

Imperfect Randomized Algorithms for the Optimal Control of Wireless Networks

Atilla Eryilmaz
Electrical and Computer Engineering
Ohio State University
Columbus, OH 43210
Email: eryilmaz@ece.osu.edu

Asuman Ozdaglar, Devavrat Shah
Electrical Eng. and Computer Science
Massachusetts Institute of Technology
Cambridge, MA, 02139
Emails: {asuman, devavrat}@mit.edu

Eytan Modiano
Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, MA, 02139
Email: modiano@mit.edu

Abstract— We consider a joint randomized scheduling, congestion control mechanism for general wireless networks. We allow for a set of imperfections in the operation of the randomized scheduler to account for potential errors in its operation. These imperfections enable the design of a large class of low-complexity and distributed implementations for different interference models. We study the effect of such imperfections on the stability and fairness characteristics of the system, and explicitly characterize the degree of fairness achieved as a function of the level of imperfections. Our results also reveal the relative importance of different types of errors on the performance of the system, and provide valuable insight to the design of distributed controllers with favorable fairness characteristics.

I. INTRODUCTION

There has been considerable recent interest in developing network protocols to achieve the multiple objectives of throughput maximization and fair allocation of resources among competing users. Much of the work in wireless communication networks has focused on centralized control and has developed *throughput-optimal*¹ policies (e.g. [23], [17], [6]). However, these policies do not directly lend themselves to distributed implementation, which is essential in practice. In this paper, we provide a class of imperfect randomized scheduling and flow control algorithms that are amenable to distributed implementation, and study their stability and fairness characteristics. In particular, we explicitly characterize the degree of fairness achieved by such imperfect policies as a function of the level of imperfections.

Queue-length-based dynamic scheduling and routing strategies has proven to be very successful in effectively utilizing the system resources. In particular, the class of *back-pressure policies* introduced in [23] and later extended in more recent works (e.g. [1], [19], [17], [8]) have been shown to be throughput-optimal. While these schedulers can be implemented distributively in wireline networks, the implementation requires centralized and high-complexity operations in wireless networks due the coupling effect of the interference between concurrent transmissions.

Recent work (e.g. [12], [24], [2], [25]) focused on developing distributed implementation at the cost of sacrificing a

portion of the capacity region, where the level of sacrifice depends on the network topology and the interference model. It is observed in [3] that the level of sacrifice can be significant for some network models.

In another line of work, a class of low-complexity schedulers are proposed which preserves the throughput-optimality characteristics. These schedulers are based on [22], where an iterative algorithm is proposed to *gradually* improve the scheduler performance rather than searching for the best scheduler in every step. This reduces the complexity of the implementation considerably at the possible cost of a rate-of-convergence degradation. Also, the associated implementation was not distributed. In recent works ([15], [4], [18]) the framework of [22] has been improved to propose distributed implementations for specific interference models. In particular, the work of [15] studies the throughput performance of a class of randomized schedulers with imperfections in their operation. The incorporation of such imperfections may be useful in capturing potential errors in the distributed implementation.

In addition to throughput characteristics, another important performance criteria, especially in the presence of elastic traffic such as data transmissions, is *fairness*. We say that a rate allocation is *fair* or *optimal* if it maximizes, over all feasible rates, the sum of a set of utility functions that model the preferences of different flows (see e.g. [9], [10], [13], [26], [20]). Congestion controllers have been developed in numerous studies (see [20] for an overview) to achieve the optimal allocation via decentralized methods. These controllers are operated at the sources to regulate the traffic of each flow based on the congestion level experienced at the time.

More recently, the routing, scheduling component is combined with the congestion controller to develop a joint mechanism ([11], [21], [5], [16], [20], [7]) with optimality properties. However, the implementation difficulties of the optimal scheduler for wireless networks persists in this joint mechanism.

The goal of this work is to study the proximity of the achieved rate allocation to the optimal allocation under a large class of imperfect randomized schedulers that are amenable to distributed implementation. By modeling a variety of imperfections that can occur due computational or communication errors, our work creates a flexible framework for the design of distributed schedulers under general interference models with

¹A throughput optimal policy stabilizes any set of flow rates that is stabilizable by any other policy.

other attractive features such as robustness, low-complexity, high convergence rate, and low-overhead.

In a related work [12], authors consider one particular imperfection in the scheduler operation and study its effect on the optimality. However, the type of imperfection considered [12] cannot model the attractive class of schedulers proposed in [22], [15], [4], [18]. In a more recent work [4], the authors study a particular class of distributed randomized schedulers for the secondary interference model with the iterative nature of [22]. However, the model assumes error-free operations and hence may not be able applicable to cases where errors are unavoidable and/or helpful for lower-complexity implementation.

In this work, we extend the framework of [4] by modeling various potential errors that can occur in the operation of the scheduler, and by investigating their effect on the rate allocation under our proposed joint mechanism. Our results also reveal the critical components in the design of high-performance distributed algorithms in wireless networks.

The rest of the paper is organized as follows. In Section II, we provide the system model together with our objective. In Section III, we describe a class randomized scheduling mechanisms that contains many of the existing schedulers and also allows for many others that can be implemented distributively. In Section IV, we describe a flow-control mechanism that will operate concurrently with the randomized scheduler. In Section V, we analyze the optimality characteristics of the joint scheduling, flow-control mechanism and explicitly characterize the proximity of the achieved throughput to the optimal in terms of the parameters of the randomized policy. We will provide some concluding remarks in Section VI.

II. SYSTEM MODEL AND OBJECTIVES

A fixed wireless network can be modeled by a directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{L})$ with a node set \mathcal{N} (with cardinality N), a link set \mathcal{L} (with cardinality L). We assume that the nodes are synchronized to a common clock and that they operate in a time-slotted fashion with each slot long enough to accommodate a single *packet* transmission. Each directed link corresponds to a potential transmission from the head to the tail of the link.

A. Interference Model and Stability Region

Due to the nature of the wireless medium, only links that do not interfere with each other can successfully convey information simultaneously. In particular, we let $\mathcal{I}(l)$ denote the set of links, other than l itself, that interfere with link $l \in \mathcal{L}$. Each link is either *active* or *inactive* in every time slot. We assume that each link that does not interfere with any other link that is active in the same slot, can convey a single packet within one slot.

We let \mathcal{S} denote the set of (*feasible*) *link activations* or (*feasible*) *schedules*. For each feasible schedule $\mathbf{S} \in \mathcal{S}$, we have: $\mathbf{S} = (S_l)_{l \in \mathcal{L}}$ with $S_l \in \{0, 1\}$ for each $l \in \mathcal{L}$; and

$$S_l + \sum_{e \in \mathcal{I}(l)} S_e \leq 1,$$

for all $l \in \mathcal{L}$. The latter condition says that no two interfering links of a feasible schedule can be simultaneously active. We let $\mathbf{S}[t]$ denote the schedule that is picked at Slot- t . This is a general model that captures all k^{th} - order interference scenarios whereby any two links within k hops from each other interfere. In particular: when $k = 1$, the set of feasible schedules corresponds to the set of all matchings, which is a good model for CDMA communication; when $k = 2$, no two links within a two hop neighborhood can communicate simultaneously, which is a good model for IEEE 802.11 networks.

We also define the *Stability Region* of the network, denoted by $\overline{\mathcal{S}}$, as

$$\begin{aligned} \overline{\mathcal{S}} &\triangleq \text{Convex Hull } \{\mathcal{S}\} \\ &= \left\{ \mathbf{S}' : \exists \text{ a prob. dist'n } \pi \text{ with } \mathbf{S}' = \sum_{\mathbf{S} \in \mathcal{S}} \pi \cdot \mathbf{S} \right\}, \end{aligned}$$

where $\pi \cdot \mathbf{S}$ denotes component-wise multiplication.

B. Traffic Model and Optimal Rate Allocation

In this work we assume that each flow traverses a single link. This is in contrast to the more realistic scenario of flows with multi-hop routes. The single-hop scenario is sufficient to capture the essential components of the mechanism while avoiding the additional complications associated with the routing mechanism. This allows us to present the analysis more clearly. We will later comment on the generalization to the multi-hop routing scenario.

Different flows have different demands from the network, some valuing higher data rate than others. It is observed in the literature that utility functions can be used to capture such preferences in a mathematical framework (e.g. [10], [13], [14], [20], [5]). Therefore, under the aforementioned traffic model, we associate a utility function $U_l(\cdot)$ to each link $l \in \mathcal{L}$ that is a function of the mean throughput of link l . We further assume that $U_l(\cdot)$ is a non-negative, non-decreasing, differentiable and concave function of its parameter.

Interpreting the utility functions as measures of satisfaction obtained over links, we define the *optimal rate allocation vector* $\mathbf{x}^* = (x_l^*)_l$ as the feasible mean rate that yields the maximum total satisfaction over all the rates within the stability region. In more mathematical terms, we define

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \overline{\mathcal{S}}} \sum_{l \in \mathcal{L}} U_l(x_l). \quad (1)$$

C. System Evolution

Packets are generated for each link based on a flow/congestion control mechanism that is to be designed. We let $X_l[t]$ denote the number of packets (not necessarily an integer) generated during Slot- t over link l . Also, associated with each link $l \in \mathcal{L}$, there exists a virtual queue, denoted by Q_l that holds the packets associated with link l . In particular, $Q_l[t]$ denotes the number of packets that Queue- l contains at the beginning of Slot- t .

Then, the evolution of the queue-lengths vector $\mathbf{Q}[t] = (Q_l[t])_l$ can be given by

$$\begin{aligned}\mathbf{Q}[t+1] &= (\mathbf{Q}[t] - \mathbf{S}[t])^+ + \mathbf{X}[t] \\ &= \mathbf{Q}[t] - \mathbf{S}[t] + \mathbf{X}[t] + \mathbf{U}[t]\end{aligned}\quad (2)$$

where $(y)^+ \triangleq \max(0, y)$, and $\mathbf{U}[t] \triangleq \max(0, \mathbf{S}[t] - \mathbf{Q}[t])$ denotes the amount of *unused service* over link l in Slot- t . We say that the queue Q_l is said to be *stable* if

$$\limsup_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{t=0}^{\tau-1} \mathbb{E}[Q_l[t]] \leq \mathbb{Q}$$

for some $\mathbb{Q} < \infty$. We say the *network is stable* if all queues are stable. Also, we define

$$\bar{x}_l \triangleq \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{t=0}^{\tau-1} \mathbb{E}[X_l[t]],$$

when it exists, as the *mean flow rate over Link- l* .

D. Objective

There are two operations, namely scheduling and flow control, that need to be designed in the context of the network described so far. On the one hand, the choice of the feasible schedule needs to be determined in each slot subject to the interference constraints, and on the other hand, the amount of traffic to be generated by each flow needs to be determined for each link. These two mechanisms interact to jointly determine the performance of the system.

The goal of the joint mechanism is to achieve mean link rates that get arbitrarily close to the optimal rate defined in (1). Moreover, the resulting mechanism must be amenable to distributed implementation. In this work, we study the optimality characteristics of a class of randomized controllers that are amenable to distributed implementation.

III. PARAMETRIC CLASS OF RANDOMIZED SCHEDULERS

In this section, we describe a generic class of imperfect randomized schedulers that we will consider. The class, introduced in [15], extends the PICK-and-UPDATE Scheduler of Tassiulas [22] to include errors in the implementation. Such an extension may be critical in the distributed operation of the system since errors and imperfections are typical consequences of distributed computations. The general idea behind the algorithm is to have the scheduler evolve towards the optimal solution by an iterative process of: randomly selecting a candidate feasible schedule (PICK) and replacing the existing one in case the candidate offers better performance (UPDATE).

Before we start describing the scheduler, let us define the *optimal schedule* $\mathbf{S}^*[t]$ as

$$\mathbf{S}^*[t] = \arg \max_{\mathbf{S} \in \mathcal{S}} \langle \mathbf{S}, \mathbf{Q}[t] \rangle,$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product of two vectors. We are now ready to describe the scheduler that is parametrized by three variables (δ, γ, ψ) , each capturing a different imperfection in its operation.

Definition 1 (RANDOMIZED SCHEDULER (δ, γ, ψ)): In each slot, say t , the scheduler performs two steps:

- **PICK**: Pick a random feasible schedule, $\mathbf{R}[t]$, such that

$$\mathbb{P}(\mathbf{R}[t] = \mathbf{S}^*[t]) > \delta,$$

for some $\delta > 0$.

- **UPDATE**: Set the schedule $\mathbf{S}[t]$ such that

$$\mathbb{P}(\langle \mathbf{S}[t], \mathbf{Q}[t] \rangle \geq \max\{\langle \mathbf{S}[t-1], \mathbf{Q}[t] \rangle, (1-\gamma)\langle \mathbf{R}[t], \mathbf{Q}[t] \rangle\}) \geq 1-\psi,$$

for some γ and ψ in $[0, 1]$.

The PICK step randomly selects a feasible schedule with the mild limitation that there is a nonzero probability of being the optimal schedule as measured by the parameter δ . The UPDATE step actually encapsulates two operations: **Compute**, where the weights of $\mathbf{R}[t]$ and $\mathbf{S}[t-1]$ are computed with respect to the current queue-length vector $\mathbf{Q}[t]$; **Compare**, where the maximum of the weights are chosen. The γ parameter allows for the possibility of error in the **Compute** operation, while the ψ parameter allows for the possibility of error in the **Compare** operation. These parameters help us capture the imperfections of distributed or approximate implementations.

Several existing schedulers fall into this class of schedulers. In particular: the schedulers of [22] and [4] are the cases when $\delta > 0, \gamma = \psi = 0$; the schedulers of [15] contains examples of $\delta, \gamma, \psi > 0$; the scheduler of [18] is the case when $\delta > 0, \psi = 0, \gamma = 1/m$, for a design parameter $m > 0$. Thus, the framework is flexible enough to accommodate many other schedulers with varying features.

Next, we repeat the main result of [15], which describes the set of rates that are achievable by a Randomized Scheduler.

Theorem 1 ([15]): The Randomized Scheduler (δ, γ, ψ) stabilizes the network for any mean link flow rate vector \mathbf{x} satisfying

$$\mathbf{x} \in \left\{ \mathbf{y} \geq \mathbf{0} : \left(1 - \gamma - 2\sqrt{\frac{\psi}{\delta}} \right) \mathbf{y} \in \bar{\mathcal{S}} \right\}.$$

Therefore, the impact of the imperfections of the randomized scheduler results in the shrinkage of the stability region by a factor of $(1 - \gamma - 2\sqrt{\psi/\delta})$. Note that when $\psi = \gamma = 0$, and $\delta > 0$, the stability region is undisturbed. The questions we are concerned with in this work are regarding the effect of congestion control upon this class of schedulers, and the optimality performance of the overall scheme.

IV. FLOW/CONGESTION CONTROLLER

The congestion control mechanism that determines $\mathbf{X}[t]$ as a function of $\mathbf{Q}[t]$ is provided in this Section. Each link utilizes its queue-length information in the decision process by interpreting it as a measure of congestion. Then, depending on its own utility function, each link determines the amount of new packets it will generate. Such mechanisms have been shown to provide attractive performance characteristics (see for example [5], [7], [16]).

Definition 2 (CONGESTION CONTROL MECHANISM): For every link $l \in \mathcal{L}$, given $Q_l[t]$, the congestion control mechanism picks $X_l[t]$ as

$$X_l[t] = \min \left\{ M, U^{r-1} \left(\frac{Q_l[t]}{K} \right) \right\}, \quad (3)$$

where $M \in [1, \infty)$ and $K \in (0, \infty)$ are finite constants.

M is included in (3) with practical considerations in mind to bound the average number of packets generated in a single time slot. K is a design parameter critically affecting the performance of the whole mechanism. Intuitively, K can be interpreted as an aggressiveness measure of the congestion control mechanism: the greater the K the more packets are generated by the congestion control mechanism for the same level of the queues. The impact of K on the performance will be made explicit in our analysis.

It is not difficult to see that (3) implies that

$$\mathbf{X}[t] = \arg \max_{\mathbf{y} \in [0, M]^L} \mathbb{E} \left[K \sum_{l \in \mathcal{L}} U_l(y) - \langle \mathbf{y}, \mathbf{Q}[t] \rangle \mid \mathbf{Q}[t] \right], \quad (4)$$

which is a compact form of (3). While (3) reveals the decentralized nature of the congestion controller, (4) is useful in the analysis, which is the subject of the next section.

V. ANALYSIS OF THE JOINT SCHEDULER, CONGESTION CONTROLLER MECHANISM

In this section, we investigate the stability and optimality characteristics of the joint scheduling, congestion-control mechanism. In particular, we show that the network will be stable under the joint mechanism, and we explicitly characterize the proximity of the achieved mean link rates under the joint mechanism to the optimal rate allocation \mathbf{x}^* in terms of the randomized scheduler parameters (δ, γ, ψ) . In a relevant work [4], this question is answered for the case when $\delta > 0$, and $\gamma = \psi = 0$. Here, we provide the non-trivial extension to the case when $\delta, \gamma, \psi > 0$.

To that end, let us first define ε -relaxed stability region for an arbitrary $\varepsilon \in (0, 1)$, denoted by $\overline{\mathcal{S}}(\varepsilon)$, as

$$\overline{\mathcal{S}}(\varepsilon) \triangleq \{(1 - \varepsilon)\mathbf{x} : \mathbf{x} \in \overline{\mathcal{S}}\}.$$

Thus $\overline{\mathcal{S}}(\varepsilon)$ is a compressed version of the stability region $\overline{\mathcal{S}}$ by a factor of $(1 - \varepsilon)$. We also define the ε -optimal rate allocation $\mathbf{x}^*(\varepsilon)$ as

$$\mathbf{x}^*(\varepsilon) = \arg \max_{\mathbf{x} \in \overline{\mathcal{S}}(\varepsilon)} \sum_{l \in \mathcal{L}} U_l(x_l).$$

Notice that we have $\lim_{\varepsilon \rightarrow 0} \mathbf{x}^*(\varepsilon) = \mathbf{x}^*$. Since $M \geq 1$, we have $\overline{\mathcal{S}} \subset [0, M]^L$, and therefore $\mathbf{x}^*(\varepsilon) \in [0, M]^L$.

The next theorem proves the stability of the network under the joint mechanism, and quantifies the proximity of the achieved rate allocation to the optimal allocation in terms of the parameters (δ, γ, ψ) of the imperfect randomized scheduler.

Theorem 2: The network operating under the joint Randomized Scheduler (δ, γ, ψ) and Congestion Controller yields

$$\lim_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{t=0}^{\tau} \mathbb{E} [\|\mathbf{Q}[t]\|_1] \leq KB_1 \quad (5)$$

for some bounded B_1 , where $\|\mathbf{Q}\|_1 \triangleq \sum_{l \in \mathcal{L}} Q_l$.

Moreover, the proximity of the mean flow rates $\bar{\mathbf{x}} = (\bar{x}_l)_l$ to the ε -optimal rate $\mathbf{x}^*(\varepsilon)$ is characterized as

$$\sum_{l \in \mathcal{L}} U_l(\bar{x}_l) \geq \sum_{l \in \mathcal{L}} U_l(x_l^*(\varepsilon)) - \frac{B_2}{K}, \quad (6)$$

for any $\varepsilon > \left(\gamma + 2\sqrt{\frac{\psi}{\delta}} \right)$, where B_2 is a bounded constant.

Proof: We provide an analysis of the joint congestion control-imperfect randomized scheduler mechanism via a Lyapunov analysis. Let $\mathbf{Y} = (\mathbf{Q}, \mathbf{S})$ be a new state. It is not difficult to see that $\{\mathbf{Y}[t]\}$ under the joint mechanism describes a Markov Chain.

Consider the quadratic Lyapunov function

$$V(\mathbf{Y}) \triangleq \frac{1}{2} \sum_{l \in \mathcal{L}} Q_l^2 = \frac{1}{2} \|\mathbf{Q}\|_2^2, \quad (7)$$

where $\|\mathbf{Q}\|_2^2 \triangleq \langle \mathbf{Q}, \mathbf{Q} \rangle$.

For notational convenience, let us define the W -step conditional mean drift of the Lyapunov function starting at time t as

$$\Delta_t^{(W)} V \triangleq \mathbb{E}[V(\mathbf{Y}[t+W]) - V(\mathbf{Y}[t]) \mid \mathbf{Y}[t]],$$

for each $W \in \{1, 2, \dots\}$. We first bound the single-step conditional mean of $V(\cdot)$

$$\begin{aligned} \Delta_t^{(1)} V &= \frac{1}{2} \|\mathbf{Q}[t+1]\|_2^2 - \|\mathbf{Q}[t]\|_2^2 \\ &= \frac{1}{2} \|\mathbf{Q}[t] - \mathbf{S}[t] + \mathbf{X}[t] + \mathbf{U}[t]\|_2^2 - \|\mathbf{Q}[t]\|_2^2 \\ &= \langle \mathbf{Q}[t], \mathbf{X}[t] - \mathbf{S}[t] \rangle + \frac{1}{2} \|\mathbf{X}[t] - \mathbf{S}[t]\|_2^2 + \langle \mathbf{U}[t], \mathbf{X}[t] \rangle \\ &\quad + \langle \mathbf{U}[t], \mathbf{Q}[t] - \mathbf{S}[t] \rangle + \frac{1}{2} \|\mathbf{U}[t]\|_2^2 \\ &\leq \langle \mathbf{Q}[t], \mathbf{X}[t] - \mathbf{S}[t] \rangle + \frac{1}{2} \|\mathbf{X}[t] - \mathbf{S}[t]\|_2^2 + \langle \mathbf{U}[t], \mathbf{X}[t] \rangle, \end{aligned}$$

where the last inequality follows from the fact that $\langle \mathbf{U}[t], \mathbf{Q}[t] - \mathbf{S}[t] \rangle = -\|\mathbf{U}[t]\|_2^2$. Also noting that $U_l[t] \leq 1$ since $S_l[t] \leq 1$, and that $X_l[t] \leq M$ from (3), we can further upper-bound the mean drift as

$$\Delta_t^{(1)} V \leq \langle \mathbf{Q}[t], \mathbf{X}[t] - \mathbf{S}[t] \rangle + \left(\frac{M^2}{2} + M \right) L.$$

We can extend the same analysis to W -step mean drift to get

$$\begin{aligned}
& \Delta_t^{(W)} V \\
& \leq \sum_{w=0}^{W-1} \mathbb{E} [\langle \mathbf{Q}[t+w], \mathbf{X}[t+w] - \mathbf{S}[t+w] \rangle | \mathbf{Y}[t]] \\
& \quad + \left(\frac{M^2}{2} + M \right) WL \\
& \stackrel{(a)}{=} \sum_{w=0}^{W-1} \mathbb{E} \left[K \sum_{l \in \mathcal{L}} U_l(X_l[t+w]) | \mathbf{Y}[t] \right] \\
& \quad - \sum_{w=0}^{W-1} \mathbb{E} [\mathbb{E} [K \sum_{l \in \mathcal{L}} U_l(X_l[t+w]) - \\
& \quad \quad \langle \mathbf{Q}[t+w], \mathbf{X}[t+w] \rangle | \mathbf{Y}[t+w]] | \mathbf{Y}[t]] \quad (8) \\
& \quad - \sum_{w=0}^{W-1} \mathbb{E} [\langle \mathbf{Q}[t+w], \mathbf{S}[t+w] \rangle | \mathbf{Y}[t]] \\
& \quad + \left(\frac{M^2}{2} + M \right) WL
\end{aligned}$$

where (a) follows from adding and subtracting $K \sum_{l \in \mathcal{L}} U_l(X_l[t+w])$ to the previous expression, and from conditioning in (8). Notice from (4) that our congestion controller is designed to optimize the inner expectation in (8). Thus, we can get a lower bound on this inner expectation by noting that $\mathbf{x}^*(\varepsilon) \in [0, M]^L$, and therefore (4) implies that

$$\begin{aligned}
& \mathbb{E} \left[K \sum_{l \in \mathcal{L}} U_l(X_l[t+w]) - \langle \mathbf{Q}[t+w], \mathbf{X}[t+w] \rangle | \mathbf{Y}[t+w] \right] \\
& \quad \geq K \sum_{l \in \mathcal{L}} U_l(x_l^*(\varepsilon)) - \langle \mathbf{Q}[t+w], \mathbf{x}^*(\varepsilon) \rangle.
\end{aligned}$$

We substitute this as a bound to (8) and organize the terms to obtain

$$\begin{aligned}
\Delta_t^{(W)} V & \leq \sum_{w=0}^{W-1} \mathbb{E} \left[K \sum_{l \in \mathcal{L}} U_l(X_l[t+w]) | \mathbf{Y}[t] \right] \\
& \quad + \sum_{w=0}^{W-1} \mathbb{E} [\langle \mathbf{Q}[t+w], \mathbf{x}^*(\varepsilon) - \mathbf{S}[t+w] \rangle | \mathbf{Y}[t]] \quad (9) \\
& \quad - WK \sum_{l \in \mathcal{L}} U_l(x_l^*(\varepsilon)) + \left(\frac{M^2}{2} + M \right) WL
\end{aligned}$$

Using Theorem 1, after replacing the exogenous arrival rate with $\mathbf{x}^*(\varepsilon)$ and assuming that $\rho \triangleq \varepsilon - \gamma + 2\sqrt{\frac{\psi}{\delta}} > 0$, we can write

$$\begin{aligned}
\Delta_t^{(W)} V & \leq \sum_{w=0}^{W-1} \mathbb{E} \left[K \sum_{l \in \mathcal{L}} U_l(X_l[t+w]) | \mathbf{Y}[t] \right] \\
& \quad - WK \sum_{l \in \mathcal{L}} U_l(x_l^*(\varepsilon)) \\
& \quad - \rho W \|\mathbf{Q}[t]\|_1 + \left(\frac{M^2}{2} + M \right) WL.
\end{aligned}$$

Then, we take the expectation of both sides of the previous inequality to eliminate the conditioning, and then take the

telescoping sum of P such consecutive W -step drifts to obtain $\mathbb{E} [V(\mathbf{Y}[PW]) - V(\mathbf{Y}[0])]$

$$\begin{aligned}
& \leq \sum_{k=0}^{PW-1} \mathbb{E} \left[K \sum_{l \in \mathcal{L}} U_l(X_l[k]) \right] - KPW \sum_{l \in \mathcal{L}} U_l(x_l^*(\varepsilon)) \\
& \quad + \left(\frac{M^2}{2} + M \right) WLP - \rho \sum_{p=0}^{P-1} \mathbb{E} [\|\mathbf{Q}(pW)\|_1]. \quad (10)
\end{aligned}$$

Noting that $V(\mathbf{y}) \geq 0$ for all $\mathbf{y} \geq \mathbf{0}$, and re-arranging and bounding some terms in this expression, we can obtain

$$\begin{aligned}
& \frac{1}{P} \sum_{p=0}^{P-1} \mathbb{E} [\|\mathbf{Q}(pW)\|_1] \\
& \leq \frac{K \sum_{l \in \mathcal{L}} U_l(M) + \left(\frac{M^2}{2} + M \right) WL + \frac{\mathbb{E}[V(\mathbf{Y}[0])]}{PW}}{\rho}.
\end{aligned}$$

Also noting that $Q_l[t+k] \leq Q_l[t] + kM$ for each $l \in \mathcal{L}$, we can write

$$\begin{aligned}
\frac{1}{W} \sum_{k=0}^{W-1} \mathbb{E} [\|\mathbf{Q}(pW+k)\|_1] & \leq \mathbb{E} [\|\mathbf{Q}(pW)\|_1] \\
& \quad + \frac{M(W-1)L}{2},
\end{aligned}$$

which, when combined the previous inequality, yields

$$\begin{aligned}
& \lim_{P \rightarrow \infty} \frac{1}{PW} \sum_{k=0}^{PW-1} \mathbb{E} [\|\mathbf{Q}(k)\|_1] \\
& \leq \lim_{P \rightarrow \infty} \frac{1}{P} \sum_{p=0}^{P-1} \mathbb{E} [\|\mathbf{Q}(pW)\|_1] + \frac{M(W-1)L}{2} \\
& \leq \frac{\left(K \sum_{l \in \mathcal{L}} U_l(M) + \left(\frac{M^2}{2} + M \right) WL \right)}{\rho} + \frac{M(W-1)L}{2} \\
& \leq B_1 K,
\end{aligned}$$

for the appropriate choice of B_1 , which proves (5).

Next, we re-organize the terms of (10) in a different way to obtain

$$\begin{aligned}
& \frac{1}{WP} \sum_{k=0}^{PW-1} \mathbb{E} \left[\sum_{l \in \mathcal{L}} U_l(X_l[k]) \right] \\
& \quad \geq \sum_{l \in \mathcal{L}} U_l(x_l^*(\varepsilon)) - \frac{L}{K} \left(\frac{M^2}{2} + M \right)
\end{aligned}$$

Also note that

$$\begin{aligned}
& \frac{1}{WP} \sum_{k=0}^{PW-1} \mathbb{E} \left[\sum_{l \in \mathcal{L}} U_l(X_l[k]) \right] \\
& \leq \sum_{l \in \mathcal{L}} U_l \left(\frac{1}{WP} \sum_{k=0}^{PW-1} \mathbb{E} [X_l[k]] \right)
\end{aligned}$$

by Jensen's inequality.

Since $\bar{x}_l = \lim_{P \rightarrow \infty} \frac{1}{WP} \sum_{k=0}^{PW-1} \mathbb{E}[X_l[k]]$, in the limit as $P \rightarrow \infty$, we can write

$$\sum_{l \in \mathcal{L}} U_l(\bar{x}_l) \geq \sum_{l \in \mathcal{L}} U_l(x_l^*(\varepsilon)) - \frac{L}{K} \left(\frac{M^2}{2} + M \right),$$

which proves (6) with $B_2 := L \left(\frac{M^2}{2} + M \right)$. ■

This shows that our joint mechanism, despite the imperfections of the randomized scheduler, is guaranteed to achieve provably good performance, where the degree of optimality is explicitly described as functions of the imperfections of the underlying scheduling mechanism. Since M and K are design parameters, they can be picked to make the utility achieved by our mechanism be at least as good as $\mathbf{x}^*(\varepsilon)$.

Note that the imperfection measured by ψ (associated with the `Compare` operation) can be much more detrimental to the optimality as opposed to the imperfection measured by γ (associated with the `Compute` operation). This is reasonable since an error in the `Compare` operation may result in a completely undesirable scheduler in which case it will take a long time for the iterative scheduling mechanism to recover from such a mistake. In comparison, an error in the `Compute` operation will result in a scheduler with a performance that is guaranteed to be within a $(1 - \gamma)$ fraction of the current scheduler. Thus, its influence on the performance is milder. This suggests that in the design of distributed algorithms within this framework, one should be considerably more careful against errors in the `Compare` operation as opposed to errors in the `Compute` operation.

In the case of multi-hop traffic, the joint scheduling, congestion control mechanism needs to be extended by adding a backpressure routing scheme (see e.g. [23], [17], [21], [12], [7]) which utilizes queue-length levels to dynamically route multi-hop traffic towards a direction of less congestion. The results of Theorem 2 can be generalized to this scenario.

VI. CONCLUSIONS

We investigated the stability and optimality properties of a class of joint scheduling, congestion control mechanisms. The randomized and imperfect nature of the schedulers enables distributed implementations. The imperfections allow for errors in the computation and communication operations that need to be performed during scheduling. We explicitly characterized the impact of these imperfections on the optimality characteristics of the joint mechanism. Our results revealed the sensitivity of the framework to different types of errors, which will be useful in the design of new distributed algorithms.

REFERENCES

- [1] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, R. Vijayakumar, and P. Whiting. Scheduling in a queueing system with asynchronously varying service rates, 2000. Bell Laboratories Technical Report.
- [2] L. Bui, A. Eryilmaz, R. Srikant, and X. Wu. Joint asynchronous congestion control and distributed scheduling for wireless networks. *Proceedings of IEEE Infocom* 2006.
- [3] P. Chaporkar, K. Kar, and S. Sarkar. Throughput guarantees through maximal scheduling in wireless networks. In *Proceedings of the Allerton Conference on Control, Communications and Computing*, 2005.
- [4] A. Eryilmaz, Ozdaglar A, and E. Modiano. Polynomial complexity algorithms for full utilization of multi-hop wireless networks. In *Proceedings of IEEE Infocom*, 2007.
- [5] A. Eryilmaz and R. Srikant. Fair resource allocation in wireless networks using queue-length based scheduling and congestion control. In *Proceedings of IEEE Infocom*, volume 3, pages 1794–1803, Miami, FL, March 2005.
- [6] 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*, 14:1514–1524, August 2006.
- [7] A. Eryilmaz and R. Srikant. Resource allocation of multi-hop wireless networks. In *Proceedings of International Zurich Seminar on Communications*, February 2006.
- [8] A. Eryilmaz, R. Srikant, and J. R. Perkins. Stable scheduling policies for fading wireless channels. *IEEE/ACM Transactions on Networking*, 13:411–425, April 2005.
- [9] F. P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.
- [10] F. P. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [11] 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.
- [12] X. Lin and N. Shroff. The impact of imperfect scheduling on cross-layer rate control in multihop wireless networks. In *Proceedings of IEEE Infocom*, Miami, FL, March 2005.
- [13] S. H. Low and D. E. Lapsley. Optimization flow control, I: Basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, 7:861–875, December 1999.
- [14] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556–567, October 2000.
- [15] E. Modiano, D. Shah, and G. Zussman. Maximizing throughput in wireless networks via gossiping. In *ACM SIGMETRICS/IFIP Performance*, 2006.
- [16] M.J. Neely, E. Modiano, and C. Li. Fairness and optimal stochastic control for heterogeneous networks. In *Proceedings of IEEE Infocom*, pages 1723–1734, Miami, FL, March 2005.
- [17] M.J. Neely, E. Modiano, and C.E. Rohrs. Dynamic power allocation and routing for time varying wireless networks. In *Proceedings of IEEE Infocom*, pages 745–755, April 2003.
- [18] S. Sanghavi, L. Bui, and R. Srikant. Distributed link scheduling with constant overhead. In *ACM Sigmetrics*, 2007.
- [19] S. Shakkottai and A. Stolyar. Scheduling for multiple flows sharing a time-varying channel: The exponential rule. *Translations of the AMS, Series 2*, A volume in memory of F. Karpelevich, 207:185–202, 2002.
- [20] R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhäuser, Boston, MA, 2004.
- [21] A. Stolyar. Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm. *Queueing Systems*, 50(4):401–457, 2005.
- [22] L. Tassiulas. Linear complexity algorithms for maximum throughput in radio networks and input queued switches. In *Proceedings of IEEE Infocom*, pages 533–539, 1998.
- [23] 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*, 36:1936–1948, December 1992.
- [24] X. Wu and R. Srikant. Regulated maximal matching: A distributed scheduling algorithm for multi-hop wireless networks with node-exclusive spectrum sharing. In *Proceedings of IEEE Conference on Decision and Control*, 2005.
- [25] X. Wu and R. Srikant. Bounds on the capacity region of multi-hop wireless networks under distributed greedy scheduling. In *Proceedings of IEEE Infocom*, 2006.
- [26] H. Yaiche, R. R. Mazumdar, and C. Rosenberg. A game-theoretic framework for bandwidth allocation and pricing in broadband networks. *IEEE/ACM Transactions on Networking*, 8(5):667–678, October 2000.