol{Q}[t]$ to achieve throughput-optimality, but also include $\boldsymbol{T}[t]$ to improve service regularity. However, as we shall see in Section IV-A, the involvement of $\boldsymbol{T}[t]$ with its unique dynamics makes traditional Foster-Lyapunov arguments inapplicable, and necessitates the development of a novel approach to establish the stability of the proposed policy.

The evolution of T_i is tightly related to the inter-service time I_i, where I_i is the time between two consecutive instances when T_i hits zero, as shown in Fig. 2. In fact, we have the following lemma relating the two in steady-state.

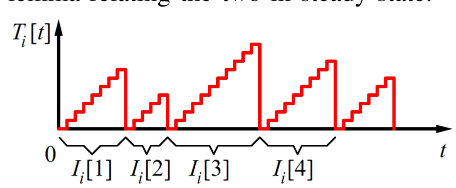

Fig. 2. A sample trajectory of $I_i[m]$ and $T_i[t]$, where the red curve shows the evolution of $T_i[t]$.

Lemma 1: For any stationary and ergodic policy with $\overline{\boldsymbol{T}}$ and $\overline{\boldsymbol{I}}$ describing its steady-state TSLS and inter-service times, respectively, we have

$$\mathbb{E}[\overline{T}_i] = \frac{1}{2}\left(\frac{1}{\mathbb{E}[\overline{I}_i]}\mathbb{E}[\overline{I}_i^2] - 1\right), \tag{4}$$

for each user i. ◇

Proof: The detailed proof is provided in [17]. ■

Lemma 1 reveals the connection between the second moment of the inter-service time $\overline{I}_i$ and the mean of the TSLS $\overline{T}_i$ in steady-state. This can be intuitively seen in Fig. 2, where the area of each "triangle" under the trajectory of $T_i[t]$ is roughly $\frac{1}{2}I_i^2$. We will utilize this useful relationship to investigate the second moment behavior of the inter-service time by studying the evolution of $T_i[t]$.

Objective: Given the above model, in this work, we aim to design a scheduling policy that is not only throughput-optimal, but also yields provably good characteristics in the service regularity as measured through the second moment of the inter-service time.

We achieve this dual objective by developing a parametric class of schedulers (cf. Section III-B) that utilize a combination of $\boldsymbol{Q}[t]$ and $\boldsymbol{T}[t]$ in its decisions. Our policy is shown to be throughput-optimal (cf. Section IV-A) through non-traditional arguments, and guarantees a constant ratio (as a function of the arrival rates) in its service regularity with respect to a fundamental lower bound (cf. Section IV-B).

III. Policy Design for Regulated Service

In this section, we first (cf. Section III-A) recall two baseline scheduling policies each with a favorable characteristic in either its throughput or service regularity performance. We then (cf. Section III-B) propose our Regulated Throughput-Optimal policy which will later (cf. Section IV) be shown to possess the advantages of both.

A. Baseline Schedulers

In this subsection, we describe two well-known scheduling policies, namely the Maximum Weight (MW) policy and the Round Robin (RR) policy, that will be used as a baseline for our policy. We selected these policies, since the MW policy possesses the throughput-optimality-characteristics without any guarantees on service regularity, and the RR policy provides regular service to the users without guarantees on system stability.

Definition 2 (MW Policy): Under our model, the Maximum Weight (MW) policy schedules the user $i^{MW}[t]$ with the maximum queue-length, i.e., it chooses

$$i^{MW}[t] \in \underset{1\le i\le N}{\operatorname{argmax}}\left(Q_i[t]\right),$$

and sets $S_{i^{MW}}[t]=1$. ◇

The MW policy is known to be throughput-optimal (e.g., [1], [5], [18], [19]), i.e., it stabilizes the network for any arrival rate vector $\boldsymbol{\lambda}$ that lies within the capacity region Λ. In our setup, the MW policy can be expected to have close-to-lower-bound average delay performance ([20]). It has also been shown to have heavy-traffic optimality (see [21], [14]), in the sense that the expected steady-state queue-length coincides with the lower-bound under heavy-traffic conditions.

However, despite its throughput optimality and a number of favorable properties on the delay performance, the MW policy may result in poor performance in terms of service regularity. This can be observed when the MW policy serves a set of flows with heterogeneous arrival statistics. In Fig. 3, the blue line shows a scenario where the i^{th} user has a Bernoulli arrival with rate 2^{-i} for $i \in \{1,\dots,8\}$. In this case, we observe that the variance of the inter-service time increases exponentially as the arrival rate of the user reduces. The red curve illustrates a different scenario where all 8 users have the same mean, but increasing variances (i.e., burstiness) in their arrivals, where we observe that the user with more bursty arrivals suffers from higher variance in its inter-service time.

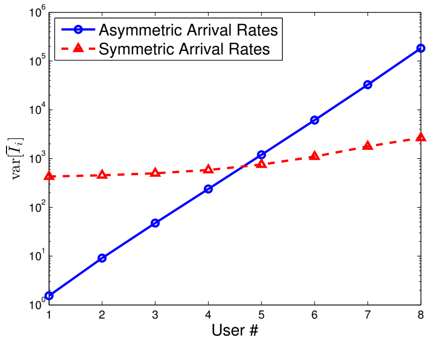

Fig. 3. The variance of the inter-service under the MW policy for users with different arrival processes. The users with smaller rates or more bursty arrivals suffer from high variance.

Next, we turn to the definition of the round-robin policy that schedules the users in turn and periodically. In our model, its operation can be described as follows.

Definition 3 (RR Policy): Under our model, the Round-Robin (RR) policy schedules the user $i^{RR}[t]$ with the maximum time-since-last-service, i.e., it chooses

$$i^{RR}[t] = \underset{1\leq i\leq N}{\text{argmax}} \left(T_i[t]\right),$$

and sets $S_{i^{RR}}[t] = 1$. $\diamond$

The RR policy, thus, serves the users in a fixed order whenever the channel is ON. As such, it provides highly regular service to the users. In fact, it will be revealed in Section IV-B that it minimizes our service regularity metric $\sum_{i=1}^{N} \mathbb{E}[\overline{T}_i]$ under our channel model. However, unlike the MW policy, the RR policy does not possess the throughput optimality property. In fact, it can only stabilize the queues for arrival rates satisfying $\frac{\lambda_i}{r_i} < \frac{p}{N}$ for each user i.

B. The Regulated Throughput-Optimal (RTO) Policy

As discussed above, the MW policy is throughput-optimal, and the RR policy has favorable service regularity performance, and can both be written in the same maximization form with different quantities as weights. Inspired by these two policies, we propose the following parametrized policy which is later revealed to possess the characteristics of throughput optimality and service regularity.

Definition 4 (Regulated Throughput-Optimal (RTO) Policy): In each time slot t,

- **Weight Calculation:** For each user i, compute its weight:
$$w_i[t] = \alpha_i Q_i[t] + \gamma\beta_i T_i[t], \qquad i = 1, \ldots, N$$
where $\alpha_i > 0$, $\beta_i \geq 0$ and $\gamma \geq 0$ are fixed control parameters.
- **Scheduling:** The scheduler chooses the user $i^*[t]$ with the maximum weight for transmission, i.e., select
$$i^*[t] \in \underset{1\leq i\leq N}{\text{argmax}} \left(w_i[t]\right),$$
set $S_{i^*}[t] = 1$ and $S_i[t] = 0$ for all $i \neq i^*[t]$.
- **Queue Evolution:** Each user i updates its queue-length $Q_i[t]$ as in (1).
- **Timer Evolution:** The timer $T_i[t]$ associated with each user i updates its value as in (3). $\diamond$

We note that there are three sets of control parameters in the RTO policy and they affect different behaviors of the policy. Yet, it will be revealed later that none of them affects its throughput optimality.

The parameters α_i are weighing factors for the queue-lengths, where a larger α_i will result in a smaller average queue-length. Some examples of possible choices of α_i can be let they be all equal, where the resulting average queue-length will be roughly the same for each user, or let $\alpha_i = \frac{1}{r_i}$ for each user, under which the average queueing delay will be roughly equal for all users.

The parameters β_i weigh $T_i[t]$ differently for each user i, with γ being a common scaling factor for all users. It will be revealed in Section IV-B that with proper choice of β_i, we can derive the upper bound for the service regularity performance, and the scaling factor γ can reduce the second moment of the inter-service time as it increases.

Also note that when $\gamma = 0$, our policy coincides with the MW policy. When $\gamma > 0$, with the addition of $T_i[t]$ terms in the weight of each user, it is clear that our policy may schedule a user with no packets waiting. Despite of this, we can still show that our policy possesses the throughput optimality property.

IV. Performance Analysis

In this section, we study the performance of our proposed RTO policy analytically. We first show that our policy is throughput-optimal, and hence attains a steady-state distribution. This, then, enables us to study the second moment of the inter-service time under the steady-state distribution of the system using the novel Lyapunov type approach developed in [14]. In particular, we derive explicit lower and upper bounds on the service regularity as a function of system and design parameters. These investigations reveal that the service regularity performance of our RTO policy can be guaranteed to remain within a constant factor of the lower bound. This ratio is expressed as a function of the arrival processes and the design parameters, and can be as low as 2.

A. Throughput Optimality

As discussed in Section III-B, our policy may schedule a user that has not enough packets waiting in its queue, and potentially wastes some service even if there are other users in the system with enough packets to send. Hence it is not immediately clear that whether our policy is throughput-optimal.

We study the stability of the system under a fixed supportable arrival rate vector using a Lyapunov drift argument by looking at the behavior of the total normalized queue-length defined by $Q_\Sigma[t] \triangleq \sum_{i=1}^N \frac{Q_i[t]}{r_i}$. Based on the queue evolution given by (1), its evolution can be written as

$$Q_\Sigma[t+1] \begin{cases} = Q_\Sigma[t] + A_\Sigma[t] - C[t] & \text{if } Q_{i^*}[t] \geq r_{i^*} \\ \leq Q_\Sigma[t] + A_\Sigma[t] & \text{if } Q_{i^*}[t] < r_{i^*} \end{cases}, \quad (5)$$

where $A_\Sigma \triangleq \sum_{i=1}^N \frac{A_i[t]}{r_i}$ is the normalized total arrival rate, and $Q_{i^*}[t]$ denotes the queue-length of the selected user at time slot t. The two different conditions on the right hand side of (5) correspond to the cases where no service is wasted (the upper equation), and the service is (partially) wasted (the lower equation). We first prove the following lemma which is critical in the establishing of the throughput optimality.

Lemma 2: If $Q_{i^*}[t] < r_{i^*}$ for any time slot t, then we have

$$\mathbb{E}\left[Q_\Sigma[t]\mathbb{1}\{Q_{i^*}[t] < r_{i^*}\}\right] \leq B_1(\gamma), \quad (6)$$

where $B_1(\gamma)$ is a finite constant (derived explicitly in the proof in terms of the system parameters) that is independent of the chosen user $i^*[t]$, and scales linearly with the control parameter γ. ◇

Proof: The detailed proof is provided in Appendix A. ■

Lemma 2 reveals the important fact that whenever the service is wasted due to choosing a user with too few packets in its queue, the expected total normalized queue-length at that time slot is bounded by a constant. Utilizing this lemma, we have the following proposition:

Proposition 1: For any arrival rate vector $\boldsymbol{\lambda} \in \Lambda$, our policy stabilizes the system in the sense that the Markov chain $(\boldsymbol{Q}[t], \boldsymbol{T}[t])$ is positive recurrent, with

$$\limsup_{K\to\infty} \frac{1}{K} \sum_{t=0}^{K-1} \mathbb{E}[Q_\Sigma[t]] \leq \frac{B_1(\gamma) + B_2}{\epsilon}.$$

where B_2 is some finite constant related to the second moment of the arrival processes, channel statistics and the design parameters. ◇

Proof: The detailed proof is provided in Appendix B. ■

Proposition 1 establishes the throughput optimality of the RTO policy, thus $Q_i[t]$ and $T_i[t]$ will converge in distribution to $\overline{Q}_i^*$ and $\overline{T}_i^*$, which attain steady-state distribution under our policy. Proposition 1 also gives an upper bound for the expected total queue-length under the steady-state, which increases linearly with the control parameter γ. It will be revealed later that γ controls the tradeoff between the average total queue-length, and the service regularity performance.

Note that our policy coincides with the MW policy when $\gamma = 0$, and the upper bound given above also coincides with that of the MW policy which can be derived via the traditional Foster-Lyapunov type of argument. The connection between the RTO and the RR policy will be revealed in the next subsection, where we study the behavior of the second moment of the inter-service time.

B. Second Moment of the Inter-Service Time

Having established the throughput optimality of the RTO policy, in this section, we utilize it to derive bounds on the expected value of $\overline{T}_\Sigma \triangleq \sum_{i=1}^N \overline{T}_i$, which is related to the second moment of the inter-service time by Lemma 1. We assume the parameter $\gamma > 0$ throughout this subsection.

1) Lower Bound on the Mean of TSLS: Here, we first derive a lower bound based on a Lyapunov drift argument inspired by the technique used in [14], which is more generally applicable and easier to extend. Then we study the operation of the RR policy under our channel model and show it is optimal in the sense of minimizing $\mathbb{E}[\overline{T}_\Sigma]$. The latter yields a tighter lower bound, however it relies on the particular channel fading model we are using in this work, and is difficult to extend to a more general setup where each user has independent fading channel.

To study the lower bound of the second moment of the inter-service time by the Lyapunov drift argument, we look at a class of policies, called Π, that can guarantee that the Markov chain $(\boldsymbol{Q}[t], \boldsymbol{T}[t])$ is positive recurrent. Note that our proposed policy, as well as the MW policy, falls into this class by Proposition 1. For such class of policies, we have the following lemma:

Lemma 3: For any policy $\pi \in \Pi$, we have

$$\mathbb{E}\left[\overline{T}_\Sigma^\pi\right] \geq \frac{N(N-p)}{2p}, \quad (7)$$

where $\overline{T}_\Sigma^\pi \triangleq \sum_{i=1}^N \overline{T}_i^\pi$ attains the steady-state distribution under policy π. ◇

Proof Sketch: To prove Lemma 3, we investigate the Lyapunov function given by $V_{\boldsymbol{T}}(\boldsymbol{Q}[t], \boldsymbol{T}[t]) = \sum_{i=1}^N T_i^2[t]$. and use the fact that its mean drift is zero under the steady-state operation of π. The detailed proof is provided in [17]. ■

We then study the performance of the RR policy under our model, which intuitively has the best service regularity when the channel is always ON, since in that case, it serves each user with a fixed inter-service time. In the following proposition, we establish a better lower bound for $\mathbb{E}[\overline{T}_\Sigma]$ by showing the optimality of the RR policy and analyzing its performance.

Proposition 2: Under our model where all users shares a common fading channel, the RR policy minimizes $\mathbb{E}[\overline{T}_\Sigma]$ over all policies. More precisely, it yields

$$\mathbb{E}[\overline{T}_\Sigma^{RR}] = \frac{N(N+1-2p)}{2p}, \quad (8)$$

where $\overline{T}_\Sigma^{RR}$ is the steady-state distribution of the total TSLS for all users under RR policy. ◇

Proof Sketch: The optimality of the RR policy is established by a sample-path argument. We first look at the case where channel is always in its ON state. In this scenario, the idea is that starting from any initial state of $\boldsymbol{T}[1]$, for any policy π_0, we can construct another policy π_1, such that π_1 agrees with the RR policy in the first time slot, and couples with policy π_0 in all subsequent time slots such that π_1 has a $T_\Sigma[t]$ value no larger than policy π_0 for all t. Repeat this construction recursively we get our sample path optimality result for the RR policy. Since no policy can do anything when the channel is OFF, the above result holds for the common fading channel with general p values.

To calculate the value of $\mathbb{E}[\overline{T}_\Sigma]$, we argue that the TSLS vector $\boldsymbol{T}^{RR}$ under RR policy is always a permutation of $\{\tau - 1, 2\tau - 1, \ldots, N\tau - 1\}$ when the channel is ON, where τ is the geometric random variable denoting the number of time slots between two consecutive ON state of the channel. Take expectation and we get the desired result.

The detailed proof is provided in [17]. ■

Note that $\frac{N(N-p)}{2p} \leq \frac{N(N+1-2p)}{2p}$, where the equality holds when $p=1$. Since the RR policy minimizes $\mathbb{E}[\overline{T}_\Sigma]$, the lower bound provided in Proposition 2 is tighter than the one in Lemma 3. However, the approach used in deriving the lower bound in Lemma 3 is more generally applicable to different channel fading model, while the above proof only works under the particular channel fading model used in this work.

We also would like to point out that the RR policy is not throughput-optimal. Thus, for an arrival rate vector $\boldsymbol{\lambda}$ that can not be supported by the RR policy, we do not expect a throughput-optimal policy to approach the above lower bound when serving it. However, for the arrival rate vectors that can be supported by the RR policy, we shall see in our numerical results that the performance of our policy can approach this lower bound when we increase the scaling parameter γ.

2) Upper Bound on the Mean of TSLS: Now we derive the performance upper bound for our policy using the same technique as in the derivation of Lemma 3. Recall that our policy is throughput-optimal by Proposition 1, thus for all $\boldsymbol{\lambda} \in \Lambda$, the Markov chain $(\boldsymbol{Q}[t], \boldsymbol{T}[t])$ is positive recurrent. We established the following upper bound for our policy:

Proposition 3: Under the steady-state operation of our policy, we have

$$\sum_{i=1}^{N} \frac{\lambda_i \beta_i}{r_i} \mathbb{E}[\overline{T}_i^*] \leq \frac{p-\epsilon}{2\gamma p}\left(\sum_{i=1}^{N} \mathbb{E}\left[\frac{\alpha_i A_i^2}{r_i}\right] + p\right) + \frac{(p-\epsilon)}{p}\sum_{i=1}^{N}(\beta_i + p\beta_i \mathbb{E}[\overline{S}_i^*]),$$

where $\overline{S}_i^*$ represents the limiting random variables of the scheduling decisions achieved in steady-state under our RTO policy. Under the particular choice of $\alpha_i = \frac{1}{r_i}$, $\beta_i = \frac{\lambda_i}{r_i}$, the above equation gives the following upper bound for the total expected TSLS:

$$\mathbb{E}[\overline{T}_\Sigma^*] \leq \frac{p-\epsilon}{2\gamma p}\left(\sum_{i=1}^{N} \mathbb{E}\left[\frac{A_i^2}{r_i^2}\right] + p\right) + \frac{(p-\epsilon)}{p}\sum_{i=1}^{N}\left(\frac{r_i}{\lambda_i} + \frac{p r_i}{\lambda_i}\mathbb{E}[\overline{S}_i^*]\right). \tag{9}$$

If we further assume that the arrival rates are symmetric, i.e., $\frac{\lambda_i}{r_i} = \frac{p-\epsilon}{N}$ for each user i, then (9) can be rewritten as

$$\mathbb{E}[\overline{T}_\Sigma^*] \leq \frac{N}{2\gamma p}\left(\sum_{i=1}^{N} \mathbb{E}[A_i^2] + p\right) + \frac{N(N-p)}{p}. \tag{10}$$

◇

Proof Sketch: To derive the upper bound, we study the drift of the Lyapunov function given by $V_{\boldsymbol{Q}}(\boldsymbol{Q}[t], \boldsymbol{T}[t]) = \frac{1}{2}\sum_{i=1}^{N} \frac{\alpha_i}{r_i} Q_i^2[t]$, and utilizing the fact that its mean drift is zero under steady-state operation. We also need to use the fact that our policy schedules the user $i^*[t]$ at time slot t such that $\alpha_i Q_i[t] + \gamma \beta_i T_i[t]$ is maximized. The detailed proof is provided in [17]. ■

Note that the first term of the right hand side of (9) captures various random effects in the network: the burstiness of the arrival processes (which is captured by the second moment) and the channel variations. Under our policy these effects diminish as the scaling factor γ goes to infinity. Hence, together with Proposition 1 in Section IV-A, Proposition 3 reveals a tradeoff: when increasing γ, the upper bound on the total queue-length increases linearly with γ, but the upper bound for $\mathbb{E}[\overline{T}_\Sigma]$ decreases.

As γ goes to infinity, the upper bound for the symmetric arrival rate case given in (10) becomes $\frac{N(N-p)}{p}$, which is always within twice the value of the lower bound derived in Proposition 2. In the more general case, the upper bound converges to a constant that is determined by the arrival rates as γ goes to infinity. Moreover, we shall see in the numerical results presented in Section V-B that as γ increases, $\mathbb{E}[\overline{T}_\Sigma]$ actually converges to the lower bound shown in Proposition 2 under the symmetric arrival rate scenario.

3) Service Regularity Bounds under Symmetric Arrival Rates: Combining the results obtained in Proposition 2 and 3, together with the relationship between $\mathbb{E}[\overline{T}_i]$ and $\mathbb{E}[\overline{I}_i^2]$ established in Lemma 1, we can explicitly express the upper and lower bounds for $\sum_{i=1}^{N} \mathbb{E}[\overline{I}_i^2]$.

Note that since the arrival rate vector is supportable, under the steady-state operation we have $\frac{1}{\mathbb{E}[\overline{I}_i]} = \mu_i$, where μ_i is the fraction of time that user i is scheduled. The system is stable, hence $\mu_i \geq \frac{\lambda_i}{r_i} = \frac{p-\epsilon}{N}$ for all user i, which further implies $\mu_i \leq \frac{p-\epsilon}{N} + \epsilon$ for all users since $\sum_{i=1}^{N} \mu_i = 1$. Substitute the above relationships into (4) yields

$$\sum_{i=1}^{N} \mathbb{E}[\overline{T}_i] \leq \frac{p+(N-1)\epsilon}{2N}\sum_{i=1}^{N} \mathbb{E}[\overline{I}_i^2] - \frac{N}{2}$$

$$\sum_{i=1}^{N} \mathbb{E}[\overline{T}_i] \geq \frac{p-\epsilon}{2N}\sum_{i=1}^{N} \mathbb{E}[\overline{I}_i^2] - \frac{N}{2}$$

under symmetric arrival rates. By substituting the above upper and lower bounds into (8) and (10) respectively, we have

$$\frac{N^2(N+1-p)}{p^2+p(N-1)\epsilon} \leq \sum_{i=1}^{N} \mathbb{E}[\overline{I}_i^2] \leq \frac{N^2}{\gamma p(p-\epsilon)}\left(\sum_{i=1}^{N} \mathbb{E}[A_i^2] + p\right) + \frac{2N^3 - N^2 p}{p(p-\epsilon)}.$$

As γ goes to infinity, the upper bound converges to a constant factor with respect to the lower bound, similar to the upper bound on $\mathbb{E}[\overline{T}_i]$ discussed above. Moreover, as we shall see in our numerical investigations, that our policy performs much better than the upper bound, giving a close-to-lower-bound performance when γ increases.

V. Numerical Results

In this section, we provide simulation results for our proposed RTO policy and compare its performance to the baseline policies and bounds. In addition to investigating the throughput (cf. Section V-A) and service regularity (cf. Section V-B) performances of our policy, we also look at the service rate distribution (cf. Section V-C) among the users to reveal one of the reasons why our policy has better service regularity than the MW policy. In all of our simulations, we assume $p = 0.8$, $r_i = 1$ and $\alpha_i = 1$ for all users, unless otherwise specified.

A. Throughput Performance

To illustrate the throughput performance, we use a set of flows and let their total rate increase to the boundary of the capacity region. In Fig. 4, we compare the total queue-length under the MW policy, as well as our policy with different γ values.

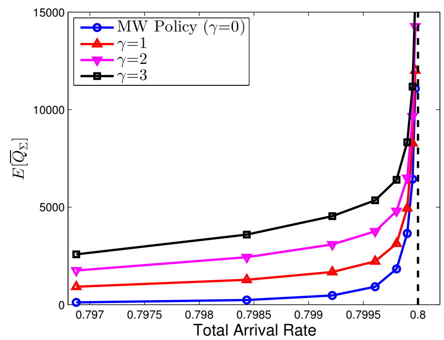

Fig. 4. The total queue-length of MW policy and our policy with different γ values.

It can be observed in Fig. 4 that the total queue-length remains finite for all arrival rates that lies within the capacity region given by $\sum_{i=1}^{N} \frac{\lambda_i}{r_i} < p = 0.8$, which confirms that our policy is indeed throughput-optimal. It also can be observed that the total queue-length of our policy increases with the parameter γ. This is expected since as γ increases, it becomes more likely for the RTO policy to choose a queue with less packet to serve, potentially wasting some service while improving the service regularity, as we shall see next.

B. Service Regularity Performance

In this subsection, we investigate the service regularity performance of our RTO policy, as well as illustrate the tradeoff between the total queue-length and the service regularity. We presents our results in two scenarios, namely the symmetric arrival rates, and the asymmetric ones. We choose the scaling parameter γ to be the powers of 2, ranging from 2^{-7} to 2^{7}.

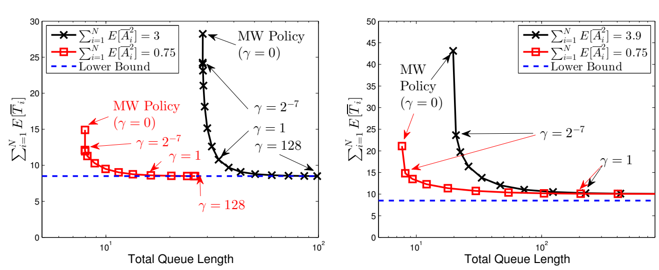

(a) (b)

Fig. 5. The lower bound and the relationship between the expected total queue-length and the expected total TSLS value for (a) symmetric arrival rates with different burstiness and (b) asymmetric arrival rates.

1) Symmetric Arrival Rates: To investigate the service regularity performance of our policy, we run the simulation with a set of 4 users with symmetric arrival rates, with $\epsilon = \frac{1}{16}$. All users have different arrival distributions such that their second moment is different from each other.

Fig. 5(a) shows the relationship between the expected total queue-length and the expected total TSLS value for 2 different set of arrivals with different second moments. The tradeoff between the service regularity and the total queue-length can be clearly seen: as γ increases, the service regularity improves while the total queue-length also increases. It can be observed that in both cases, the simulated values converge to the lower bound achieved by the RR policy with relatively small γ values.

2) Asymmetric Arrival Rates: In this setup, we have 4 flows with rates $\frac{12}{40}$, $\frac{9}{40}$, $\frac{6}{40}$ and $\frac{3}{40}$, respectively, and the parameter β_i for user i is chosen to be $\frac{1}{\lambda_i}$. Since this arrival rate is outside the stability region of the RR policy, we do not expect our RTO policy to converge to the lower bound provided by it. Indeed, it can be observed in Fig. 5(b) that under both set of the arrival processes, the total expected TSLS value converges to a point that is larger than the lower bound.

We can also observe the effect of the inclusion of the TSLS on the service regularity under the asymmetric arrival rates scenario. Even with very small γ values (e.g., 2^{-7}), our RTO policy significantly improves the service regularity, while introducing negligible increase in the total queue-length.

C. Distribution of Services among the Users

In this subsection, we study the distribution of the services among the users under non-fading conditions to illustrate one of the reasons why our policy outperforms MW policy in terms of service regularity. We consider a set of 8 users with symmetric arrival rates. Yet, all users have different arrival distributions, with the arrival of the i^{th} user being either 2^{i-1} with probability $\frac{1-\epsilon}{2^{i-1}N}$ or 0. Thus the second moment of the arrival process for user i is $\frac{2^{i-1}(1-\epsilon)}{N}$.

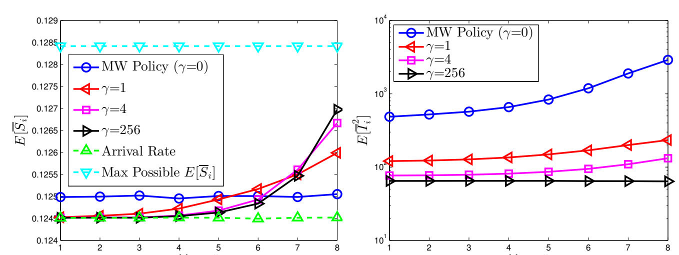

(a) (b)

Fig. 6. (a) The service rate distribution among different users. (b) $\mathbb{E}[\overline{I}_i^2]$ values for individual users.

Fig. 6(a) illustrates the service rate distribution among different users in the system. The MW policy, which guarantees only stability, allocates roughly the same service rate ($\approx \frac{1}{8}$) to each user, since the arrival rates are symmetric. Our policy intelligently adapts to the second moment of the arrival processes, and allocates more service to the users with higher second moment, while maintaining the stability of the system. This particular service rate distribution leads to the increasing in the total queue-length which is revealed by the queue-length upper bound given in Proposition 1, and also illustrated in the previous simulation results.

We also examine the expected second moment of the inter-service time for each individual user, which is shown in Fig. 6(b). As a result of the service rate distribution shown in Fig. 6(a), the users with higher second moment of the arrival process suffer in the performance of service regularity under the MW policy. Under our policy, the $\mathbb{E}[\overline{I}_i^2]$ value becomes close to each other and much smaller than that of the MW policy for all users, indicating a better service regularity for each user. This performance improves further when we increase the value of γ. It can be also observed that the MW policy has a much higher value (note the logarithmic scale on the y axis) of $\sum_{i=1}^{N}\mathbb{E}[\overline{I}_i^2]$ than our policy.

VI. Conclusion

In this work, we investigated the problem of designing a scheduling policy that is both throughput-optimal and possesses favorable service regularity characteristics. We introduced a new parameter of time-since-last-service, and proposed a novel scheduling policy that combines this parameter with the queue-lengths in its weight. After establishing the throughput optimality of our policy, we showed that it also has provable service regularity performance. In particular, the service regularity of our policy can be guaranteed to remain within a constant factor distance of a fundamental lower bound. We explicitly expressed this constant factor as a function of the mean arrival rates and the system parameters. We performed extensive numerical studies to illustrate the significant gains achieved by our policy over the traditional queue-length-based policies. Our results show the significance of utilizing the time-since-last-service in improving the service regularity performance of throughput-optimal policies.

In our future work, we will extend the analysis of our promising policy to more general channel fading and network models. Another interesting direction is the study of the heavy-traffic characteristics of this class of policies.

References

[1] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Transactions on Automatic Control*, vol. 36, pp. 1936–1948, December 1992.

[2] A. Eryilmaz and R. Srikant, "Fair resource allocation in wireless networks using queue-length based scheduling and congestion control," in *Proceedings of IEEE Infocom*, vol. 3, Miami, 2005, pp. 1794–1803.

[3] X. Lin and N. Shroff, "Joint rate control and scheduling in multihop wireless networks," in *Proceedings of IEEE Conference on Decision and Control*, Paradise Island, Bahamas, December 2004.

[4] M. Neely, E. Modiano, and C. Li, "Fairness and optimal stochastic control for heterogeneous networks," in *Proceedings of IEEE Infocom*, Miami, FL, March 2005, pp. 1723–1734.

[5] X. Lin, N. B. Shroff, and R. Srikant, "A tutorial on cross-layer optimization in wireless networks," *IEEE Journal on Selected Areas in Communications, special issue on Nonlinear Optimization of Communication Systems*, vol. 14, no. 8, Aug 2006.

[6] L. Bui, R. Srikant, and A. Stolyar, "Novel architectures and algorithms for delay reduction in back-pressure scheduling and routing," in *Proceedings of IEEE Infocom*, Rio de Janeiro, Brazil, 2009, pp. 2936–2940.

[7] L. Ying, S. Shakkottai, and A. Reddy, "On combining shortest-path and back-pressure routing over multihop wireless networks," in *Proceedings of IEEE Infocom*, Rio de Janeiro, Brazil, April 2009, pp. 1674–1682.

[8] H. Xiong, R. Li, A. Eryilmaz, and E. Ekici, "Delay-aware cross-layer design for network utility maximization in multi-hop networks," *Selected Areas in Communications, IEEE Journal on*, vol. 29, no. 5, pp. 951 – 959, May 2011.

[9] M. Neely, "Delay-based network utility maximization," in *INFOCOM, 2010 Proceedings IEEE*, 2010, pp. 1 –9.

[10] I.-H. Hou, V. Borkar, and P. R. Kumar, "A theory of qos for wireless," in *INFOCOM 2009, IEEE*, 2009, pp. 486 –494.

[11] I.-H. Hou and P. R. Kumar, "Admission control and scheduling for qos guarantees for variable-bit-rate applications on wireless channels," in *MobiHoc '09: Proceedings of the tenth ACM international symposium on Mobile ad hoc networking and computing*, 2009, pp. 175–184.

[12] ——, "Scheduling heterogeneous real-time traffic over fading wireless channels," in *INFOCOM, 2010 Proceedings IEEE*, 2010, pp. 1 –9.

[13] J. Jaramillo and R. Srikant, "Optimal scheduling for fair resource allocation in ad hoc networks with elastic and inelastic traffic," in *INFOCOM, 2010 Proceedings IEEE*, 2010, pp. 1 –9.

[14] A. Eryilmaz and R. Srikant, "Asymptotically tight steady-state queue length bounds implied by drift conditions," *to appear in QUESTA*, 2012. [Online]. Available: arxiv.org/abs/1104.0327

[15] R. Urgaonkar and M. Neely, "Opportunistic scheduling with reliability guarantees in cognitive radio networks," in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, 2008, pp. 1301 –1309.

[16] X. Li, Q. Zhao, X. Guan, and L. Tong, "Optimal cognitive access of markovian channels under tight collision constraints," *Selected Areas in Communications, IEEE Journal on*, vol. 29, no. 4, pp. 746 –756, 2011.

[17] R. Li, A. Eryilmaz, and B. Li, "Throughput-optimal scheduling with regulated inter-service times technical report." [Online]. Available: http://www2.ece.ohio-state.edu/~eryilmaz/RegulatedScheduling_report.pdf

[18] M. Neely, E. Modiano, and C. Rohrs, "Dynamic power allocation and routing for time varying wireless networks," in *Proceedings of IEEE Infocom*, April 2003, pp. 745–755.

[19] A. Eryilmaz and R. Srikant, "Joint congestion control, routing and mac for stability and fairness in wireless networks," *IEEE Journal on Selected Areas in Communications, special issue on Nonlinear Optimization of Communication Systems*, vol. 14, pp. 1514–1524, August 2006.

[20] G. Gupta and N. Shroff, "Delay analysis for multi-hop wireless networks," in *INFOCOM 2009. The 28th Conference on Computer Communications. IEEE*, April 2009, pp. 2356–2364.

[21] A. L. Stolyar, "Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic," *The Annals of Applied Probability*, vol. 14, no. 1, pp. pp. 1–53, 2004.

Appendix A
Proof of Lemma 2

Recall that our policy schedules the user $i^*[t] = \text{argmax}_i(\alpha_i Q_i[t] + \gamma\beta_i T_i[t])$ at each time slot t. Hence, if $Q_{i^*}[t] < r_{i^*}$, the following inequality holds:

$$\begin{aligned}&\alpha_i Q_i[t] \le \alpha_i Q_i[t] + \gamma\beta_i T_i[t]\\ &\quad\le \quad \alpha_{i^*} Q_{i^*}[t] + \gamma\beta_{i^*} T_{i^*}[t] \le \alpha_{i^*} r_{i^*} + \gamma\beta_{i^*} T_{i^*}[t],\end{aligned}$$

for all $i \neq i^*[t]$. Consequently, we have

$$\begin{aligned}Q_\Sigma[t] \quad &= \quad \sum_{i=1}^{N}\frac{Q_i[t]}{r_i} \le \sum_{i=1}^{N}\frac{\alpha_{i^*} r_{i^*} + \gamma\beta_{i^*} T_{i^*}[t]}{\alpha_i r_i}\\ &\le \quad \frac{N\max_i(\alpha_i r_i)}{\min_i(\alpha_i r_i)} + \frac{N\gamma\beta_{max}}{\min(\alpha_i r_i)}T_{i^*}[t],\end{aligned}$$

where $\beta_{max} = \max_i(\beta_i)$. Thus, we can write the following expectation:

$$\begin{aligned}&\mathbb{E}\left[Q_\Sigma[t]\mathbb{1}\{Q_{i^*}[t] < r_{i^*}\}\right]\\ &\le \mathbb{E}\left[\frac{N\gamma\beta_{max}}{\min(\alpha_i r_i)}T_{i^*}[t]\mathbb{1}\{Q_{i^*}[t] < r_{i^*}\}\right] + \frac{N\max_i(\alpha_i r_i)}{\min_i(\alpha_i r_i)}\\ &= \frac{N\gamma\beta_{max}}{\min(\alpha_i r_i)}\sum_{m=1}^{\infty} mP(T_{i^*}[t] = m, Q_{i^*}[t] < r_{i^*}) + B_0, \quad (11)\end{aligned}$$

where $B_0 \triangleq N\max_i(\alpha_i r_i)/\min_i(\alpha_i r_i)$.

By the definition of $T_{i^*}[t]$, $T_{i^*}[t] = m$ is equivalent to say that user i has not been served in the past m time-slots. Thus, we have the following relationship between the events:

$$\begin{aligned}
&\{T_{i^*}[t] = m, Q_{i^*}[t] = 0\} \\
&= \left\{ \begin{array}{l} S_{i^*}[t-m]C[t-m] = \cdots = S_{i^*}[t-1]C[t-1] = 0, \\ Q_{i^*}[t] < r_{i^*} \end{array} \right\} \\
&\overset{(a)}{=} \left\{ \begin{array}{l} S_{i^*}[t-m]C[t-m] = \cdots = S_{i^*}[t-1]C[t-1] = 0, \\ Q_{i^*}[t] < r_{i^*},\ \sum_{k=t-m}^{t-1} A_{i^*}[k] < r_{i^*} \end{array} \right\} \\
&\subseteq \left\{ \sum_{k=t-m}^{t-1} A_{i^*}[k] \leq r_{max} \right\}, \qquad (12)
\end{aligned}$$

where $r_{max} \triangleq \max_i(r_i)$, and (a) holds due to the nature of the queue evolution, that if there has been no service during the past m slots while $Q_{i^*}[t] < r_{i^*}$, then it must imply that the total arrivals during these m slots is less than r_{i^*}. As a consequence of (12), we have

$$P\{T_{i^*}[t] = m, Q_{i^*}[t] < r_{i^*}\} \leq P\left\{ \sum_{k=t-m}^{t-1} A_{i^*}[k] \leq r_{max} \right\}.$$

We can then bound the above probability by using a large deviation argument. Let $0 < a < \min_i \lambda_i$ be a constant. Choose and fix an M large enough such that for each i, we have $\frac{r_{max}}{M} - \lambda_i \leq -a$. Then for $m > M$, the following holds:

$$\begin{aligned}
&P\left\{ \sum_{k=t-m}^{t-1} A_{i^*}[k] \leq r_{max} \right\} \\
&\leq \quad P\left\{ \frac{1}{m} \sum_{k=t-m}^{t-1} A_{i^*}[k] - \lambda_{i^*} \leq a \right\} \\
&\leq \quad e^{-mI_{i^*}(a)} \leq e^{-mI_{min}(a)}, \qquad (13)
\end{aligned}$$

where $I_i(a)$ is the rate function for user i, given by

$$I_i(a) = \sup_\theta \left(\theta a - \log \mathbb{E}\left[e^{\theta A_i} \right] \right),$$

and $I_{min}(a) = \min_i I_i(a)$. Note that under our i.i.d and finite support assumption of the arrival process for each user, the rate function $I_i(a)$ is well defined for each user i.

Substitute (13) into (11), we get

$$\begin{aligned}
&\mathbb{E}\left[Q_\Sigma \mathbb{1}\{Q_{i^*}[t] < r_{i^*}\} \right] \\
&\leq \quad \frac{N\gamma\beta_{max}}{\min(\alpha_i r_i)} \sum_{m=1}^{M} mP(T_{i^*}[t] = m, Q_{i^*}[t] < r_{i^*}) \qquad (14) \\
&\quad + \frac{N\gamma\beta_{max}}{\min(\alpha_i r_i)} \sum_{m=M+1}^{\infty} me^{-mI_{min}(a)} + B_0 \qquad (15) \\
&\leq \quad B_1(\gamma),
\end{aligned}$$

where (14) is a finite sum of finite terms for all user i, and (15) is finite due to the fact that $\sum_{m=1}^{\infty} mx^m = \frac{x}{(1-x)^2},\ \forall |x| < 1$.

Appendix B
Proof of Proposition 1

We look at the Lyapunov function defined by

$$W(\boldsymbol{Q}[t], \boldsymbol{T}[t]) \triangleq \frac{1}{2} Q_\Sigma^2[t],$$

and study its mean drift given by

$$\begin{aligned}
&\triangle W(\boldsymbol{Q}[t], \boldsymbol{T}[t]) \\
&= \mathbb{E}\left[W(\boldsymbol{Q}[t+1], \boldsymbol{T}[t+1]) - W(\boldsymbol{Q}[t], \boldsymbol{T}[t]) | \boldsymbol{Q}[t], \boldsymbol{T}[t] \right] \\
&= \mathbb{E}\Big[\frac{Q_\Sigma^2[t+1] - Q_\Sigma^2[t]}{2} \mathbb{1}\{Q_{i^*}[t] < r_{i^*}\} | \boldsymbol{Q}[t], \boldsymbol{T}[t] \Big] \qquad (16) \\
&\quad + \mathbb{E}\Big[\frac{Q_\Sigma^2[t+1] - Q_\Sigma^2[t]}{2} \mathbb{1}\{Q_{i^*}[t] \geq r_{i^*}\} | \boldsymbol{Q}[t], \boldsymbol{T}[t] \Big] \qquad (17)
\end{aligned}$$

By substitute (5) into (16), and omit the time index t for brevity, we have

$$\begin{aligned}
(16) \quad &= \quad \mathbb{E}\left[(Q_\Sigma A_\Sigma + \frac{1}{2} A_\Sigma^2) \mathbb{1}\{Q_{i^*} < r_{i^*}\} \middle| \boldsymbol{Q}, \boldsymbol{T} \right] \\
&\leq \quad Q_\Sigma (1 - \epsilon) \mathbb{1}\{Q_{i^*} < r_{i^*}\} + \frac{N^2 A_{max}^2}{2r_{min}^2}, \qquad (18)
\end{aligned}$$

where the last inequality above holds because we assume the arrival process is supportable and has a finite support, thus $\mathbb{E}[A_\Sigma] = \sum_{i=1}^{N} \frac{\lambda_i}{r_i} = p - \epsilon$, and $\mathbb{E}[A_\Sigma^2] \leq \frac{N^2 A_{max}^2}{r_{min}^2}$. Similarly, we can write (17) as:

$$\begin{aligned}
(17) &= \mathbb{E}\left[(Q_\Sigma (A_\Sigma - C) + \frac{1}{2} (A_\Sigma - C)^2) \mathbb{1}\{Q_{i^*} \geq r_{i^*}\} \middle| \boldsymbol{Q}, \boldsymbol{T} \right] \\
&\overset{(a)}{\leq} - \epsilon Q_\Sigma \mathbb{1}\{Q_{i^*} \geq r_{i^*}\} + \frac{N^2 A_{max}^2}{2r_{min}^2} + \frac{p}{2} \\
&= - \epsilon Q_\Sigma (1 - \mathbb{1}\{Q_{i^*} < r_{i^*}\}) + \frac{N^2 A_{max}^2}{2r_{min}^2} + \frac{p}{2} \\
&= - \epsilon Q_\Sigma + \epsilon Q_\Sigma \mathbb{1}\{Q_{i^*} < r_{i^*}\} + \frac{N^2 A_{max}^2}{2r_{min}^2} + \frac{p}{2}, \qquad (19)
\end{aligned}$$

where (a) holds because the channel condition $C[t]$ is independent of $(\boldsymbol{Q}[t], \boldsymbol{T}[t])$, thus it can be pulled out of the expectation directly.

Combine (18) and (19), we get

$$\triangle W(\boldsymbol{Q}[t], \boldsymbol{T}[t]) \leq Q_\Sigma[t] \mathbb{1}\{Q_{i^*}[t] < r_{i^*}\} - \epsilon Q_\Sigma[t] + B_2,$$

where $B_2 = \frac{N^2 A_{max}^2}{r_{min}^2} + \frac{p}{2}$ is a finite constant. Taking expectation over $(\boldsymbol{Q}[t], \boldsymbol{T}[t])$ on both sides yields:

$$\begin{aligned}
&\frac{1}{2} \mathbb{E}\left[Q_\Sigma^2[t+1] - Q_\Sigma^2[t] \right] \\
&\quad \leq \quad \mathbb{E}\left[Q_\Sigma[t] \mathbb{1}\{Q_{i^*}[t] < r_{i^*}\} \right] - \epsilon \mathbb{E}[Q_\Sigma[t]] + B_2 \\
&\quad \overset{(a)}{\leq} \quad -\epsilon \mathbb{E}[Q_\Sigma[t]] + B_1(\gamma) + B_2,
\end{aligned}$$

where (a) holds due to Lemma 2. Sum the above expectation of the drift over $t = 0$ through $K - 1$, we get

$$\frac{\mathbb{E}\left[Q_\Sigma^2[K] - Q_\Sigma^2[0] \right]}{2} \leq -\epsilon \sum_{t=0}^{K-1} \mathbb{E}[Q_\Sigma[t]] + K\left(B_1(\gamma) + B_2 \right),$$

which can be simplified to

$$\frac{\epsilon}{K} \sum_{t=0}^{K-1} \mathbb{E}[Q_\Sigma[t]] \leq \frac{\mathbb{E}\left[Q_\Sigma^2[0] \right]}{2K} + B_1(\gamma) + B_2.$$

Take the limit as K goes to infinity yields:

$$\limsup_{K \to \infty} \frac{1}{K} \sum_{t=0}^{K-1} \mathbb{E}[Q_\Sigma[t]] \leq \frac{B_1(\gamma) + B_2}{\epsilon}.$$