Stable scheduling policies for fading wireless channels[*]

Atilla Eryilmaz, R. Srikant
University of Illinois
{eryilmaz,rsrikant}@uiuc.edu

James R. Perkins
Boston University
perkins@bu.edu

Abstract

We study the problem of stable scheduling for a class of wireless networks. The goal is to stabilize the queues holding information to be transmitted over a fading channel. Few assumptions are made on the arrival process statistics other than the assumption that their mean values lie within the capacity region and that they satisfy a version of the law of large numbers. We prove that, for any mean arrival rate that lies in the capacity region, the queues will be stable under our policy. Moreover, we show that it is easy to incorporate imperfect queue length information and other approximations that can simplify the implementation of our policy.

1 Introduction

Consider a wireless network, where data collected in N separate queues are to be transmitted over a common medium that is time varying. Several well-known models fit into this definition. Two such examples would be the downlink and the uplink scenarios of a cellular environment.

A *scheduling policy* is an allocation of service rates to the various queues, under the constraint that, at each time instant and each channel state, the set of allocated rates lies within some allowable set of rates. The set of allowable rates for each channel state is assumed to be a convex region. Our goal is to find a scheduling strategy which stabilizes the system using only queue length information and the current channel state (i.e., without knowing channel or arrival statistics).

Stable scheduling policies for wireless systems without time-varying channels were first studied in [16]. In fact the model in [16] can also be thought of as a model for a high-speed input-queued switch. Systems with time-varying channels, but limited to the case of ON and OFF channels were studied in [17, 14]. More general channel models have been studied by others recently [1, 13, 12, 7]. We generalize the class of scheduling policies considered in [17, 14, 1, 13]. Further, we allow imperfect queue length information and prove the stability of policies that reduce computational complexity. These class of policies for wireless networks are natural extensions of those studied in [15, 5] for high-speed switches. Our proof uses a quadratic Lyapunov

1

function argument along the lines of the proofs in [16, 6]. We also refer the reader to [2, 11] for a geometric approach to scheduling problems.

In the context of time-varying wireless channels with many users, our work is an example of exploiting multiuser diversity to maximize the capacity of the system. Here, we try to maximize the throughput of the system without the knowledge of system statistics. Alternatively, one can formulate a fair resource allocation problem where each user is allocated a certain fraction of the system resources according to some criterion [18, 9]. The approaches in [18, 9] are not throughput optimal, but are *fair* according to some appropriate notion of fairness. Our work and other related work assume that the channel is time-varying and attempt to exploit this feature. In [19], an interesting technique to induce time variations in channels which may not be inherently time-varying is discussed.

The rest of the paper is organized as follows. Section 2 describes the system model and presents a statement of the problem we consider in this paper, the scheduling policy and assumption on the arrival and channel processes. We state the main theorem, which establishes the stability of the system, in Section 3. Section 4 gives several useful applications of the policy operating both in uplink and downlink scenarios. Several properties of the set of scheduling policies are illustrated through simulations in Section 5. Conclusions and further directions are provided in Section 6. And finally, the proofs of the theorems are collected in Section 7.

2 System model

Consider a wireless network where N data streams are to be transmitted over a single fading channel. An example of such a network can be a single transmitter sending data to N receivers (the downlink in a cellular system) as depicted in Figure 1 or N transmitters sending data to a single receiver (the uplink) as shown in Figure 2. We assume that the arriving bits are stored in N separate queues, one for each data stream. Assuming that time is slotted, the evolution of the i^{th} queue is described by the following equation:

$$x_i[k + 1] = (x_i[k] + a_i[k] - \eta_i[k])^+, \tag{1}$$

where $a_i[k]$ is the number of bits arriving to Queue i at time k and $\eta_i[k]$ is the number of bits from Queue i that are served at time k. One can also ensure that $x_i[k]$ takes on only discrete values in our model. This would be more realistic when transmission can take place only in units of packets, for example. In such a case, there may be wasted service even when $x_i[k] > 0$. To allow for this, we can rewrite the evolution of the queue lengths given by (1) as

$$x_i[k + 1] = x_i[k] + a_i[k] - \eta_i[k] + u_i[k],$$

where $u_i[k]$ is a positive quantity, which denotes the wasted service provided to the i^{th} queue during slot k.

The state of the channel is assumed to be fixed within a time slot, but is allowed to vary from one slot to the next. Let J be the number of possible channel states. Suppose that the channel is in State j at time k, then $\{\eta_i[k]\}$ is constrained to be in some region $\mathcal{S}_j$. Thus, $\mathcal{S}_j$ identifies the allowable set of rates at which the queues can be
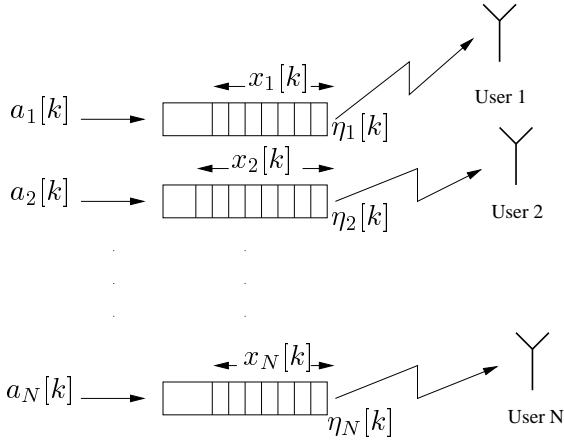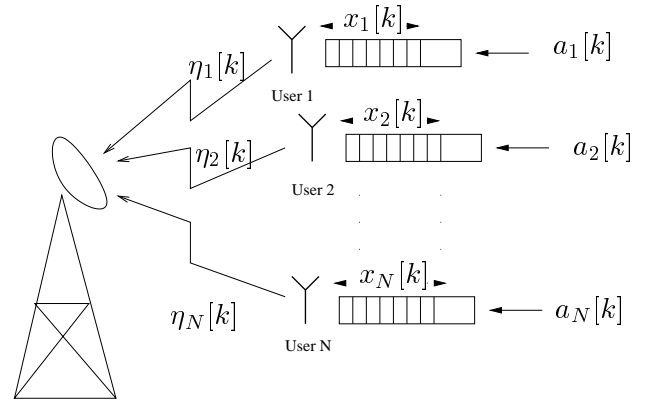
Figure 1: Downlink model

Figure 2: Uplink model

drained when the channel is in State j. For now, we can simply visualize $\mathcal{S}_j$ to be a bounded, convex region such as the broadcast channel capacity region [8]. Precise conditions on the allowable set of rates, the channel state process and arrival processes will be given later.

In this paper, we consider the following class of scheduling policies: at any time k, given the current channel state $s[k]$, the scheduler chooses a service rate vector $\vec{\mu} = (\mu_1, \cdots, \mu_N)' \in \mathcal{S}_{s[k]}$ that satisfies:

$$\vec{\mu} = \arg \max_{\vec{\eta} \in \mathcal{S}_{s[k]}} \sum_{i=1}^{N} f_i(x_i[k]) \eta_i, \tag{2}$$

where $f_i : [0, \infty) \to [0, \infty)$ are functions that satisfy the following conditions:

- $f_i(x)$ is a nondecreasing, continuous function with $\lim_{x \to \infty} f_i(x) = \infty$.

- Given any $M_1 > 0$, $M_2 > 0$ and $0 < \epsilon < 1$, there exists an $\mathcal{X} < \infty$, such that for all $x > \mathcal{X}$, we have

$$(1 - \epsilon) f_i(x) \leq f_i(x - M_1) \leq f_i(x + M_2) \leq (1 + \epsilon) f_i(x), \qquad \forall i. \tag{3}$$

Examples of the the functions $f_i(\cdot)$ that satisfy (3) are $f_i(x) = (K_i x)^\alpha$ for any $K_i \in [0, \infty)$ and $\alpha \in [0, \infty)$, or $f_i(x) = e^{\sqrt{x}} - 1$. Note that the exponential function $f_i(x) = e^{\alpha x} - 1$ for any fixed $\alpha > 0$ does not satisfy (3).

As we will see later, for various reasons, it may be difficult to implement the policy (2). For example, the queue length information may be delayed, the maximization involved may be too complex or one may wish to use waiting times, instead of queue lengths, to choose the service rates. We will show later that, in all such cases, the scheduling policy will satisfy the following property.

Property 1. *Given any ρ and ζ such that $0 \leq \rho, \zeta < 1$, there exists a $B > 0$ such that the scheduling policy satisfies the following condition: at any time k, with probability greater than $(1 - \rho)$, the scheduler chooses a service rate vector $\vec{\mu} = (\mu_1, \cdots, \mu_N)' \in \mathcal{S}_{s[k]}$ that satisfies:*

$$\sum_{i=1}^{N} f_i(x_i[k]) \mu_i \geq (1 - \zeta) \max_{\vec{\eta} \in \mathcal{S}_{s[k]}} \sum_{i=1}^{N} f_i(x_i[k]) \eta_i \tag{4}$$

3

whenever $\|\vec{x}[k]\| > B$, where $\vec{x}[k] := (x_1[k], \cdots, x_N[k])'$, and $s[k] \in \{1, \cdots, J\}$ is the channel state in time slot k. ◇

Thus, for the purpose of establishing stability, we will consider scheduling policies that satisfy the above property.

In the following subsections, we state the various assumptions that we make on the arrival and channel processes, the admissible rate regions, and state a fact about the scheduling policy which will be useful for later proofs.

2.1 The channel state process

1) The channel state process has a stationary distribution, where the stationary probability of being in state j, $j \in \{1, \cdots, J\}$, is denoted by π_j. Further, we assume $\pi_j > 0 \; \forall j$.

2) Let us denote the state of the channel at time n by $s[n]$. Note that $s[n]$ can take any one of the J possible values. Given any $\epsilon > 0$, there exists a positive integer $\bar{M}$ such that $\forall M > \bar{M}$

$$E\left[\left\|\pi_j - \frac{1}{M}\sum_{n=k}^{k+M-1} I_{s[n]=j}\right\|\right] < \epsilon, \tag{5}$$

for any $k > 0$ and $j = 1, 2, \ldots, J$.

2.2 The achievable rate regions $\{\mathcal{S}_j\}$

1) Consider any region $\mathcal{S}_j$ and any $\vec{\eta} \in \mathcal{S}_j$. There exists an $\hat{\eta}$ such that $\eta_i \leq \hat{\eta}$. In other words, each of the regions $\mathcal{S}_j$ is bounded.

2) Each of the regions $\mathcal{S}_j$, $j = 1, 2, \ldots, J$ is convex.

3) For each $j = 1, 2, \ldots, J$, the following is true: if $\{\eta_1, \eta_2, \ldots, \eta_i, \ldots, \eta_N\} \in \mathcal{S}_j$, then $\{\eta_1, \eta_2, \ldots, 0, \ldots, \eta_N\} \in \mathcal{S}_j$ for all $i = 1, 2, \ldots, N$.

4) Fix a channel state j. Given any fixed $A \in \Re^+$, $\forall \epsilon > 0$, $\exists R < \infty$ such that, for any $\vec{x} \in \left(\Re^N\right)^+$, $\vec{y} \in \left(\Re^N\right)^+$ satisfying $|x_i - y_i| < A$, $i = 1, \cdots, N$ and $\|\vec{x}\|, \|\vec{y}\| > R$, we have

$$\left|\sum_{i=1}^{N} f_i(x_i)\mu_i(j, \vec{x}) - \sum_{i=1}^{N} f_i(y_i)\mu_i(j, \vec{y})\right| < \epsilon \sum_{i=1}^{N} f_i(x_i)\mu_i(j, \vec{x}),$$

where $\vec{\mu}$ is determined according to our scheduling policy. (We use the notation $\vec{\mu}(j, \vec{x})$ to denote the vector of service rates when the current channel state is j and the current queue length vector is $\vec{x}$.)

2.3 The arrival processes

1) The arrivals to each Queue i form a stationary process, with a mean denoted by $\lambda_i := E[a_i[1]]$.

2) Define

$$\bar{\mathcal{S}} = \{\vec{\eta} : \vec{\eta} = \sum_{j=1}^{J} \pi_j \vec{\eta}^{(j)} \text{ and } \vec{\eta}^{(j)} \in \mathcal{S}_j, \forall j\}.$$

The vector of mean arrival rates $\vec{\lambda}$ is such that there exists $\vec{\nu} \in \bar{\mathcal{S}}$ satisfying $\nu_i > \lambda_i \, \forall i$.

3) Given any $\epsilon > 0$, there exists a positive integer $\bar{M}$ such that $\forall M > \bar{M}$

$$E\left[\left|\frac{1}{M}\sum_{n=k}^{k+M-1} a_i[n] - \lambda_i\right|\right] < \epsilon, \qquad \forall i. \tag{6}$$

4) Finally, f_i and $a_i[k]$ should satisfy

$$\lim_{A \to \infty} \sum_{i=1}^{N} A f_i(A) P(a_i[1] > A) \;=\; 0.$$

2.4 Observation on the Scheduling Policy

Claim 1. *If the scheduling policy satisfies Property 1, then with probability greater than* $(1-\rho)$, $\vec{\bar{\mu}} := \sum_{j=1}^{J} \pi_j \vec{\mu}(j, \vec{x})$

satisfies

$$\sum_{i=1}^{N} f_i(x_i) \bar{\mu}_i \geq (1-\zeta) \max_{\vec{\eta} \in \bar{\mathcal{S}}} \sum_{i=1}^{N} f_i(x_i) \eta_i \tag{7}$$

for all $\|\vec{x}\| > B$.

Proof: See Section 7.1 of the Appendix for the proof. $\diamond$

3 Stability of the stochastic model

We state the main result of the paper in the following theorem.

Theorem 1. *For sufficiently small values of* $\zeta, \rho \geq 0$, *the system is stable in the mean under the policy described in Section 2, i.e.,*

$$\limsup_{p \to \infty} \frac{1}{p} \sum_{k=0}^{p-1} E\left[\|\vec{f}(\vec{x}[k])\|_2\right] < \infty, \tag{8}$$

where $\|\vec{f}(\vec{x})\|_2 := \left(\sum_{i=1}^{N} f_i^2(x_i)\right)^{1/2}$.

Proof: The proof of the theorem is in Appendix 7.2. $\diamond$

In addition to the assumptions presented in Section 2, if we further assume that the queue lengths $\{x_i[k]\}$ can only take values in $\{0, 1, 2, \ldots\}$, and that the arrival and channel state processes make the queueing system an aperiodic Markov chain with a single communicating class, then the stability-in-the-mean property further implies that the Markov chain is *positive recurrent* [6].

An example of a system that is positive recurrent is one where the arrival and channel state processes satisfy the following conditions:

- *The arrival process to each queue is a Markov-modulated Poisson process.* In other words, the arrival process is one of many states, the stochastic process describing the evolution of these states is a countable state, aperiodic Markov chain with a single communicating class. Further, in each arrival state, the number of arrivals generated is a Poisson random variable. The mean of the Poisson random variable can be state dependent.

- *The channel state process is a countable state, aperiodic Markov chain with a single communicating class.*

Under the above conditions, if we enlarge the definition of the state to be

$$(\text{channel state, states of the arrival processes, queue lengths}),$$

then the state transition process is a Markov chain. Further, due to the Poisson nature of the arrivals, it is easy to see that the queue lengths can empty from any initial state with non-zero probability, and that from any state with empty queues, it is possible to reach any other state with non-zero probability. Thus, the Markov chain has a single communicating class. Further, it is also easy to see that the system can remain in any state with empty queues for more than one time instant with non-zero probability. Thus, the Markov chain is also aperiodic. Finally, we note that the arrival and channel state processes are short-range dependent and, thus, satisfy the law-of-large-number type conditions (5) and (6) in Section 2.

3.1 Instability

If the mean arrival rate vector $\vec{\lambda}$ lies outside the average achievable rate region $\bar{\mathcal{S}}$, then the system will be unstable. To prove this, we make use of the *Strict Separation Theorem,* [3, Proposition B. 14] which states that since $\vec{\lambda}$ is a point that does not belong to the convex set $\bar{\mathcal{S}}$, there exists a vector $\vec{\beta}$ such that

$$\sum_{i=1}^{N} \beta_i \mu_i \leq \sum_{i=1}^{N} \beta_i \lambda_i - \delta,$$

for some $\delta > 0$. Further, due to the fact that $\lambda_i \geq 0$, $\forall i$, and Assumption (3) in Section 2.2, a little thought shows that β_i can be chosen to be non-negative, with at least one β_i positive. Given this $\vec{\beta}$, we define the Lyapunov function,

$$W(\vec{x}) := \sum_{i=1}^{N} \beta_i x_i.$$

Then, from a drift analysis, we have

$$
\begin{aligned}
E\left(W(\vec{x}[k+1]) - W(\vec{x}[k]) \mid \vec{x}[k]\right) &= \sum_{i=1}^{N} \beta_i E\left(x_i[k+1] - x_i[k] \mid \vec{x}[k]\right) \\
&= \sum_{i=1}^{N} \beta_i E\left(a_i[k] - \eta_i[k] + u_i[k] \mid \vec{x}[k]\right) \\
&\geq \sum_{i=1}^{N} \beta_i \left(\lambda_i - E(\eta_i[k] \mid \vec{x}[k])\right) \\
&\geq \delta,
\end{aligned}
$$

6

which implies that $E(W(\vec{x}[k])) \to \infty$ as $k \to \infty$ and therefore, the system is not stable-in-the-mean.

3.2   Non-convex set of allowable rates

There are many practical systems where the set of rate vectors that can be used by the scheduler may not be convex. An example is a cellular downlink with a TDMA protocol. We will refer to the set of rate vectors that can actually be implemented by the scheduler as the *set of allowable rates*. Then we define the *achievable rate region* to be the convex hull of the set of allowable rates for each channel state j. Now suppose we use a policy of the form

$$\vec{\mu}[k] \in \arg \max_{\vec{\eta}[k] \in \mathcal{S}_{s[k]}} \sum_{i=1}^{N} f_i(x_i[k]) \eta_i[k]. \tag{9}$$

We claim that this policy will yield a set of optimal rate vectors, at least one element of which is in the set of allowable rate vectors.

To see that this claim is true, we first note that, from the definition of a convex hull, any rate, $\vec{\nu}$, which belongs to the convex hull can be written as a convex combination of some allowable rate vectors, $\{\vec{c}^n\}$, i.e.,

$$\vec{\nu} = \sum_{n=1}^{L} \alpha_n \vec{c}^n,$$

where $L > 0$ is an integer and $\sum_{n=1}^{L} \alpha_n = 1$ with $\alpha_n > 0 \, \forall n$. If for any state j, and some queue length vector $\vec{x}$, the set of rates which maximizes (9) does not contain any of the allowable rate vectors, then we must have at least one achievable rate vector, $\vec{\nu}$, such that it satisfies

$$\sum_{i=1}^{N} f_i(x_i) \nu_i \ > \ \sum_{i=1}^{N} f_i(x_i) c_i^n \qquad \forall n \in \{1, \ldots, L\},$$

which in turn implies

$$\sum_{n=1}^{L} \alpha_n \sum_{i=1}^{N} f_i(x_i) c_i^n \ > \ \sum_{i=1}^{N} f_i(x_i) c_i^n \qquad \forall n \in \{1, \ldots, L\}.$$

However, the last equation cannot be true since the convex combination of a set of positive numbers cannot be strictly larger than each of them. Hence, by contradiction, it follows that at least one solution to the maximization problem in (9) must belong to the set of allowable rates.

4   Applications

The scheduling policy given in (2) is a generalization of the policy examined in [16, 14, 1]. In a later section, we show through simulations that general functions of the form $f_i(\cdot)$ can be very useful in controlling queue lengths. In this section, we show that the introduction of the parameters, ρ, ζ, enables the application of the policy to scenarios where instantaneous queue length information is not available or the scheduler has computational limitations.

4.1 Infrequent or Delayed Queue Length Updates

Consider the multiple access uplink scenario, where each of the N users maintains an infinite length queue, holding information to be transmitted to the base station over a fading multiple access channel. This scenario is depicted in Figure 2. In this case, it may not be reasonable to expect the queue length to be updated at each time slot. To reduce the amount of information transferred between the transmitters and the base station, suppose that each transmitter updates the queue length only once every T time slots. Let $\hat{x}_i[k]$ denote the estimate of the queue length of the i^{th} queue at time k. In other words, $\hat{x}_i[k]$ is the last update of the queue length, prior to time k, received by the base station from Transmitter i. Further, suppose that at each time slot k, the base station allocates a service rate vector that satisfies

$$\arg \max_{\vec{\eta}[k] \in \mathcal{S}_{s[k]}} \sum_{i=1}^{N} f_i(\hat{x}_i[k]) \eta_i[k]. \tag{10}$$

In the following theorem, we show that this policy satisfies Property 1 in Section 2.

Theorem 2. *Suppose that the scheduler is only allowed to sample the queue length information once every T slots (i.e. $\hat{x}_i[nT + l] = x_i[nT]$ for $l = 0, 1, \ldots, T - 1$ and $n = 0, 1, \ldots$), and it uses this sampled value as the current queue length to determine the service rates according to (10), then the system is stable-in-the-mean.*

Proof: Since the mean arrival rate to each of the queues is finite, given any $\rho \in (0, 1)$, we can find $A < \infty$ such that

$$Prob\{a_i[k] \leq A \; \forall i\} > (1 - \rho).$$

Let us consider two sampling instants k and $k + T$. Consider any $n \in \{0, \cdots, T - 1\}$, and define the following quantities for each channel state j :

$$\vec{\mu}^*(j, \vec{x}[k + n]) \in \arg \max_{\vec{\mu} \in \mathcal{S}_j} \sum_{i=1}^{N} f_i(x_i[k + n]) \mu_i(j, \vec{x}[k + n])$$

$$\vec{\hat{\mu}}(j, \vec{\hat{x}}[k + n]) = \vec{\hat{\mu}}(j, \vec{x}[k]) \in \arg \max_{\vec{\mu} \in \mathcal{S}_j} \sum_{i=1}^{N} f_i(x_i[k]) \mu_i(j, \vec{x}[k]).$$

Observe that for any $i \in \{1, 2, \ldots, N\}$, and $n \in \{0, 1, \ldots, T - 1\}$, we have

$$x_i[k + n] - TA \leq \hat{x}_i[k + n] = x_i[k] \leq x_i[k + n] + T\hat{\eta} \qquad \text{w.p. } (1 - \rho).$$

Moreover, due to Assumption (4) of Section 2.2, given any $\zeta \in (0, 1)$, we can find a bounded region around the origin, outside of which the following inequality holds

$$\sum_{i=1}^{N} f_i(\hat{x}_i[k + n]) \hat{\mu}_i(j, \vec{\hat{x}}[k + n]) = \sum_{i=1}^{N} f_i(x_i[k]) \hat{\mu}_i(j, \vec{x}[k])$$

$$\geq (1 - \zeta) \sum_{i=1}^{N} f_i(x_i[k + n]) \mu_i^*(j, \vec{x}[k + n]) \qquad \text{w.p. } (1 - \rho).$$

8

Therefore, this policy satisfies Property 1. ◇

There are alternative ways to update the queue length information instead of periodic sampling. For example, the scheduler may sample each queue with some probability at each time instant. In this case, given any $\epsilon > 0$, we can find a T such that the probability that all queues have been updated at least once in the past T slots is greater than $1 - \epsilon$. By making ϵ arbitrarily small and following the lines of the proof of previous theorem, we can again prove the stability of the system.

While periodic sampling and random sampling would ensure stability, they may result in poor delay performance. An alternative sampling technique which may be particularly useful with bursty arrivals, is to update the queue length information for each queue whenever the absolute value of the difference between the current length and the last update exceeds some threshold. Along the lines of the proof of the previous theorem, we can again show that this policy is stable. However, we will show through simulations later that this update mechanism reduces the mean queueing delay as compared to random or periodic sampling.

Finally, we note that delayed queue length updates can also be cast in the same framework as above.

4.2 Reducing computational complexity

Typically, the allowed set of power levels at a mobile or a base station is a finite set. Consequently, the set of allowable rates will be finite for each channel state. In this case, as discussed earlier, the achievable rate region in each state is the convex hull of the set of allowable rates in the state. The convex hull would be a convex polyhedron and a policy of the form

$$\arg \max_{\vec{\eta}[k] \in \mathcal{S}_{s[k]}} \sum_{i=1}^{N} f_i(x_i[k]) \eta_i[k]$$

would involve an optimization over the vertices of the convex polygon. The complexity issues arising due to this has been addressed in the context of high-speed switches in [15, 5].

In this section, we show that the solutions proposed in [15, 5] for high-speed switches are also applicable to wireless networks with time-varying connectivity and more general functions $f_i(x_i)$ than the ones considered in [15, 5]. The basic idea behind the solution in [5] is to perform a Hamiltonian walk over the set of allowable rates (or more simply, over the vertices of the convex polygon of allowable rates) for each state, and store in memory, the *best* schedule so far in each channel state. (In our context, performing a Hamiltonian walk corresponds to maintaining a list of allowable rates and visiting each possible rate vector in a fixed order. Once all the rate vectors are visited, the list is again scanned from the beginning.) This way, at each step we only need to compare two values, which is a significant reduction of complexity. In the following, we present the algorithm and prove its stability.

Algorithm A: Assume that the current channel state is $s[k] = j$ and let $k^{(j)}_{[-d]}$ denote the d^{th} time slot before k when the channel was in State j. Let $\mathcal{L}^{(j)}$ denote the number of available rate vectors we need to choose from when the channel state is j, and $\vec{\eta}[k^{(j)}_{[-d]}]$ denote the rate vector at time $k^{(j)}_{[-d]}$. (Note that $k^{(j)}_{[-1]}$ denotes the last time before k when the channel state was j. In the following, we will omit the subscript $[-d]$ when $d = 1$.) Then the

algorithm is comprised of the repetition of the following steps:

(1) $\vec{h}[k]$ = next rate vector visited by the Hamiltonian walk in the current channel state $s[k] = j$.

(2)

$$\vec{\eta}[k] = \arg \max_{\vec{\nu} \in \{\vec{\eta}[k^{(j)}], \vec{h}[k]\}} \sum_{i=1}^{N} f_i(x_i[k]) \nu_i.$$

Remark: Even if Step (1) of the algorithm is modified to choosing a rate vector *randomly* from the set of possible rate vectors as in [15], the following theorem will continue to hold. Since the proof is essentially the same, only *Algorithm A* is considered here.

Theorem 3. *The policy defined by* Algorithm A *satisfies Property 1 of Section 2 and hence Theorem 1 continues to hold.*

Proof: For any $\rho_1 > 0$, we can find an $A < \infty$ such that $Prob\{a_i[k] \leq A \; \forall i\} > (1 - \rho_1)$. Then, note that for any $k, n \geq 0$, we have

$$x_i[k] - n\hat{\eta} \leq x_i[k+n] \leq x_i[k] + nA \qquad \forall i, \qquad \text{w.p.} \quad (1 - \rho_1)^n,$$

which in turn implies that $\forall \epsilon_1 > 0$ and any allowable rate vector $\vec{\nu}$, we can find a large enough bounded region, outside of which, the following holds:

$$\left| \sum_{i=1}^{N} f_i(x_i[k+n]) \nu_i - \sum_{i=1}^{N} f_i(x_i[k]) \nu_i \right| \leq \epsilon_1 \sum_{i=1}^{N} f_i(x_i[k]) \nu_i, \qquad \text{w.p.} \quad (1 - \rho_1)^n. \tag{11}$$

The assumptions on the channel state process imply that the probability of not visiting a state j within M slots goes to zero as M tends to infinity. Therefore, for any $\rho_2 > 0$, we can find a finite M, such that the probability of not visiting a state j is less than ρ_2, and this is true for any $j \in \{1, \cdots, J\}$.

Consider any slot m and, without loss of generality, assume that the channel state at that slot is j. Also let us define $\mathcal{L} := \max_j \mathcal{L}^{(j)} < \infty$. Let $\vec{\eta}^*[m]$ be a rate vector that satisfies the following at time m :

$$\vec{\eta}^*[m] \in \arg \max_{\vec{\eta}[k] \in \mathcal{S}_j} \sum_{i=1}^{N} f_i(x_i[m]) \eta_i[m].$$

Then, as observed in [5], due to the nature of the Hamiltonian walk, there exists a time slot $m' \in [m_{[-\mathcal{L}]}^{(j)}, m]$ for which the channel state satisfies $s[m'] = j$, and the rate vector visited by the Hamiltonian walk at that time is $\vec{h}[m'] = \vec{\eta}^*[m]$. In other words, we can write $m' = m_{[-t]}^{(j)}$ for some $t \in \{0, \ldots, \mathcal{L}\}$.

Moreover, repeating the argument that the channel state process visits state j at least once in M slots with probability $(1 - \rho_2)$, we have

$$m - M\mathcal{L} \leq m_{[-\mathcal{L}]}^{(j)} \leq m \qquad \text{w.p.} (1 - \rho_2)^{\mathcal{L}}. \tag{12}$$

10

Combining this observation with the properties of m' we have $m' \in [m - M\mathcal{L}, m]$ with probability $(1 - \rho_2)^{\mathcal{L}}$. Then, Step (2) of the algorithm enables us to write

$$\sum_{i=1}^{N} \eta_i[m'] f_i(x_i[m']) \geq \sum_{i=1}^{N} \eta_i^*[m] f_i(x_i[m']) \tag{13}$$

$$\geq (1 - \epsilon_1) \sum_{i=1}^{N} \eta_i^*[m] f_i(x_i[m]) \qquad \text{w.p. } (1 - \rho_2)^{\mathcal{L}} (1 - \rho_1)^{M\mathcal{L}}, \tag{14}$$

where the last inequality follows from (11) and (12).

Also note that, for any n and any $\epsilon_2 > 0$, we can find a bounded region around the origin, outside of which we have: with probability $(1 - \rho_1)^M (1 - \rho_2)$,

$$(1 - \epsilon_2) \sum_{i=1}^{N} f_i(x_i[n^{(s[n])}]) \eta_i[n^{(s[n])}] \leq \sum_{i=1}^{N} f_i(x_i[n]) \eta_i[n^{(s[n])}] \leq \sum_{i=1}^{N} f_i(x_i[n]) \eta_i[n], \tag{15}$$

where the first inequality follows from (11), and the second inequality is due to Step (2) of Algorithm A. We continue as follows:

$$\sum_{i=1}^{N} f_i(x_i[m]) \eta_i[m]$$

$$\geq (1 - \epsilon_2)^t \sum_{i=1}^{N} f_i(x_i[m']) \eta_i[m'] \qquad \text{w.p. } (1 - \rho_1)^{tM} (1 - \rho_2)^t \tag{16}$$

$$\geq (1 - \epsilon_1)(1 - \epsilon_2)^{\mathcal{L}} \sum_{i=1}^{N} f_i(x_i[m]) \eta_i^*[m] \qquad \text{w.p. } (1 - \rho_1)^{2M\mathcal{L}} (1 - \rho_2)^{2\mathcal{L}}. \tag{17}$$

In the previous steps, (16) follows from (15), and step (17) follows from (14) and the fact that $|t| \leq \mathcal{L}$. Hence, given any $\zeta > 0$, and $\rho > 0$, we can find $\rho_1, \rho_2 > 0$ satisfying $(1 - \rho_1)^{M\mathcal{L}}(1 - \rho_2)^{2\mathcal{L}} \geq (1 - \rho)$ and find parameters $\epsilon_1, \epsilon_2 > 0$ satisfying $(1 - \epsilon_1)(1 - \epsilon_2)^{\mathcal{L}} \geq (1 - \zeta)$, which in turn yields

$$\sum_{i=1}^{N} f_i(x_i[m]) \eta_i[m] \geq (1 - \zeta) \sum_{i=1}^{N} f_i(x_i[m]) \eta_i^*[m] \qquad \forall m, \qquad \text{w.p. } (1 - \rho),$$

outside a closed, bounded region around the origin. Therefore, we have proved that Algorithm A satisfies Property 1. $\diamond$

We note that, while *Algorithm A* lowers the computational complexity considerably, it adds a memory requirement non-existent previously. To see this, first observe that the algorithm keeps track the rate vectors for each fading state. Since the number of fading states increases exponentially with the number of users, so will the memory requirement. Evidently such a memory requirement is necessary to assure stability if computational complexity is an issue. Alternatively, one can tradeoff between memory and computational complexity requirements, by using Algorithm A only in some channel states and performing exact computations in other channel states. For example, in some channel states, the SNR may be too low so that a packet cannot be transmitted

unless the transmit power is above some threshold. This may limit the number of candidate optimal solutions, thus automatically reducing the computation required. For these states, one can use exact computation, whereas, for other states, reduced complexity algorithms could be used.

4.3 Downlink

In the downlink scenario, a single transmitter maintains N infinite length queues, one for each receiver and sends this information over a fading channel as depicted in Figure 1. Hence, the scheduler at the transmitter has immediate access to the queue length values at any time, and we assume that it knows the current channel state. Then at the beginning of each time slot, say k, it chooses the service rate vector $\vec{\mu}[k]$, such that

$$\vec{\mu}[k] \in \arg \max_{\vec{\eta}[k] \in \mathcal{S}_{s[k]}} \sum_{i=1}^{N} f_i(x_i[k])\eta_i[k], \tag{18}$$

whenever $\|\vec{x}\| > B$ for any fixed value $B < \infty$. Then the results of Section 3 hold for this system. This is a generalization of the result in [1] where the result has been proved for the case $f_i(x)$ of the form $(K_i x)_i^{\alpha}$. As we will see through simulations later, our generalization allows for better queue length performance.

4.4 Waiting Times

Instead of the current queue length information, the scheduler may alternatively use the delay experienced by the packets within the queues as its input. To incorporate this into our model, we first let $w_i[k]$ denote the waiting time of the head of the line (H.O.L.) packet in the i^{th} queue at time k. In the following, we consider a policy that chooses the service rate vector $\vec{\mu}[k]$, such that

$$\vec{\mu}[k] \in \arg \max_{\vec{\eta}[k] \in \mathcal{S}_{s[k]}} \sum_{i=1}^{N} h_i(w_i[k])\eta_i[k], \tag{19}$$

for continuous, non-decreasing functions $\{h_i(\cdot)\}$, satisfying $h_i(0) = 0$ and $\lim_{x \to \infty} h_i(x) = \infty$.

Observe that given any $\mathcal{W} < \infty$ and $\epsilon \in (0, 1)$, we can find a value $\mathcal{X}$ such that for all $x_i[k] > \mathcal{X}$, we have $w_i[k] \geq \mathcal{W}$ with probability greater than or equal to $(1 - \epsilon)$. In this subsection alone, we make stronger assumptions on the arrival process than we had previously used. Let us assume that given any $\delta > 0$, we can find a $\Gamma < \infty$ so that, with probability greater than $(1 - \delta)$, we have

$$\lambda_i w_i[k] - \Gamma \sqrt{w_i[k]} \leq x_i[k] \leq \lambda_i w_i[k] + \Gamma \sqrt{w_i[k]} \qquad \forall i, k.$$

Thus, we assume that the arrival process obeys a central limit theorem (CLT). Conditions on the arrival process under which it obeys a CLT are given, for example, in [4].

From the CLT assumption, it is easy to see that $w_i[k]$ can be upper and lower bounded as follows:

$$\alpha_i x_i[k] - K\sqrt{x_i[k]} \leq w_i[k] \leq \alpha_i x_i[k] + K\sqrt{x_i[k]} \qquad \forall i, k,$$

for appropriate values of $\{\alpha_i\}$ and K. Then, we have

$$h_i(\alpha_i x_i[k] - K\sqrt{x_i[k]}) \le h_i(w_i[k]) \le h_i(\alpha_i x_i[k] + K\sqrt{x_i[k]}) \qquad \forall i, k.$$

We assume that, given any $\epsilon > 0$ and a finite K, we can find a bounded region around the origin, outside of which the following holds:

$$(1 - \epsilon)h_i(\alpha_i x_i) \le h_i(\alpha_i x_i - K\sqrt{x_i}) \le h_i(\alpha_i x_i + K\sqrt{x_i}) \le (1 + \epsilon)h_i(\alpha_i x_i), \qquad \forall i. \qquad (20)$$

Now, if we define $f_i(x_i[k]) := h_i(\alpha_i x_i[k])$, then (20) implies that outside a bounded region around the origin,

$$(1 - \epsilon)f_i(x_i[k]) \le h_i(w_i[k]) \le (1 + \epsilon)f_i(x_i[k]) \qquad \forall i.$$

This property enables us to bound (19) as follows

$$\max_{\vec{\eta}[k] \in \mathcal{S}_{s[k]}} \sum_{i=1}^{N} h_i(w_i[k])\eta_i[k] \ \ge \ (1 - \epsilon) \max_{\vec{\eta}[k] \in \mathcal{S}_{s[k]}} \sum_{i=1}^{N} f_i(x_i[k])\eta_i[k],$$

which shows that (4) is satisfied.

An example of $h_i(\cdot)$ that satisfy (20) is $h_i(w) := K_i w^{\alpha_i} \, \forall \alpha_i > 0$. To see that this is in fact the case, we can write

$$\begin{aligned} h_i(x_i \pm K\sqrt{x_i}) &= K_i(x_i \pm K\sqrt{x_i})^{\alpha_i} \\ &= K_i x_i^{\alpha_i}\left(1 \pm \frac{K}{\sqrt{x_i}}\right)^{\alpha_i}, \end{aligned}$$

where the term in the parenthesis can be made arbitrarily close to 1 by choosing x_i large enough. Hence, (20) holds.

Another example of $h_i(\cdot)$ that satisfies (20) is $h_i(w) := \exp(w^{\alpha_i}) \, \forall \alpha_i \in (0, 0.5)$. To justify this, we proceed as follows:

$$\begin{aligned} h_i(x_i \pm K\sqrt{x_i}) &= \exp((x_i \pm K\sqrt{x_i})^{\alpha_i}) \\ &= \exp(x_i^{\alpha_i}(1 \pm \frac{K}{\sqrt{x_i}})^{\alpha_i}) \\ &\approx \exp(x_i^{\alpha_i}(1 \pm \frac{K\alpha_i}{\sqrt{x_i}})) \\ &= \exp(x_i^{\alpha_i})\exp(\pm K\alpha_i x_i^{(\alpha_i - 0.5)}), \end{aligned}$$

where the second exponent can be made arbitrarily small by choosing x_i large enough if $\alpha_i \in (0, 0.5)$.

In the previous example of $h_i(\cdot)$, if $\alpha_i \in [0.5, 1)$, then the system described above is not necessarily stable in the mean. To guarantee stability, we have to strengthen the conditions on the arrival process. Suppose that we consider leaky bucket type of arrivals, i.e., the number of arrivals between time s and t, denoted by $A(s, t)$, satisfies $A(s, t) \le \rho(t - s) + \sigma \, \forall 0 \le s < t$ with positive constants ρ, σ. There are many examples of stationary

stochastic processes that satisfy such a constraint when the arrival process is further peak-rate constrained. We refer the reader to [10] for one such example. The leaky-bucket constraint limits the burstiness of the arrivals, which in turn enables us to upper-bound the difference between x_i and $\lambda_i w_i$, with high probability, by a large enough constant. Hence, with probability greater than $(1 - \delta)$, we have

$$\beta_i x_i[k] - K \leq w_i[k] \leq \beta_i x_i[k] + K \qquad \forall i, k,$$

for appropriate values of $\{\beta_i\}$ and K. Then, we have

$$h_i(\beta_i x_i[k] - K) \leq h_i(w_i[k]) \leq h_i(\beta_i x_i[k] + K) \qquad \forall i, k.$$

If we define $f_i(x) := h_i(\beta_i x)$, and assume that $\{f_i(\cdot)\}$ satisfies (3), then the previous set of inequalities holds. Then it is easy to see that $\{h_i(\cdot)\}$ of the form $h_i(w) := \exp(w^{\alpha_i}) \, \forall \alpha_i \in [0.5, 1)$ satisfies (20).

5 Simulations

In this section, the performance of the class of scheduling policies described in Section 2 is illustrated through simulations. For ease of exposition, all the simulations consider the case of two users.

Experiment 1: In this experiment, we illustrate the effect of using different sets of functions $\{f_i(x)\}$ on the queue length evolution of the system. The average arrival rates to the two queues are $\lambda_1 = 50$ and $\lambda_2 = 50$. The channel is in one of five states and the achievable rates η_{1s} and η_{2s} for the two queues when the channel is in State s satisfy the following equation:

$$\eta_{1s}^2 + \eta_{2s}^2 \leq r_s \sqrt{\lambda_1^2 + \lambda_2^2}.$$

The values for r_s were chosen to be 0.3, 0.7, 1, 1.3 and 1.7. The channel state process is a discrete-time Markov chain, such that, given that the Markov chain is in a particular state, the probability of a transition to any other state (including itself) is 0.2. The number of arrivals to each queue in each time slot has a Poisson distribution. The arrivals to the two queues are independent from time slot to slot, and are independent of each other and of the channel state process.

The functions we used in the simulations are of the form $f_i(x_i) = (K_i x_i)^\alpha$. Using functions of this form, the queue length evolutions are illustrated in Figure 3. Observe that as α increases, the queue length vector gets closer to the line $K_1 x_1 = K_2 x_2$ from any initial condition, after which it stays around it and moves toward the origin. Such a behavior empirically shows that we can choose the functions $\{f_i(\cdot)\}$ so that priorities may be assigned to different queues without sacrificing stability. Moreover, this analysis justifies the fairness property inherent in the policy, again through empirical methods. We note that the rule in [13] drives the system to the line where the $K_i x_i$ for all the users are equal and is shown to be pathwise optimal in [12].
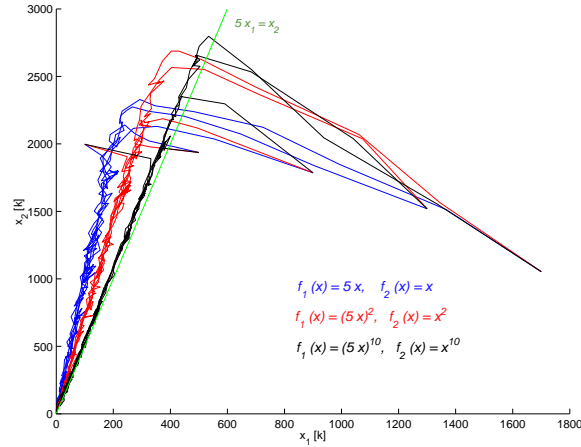
Figure 3: Queue length evolutions in the stochastic model.

Experiment 2: In this experiment, the channel setting is kept the same as in Experiment 1, but the arrivals to both of the queues are chosen to be independent, Bernoulli distributed random variables having mean λ_i for Queue i with a peak value of 500 packets per slot. In the case of such bursty arrivals, this experiment compares the performance of two queue length update mechanisms:

- periodically updating the queue length information (we refer to this policy as the *Periodic Update Policy*) and

- updating it either when the number of arrivals exceeds a certain limit since the last update or if the time since the last update has exceeded a threshold (we refer to this policy as the *Enhanced Update Policy*).

The stability analysis of such systems was done in Section 4.1.

In the Periodic update policy, the values of the queue lengths are updated once in every 200 slots in our simulations. When the arrivals are bursty, such a strategy does not track the queue length values very closely. Even though, we have proved that the system will be stable in the mean, the packets might experience large delays.

On the other hand, if we instead use the Enhanced update strategy, which guarantees that the queue length information is updated at least once in every 200 slots and also whenever current queue length differs from the most recent update by more than a certain threshold (50 in our example), then we get better performance under a bursty traffic, since we can track the actual queue length values more closely.

We note that the average achievable rate region is a quarter circle of radius 50. We define *traffic intensity* to be the ratio $\sqrt{\lambda_1^2 + \lambda_2^2}/50$. Figure 4 examines the effect of varying the traffic intensity for the two update policies, where the sampling time for the periodic update policy and the bound for the threshold update policy are both taken to be 50. It is seen that under heavy load, the Enhanced update policy yields much better average delay performance.
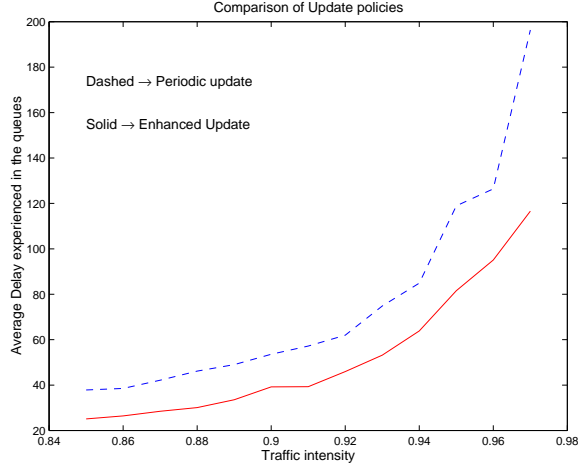
Figure 4: Delay characteristics of the two queue length update strategies defined in Experiment 2, with varying load.

Experiment 3: In this experiment, our goal is to study the ability of our class of policies to minimize buffer overflow. For this purpose, we consider the following measure of performance:

$$P(x_1 > B_1) + P(x_2 > B_2),$$

where B_1 and B_2 are both taken to be 5000. In other words, the objective is the sum of the overflow probabilities in the two queues. We wish to study the impact of the choice of $\{f_i(\cdot)\}$ on the above performance measure.

We use the following heuristic to choose $\{f_i(\cdot)\}$. From Markov's inequality, we have $P(x_i > B_i) \leq \frac{g_i(x_i/B_i)}{g_i(1)}$ for any positive, increasing function $g_i(\cdot)$. Since we do not have expressions for the overflow probability, we choose functions $\{f_i(\cdot)\}$ that we expect would minimize the above upper bounds on the overflow probability. To do this, we choose $f_i(x_i) = g_i'(x_i)$. The heuristic behind this is that, in the fluid model (see proof of Theorem 1 in the Appendix), at each instant, we attempt to minimize the time derivative of $\sum_i g_i(x_i)$. Thus, it is natural to choose $\{f_i(\cdot)\}$ to be the derivatives of the upper bound expressions.

1. The first policy chooses $\vec{\mu}$ such that

$$\sum_{i=1}^{n} \mu_i \frac{x_i}{B_i^2}$$

   is maximized over all $\vec{\mu}$ within the current achievable rate region. This corresponds to $g_i(x_i) = (x_i/B_i)^2$.

2. The second policy chooses $\vec{\mu}$ such that

$$\sum_{i=1}^{n} \mu_i \exp(\frac{x_i}{B_i})^{0.5}$$

   is maximized. This corresponds to

$$g_i(x_i) = \frac{1}{B_i} \int_0^{x_i} e^{\sqrt{y}} dy.$$

16

3. For comparison, we also study the performance of the $EXP-Q$ rule which was shown to be throughput-optimal in [13], and has recently been shown to be pathwise optimal in the heavy traffic regime [12]. This policy chooses $\vec{\mu}$ such that

$$\sum_{i=1}^{n} \mu_i e^{\left(\frac{\frac{x_i}{B_i}}{1+\sqrt{\frac{1}{n}\left(\frac{x_1}{B_1}+\frac{x_2}{B_2}+\cdots+\frac{x_n}{B_n}\right)}} \right)}$$

is maximized over all $\vec{\mu}$ within the current achievable rate region.

The channel state process is allowed to vary among five equiprobable states as in Experiments 1 and 2. The initial queue length values are chosen as $(x_1[0], x_2[0]) = (1000, 1000)$. Note that this choice is arbitrary and we ran the simulations for 10 million iterations so that the transient effects will be negligible. The arrival rates are chosen as $(\lambda_1, \lambda_2) = (50, 50)$. In Figure 5, the performance of the three policies are compared as a function of traffic intensity. The range of traffic intensities for which the fraction of overflow duration is on the order of 10^{-3} to 10^{-2} is shown in the figure. It can be seen that with increasing traffic intensity, the second policy, which uses an exponential function to determine the rates, enables a 10% to 20% reduction in the overflow probability compared to the first policy. Somewhat surprisingly, the second policy has a 5% to 10% smaller overflow probability compared to the EXP-Q rule. Even though the EXP-Q has been proved to be pathwise optimal in [12], it is an asymptotic result in the heavy-traffic regime and so it is quite possible that another scheme could perform slightly better at traffic intensities within the boundary of the capacity region. This illustrates the fact that by suitably choosing of functions $f_i(x_i)$, system performance can be improved. However, we do not have a theoretical handle on how these $f_i(x_i)$ should be chosen given a requirement on the overflow probability. This is a subject of future research.
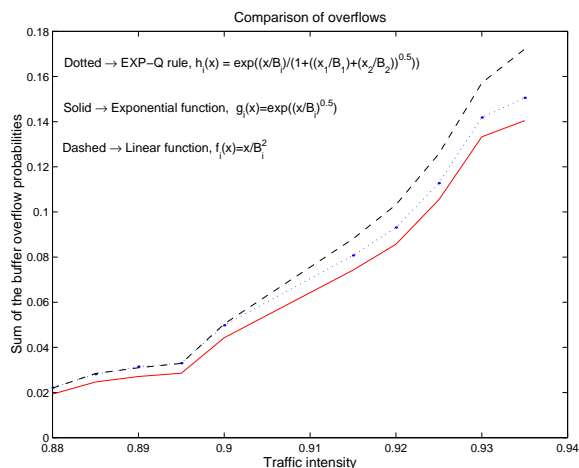

Figure 5: Comparison of the policies with increasing traffic intensity.
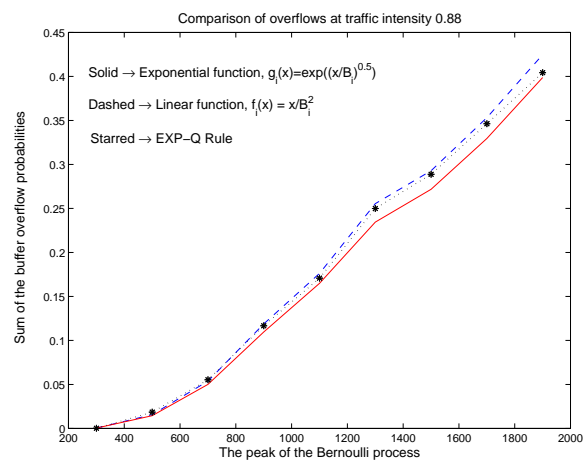

Figure 6: Comparison of the policies with increasing burstiness of the Bernoulli arrivals.

Figure 6 shows the effect of increasing the burstiness of the arrivals on the overflows of the two policies. We increase the burstiness by increasing the peak value, M, of the Bernoulli arrivals while keeping the mean

unchanged. Although the figure is plotted for the traffic intensity of 0.88, it is representative of other traffic intensities. Again, the exponential function gives a better performance than the linear function.

6 Conclusions

We have presented conditions on scheduling policies that guarantee stability for a large class of arrival and channel models. We have shown that the conditions are satisfied for a variety of policies that use probabilistic, periodic or otherwise scheduled queue length updates, policies that result in computational reduction and policies that use head-of-the-line waiting times. A line of future research would be the study of the interaction of such scheduling mechanisms with congestion control. Currently, we have assumed that the mean arrival rates lies within the achievable rate region. Congestion control naturally provides a mechanism to move the mean arrival rates within the achievable rate region. Thus, it would be natural to study the combination of scheduling and congestion control.

7 Appendix: Proofs

7.1 Proof of Claim 1 in Section 2.4

Note that solving the maximization in (7) is equivalent to

$$
\begin{aligned}
\max_{\{\vec{\eta} \in \mathcal{S}_j\}_{j=1}^J} \sum_{i=1}^N f_i(x_i) \sum_{j=1}^J \pi_j \eta_i \;=\;& \max_{\{\vec{\eta} \in \mathcal{S}_j\}_{j=1}^J} \sum_{j=1}^J \pi_j \sum_{i=1}^N f_i(x_i)\eta_i \\
\leq\;& \sum_{j=1}^J \pi_j \max_{\vec{\eta}(j) \in \mathcal{S}_j} \sum_{i=1}^N f_i(x_i)\eta_i(j) \\
\leq\;& \frac{1}{1-\zeta} \sum_{j=1}^J \pi_j \sum_{i=1}^N f_i(x_i)\mu_i(j).
\end{aligned}
$$

But $\vec{\mu} \in \bar{\mathcal{S}}$ and hence the previous upper bound is in fact achievable by $\vec{\mu}$. ⋄

7.2 Proof of Theorem 1

The stability of the class of scheduling policies is proved in several steps. We first consider a continuous-time model with constant arrival rates and a deterministic channel, and show that the system evolves towards a closed region around the origin in the state space (i.e., the space of queue length vectors). As we will see, this suggests a natural Lyapunov function to analyze the stability of the original discrete-time stochastic system. However, before we consider the stochastic system, we study a deterministic discrete-time system and show that the Lyapunov function decreases, except in a bounded region around the origin of the state space. Then, we consider the evolution of the Lyapunov function at time instants that are M steps apart, for some large M. This allows us

to use law-of-large-numbers type assumptions to view this system as being nearly deterministic and apply the results of the deterministic model to complete the proof of stability.

7.2.1 Deterministic model of the system

In this subsection, we assume that the arrival process to the i^{th} queue is deterministic and constant at each time slot, with the constant equal to the mean, λ_i, of the corresponding stochastic arrival processes, $a_i[k]$. Further, the evolution of each of the queues is assumed to be

$$x_i[k+1] \quad = \quad x_i[k] + \lambda_i - \bar{\eta}_i[k] + u_i[k], \tag{21}$$

where $\vec{\bar{\eta}}[k] = \vec{\bar{\mu}}(\vec{x}[k]), \; \vec{\bar{\mu}}(\vec{x}[k]) := \sum_{j=1}^{J} \pi_j \vec{\mu}(j, \vec{x}[k])$ and $u_i[k]$ is an upper-bounded, positive quantity, which denotes the wasted service provided to the i^{th} queue. Thus, $\vec{\bar{\mu}}(\vec{x}[k])$ can be interpreted as the average service provided to Queue i when the queue state is $\vec{x}[k]$, where the averaging is performed over the channel state process. In the following subsection, we state two lemmas, which will be used in the proof of Theorem 1.

Continuous-time model

In this model, time is no longer discrete, but is continuous, and the evolution of the queue lengths is governed by the following differential equation

$$\dot{x}_i(t) = \begin{cases} (\lambda_i - \bar{\mu}_i(t)), & x_i(t) > 0 \\ (\lambda_i - \bar{\mu}_i(t))^+, & x_i(t) = 0 \end{cases} \qquad i = 1, 2, \ldots, N. \tag{22}$$

Using the above facts, we will now show that we can find a Lyapunov function for the system (22), such that its derivative is negative.

Lemma 1. *Suppose at any time instant t, the service rate vector $\vec{\mu}(t)$ is chosen such that it satisfies*

$$\sum_{i=1}^{N} f_i(x_i(t))\mu_i(t) \quad \geq \quad (1-\zeta) \max_{\vec{\eta} \in \bar{\mathcal{S}}} \; \sum_{i=1}^{N} f_i(x_i(t))\eta_i,$$

where an upper bound on the parameter ζ is provided in the proof of this lemma. Consider the following Lyapunov function:

$$V(\vec{x}) = \sum_{i=1}^{N} g_i(x_i), \tag{23}$$

where $g_i'(x) = f_i(x)$. Then for some $\delta > 0$, we have

$$\dot{V}(\vec{x}) \leq -\|\vec{f}(\vec{x})\|_2 \delta, \tag{24}$$

holding for all $\vec{x}$.

19

Proof:

Consider

$$\dot{V}(\vec{x}) \;=\; \sum_{i=1}^{N} f_i(x_i)(\lambda_i - \bar{\mu}_i). \tag{25}$$

Let us denote $||\vec{f}(\vec{x})||_2 := (\sum_{i=1}^{N} f_i^2(x_i))^{1/2}$ and define $\cos(\theta_i) = \frac{f_i(x_i)}{||\vec{f}(\vec{x})||_2}$. Then using (4), (25) can be rewritten as

$$\dot{V}(\vec{x}) \;=\; ||\vec{f}(\vec{x})||_2 \left(\sum_{i=1}^{N} \cos(\theta_i)\lambda_i - \sum_{i=1}^{N} \cos(\theta_i)\bar{\mu}_i \right)$$

$$\leq\; ||\vec{f}(\vec{x})||_2 \left(\underbrace{\sum_{i=1}^{N} \cos(\theta_i)\lambda_i}_{=:K_\lambda} - \underbrace{\sum_{i=1}^{N} \cos(\theta_i)\tilde{\mu}_i}_{=:K_{\tilde{\mu}}} + \zeta \sum_{i=1}^{N} \cos(\theta_i)\tilde{\mu}_i \right) \tag{26}$$

where we define

$$\vec{\tilde{\mu}} = \arg\max_{\vec{\eta} \in \bar{\mathcal{S}}} \sum_{i=1}^{N} f_i(x_i)\eta_i = \arg\max_{\vec{\eta} \in \bar{\mathcal{S}}} \sum_{i=1}^{N} \cos(\theta_i)\eta_i \tag{27}$$

since $||\vec{f}(\vec{x})||_2 \geq 0$ is constant for a fixed $\vec{x}$. Let us consider the expression in (26). The maximization amounts to finding the point on the boundary of $\bar{\mathcal{S}}$ at which a line with a certain slope (determined by $\vec{\theta}$) is tangential to the boundary. Note that since $\vec{\lambda}$ is not on the boundary, any two lines with the same slope such that one passes through $(\lambda_1, \cdots, \lambda_N)$ and the other is tangent to the boundary of $\bar{\mathcal{S}}$, will have a difference of at least $\delta_j > 0$ in its intercept with the j^{th} axis. Choose $\Delta := \min(\delta_1, \cdots, \delta_N) > 0$.

If $\theta_i = \frac{\pi}{2}$, then $K_\lambda = K_{\tilde{\mu}} = 0$. So consider any other index j such that $\theta_j < \frac{\pi}{2}$, which implies $\cos(\theta_j) \geq \gamma_j$ for some $\gamma_j > 0$. Define $\Gamma = \max(\gamma_1, \cdots, \gamma_N) > 0$. Then the j^{th} intercepts are $\frac{K_\lambda}{\cos(\theta_j)}$ and $\frac{K_{\tilde{\mu}}}{\cos(\theta_j)}$ and we can write

$$K_\lambda - K_{\tilde{\mu}} \leq -\Delta \max_j \cos(\theta_j) \leq -\Delta\Gamma := -2\delta.$$

Hence we can conclude that if $\zeta < \frac{\delta}{N\hat{\eta}}$, then we have

$$\dot{V}(\vec{x}) \leq ||\vec{f}(\vec{x})||_2(-2\delta + \zeta N\hat{\eta}) \leq -\delta||\vec{f}(\vec{x})||_2. \tag{28}$$

$$\diamond$$

In what follows, we will use the above Lyapunov function to first show the stability of the deterministic, discrete-time model and to later show that the original stochastic system is stable.

Discrete-time model

Using the result of Lemma 1, we will now show that the Lyapunov function (23) applied to the system described by (21) has a negative drift outside a bounded region.

Lemma 2. *Consider any policy that satisfies Property 1. When such a policy is applied to the discrete-time system (21) with constant arrivals and time-invariant channel, for sufficiently large M, the M-step drift of the Lyapunov function V satisfies the following inequality:*

$$E\left[\Delta V^{(M)}(\vec{x}[k])\right] \leq -M\epsilon_M||\vec{f}(\vec{x}[k])||_2\mathcal{I}_{\vec{x}[k]\in\mathcal{B}_{(M)}^c}[k] + K_{(M)}[k]\mathcal{I}_{\vec{x}[k]\in\mathcal{B}_{(M)}[k]}, \tag{29}$$

for some $\epsilon_M > 0$, where $\Delta V^{(M)}(\vec{x}[k]) := V(\vec{x}[k+M]) - V(\vec{x}[k])$, $\mathcal{B}_{(M)}$ is a bounded region around the origin and $K_{(M)}$ is a finite constant, both dependent on $\vec{x}[k]$.

Proof:

$$\begin{aligned}\Delta V^{(1)}(\vec{x}[k]) &:= V(\vec{x}[k+1]) - V(\vec{x}[k]) \\ &= \sum_{i=1}^{N} f_i(y_i[k])(\lambda_i - \bar{\mu}_i[k] + u_i[k]),\end{aligned}$$

where $y_i[k]$ lies between $x_i[k]$ and $x_i[k+1]$ from Taylor's theorem. Then we get,

$$\Delta V^{(1)}(\vec{x}[k]) = \sum_{i=1}^{N} f_i(y_i[k])(\lambda_i - \bar{\mu}_i[k]) \tag{30}$$

$$+ \sum_{i=1}^{N} f_i(y_i[k])u_i[k]. \tag{31}$$

For (31) observe that if $x_i[k] > \hat{\eta}$, then $u_i[k] = 0$ and if $x_i[k] \leq \hat{\eta}$, then $u_i[k] \leq \hat{\eta}$. Hence, using the fact that f_i is nondecreasing, and $y_i[k] \leq x_i[k] + \lambda_i$, we get

$$\sum_{i=1}^{N} f_i(y_i[k])u_i[k] \leq \sum_{i=1}^{N} f_i(\lambda_i + \hat{\eta})\hat{\eta} =: C_1 < \infty.$$

Note that (30) can be bounded as

$$\sum_{i=1}^{N} f_i(y_i[k])(\lambda_i - \bar{\mu}_i[k]) \leq \sum_{i=1}^{N} f_i(x_i[k])(\lambda_i - \bar{\mu}_i[k]) \tag{32}$$

$$+ \sum_{i=1}^{N} |f_i(y_i[k]) - f_i(x_i[k])||\lambda_i - \bar{\mu}_i[k]|. \tag{33}$$

To upper-bound (32), we will make use of Lemma 1. Note that (32) is exactly in the same form as (25), except that in this case with probability ρ, (24) may not hold. However, by Property 1, ρ can be chosen small enough by making M large. Hence, we can upper-bound (32) by $-\frac{\delta}{2}||\vec{f}(\vec{x}[k])||_2$.

Due to the properties of $\{f_i(\cdot)\}$, (33) can be upper-bounded by $\gamma||\vec{f}(\vec{x})||_2$ outside a bounded, closed region. Hence, by choosing $\epsilon_1 := \frac{\delta}{2} - \gamma > 0$, we get the following result

$$E\left[\Delta V^{(1)}(\vec{x}[k])\right] \leq -\epsilon_1||\vec{f}(\vec{x}[k])||_2\mathcal{I}_{\vec{x}[k]\in\mathcal{B}_{(1)}^c}[k] + K_{(1)}[k]\mathcal{I}_{\vec{x}[k]\in\mathcal{B}_{(1)}[k]}, \tag{34}$$

where $\mathcal{I}_A$ denotes the indicator function of the event A, $\mathcal{B}_{(1)}[k]$ is the closed and bounded region around the origin and $K_{(1)}[k] < \infty$ is appropriately chosen.

Next we extend the previous analysis to examine the M-step drift

$$
\begin{aligned}
E\left[\Delta V^{(M)}(\vec{x}[k])\right] &= E\left[V(\vec{x}[k+M]) - V(\vec{x}[k])\right] \\
&= \sum_{i=0}^{M-1} E\left[\Delta V^{(1)}(\vec{x}[k+i])\right] \\
&\leq \sum_{i=0}^{M-1}\left[-\epsilon_1 ||\vec{f}(\vec{x}[k+i])||_2 \mathcal{I}_{\vec{x}[k+i]\in\mathcal{B}_{(1)}^c[k+i]} + K_{(1)}[k]\mathcal{I}_{\vec{x}[k+i]\in\mathcal{B}_{(1)}[k+i]}\right],
\end{aligned}
\tag{35}
$$

which follows from (34). We can write (35) as

$$
E\left[\Delta V^{(M)}(\vec{x}[k])\right] \leq -\epsilon_1\left(\sum_{i=0}^{M-1}||\vec{f}(\vec{x}[k+i])||_2\right)\mathcal{I}_{\vec{x}[k]\in\mathcal{B}_{(M)}^c[k]} + K_{(M)}[k]\mathcal{I}_{\vec{x}[k]\in\mathcal{B}_{(M)}[k]},
\tag{36}
$$

where $\mathcal{B}_{(M)}[k]$, and $K_{(M)}[k] < \infty$ are $M-$step equivalents of $\mathcal{B}_{(1)}[k]$, and $K_{(1)}[k] < \infty$.

Now, consider any $j \in \{1, \cdots, N\}$ and $n \in \{0, \cdots, M-1\}$. Due to the property of f given by (3), for any $\gamma_n \in (0,1)$, we have

$$
\begin{aligned}
f_j(x_j[k+n]) &\geq f_j(x_j[k] - n\eta^*) \\
&= (1-\gamma_n)f_j(x_j[k]),
\end{aligned}
$$

for $x_j[k]$ large enough. Taking squares, summing over j, and taking square roots yields

$$
\begin{aligned}
||\vec{f}(\vec{x}[k])||_2 &= \sqrt{\sum_{j=1}^{N} f_j^2(x_j[k])} \\
&\leq \frac{1}{1-\gamma_n}||\vec{f}(\vec{x}[k+n])||_2.
\end{aligned}
$$

Let us define $\gamma := \max_{n\in\{0,\cdots,M-1\}} \gamma_n < 1$, then we can easily write

$$
\sum_{i=0}^{M-1}||\vec{f}(\vec{x}[k+i])||_2 \geq M(1-\gamma)||\vec{f}(\vec{x}[k])||_2.
$$

Hence, if we denote $\epsilon_M := \epsilon_1(1-\gamma)$, we can upper-bound the M-step drift as

$$
E\left[\Delta V^{(M)}(\vec{x}[k])\right] \leq -M\epsilon_M ||\vec{f}(\vec{x}[k])||_2 \mathcal{I}_{\vec{x}[k]\in\mathcal{B}_{(M)}^c[k]} + K_{(M)}[k]\mathcal{I}_{\vec{x}[k]\in\mathcal{B}_{(M)}[k]},
\tag{37}
$$

with $\epsilon_M > 0$. ◇

7.2.2 Stochastic model

In the following proof, we will make use of the result of Section 7.2.1 even though the arrivals are now stochastic processes and the channel state is time-varying. To facilitate this, we denote the vectors of queue length, allocated service rates and the unused services, at any time n, under the deterministic model by $\vec{x}^d[n]$, $\vec{\mu}^d[n]$, and $\vec{u}[n]$, respectively. Let $\vec{x}^d[k] = \vec{x}$. Next, we write the M-step mean drift for the stochastic model. Recall that, in Section 7.2.1, we obtained an expression for the drift of the function V assuming that the arrivals are constant and the service provided at each time instant is an average (over the channel states) of the service that would have been provided had the channel been in a particular state. Now, for the stochastic arrival and channel model,

$$\Delta W^{(M)}(\vec{x}[k]) \quad := \quad E\left[V(\vec{x}[k+M]) - V(\vec{x}^d[k+M]) \quad \mid \quad \vec{x}[k] = \vec{x}\right] \tag{38}$$

$$+ E\left[V(\vec{x}^d[k+M]) - V(\vec{x}[k]) \quad \mid \quad \vec{x}[k] = \vec{x}\right]. \tag{39}$$

Observe that (39) can be upper-bounded using (29). Next, we consider (38). Note that we can write

$$x_i[k+M] \;=\; x_i[k] + \underbrace{\sum_{n=k}^{k+M-1} a_i[n]}_{=:A_i(k,M)} - \underbrace{\sum_{n=k}^{k+M-1} \mu_i(s[n],\vec{x}[n])}_{=:C_i(k,M)} + \underbrace{\sum_{n=k}^{k+M-1} u_i[n]}_{=:U_i(k,M)}$$

$$x_i^d[k+M] \;=\; x_i[k] + M\lambda_i - \underbrace{\sum_{n=k}^{k+M-1}\sum_{j=1}^{J} \pi_j \mu_i(j,\vec{x}^d[n])}_{=:C_i^d(k,M)} + \underbrace{\sum_{n=k}^{k+M-1} u_i^d[n]}_{=:U_i^d(k,M)}.$$

Hence, we can write

$$\Sigma_{\Delta,i} \quad := \quad x_i[k+M] - x_i^d[k+M]$$

$$= \quad \underbrace{A_i(k,M) - M\lambda_i}_{=:\Delta A_i} + \underbrace{C_i^d(k,M) - C_i(k,M)}_{=:\Delta C_i} + \underbrace{U_i(k,M) - U_i^d(k,M)}_{=:\Delta U_i}.$$

Therefore, (38) can be written as

$$E\left[V(\vec{x}[k+M]) - V(\vec{x}^d[k+M]) \;\mid\; \vec{x}[k] = \vec{x}\right] \;=\; E\left[\sum_{i=1}^{N} g_i(x_i[k+M]) - g_i(x_i^d[k+M]) \;\mid\; \vec{x}[k] = \vec{x}\right]$$

$$= \quad \sum_{i=1}^{N} E\left[f_i(z_i(k,M))\Sigma_{\Delta,i} \;\mid\; \vec{x}[k] = \vec{x}\right],$$

which can be further written as

$$= \quad \sum_{i=1}^{N} E\left[f_i(z_i(k,M))(A_i(k,M) - M\lambda_i) \;\mid\; \vec{x}[k] = \vec{x}\right] \tag{40}$$

$$+ \sum_{i=1}^{N} E\left[f_i(z_i(k,M))(C_i^d(k,M) - C_i(k,M)) \;\mid\; \vec{x}[k] = \vec{x}\right] \tag{41}$$

$$+ \sum_{i=1}^{N} E\left[f_i(z_i(k,M))(U_i(k,M) - U_i^d(k,M)) \;\mid\; \vec{x}[k] = \vec{x}\right] \tag{42}$$

where $z_i(k, M) := \alpha x_i^d[k+M] + (1-\alpha)x_i[k+M]$ for some $\alpha \in [0, 1]$. To upper-bound the above expression, we will consider two events, one when the arrivals to each of the queues are upper-bounded by a finite value A and the other, the complement of this event. Let us denote the first event by

$$\mathcal{A} = \{a_i[n] \leq A : 1 \leq i \leq N, k \leq n \leq (k+M-1)\}.$$

First, let us concentrate on (40). In the case when the event $\mathcal{A}$ occurs, we can upper-bound $z_i(k, M)$ as

$$z_i(k, M) \leq x_i[k] + \alpha\lambda_i M + (1-\alpha)AM.$$

Then we can upper-bound (40) as

$$\leq \sum_{i=1}^{N} E\left[|f_i(x_i[k] + \alpha\lambda_i M + (1-\alpha)AM)| \, |A_i(k,M) - M\lambda_i| \quad | \quad \vec{x}[k] = \vec{x}\right].$$

For any fixed M, and $\zeta_1 > 0$ it is possible to find an X_i such that for all $x_i > X_i$, we have

$$f_i(x_i + \alpha\lambda_i M + (1-\alpha)AM) \leq (1+\zeta_1)f_i(x_i).$$

Next, we define $E := \{i : x_i[k] > X_i\}$, which enables us to further upper-bound (40) as

$$\leq \quad (1+\zeta_1)\sum_{i \in E} f_i(x_i[k])ME\left[\left|\frac{A_i(k,M)}{M} - \lambda_i\right| \quad | \quad \vec{x}[k] = \vec{x}\right]$$

$$+ \sum_{i \in E^c} f_i(X_i)ME\left[\left|\frac{A_i(k,M)}{M} - \lambda_i\right| \quad | \quad \vec{x}[k] = \vec{x}\right].$$

Now, using Assumption (3) in Section 2.3, for any $\epsilon_2 > 0$, we can find a large enough M such that

$$E\left[\left|\frac{A_i(k,M)}{M} - \lambda_i\right| \quad | \quad \vec{x}[k] = \vec{x}\right] < \epsilon_2 \quad \forall i.$$

This enables us to obtain

$$\sum_{i=1}^{N} E\left[f_i(z_i(k,M))(A_i(k,M) - M\lambda_i)\mathcal{I}_\mathcal{A} \quad | \quad \vec{x}[k] = \vec{x}\right] \leq \left((1+\zeta_1)M\sum_{i=1}^{N} f_i(x_i[k]) + H_1\right)\epsilon_2, \quad (43)$$

where $H_1 := \sum_{i \in E^c} f_i(X_i) < \infty$.

If instead the event $\mathcal{A}^c$ occurs, then Assumption (4) in Section 2.3 implies that for any $\epsilon_3 > 0$, we can find A large enough so that

$$\sum_{i=1}^{N} E\left[f_i(z_i(k,M))(A_i(k,M) - M\lambda_i)\mathcal{I}_{\mathcal{A}^c} \quad | \quad \vec{x}[k] = \vec{x}\right] < \epsilon_3. \quad (44)$$

Secondly, we concentrate on (41). First we write

$$\sum_{i=1}^{N} E\left[f_i(z_i(k,M))\Delta C_i \mid \vec{x}[k] = \vec{x}\right] = \sum_{i=1}^{N} E\left[\mathcal{I}_\mathcal{A} f_i(z_i(k,M))(C_i^d(k,M) - C_i(k,M)) \mid \vec{x}[k] = \vec{x}\right] \quad (45)$$

$$+ \sum_{i=1}^{N} E\left[\mathcal{I}_{\mathcal{A}^c} f_i(z_i(k,M))(C_i^d(k,M) - C_i(k,M)) \mid \vec{x}[k] = \vec{x}\right] \quad (46)$$

24

Expand (45) as follows.

$$\sum_{i=1}^{N} E[I_{\mathcal{A}}f_i(z_i(k,M)) \ (\ \sum_{n=k}^{k+M-1} \{ \ \sum_{j=1}^{J} \pi_j \mu_i(j, \vec{x}[n]) - \sum_{j=1}^{J} \pi_j \mu_i(j, \vec{x}[k])$$

$$\sum_{j=1}^{J} \pi_j \mu_i(j, \vec{x}[k]) - \mu_i(s[n], \vec{x}[k])$$

$$\mu_i(s[n], \vec{x}[k]) - \mu_i(s[n], \vec{x}[n]) \ \ \} \ ) \mid \vec{x}[k] = \vec{x}],$$

which can be further bounded as

$$\leq \ \sum_{i=1}^{N} E[I_{\mathcal{A}}f_i(z_i(k,M)) \sum_{n=k}^{k+M-1} \sum_{j=1}^{J} \pi_j (\mu_i(j, \vec{x}[n]) - \mu_i(j, \vec{x}[k])) \mid \vec{x}[k] = \vec{x}] \tag{47}$$

$$+ \sum_{i=1}^{N} E[I_{\mathcal{A}}f_i(z_i(k,M))M \left| \sum_{j=1}^{J} \mu_i(j, \vec{x}[k]) \left(\pi_j - \sum_{n=k}^{k+M-1} \frac{1}{M}\mathcal{I}_{s[n]=j} \right) \right| \mid \vec{x}[k] = \vec{x}] \tag{48}$$

$$+ \sum_{i=1}^{N} E[I_{\mathcal{A}}f_i(z_i(k,M)) \sum_{n=k}^{k+M-1} (\mu_i(s[n], \vec{x}[k]) - \mu_i(s[n], \vec{x}[n])) \mid \vec{x}[k] = \vec{x}], \tag{49}$$

Note that, under the event $\mathcal{A}$, $z_i(k,M) \in [x_i[k] - M\hat{\eta}, x_i[k] + \alpha\lambda_i M + (1-\alpha)AM]$, for some $\alpha \in [0,1]$. Then for any given M, A and $\zeta > 0$, we can find an $X_i < \infty$ such that for any $x_i > X_i$, we have $(1-\zeta)f_i(x_i) \leq f_i(y_i) \leq (1+\zeta)f_i(x_i)$.

Define the set of indices $E := \{i : x_i[k] > X_i\}$. Then using Assumption (2) of Section 2.1, along with a repetition of the argument we had for (40), given any $\epsilon_4 > 0$, we can come up with a closed, bounded region, outside of which we have

$$(48) \leq \left((1+\zeta)M \sum_{i=1}^{N} f_i(x_i[k]) + H_2 \right) \epsilon_4,$$

with $H_2 < \infty$. Since (47) and (49) follow the same reasoning, only the latter will be examined in detail. Using our earlier analysis, given any $\epsilon_5 > 0$, it is easy to upper bound the following term:

$$\sum_{i=1}^{N} E[I_{\mathcal{A}}f_i(z_i(k,M)) \sum_{n=k}^{k+M-1} \mu_i(s[n], \vec{x}[k]) \mid \vec{x}[k] = \vec{x}] \ \leq \ (1 + \frac{\epsilon_5}{2M\hat{\eta}}) \sum_{i=1}^{N} \sum_{n=k}^{k+M-1} \{f_i(x_i[k])$$

$$E[I_{\mathcal{A}}\mu_i(s[n], \vec{x}[k]) \mid \vec{x}[k] = \vec{x}]\} + H_3,$$

with $H_3 < \infty$. As for the second term, note that for all $i \in E$, we have $f_i(z_i(k,M)) \geq (1 - \frac{\epsilon_5}{2M\hat{\eta}})f_i(x_i[k])$, and for $i \in E^c$, we have $f_i(z_i(k,M)) \geq (1 - \frac{\epsilon_5}{2M\hat{\eta}})f_i(x_i[k]) - H_4$, with an appropriate choice of $H_4 < \infty$. Hence we obtain the following lower bound:

$$\sum_{i=1}^{N} E[I_{\mathcal{A}}f_i(z_i(k,M)) \sum_{n=k}^{k+M-1} \mu_i(s[n], \vec{x}[n]) \mid \vec{x}[k] = \vec{x}] \ \geq \ (1 - \frac{\epsilon_5}{2}) \sum_{i=1}^{N} \sum_{n=k}^{k+M-1} \{f_i(x_i[k])$$

$$E[I_{\mathcal{A}}\mu_i(s[n], \vec{x}[n]) \mid \vec{x}[k] = \vec{x}]\} + H_5,$$

with $H_5 < \infty$. Putting the last two bounds together, we get the following upper bound on (49):

$$(49) \quad \leq \quad \sum_{n=k}^{k+M-1} E[I_{\mathcal{A}} \sum_{i=1}^{N} (f_i(x_i[k])\mu_i(s[n], \vec{x}[k]) - f_i(x_i[k])\mu_i(s[n], \vec{x}[n])) \mid \vec{x}[k] = \vec{x}] \tag{50}$$

$$+ \frac{\epsilon_5}{2M\hat{\eta}} \sum_{i=1}^{N} f_i(x_i[k]) E[I_{\mathcal{A}} \sum_{n=k}^{k+M-1} (\mu_i(s[n], \vec{x}[k]) - \mu_i(s[n], \vec{x}[n])) \mid \vec{x}[k] = \vec{x}] \tag{51}$$

$$+ H_6$$

$$\leq \quad \epsilon_5 \sum_{i=1}^{N} f_i(x_i[k]) + H_6 \tag{52}$$

with H_6 bounded, where the last inequality follows by observing that the expectation in (51) is upper-bounded by $M\hat{\eta}$, and, given M and A, we can find the bounded region around the origin outside of which the expectation in (50) can be upper-bounded by $\frac{\epsilon_5}{2M} \sum_{i=1}^{N} f_i(x_i[k])$. Similarly, an upper bound on (47) can be obtained by choosing the bounded region large enough. Such an argument would yield the following expression

$$(47) \leq \epsilon_6 \sum_{i=1}^{N} f_i(x_i[k]) + H_7,$$

with H_7 bounded. Putting all these bounds together, we can upper-bound (45) as

$$(45) \leq \epsilon_7 M ||\vec{f}(\vec{x}[k])||_2 + H_8, \tag{53}$$

where H_8 is finite. As for (46), we can choose the parameter A large enough so that, due to Assumption (4) of Section 2.3, it converges to zero.

Thirdly, consider (42) under the event $\mathcal{A}$. Observe that for any queue, say i, and any fixed M, the sum of unused service may be nonzero over a duration of M slots only if $x_i[n] < \hat{\eta}$ for some $n \in \{k, \cdots, k + M - 1\}$. Therefore, if $U_i(k, M) > 0$, it is easy to see that $f_i(z_i(k, M)) \leq f_i(\hat{\eta} + MA)$. Similarly, $U_i^d(k, M) > 0$ only if $f_i(z_i(k, M)) \leq f_i(\hat{\eta} + M\lambda_i)$. Moreover, the cumulative unused service over M slots is upper-bounded by $M\hat{\eta}$, since within each slot, the maximum amount of possible unused service is $\hat{\eta}$. Hence, we can easily upper-bound (42) under the event $\mathcal{A}$ as

$$\sum_{i=1}^{N} E\left[\mathcal{I}_{\mathcal{A}} f_i(z_i(k, M))(U_i(k, M) - U_i^d(k, M)) \mid \vec{x}[k] = \vec{x}\right] \leq M\hat{\eta} \max_i f_i(\hat{\eta} + MA). \tag{54}$$

If the event $\mathcal{A}^c$ occurs then we can write, for any $\epsilon_8 > 0$,

$$\sum_{i=1}^{N} E\left[\mathcal{I}_{\mathcal{A}^c} f_i(z_i(k, M))(U_i(k, M) - U_i^d(k, M)) \mid \vec{x}[k] = \vec{x}\right] \leq M\hat{\eta} E\left[\mathcal{I}_{\mathcal{A}^c} f_i(z_i(k, M)) \mid \vec{x}[k] = \vec{x}\right]$$

$$< M\epsilon_8, \tag{55}$$

with $A < \infty$ chosen to be large enough, due to Assumption (4) of Section 2.3. Putting (54) and (55) together, (42) can be upper-bounded as

$$(42) \leq M\epsilon_8 + H_9 \tag{56}$$

for any $\epsilon_8 > 0$ and with $H_9 < \infty$ chosen appropriately.

Now, combining (43), (44), (53) and (56), for any $\epsilon > 0$, we can come up with a closed, bounded region around the origin outside of which we have

$$E\left[V(\vec{x}[k+M]) - V(\vec{x}^d[k+M]) \mid \vec{x}[k] = \vec{x}\right] \leq M\epsilon||\vec{f}(\vec{x}[k])||_2 + H,$$

for some $H < \infty$ chosen appropriately. So if we choose $\epsilon_M - \epsilon =: \xi > 0$, then we can come up with a closed region, denoted by $\mathcal{G}_{(M)}[k]$, outside of which $\vec{x}[k]$ is sufficiently large, and

$$\Delta W^{(M)}(\vec{x}[k]) \leq -M\xi||\vec{f}(\vec{x}[k])||_2\mathcal{I}_{\vec{x}[k]\in\mathcal{G}_{(M)}^c[k]} + K_{(M)}\mathcal{I}_{\vec{x}[k]\in\mathcal{G}_{(M)}[k]}, \tag{57}$$

with an appropriate choice of $K_{(M)} < \infty$. We can also write the previous expression as

$$E\left[W(\vec{x}[k+M]) - W(\vec{x}[k]) \mid \vec{x}[k] = \vec{x}\right] \leq -M\xi||\vec{f}(\vec{x}[k])||_2\mathcal{I}_{\vec{x}[k]\in\mathcal{G}_{(M)}^c[k]} + K_{(M)}.$$

Taking expectations on both sides, we get

$$E[W(\vec{x}[k+M])] - E[W(\vec{x}[k])] \leq -M\xi E\left[||\vec{f}(\vec{x}[k])||_2\mathcal{I}_{\vec{x}[k]\in\mathcal{G}_{(M)}^c[k]}\right] + K_{(M)}.$$

Then, for any positive integer p, we have

$$E[W(\vec{x}[(p-1)M])] - E[W(\vec{x}[0])] \leq -M\xi\sum_{k=0}^{p-1}E\left[||\vec{f}(\vec{x}[kM])||_2\mathcal{I}_{\vec{x}[kM]\in\mathcal{G}_{(M)}^c[kM]}\right] + pK_{(M)}.$$

Since $E[W(\vec{x}[(p-1)M])] > 0$, we have

$$-E[W(\vec{x}[0])] \leq -M\xi\sum_{k=0}^{p-1}E\left[||\vec{f}(\vec{x}[kM])||_2\mathcal{I}_{\vec{x}[kM]\in\mathcal{G}_{(M)}^c[kM]}\right] + pK_{(M)}.$$

Re-arranging the terms, we get

$$\sum_{k=0}^{p-1}E\left[||\vec{f}(\vec{x}[kM])||_2\mathcal{I}_{\vec{x}[kM]\in\mathcal{G}_{(M)}^c[kM]}\right] \leq \frac{E[W(\vec{x}[0])]}{M\xi} + \frac{pK_{(M)}}{M\xi}.$$

Now observe that $E\left[||\vec{f}(\vec{x}[kM])||_2\mathcal{I}_{\vec{x}[0]\in\mathcal{G}_{(M)}[kM]}\right] < \infty$ since $\mathcal{G}_{(M)}$ represents a closed bounded region around the origin and $||\vec{f}||_2$ is bounded is inside such a region. This allows us to write

$$\sum_{k=0}^{p-1}E\left[||\vec{f}(\vec{x}[kM])||_2\right] \leq \frac{E[W(\vec{x}[0])]}{M\xi} + \frac{pT_{(M)}}{M\xi},$$

for some finite $T_{(M)}$. Finally, dividing both sides by p and letting $p \to \infty$, we obtain

$$\limsup_{p\to\infty}\frac{1}{p}\sum_{k=0}^{p-1}E\left[||\vec{f}(\vec{x}[k])||_2\right] \leq \frac{T_{(M)}}{M\xi} =: \bar{l} < \infty, \tag{58}$$

as claimed in the theorem statement. $\diamond$

27

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] M. Armony and N. Bambos. Queueing dynamics and maximal throughput scheduling in switched processing systems. Technical Report Netlab-2001-09/01, Stanford University.

[3] D. Bertsekas. *Nonlinear Programming*. Athena Scientific, Belmont, MA, 1995.

[4] P. Billingsley. *Convergence of Probability Measures*. Wiley, 1968.

[5] P. Giaccone, B. Prabhakar, and D. Shah. Towards simple, high-performance schedulers for high-aggregate bandwidth switches. In *Proceedings of IEEE Infocom*, 2002.

[6] P.R. Kumar and S.P. Meyn. Stability of queueing networks and scheduling policies. *IEEE Transactions on Automatic Control*, 40:251–260, February 1995.

[7] R. Leelahakriengkrai and R. Agrawal. Scheduling in multimedia wireless networks. In *Proceedings of ITC*, Brazil, 2001.

[8] L. Li and A. J. Goldsmith. Optimal resource allocation for fading broadcast channels- Part I: Ergodic capacity. *IEEE Transactions on Information Theory*, March 2001.

[9] X. Liu, E. Chong, and N. Shroff. Opportunistic transmission scheduling with resource-sharing constraints in wireless networks. *IEEE Journal on Selected Areas in Communications*, 19(10):2053–2064, October 2001.

[10] D. Mitra and J.A. Morrison. Multiple time scale regulation and worst case processes for ATM network control. In *Proceedings of the IEEE Conference on Decision and Control*, pages 353–358, New Orleans, LA, 1995.

[11] K. Ross and N. Bambos. Projective processing schedules in queueing structures: Applications to packet scheduling in communication network switches. Technical Report Netlab-2002-05/01, Stanford University.

[12] S. Shakkottai, R. Srikant, and A. Stolyar. Pathwise optimality of the exponential scheduling rule for wireless channels. In *Proceedings of ISIT*, Lausanne, Switzerland, July 2002. Longer version available at http://www.comm.csl.uiuc.edu/~srikant.

[13] S. Shakkottai and A. Stolyar. Scheduling for multiple flows sharing a time-varying channel: The exponential rule. *Translations of the AMS,* A volume in memory of F. Karpelevich, 2002. To appear.

[14] L. Tassiulas. Scheduling and performance limits of networks with constantly varying topology. *IEEE Transactions on Information Theory*, pages 1067–1073, May 1997.

[15] L. Tassiulas. Linear complexity algorithms for maximum throughput in radio networks and input queued switches. In *Proceedings of IEEE Infocom*, 1998.

[16] 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*, pages 1936–1948, December 1992.

[17] L. Tassiulas and A. Ephremides. Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Transactions on Information Theory*, 39:466–478, March 1993.

[18] D. Tse. Forward link multiuser diversity through rate adaptation and scheduling. In preparation.

[19] P. Viswanath, D. Tse, and R. Laroia. Opportunistic beamforming using dumb antennas. *IEEE Transactions on Information Theory*, 48(6):1277–1294, June 2002.