On Wireless Scheduling With Partial Channel-State Information

Aditya Gopalan, Constantine Caramanis, *Member, IEEE*, and Sanjay Shakkottai, *Senior Member, IEEE*

***Abstract*—A time-slotted queueing system for a wireless downlink with multiple flows and a single server is considered, with exogenous arrivals and time-varying channels. It is assumed that only one user can be serviced in a single time slot. Unlike much recent work on this problem, attention is drawn to the case where the server can obtain only partial information about the instantaneous state of the channel. In each time slot, the server is allowed to specify a single subset of flows from a collection of observable subsets, observe the current service rates for that subset, and subsequently pick a user to serve. The stability region for such a system is provided. An online scheduling algorithm is presented that uses information about marginal distributions to pick the subset and the Max-Weight rule to pick a flow within the subset, and which is provably throughput-optimal. In the case where the observable subsets are all disjoint, or where the subsets and channel statistics are symmetric, it is shown that a simple scheduling algorithm—Max-Sum-Queue—that essentially picks subsets having the largest squared-sum of queues, followed by picking a user using Max-Weight within the subset, is throughput-optimal.**

***Index Terms*—Partial information, throughput optimality, wireless scheduling.**

I. Introduction

THERE has been much recent interest in scheduling over wireless cellular networks where channel state information is available at the base-station [2]–[4]. A canonical system consists of a base-station (the server) and a collection of mobile users (the queues). Time is slotted (typically of the order of a millisecond), like in the high-speed Worldwide Interoperability for Microwave Access (WiMAX) [5], Ultra Mobile Broadband (UMB), Global System for Mobile Communications (GSM)-based High-Speed Downlink Packet Access (HSDPA), and Evolution-Data Optimized (EV-DO) communications technologies. In each time-slot, the channel state, i.e., the channel quality such as the signal to interference-plus-noise ratio (SINR) or data rate that can be sustained over the time-slot to the mobile, is potentially available via a feedback channel from the mobile terminals to the base-station. Based on the load (packets queued at the base-station) as well as the channel state, the base-station schedules users for channel access at each time-slot.

As the capacity of the wireless system increases, growing numbers of users will be connected to the base-station at any given time. As a result, schemes wherein all users transmit channel state feedback to the base-station may become untenable, due to feedback bandwidth constraints. One approach to mitigate this problem is for the base-station to request channel state information from a (small) subcollection of users and make scheduling decisions based on this partial channel state information. Our goal is to understand how the base-station can intelligently decide which subsets of the users to sample to obtain partial channel state information, and how to schedule users based on this information. Furthermore, we are interested in understanding how this partial information degrades the stability region, i.e., what is the effect of partial information on the capacity of a wireless network.

We characterize the exact stability region given any set of observable subsets, and we provide an algorithm that is throughput-optimal. Unlike the full-information case studied in e.g., [2] that requires no distributional information, our algorithm requires knowledge of the marginals of the channel state distribution for the observable subsets. For the special case of symmetric flows, we provide a simpler throughput-optimal algorithm that requires no such information. We further show that the reduction in the stability region is due precisely to the inability to observe the full instantaneous state, as opposed to failure to obtain the full joint distribution of the channel state. Indeed we show that knowledge of the full distribution may not yield a larger stability region, unless the observable subsets themselves are enlarged.

A. Main Contributions

We consider a base-station system serving N users and channels, with each user generating data, and with channels which have an arbitrary joint distribution over a finite state-space (the channel is assumed to be independent across time but not across users), and the server *not having* knowledge of the channel joint distribution.

In each time-slot, the base-station is allowed to acquire channel state[1] from one among a predefined collection of subsets of channels. For example, in a ten-user system, the constraint could be that we can acquire channel state from at most three users per time-slot (we note, though, that our main

[1]At each time-slot, the complete channel state is a N dimensional vector, with the ith component of the vector corresponding to the data rate that can be sustained to the ith mobile user over the time-slot if this user is chosen by the scheduler. Correspondingly, the partial channel state corresponds to a subvector of this N dimensional vector.

Manuscript received October 30, 2009; revised December 24, 2010; accepted May 23, 2011. Date of current version nulldate. This work was partially supported by the NSF Grants CNS-0325788, CNS-0347400, CNS-0519535, CNS-0721532, EFRI-0735905 and by the DARPA ITMANET program. A preliminary version of this work was presented at the 45th Allerton Conference on Communication, Control and Computing, Urbana, IL, September 2007.

The authors are with the Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, TX **XXXXX** USA. They are also with the Wireless Communications and Networking Group (WNCG) (e-mail: gadit@mail.utexas.edu; caramanis@mail.utexas.edu, shakkott@mail.utexas.edu).

Communicated by A. Nosratinia, Communication Networks.
Digital Object Identifier 10.1109/TIT.2011.2169543

results are completely general with respect to the structure of the observable subset collection). We henceforth refer to this as a system with partial channel-state information.

The scheduling task at each time-slot is to first determine the subset of channels for which channel state will be acquired and then determine a single user to schedule from within this subset. In this paper, we characterize the stability region for this multiuser system, and develop algorithms that achieve the full stability region. Specifically, the main contributions in this paper are as follows.

1) We derive the stability region for a system with N users and an arbitrary collection of observable subsets (i.e., a collection of subsets of users for which the channel state can be simultaneously acquired), and for any joint channel distribution across users where channel realizations are independent and identically distributed over time. The stability region corresponds to the set of arrival rates that can be sustained such that the queues at the base-station are stable (positive recurrent).

   We demonstrate that the stability region with partial channel state information can be described by the convex hull of "local" stability regions for the observable user subsets. These local regions are completely characterized by a simple class of scheduling policies commonly called *Static Split Service rules* (e.g., [2]).

   A numerical example is presented that illustrates the degradation in the stability region as the amount of channel state information decreases (i.e., when there are fewer simultaneously observable channels).
2) The characterization of the stability region shows that it is completely determined by just the marginal statistics of the aggregate channel over observable subsets. It also leads to the important counterintuitive result that additional information about the joint distribution of the channel state, even if provided to the scheduler at all times, cannot help increase throughput. In other words, the degradation of the stability region is precisely due to the lack of capability to observe channel state, as opposed to lack of knowledge about how the channel state is distributed.
3) Next, we develop a queue-length based "online" scheduling policy that uses queue-length information along with knowledge of subset-marginal distributions, and which is throughput-optimal, i.e., the policy attains all rate points within the stability region. The policy consists of two stages: In each time slot, *(a)* the base-station first determines the subset of channel measurements to observe. This is done using the *expected* rates over the observable subsets weighted by the *actual* queue lengths at the base-station; and *(b)* within the chosen subset, the policy uses the *Max-Weight* rule [6], [2] which uses the product of the *actual* channel rate (received from the mobile in the chosen subset) and the *actual* queue-length to make the scheduling decision.
4) We develop a simpler online policy (the *Max-Sum-Queue rule*) that requires no distributional information. In the first stage, this policy determines the subset of users chosen by only the queue lengths and does not use the expected channel rates. The Max-Sum-Queue policy chooses that subset over which the sum of the squares of the queue-lengths is largest. The second stage is the same as before, namely, the Max-Weight policy restricted to the chosen subset. We show that if the observable subsets are disjoint or the observable subsets and channels are symmetric, this policy is throughput-optimal. Finally, we provide an example to show that in general this policy is not throughput optimal if the symmetric-channel-and-observable-subsets/disjoint-observable-subsets condition is not met.

B. Related Work

There has been much work in developing scheduling algorithms for downlink wireless systems for various performance metrics that include stability, utility maximization and probabilistic delay guarantees [6]–[13]. However, the above studies primarily focus on the case where complete channel state information is available at the base-station, and thus consider problems orthogonal to the main issues in this paper.

In the context of partial channel information, related work includes that of [14] where the authors study the problem of a server (terminal) accessing N time varying channels which are independent across users and time (e.g., a multichannel MAC). The server has a cost for (sequentially) probing channels, with a channel dependent probing cost, and gains a reward which depends on the user and the probed state, if a packet is transmitted successfully. The authors formulate the problem of minimizing the expected cost (probing cost minus reward for transmissions) where the cost functions and the channel probabilities are known to the server. They further develop constant factor (within the optimal cost) approximation algorithms that operate in polynomial time for both the saturated data case, as well as when the user (terminal) generates packets according to a Markov chain. The authors in [15] and [16] have earlier considered the special cases with equal probing costs and identically distributed channels. Recent results in this context include [17] where the authors develop structural properties of the optimal probing strategy using a dynamic programming approach, and [18] which treats the problem of optimal channel probing and scheduling for stability, for Markovian channels independent across users, in a Markov-decision-theoretic framework and derives polynomial-time computable optimal policies or approximations in certain cases.

For systems with channels that are independent across users and with infinitely backlogged data at the base-station, there has been work considering limited feedback from the mobile users to the base-station. In these studies, the mobiles use thresholds to determine if their channel quality is "good enough," and if so, send their channel state information to the base-station [19]–[23].

A closely related work that appeared subsequent to the conference version of this work is [24], which proposes empirical sampling and learning of incomplete channel state statistics in order to maximize a convex utility of rates while maintaining stability. The work in this paper is, to the best of our knowledge, the first to consider characterizing stability of these wireless networks under availability of limited channel state information, while obtaining corresponding throughput optimal efficient algorithms. In particular, the work here differs significantly from the previous work described above in the sense of investigating stability in the presence of partial channel state information. Also, we emphasize the need for efficient scheduling rules based on feedback received via queue length information.

II. System Model and Definitions

Throughout the paper, we assume a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$ which supports all random variables and random processes.

Consider a time-slotted model of $N < \infty$ users serviced by a single server across N unidirectional communication channels $\{c_1, \ldots, c_N\} \stackrel{\triangle}{=} C$. An integer number of data packets arrive at the input of every channel at the beginning of a time slot, to be serviced by the server. Packets get queued at the inputs of channels if they are not immediately transmitted. We assume that at most one of the channels can be activated for transmission in a single time slot.

Further, in any given time slot $t \in \{0, 1, 2, \ldots\}$, the set of channels C assumes a *state* $L(t)$ from a finite set of aggregate channel states $\mathcal{L} = \{l_1, \ldots, l_{|\mathcal{L}|}\}$, with the channel state remaining constant within each time slot. In each channel state $l \in \mathcal{L}$, every channel $c_i \in C$ assumes a data *service rate* of μ_i^l, i.e., a maximum of μ_i^l packets can be served from queue i (corresponding to channel c_i) when the aggregate channel is in state l. Henceforth, we identify each state $l \in \mathcal{L}$ with its N-dimensional vector of service rates $(\mu_i^l)_{i=1}^N$, and treat $L(t)$ as a random vector which can take any such value l.

The random channel state process $\mathbf{L} \stackrel{\triangle}{=} (L(t))_{t=0}^{\infty}$ is assumed to be an independent and identically distributed (i.i.d.) discrete-time random process taking values from the finite state space $\mathcal{L}$. For $l \in \mathcal{L}$, let $\pi^l \stackrel{\triangle}{=} \mathbb{P}(L(0) = l)$. Observe that the channel state process is i.i.d. across time only, and can have any joint distribution across users (i.e., across channels).

Let us denote by $A_i(t)$ the number of packets that arrive at channel c_i at time slot t, and let $A(t) \stackrel{\triangle}{=} (A_i(t))_{i=1}^N \in \mathbb{R}^N$. The packet arrival process $\mathbf{A}_i \stackrel{\triangle}{=} (A_i(t))_{t=0}^{\infty}$ at the input of each channel c_i, $i = 1, \ldots, N$, is assumed to be a nonnegative finite-state irreducible discrete-time Markov chain in its stationary distribution. We call $\mathbb{E}[A_i(0)] = \lambda_i > 0$ the *arrival rate* at channel i, $i = 1, \ldots, N$. Each arrival process is taken to be independent of all other processes.

Our channel observations are limited to a given collection of subsets of C (whose union is assumed to be C) called the collection of *observable subsets*. Let us denote this collection of observable subsets by $O = \{O_1, O_2, \ldots, O_{|O|}\}$. In the example of Section III-B, C is a set of three channels and the set O contains all subsets of size two. In a given time slot, an observable subset $\alpha = \{c_{n_1}, \ldots, c_{n_m}\} \subset C$ is said to be in a *substate* $\mu^k = (\mu_{n_1}^k, \ldots, \mu_{n_m}^k) \in \mathbb{R}^m$ if $L(t)_{n_j} = \mu_{n_j}^k$, $j = 1, \ldots, m$. Denote by $L^\alpha(t)$ the m-length substate random vector that is the projection of $L(t)$ onto coordinates $n_1, \ldots, n_m$.

Similar to the treatment in [2], we define the *state* of the system to be the random process $\mathbf{S} = (S(t))_{t=0}^{\infty}$ where $S(t) \stackrel{\triangle}{=} (Q_1(t), \ldots, Q_N(t))$, augmented by the state of the arrivals. Here, $Q_j(t)$ denotes the length of the packet queue for channel $c_j \in C$ at time slot t.

We model the system state as evolving via the action of a *scheduling policy*. A *scheduling policy* $\mathcal{P}$ is a pair of maps $(\mathcal{G}, \mathcal{H})$, where $\mathcal{G}$ is a map from the state of the system $S(t)$ to a fixed probability distribution on the set of observable subsets O, and $\mathcal{H}$ is a map which takes $S(t)$ restricted to a particular observable subset, along with its substate, into a fixed probability distribution on the channels which comprise the subset. Such a scheduling policy $\mathcal{P}$ is applied to select a transmitting channel using two steps. At every time slot t, in the first step, we pick an observable set α randomly according to the distribution $\mathcal{G}(S(t))$ after which we are able to sample the substate of the chosen observable set. Then, according the distribution $\mathcal{H}$ on the observable set α and its substate $L^\alpha(t)$, we pick a channel $c_i \in \alpha$ for transmission from α. Following this choice of channel, the queue length Q_i evolves in the standard sense as $Q_i(t+1) = \max\{0, Q_i(t) + A_i(t) - L(t)_i\}$, whereas all the other queue lengths Q_j, $j \neq i$, evolve as $Q_j(t+1) = Q_j(t) + A_j(t)$. This scheduling model differs from the one in [2] in that this is a *two-stage* procedure where the subset to be sampled in the first step is a function of *only* queue information and not the instantaneous channel state.

Under a scheduling policy $\mathcal{P}$, the state $\mathbf{S}$ is a discrete-time countable-state Markov chain, which we further assume to be irreducible and aperiodic. This can, for example, be satisfied if the arrival and channel process marginal distributions have positive probability on a finite subset of the nonnegative integer lattice $\{0, 1, \ldots, L\}$. Weaker conditions suffice by using different notions of stability, e.g., that there is a nonempty positive recurrent set of states, and an associated finite subset which is entered in finite time with probability one [2]. However, to avoid purely technical complications, we assume that the support of the arrival and channel state processes are such that the scheduling policies we consider render the Markov Chain to be aperiodic and irreducible (see also Section III in [25]).

A rate vector $\Lambda = (\lambda_1, \ldots, \lambda_N) \in \mathbb{R}^N$ is said to be *supported* by a scheduling policy $\mathcal{P}$ if the Markov chain $\mathbf{S}$ is ergodic or positive recurrent under scheduling using $\mathcal{P}$, when the arrival rates at the inputs of channels $c_1, \ldots, c_N$ are $\lambda_1, \ldots, \lambda_N$, respectively. In other words a policy supports an arrival rate vector if the input packet queues at all channels in the system remain stable under the policy. Associated with each policy $\mathcal{P}$ is its *rate region* $\mathcal{R}_{\mathcal{P}} \stackrel{\triangle}{=} \{\Lambda \in \mathbb{R}^N : \Lambda \text{ is supported by } \mathcal{P}\}$. The *achievable rate region* or *throughput region* or *stability region* $\mathcal{R}$ is then defined to be the union of the rate regions for all possible scheduling policies $\mathcal{P}$. A rate vector Λ is said to be *achievable* if it is supported by some scheduling policy. Likewise, a set or region $\mathcal{A} \subset \mathbb{R}^N$ is said to be achievable if all its elements are achievable. A scheduling policy is said to be *throughput-optimal* if it supports all vectors in the achievable rate region.

We wish to characterize the achievable rate region for the model we have described. Henceforth, we shall naturally assume that all the subsets in O are maximal with respect to set inclusion.

III. The Achievable Rate Region

In the first part of this section, we show two main results. First, we characterize the achievable rate region for any collection of observable subsets O. Next, we show that this region is attained using a special class of scheduling policies called *Static Split Service (SSS) rules* [2]. The reason they are called so is that they are independent of the queue lengths at every time slot and rely only on the channel state to make randomized scheduling decisions. We present an example in which we explicitly describe the achievable rate region for a system of three channels, under different partial information structures. The final part of this section characterizes 'good' or optimal SSS scheduling rules.

A. *Description of the Throughput Region*

Consider an observable subset $\alpha \in O$, $\alpha = \{c_{k_1}, c_{k_2}, \ldots, c_{k_m}\}$ where $k_1, \ldots, k_m \in \{1, \ldots, N\}$. Let $\mathcal{Q}(\alpha)$ denote the m-dimensional subspace of $\mathbb{R}^N$ where coordinates with indices other than $k_1, \ldots, k_m$ are zero. If only users from α are served, then any stabilizable rate must lie in $\mathcal{Q}(\alpha)$. Denote this stabilizable rate region by $\mathcal{R}(\alpha)$. Applying [2, Theorem 1] to the subset α, we can describe the achievable rate region when only α is allowed to be picked in the first scheduling step.

Lemma 1: There exists a scheduling policy $\mathcal{P}$ stabilizing a rate vector $\Lambda \in \mathcal{R}(\alpha)$ if and only if there exists a stochastic matrix ϕ^α such that

$$\lambda_i < v_i^\alpha(\phi^\alpha) \triangleq \sum_{l \in \mathcal{L}_\alpha} \pi^{l,\alpha} \phi_{li}^\alpha \mu_i^{l,\alpha}, \quad \forall c_i \in \alpha.$$

Here, $\mathcal{L}_\alpha$ is the set of substates of α, $\pi^{l,\alpha}$ is the marginal probability of the substate l and $\mu_i^{l,\alpha}$ is the service rate for channel c_i in substate l.

The matrix ϕ^α defines an SSS rule for the subset α. The rows of ϕ^α correspond to every substate of α and the columns of ϕ^α correspond to every channel in α. When α is in the substate $m = (\mu_{c_{k_1}}, \ldots, \mu_{c_{k_l}})$, the SSS rule picks channel i for transmission with probability ϕ_{mi}^α.

Lemma 1 states that the stability region for scheduling using α is the convex polytope $\mathcal{R}(\alpha)$. The following theorem establishes that the stability region for the whole system is the convex hull of such polytopes.

Theorem 1: The achievable region $\mathcal{C}$ for the whole system is the convex hull of the stabilizable regions in each subspace $\mathcal{Q}(\alpha)$, for $\alpha \in O$

$$\mathcal{C} \triangleq \mathrm{conv}\left(\{\mathcal{R}(\alpha) : \alpha \in O\}\right).$$

The theorem says that any rate vector in the stability region can be supported by timesharing across observable subsets and across users within subsets. The proof of the theorem follows from the following two lemmas which establish matching inner and outer bounds on the region $\mathcal{C}$.

Lemma 2: $\mathcal{C}$ is achievable.

Proof: Let K be the total number of observable subsets. By definition, if $\Lambda \in \mathcal{C} = \mathrm{conv}\left(\{\mathcal{R}(\alpha_i) : i = 1, \ldots, K\}\right)$, then there exist nonnegative reals p_{α_i} with $\sum_{i=1}^K p_{\alpha_i} = 1$ and $\Lambda_{\alpha_i} \in \mathcal{R}(\alpha_i)$ such that $\Lambda = \sum_{i=1}^K p_{\alpha_i} \Lambda_{\alpha_i}$. This shows that the static service split (SSS) scheduling rule which chooses each subset α_i with probability p_{α_i} and each user in α_i with a suitable probability to ensure a mean service rate of Λ_{α_i} stabilizes the system. ■

Next, we show that no more rate vectors are achievable.

Lemma 3: If $\Lambda \in \mathbb{R}^N$ is achievable, then $\Lambda \in \mathcal{C}$. In particular, Λ can be achieved by a global SSS scheduling rule parametrized by a stochastic matrix ϕ of the form

$$\phi = \sum_{\alpha \in O} p_\alpha \phi^\alpha \tag{1}$$

where ϕ^α are stochastic matrices as described above, and p_α is a probability distribution on the maximal observable subsets, O. Similar to the notion of an SSS rule for a maximal observable subset, the matrix ϕ above defines a *global SSS rule* for our system. A scheduling policy implementing this global SSS rule selects a subset α in the first step with probability p_α and subsequently uses the subset SSS rule ϕ^α to pick a queue in α. The (long-term) service rate such a rule provides to queue i is

$$v_i \triangleq \sum_{\alpha \in O} p_\alpha v_i^\alpha(\phi^\alpha) = \sum_{\alpha \in O} p_\alpha \sum_{l \in \mathcal{L}_\alpha} \pi^{l,\alpha} \mu_i^{l,\alpha} \phi_{li}^\alpha \tag{2}$$

and the throughput region $\mathcal{C}$ is essentially the set of all $(v_1, \ldots, v_N)$ as p_α and ϕ_{li}^α range from 0 to 1 with $\sum_\alpha p_\alpha = 1$ and $\sum_i \phi_{li}^\alpha$, for each $\alpha \in O$, $l \in \mathcal{L}_\alpha$ and $i \in \alpha$.

See Appendix A for the proof of Lemma 3.

Implications of the result: According to Theorem 1:

- The rate region $\mathcal{C}$ is a function of the service rates of the channels and *marginal probabilities over the observable subsets* only, and does not explicitly depend upon the overall joint probability distribution of all the channels. In other words, *two systems of channels with different overall joint distributions but with identical marginal distributions $\pi^{l,\alpha}$ on all observable subsets α are indistinguishable to scheduling policies which use partial information, from the point of view of long-term service rates that can be achieved.*
- Suppose a scheduling policy with partial channel state information is aided by a 'genie' which furnishes the policy with the joint probability distribution of all the channel states. Theorem 1 says that this *additional joint distribution information cannot help the scheduler enlarge the throughput region*. Intuitively, this can be understood in two ways—
  1) The scheduler's action of observing the channel states of only a subset of channels (in a time slot) forces the scheduler to work (in the sense of service rates) in the subspace corresponding to that subset in N-dimensional space. Coupled with the fact that only one subset can be observed per time slot, scheduling in this case reduces to timesharing between service rates attainable within observable subsets. This holds even when the entire joint distribution of channel states is known in advance, and is the reason why a scheduler with knowledge of the joint distribution cannot improve the throughput region outside the convex hull of the throughput regions of the observable subsets.
  2) Being able to observe the entire set of channel state realizations and directly schedule a channel allows for

TABLE I
PROBABILITY ASSIGNMENTS FOR THREE-CHANNEL SYSTEM

Channel \ State	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8
c_1	a	a	a	a	b	b	b	b
c_2	a	a	b	b	a	a	b	b
c_3	a	b	a	b	a	b	a	b
State probability	π_1	π_2	π_3	π_4	π_5	π_6	π_7	π_8

more global SSS rules compared to the restricted set of SSS rules that can be achieved by picking a subset of channels to observe and scheduling a channel within the subset. This limitation on the available space of static rules in the case of reduced instantaneous channel state information leads to the diminishment of the throughput region.

B. Example: Rate Region for Three Symmetric Channels

In this part of the section we derive the throughput region for a system of three channels by applying Theorem 1. We consider three subset structures—completely observable, pairwise observable and singleton observable - and demonstrate how the throughput region shrinks with reduction in the available partial information.

Consider a system of three channels $C_3 = \{c_1, c_2, c_3\}$ in which the system can take one of eight possible states $\{l_1, \ldots, l_8\}$ (Table I), and where each of the channels c_i takes a rate of either a or $b (a < b)$ in every state. We denote the 8 values that specify the joint distribution of all three channels by $\pi_1, \pi_2, \ldots, \pi_8$ as shown in the table. Further, let us assume that $\pi_1 = \pi_2 = \cdots = \pi_8 = 1/8$ which corresponds to an i.i.d. system of channels. We compute the throughput region for the following channel state information structures.

1) Complete Channel State Information: Let $O = \{\{c_1, c_2, c_3\}\}$, i.e., all channels are simultaneously observable. For this lone observable subset $\alpha = \{c_1, c_2, c_3\}$, we have

- $\mathcal{L}_\alpha = \{l_1, \ldots, l_8\}$
- $\pi^{l_1,\alpha} = \pi_1, \ldots, \pi^{l_8,\alpha} = \pi_8$
- $\mu_1^{l_1,\alpha} = a, \mu_2^{l_1,\alpha} = a, \mu_3^{l_1,\alpha} = a,\ \mu_1^{l_2,\alpha} = a, \mu_2^{l_2,\alpha} = a, \mu_3^{l_2,\alpha} = b$, etc.

In this case, $\mathcal{C} = \mathcal{R}(\{c_1, c_2, c_3\})$ where $\mathcal{R}(\alpha)$ is the set of all rate vectors $(\lambda_1, \lambda_2, \lambda_3) \in \mathbb{R}^3$, $\lambda_1 \geq 0, \lambda_2 \geq 0, \lambda_3 \geq 0$, for which there exists a stochastic 8×3 matrix $\phi \equiv \phi^{\{c_1,c_2,c_3\}}$ such that

$$\begin{aligned}
&\phi_{11}\pi_1 a + \phi_{21}\pi_2 a + \phi_{31}\pi_3 a + \cdots + \phi_{81}\pi_8 b > \lambda_1 \\
&\phi_{12}\pi_1 a + \phi_{22}\pi_2 a + \phi_{32}\pi_3 b + \cdots + \phi_{82}\pi_8 b > \lambda_2, \quad \text{and} \\
&\phi_{13}\pi_1 a + \phi_{23}\pi_2 b + \phi_{33}\pi_3 a + \cdots + \phi_{83}\pi_8 b > \lambda_3.
\end{aligned}$$

With $\pi_i = 1/8$ for all i, we get the three-dimensional throughput region shown in Fig. 1.

2) Pairwise Channel State Information: Let $O = \{\{c_1, c_2\}, \{c_2, c_3\}, \{c_3, c_1\}\}$, i.e., at most a pair of channels is simultaneously observable. Recalling the notation used in the system model, for the observable subset $\alpha = \{c_1, c_2\}$ we have

- $\mathcal{L}_\alpha = \{(a, a), (a, b), (b, a), (b, b)\}$
- $\pi^{(a,a),\alpha} = \pi_1 + \pi_2, \pi^{(a,b),\alpha} = \pi_3 + \pi_4, \pi^{(b,a),\alpha} = \pi_5 + \pi_6, \pi^{(b,b),\alpha} = \pi_7 + \pi_8$

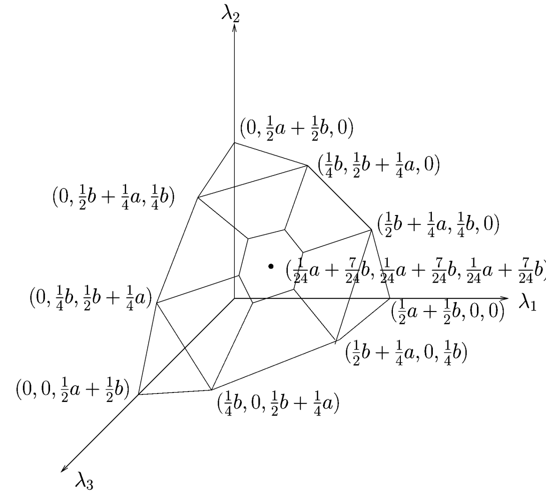

Fig. 1. Rate region for 3 channels with complete channel state information.

- $\mu_1^{(a,a),\alpha} = a, \mu_2^{(a,a),\alpha} = a, \mu_1^{(a,b),\alpha} = a, \mu_2^{(a,b),\alpha} = b$, etc.

and similarly for the other observable subsets $\beta = \{c_2, c_3\}$ and $\gamma = \{c_3, c_1\}$. In this case

$$\mathcal{C} = \text{conv}(\mathcal{R}(\{c_1, c_2\}), \mathcal{R}(\{c_2, c_3\}), \mathcal{R}(\{c_3, c_1\})).$$

The subset throughput region $\mathcal{R}(\{c_1, c_2\})$, say, is the set of all rate vectors $(\lambda_1, \lambda_2, 0) \in \mathbb{R}^3$, $\lambda_1 \geq 0, \lambda_2 \geq 0$, for which there exists a stochastic 4×2 matrix $\phi \equiv \phi^{\{c_1,c_2\}}$ such that

$$\begin{aligned}
&\phi_{11}(\pi_1 + \pi_2)a + \phi_{21}(\pi_3 + \pi_4)a + \phi_{31}(\pi_5 + \pi_6)b \\
&\quad + \phi_{41}(\pi_7 + \pi_8)b > \lambda_1, \quad \text{and} \\
&\phi_{12}(\pi_1 + \pi_2)a + \phi_{22}(\pi_3 + \pi_4)b + \phi_{32}(\pi_5 + \pi_6)a \\
&\quad + \phi_{42}(\pi_7 + \pi_8)b > \lambda_2.
\end{aligned}$$

In general, for the subset $\{c_i, c_j\} = \{c_1, c_2, c_3\} \backslash \{c_k\}$ with $i, j, k \in \{1, 2, 3\}$ and $\pi_n = 1/8$ for all n, the orthogonal projection of $\mathcal{R}(\{c_i, c_j\})$ onto the plane $\lambda_k = 0$ is as shown in Fig. 2. Accordingly, the throughput region $\mathcal{C}$ for the system is depicted in Fig. 3. Observe that:

- The throughput region $\mathcal{C}$ is now a function only of the marginal probabilities $(\pi_1 + \pi_2) = \mathbb{P}(L(t)_1 = a, L(t)_2 = a)$, etc.
- The throughput region of Fig. 3 has shrunk compared to the region in Fig. 1 due to the pairwise observability constraint.

3) Singleton Channel State Information: Let $O = \{\{c_1\}, \{c_2\}, \{c_3\}\}$, i.e., only the state of one channel can be observed in the first scheduling step. In this case

$$\mathcal{C} = \text{conv}(\mathcal{R}(\{c_1\}), \mathcal{R}(\{c_2\}), \mathcal{R}(\{c_3\})).$$

Each observable subset now has only two substates with corresponding rates a and b; for instance, for the observable subset $\alpha = \{c_1\}$:

- $\mathcal{L}_\alpha = \{(a), (b)\}$
- $\pi^{(a),\alpha} = \pi_1 + \pi_2 + \pi_3 + \pi_4, \pi^{(b),\alpha} = \pi_5 + \pi_6 + \pi_7 + \pi_8$
- $\mu_1^{(a),\alpha} = a, \mu_1^{(b),\alpha} = b$

and similarly for the other observable subsets $\beta = \{c_2\}$ and $\gamma = \{c_3\}$. The subset throughput region $\mathcal{R}(\{c_1\})$, say, is the

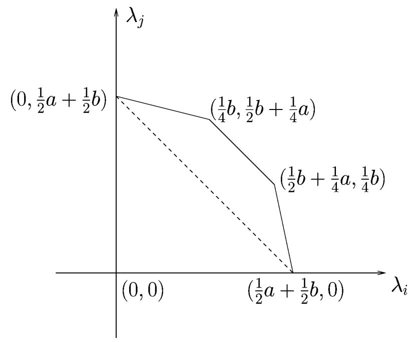

Fig. 2. Rate region for 2 channels c_i and c_j, with complete channel state information.

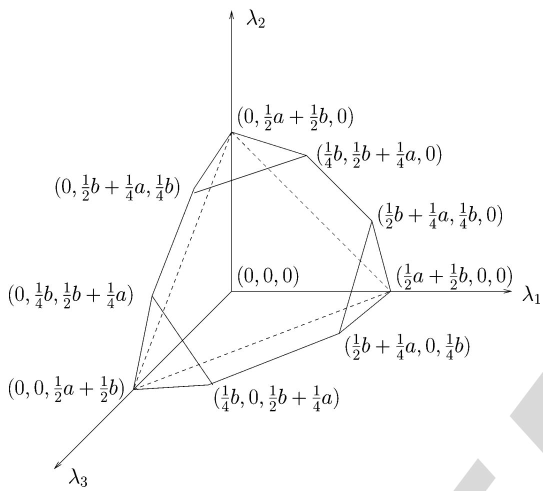

Fig. 3. Rate region for 3 channels with pairwise and singleton channel state information.

set of all rate vectors $(\lambda_1, 0, 0) \in \mathbb{R}^3$, $\lambda_1 \geq 0$, for which there exists a stochastic 2×1 matrix $\phi = \phi^{\{c_1\}}$ such that

$$\phi_{11}(\pi_1 + \pi_2 + \pi_3 + \pi_4)a + \phi_{21}(\pi_5 + \pi_6 + \pi_7 + \pi_8)b > \lambda_1.$$

Using $\pi_n = 1/8$ for all n, we get that $\mathcal{R}(\{c_1\})$ is just the line segment joining (0,0,0) and $((a+b)/2, 0, 0)$, and likewise for $\mathcal{R}(\{c_2\})$ and $\mathcal{R}(\{c_3\})$. Thus the throughput region $\mathcal{C}$ is the dotted simplex which is shown in Fig. 2. Observe that:

- The throughput region $\mathcal{C}$ is now a function only of the marginal probabilities $(\pi_1 + \pi_2 + \pi_3 + \pi_4) = \mathbb{P}(L(t)_1 = a)$, etc.
- The simplex $\mathcal{C}$ is strictly smaller than the throughput region with pairwise channel state information, due to the singleton observability constraint.

C. The "Regret" of a Partial Information Scheduler

We have seen that the throughput region of a system with partial channel state information depends only on the marginal channel state distributions over observable subsets. Let a collection of observable subsets be fixed. Given a joint channel state distribution that induces marginals over the observable subsets, the set of rate vectors that belong to the throughput region with complete channel state information *exclusive of* the throughput region with partial channel state information is a measure of how much a partial information scheduler which 'knows' the joint channel distribution would 'regret' not being able to observe the full instantaneous channel state.

However, given only the marginals over observable subsets, there are, in general, many joint distributions that are consistent with the marginals. In this situation, a natural measure of how much a partial information scheduler would 'regret' not being able to observe the full instantaneous channel state is the set of rate vectors that belong to the throughput region for *every* joint distribution consistent with the given marginals on the observable subsets *exclusive of* the throughput region with partial channel state information. In other words, this 'regret region' is the intersection of the throughput regions for all systems with a consistent joint channel state distribution, excluding the throughput region with partial channel state information over the observable subsets. In this section, we present two examples—the first example demonstrating that the regret region is empty and the second example showing that the regret region can be nonempty (i.e., *any* scheduling policy with complete channel state information can *guaranteeably* support more rates than all policies with partial channel state information).

1) Consider the example of the previous section with pairwise channel state information, i.e., $O = \{\{c_1, c_2\}, \{c_2, c_3\}, \{c_3, c_1\}\}$. Suppose we know the pairwise marginals to be as follows: $\mathbb{P}(L(t)_i = \mu_i, L(t)_j = \mu_j) = \frac{1}{4}, i, j \in \{1, 2, 3\}, i \neq j, \mu_i, \mu_j \in \{a, b\}$. Note that the i.i.d. joint distribution $\pi_1 = \cdots = \pi_8 = 1/8$ used earlier agrees with these pairwise marginals.
These pairwise constraints give us a feasible set of possible *joint* channel distributions: it is the set of vectors $(\pi_1, \ldots, \pi_8)$ in the simplex that satisfy the equations $\pi_1 + \pi_2 = \frac{1}{4}, \pi_1 + \pi_3 = \frac{1}{4}, \pi_1 + \pi_4 = \frac{1}{4}, \pi_2 + \pi_5 = \frac{1}{4}$, etc. In matrix form, these constraints along with the simplex constraints become

$$\begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 \\ & & & \vdots & & & & \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} \pi_1 \\ \pi_2 \\ \pi_3 \\ \pi_4 \\ \pi_5 \\ \pi_6 \\ \pi_7 \\ \pi_8 \end{pmatrix} = \begin{pmatrix} 1/4 \\ 1/4 \\ 1/4 \\ 1/4 \\ \vdots \\ 1/4 \\ 1 \end{pmatrix}$$

with $\pi_i \geq 0$ for all i. The set of solutions for the vector $\vec{\pi} = (\pi_1 \pi_2 \ldots \pi_8)^T$ is the set of convex combinations of the vectors $\vec{\pi}_{(1)} = (1/4 \quad 0 \quad 0 \quad 1/4 \quad 0 \quad 1/4 \quad 1/4 \quad 0)^T$ and $\vec{\pi}_{(2)} = (0 \quad 1/4 \quad 1/4 \quad 0 \quad 1/4 \quad 0 \quad 0 \quad 1/4)^T$, i.e.

$$\vec{\pi} \in \Pi \triangleq \{\eta \vec{\pi}_{(1)} + (1 - \eta)\vec{\pi}_{(2)} : 0 \leq \eta \leq 1\}.$$

The i.i.d. joint distribution $\pi_1 = \cdots = \pi_8 = 1/8$ corresponds to $\eta = 1/2$. Let $\mathcal{C}_{\vec{\pi}}$ denote the throughput region with complete channel state information when the joint distribution of channel states is $\vec{\pi} \in \Pi$. As before $\mathcal{C}$ denotes the throughput region with pairwise partial information, as in Fig. 3. Since $\mathcal{C} \subset \mathcal{C}_{\vec{\pi}} \forall \vec{\pi} \in \Pi$ due to the joint distributions agreeing with the marginals, we must have

$$\mathcal{C} \subset \bigcap_{\vec{\pi} \in \Pi} \mathcal{C}_{\vec{\pi}}.$$

The hexagonal face of $\mathcal{C}$ in Fig. 3 represents the maximum sum rate that can be supported, and is described by $\lambda_1+\lambda_2+\lambda_3 = \frac{1}{4}a+\frac{3}{4}b$. We observe that for the rate region $\mathcal{C}_{\vec{\pi}_{(1)}}$, the sum $\lambda_1+\lambda_2+\lambda_3$ can be at most $\frac{1}{4}a+\frac{3}{4}b$, showing that $\mathcal{C}_{\vec{\pi}_{(1)}} = \mathcal{C}$. Thus we get $\mathcal{C} = \bigcap_{\vec{\pi}\in\Pi}\mathcal{C}_{\vec{\pi}}$. This shows that $\bigcap_{\vec{\pi}\in\Pi}\mathcal{C}_{\vec{\pi}}$—the set of rates which can guaranteeably be supported by scheduling policies with complete state information given pairwise marginals—is no more than $\mathcal{C}$—the set of rates which can be supported by policies with partial channel state information.

2) Our next example illustrates that $\mathcal{C} \subsetneq \bigcap_{\vec{\pi}\in\Pi}\mathcal{C}_{\vec{\pi}}$ in general. Consider two channels c_1 and c_2 which take two states each—rate 1 and rate 2. The aggregate channel thus takes one out of four states in each time slot, with the corresponding rate pairs (μ_1, μ_2) being (1,1), (1,2), (2,1), and (2,2). Let the (joint) probabilities of these states be denoted by π_1, π_2, π_3, and π_4, respectively. We denote the (singleton) observable subsets by $\alpha = \{c_1\}$ and $\beta = \{c_2\}$. Let us constrain the distribution $(\pi_i)_{i=1}^4$ by insisting that the marginals be as follows:

$$\begin{aligned}
\pi^{1,\alpha} &= \mathbb{P}(L(t)_1 = 1) = \pi_1 + \pi_2 = 0.7,\\
\pi^{2,\alpha} &= \mathbb{P}(L(t)_1 = 2) = \pi_3 + \pi_4 = 0.3,\\
\pi^{1,\beta} &= \mathbb{P}(L(t)_2 = 1) = \pi_1 + \pi_3 = 0.4, \quad \text{and}\\
\pi^{2,\beta} &= \mathbb{P}(L(t)_2 = 2) = \pi_2 + \pi_4 = 0.6.
\end{aligned}$$

These are verified to be valid marginals; for instance, the joint probability distributions $\vec{\pi}_{(1)} = (0.1, 0.6, 0.3, 0)$ and $\vec{\pi}_{(2)} = (0.4, 0.3, 0, 0.3)$ induce these marginals. In fact, we can parametrize the set Π of all valid joint distributions which yield these marginals by

$$\Pi = \{\eta\vec{\pi}_{(1)} + (1-\eta)\vec{\pi}_{(2)} : 0 \le \eta \le 1\}.$$

From the marginal distribution, we get $\mathbb{E}[\mu_1] = 1.3$ and $\mathbb{E}[\mu_2] = 1.6$, hence the achievable rate region with partial (singleton) channel state information is as in Fig. 4(a). However, the full channel state information rate region assuming the 'extreme-case' joint distributions $\vec{\pi}_{(1)}$ and $\vec{\pi}_{(2)}$ is as depicted in Fig. 4(b) and (c), respectively. We observe that

$$\bigcap_{\vec{\pi}\in\Pi}\mathcal{C}_{\vec{\pi}} = \mathcal{C}_{\vec{\pi}_{(2)}} \supsetneq \mathcal{C}.$$

Thus, in this case, given the singleton marginals, a scheduler with complete channel state information can support a strictly higher rate guaranteeably over all joint distributions (e.g., the rate (1, 0.6)) than a scheduler with partial channel state information.

D. The Structure of 'Good' SSS Rules

We conclude the section with a theorem which provides a characterization of *maximal* global SSS rules. We call a global SSS rule maximal if no vector in $\mathcal{C}$ dominates its vector of service rates $(v_i)_{i=1}^N$, where a vector $x \in \mathbb{R}^N$ *dominates* a vector $y \in \mathbb{R}^N$ if $x_i \le y_i$ for all i, and $x_j < y_j$ holds for at least one j. The result says that a maximal or optimal global SSS rule chooses the subset that gives the highest expected value of maximum weighted service rate for a subset, and further picks that user to serve that gives the maximum weighted observed rate.

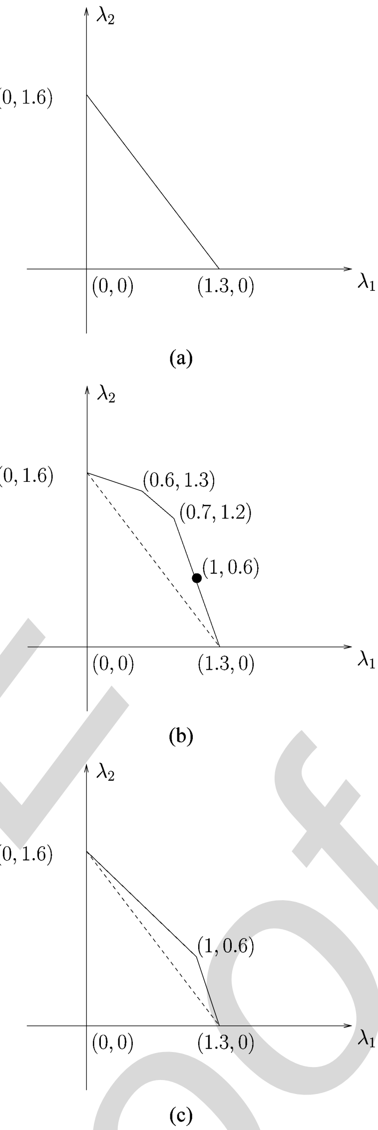

Fig. 4. (a) Rate region with singleton channel state information for 2 channels. (b) Rate region with full channel state information for joint distribution $\pi_{(1)}$. (c) Rate region with full channel state information for joint distribution $\pi_{(2)}$.

Theorem 2: Consider a maximal global SSS rule associated with SSS rules $\{\phi^{*\alpha} : \alpha \in O\}$ and a distribution $\{p^*_\alpha : \alpha \in O\}$ over subsets. Then, there exists a set of strictly positive constants ν_i, $i = 1, \ldots, N$ such that for any l, i and α

$$p^*_\alpha > 0,\ \phi^{*\alpha}_{li} > 0 \Rightarrow i \in \arg\max_{j\in\alpha} \nu_j \mu_j^{l,\alpha}, \quad \text{and} \tag{3}$$

$$p^*_\alpha > 0 \Rightarrow \alpha \in \arg\max_{\beta\in O} \sum_{l\in\mathcal{L}_\beta} \pi^{l,\beta}\left(\max_{j\in\beta}\nu_j\mu_j^{l,\beta}\right). \tag{4}$$

According to Theorem 2, at time t, in the first scheduling step, a maximal global SSS rule chooses a subset α for which

$\sum_{l\in\mathcal{L}_\alpha}\pi^{l,\alpha}(\max_{j\in\alpha}\nu_j\mu_j^{l,\alpha})$ is maximized, and further picks queue i in α which maximizes $\nu_i\mu_i^{l(t),\alpha}$, where $l(t)$ is the observed substate of subset α. We refer the reader to Appendix B for the proof of Theorem 2.

IV. A Throughput-Optimal Scheduling Algorithm

Motivated by the form of the result in Theorem 2, we present a scheduling algorithm which, for a system having arrival rates in the described achievable region, takes as input only the state of the system at each time slot and decides which (maximal) subset to observe and ultimately, which channel in that subset to schedule. Knowledge of the arrival rates is not assumed in such a case. However, it is presumed that the marginal probabilities $\pi^{l,\alpha}$ of the subset α being in the substate l are known.

Algorithm 1

At each time slot t,

- *Step 1:* Select a set $\delta\in O$, given by

$$\delta\in\arg\max_{\alpha\in O}\sum_{l\in\mathcal{L}_\alpha}\pi^{l,\alpha}\left(\max_{i\in\alpha}Q_i(t)\mu_i^{l,\alpha}\right)$$

  where the symbols O, $\mathcal{L}_\alpha$ and $\mu_i^{l,\alpha}$ have the same meaning as in the proof of Lemma 3 and $Q_i(t)$ represents the length of the ith queue at the beginning of time slot t.
- *Step 2:* Let the observed substate of δ be $s\in\mathcal{L}_\delta$. Schedule channel $j\in\delta$ using the *Max-Weight* rule (also known as the *Modified Largest-Weighted-Work-First (M-LWWF)* rule [2], [6]), i.e.

$$j\in\arg\max_{i\in\delta}Q_i(t)\mu_i^{s,\delta}.$$

Note: A suitable rule to break ties in each case is assumed.

The following result provides an important equivalent characterization of the above algorithm in terms of knowing the extreme points of the achievable rate region $\mathcal{C}$. This fact is the basis for the throughput-optimality property of the algorithm, shown by Theorem 3.

Lemma 4: Let $\mathcal{E}$ be the (finite) set of extreme points for the achievable rate region $\mathcal{C}$. If subset δ is chosen in *Step 1* of Algorithm 1 at time t, then

$$\mathcal{R}(\delta)\cap\arg\max_{v\in\mathcal{E}}\langle v,Q(t)\rangle\neq\emptyset.$$

That is, the algorithm selects any subset whose rate region contains an extreme point maximizing the inner product $\langle x,Q(t)\rangle$ over all $x\in\mathcal{E}$ and hence a point maximizing $\langle y,Q(t)\rangle$ over all $y\in\mathcal{C}$. Refer to Appendix C for the proof of Lemma 4.

The chief result in this section is the following theorem, which says that the scheduling policy defined above is throughput-optimal for scheduling with partial channel-state information.

Theorem 3: Algorithm 1 makes the system stable if the vector of arrival rates lies in the achievable region.

The proof of stability uses fluid limit machinery. Roughly, by scaling and "compressing" time and concurrently scaling down the magnitude of the queue length process, the discrete and random queue length process "looks like" a deterministic fluid process which is driven by a (vector) constant rate fluid arrival process (the components corresponding to the mean arrival rates to each of the users), and whose service rate corresponds to the "average" service rate under the scheduling algorithm. For the system we are considering, showing that such a limiting fluid queue length trajectory has negative drift is sufficient to prove that the discrete-time stochastic queue length process is stable (positive recurrent) [26], [2].

The full technical details are deferred to the Appendix, and here we give only the key Lyapunov function idea for proving negative drift. Unlike the proof used for Theorem 3 of [2], here we face the additional difficulty of assuring that we pick the correct observation subset $\delta\in O$, in addition to picking the correct queue to serve in δ. We show that maximizing the negative drift of our Lyapunov function is exactly the problem of maximizing the inner product $\langle y,Q(t)\rangle$ over all $y\in\mathcal{C}$. If we pick the "wrong" subset, then maximizing the linear function above becomes impossible. To side-step this problem, we rely on Lemma 4, which guarantees that the chosen subset will indeed be one with an extreme point maximizing the linear function.

We use the quadratic Lyapunov function

$$L_1(y)=\frac{1}{2}\sum_{i=1}^{N}y_i^2 \tag{5}$$

for a vector $y=(y_1,\ldots,y_N)$. Let $q(t)$ denote the queue-length component of a fluid limit of the system, which in turn is an appropriate collection of almost-sure limits of scaled system processes under uniform convergence over compact sets (see Appendix D.A for details). The following property establishes negative drift of $q(\cdot)$, and (as in [2]) along with a result from [26] implies Theorem 3.

Lemma 5: Under Algorithm 1, for any $\delta_1>0$, there exists $\delta_2>0$ such that almost surely, at any regular point t of the fluid limit

$$L_1(q(t))\geq\delta_1\quad\Rightarrow\quad\frac{d}{dt}L_1(q(t))\leq-\delta_2<0.$$

The proof of this lemma relies on Lemma 4, and can be found in Appendix D-A.

V. The Max-Sum-Queue Algorithm

The throughput-optimal scheduling algorithm in the previous section requires knowledge of both the instantaneous queue lengths and marginal statistics of the channel. In this section, we present a 'simpler' scheduling policy which only uses queue-length information to pick the subset to observe:

Algorithm 2 (Max-Sum-Queue)

At each time slot t,

- *Step 1:* Select a set $\delta \in O$, given by

$$\delta = \arg\max_{\alpha \in O} \sum_{i \in \alpha} Q_i^2(t),$$

where $Q_i(t)$ denotes the length of the ith queue at the beginning of time slot t.

- *Step 2:* Let the observed substate of δ be $s \in \mathcal{L}_\delta$. Schedule channel $j \in \delta$ using the *Max-Weight* rule, i.e.

$$j = \arg\max_{i \in \delta} Q_i(t)\mu_i^{s,\delta}.$$

Note: A suitable rule to break ties in each case is assumed.

In this section, we show that the *Max-Sum-Queue* algorithm is throughput-optimal in two cases of interest: (i) when the subsets in O are disjoint; and (ii) when the channel is symmetric in the users. In the next section, we prove by example that *Max-Sum-Queue* is not throughput-optimal in general.

A. Max-Sum-Queue for Disjoint Subsets

The following result shows that when the collection of observable subsets is mutually disjoint, Max-Sum-Queue is throughput-optimal.

Theorem 4: Under the assumption that every pair of maximal observable subsets is disjoint, the Max-Sum-Queue scheduling algorithm makes the system stable if the vector of arrival rates lies in the achievable region.

To prove Theorem 4, we follow a similar route as in the previous section, defining fluid limits and proving that a suitably defined Lyapunov function has negative drift. The Lyapunov function we use here is

$$L_2(y) = \max_{\beta \in O} h_\beta(y),$$

where

$$h_\beta(y) = \frac{1}{2}\sum_{i \in \beta} y_i^2.$$

The following key lemma is used to establish the negative drift of the Lyapunov function, and is the analog of Lemma 5.

Lemma 6: Under Max-Sum-Queue scheduling, for any $\delta_1 > 0$, there exists $\delta_2 > 0$ such that almost surely, at any regular point t of a fluid limit

$$L_2(q(t)) \geq \delta_1 \quad \Rightarrow \quad \frac{d}{dt}L_2(q(t)) \leq -\delta_2.$$

We refer the reader to Appendix D-B for the details of the proof of the lemma. There is an intuitive geometric explanation for this result. It is based on two observations: first, due to the disjoint subset assumption and the Max-Sum-Queue algorithm, if any queue is unstable, all queues are unstable; next, given an extreme point x_α in each set $\mathcal{R}(\alpha)$, the convex hull of those extreme points will always lie on an exposed face of $\mathcal{C}$. Note that this is not true in the general case.

B. Max-Sum-Queue for Symmetric Channels

It is instructive to note that the reason that the presented scheduling policies work in their respective cases is that at any point $t \in [0, \infty)$, they maximize the linear objective function $\langle q(t), u\rangle$ over all u in the convex polytope $\mathcal{C}$ which represents the achievable rate region. The drift of the sum-of-squares Lyapunov function defined by (5) happens to be precisely the difference between $\langle q(t), \lambda\rangle$ and $\max_{u \in \mathcal{C}}\langle q(t), u\rangle$. This geometric interpretation allows us to prove the useful result that Max-Sum-Queue is actually throughput-optimal for systems of *symmetric channels and subsets*. A symmetric system is where the distribution of the aggregate channel state $L(t)$ is such that *every permutation of a given aggregate channel state occurs with the same probability*. For instance, for a system of three channels $\{c_1, c_2, c_3\}$, the aggregate channel state distributed as (1,0,0), (0,1,0), and (0,0,1) equally likely *qualifies* as a symmetric system, whereas the system with the aggregate channel state distributed as (1,0,0), (0,1,0), (0,0,1), and (0,0,0) equally likely *does not qualify* as a symmetric system.

Theorem 5: Consider a symmetric system, i.e., where the observable subsets are all the subsets of a fixed cardinality K. For such a system, Max-Sum-Queue is throughput-optimal.

Proof: Let λ be the vector of arrival rates to the system of N channels represented by $S = \{1, \ldots, N\}$, such that $\lambda \in \operatorname{int} C$. As before, we consider the drift of the sum-of-squares Lyapunov function defined by (5)

$$\begin{aligned}\frac{d}{dt}L_1(q(t)) &= \sum_{i=1}^{N} q_i(t)(\lambda_i - \hat{f}_i(t)) \\ &= \langle q(t), \lambda\rangle - \langle q(t), \hat{f}(t)\rangle\end{aligned}$$

where $\hat{f}(t) \equiv (\hat{f}_i(t))_{i=1}^N$ is the instantaneous vector of service rates chosen by Max-Sum-Queue at time t in the fluid time scale. We show that $\hat{f}(t) \in \mathcal{C}$ maximizes the inner product $\langle q(t), x\rangle$ over all $x \in \mathcal{C}$ or equivalently over all the extreme points of $\mathcal{C}$; this establishes that the drift of $L_1(q(t))$ is strictly negative and bounded away from zero and hence Max-Sum-Queue is throughput-optimal.

The subsets which Max-Sum-Queue picks for scheduling at t are the ones that contain the top K queues in the system. Without loss of generality, let $q_1(t) \geq q_2(t) \geq \ldots \geq q_N(t)$, and let

$$A = \arg\max_{\beta \subset S, |\beta| = K} \sum_{i \in \beta} q_i^2(t).$$

Every set $\alpha \in A$ is picked by Max-Sum-Queue in the fluid timescale, and thus has the same queue values ordered in descending order. Further, since the channels are symmetric, every subset rate region $\mathcal{R}(\beta)$ for $\beta \subset S$, $|\beta| = K$, is identical up to a permutation of indices. It follows that the extreme points of $\mathcal{C}$ maximizing $\langle \cdot, q(t)\rangle$ must lie in the rate regions $\mathcal{R}(\alpha)$ where $\alpha \in A$, since only the K heaviest queues can maximize this inner product over all permutations of extreme points.

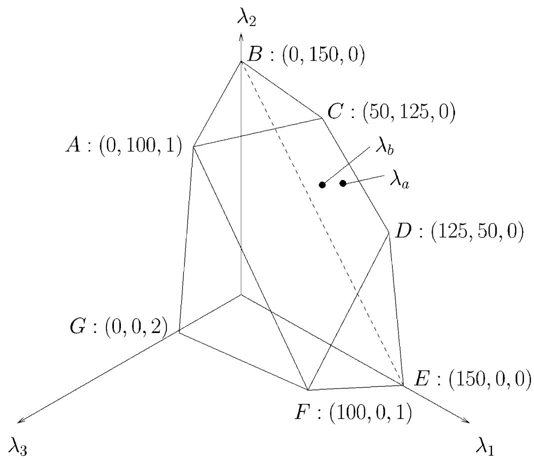

Fig. 5. Rate region for described 3-channel system.

Since these extreme points are precisely the ones picked by Max-Sum-Queue in each subset, and that $\hat{f}(t)$ lies in the convex hull of these extreme points, $\hat{f}(t)$ maximizes the inner product $\langle q(t), x\rangle$ over all $x \in \mathcal{C}$, and we are done. ∎

For an alternative view of why the Max-Sum-Queue policy works for symmetric channels, refer to Appendix E.

VI. Max-Sum-Queue Applied to Arbitrary Subsets

In this section, we show that the simple *Max-Sum-Queue* scheduling algorithm is not throughput-optimal in general. An intuitive fluid argument is presented first, followed by a formal proof. Consider a system of three channels c_1, c_2 and c_3. The system assumes four possible states S_1, S_2, S_3 and S_4 with the corresponding channel rates, expressed by (rate of c_1, rate of c_2, rate of c_3), being (100,100,2), (100,200,2), (200,100,2), and (200,200,2), respectively. Further, each state occurs with probability $\frac{1}{4}$. The maximal observable subsets are $\alpha = \{c_1, c_2\}$, $\beta = \{c_2, c_3\}$ and $\gamma = \{c_3, c_1\}$, i.e., all pairs of channels. The achievable rate region for the system is shown in Fig. 5.

Set the vector of arrival rates to be $\lambda_b \equiv (\lambda_{1b}, \lambda_{2b}, \lambda_{3b}) = (\frac{175}{2}, \frac{175}{2}, 0) - \epsilon(1,1,0) + \delta(0,0,1)$, with $\epsilon = \frac{1}{2}$ and $0 < \delta = \frac{1}{100} < \frac{1}{75}$ (shown in Fig. 5). It is easily verifiable that λ lies in the interior of the rate region. We show that a regular point $t \in [0,\infty)$ can exist with the fluid-limit queue-length process satisfying $q_1(t) = q_2(t) = q_3(t) > 0$, and with $\dot{q}_1(t) = \dot{q}_2(t) = \dot{q}_3(t) > 0$. In such a case, the queue fluid levels $q_1(t)$, $q_2(t)$ and $q_3(t)$ increase (linearly) at a constant rate.

Let us hypothesize that t is a regular point in $[0,\infty)$ satisfying $q_1(t) = q_2(t) = q_3(t) > 0$. Since all the $q_i(t)$ are equal, the system must 'serve' all three subsets with some timesharing probabilities p_α, p_β and p_γ which must be strictly positive. The regularity hypothesis now implies $\dot{q}_1(t) = \dot{q}_2(t) = \dot{q}_3(t)$, and hence

$$\begin{aligned}
\Rightarrow \quad & \lambda_{1b} - 150p_\gamma - \frac{175}{2}p_\alpha \\
& = \lambda_{2b} - 150p_\beta - \frac{175}{2}p_\alpha \\
& = \lambda_{3b} - 0 = \delta \\
\Rightarrow \quad & p_\gamma = p_\beta, \quad \text{and}
\end{aligned}$$

$$150p_\beta + \frac{175}{2}p_\alpha = \lambda_{2b} - \delta = 86.99.$$

Together with $p_\alpha + p_\beta + p_\gamma = 1$, we get $p_\beta = p_\gamma \approx 0.02$ and $p_\alpha \approx 0.96$ which is the unique timesharing solution between the subsets α, β and γ. Hence t is indeed a regular point, all the queue fluid limits are equal, and increase linearly at the same rate $\delta > 0$.

Remarks:

1) We observe that the (mutually exclusive) conditions $q_1(t) = q_2(t) > q_3(t)$ and $q_1(t) = q_2(t) < q_3(t)$ lead to all the q_i becoming equal within finite time. Hence the state $q_1(t) = q_2(t) = q_3(t)$ is an 'unstable attractor' for the fluid limits in this sense.
2) For the arrival rate vector $\lambda_a = (87, 87, 0)$ (shown in Fig. 5), we can similarly show that starting from $q_1(0) = q_2(0) = q_3(0) = c > 0$ implies that $q_1(t) = q_2(t) = q_3(t) = c$ at all times $t \in [0,\infty)$.

Next, as a consequence of the linear growth, we show that the Markov chain describing the state of the system is transient, which implies that all the queues grow without bound almost surely. This is accomplished by demonstrating two crucial properties:

- with high probability the aggregate state of the system of channels is distributed according to the invariant distribution π of the Markov chain describing its evolution, and
- whenever the channel states are typically distributed thus, the smallest queue always grows with a rate bounded away from zero.

Theorem 6: The three-channel system considered is unstable under the Max-Sum-Queue scheduling policy. Furthermore, the Markov chain describing the evolution of its state is transient. The proof is deferred to Appendix F.

VII. Conclusion and Future Work

The *Max-Weight* rule is a striking example of a simple feedback based scheduling policy that is throughput-optimal. Likewise, with partial channel state information, the algorithm we presented which uses queue lengths and expected channel states is throughput-optimal. Under constraints like disjoint set observations or symmetric channels, just looking at the heaviest queues suffices for stability. Both the scheduling algorithms we studied for the partial information case can be viewed as extensions of Max-Weight, which inherit its property of throughput-optimality.

Possible directions for future work include extensions to network-wide scheduling with partial observability of channels. Can the presented scheduling policies be extended to network-wide policies as with *Max-Weight* and the Back Pressure algorithm [27]?

Another line of research would be to study what happens when the channel is correlated across time and when the scheduler is allowed to use the whole past history to make service decisions. For instance, allowing the channel to be Markovian in time leads to a Partially Observable Markov Decision Process (POMDP) problem, and it is interesting to investigate the stability region and the existence of throughput-optimal scheduling policies. Is scheduling based on queue lengths and expected channel states still optimal?

A different direction to pursue would be to examine the delay tails of such scheduling policies. One could also examine the large deviations of the queue lengths arising from these policies.

APPENDIX A
PROOF OF LEMMA 3

Lemma 3: If $\Lambda \in \mathbb{R}^N$ is achievable, then $\Lambda \in \mathcal{C}$. In particular, Λ can be achieved by a global SSS scheduling rule parametrized by a stochastic matrix ϕ of the form

$$\phi = \sum_{\alpha \in O} p_\alpha \phi^\alpha$$

where ϕ^α are stochastic matrices as described above, and p_α is a probability distribution on the maximal observable subsets, O.

Proof: Let $\Lambda = (\lambda_1, \ldots, \lambda_N)$ be supported under the scheduling policy $\mathcal{P} = (\mathcal{G}, \mathcal{H})$. Note that for a maximal observable subset α, the SSS matrix ϕ^α introduced earlier corresponds to a global SSS matrix ϕ where for a row m of ϕ, i.e., a global system state $l \in \mathcal{L}$, columns representing channels in α take the same values as the substate of α induced by l. Other columns are identically zero. Henceforth, by the matrix ϕ^α we will mean the (global) SSS matrix obtained by such an embedding procedure.

Since the discrete-time Markov chain $\mathbf{S}$ representing the evolution of the system is assumed to be ergodic under the policy $\mathcal{P}$, let p_α denote the long term fraction of time in which the maximal observable subset α is chosen in the first scheduling stage, with $\sum_{\alpha \in O} p_\alpha = 1$. Due to the nature of the map $\mathcal{G}$ and the fact that the channel-state process is i.i.d. across time slots, we can write

$$\phi_{li} = \sum_{\alpha \in O} p_\alpha \phi^\alpha_{li}$$

where ϕ_{li} represents the probability with which channel c_i is picked for scheduling in the global system state l. Accordingly, the service rate seen by channel c_i can be written as

$$\begin{aligned}
\lambda_i &= \sum_{l \in \mathcal{L}} \pi^l \mu_i^l \phi_{li} \\
\Rightarrow \lambda_i &= \sum_{l \in \mathcal{L}} \pi^l \mu_i^l \sum_{\alpha \in O} p_\alpha \phi^\alpha_{li} \\
&= \sum_{\alpha \in O} p_\alpha \sum_{l \in \mathcal{L}} \pi^l \mu_i^l \phi^\alpha_{li} \\
&= \sum_{\alpha \in O} p_\alpha \sum_{l=(l_1,l_2) \in \mathcal{L}} \pi^{(l_1,l_2)} \mu_i^{(l_1,l_2)} \phi^\alpha_{(l_1,l_2)i}
\end{aligned}$$

where $l \in \mathcal{L}$ is written (with respect to the maximal observable subset α) as the pair $(l_1.l_2)$ with l_1 denoting the substate of α and l_2 the substate of $\alpha^C = C - \alpha$. We note that if $c_i \notin \alpha$ then $\phi^\alpha_{(l_1,l_2)i} = 0$, and if $c_i \in \alpha$ then $\phi^\alpha_{(l_1,l_2)i} = \phi^\alpha_{l_1 i}$ which is independent of the substate l_2 of α^C. Also, when $c_i \in \alpha$, we denote $\mu_i^{(l_1,l_2)}$ by $\mu_i^{l_1,\alpha}$, the rate of channel c_i in α. For $\alpha \subset C$, we let $\mathcal{L}_\alpha$ denote the set of all possible substates of α. Hence we have

$$\begin{aligned}
\lambda_i &= \sum_{\{\alpha \in O : i \in \alpha\}} p_\alpha \sum_{l_1 \in \mathcal{L}_\alpha} \sum_{l_2 \in \mathcal{L}_{\alpha^C}} \pi^{(l_1,l_2)} \mu_i^{l_1,\alpha} \phi^\alpha_{l_1 i} \\
&= \sum_{\{\alpha \in O : i \in \alpha\}} p_\alpha \left\{ \sum_{l_1 \in \mathcal{L}_\alpha} \left[\sum_{l_2 \in \mathcal{L}_{\alpha^C}} \pi^{(l_1,l_2)} \right] \mu_i^{l_1,\alpha} \phi^\alpha_{l_1 i} \right\}.
\end{aligned} \tag{6}$$

The quantity in square brackets is just the probability $\pi^{l_1,\alpha}$ of the maximal observable subset α being in substate l_1, hence the expression in curly brackets can be labeled λ_i^α: the service rate to channel c_i when only subset α is being observed. Hence

$$\begin{aligned}
\lambda_i &= \sum_{\{\alpha \in O : i \in \alpha\}} p_\alpha \lambda_i^\alpha \\
\Rightarrow \Lambda &= \sum_{\{\alpha \in O : i \in \alpha\}} p_\alpha \Lambda_i^\alpha
\end{aligned}$$

where $\Lambda^\alpha = (\lambda_1^\alpha, \ldots, \lambda_N^\alpha)$ ($\lambda_j^\alpha = 0$ if $c_j \notin \alpha$). Notice that Λ^α is achievable using the trivial distribution on α and the SSS rule $(\phi^\alpha_{l_1 i} : l_1 \in \mathcal{L}_\alpha, c_i \in \alpha)$, hence $\Lambda^\alpha \in \mathcal{R}(\alpha)$. Therefore $\Lambda \in \mathcal{C}$. ■

APPENDIX B
PROOF OF THEOREM 2

Theorem 2: Consider a maximal global SSS rule associated with SSS rules $\{\phi^{*\alpha} : \alpha \in O\}$ and a distribution $\{p^*_\alpha : \alpha \in O\}$ over subsets. Then, there exists a set of strictly positive constants ν_i, $i = 1, \ldots, N$ such that for any l, i and α

$$p^*_\alpha > 0,\ \phi^{*\alpha}_{li} > 0 \Rightarrow i \in \arg\max_{j \in \alpha} \nu_j \mu_j^{l,\alpha}, \quad \text{and} \tag{7}$$

$$p^*_\alpha > 0 \Rightarrow \alpha \in \arg\max_{\beta \in O} \sum_{l \in \mathcal{L}_\beta} \pi^{l,\beta} (\max_{j \in \beta} \nu_j \mu_j^{l,\beta}). \tag{8}$$

For proving Theorem 2, we use the following lemma to first characterize what is meant by a vector of rates being maximal in the rate region.

Lemma 7: For a maximal global SSS rule corresponding to the vector $v = (v_i)_{i=1}^N \in \mathcal{C}$ of service rates, there exist positive constants $\nu_1, \ldots, \nu_N$ such that v solves $\max_{u \in \mathcal{C}} \sum_{i=1}^N \nu_i u_i$.

Remark: The proof of the lemma is an application of the Arrow-Barankin-Blackwell theorem [28]. For the sake of clarity, however, we give the full proof here.

Proof: Let $\{e_1, \ldots, e_M\}$ be all the extreme points of $\mathcal{C}$, with $v = \sum_{j=1}^M p_j e_j$ and $\sum_{j=1}^M p_j = 1$, $p_j \geq 0$ $j = 1, \ldots, M$. Consider the following linear program:

$$\max_{\zeta, \{q_j\}} \zeta$$

subject to

$$\sum_{j=1}^{M} q_j e_{ji} \geq \zeta v_i, \quad \forall i = 1, \ldots, N$$
$$\sum_{j=1}^{M} q_j = 1, \quad 0 \leq q_j \leq 1, \quad \forall j = 1, \ldots, M. \tag{9}$$

where e_{ji} denotes the ith coordinate of e_j. We know that $\zeta = 1$ and $\{q_j\} = \{p_j\}$ solve this linear program with constraints (9) satisfied as equalities. Then, by the Kuhn-Tucker theorem [29], there exists a set of nonnegative Lagrange multipliers $\nu_0, \ldots, \nu_N$ such that $\zeta = 1$ and $\{q_j\} = \{p_j\}$ also solve the following linear program (with the same value of the maximum):

$$\max_{\zeta, \{q_j\}} \nu_0 \zeta + \sum_{i=1}^{N} \nu_i \left(\sum_{j=1}^{M} q_j e_{ji} - \zeta v_i \right) \tag{10}$$

subject to

$$\sum_{j=1}^{M} q_j = 1, \quad 0 \leq q_j \leq 1, \quad \forall j = 1, \ldots, M.$$

We note that every ν_i must be strictly positive owing to the tightness in (9), and $\nu_0 = 1$. Rewriting (10), we get that $\{q_j\} = \{p_j\}$ maximizes

$$\sum_{i=1}^{N} \alpha_i \sum_{j=1}^{M} q_j e_{ji}$$

over all distributions $\{q_j\}$, i.e., $v = (v_i)_{i=1}^N$ maximizes $\sum_{i=1}^{N} \nu_i u_i$ over all $u \in \mathcal{C}$. ■

Proof of Theorem 2: Let v^* be the vector of long-term service rates for a maximal global SSS rule parametrized by the distributions $\{\phi^{*\alpha} : \alpha \in O\}$ and $\{p^*_\alpha : \alpha \in O\}$. We have by Lemma 7 that there exist positive constants $\nu_1, \ldots, \nu_N$ such that v^* solves

$$\begin{aligned} &\max_{v \in \mathcal{C}} \sum_{i=1}^{N} \nu_i v_i \\ &= \max_{\{p_\nu\}, \{\phi^\nu\}} \sum_{i=1}^{N} \nu_i \sum_{\nu} p_\nu \sum_{l \in \mathcal{L}_\nu} \pi^{l,\nu} \mu_i^{l,\nu} \phi_{li}^{\nu} \\ &= \max_{\{p_\nu\}, \{\phi^\nu\}} \sum_{\nu} p_\nu \sum_{l \in \mathcal{L}_\nu} \pi^{l,\nu} \sum_{i=1}^{N} \nu_i \mu_i^{l,\nu} \phi_{li}^{\nu}. \end{aligned} \tag{11}$$

Equivalently, $\{p_\alpha\} = \{p^*_\alpha\}$ and $\{\phi^\alpha\} = \{\phi^{*\alpha}\}$ solves (11), and (7) and (8) of the theorem follow, since otherwise the maximum in (11) would not be attained. ■

Appendix C
Proof of Lemma 4

Lemma 4: Let $\mathcal{E}$ be the (finite) set of extreme points for the achievable rate region $\mathcal{C}$. If subset δ is chosen in *Step 1* of Algorithm 1 at time t, then

$$\mathcal{R}(\delta) \cap \arg\max_{v \in \mathcal{E}} \langle v, Q(t) \rangle \neq \emptyset.$$

Proof: If $v \in \mathcal{C}$, then

$$\begin{aligned} \langle v, Q(t) \rangle &= \sum_i Q_i(t) v_i \\ &= \sum_i Q_i(t) \sum_{\alpha \in O} p_\alpha \sum_{l \in \mathcal{L}_\alpha} \pi^{l,\alpha} \mu_i^{l,\alpha} \phi_{li}^{\alpha} \\ &\quad \text{(from Theorem 1 and (2))} \\ &= \sum_{\alpha \in O} p_\alpha \sum_{l \in \mathcal{L}_\alpha} \pi^{l,\alpha} \sum_{i \in \alpha} Q_i(t) \mu_i^{l,\alpha} \phi_{li}^{\alpha} \qquad (12) \\ &\leq \sum_{\alpha \in O} p_\alpha \sum_{l \in \mathcal{L}_\alpha} \pi^{l,\alpha} \left(\max_{i \in \alpha} Q_i(t) \mu_i^{l,\alpha} \right) \\ &\leq \max_{\alpha \in O} \sum_{l \in \mathcal{L}_\alpha} \pi^{l,\alpha} \left(\max_{i \in \alpha} Q_i(t) \mu_i^{l,\alpha} \right) \\ &= \sum_{l \in \mathcal{L}_\delta} \pi^{l,\delta} \left(\max_{i \in \delta} Q_i(t) \mu_i^{l,\delta} \right) \\ &\quad \text{(by hypothesis and definition of Algorithm 1).} \end{aligned}$$

Let $k(l) \triangleq \arg\max_{i \in \delta} Q_i(t) \mu_i^{l,\delta}$, $l \in \mathcal{L}_\delta$, with ties broken according to a fixed precedence rule among $i \in \delta$. Define a SSS rule $\phi^{*\delta}$ which serves only subset δ, and for which $\phi^{*\delta}_{li} = i$ if $i = k(l)$ and 0 otherwise, $l \in \mathcal{L}_\delta$. If $u \equiv (u_1, \ldots, u_N)$ is the vector of long term service rates for $\phi^{*\delta}$, then we have $u \in \mathcal{R}(\delta)$, and

$$\begin{aligned} \langle u, Q(t) \rangle &= \sum_i Q_i(t) u_i \\ &= \sum_{l \in \mathcal{L}_\delta} \pi^{l,\delta} \sum_{i \in \delta} Q_i(t) \mu_i^{l,\delta} \phi_{li}^{*\delta} \quad \text{(following(12))} \\ &= \sum_{l \in \mathcal{L}_\delta} \pi^{l,\delta} \left(\max_{i \in \delta} Q_i(t) \mu_i^{l,\delta} \right) \quad \text{(definition of } \phi^{*\delta}\text{).} \end{aligned}$$

Since $\mathcal{R}(\delta)$ is a convex polytope, it must contain an extreme point w such that $\max_{v \in \mathcal{E}} \langle v, Q(t) \rangle = \langle w, Q(t) \rangle = \langle u, Q(t) \rangle$. Thus $w \in \mathcal{R}(\delta) \cap \arg\max_{v \in \mathcal{E}} \langle v, Q(t) \rangle$ which proves the lemma. ■

Appendix D

The following two Appendices (D.A and D.B) are provided for completeness. The setup in these sections parallels that in [2], and uses the machinery of fluid limits to establish stability of Algorithm 1 and Algorithm 2 (Max-Sum-Queue) via drift properties.

A. Proof of Lemma 5

Lemma 5: Under Algorithm 1, for any $\delta_1 > 0$, there exists $\delta_2 > 0$ such that almost surely, at any regular point t of the fluid limit,

$$L_1(q(t)) \geq \delta_1 \quad \Rightarrow \quad \frac{d}{dt} L_1(q(t)) \leq -\delta_2 < 0.$$

To prove the lemma, we set up fluid limit processes for the system dynamics following the development in [2]. For this purpose, we first define the "norm" of the system state $S(t)$ as $\|S(t)\| \triangleq \sum_1^N Q_i(t)$. Consider a sequence of queueing systems, indexed by $n = 1, 2, \ldots$, with corresponding system state processes $S^{(n)}(t), t = 0, 1, \ldots$ such that $\|S^{(n)}(0)\| = n$. We define the following discrete time random processes for $t = 0, 1, \ldots$:

$$
\begin{aligned}
F_i^{(n)}(t) &\triangleq \#\text{ of packets to queue } i \text{ that arrived by time } t,\\
\hat{F}_i^{(n)}(t) &\triangleq \#\text{ of queue } i'\text{s packets that got served by time } t,\\
C_\alpha^{(n)}(t) &\triangleq \#\text{ of time slots before } t \text{ when subset } \alpha \text{was}\\
&\quad\ \text{chosen for scheduling},\\
G_l^{\alpha,(n)}(t) &\triangleq \#\text{ of time slots before } t \text{ when subset } \alpha \text{was}\\
&\quad\ \text{picked and its substate was } l,\\
\hat{G}_{li}^{\alpha,(n)} &\triangleq \#\text{ of time slots before } t \text{ when subset } \alpha \text{was}\\
&\quad\ \text{picked its observed substate was } l \text{ and queue } i\\
&\quad\ \text{was scheduled for service.}
\end{aligned}
$$

As in [2], for a canonical discrete time process $Z^{(n)}(t)$, define its corresponding *scaled* (by n in space and time) process in continuous time by

$$z^{(n)}(t) \triangleq \frac{1}{n} Z^{(n)}(\lfloor nt \rfloor),\ t \geq 0$$

following which we get the scaled versions of our system processes:

$$
\begin{aligned}
f^{(n)} &= \left(f_i^{(n)}(t),\ t \geq -1,\ i = 1, 2, \ldots, N\right)\\
\hat{f}^{(n)} &= \left(\hat{f}_i^{(n)}(t),\ t \geq 0,\ i = 1, 2, \ldots, N\right)\\
c^{(n)} &= \left(c_\alpha^{(n)}(t),\ t \geq 0,\ \alpha \in O\right)\\
g^{(n)} &= \left(g_l^{\alpha,(n)}(t),\ t \geq 0,\ l \in \mathcal{L}_\alpha,\ \alpha \in O\right)\\
\hat{g}^{(n)} &= \left(\hat{g}_{li}^{\alpha,(n)}(t),\ t \geq 0,\ l \in \mathcal{L}_\alpha,\ \alpha \in O,\ i = 1, 2, \ldots, N\right)\\
q^{(n)} &= \left(q_i^{(n)}(t),\ t \geq 0,\ i = 1, 2, \ldots, N\right).
\end{aligned}
$$

For the sake of completeness, we reproduce (with minor modifications) the following lemma from [2], which establishes convergence of these scaled processes to the corresponding *fluid limit* processes. These fluid limit processes have desirable properties like being absolutely continuous and thus differentiable almost everywhere, non-decreasing and time-conserving.

[2, (Lemma 1)]. The following statements hold with probability 1. For any sequence of processes $X^{(n)}$, there exists a subsequence $X^{(k)}$, $k \subseteq n$, such that for each i, $1 \leq i \leq N$, $\alpha \in O$ and $l \in \mathcal{L}_\alpha$

$$
\begin{aligned}
\left(f_i^{(k)}(t), t \geq -1\right) &\Rightarrow (f_i(t), t \geq -1)\\
\left(f_i^{(k)}(t), t \geq 0\right) &\to (f_i(t), t \geq 0) \quad u.o.c.\\
\left(\hat{f}_i^{(k)}(t), t \geq 0\right) &\to (\hat{f}_i(t), t \geq 0) \quad u.o.c.\\
\left(q_i^{(k)}(t), t \geq 0\right) &\to (q_i(t), t \geq 0) \quad u.o.c.\\
\left(g_l^{\alpha,(k)}(t), t \geq 0\right) &\to (g_l^\alpha(t), t \geq 0) \quad u.o.c.\\
\left(\hat{g}_{li}^{\alpha,(k)}(t), t \geq 0\right) &\to (\hat{g}_{li}^\alpha(t), t \geq 0) \quad u.o.c.\\
\left(c_\alpha^{(k)}(t), t \geq 0\right) &\to (c_\alpha(t), t \geq 0) \quad u.o.c.
\end{aligned}
$$

where the functions f_i are nonnegative, nondecreasing, and right-continuous with left limits (RCLL) in $[-1, \infty)$, the functions f_i, $\hat{f}_i$, g_l^α, $\hat{g}_{li}^\alpha$, c_α are nonnegative nondecreasing Lipschitz-continuous in $[0, \infty)$, functions q_i are continuous in $[0, \infty)$, "$\Rightarrow$" signifies convergence at continuity points of the limit, and "u.o.c." means uniform convergence on compact sets, as $k \to \infty$. The limiting set of functions

$$x = (f, \hat{f}, g, \hat{g}, q, c)$$

also satisfies the following properties, for all i, $1 \leq i \leq N$, $\alpha \in O$ and $l \in \mathcal{L}_\alpha$:

$$
\begin{aligned}
f_i(t) - f_i(0) &= \lambda_i t,\ t \geq 0,\\
\hat{f}_i(0) &= 0,\\
\hat{f}_i(t) &\leq f_i(t),\ t \geq 0,\\
\sum_{i \in \alpha} \hat{g}_{li}^\alpha(t) &= g_l^\alpha(t),\ t \geq 0,\\
g_l^\alpha(t) &= \pi^{l,\alpha} c_\alpha(t),\ t \geq 0,\\
\sum_{\alpha \in O} c_\alpha(t) &= t,\ t \geq 0
\end{aligned}
$$

for any interval $[t_1, t_2] \subset [0, \infty)$

$$\hat{f}_i(t_2) - \hat{f}_i(t_1) \leq \sum_{\alpha \in O} \sum_{l \in \mathcal{L}_\alpha} \mu_i^{l,\alpha} \left(\hat{g}_{li}^\alpha(t_1) - \hat{g}_{li}^\alpha(t_2)\right)$$

if $q_i(t) > 0$ for $t \in [t_1, t_2] \subset [0, \infty)$, then

$$\hat{f}_i(t_2) - \hat{f}_i(t_1) = \sum_{\alpha \in O} \sum_{l \in \mathcal{L}_\alpha} \mu_i^{l,\alpha} \left(\hat{g}_{li}^\alpha(t_1) - \hat{g}_{li}^\alpha(t_2)\right).$$

Analogous to Lemma 2 in [2], Algorithm 1 can be shown to induce the following properties on the fluid limit processes via the corresponding pre-limit processes:

1) If, for some regular point $t \geq 0$ and some i, l and α

$$\mu_i^{l,\alpha} q_i(t) < \max_j \mu_j^{l,\alpha} q_j(t)$$

then

$$\hat{g}_{li}^{'\alpha}(t) = 0.$$

2) If, for some regular point $t \geq 0$ and some η

$$\sum_{l \in \mathcal{L}_\eta} \pi^{l,\eta} \left(\max_{i \in \eta} q_i(t) \mu_i^{l,\eta}\right) < \max_{\alpha \in O} \sum_{l \in \mathcal{L}_\alpha} \pi^{l,\alpha} \left(\max_{i \in \alpha} q_i(t) \mu_i^{l,\alpha}\right)$$

then

$$c_\eta'(t) = 0.$$

Proof of Lemma 5: Since the system is assumed to be feasible, its rate vector is a convex combination of feasible rate vectors of its maximal observable subsets, by Lemma 3. Hence there must exist a fixed distribution $\{p_\alpha\}_{\alpha\in O}$ together with subset SSS rules $\{\phi^\alpha\}_{\alpha\in O}$ such that, using Theorem 1 of [2] and (6), we have

$$\lambda_i < v_i(\{p_\alpha\},\{\phi^\alpha\}) := \sum_{\alpha\in O} p_\alpha \sum_{l\in\mathcal{L}_\alpha} \pi^{l,\alpha}\mu_i^{l,\alpha}\phi_{li}^\alpha.$$

For any regular $t\geq 0$ such that $L_1(q(t))>0$, the derivative of $L_1(q(t))$ can be written as follows:

$$\begin{aligned}\frac{d}{dt}L_1(q(t)) &= \sum_{i=1}^N q_i(t)(\lambda_i - \hat{f}_i'(t))\\ &= \sum_{i=1}^N q_i(t)(\lambda_i - v_i(\{p_\alpha\},\{\phi^\alpha\})) + K(\{p_\alpha\},\{\phi^\alpha\},q(t))\\ &\quad - K(\{c_\alpha'(t)\},\{\hat{\phi}^\alpha\},q(t)) \end{aligned}\tag{13}$$

where we use the notation

$$K(\{r_\alpha\},\{\psi^\alpha\},y) \equiv \sum_i y_i \sum_\alpha r_\alpha \sum_{l\in\mathcal{L}_\alpha}\pi^{l,\alpha}\mu_i^{l,\alpha}\psi_{li}^\alpha,$$
$$\hat{\phi}_{li}^\alpha \equiv \frac{\hat{g}_{li}'^{\alpha}(t)}{\pi^{l,\alpha}c_\alpha'(t)}$$

and we use the fact, following from properties of the fluid limits, that

$$\hat{f}_i'(t) = \sum_{\alpha\in O}\sum_{l\in\mathcal{L}_\alpha}\mu_i^{l,\alpha}\hat{g}_{li}'^{\alpha}(t).$$

We can always choose $\delta_3>0$ such that $L_1(y)\geq\delta_1$ implies $\max_i y_i\geq\delta_3$. Then the first sum in (13) is bounded as follows:

$$\begin{aligned}&\sum_{i=1}^N q_i(t)(\lambda_i - v_i(\{p_\alpha\},\{\phi^\alpha\}))\\ &\quad\leq -\delta_3\min_i(v_i(\{p_\alpha\},\{\phi^\alpha\})-\lambda_i)\equiv-\delta_2.\end{aligned}$$

It remains to show that

$$K(\{p_\alpha\},\{\phi^\alpha\},q(t)) \leq K(\{c_\alpha'(t)\},\{\hat{\phi}^\alpha\},q(t)).\tag{14}$$

Using Properties 1 and 2 of the fluid limit processes under Algorithm 1, we have

$$\begin{aligned}&K(\{c_\alpha'(t)\},\{\hat{\phi}^\alpha\},q(t))\\ &\quad=\sum_i q_i(t)\sum_\alpha c_\alpha'(t)\sum_{l\in\mathcal{L}_\alpha}\pi^{l,\alpha}\mu_i^{l,\alpha}\hat{\phi}_{li}^\alpha\\ &\quad=\sum_\alpha c_\alpha'(t)\sum_{l\in\mathcal{L}_\alpha}\pi^{l,\alpha}\sum_i q_i(t)\mu_i^{l,\alpha}\hat{\phi}_{li}^\alpha\\ &\quad=\max_\alpha\sum_{l\in\mathcal{L}_\alpha}\pi^{l,\alpha}\left(\max_i q_i(t)\mu_i^{l,\alpha}\right)\\ &\quad=\max_\alpha\sum_{l\in\mathcal{L}_\alpha}\pi^{l,\alpha}\sum_i q_i(t)\mu_i^{l,\alpha}\hat{\phi}_{li}^\alpha\\ &\quad\geq\sum_\alpha p_\alpha\sum_{l\in\mathcal{L}_\alpha}\pi^{l,\alpha}\left(\max_i q_i(t)\mu_i^{l,\alpha}\right)\\ &\quad\geq\sum_\alpha p_\alpha\sum_{l\in\mathcal{L}_\alpha}\pi^{l,\alpha}\sum_i q_i(t)\mu_i^{l,\alpha}\phi_{li}^\alpha\\ &\quad=K(\{p_\alpha\},\{\phi^\alpha\},q(t)).\end{aligned}$$

This proves Lemma 5. ■

B. Proof of Lemma 6

Lemma 6: Under Max-Sum-Queue scheduling, for any $\delta_1>0$, there exists $\delta_2>0$ such that almost surely, at any regular point t of a fluid limit

$$L_2(q(t))\geq\delta_1 \quad\Rightarrow\quad \frac{d}{dt}L_2(q(t))\leq-\delta_2.$$

To prove Lemma 6, we use the same framework of fluid limits as before (Section D.A), this time working with a new Lyapunov function $L_2(\,\cdot\,)$ where

$$L_2(y)\triangleq\max_{\beta\in O}h_\beta(y),$$
and
$$h_\beta(y)\triangleq\frac{1}{2}\sum_{i\in\beta}y_i^2.$$

The following properties of the fluid limits under Max-Sum-Queue scheduling follow from the properties of the pre-limit processes, in a manner similar to [2, Lemma 2]:

1) If, for some regular point $t\geq 0$ and some i, l and α

$$\mu_i^{l\cdot\alpha}q_i(t)<\max_j\mu_j^{l,\alpha}q_j(t)$$

then

$$\hat{g}_{li}'^{\alpha}(t)=0.$$

2) If, for some regular point $t\geq 0$ and some η

$$h_\eta(q(t))<\max_{\alpha\in O}h_\alpha(q(t))$$

then

$$c_\eta'(t)=0.$$

Proof of Lemma 6: Let $\arg\max_{\alpha\in O}h_\alpha(q(t))$ be comprised of $\beta_1,\ldots,\beta_m$ with each $\beta_i\in O$, where β_1 is chosen by a fixed precedence rule among subsets in O; thus $L_2(q(t))=\frac{1}{2}\sum_{i\in\beta_1}q_i^2(t)$. For a regular point $t\in[0,\infty)$, we have

$$\frac{d}{dt}L_2(q(t))=\frac{d}{dt}h_{\beta_1}(t)=\sum_{i\in\beta_1}q_i(t)(\lambda_i-\hat{f}_i'(t))\tag{15}$$

$$\begin{aligned}&=\sum_{i\in\beta_1}q_i(t)\left(\lambda_i-\sum_{l\in\mathcal{L}_{\beta_1}}\mu_i^{l,\beta_1}\hat{g}_{li}'^{\alpha}(t)\right)\\ &=\sum_{i\in\beta_1}q_i(t)\lambda_i-c_{\beta_1}'(t)\sum_{l\in\mathcal{L}_{\beta_1}}\pi^{l,\beta_1}\sum_{i\in\beta_1}q_i(t)\mu_i^{l,\beta_1}\hat{\phi}_{li}^{\beta_1}\end{aligned}$$

$$
\begin{aligned}
&= \sum_{i\in\beta_1} q_i(t)\lambda_i - c'_{\beta_1}(t) \sum_{l\in\mathcal{L}_{\beta_1}} \pi^{l,\beta_1}(\max_i q_i(t)\mu_i^{l,\beta_1}) \\
&= \sum_{i\in\beta_1} q_i(t)\lambda_i - c'_{\beta_1}(t) K_{\beta_1}(q(t)) \\
&= \langle q(t),\lambda\rangle_{\beta_1} - c'_{\beta_1}(t)K_{\beta_1}(q(t)) = w, \quad \text{say} \qquad (16)
\end{aligned}
$$

where $\hat{\phi}^{\alpha}_{li} \triangleq \frac{\hat{g}'^{\alpha}_{li}(t)}{\pi^{l,\alpha} c'_{\alpha}(t)}$, $K_\alpha(y) \triangleq \sum_{l\in\mathcal{L}_\alpha} \pi^{l,\alpha}(\max_{i\in\alpha} y_i\mu_i^{l,\alpha})$ for $\alpha \in O$, $\langle x,y\rangle_\alpha \triangleq \sum_{i\in\alpha} x_i y_i$, and $\sum_{j=1}^{m} c'_{\beta_j}(t) = 1$. More generally, as a result of the above, we have

$$
\frac{d}{dt}h_{\beta_j}(t) = \langle q(t),\lambda\rangle_{\beta_j} - c'_{\beta_j}(t)K_{\beta_j}(q(t)) \quad \forall j = 1,\ldots,m. \qquad (17)
$$

Since t is a regular point, we have, using (17),

$$
\begin{aligned}
&\frac{d}{dt}h_{\beta_j}(q(t)) = \frac{d}{dt}h_{\beta_1}(q(t)), \quad \forall j = 1,\ldots,m \\
\Rightarrow\ &\langle q(t),\lambda\rangle_{\beta_j} - c'_{\beta_j}(t)K_{\beta_j}(q(t)) \\
&\quad = \langle q(t),\lambda\rangle_{\beta_1} - c'_{\beta_1}(t)K_{\beta_1}(q(t)) \quad \forall j = 1,\ldots,m.
\end{aligned}
$$

Define

$$
\tilde{q}_i(t) \triangleq \frac{q_i(t)}{K_{\beta_j}(q(t))} \quad \text{if} \quad i \in \beta_j \text{ for some } j, \text{and } 0 \text{ otherwise}
$$

and let $\tilde{q}(t) \triangleq (\tilde{q}_1(t),\ldots,\tilde{q}_N(t))$. $\tilde{q}(t)$ is well-defined since a queue i belongs to at most one of the (disjoint) β_j. Consider

$$
\begin{aligned}
\langle \tilde{q}(t), \hat{f}'(t)\rangle &= \sum_{j=1}^{m} \langle \tilde{q}(t), \hat{f}'(t)\rangle_{\beta_j} \\
&= \sum_{j=1}^{m}\sum_{i\in\beta_j} \tilde{q}_i(t)\hat{f}'_i(t) \\
&= \sum_{j=1}^{m} [K_{\beta_j}(q(t))]^{-1} \sum_{i\in\beta_j} q_i(t)\hat{f}'_i(t) \\
&= \sum_{j=1}^{m} [K_{\beta_j}(q(t))]^{-1} c'_{\beta_j}(t) K_{\beta_j}(q(t)) \\
&\quad (\text{using}(15)-(16)) \\
&= \sum_{j=1}^{m} c'_{\beta_j}(t) = 1.
\end{aligned}
$$

Note that due to Step 2 of the Max-Sum-Queue policy, $\hat{f}'(t)|_{\beta_j}$ maximizes $\langle q(t), x\rangle_{\beta_j}$ over all $x \in \mathcal{C}$. Hence by the above we have that $\hat{f}'(t)$ maximizes $\langle \tilde{q}(t), x\rangle_{\cup_j \beta_j}$ over all $x \in \mathcal{C}$. Since λ lies in the interior of $\mathcal{C}$, there exists $\epsilon > 0$ such that

$$
\begin{aligned}
&\langle \tilde{q}(t),\lambda\rangle - \langle \tilde{q}(t), \hat{f}(t)\rangle \le -\epsilon. \\
\Rightarrow\ &\langle \tilde{q}(t),\lambda\rangle - \sum_{j=1}^{m} c'_{\beta_j}(t) \le -\epsilon \\
\Rightarrow\ &\sum_{j=1}^{m} \langle \tilde{q}(t),\lambda\rangle_{\beta_j} - \sum_{j=1}^{m} c'_{\beta_j}(t) \le -\epsilon \\
\Rightarrow\ &\sum_{j=1}^{m} \frac{\langle q(t),\lambda\rangle_{\beta_j}}{K_{\beta_j}(q(t))} - \sum_{j=1}^{m} c'_{\beta_j}(t) \le -\epsilon \\
\Rightarrow\ &\sum_{j=1}^{m} \left(\frac{\langle q(t),\lambda\rangle_{\beta_j} - c'_{\beta_j}(t)K_{\beta_j}(q(t))}{K_{\beta_j}(q(t))} \right) \le -\epsilon \\
\Rightarrow\ &w \sum_{j=1}^{m} K^{-1}_{\beta_j}(q(t)) \le -\epsilon. \\
\therefore\ &w \le -\epsilon \left[\sum_{j=1}^{m} K^{-1}_{\beta_j}(q(t)) \right]^{-1} =: -\delta_2
\end{aligned}
$$

whenever $L_2(q(t)) \ge \delta_1$, since $[\sum_{j=1}^{m} K^{-1}_{\beta_j}(q(t))]^{-1}$ is monotone increasing in $q(t)$.

This establishes the strictly negative drift and concludes the proof. ■

Appendix E
Proof of Theorem 5

Theorem 5: Consider a symmetric system, where the observable subsets are all subsets of a fixed cardinality K. For such a system, Max-Sum-Queue is throughput-optimal.

Proof: We show that Max-Sum-Queue is equivalent to the throughput-optimal rule defined in Section 4 for a symmetric system. The throughput-optimal algorithm picks a subset $\gamma \in O$ such that

$$
\gamma \in \arg\max_{\alpha\in O} \sum_{l\in\mathcal{L}_\alpha} \pi^{l,\alpha} \left(\max_{i\in\alpha} Q_i(t)\mu_i^{l,\alpha} \right) \qquad (18)
$$

while Max-Sum-Queue picks a subset $\delta \in O$ such that

$$
\delta \in \arg\max_{\alpha\in O} \sum_{i\in\alpha} Q_i^2(t)
$$

i.e., δ contains the top K queues at time t. We claim that δ belongs to the set in the right hand side of (18). For if not, there exists a subset $\eta \neq \delta$ such that η does not contain the top K queues and

$$
\sum_{l\in\mathcal{L}_\eta} \pi^{l,\eta} \left(\max_{i\in\eta} Q_i(t)\mu_i^{l,\eta} \right) > \sum_{l\in\mathcal{L}_\delta} \pi^{l,\delta} \left(\max_{j\in\delta} Q_j(t)\mu_j^{l,\delta} \right).
$$

Assume without loss of generality that the queues $\{Q_i\}_{i\in\eta}$ and $\{Q_j\}_{j\in\delta}$ are ordered in descending order within subsets η and δ respectively. Since the system is symmetric, each subset sees identical substates with identical distributions on them, so we can assume that $\mathcal{L}_\eta$ and $\mathcal{L}_\delta$ are also identical, along with the corresponding sets of $\pi^{l,\eta}$ and $\pi^{l,\delta}$. Hence we have

$$
\sum_{l\in\mathcal{L}_\eta} \pi^{l,\eta} \left[\left(\max_{i\in\eta} Q_i(t)\mu_i^{l,\eta} \right) - \left(\max_{j\in\delta} Q_j(t)\mu_j^{l,\delta} \right) \right] > 0
$$

which is a contradiction as each queue in δ is at least as large as its corresponding queue in η. This proves the theorem. ■

Appendix F
Proof of Theorem 6

Theorem 6: The three-channel system considered in Section VI is unstable under the Max-Sum-Queue scheduling policy. Furthermore, the Markov chain describing the evolution of its state is transient.

We will need technical preliminaries similar to [30] to prove Theorem 6. As stated earlier, fix the vector of arrival rates to be $\lambda = (87, 87, 0.01)$. For each n we split the nonnegative real line $[0, \infty)$ into equal contiguous intervals of size nT each. Let the kth interval $[(k-1)nT, knT)$ be denoted by C_k. We divide every interval uniformly into $P_T^n \triangleq nT/n^{1/4}$ equal contiguous subintervals of size $n^{1/4}$. Define:

- $A_j^{i,n,k}$: Number of arrivals from flow j in the ith subinterval of the kth interval, i.e., in $[(k-1)nT + (i-1)n^{1/4}, (k-1)nT + in^{1/4})$, and
- $B_m^{i,n,k}$: Number of time slots that the channel is in state m in the ith subinterval of the kth interval, i.e., in $[(k-1)nT + (i-1)n^{1/4}, (k-1)nT + in^{1/4})$.

We define the following arrival process and channel process deviation events:

$$E_j^{n,k}(T,\nu) = \bigcup_{1 \le i \le P_T^n} \left\{ \left| \frac{A_j^{i,n,k}}{n^{\frac{1}{4}}} - \lambda_j \right| > \nu \right\},$$

$$G_m^{n,k}(T,\nu) = \bigcup_{1 \le i \le P_T^n} \left\{ \left| \frac{B_m^{i,n,k}}{n^{\frac{1}{4}}} - \pi^m \right| > \nu \right\}.$$

Note that by hypothesis, the events $E_j^{n,k}(T,\nu)$ and $G_m^{n,k}(T,\nu)$ are independent for any n, m, j, k, $T > 0$ and $\nu > 0$. For positive integers n and k and real numbers $T > 0$ and $\nu > 0$ we define the following error event, corresponding to at least one of the channel service rates or input flows being 'atypical' in its empirical distribution in the kth time interval:

$$F^{n,k}(T,\nu) = \left(\bigcup_j E_j^{n,k}(T,\nu) \right) \bigcup \left(\bigcup_m G_m^{n,k}(T,\nu) \right).$$

The following lemma bounds from above the probability of this error event.

Lemma 8: Fix $T > 0$, $\nu > 0$ and $\epsilon > 0$. Then, there exists $n_0 = n_0(T,\nu,\epsilon) > 0$ such that for all $n > n_0$, for any fixed positive integer k, and uniformly over $a_1, \ldots, a_{(k-1)nT-1}$, we have

$$\mathbb{P}(F^{n,k}(T,\nu) | A(1) = a_1, \ldots, A((k-1)nT-1) = a_{(k-1)nT-1}) < \epsilon.$$

Proof: Since by hypothesis there are only finitely many channels and aggregate channel states, it suffices to show that

$$\mathbb{P}\left(E_j^{n,k}(T,\nu) | A(1) = a_1, \ldots, A((k-1)nT-1) = a_{(k-1)nT-1}\right) < \epsilon, \tag{19}$$

and

$$\mathbb{P}\left(G_m^{n,k}(T,\nu) | A(1) = a_1, \ldots, A((k-1)nT-1) = a_{(k-1)nT-1}\right) < \epsilon \tag{20}$$

for all k, j and m, and n large enough.

By [31, Theorem 3.1.2], since A_i is a finite-state irreducible discrete-time Markov chain for every i, the empirical mean $A_j^{i,n,k}/n^{\frac{1}{4}}$ obeys a large deviations principle with a convex, good rate function. This means that $A_j^{i,n,k}/n^{\frac{1}{4}} \to \lambda_j$ in probability for every j at a uniformly exponential rate. There are only a polynomial $(Tn^{\frac{3}{4}})$ number of subintervals in every interval of size nT, hence (19) follows. (20) is obtained in a similar manner since by Cramér's Theorem [31], the empirical mean of i.i.d. random variables obeys a large deviations principle with a convex, good rate function. ■

The lemma basically lets us assume that the empirical measures of the channel service rate and arrival processes look like their true measures, with very high probability.

Let $Q_S(k,t)$ denote the smallest queue length in the three queue lengths $Q_i(k,t)$, $i = 1, 2, 3$, at the beginning of the tth subinterval in the kth interval of time. The following lemma is crucial to the proof of instability and describes how the queueing system behaves in a typical interval.

Lemma 9: Fix $T > 0$. There exists $n_1 \in \mathbb{Z}^+$, $\nu > 0$ and $r > 0$ such that for any $k \in \mathbb{Z}^+$ and $n > n_1$, conditioned on the fact that the event $F^{n,k}(T,\nu)$ has not occurred, the following happens. If

$$\min\{Q_1(k,l), Q_2(k,l), Q_3(k,l)\} > n$$

for some $l \in \{1, \ldots, P_T^n\}$, then

$$\frac{Q_S(k, l+l') - Q_S(k,l)}{l'} \ge r \quad \forall l' \in \{1 \cdot \ldots, P_T^n - l\}.$$

This lemma essentially tells us that in any typical interval, the lowest of the three queues strictly increases with a uniform minimum rate, provided that the queues are sufficiently large to start with.

Proof: Recall that a subinterval consists of $n^{\frac{1}{4}}$ time slots. We can choose $\nu > 0$ to be much smaller than all the λ_i. If we denote by $\Delta(< \infty)$ the maximum possible channel service rate in the system and by Λ the maximum of the three arrival rates $\lambda_1, \lambda_2, \lambda_3$, the change in any queue within an interval is at most $\Gamma n^{\frac{1}{4}}$ where $\Gamma = |\Lambda + \nu - \Delta|$. We can pick $\delta > 0$ sufficiently small, and $n \in \mathbb{Z}^+$ sufficiently large such that this quantity is negligible compared to a queue length error of δn and such that the second step of the Max-Sum-Queue policy is immune to queue length errors of up to δn if all queues are at least of length n.

Let $S^{\delta n}(k,l)$ denote the subset of queues whose lengths are within δn of the smallest queue length $Q_L(k,l)$, at the beginning of the lth subinterval in the kth interval. Without loss of generality we can assume that $S^{\delta n}(k,\cdot)$ does not change from l to l', since if otherwise we can partition the set of intervals from l to l' into contiguous subsets with this property and obtain the result using the uniform bound r.

The proof proceeds by considering various cases for the number of queues in $S^{\delta n}(k,l)$.

Case 1—$S^{\delta n}(k,l)$ contains exactly one element: In this case, the single element is the unique smallest queue in the system, and remains the unique smallest queue throughout, until subinterval l'. Hence the only subset picked by the first step of the Max-Sum-Queue policy is that consisting of the other two queues. Consequently this queue is never served at all, and must increase at a rate at least $(\min\{\lambda_1, \lambda_2, \lambda_3\} - \nu) > 0$.

Case 2—$S^{\delta n}(k,l)$ contains exactly two elements: Here we need to consider three further subcases:

2(a)—$S^{\delta n}(k,l) = \{1,2\}$: In such a situation only the subset $\{2,3\}$ or $\{1,3\}$ can be picked by the first step of Max-Sum-Queue. Let us bound from above the maximum rate at which Q_1 and Q_2 together are served. If we assume (in the best case) that Q_3 is never picked for service in the second step of Max-Sum-Queue, then Q_1 and Q_2 share the service time, and the total service rate to them is at most $200(\frac{1}{2}+\nu)+150(\frac{1}{2}-\nu) = 150+50\nu$. The service discipline reduces to serving the longest of Q_1 and Q_2.

We claim that in this time, the difference $|Q_1 - Q_2|$ cannot exceed a constant amount, say 10Γ, since if it did, then there was a last previous time when the order of the queues was the same and the difference was under 5Γ. This implies that the difference grew under the longest-queue policy, a contradiction.

The total arrival rate to Q_1 and Q_2, however, is at least $2(87-\nu) = 174 - 2\nu$, hence $Q_1 + Q_2$ increases with a net rate of at least $(174-2\nu)-(150+50\nu) = 24-52\nu$. Hence their average increases with a net rate of at least $(24-52\nu)/2 = 12-26\nu$, and since Q_1 and Q_2 remain within 10Γ of each other throughout and n can be chosen large enough, the lowest queue increases with rate at least (arbitrarily close to) $12 - 26\nu > 0$ for small enough $\nu > 0$.

2(b)—$S^{\delta n}(k,l) = \{1,3\}$: Here, the only possible subsets which can be picked are $\{1,2\}$ and $\{2,3\}$. Note that Q_3 can never be served when the first subset is picked. If the second subset is picked, since $Q_2 > Q_1$ always, the only state in which Q_1 is served is when its rate is 200 and Q_2's rate is 100. Hence Q_3 increases with a rate at least $0.01 - \nu$, while Q_1 increases with a rate at least $(87-\nu) - (\frac{1}{4}+\nu)200 = 37 - 201\nu > 0$ for small enough $\nu > 0$.

2(c)—$S^{\delta n}(k,l) = \{2,3\}$: Similar to the previous case by symmetry.

Case 3—$S^{\delta n}(k,l)$ contains exactly three elements: In this case, all three subsets are capable of being chosen in the first scheduling step. Q_3 is never served, hence its length increases at rate at least $0.01 - \nu > 0$. Partition the total time into sections where: i) Q_3 is the smallest queue; ii) Q_3 is between the other two queues; and iii) Q_3 is the largest queue. For i) , the smallest queue Q_3 clearly increases with rate bounded away from zero. For ii), only the subset consisting of the top two queues is picked, hence the smallest queue increases with rate at least $87-\nu > 0$. For iii), we can use the same argument as with case *2(a)* to see that the smallest queue increases with a strictly positive rate. This completes the proof. ∎

Proof of Theorem 6: Fix any $T > 0$. Lemma 9 gives us $n_1, \nu > 0$ and $r > 0$ such that in a typical interval (where the interval size is nT subintervals with $n > n_1$), if all queues are greater than n, then the lowest queue always increases with a rate at least r. Let $R > 0$ be the maximum possible service rate in the system (in this case $R = 200$), and let $K = \lceil \frac{R}{r} \rceil + 1$ where $\lceil x \rceil$ denotes the smallest integer at least x. Choose $\epsilon \in (0,1)$ small enough so that $\epsilon 2^K < 1$, and furthermore, so that

$$\frac{(K-1)\epsilon 2^K}{(1-\epsilon 2^K)^2} + \frac{2\epsilon 2^K}{1-\epsilon 2^K} < 1. \tag{21}$$

Lemma 8 now gives us n_0 such that for $n > n_0$ and any k,

$$\mathbb{P}(F^{n,k}(T,\nu) | A(1) = a_1, \ldots, A((k-1)nT - 1) \tag{22}$$
$$= a_{(k-1)nT-1}) < \epsilon \tag{23}$$

uniformly over all $a_1, \ldots, a_{(k-1)nT-1}$. Fix n to be any integer greater than n_0 and n_1.

Using $Q_S(k,t)$ introduced earlier to mean the length of the smallest queue at the beginning of the tth time subinterval in the kth time interval, let $X = (X_s : s = 1,2,3,\ldots)$ be the random process denoting the size of the smallest queue at the beginning of every interval of time: $X_s = Q_S(s,1)\, \forall s = 1,2,3,\ldots$ Let m be an integer such that $m > n$ and let all queues start with initial state m: $Q_1(1,1) = Q_2(1,1) = Q_3(1,1) = m$. Define the (time-valued) random variable τ_m to be the first time after starting that X drops below m: $\tau_m \triangleq \min\{s > 1 : X_s \le m\}$. We show that $\mathbb{P}(\tau_m < \infty) < 1$, implying that the smallest queue (and hence every queue) grows without bound with a nonzero probability and establishing transience of the Markov chain describing the evolution of the system state $S(t)$. We can write

$$\mathbb{P}(\tau_m < \infty) = \sum_{l=0}^{\infty} \mathbb{P}(lK \le \tau_m < (l+1)K). \tag{24}$$

Let B_l be the random variable which counts the number of atypical intervals of time up to interval lK

$$B_l \triangleq \sum_{s=1}^{lK} \chi_{F^{n,s}(T,\nu)}.$$

We claim that $lK \le \tau_m < (l+1)K$ implies $B_{l+1} \ge l$, for otherwise $B_{l+1} < l$, and by Lemma 9, for $lK \le t < (l+1)K$, we have

$$\begin{aligned} X_t &\ge m + [t - B_{l+1}]rnT - B_{l+1}RnT \\ &\ge m + [lK - B_{l+1}]rnT - B_{l+1}RnT \\ &= m + nT[lKr - B_{l+1}(r+R)] \\ &> m + lnT[(K-1)r - R] \\ &\ge m \end{aligned}$$

a contradiction to $lK \le \tau_m < (l+1)K$.

From (24), we can write

$$\begin{aligned} \mathbb{P}(\tau_m < \infty) &\le \sum_{l=1}^{\infty} \mathbb{P}(B_l \ge l-1) \\ &= \sum_{l=1}^{\infty} \sum_{l'=l-1}^{lK} \mathbb{P}(B_l = l') \\ &\overset{(a)}{\le} \sum_{l=1}^{\infty} \sum_{l'=l-1}^{lK} \binom{lK}{l'} \epsilon^{l'} \end{aligned}$$

$$
\begin{aligned}
&\le \sum_{l=1}^{\infty} \sum_{l'=l-1}^{lK} 2^{lK} \epsilon^{l} \\
&= \sum_{l=1}^{\infty} ((K-1)l+2)(\epsilon 2^{K})^{l} \\
&= (K-1)\sum_{l=1}^{\infty} l(\epsilon 2^{K})^{l} + 2\sum_{l=1}^{\infty}(\epsilon 2^{K})^{l} \\
&= \frac{(K-1)\epsilon 2^{K}}{(1-\epsilon 2^{K})^{2}} + \frac{2\epsilon 2^{K}}{1-\epsilon 2^{K}} \\
&\overset{(b)}{<} 1.
\end{aligned}
$$

Here (a) is by applying (22) to

$$
\begin{aligned}
\mathbb{P}(B_l = l') &= \mathbb{P}\left(\sum_{s=1}^{lK} \chi_{F^{n,s}(T,\nu)} = l'\right) \\
&= \sum_{1\le i_1 < \cdots < i_{l'} \le lK} \mathbb{P}(\chi_{F^{n,i_1}(T,\nu)} = 1, \ldots, \chi_{F^{n,i_{l'}}(T,\nu)} = 1, \\
&\qquad\qquad \chi_{F^{n,i}(T,\nu)} = 0, i \notin \{i_1, \ldots, i_{l'}\}) \\
&\le \sum_{1\le i_1 < \cdots < i_{l'} \le lK} \epsilon^{l'} \\
&= \binom{lK}{l'} \epsilon^{l'}
\end{aligned}
$$

and (b) follows from (21). This completes the proof. ■

ACKNOWLEDGMENT

The authors are grateful to the anonymous reviewers for their suggestions and feedback towards improving the presentation of the paper.

REFERENCES

[1] A. Gopalan, C. Caramanis, and S. Shakkottai, "On wireless scheduling with partial channel state information," in *Proc. Allerton Conf. Commun., Contr. Comput.*, 2007.

[2] M. Andrews, K. Kumaran, K. Ramanan, A. L. Stolyar, R. Vijayakumar, and P. Whiting, "Scheduling in a queueing system with asynchronously varying service rates," *Probabil. Eng. Inform. Sci.*, vol. 14, pp. 191–217, 2004.

[3] R. Berry and R. Gallager, "Communication over fading channels with delay constraints," *IEEE Trans. Inf. Theory*, vol. 48, pp. 1135–1149, May 2002.

[4] X. Liu, E. K. P. Chong, and N. B. Shroff, "A framework for opportunistic scheduling in wireless networks," *Comput. Netw. J.*, vol. 41, no. 4, 2003.

[5] J. G. Andrews, A. Ghosh, and R. Muhamed, *Fundamentals of WiMAX: Understanding Broadband Wireless Networking*. Englewood Cliffs, NJ: Prentice-Hall, 2007.

[6] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Trans. Inf. Theory*, vol. 39, pp. 466–478, Mar. 1993.

[7] X. Liu, E. P. K. Chong, and N. Shroff, "Opportunistic transmission scheduling with resource-sharing constraints in wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 19, no. 10, pp. 2053–2064, Oct. 2001.

[8] S. Shakkottai, R. Srikant, and A. L. Stolyar, "Pathwise optimality of the exponential scheduling rule for wireless channels," *Adv. Appl. Probabil.*, vol. 36, no. 4, pp. 1021–1045, 2004.

[9] X. Lin, N. Shroff, and R. Srikant, "A tutorial on cross-layer optimization in wireless networks," *IEEE J. Sel. Areas Commun.*, pp. 1452–1463, Aug. 2006.

[10] A. Eryilmaz, R. Srikant, and J. R. Perkins, E. K. P. Chong, Ed., "Throughput optimal scheduling for broadcast channels," in *Proc. SPIE Modeling and Design of Wireless Netw.* , Denver, CO, 2001, vol. 4531, pp. 70–78, (editor).

[11] M. Neely, E. Modiano, and C. Rohrs, "Power allocation and routing in multi-beam satellites with time varying channels," *IEEE/ACM Trans. Netw.*, vol. 11, no. 1, pp. 138–152, Feb. 2003.

[12] M. Chiang, "Balancing transport and physical layers in wireless multihop networks: Jointly optimal congestion control and power control," *IEEE J. Sel. Areas Commun.*, vol. 23, pp. 104–116, Jan. 2005.

[13] A. Stolyar, "Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic," *Ann. Appl. Probabil.*, vol. 14, no. 1, pp. 1–53, 2004.

[14] S. Guha, K. Munagala, and S. Sarkar, Performance Guarantees Through Partial Information Based Control in Multichannel Wireless Networks Univ. Penn., Tech. Rep., 2006 [Online]. Available: http://www.seas.upenn.edu/~swati/report.pdf

[15] Z. Ji, Y. Yang, J. Zhou, M. Takai, and R. Bagrodia, "Exploiting medium access diversity in rate adaptive wireless LANs," in *Proc. ACM MOBICOM*, 2004.

[16] A. Sabharwal, A. Khoshnevis, and E. Knightly, "Opportunistic spectral usage: Bounds and a multi-band CSMA/CA protocol," *IEEE/ACM Trans. Netw.*, 2006.

[17] N. Chang and M. Liu, "Optimal channel probing and transmission scheduling for opportunistic spectrum access," in *Proc. ACM Int. Conf. Mobile Comput. Netw. (MobiCom)*, Montreal, Canada, Sep. 2007.

[18] P. Chaporkar, A. Proutiere, H. Asnani, and A. Karandikar, "Scheduling with limited information in wireless systems," in *Proc. 10th ACM Int. Symp. Mobile Ad-Hoc Netw. Comput.*, 2009, pp. 75–84.

[19] D. Gesbert and M. Slim-Alouini, "How much feedback is multi-user diversity really worth?," in *Proc. Int. Conf. Commun.*, 2004, pp. 234–238.

[20] X. Qin and R. Berry, "Opportunistic splitting algorithms for wireless networks," in *Proc. IEEE INFOCOM*, Mar. 2004.

[21] X. Qin and R. Berry, "Opportunistic splitting algorithms for wireless networks with heterogenous users," in *Proc. Conf. Inf. Sci. Syst. (CISS)*, Mar. 2004.

[22] T. Tang and R. W. Heath, "Opportunistic feedback for downlink multiuser diversity," *IEEE Commun. Lett.*, vol. 9, pp. 948–950, 2005.

[23] S. Patil and G. de Veciana, "Reducing feedback for opportunistic scheduling in wireless systems," *IEEE Trans. Wireless Commun.*, vol. 12, no. 6, pp. 4227–4238, Dec. 2007.

[24] M. Neely, "Max weight learning algorithms with application to scheduling in unknown environments," Feb. 2009 [Online]. Available: http://arxiv.org/abs/0902.0630

[25] A. Eryilmaz, R. Srikant, and J. R. Perkins, "Stable scheduling policies for fading wireless channels," *IEEE/ACM Trans. Netw.*, vol. 13, pp. 411–424, 2005.

[26] V. A. Malyshev and M. V. Menshikov, "Ergodicity, continuity and analyticity of countable Markov chains," *Trans. Moscow Math. Soc.*, vol. 39, pp. 3–48, 1979.

[27] L. Tassiulas, "Adaptive back-pressure congestion control based on local information," *IEEE Trans. Autom. Control*, vol. 40, pp. 236–250, Feb. 1995.

[28] K. J. Arrow, E. W. Barankin, and D. Blackwell, "Admissible points of convex sets," in *Contributions to the Theory of Games*, ser. Ann. Math. Studies. Princeton, NJ: Princeton Univ. Press, 1953, vol. 2, pp. 87–91 [Online]. Available: http://www.ams.org/mathscinet-getitem?mr=54919, No. 28, [Online]. Available

[29] P. E. Gill and W. Murray, *Numerical Methods for Constrained Optimization*. London, U.K.: Academic , 1974.

[30] S. Shakkottai and A. Stolyar, , F. Karpelevich and Y. M. Suhov, Eds., "Scheduling for multiple flows sharing a time-varying channel: The exponential rule," in *Amer. Math. Soc. Transl.*, ser. 2. : , 2002, vol. 207, a volume in memory of Editor.

[31] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*. New York: Jones and Bartlett, 1993.

Aditya Gopalan received the B.Tech. and M.Tech. degrees in electrical engineering from the Indian Institute of Technology (IIT) Madras, Chennai, in 2006.

He was a summer intern at the Corporate Research and Development Center, Qualcomm Inc., in 2009. He is currently working toward the Ph.D. degree in electrical engineering at The University of Texas at Austin. His current research interests include network algorithms, stochastic control, and machine learning.

Constantine Caramanis (M'06) received the Ph.D. degree in EECS from the Massachusetts Institute of Technology, Cambridge, in 2006.

Since then, he has been on the faculty of Electrical and Computer Engineering, The University of Texas at Austin. He received the NSF CAREER

award in 2011. His current research interests include robust and adaptable optimization, machine learning and high-dimensional statistics, with applications to large scale networks.

Sanjay Shakkottai (M'02–SM'11) received the Ph.D. degree from the ECE Department, University of Illinois at Urbana-Champaign, in 2002.

He is with The University of Texas at Austin, where he is currently an Associate Professor and the Engineering Foundation Centennial Teaching Fellow with the Department of Electrical and Computer Engineering. His current research interests include network architectures, algorithms, and performance analysis for wireless and sensor networks.

Dr. Shakkottai received the NSF CAREER award in 2004.