On the Modeling and Optimization of Short-Term Performance for Real-Time Wireless Networks

I-Hong Hou
Department of ECE
Texas A&M University
College Station, TX 77843, USA
Email: ihou@tamu.edu

***Abstract*—This paper studies wireless networks consisting of multiple real-time flows that impose hard delay bounds for all packets. In contrast to most current studies that focus on the long-term average rate of timely deliveries, we aim to model and optimize the short-term performance of each real-time flow, which is vital for most safety-critical applications. We propose to define the instantaneous performance of a flow by a moving average within a short window in the past. Each flow incurs some penalty if its moving average is below some specified requirement, and we aim to minimize the overall penalty of the system.**

We approximate the system by Brownian motions, and formulate an optimization problem for minimizing the overall penalty. While the optimization problem is not convex, we establish a low-complexity algorithm for optimally solving it by leveraging inherent structures of real-time wireless networks. We also propose a simple online packet scheduling policy and prove that it achieves the minimum overall penalty. Simulation results show that Brownian approximations are accurate in capturing short-term performance, and our policies achieve much better performance than other policies. Moreover, they also demonstrate that policies with optimal long-term average delivery rates can actually have poor short-term performance.

I. Introduction

Wireless networks will be increasingly employed in safety-critical systems, such as smartgrid communications, smart manufacturing systems, and other cyber-physical systems. In order to ensure safety, these systems usually require stringent delay bounds for all packets. Therefore, most existing studies on wireless networks that aim to optimize traditional quality of service (QoS) metrics, such as throughput, average delay, delay jitters, etc., can be insufficient to meet the performance requirements of these systems.

Recently, there have been some advancements in real-time wireless networks. In real-time wireless networks, a strict deadline is imposed for each packet. Packets that fail to be delivered before their deadlines are dropped, and packet drops cause glitches in the corresponding application. Most current studies model the performance of real-time wireless networks by the timely-throughput, or, equivalently, the long-term average packet delivery ratio, of each flow. While the study of timely-throughput reveals many important features of real-time wireless networks, it remains insufficient to fully capture the behavior of safety-critical systems. A major reason is that timely-throughput only focuses on long-term average performance and ignores short-term fluctuations. For example, consider the following two scenarios for the same flow: In the first scenario, the first packet is dropped, the next 9 packets are delivered, the 11-th packet is dropped, the next 9 packets are delivered, and so on. In the second scenario, the first 1000 packets are dropeed, the next 9000 packets are delivered, and so on. Both scenarios deliver 90% of packets on time. However, it is clear that the first scenario is more desirable in most practical applications.

In this paper, we propose a different approach to model the performance of real-time wireless networks. In this approach, the instantaneous performance of a flow only depends on a short time horizon in the past. Therefore, events occurred prior to the horizon have no influence on the instantaneous performance. At any point of time, a flow suffers from some penalty if its instantaneous performance fails to meet some specified requirement. We then aim to minimize the long-term average total penalty over all flows. This approach captures system performance on all time scales: the per-packet deadline is on the order of milliseconds, the time horizon can be chosen around hundreds of milliseconds, and the total penalty can be calculated on a per-second basis.

By employing Brownian motion approximation, we formulate an optimization problem that characterizes a lower-bound for the long-term average total penalty. While this optimization problem is not convex and may appear to be complicated, we establish a simple algorithm that finds its optimal solution by leveraging some inherent characteristics of real-time wireless networks. Further, we propose a simple online scheduling policy and prove that it achieves the optimal solution, and thus minimizes the total penalty.

Our theoretical analysis is evaluated by extensive simu-

This material is based upon work supported by, or in part by, the U. S. Army Research Laboratory and the U. S. Army Research Office under contract/grant number W911NF-15-1-0279.

lations. We first demonstrate that Brownian motion is indeed an accurate approximation to the short-term performance of real-time wireless networks. We then compare our online scheduling policy against other policies, including one that has been shown to achieve the optimal performance in terms of long-term timely-throughput. Simulation results show that our policy achieves significantly smaller total penalty. In addition to showing that our policy is indeed optimal, these results also demonstrate that policies focusing on long-term timely-throughput can have poor short-term performance.

The rest of the paper is organized as follows: Section II reviews some existing studies on queueing networks and real-time wireless networks. Section III presents our model for studying the short-term performance of real-time wireless networks. Section IV employs Brownian approximation to approximate the system and formulate an optimization problem that characterize a lower-bound of total penalty. Section V derives an algorithm that solves the optimization problem and obtains the lower-bound. Section VI proposes a simple online scheduling policy that minimizes total penalty. Section VII presents our simulation results. Finally, Section VIII concludes the paper.

II. Related Work

Quantifying and optimizing short-term performance for communication networks have been widely studied. Typically, these studies focus on studying the distribution or moments of total queue lengths in the heavy-traffic regime, as they are correlated to packet delays and delay jitters. Laws [1] employs Brownian approximation to derive a dynamic routing algorithm for reducing total queue lengths. Kang et al. [2] studies the diffusion approximation of queue lengths for document transfers in networks with proportional fair policy. Bhardwaj and Williams [3] studies diffusion approximation when base stations can cooperate to form a MIMO system. Destounis et al. [4] proposes a power allocation strategy for dense networks to limit the probability that a queue length exceeds some threshold. Verloop and Núnez-Queija [5] studies the optimal resource allocation for a two-class network. Eryilmaz and Srikant [6] derives an asymptotically tight bound for queue length as the system approaches the heavy-traffic regime. Wang et al. [7] studies heavy-traffic-optimal policies for MapReduce. Li, Li, and Eryilmaz [8] proposes a scheduling policy that provides both service regularity and heavy-traffic-optimality. These studies lack hard per-packet delay guarantee, which is essential for many real-time applications.

Hou, Borkar, and Kumar [9] establishes a framework for studying real-time wireless networks. This framework enforces a hard deadline for each packet by dropping packets that exceed their deadlines. This framework has been extended in several directions, including systems with both real-time and non-real-time flows [10], multi-hop networks [11], systems with heterogeneous deadline constraints [12], [13], and fading wireless channels [14], [15]. While all these studies address hard per-packet deadlines, they measure the per-flow performance by the long-term average delivery ratio, and therefore neglect short-term performance fluctuations. Singh, Hou, and Kumar [16] studies the asymptotic behavior of performance fluctuations by employing law of iterated logarithm.

III. System Model

We extend the model in [9] to account for the short-term performance of real-time flows. Consider a system where one AP serves $\mathbb{J}$ clients, each with one real-time flow. Time is slotted and the AP can schedule at most one transmission in each time slot. Time slots are further grouped into *intervals* with length τ slots. At the beginning of each interval, each real-time flow generates one packet. Each flow requires a stringent per-packet deadline of τ slots, i.e., packets generated at the beginning of an interval need to be delivered by the end of that interval. Packets that are not delivered on time are dropped from the system. We further consider the unreliable wireless transmissions by assuming that each transmission is successful with probability p, and the AP has instant knowledge on whether a transmission is successful. We note that, in contrast to [9], we assume that transmissions to all clients have the same success probability and do not consider the different channel qualities of different clients. Such an assumption holds for many practical systems. For example, we can consider a system where one wireless link carries multiple real-time flows. We then treat each flow as one wireless client, and all transmissions have the same success probability as they actually use the same link. We can also consider wireless systems where the primary cause of failed transmission is the interference from nearby networks, rather than the shadowing/fading of wireless channels. As all clients experience similar interference, they have similar success probability.

Let $X_i(t)$ be the number of packets delivered for client i in the first t intervals, or, equivalently, in the first $t\tau$ slots. We then have $X_i(t) - X_i(t-1) = 1$, if there is a packet delivered for client i in the t-th interval, and $X_i(t) - X_i(t-1) = 0$, otherwise.

We model the short-term performance of a client as follows: At the end of an interval t, each client counts the number of packets it receives in the last $\mathbb{T}$ intervals, which can be expressed as $X_i(t) - X_i(t-\mathbb{T})$. Each client i requires to receive $q_i\mathbb{T}$ packets in every $\mathbb{T}$ consecutive intervals, that is, $X_i(t) - X_i(t-\mathbb{T}) \geq q_i\mathbb{T}$, for all t. If $X_i(t) - X_i(t-\mathbb{T}) < q_i\mathbb{T}$, then the system suffers from a penalty of $C(q_i\mathbb{T} - [X_i(t) - X_i(t-\mathbb{T})])$, where we assume that $C(x)$ is a strictly increasing, strictly convex, and differentiable function over $x \in [0, +\infty)$ with $C(0) = 0$ and $C'(0) = 0$. To simplify notation, we define $C(x) = 0$, for all $x \leq 0$. Examples of $C(x)$ include $C(x) = x^\alpha$, if $x \geq 0$, for any $\alpha > 1$, and $C(x) = e^x - x - 1$, if $x \geq 0$.

We aim to design scheduling policies that minimizes the long-term average penalty incurred by all clients, which can be expressed as $\limsup_{T\to\infty} \frac{\sum_{t=1}^{T}\sum_i C(q_i\mathbb{T}-[X_i(t)-X_i(t-\mathbb{T})])}{T}$. We note that our model covers the whole spectrum of different time scales: The per-packet deadline is usually chosen to be on the order of several ms. By choosing $\mathbb{T}$ to be on the order of hundreds, the short-term performance is measured over a window of hundreds of ms. Finally, the long-term average penalty is defined over the slowest time scale, and is calculated every second in practice.

IV. An Optimization Formulation for the Penalty Lower-Bound

In this section and the next, we derive a lower bound of the long-term average total penalty. We will further establish an online scheduling policy that achieves the lower bound in Section VI.

Apparently, the vector $\{X_i(t)-X_i(t-1)\}$[1] is a random variable whose distribution is determined by the employed scheduling policy. To simplify discussions, we will focus on *ergodic* scheduling policies under which $\{X_i(t)-X_i(t-1)\}$ can be modeled by a positive recurrent Markov chain. Then, by the law of large numbers and central limit theorem of Markov chains [17], $\bar{X}_i := \lim_{T\to\infty}\frac{X_i(T)}{T}$ exists and is positive, and $\hat{X}_i(t) := \lim_{T\to\infty}\frac{X_i(tT)-tT\bar{X}_i}{\sqrt{T}}$ is a driftless Brownian motion with variance σ_i^2, for some $\sigma_i \geq 0$. In other words, $\hat{X}_i(t_1)-\hat{X}_i(t_2)$ is a Gaussian random variable with mean 0 and variance $\sigma_i^2(t_1-t_2)$, for any $t_1 > t_2$. Let $\Phi(x)$ be the cumulative distribution function (CDF) of a normal random variable. The CDF of $\hat{X}_i(t_1)-\hat{X}_i(t_2)$ is then $\Phi(\frac{x}{\sqrt{\sigma_i^2(t_1-t_2)}})$.

Many studies have established that $\hat{X}_i(t)$ is a good approximation to $X_i(t)-t\bar{X}_i$ under various mild assumptions. Therefore, we will assume that $X_i(t)-t\bar{X}_i$ has the same distribution as $\hat{X}_i(t)$ for sufficiently large t throughout this paper. In Section VII, we will demonstrate by simulations that $\hat{X}_i(t)$ is indeed a good approximation to $X_i(t)-t\bar{X}_i$ under our proposed online policy.

With the assumption of approximation, we can rewrite the long-term average total penalty as follows:

$$\begin{aligned}
&\limsup_{T\to\infty}\frac{\sum_{t=1}^{T}\sum_i C(q_i\mathbb{T}-[X_i(t)-X_i(t-\mathbb{T})])}{T}\\
=&\lim_{T\to\infty}\sum_i E[C(q_i\mathbb{T}-[X_i(T)-X_i(T-\mathbb{T})])]\\
\approx&\sum_i E[C((q_i-\bar{X}_i)\mathbb{T}-\hat{X}_i(\mathbb{T})]\\
=&\sum_i\int_z C(\sqrt{\sigma_i^2\mathbb{T}}z-(\bar{X}_i-q_i)\mathbb{T})d\Phi(z). \qquad (1)
\end{aligned}$$

The last equation holds because $\hat{X}_i(\cdot)$ has the same distribution as $-\hat{X}_i(\cdot)$.

[1] In this paper, we use $\{x_i\}$ to indicate the vector consists of x_i for all $1 \leq i \leq \mathbb{J}$.

With this approximation, minimizing the long-term average total penalty becomes finding $\{\bar{X}_i\}$ and $\{\sigma_i\}$ that minimizes (1), with the condition that $\{\bar{X}_i\}$ and $\{\sigma_i\}$ can be achieved by some scheduling policy. Next, we will establish some constraints on $\{\bar{X}_i\}$ and $\{\sigma_i\}$.

Let $\mathcal{F}(n) := E[\min\{\tau, \sum_{m=1}^{n}\gamma_m\}]$, where $\gamma_1, \gamma_2, \dots$ is a sequence of independent geometric random variables with mean $\frac{1}{p}$. Recall that τ is the number of slots in an interval. [9] has established the following condition for $\{\bar{X}_i\}$:

Theorem 1: There exists a policy that achieves $\{\bar{X}_i\}$ if and only if $\sum_{i\in S}\bar{X}_i \leq p\mathcal{F}(|S|)$, for all subsets S of clients. Further, $\sum_{i=1}^{\mathbb{J}}\bar{X}_i = p\mathcal{F}(\mathbb{J})$ under all *work-conserving* policies, which always schedule transmissions as long as there are packets available to transmit. □

Since a policy cannot increase its penalties by making more transmissions, we can focus on work-conserving policies and assume that $\sum_{i=1}^{\mathbb{J}}\bar{X}_i = p\mathcal{F}(\mathbb{J})$. Given $\{\bar{X}_i\}$ that satisfies the constraints in Theorem 1, we say that a subset S is *tight* if $\sum_{i\in S}\bar{X}_i = p\mathcal{F}(|S|)$. The lemma below shows that there is a nested structure among tight subsets.

Lemma 1: If S_1 and S_2 are two tight subsets, then either $S_1 \subseteq S_2$, or $S_2 \subseteq S_1$.

Proof: With the definition of $\mathcal{F}(\cdot)$, it is easy to verify that, for any positive integers a, b and non-negative integer c, $\mathcal{F}(a+c)-\mathcal{F}(c) \geq \mathcal{F}(a+b+c)-\mathcal{F}(b+c)$, with equality holds if and only if $\mathcal{F}(a+c)-\mathcal{F}(c) = 0$.

We prove this lemma by contradiction. Suppose the lemma does not hold, and therefore $\Delta_1 := |S_1\backslash S_2| > 0$ and $\Delta_2 := |S_2\backslash S_1| > 0$. We also define $\wedge := |S_1\cap S_2| \geq 0$. Since S_1 and S_2 are tight, we have $\sum_{i\in S_1}\bar{X}_i = p\mathcal{F}(\Delta_1+\wedge)$, $\sum_{i\in S_2}\bar{X}_i = p\mathcal{F}(\Delta_2+\wedge)$, and therefore $\sum_{i\in S_1\cup S_2}\bar{X}_i = p\mathcal{F}(\Delta_1+\wedge)+p\mathcal{F}(\Delta_2+\wedge)-\sum_{i\in S_1\cap S_2}\bar{X}_i$.

Suppose $S_1\cap S_2$ is not tight, that is, $\sum_{i\in S_1\cap S_2}\bar{X}_i < p\mathcal{F}(\wedge)$. We then have

$$\begin{aligned}
\sum_{i\in S_1\cup S_2}\bar{X}_i &> p\mathcal{F}(\Delta_1+\wedge)+p\mathcal{F}(\Delta_2+\wedge)-p\mathcal{F}(\wedge)\\
&\geq p\mathcal{F}(\Delta_1+\Delta_2+\wedge) = p\mathcal{F}(|S_1\cup S_2|),
\end{aligned}$$

which violates the constraints in Theorem 1.

On the other hand, if $\sum_{i\in S_1\cap S_2}\bar{X}_i = p\mathcal{F}(\wedge)$, then $\sum_{i\in S_1\backslash S_2}\bar{X}_i = p\mathcal{F}(\Delta_1+\wedge)-p\mathcal{F}(\wedge) > 0$, and we have

$$\begin{aligned}
\sum_{i\in S_1\cup S_2}\bar{X}_i &= p\mathcal{F}(\Delta_1+\wedge)+p\mathcal{F}(\Delta_2+\wedge)-p\mathcal{F}(\wedge)\\
&> p\mathcal{F}(\Delta_1+\Delta_2+\wedge) = p\mathcal{F}(|S_1\cup S_2|),
\end{aligned}$$

which also violates the constraints in Theorem 1, and the proof is complete. ■

Based on this lemma, we can number tight subsets such that $\phi = S_0 \subset S_1 \subset S_2 \subset \dots \subset S_K = \{1, 2, \dots, \mathbb{J}\}$.

We now turn to the constraints of $\{\sigma_i\}$. Let $R_n(t)$ be the indicator function that at least n packets are delivered in the t-th interval. The probability distribution of $R_n(t)$ is the same under all work-conserving policies, and $R_n(t)$ is independent across t. Using a similar argument as that used in [9] to establish Theorem 1, we can show that

$R_n(t)$ is a Bernouli random variable with mean $\bar{R}_n := p\mathcal{F}(n)-p\mathcal{F}(n-1)$. By the functional central limit theorem, we have $\hat{R}_n(k) := \lim_{T\to\infty}\frac{\sum_{t=1}^{kT}R_n(t)-kT\bar{R}_n}{\sqrt{T}}$ is a driftless Brownian motion with variance $\bar{R}_n(1-\bar{R}_n)$. Further, for any $n_1 < n_2$, we have $E[R_{n_1}(t)R_{n_2}(t)] = E[R_{n_2}(t)] = \bar{R}_{n_2}$, since it is not possible to deliver n_2 packets without delivering n_1 packets first. We then have

$$\begin{aligned}
&\upsilon^2_{n_1,n_2} := E[(\sum_{m=n_1}^{n_2} R_m(t))^2] - E[\sum_{m=n_1}^{n_2} R_m(t)]^2\\
&=\sum_{m=n_1}^{n_2} E[R_m(t)^2] + 2\sum_{m_1=n_1}^{n_2}\sum_{m_2=m_1+1}^{n_2} E[R_{m_1}(t)R_{m_2}(t)]\\
&\quad-\sum_{m=n_1}^{n_2}\bar{R}_m^2 + 2\sum_{m_1=n_1}^{n_2}\sum_{m_2=m_1+1}^{n_2}\bar{R}_{m_1}\bar{R}_{m_2}\\
&=\sum_{m=n_1}^{n_2}\bar{R}_m(1-\bar{R}_m) + 2\sum_{n_1\le m_1<m_2\le n_2}\bar{R}_{m_2}(1-\bar{R}_{m_1}).
\end{aligned} \tag{2}$$

Therefore, $\sum_{m=n_1}^{n_2}\hat{R}_m(t)$ is a driftless Brownian motion with variance $\upsilon^2_{n_1,n_2}$.

Consider a tight set S_k with $0 \le k \le K$. Since there are only $|S_k|$ clients in the subset S_k, the number of packets delivered to clients in S_k in the t-th interval is no larger than $\sum_{m=1}^{|S_k|}R_m(t)$. Therefore, we have $\sum_{t=1}^{T}\sum_{m=1}^{|S_k|}R_m(t) \ge \sum_{i\in S_k}X_i(T)$ for all T, on every sample path. Further, since $\sum_{m=1}^{|S_k|}\bar{R}_m = \sum_{i\in S_k}\bar{X}_i = p\mathcal{F}(|S_k|)$, $\sum_{m=1}^{|S_k|}\hat{R}_m(t) \ge \sum_{i\in S_k}\hat{X}_i(t)$ for all t, on every sample path. On the other hand, we have $E[\sum_{m=1}^{|S_k|}\hat{R}_m(t)] = E[\sum_{i\in S_k}\hat{X}_i(t)] = 0$, for all t. Therefore, $\sum_{m=1}^{|S_k|}\hat{R}_m(t) = \sum_{i\in S_k}\hat{X}_i(t)$ almost surely, for all k, which also implies that $\sum_{m=|S_{k-1}|+1}^{|S_k|}\hat{R}_m(t) = \sum_{i\in S_k\backslash S_{k-1}}\hat{X}_i(t)$, for all $k\ge 1$. We now have

$$\begin{aligned}
\upsilon^2_{|S_{k-1}|+1,|S_k|} &= E[(\sum_{m=|S_{k-1}|+1}^{|S_k|}\hat{R}_m(1))^2]\\
&= E[(\sum_{i\in S_k\backslash S_{k-1}}\hat{X}_i(t))^2] \le (\sum_{i\in S_k\backslash S_{k-1}}\sigma_i)^2,
\end{aligned}$$

and $\sum_{i\in S_k\backslash S_{k-1}}\sigma_i \ge \upsilon_{|S_{k-1}|+1,|S_k|}$, for $1\le k\le K$.

Based on the above analysis, the problem of minimizing the long-term average total penalty can be formulated as the following optimization problem:

LOWER-BOUND:

$$\text{Min } L := \sum_i\int_z C(\sqrt{\sigma_i^2\mathbb{T}}z - (\bar{X}_i - q_i)\mathbb{T})d\Phi(z), \tag{3}$$

$$\text{s.t. } \sum_{i\in S}\bar{X}_i = p\mathcal{F}(|S|), \forall S\in\{S_1,S_2,\dots,S_K\}, \tag{4}$$

$$\sum_{i\in S}\bar{X}_i < p\mathcal{F}(|S|), \forall S\notin\{S_1,S_2,\dots,S_K\}, \tag{5}$$

$$\sum_{i\in S_k\backslash S_{k-1}}\sigma_i \ge \upsilon_{|S_{k-1}|+1,|S_k|}, \forall 1\le k\le K, \tag{6}$$

$$S_1\subset S_2\subset\cdots\subset S_K = \{1,2\dots,\mathbb{J}\}. \tag{7}$$

This optimization problem consists of three parts: tight sets $\{S_1,S_2,\dots,S_K\}$, $[\bar{X}_i]$, and $[\sigma_i]$. We note that the constraints (4)–(7) are necessary, but may not be sufficient. In other words, it may be possible to find tight sets, $[\bar{X}_i]$, and $[\sigma_i]$ that satisfy the constraints, but cannot be achieved by any scheduling policy. Therefore, this optimization problem characterizes a lower bound on average total penalty. We also note that, since constraints (4)–(6) depend on the choice of tight sets, this problem is not a convex optimization problem.

V. Deriving the Penalty Lower-Bound

A. Necessary Conditions for the Optimal Solution

We derive the solution for the optimization problem (3)–(7) in this section. We first establish some conditions for the optimal solution.

Theorem 2: Assume that $\{S_1^*,S_2^*,\dots\}$, $[\bar{X}_i^*]$, and $[\sigma_i^*]$ form an optimal solution to (3)–(7). Consider two clients i and j that belong to the same tight sets $\{S_1^*,\dots,S_k^*\}$, that is, $i,j\in S_k^*\backslash S_{k-1}^*$. We then have $\bar{X}_i^*-q_i = \bar{X}_j^*-q_j$, and $\sigma_i^*=\sigma_j^*$. Further, if a client i belongs to $S_k^*\backslash S_{k-1}^*$, then $\bar{X}_i^* = q_i + \frac{p\mathcal{F}(|S_k^*|)-p\mathcal{F}(|S_{k-1}^*|)+\sum_{i\in S_k^*\backslash S_{k-1}^*}q_i}{|S_k^*\backslash S_{k-1}^*|}$ and $\sigma_i^* = \frac{\upsilon_{|S_{k-1}^*+1|,|S_k|^*}}{|S_k^*\backslash S_{k-1}^*|}$. □

We introduce a lemma that is needed to prove Theorem 2.

Lemma 2: Suppose there exists a_1,a_2,b_1,b_2 such that $a_1,a_2>0$, $\int_z C'(a_1z-b_1)d\Phi(z) = \int_z C'(a_2z-b_2)d\Phi(z)$, and $\int_z zC'(a_1z-b_1)d\Phi(z) = \int_z zC'(a_2z-b_2)d\Phi(z)$. Then, $a_1=a_2$ and $b_1=b_2$.

Proof: Since $C(z)$ is a convex function and is strictly convex over $z>0$, $C'(z)$ is a non-decreasing function and is strictly increasing over $z>0$. We establish this lemma by considering two cases: $a_1=a_2$ and $a_1\ne a_2$.

If $a_1=a_2$, and $b_1\ne b_2$. Without loss of generality, assume that $b_1>b_2$. Then $a_1z-b_1 < a_2z-b_2$, for all z. Since $\frac{d\Phi(z)}{dz}>0$ for all z, we have $\int_z C'(a_1z-b_1)d\Phi(z) < \int_z C'(a_2z-b_2)d\Phi(z)$, which contradicts with the assumption that $\int_z C'(a_1z-b_1)d\Phi(z) = \int_z C'(a_2z-b_2)d\Phi(z)$. Therefore, if $a_1=a_2$, we have $b_1=b_2$.

Next, consider that $a_1\ne a_2$. Without loss of generality, assume that $a_1>a_2$. Then, for any b_1,b_2, there exists some z_0 such that $a_1z_0-b_1 = a_2z_0-b_2$. We also have $a_1z-b_1 > a_2z-b_2$ for all $z>z_0$, and $a_1z-b_1 < a_2z-b_2$

for all $z < z_0$. Since $\int_z C'(a_1 z - b_1)d\Phi(z) = \int_z C'(a_2 z - b_2)d\Phi(z)$, we have

$$\int_{z=z_0}^{\infty} (C'(a_1 z - b_1) - C'(a_2 z - b_2))d\Phi(z)$$
$$= \int_{z=-\infty}^{z_0} (C'(a_2 z - b_2) - C'(a_1 z - b_1))d\Phi(z) > 0,$$

and therefore

$$\int_{z=z_0}^{\infty} z(C'(a_1 z - b_1) - C'(a_2 z - b_2))d\Phi(z)$$
$$> \int_{z=z_0}^{\infty} z_0(C'(a_1 z - b_1) - C'(a_2 z - b_2))d\Phi(z)$$
$$= \int_{z=-\infty}^{z_0} z_0(C'(a_2 z - b_2) - C'(a_1 z - b_1))d\Phi(z)$$
$$\geq \int_{z=-\infty}^{z_0} z(C'(a_2 z - b_2) - C'(a_1 z - b_1))d\Phi(z),$$

which contradicts with the assumption that $\int_z zC'(a_1 z - b_1)d\Phi(z) = \int_z zC'(a_2 z - b_2)d\Phi(z)$. ■

We are now ready to prove Theorem 2.

Proof of Theorem 2: We first show that, for any two clients $i, j \in S_k^* \backslash S_{k-1}^*$, $\frac{\partial L}{\partial \bar{X}_i}\Big|_{\bar{X}_i = \bar{X}_i^*} = \frac{\partial L}{\partial \bar{X}_j}\Big|_{\bar{X}_j = \bar{X}_j^*}$ and $\frac{\partial L}{\partial \sigma_i}\Big|_{\sigma_i = \sigma_i^*} = \frac{\partial L}{\partial \sigma_j}\Big|_{\sigma_j = \sigma_j^*}$ by contradiction. Suppose $\frac{\partial L}{\partial \bar{X}_i}\Big|_{\bar{X}_i = \bar{X}_i^*} < \frac{\partial L}{\partial \bar{X}_j}\Big|_{\bar{X}_j = \bar{X}_j^*}$. Set $\bar{X}_i = \bar{X}_i^* + \epsilon$, $\bar{X}_j = \bar{X}_j^* - \epsilon$, and $\bar{X}_k = \bar{X}_k^*$, for all $k \neq i, j$. Since any set that contains i and not j is not tight, there exists a sufficiently small ϵ such that the resulting $[\bar{X}_i]$ still satisfies constraints (4) and (5), and achieves smaller L. Similarly, if $\frac{\partial L}{\partial \sigma_i}\Big|_{\sigma_i = \sigma_i^*} < \frac{\partial L}{\partial \sigma_j}\Big|_{\sigma_j = \sigma_j^*}$, we can choose a sufficiently small ϵ such that increasing σ_i^* by ϵ and decreasing σ_j^* by ϵ result in a smaller L.

We now have

$$\frac{\partial L}{\partial \bar{X}_i}\Big|_{\bar{X}_i = \bar{X}_i^*} = \int_z -\mathbb{T}C'(\sigma_i^* \sqrt{\mathbb{T}} z - (\bar{X}_i^* - q_i)\mathbb{T})d\Phi(z)$$
$$= \int_z -\mathbb{T}C'(\sigma_j^* \sqrt{\mathbb{T}} z - (\bar{X}_j^* - q_j)\mathbb{T})d\Phi(z) = \frac{\partial L}{\partial \bar{X}_j}\Big|_{\bar{X}_j = \bar{X}_j^*},$$

and

$$\frac{\partial L}{\partial \sigma_i}\Big|_{\sigma_i = \sigma_i^*} = \int_z \sqrt{\mathbb{T}} z C'(\sigma_i^* \sqrt{\mathbb{T}} z - (\bar{X}_i^* - q_i)\mathbb{T})d\Phi(z)$$
$$= \int_z \sqrt{\mathbb{T}} z C'(\sigma_j^* \sqrt{\mathbb{T}} z - (\bar{X}_j^* - q_j)\mathbb{T})d\Phi(z) = \frac{\partial L}{\partial \sigma_j}\Big|_{\sigma_j = \sigma_j^*}.$$

The first part of the theorem then follows by Lemma 2. The second part is a direct result of the first condition and the fact that L is increasing with $[\sigma_i]$. ■

Recall that σ_i^2 is the variance of $\hat{X}_i(\cdot)$, and that L is the expected value of $\sum_i C(\hat{X}_i(\mathbb{T}) - (\bar{X}_i - q_i)\mathbb{T})$. Theorem 2 then implies the following corollary:

Corollary 1: Given an optimal tight sets $\{S_1^*, S_2^*, \dots\}$, L is minimized by choosing $\bar{X}_i$ as in Theorem 2 and $\hat{X}_i(t) = \frac{\sum_{m=|S_{k-1}|+1}^{S_k} \hat{R}_m(t)}{|S_k \backslash S_{k-1}|}$, if $i \in S_k \backslash S_{k-1}$. □

Given an optimal tight sets $\{S_1^*, S_2^*, \dots\}$, (3)–(7) can be solved by Corollary 1. Next, we discuss some necessary conditions for the optimal tight sets $\{S_1^*, S_2^*, \dots\}$.

Lemma 3: Given an optimal solution to (3)–(7), for any two clients $i \in S_{k-1}^* \backslash S_{k-2}^*$ and $j \in S_k^* \backslash S_{k-1}^*$, we have $q_i \geq q_j$.

Proof: We prove this lemma by contradiction. Suppose $q_i < q_j$, for some $i \in S_{k-1}^* \backslash S_{k-2}^*$ and $j \in S_k^* \backslash S_{k-1}^*$. We construct a different solution η by swapping i and j, that is, we set $\bar{X}_i^\eta = \bar{X}_j^*, \bar{X}_j^\eta = \bar{X}_j^*, \hat{X}_i^\eta = \hat{X}_j^*, \hat{X}_i^\eta = \hat{X}_j^*$, $S_{k-1}^\eta = S_{k-1}^* \cup \{j\} \backslash \{i\}$, while keeping the rest of η to be the same as the optimal solution. Obviously, η satisfies all constraints (4)–(7), and all clients other than i and j have the same amount of penalties under both policies.

Under the optimal policy, $\sum_{m=1}^{|S_{k-1}^*|} \bar{R}_m(t) = \sum_{l \in S_{k-1}^*} \bar{X}_l^*(t)$. Therefore, in almost all intervals, j is not served until the packet for i is delivered, and $X_i^*(t + \mathbb{T}) - X_i^*(t) \geq X_j^*(t + \mathbb{T}) - X_j^*(t)$ for sufficiently large t almost surely. The penalty of i and j at time $t + \mathbb{T}$ is $C(X_i^*(t+\mathbb{T}) - X_i^*(t) - q_i\mathbb{T}) + C(X_j^*(t+\mathbb{T}) - X_j^*(t) - q_j\mathbb{T})$. If the penalty of i and j is larger than 0, it is strictly larger than the penalty by swapping i and j, $C(X_i^*(t+\mathbb{T}) - X_i^*(t) - q_j\mathbb{T}) + C(X_j^*(t+\mathbb{T}) - X_j^*(t) - q_i\mathbb{T})$, since $q_i < q_j$ and $C(\cdot)$ is convex. As we assume that $X_i(t) - t\bar{X}_i$ has the same distribution as $\hat{X}_i(t)$ for sufficiently large t, η achieves a smaller L than the optimal solution. ■

Theorem 3: Given an optimal solution to (3)–(7), for any two clients $i \in S_{k-1}^* \backslash S_{k-2}^*$ and $j \in S_k^* \backslash S_{k-1}^*$, we have $\bar{X}_i^* - q_i \leq \bar{X}_j^* - q_j$.

Proof: We prove this theorem by contradiction. Recall that $\bar{X}_i^* - q_i$ is the same for all clients $i \in S_k \backslash S_{k-1}$, for all k. Suppose there exists some k such that $\bar{X}_i^* - q_i > \bar{X}_j^* - q_j$ for $i \in S_{k-1}^* \backslash S_{k-2}^*$ and $j \in S_k^* \backslash S_{k-1}^*$. Let Δ be the average of $\bar{X}_l^* - q_l$, and σ be the average of σ_l for all $l \in S_k^* \backslash S_{k-2}^*$. We construct a new solution η by choosing $\bar{X}_l^\eta = \Delta + q_l$, $\sigma_l^\eta = \sigma$, removing S_{k-1}^* from the list of tight sets, and keeping the rest of the solution the same as the optimal solution. Since $C(\cdot)$ is a convex function, it is obvious that η achieves a smaller L.

It remains to show that η satisfies all constraints (4)–(7). Since $\bar{X}_i^* - q_i > \bar{X}_j^* - q_j$ for $i \in S_{k-1}^* \backslash S_{k-2}^*$ and $j \in S_k^* \backslash S_{k-1}^*$, we have $\bar{X}_i^\eta < \bar{X}_i^*$. Further, by Lemma 3, $q_i \geq q_j$, and hence $\bar{X}_i^\eta \geq \bar{X}_j^\eta$, for $i \in S_{k-1}^* \backslash S_{k-2}^*$ and $j \in S_k^* \backslash S_{k-1}^*$. Number the clients in $S_k^* \backslash S_{k-2}^*$ so that $\bar{X}_1^\eta \geq \bar{X}_2^\eta \geq \dots \geq \bar{X}_{|S_k^* \backslash S_{k-2}^*|}^\eta$. We now have, for every $1 \leq l \leq |S_k^* \backslash S_{k-2}^*|$, $\sum_{i=1}^l \bar{X}_i^\eta \leq \sum_{i=1}^l \bar{X}_i^*$, with equality holds if and only if $l = |S_k^* \backslash S_{k-2}^*|$. Therefore, constraints (4) and (5) hold.

Finally, we check constraint (6). Let $s_1 := |S_{k-2}|$, $s_2 := |S_{k-1}|$, and $s_3 := |S_k|$. We have, by (2),

$$\begin{aligned}
&(\sum_{l\in S_k\setminus S_{k-2}} \sigma_l^\eta)^2 = (\upsilon_{s_1+1,s_2}+\upsilon_{s_2+1,s_3})^2\\
&\geq \sum_{m=s_1+1}^{s_3} \bar{R}_m(1-\bar{R}_m) + 2\sum_{s_1+1\leq n<m\leq s_2} \bar{R}_m(1-\bar{R}_n)\\
&+2\sum_{s_2+1\leq n<m\leq s_3} \bar{R}_m(1-\bar{R}_n)\\
&+2[\sum_{n=s_1+1}^{s_2}\sum_{m=s_2+1}^{s_3} \bar{R}_n\bar{R}_m(1-\bar{R}_n)(1-\bar{R}_m)\\
&+2\sum_{s_1<n_1<n_2\leq s_2<m_1<m_2\leq s_3} \bar{R}_{n_2}\bar{R}_{m_2}(1-\bar{R}_{n_1})(1-\bar{R}_{m_1})]^{1/2}\\
&\geq \upsilon_{s_1+1,s_3}^2,
\end{aligned}$$

and η satisfies constraint (6). ■

B. An Algorithm for Finding the Optimal Solution

We now introduce an algorithm that finds a solution that satisfies the necessary conditions in Theorems 2 and 3. We further prove that the solution found by our algorithm is the only solution that satisfy the necessary conditions, and hence it is the optimal solution for (3)–(7).

Define a function $\bar{X}_i(\Delta) := \max\{0, \Delta - q_i\}$. The basic procedure of our algorithm is as follows: First, set Δ to be a sufficiently small value so that $\sum_{i\in S}\bar{X}_i(\Delta) < p\mathcal{F}(|S|)$, for all S. Next, increase the value of Δ until we can find a subset S_1 with $\sum_{i\in S_1}\bar{X}_i(\Delta) = p\mathcal{F}(|S_1|)$. Let S_1 be the first tight set and fix $\bar{X}_i = \bar{X}_i(\Delta)$, for all $i \in S_1$. Continue to increase the value of Δ until there is a subset S_2 with $\sum_{i\in S_2}\bar{X}_i(\Delta) = p\mathcal{F}(|S_2|)$. Let S_2 be the second tight set and fix $\bar{X}_i = \bar{X}_i(\Delta)$, for all $i \in S_2\setminus S_1$. Repeat the above procedure until $\bar{X}_i$ are fixed for all clients. Finally, set $[\sigma_i]$ as specified in Theorem 2. Details can be found in Algorithm 1, where we streamline some of the steps. It is easy to verify that the complexity of Algorithm 1 is only $O(\mathbb{J}^2)$.

Algorithm 1 Waterfilling Algorithm

```
Sort clients so that q_1 ≥ q_2 ≥ ··· ≥ q_J
i ← 0, u ← 1
while i < J do
    delta ← min_{i+1≤n≤J} (pF(n) − pF(i) − Σ_{i+1≤k≤n} q_k) / (n−i)
    index ← arg min_{i+1≤n≤J} (pF(n) − pF(i) − Σ_{i+1≤k≤n} q_k) / (n−i)
    for j = i+1 to index do
        X̄_j ← q_j + delta
        σ_i ← υ_{i+1,index} / (index−i)
    S_u ← {1, 2, ..., index}
    i ← index, u ← u+1
```

It is obvious that the solution produced by Algorithm 1 satisfies the conditions in Theorems 2 and 3. Next, we show that it is the only solution that satisfies the conditions.

Theorem 4: The solution produced by Algorithm 1 is the only solution that satisfies the conditions in Theorems 2 and 3.

Proof: We prove this theorem by contradiction. Let $\{S_1^*, S_2^*, \dots\}$ and $[\bar{X}_i^*]$ be the solution by Algorithm 1, and $\{S_1^\eta, S_2^\eta, \dots\}$ and $[\bar{X}_i^\eta]$ be another solution that satisfies the conditions in Theorems 2 and 3. Since Theorem 2 uniquely determines the values of $[\bar{X}_i]$ and $[\sigma_i]$ when tight sets are given, we have $\{S_1^*, S_2^*, \dots\} \neq \{S_1^\eta, S_2^\eta, \dots\}$.

Let k be the smallest integer such that $S_k^* \neq S_k^\eta$. By the design of Algorithm 1, we have $\bar{X}_i^\eta - q_i > \bar{X}_j^* - q_j$, for every $i \in S_k^\eta\setminus S_{k-1}^\eta$ and $j \in S_k^*\setminus S_{k-1}^*$. Otherwise, S_k^η would have been selected as a tight set by Algorithm 1. Now, none of the clients in $S_k^*\setminus S_{k-1}^*$ are in $S_{k-1}^\eta = S_{k-1}^*$. By Theorem 3, we have $\bar{X}_j^\eta - q_j \geq \bar{X}_i^\eta - q_i > \bar{X}_j^* - q_j$, and hence $\bar{X}_j^\eta > \bar{X}_j^*$, for every $i \in S_k^\eta\setminus S_{k-1}^\eta$ and $j \in S_k^*\setminus S_{k-1}^*$. We then have $\sum_{j\in S_k^*}\bar{X}_j^\eta > \sum_{j\in S_k^*}\bar{X}_j^* = p\mathcal{F}(|S_k^*|)$, which violates (5). ■

In summary, the solution produced by Algorithm 1 is the unique optimal solution to the problem (3)–(7).

VI. An Optimal Online Scheduling Policy

While Algorithm 1 derives the optimal solution to the problem (3)–(7), it remains unclear if there indeed exists any scheduling policy that can achieve the optimal solution. In this section, we introduce a simple online scheduling policy that achieves the solution given by Algorithm 1. Our policy is based on the concept of *debt*:

Definition 1: Define $d_i(t) := \bar{X}_i t - X_i(t)$, where $\bar{X}_i$ is derived from Algorithm 1, and $X_i(t)$ is the total number of packets delivered for client i in the first t intervals. We call $d_i(t)$ the *debt* of client i. □

We now introduce our policy, which is called the *Tiered Largest Debt First* (TLDF) policy. TLDF first runs Algorithm 1 to obtain $\{S_1, S_2, \dots\}$ and $[\bar{X}_i]$. At the beginning of each interval $t+1$, TLDF assigns priorities to clients according to the following two rules: First, clients in S_k have higher priorities than clients not in S_k, for all k. Second, for all clients in $S_k\setminus S_{k-1}$, clients with higher debts receive higher priorities. For example, consider a system with 4 clients with $S_1 = \{1,3\}, S_2 = \{1,2,3,4\}, d_1(t) = 1, d_2(t) = 2, d_3(t) = 3$, and $d_4(t) = 4$. TLDF gives client 3 the highest priority, followed by clients 1, 4, and 2.

TLDF schedules clients according to their priorities. Specifically, in each time slot within the interval, TLDF schedules the client with the highest priority among those who have an undelivered packet.

We now analyze the performance of TLDF. Recall that $R_n(t)$ is defined to be the indicator function that at least n packets are delivered in the t-th interval. Section IV has shown that the mean of $R_n(t)$ is $\bar{R}_n = p\mathcal{F}(n) - p\mathcal{F}(n-1)$. By the design of TLDF, we have, for all k and t, $\sum_{i\in S_k\setminus S_{k-1}} X_i(t) - X_i(t-1) = \sum_{n=|S_{k-1}|+1}^{|S_k|} R_n(t)$. Also, by the design of Algorithm 1, $\sum_{i\in S_k\setminus S_{k-1}}\bar{X}_i = \sum_{n=|S_{k-1}|+1}^{|S_k|}\bar{R}_n$, for all k.

Since TLDF assigns priorities for clients in $S_k \backslash S_{k-1}$ solely based on clients' debts, we can construct a Markov process whose state is $[d_i(t) - \frac{\sum_{j\in S_k\backslash S_{k-1}} d_j(t)}{|S_k\backslash S_{k-1}|} | i \in S_k\backslash S_{k-1}]$. To simplify notation, we let $D_k(t) := \frac{\sum_{j\in S_k\backslash S_{k-1}} d_j(t)}{|S_k\backslash S_{k-1}|}$.

Lemma 4: Under TLDF, the Markov process with state $[d_i(t) - D_k(t) | i \in S_k\backslash S_{k-1}]$ is positive recurrent, for all k.

Proof: Define a Lyapunov function

$$L(t) := \frac{1}{2} \sum_{i\in S_k\backslash S_{k-1}} [d_i(t) - D_k(t)]^2.$$

Define the drift of the Lyapunov function by: $\Delta L(t) := E[L(t+1) - L(t)|\text{state at } t]$

In order to simplify notations, we renumber the clients so that the subset $S_k\backslash S_{k-1}$ consists of clients $\{1, 2, \dots, \mathbb{K}\}$, and $d_1(t) \geq d_2(t) \geq \dots$. By the design of TLDF, client i has the $(|S_{k-1}| + i)$-highest priority, and we have $X_i(t+1) - X_i(t) = R_{|S_{k-1}|+i}(t)$. Moreover, by the design of Algorithm 1, we have $\sum_{i=1}^{\mathbb{K}} \bar{X}_i = p\mathcal{F}(|S_k|) - p\mathcal{F}(|S_{k-1}|)$, and $\sum_{i=1}^{j} \bar{X}_i < p\mathcal{F}(|S_{k-1}|+j) - p\mathcal{F}(|S_{k-1}|)$, for all $j < \mathbb{K}$.

Let $\Delta d_i(t) := d_i(t+1) - d_i(t) = \bar{X}_i - R_{|S_k|+i}(t)$. We have $E[\sum_{i=1}^{\mathbb{K}} \Delta d_i(t)] = \sum_{i=1}^{\mathbb{K}} \bar{X}_i - \sum_{i=1}^{\mathbb{K}} \bar{R}_{|S_{k-1}|+i} = 0$. Also, by condition (5), there exists some $\delta > 0$ such that $E[\sum_{i=1}^{j} \Delta d_i(t)] = \sum_{i=1}^{j} \bar{X}_i - \sum_{i=1}^{j} \bar{R}_{|S_{k-1}|+i} < -\delta$, for all $j < \mathbb{K}$. We now have

$$\begin{aligned}
&\Delta L(t) = E[L(t+1) - L(t)|\text{state at } t] \\
\leq &\sum_{i=1}^{\mathbb{K}} E[\Delta d_i(t)][d_i(t) - D_k(t)] + \beta \\
= &\sum_{i=1}^{\mathbb{K}-1} E[\sum_{j=1}^{i} \Delta d_j(t)](d_i - d_{i+1}) \\
&+ E[\sum_{i=1}^{\mathbb{K}} \Delta d_i(t)](d_{\mathbb{K}}(t) - D_k(t)) + \beta \\
< &-\delta(d_1 - d_{\mathbb{K}}) + \beta,
\end{aligned}$$

where β is a bounded positive number. Note that $d_1(t)$ and $d_{\mathbb{K}}(t)$ are the largest and smallest debt for clients $\{1, 2, \dots, \mathbb{K}\}$, respectively. Also, $D_k(t)$ is the average of $d_i(t)$. Hence, $|d_i(t) - D_k(t)| < d_1(t) - d_{\mathbb{K}}(t)$, for all i. The inequality above then implies

$$\Delta L(t) < -\delta, \text{if } |d_i(t) - D_k(t)| > \tfrac{\beta}{\delta} + 1, \text{ for some } i,$$

and

$$\Delta L(t) < \beta, \text{otherwise}.$$

By the Foster-Lyapunov Theorem, the Markov process is positive recurrent. ∎

We are now ready to prove that TLDF achieves the smallest penalty.

Theorem 5: Under TLDF, $\lim_{T\to\infty} \frac{X_i(T)}{T} = \bar{X}_i$, and $\lim_{T\to\infty} \frac{X_i(tT) - tT\bar{X}_i}{\sqrt{T}}$ is a driftless Brownian motion with variance σ_i^2, where $\bar{X}_i$ and σ_i^2 are derived from Algorithm 1.

Proof: Since the Markov process with state $[d_i(t) - D_k(t) | i \in S_k\backslash S_{k-1}]$ is positive recurrent, we have $\lim_{T\to\infty} \frac{d_i(T) - D_k(T)}{T} = 0$. Further, by definition,

$$\begin{aligned}
\lim_{T\to\infty} \frac{D_k(T)}{T} &= \lim_{T\to\infty} \frac{\sum_{i\in S_k\backslash S_{k-1}} d_i(T)}{|S_k\backslash S_{k-1}| T} \\
&= \frac{1}{|S_k\backslash S_{k-1}|} (\sum_{i\in S_k\backslash S_{k-1}} \bar{X}_i - \sum_{n=|S_{k-1}|+1}^{|S_k|} \bar{R}_n) = 0.
\end{aligned}$$

Hence, $\lim_{T\to\infty} \frac{d_i(T)}{T} = \bar{X}_i - \lim_{T\to\infty} \frac{X_i(T)}{T} = 0$. This establishes the first part of the theorem.

The positive recurrence of the Markov process also implies that $\lim_{T\to\infty} \frac{d_i(tT) - D_k(tT)}{\sqrt{T}}$ converges to 0 in probability. We also have

$$\begin{aligned}
&\lim_{T\to\infty} \frac{D_k(tT)}{\sqrt{T}} \\
= &\lim_{T\to\infty} \frac{\sum_{n=|S_{k-1}|+1}^{|S_k|} (\bar{R}_n tT - \sum_{u=1}^{tT} R_n(u))}{|S_k\backslash S_{k-1}| \sqrt{T}} \\
= &-\sum_{n=|S_{k-1}|+1}^{|S_k|} \hat{R}_n(t) / |S_k\backslash S_{k-1}|,
\end{aligned}$$

and

$$\begin{aligned}
\lim_{T\to\infty} \frac{d_i(tT)}{\sqrt{T}} &= \lim_{T\to\infty} \frac{\bar{X}_i tT - X_i(tT)}{\sqrt{T}} \\
&= -\hat{X}_i(t).
\end{aligned}$$

Therefore, $\hat{X}_i(t) = \sum_{n=|S_{k-1}|+1}^{|S_k|} \hat{R}_n(t) / |S_k\backslash S_{k-1}|$, which is a driftless Brownian motion with variance $\frac{v^2_{|S_{k-1}|+1, |S_k|}}{|S_k\backslash S_{k-1}|^2}$, as shown in Section IV. This proves the second part of the theorem. ∎

VII. Simulation Results

We implement our TLDF policy in ns-2 and compare it against other policies. We consider a system with one AP and 12 clients. Simulation shows that the time needed for a transmission, including the time to transmit ACK and other overheads, is about $0.49ms$. We hence set the duration of one time slot as $0.5ms$. We set the duration of an interval as $10ms$, or 20 time slots. The channel reliability p is set to be 50%. All simulations presented in this section are the average of 1000 simulation runs.

A. Accuracy of the Approximation

We assume that $X_i(t) \approx \bar{X}_i t + \hat{X}_i(t)$ throughout the paper. Hence, our first simulation is to evaluate whether this is indeed an accurate approximation. We set $q_i \equiv 0.8$, for all i. By Algorithm 1, we have $\bar{X}_i \equiv 0.8152$. We run TLDF with this setting. We record the value of $[X_i(t) - X_i(t-\mathbb{T}) - \mathbb{T}\bar{X}_i]^2$ of client 1 every $250ms$. We then compare this value against the theoretical value of the variance, which equals $0.0251\mathbb{T}$ under our setting, for the cases $\mathbb{T} = 20, 50$, and 100. Since $X_i(t) - X_i(t-\mathbb{T})$ cannot

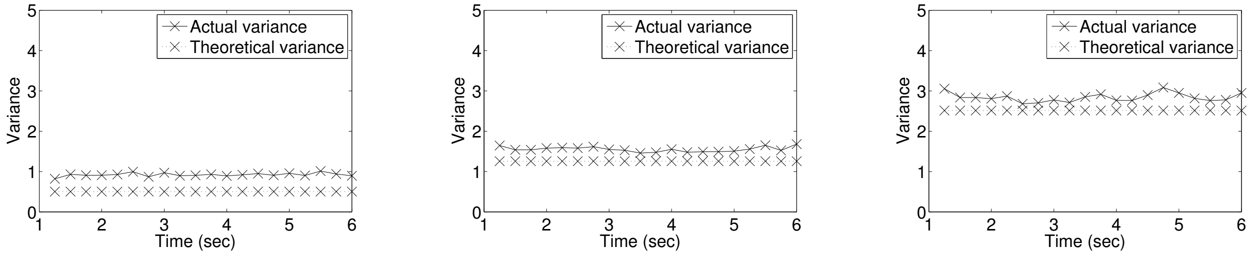

(a) $\mathbb{T} = 20$ (b) $\mathbb{T} = 50$ (c) $\mathbb{T} = 100$

Fig. 1. Actual variance and theoretical variance.

be defined in the first $\mathbb{T}$ intervals, we start recording values after one second.

Simulation results are shown in Fig. 1. It can be shown that the actual variance is very close to the theoretical variance, with difference smaller than 0.4 for most data points. This suggests that the approximation $X_i(t) \approx \bar{X}_i t + \hat{X}_i(t)$ is very accurate, even when $\mathbb{T}$ is as small as 20. Further, a significant part of the difference may be caused by rounding issues: Since $X_i(t)$ is an integer, the absolute difference between $X_i(t) - X_i(t - \mathbb{T})$ and $\mathbb{T}\bar{X}_i$ is at least $\min\{\mathbb{T}\bar{X}_i - \lfloor \mathbb{T}\bar{X}_i \rfloor, \lceil \mathbb{T}\bar{X}_i \rceil - \mathbb{T}\bar{X}_i\}$. Under our setting, the values of $\min\{\mathbb{T}\bar{X}_i - \lfloor \mathbb{T}\bar{X}_i \rfloor, \lceil \mathbb{T}\bar{X}_i \rceil - \mathbb{T}\bar{X}_i\}$ for $\mathbb{T} = 20, 50,$ and 100 are 0.497, 0.257, and 0.486, respectively. This partly explains why the difference is the smallest when $\mathbb{T} = 50$. Finally, we note that the actual variance does not change much with time. This suggests that TLDF converges very fast.

B. Comparisons against Other Policies

We compare the short-term performance of TLDF against other policies. We consider two policies: the first policy is the largest time-based debt first policy proposed by Hou, Borkar, and Kumar [9]. This policy is shown to achieve $\lim_{t\to\infty} \frac{X_i(t)}{t} = \bar{X}_i$. The other policy, which we call the tiered-random (RandT) policy, first runs Algorithm 1 to obtain the optimal choice of $\{S_1, S_2, \dots\}$. In each interval, RandT assigns priorities to clients so that clients in S_k have higher priorities than those not in S_k. The priorities between clients in the same set $S_k \backslash S_{k-1}$ are assigned randomly.

We consider two different systems. The first system is a homogeneous one where all clients have $q_i = 0.8$. The second system is heterogeneous as different clients have different q_i. Specifically, we choose $q_1 = q_2 = q_3 = q_4 = 0.97$, $q_5 = q_6 = q_7 = q_8 = 0.95$, $q_9 = 0.51, q_{10} = 0.49, q_{11} = 0.47$, and $q_{12} = 0.45$. We consider the quadratic penalty function $C(x) = x^2$. The cases $\mathbb{T} = 20, 50,$ and 100 are evaluated. In each simulation run, we record the total amount of penalty incurred in each second.

Simulation results for the homogeneous system and the heterogeneous one are shown in Fig. 2 and 3, respectively. It can be shown that TLDF outperforms the other two policies greatly. Moreover, the largest time-based debt first policy has very high penalty in all settings, even much higher than RandT. Note that the largest time-based debt first policy achieves the optimal long-term performance in the sense that $\lim_{t\to\infty} \frac{X_i(t)}{t} = \bar{X}_i$, for all i. Our results therefore demonstrate that policies that only optimize long-term performance can have very poor short-term performance, which is actually crucial for most practical safety-critical applications.

VIII. Conclusion

This paper proposes a new approach to model the performance of real-time wireless networks so as to capture both short-term and long-term behaviors. In this approach, the instantaneous penalty incurred by a flow only depends on a small time horizon in the past. We then study the problem of minimizing the long-term average total penalty in the system. By studying some inherent features of real-time wireless networks, we formulate an optimization problem that characterizes a lower-bound of total penalty. Although the optimization problem is not convex, we propose an algorithm that finds its unique optimal solution. We then establish a simple online scheduling policy that converges to the unique optimal solution. The performance of this scheduling policy is extensively evaluated by simulations. Simulation results not only show that our policy outperforms others greatly, but also demonstrate that policies aiming to optimize long-term performance can have poor short-term behaviors. By explicitly taking short-term performance into account, our approach offers a desirable solution to safety-critical applications.

References

[1] C. N. Laws, "Resource pooling in queueing networks with dynamic routing," *Advances in Applied Probability*, pp. 699–726, 1992.

[2] W. Kang, F. P. Kelly, N. H. Lee, and R. J. Williams, "State space collapse and diffusion approximation for a network operating under a fair bandwidth sharing policy," *The Annals of Applied Probability*, pp. 1719–1780, 2009.

[3] S. Bhardwaj and R. J. Williams, "Diffusion approximation for a heavily loaded multi-user wireless communication system with cooperation," *Queueing Systems*, vol. 62, no. 4, pp. 345–382, 2009.

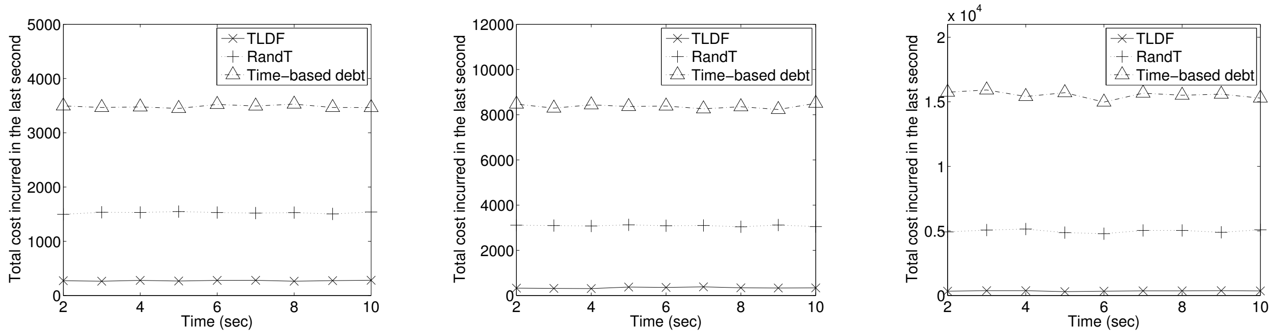

(a) $\mathbb{T} = 20$ (b) $\mathbb{T} = 50$ (c) $\mathbb{T} = 100$

Fig. 2. Performance comparison for the homogeneous system.

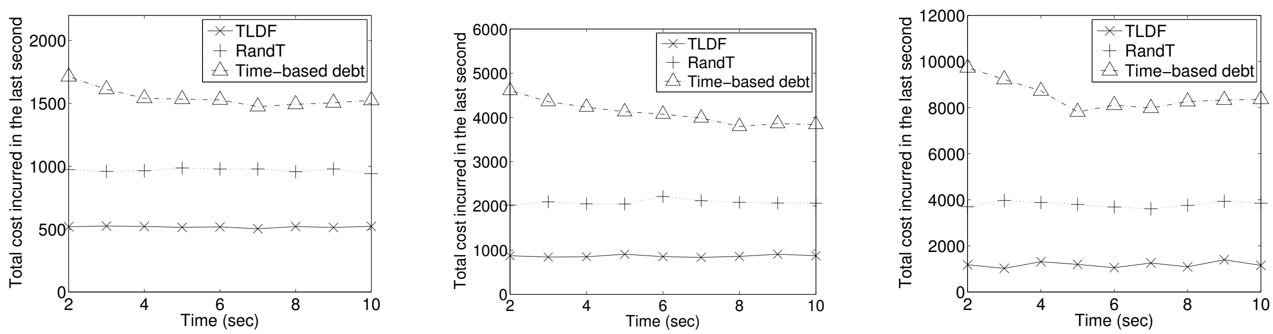

(a) $\mathbb{T} = 20$ (b) $\mathbb{T} = 50$ (c) $\mathbb{T} = 100$

Fig. 3. Performance comparison for the heterogeneous system.

[4] A. Destounis, M. Assaad, M. Debbah, B. Sayadi, and A. Feki, "On queue-aware power control in interfering wireless links: Heavy traffic asymptotic modelling and application in qos provisioning," *Mobile Computing, IEEE Transactions on*, vol. 13, no. 10, pp. 2345–2356, 2014.

[5] I. M. Verloop and R. Núnez-Queija, "Asymptotically optimal parallel resource assignment with interference," *Queueing Systems*, vol. 65, no. 1, pp. 43–92, 2010.

[6] A. Eryilmaz and R. Srikant, "Asymptotically tight steady-state queue length bounds implied by drift conditions," *Queueing Systems*, vol. 72, no. 3-4, pp. 311–359, 2012.

[7] W. Wang, K. Zhu, L. Ying, J. Tan, and L. Zhang, "Map task scheduling in mapreduce with data locality: Throughput and heavy-traffic optimality," in *INFOCOM, 2013 Proceedings IEEE*, pp. 1609–1617, IEEE, 2013.

[8] B. Li, R. Li, and A. Eryilmaz, "Heavy-traffic-optimal scheduling with regular service guarantees in wireless networks," in *Proceedings of the fourteenth ACM international symposium on Mobile ad hoc networking and computing*, pp. 79–88, ACM, 2013.

[9] I.-H. Hou, V. Borkar, and P. R. Kumar, "A theory of QoS for wireless," in *INFOCOM, 2010 Proceedings IEEE*, 2009.

[10] J. 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*, pp. 1–9, IEEE, 2010.

[11] R. Li and A. Eryilmaz, "Scheduling for end-to-end deadline-constrained traffic with reliability requirements in multihop networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 20, no. 5, pp. 1649–1662, 2012.

[12] X. Kang, W. Wang, J. J. Jaramillo, and L. Ying, "On the performance of largest-deficit-first for scheduling real-time traffic in wireless networks," in *Proceedings of the fourteenth ACM international symposium on Mobile ad hoc networking and computing*, pp. 99–108, ACM, 2013.

[13] J. J. Jaramillo, R. Srikant, and L. Ying, "Scheduling for optimal rate allocation in ad hoc networks with heterogeneous delay constraints," *Selected Areas in Communications, IEEE Journal on*, vol. 29, no. 5, pp. 979–987, 2011.

[14] I.-H. Hou and P. R. Kumar, "Scheduling heterogeneous real-time traffic over fading wireless channels," in *INFOCOM, 2010 Proceedings IEEE*, pp. 1–9, IEEE, 2010.

[15] H. Gangammanavar and A. Eryilmaz, "Dynamic coding and rate-control for serving deadline-constrained traffic over fading channels," in *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*, pp. 1788–1792, IEEE, 2010.

[16] R. Singh, I.-H. Hou, and P. R. Kumar, "Fluctuation analysis of debt based policies for wireless networks with hard delay constraints," in *INFOCOM, 2014 Proceedings IEEE*, pp. 2400–2408, IEEE, 2014.

[17] G. L. Jones, "On the markov chain central limit theorem," *Probability surveys*, vol. 1, no. 5–1, pp. 299–320, 2004.