Wireless Scheduling with Heterogeneously Delayed Network-State Information

(Invited Paper)

Akula Aneesh Reddy, Siddhartha Banerjee, Aditya Gopalan, Sanjay Shakkottai
The University of Texas at Austin
Austin, Texas 78712
Email: {aneesh, sbanerjee, gadit, shakkott}@mail.utexas.edu

Lei Ying
Iowa State University
Ames, Iowa 50011
Email: leiying@iastate.edu

***Abstract*—We study the problem of distributed scheduling in wireless networks, where each node makes individual scheduling decisions based on heterogeneously delayed network state information (NSI). This leads to inconsistency in the views of the network across nodes, which, coupled with interference, makes it challenging to schedule for high throughputs.**

We characterize the network throughput region for this setup, and develop optimal scheduling policies to achieve the same. Our scheduling policies have a threshold-based structure and, moreover, require the nodes to use only the "smallest critical subset" of the available delayed NSI to make decisions. In addition, using Markov Chain mixing techniques, we quantify the impact of delayed NSI on the throughput region. This not only highlights the value of extra NSI for scheduling, but also characterizes the loss in throughput incurred by lower complexity scheduling policies which use homogeneously delayed NSI.

I. Introduction

An increasing proportion of modern-day data networks is deployed wirelessly. Examples of such implementations include cellular networks for mobile users, citywide mesh networks and mobile ad-hoc networks (MANETs). Managing data in wireless networks, as opposed to traditional wireline networks, is complicated by two effects unique to wireless systems – channel fading and interference. Good scheduling/routing algorithms are crucial for adapting to these effects in order to improve system throughput.

There is a large body of work on wireless scheduling with instantaneous network state information (NSI: the channel and queue state information of the whole network), primarily centered around the celebrated Backpressure/Max-Weight scheduling algorithms [1], [2], [3], [4], [5], [6], [7], [8]. However, in practice, instantaneous queue and channel-state information in a wireless network could be difficult to obtain, due to *(i)* prohibitive overheads in measurement and feedback of state information and/or *(ii)* channel quality fluctuations occurring faster than the time required to convey them to the scheduler, leading to delayed channel-state feedback. Hence, much recent work has treated the case when the available NSI is partial, delayed or degraded. This includes centralized approaches to scheduling with homogeneously delayed NSI [9], heterogeneously delayed NSI/topology undertainty [10], [11], instantaneous NSI from subsets of users [12], joint channel sampling/scheduling strategies [13], and distributed scheduling approaches such as Queue-Back-Pressure Random Access in a non-fading setting [14], channel-aware ALOHA [15], and scheduling with local instantaneous NSI coupled with homogeneous, delayed global NSI [10].

In this work, we consider the problem of distributed wireless scheduling in the presence of arbitrary interference set constraints and Markovian channel fading, where each transmitter knows other transmitters' NSI with possibly *heterogeneous* delays. This disparity in the delays of NSI available to the transmitters can potentially result in inconsistent views of the global current network state, causing conflicting/poor local scheduling decisions among the transmitters. Given such a NSI structure, how can all the transmitters in the network use their (possibly inconsistent) individual information to make scheduling decisions for good overall throughout? Our main contributions in this regard are –

1) We characterize the network throughput region with each transmitter having instantaneous local (individual) NSI and heterogeneous delayed NSI from other transmitters. It turns out that the network throughput region can be parameterized by a special class of queue-length-oblivious scheduling policies which we term *static-service-split* (SSS) policies. These SSS policies are distinguished in that they need each transmitter to use only the *smallest critical* (in a suitable sense) delayed CSI received from other nodes for the purpose of achieving maximal network throughput.
2) We develop decentralized, threshold-based, throughput-optimal scheduling policies for the network that can support all data rates within the network throughput region. In these throughput-optimal policies, each transmitter uses a combination of (delayed) network queue length information and critical delayed channel state information from other transmitters to compute a local threshold, and compares its own channel state with the threshold to make its scheduling decision.
3) With respect to the canonical heterogeneous NSI setting, we quantify the loss (gain) in throughput that results from all transmitters having the maximum (minimum) possible homogeneously delayed NSI from other transmitters. This is accomplished using techniques from

mixing of Markov Chains.

II. Scheduling with Heterogeneously Delayed NSI: An Example

Let us consider an illustrative example to help understand the essential difficulties and challenges in scheduling when the NSI available to each user is delayed in a heterogeneous fashion. Suppose we have three wireless users A, B and C who attempt to send packet-based queued data to a receiver in a time-slotted manner – for instance, this could model three users on a cellular uplink. We assume that the users are located sufficiently close to each other so as to make their transmissions interfere, i.e., if the number of users attempting to transmit in a time slot is more than one, no packets reach the receiver. The channel between each user and the receiver is time-varying, and in the event of a successful transmission as above, the channel state/rate of the single attempting user specifies how many packets can be sent to the receiver in that time slot.

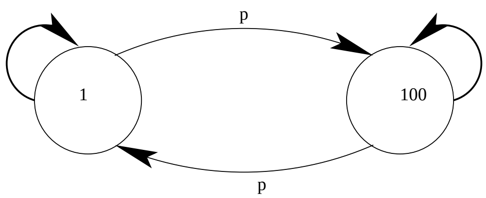

Fig. 1. User C channel Markov chain

Each user possesses instantaneous channel (and queue backlog) state information about its own channel and receives delayed channel (and queue backlog) state information from other users for the purpose of making transmit/no-transmit decisions. Let us assume for simplicity that the channels for users A and B take rates 1 or 100 (packets per time slot) each with probability $\frac{1}{2}$ *independently* in each time slot; however user C's channel state evolves as a *Markov Chain* between rates 1 and 100 with crossover probability $p = \frac{1}{4}$ (Fig. 1). User A gets channel state information from users B and C delayed by 1 time slot, user B gets channel state information from users A and C delayed by 1 and 2 time slots respectively, and user C gets channel state information from users A and B delayed by 1 time slot. Fig. 2 depicts this NSI structure at time t – a circle in the row of Tx A at time $t-1$ indicates that it is the latest information B has about A's channel state, and so on.

Note that due to this information structure, at each time users A and B have different "views" of user C's current channel state owing to disparate channel state information delays. For instance, if user C's channel two time slots ago was at rate 100 and one time slot ago was at rate 1, user A is led to believe that user C's current channel is very likely to have rate 1, whereas user B's belief would be that user C's channel is most probably at rate 100. In such events, how must the users act so that they can avoid excessive collision and achieve desired data transmission rates? It turns out, as we

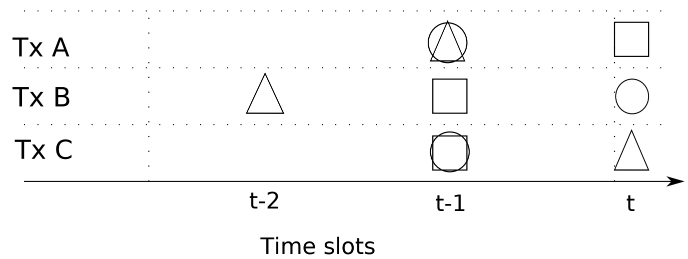

Fig. 2. Heterogeneous NSI for the 3-user network: Squares, circles and triangles represent the *most recent* channel state information available to user A, B and C respectively.

show later on, that the following *threshold*-based transmission rule, for each user, is a throughput-*optimal* scheduling strategy: In every time slot,

1) All the three users compute individual "threshold values" (to be used later) as functions of their respective delayed queue length information and certain "critical" subsets of their available delayed channel state information – user A works out a threshold value as a function of the one-step delayed channel states of user B and user C, and so on.
2) Each user looks at the value(s) of its critical set of delayed NSI, compares the corresponding threshold value and its own current channel state, and *attempts transmission only if its current channel state exceeds the threshold*.

Now, consider the case when *both* user A and user B have user C's channel state information with a delay of 2 time slots. Compared to the earlier set of delays, user A has one step "coarser" channel state information about user C, so we expect a degradation in the overall set of achievable data rates that all the users can support. As a matter of fact,

1) *The best average sum rate achievable in the latter system is* 56.69 *packets/time slot*, whereas
2) *The best average sum rate achievable in the former system is* 62.88 *packets/time slot* – an increase of about 11% in the sum rate with one additional step of channel state information.

An important question, thus, is – how can we find the throughput region for the system under a given structure of information delays?

In general, we provide a theory that answers the following useful questions:

- What are all the long-term average rates (throughput region) that such a wireless system with an arbitrary delayed NSI structure can support?
- How can each user make scheduling (transmission) decisions – just based on its limited amount of delayed information about other users' channel states – to be able to support *any* given feasible data rate? Moreover, which are the time slots whose channel state information is "crucial" or "essential" for making throughput-optimal decisions?

- By how much does the rate region change with better or worse delayed channel state information?

III. System Model

In this section, we first describe the network model, traffic model and interference model used in this paper. We describe the heterogeneous NSI structure. We then introduce the notation used in this paper along with definitions of stability and throughput optimality.

Network Model: We consider a wireless network consisting of L transmitter-receiver pairs denoted by $\mathcal{L}$. We model the (time-varying) capacity of each link l using a discrete time Markov chain, denoted by $\{C_l[t]\}$, on the finite state space $\mathcal{C}$ $= \{c_1, c_2, c_3, ..., c_M\}$, where $c_1 \leq \cdots \leq c_M$. Furthermore, we require that the Markov chain representing any link's capacity is independent and identically distributed, with transition probabilities $P_{ij} := \Pr[C_l[t+1] = c_j | C_l[t] = c_i]$. The above channel model is assumed for notational simplicity and our results hold even for the case of networks where each link can be modeled by a separate Markov chain (different state space and different transition probabilities). The only condition for our results to hold is that channels are independent across various transmitter-receiver pairs (users).

We assume that the channel-state Markov chain parameterized by the transition probabilities $\{P_{ij}\}_{i,j}$ is irreducible and aperiodic. Thus the channel state process has a stationary distribution and we denote the stationary probability of being in a state $c_j, j \in \{1, 2, 3..., M\}$ by π_j.

Interference Model: We consider collision interference model. We denote by I_l, the set of links that interfere with link l. We say that the transmission on link l at time t is successful, if no link in the interference set transmits during the same time t and we can potentially transmit $C_l[t]$ packets.

Traffic Model: We assume single hop flows in the network, and that each node does not have multiple simultaneous connections. Each link in the network has a traffic denoted by $A_l[t]$, that describes the number of packets that arrive at sender node of link at time t. For every link l, we assume that $A_l[t]_t$ is a bounded stationary process with $\lambda_l := E[A_l[t]] < \infty$.

A. NSI Structure

We assume that each transmitter can have varying levels of network state information. At time t, transmitter l has channel and queue state information history of link l till time t but has only delayed channel state information and queuing history of other links in the network. Let $\tau_l(h)$ denote the delay incurred in communicating the channel and queue state information of link h to transmitter node of link l. Thus each transmitter node l has a vector $\vec{\tau_l}$ that characterizes the available delayed NSI. We denote by τ_{min} and τ_{max} the minimum and maximum channel (and queue) state information delay across the network, i.e.,

$$\tau_{min} = \min_{l,h \in \mathcal{L}: l \neq h} \tau_l(h); \tau_{max} = \max_{l,h \in \mathcal{L}: l \neq h} \tau_l(h).$$

We denote the set $\{C_l[t-\tau], C_l[t-\tau+1], ..., C_l[t]\}$ by $C_l[t](0:\tau)$ and the set $\{C_l[t]\}_{l \in \mathcal{L}}$ by $\mathbf{C}[t]$. We denote the information available at transmitter l by $\{\mathcal{P}_l(\mathbf{C}[t](0 : \tau_{max})), \mathcal{P}_l(\mathbf{Q}[t](0 : \tau_{max}))\}$, where

$$\mathcal{P}_l(\mathbf{C}[t](0:\tau_{max})) := \{\vec{\mathcal{P}}_{lm}(\mathbf{C}[t](0:\tau_{max}))\}_{m \in \mathcal{L}},$$

$$\vec{\mathcal{P}}_{lm}(\mathbf{C}[t](0:\tau_{max})) := \{C_m[t-\tau]\}_{\tau=\tau_{max}}^{\tau_l(m)}.$$

B. Stability

Given the arrival traffic rate $\{\lambda_l\}_{l \in \mathcal{L}}$ and scheduling policy, we say that the network is stochastically stable if the mean of the sum of queue lengths is bounded. We say that an arrival rate vector $\{\lambda_l\}_{l \in \mathcal{L}}$ is supportable if there exists any scheduling policy that can make the network stochastically stable.

IV. Distributed Scheduling with Heterogeneous delayed NSI

In this section, we first characterize the throughput region with given delayed information structure and distributed nature of scheduling. We then identify the "critical set" of available NSI for every user, and show that this critical set NSI is sufficient to make throughput-optimal scheduling decisions – the Markovian nature of the wireless channels facilitates tracking only the most recent channel states between users, which is suitably formalized in this section.

A. Throughput Characterization

Let us define a set of functions $\{f_l\}_{l \in \mathcal{L}}$ defined from $\mathcal{P}_l(\mathbf{C}[t](0:\tau_{max}))$ on to $\{0,1\}$. The function f_l essentially is a scheduling policy, independent of queue states of system and of all the channel state information past τ_{max}, as follows – at each time t, every link l computes the binary value $f_l(\mathcal{P}_l(\mathbf{C}[t](0:\tau_{max})))$ and attempts to transmit (i.e., schedule itself) whenever this value is 1. These functions are similar in spirit to stationary scheduling policies (or SSS rules) used in [1], [4], [10], [12].

Given the delayed channel state information $\mathbf{C}[t-\tau_{max}] =$ $\mathbf{c}$, we define $\mathbf{S}(\mathbf{c}, \mathbf{f}) = \{S_l(\mathbf{c}, \mathbf{f})\}_{l \in \mathcal{L}}$ as follows:

$$S_l(\mathbf{c}, \mathbf{f}) = E\Big[C_l[t] f_l(\mathcal{P}_l(.)) \prod_{m \in I_l} (1 - f_m(\mathcal{P}_m(.))) \;|\; \mathbf{C}[t-\tau_{max}] = \mathbf{c}\Big],$$

where $\mathcal{P}_l(.) = \mathcal{P}_l(\mathbf{C}[t](0:\tau_{max}))$. We now define $\eta(\mathbf{c})$ as follows,

$$\eta(\mathbf{c}) = \mathcal{CH}_f\big(S(\mathbf{c}, \mathbf{f})\big).$$

Therefore, $\eta(\mathbf{c})$ is the convex hull of all possible average data transmission rates achieved by stationary scheduling policies, given that the common NSI upto time $t - \tau_{max}$ is c. We now define the throughput region Λ as follows

$$\Lambda = \big\{\lambda : \lambda = \sum_{\mathbf{c} \in \mathcal{C}^L} \pi(\mathbf{c}) \eta(\mathbf{c})\big\}.$$

Lemma 4.1: Given the above described NSI structure, traffic $A[t]$ is supportable if and only if $(1+\epsilon)E[A[t]] \in \Lambda$ for some $\epsilon > 0$.

Proof: See the Appendix for proof. ■

B. Critical NSI

As defined in section-I, $\tau_l(h)$ represents the delay with which the latest queue state and channel state information of link h is available at link l. We expect that for link l at time slot t, all the latest delayed channel state information from other users (i.e., $\{C_k[t-\tau_k(l)] : k \in \mathcal{L}, k \neq l\}$) is the information most useful with regard to the current channel states of the other users. In what follows, we introduce the important concept of critical NSI for the network – essentially *all the latest delayed channel state information observed by every user in the network* – which is later used to develop a throughput-optimal scheduling policy in which each user makes scheduling decisions just based on the critical NSI available to itself.

Given $\mathbf{C}[t](0:\tau_{max})$, the critical set of information related to link l is defined as the the channel state information at times $\{t-\tau_k(l)\}_{k\in\mathcal{L}:k\neq l}$. Let us denote the critical NSI of the network at time t as $\mathcal{CS}(.)$ and can be expressed mathematically as follows

$$\mathcal{CS}(\mathbf{C}[t](0:\tau_{max})) := \{\{C_l[t-\tau_k(l)]\}_{k\in\mathcal{L}:k\neq l}\}_{l\in\mathcal{L}}.$$

For every $l \in \mathcal{L}$, we define the critical NSI available at transmitter l as follows:

$$\mathcal{CS}_l(\mathbf{C}[t](0:\tau_{max})) := \mathcal{CS}(\mathbf{C}[t](0:\tau_{max})) \bigcap \mathcal{P}_l(\mathbf{C}[t](0:\tau_{max})).$$

Going back to the example in Section-II, we have the $\tau_{max} = 2$, and critical set at time t is $\{C_A[t-1], C_B[t-1], C_C[t-1], C_C[t-2]\}$. Thus at time t, the critical set available at transmitter A is $\{C_A[t-1], C_B[t-1], C_C[t-1], C_C[t-2]\}$, at B is $\{C_A[t-1], C_B[t-1], C_C[t-2]\}$, and at C is $\{C_A[t-1], C_B[t-1], C_C[t-1], C_C[t-2]\}$.

We now describe the queue dynamics at each transmitter node. Each transmitter maintains a queue of packets corresponding to its destination. Once a packet is sent, this node does not flush the packets from its queues until an acknowledgment is received indicating successful reception. This acknowledgment (ack) is received with some delay, and this delay is consistent with the critical channel state information delays. By this, we mean that the information contained in the acknowledgment, either explicitly (in the header) or implicitly (via the observation that presence of the ack/nack "encodes" the interfering links' critical NSI) does not contain additional NSI as compared to the nodes' critical NSI. This is to ensure that by learning based on queue-lengths and acks, nodes cannot get more NSI that the critical NSI. This consistency of ack "state" information can be characterized explicitly where each transmitter node has potentially a different ack delay, which is "naturally" consistent with the critical NSI in the system. However, in this paper, for notational simplicity, we assume that the acknowledgment is received only after τ_{max} time slots (thus trivially ensuring that the ack information is consistent with the critical NSI). The queue dynamics therefore is represented as follows,

$$Q_l[t+1] = (Q_l[t] + A_l[t] - S_l[t-\tau_{max}])^+,$$

where $S_l[t]$ denotes the number of packets successfully transmitted at time t.

C. Threshold-based Scheduling Algorithm:

In this section, we present a threshold-based decentralized scheduling algorithm, which can stabilize the network for all arrival rates in the interior of the throughput region Λ.

The algorithm we propose consists of two steps. At each time slot,

Step 1: All the transmitters compute threshold functions based on common NSI available at all transmitters. These threshold functions, one for each transmitter, map the respective transmitter's critical NSI to a corresponding threshold value, and are computed by solving the following optimization problem:

$$\arg\max_{\mathbf{T}} \sum_{l\in\mathcal{L}} Q_l(t-\tau_{max}) R_{l,\tau_{max}}(\mathbf{T}),$$

where

$$R_{l,\tau}(\mathbf{T}) := E\big[C_l[t] 1_{C_l[t]\geq T_l(.)} \prod_{m\in I_l} 1_{C_m[t]<T_m(.)} |\mathbf{C}[t-\tau]\big],$$

and $T_l(.) := T_l(\mathcal{CS}_l(\mathbf{C}[t](0:\tau_{max})))$.

Step 2: Each transmitter observes its current critical NSI, evaluates its threshold function (found in Step 1) at this critical NSI, and attempts to transmit if and only if its current channel rate exceeds the threshold value, i.e., when

$$C_l[t] \geq T_l(\mathcal{CS}_l(\mathbf{C}[t](0:\tau_{max}))).$$

We can now state our first main result:

Theorem 4.2: The proposed algorithm is throughput optimal

Proof outline: We briefly describe the proof technique here – the detailed proof is provided in the Appendix. The following lemma shows that solving an optimization problem locally in each time slot results in (globally) throughput-optimal scheduling.

Lemma 4.3: Consider the optimization problem

$$\arg\max_{\mathbf{F}(.)} \sum_{l\in\mathcal{L}} Q_l(t-\tau_{max}) R_{l,\tau_{max}}(\mathbf{F}(.)), \quad (1)$$

where

$$R_{l,\tau}(\mathbf{F}(.)) := E\big[C_l[t] F_l(.) \prod_{m\in I_l} (1-F_m(.)) |\mathbf{C}[t-\tau]\big],$$

and $F_l(.) := F_l(\mathcal{P}_l(\mathbf{C}[t](0:\tau_{max}))) \in \{0,1\}$ for each $l \in \mathcal{L}$. If each transmitter l at time t is scheduled to transmit whenever the optimizing $F_l^*(\mathcal{P}_l(\mathbf{C}[t](0:\tau_{max}))) = 1$, then any $\vec{\lambda}$ that satisfies $(1+\epsilon)\vec{\lambda} \in \Lambda$ for $\epsilon > 0$ is supportable.

Next, we show that the optimizing solution (i.e., the functions $F_l^*(.)$ of the individual NSI for all $l \in \mathcal{L}$)

1) Satisfies a *threshold* property, i.e.,

$$F_l^*(\mathcal{P}_l(\mathbf{C}[t](0:\tau_{max}))) = 1_{C_l[t]\geq T_l^*(\mathcal{P}_l(\mathbf{C}[t](0:\tau_{max})))},$$

2) Depends *only on the critical set* of NSI for each $l \in \mathcal{L}$, i.e.,

$$T_l^*(\mathcal{P}_l(\mathbf{C}[t](0:\tau_{max}))) = T_l^*(\mathcal{CS}_l(\mathbf{C}[t](0:\tau_{max}))).$$

The proof is completed by noting that the proposed algorithm finds the best threshold-based scheduling decisions where each transmitter's thresholds are based only on its currently available NSI. ■

V. Impact of delayed NSI on throughput Region

In this section, we are interested in characterizing the throughput region of the network with heterogeneous delayed NSI vis-à-vis the throughput region with homogeneous delayed NSI. This characterization provides the trade-off between the throughput achieved by scheduling algorithms using delayed NSI and the computational/storage complexity required to implement the algorithms. Let us denote the throughput region with NSI delays $\{\vec{\tau_l}\}_{l\in\mathcal{L}}$ by Λ. For an integer $\tau \geq 0$, let Λ_τ denote the throughput region assuming that each link has its own instantaneous NSI and knows the NSI of other links in the network with a fixed delay of τ. We note that

$$\Lambda_{\tau_{max}} \subseteq \Lambda \subseteq \Lambda_{\tau_{min}}.$$

The following theorem – our second main result – quantifies the loss (gain) in the throughput region by using the minimum (maximum) homogeneous delayed NSI compared to the canonical heterogeneous case.

Theorem 5.1: For integers $\tau_1, \tau_2 \geq 0$, let

$$\alpha(\tau_1,\tau_2) := \frac{2Lk_o\beta(\tau_1,\tau_2)}{\sum_j c_j \min_i P_{ij}^{\tau_1}}, \quad (2)$$

where $k_o = (1+M|I|)(\sum c_i)$, $\beta(\tau_1,\tau_2) = \max|P_{ij}^{\tau_1} - P_{kj}^{\tau_2}|$ and $|I|$ denotes the maximum size of an interfering set of transmitters. Then,

$$(1-\alpha(.))\Lambda_{\tau_{min}} \subseteq \Lambda \subseteq (1-\alpha(.))^{-1}\Lambda_{\tau_{max}},$$

with $\alpha(.) = \alpha(\tau_{min}, \tau_{max})$.

Proof: We prove a more general result which implies the above theorem: Given τ_1 and τ_2 such that $\tau_1 \leq \tau_2$, we have $\Lambda_{\tau_2} \supseteq (1-\alpha(\tau_1,\tau_2))\Lambda_{\tau_1}$.

For a NSI structure where each transmitter knows its current information and delayed information (by τ_1) of other links in the network, we have a scheduling policy based on thresholds (from Theorem 4.2) that is throughput optimal. We will need the following useful lemma [6].

Lemma 5.2: At any time t, given the common NSI $(\mathbf{Q}[\mathbf{t}](\tau_1 : t), \mathbf{C}[\mathbf{t}](\tau_1 : t))$, let $\mathbf{T_1^*}$ be the optimal set of thresholds calculated using the proposed algorithm and $\mathbf{T_2}$ be set of thresholds computed using a scheduling policy S_ρ such that the following condition holds (for some $\rho \in [0,1]$):

$$\sum_{l\in\mathcal{L}} Q_l(t-\tau_2)R_{l,\tau_1}(\mathbf{T_2}) \geq (1-\rho)\sum_{l\in\mathcal{L}} Q_l(t-\tau_2)R_{l,\tau_1}(\mathbf{T_1^*}).$$

Then, the scheduling policy S_ρ can stabilize any arrival rate $\vec{\lambda} \in (1-\rho)\Lambda_{\tau_1}$.

Let $\mathbf{T_2^*}$ be the set of thresholds computed using the proposed algorithm with the "degraded" NSI $(\mathbf{Q}[\mathbf{t}](\tau_2 : t), \mathbf{C}[\mathbf{t}](\tau_2 : t))$. Thus, $\mathbf{T_2^*}$ need not be an optimal set of thresholds for scheduling with the "non-degraded" partial NSI $(\mathbf{Q}[\mathbf{t}](\tau_1 : t), \mathbf{C}[\mathbf{t}](\tau_1 : t))$. Also, the proposed algorithm which uses only degraded partial NSI (τ_2 instead of τ_1) can stabilize the system for all arrival rates $\vec{\lambda} \in \Lambda_{\tau_2}$. We can write

$$R_{l,\tau_1}(\mathbf{T_2^*}) = E\left[C_l[t]1_{C_l[t]\geq T^*_{2,l}}\prod_{m\in I_l} 1_{C_m[t]<T^*_{2,m}}|\mathbf{C}[t-\tau_1]\right].$$

Since the random variables $C_l[t]$ and $C_m[t]$ are independent, we can rewrite the above expression as

$$R_{l,\tau_1}(\mathbf{T_2^*}) = E\big[C_l[t]1_{C_l[t]\geq T^*_{2,l}}|C_l[t-\tau_1] \times \prod_{m\in I_l} E\big[1_{C_m[t]<T^*_{2,m}}|C_m[t-\tau_1].$$

Let P_{ij}^τ denote the τ-step transition probability of the channel state Markov chain from state c_i to state c_j. Rewriting the above expression in terms of P_{ij}^τ, we have

$$R_{l,\tau_1}(\mathbf{T_2^*}) = \left(\sum_{i=1}^{M} c_i P_{.i}^{\tau_1} 1_{c_i\geq T^*_{2,l}}\right)\prod_{m\in I_l}\left(\sum_{i=1}^{M} P_{.i}^{\tau_1}1_{c_m\geq T^*_{2,l}}\right). \quad (3)$$

We now state another lemma that bounds the difference between $R_{l,\tau_1}(\mathbf{T_2^*})$ and $R_{l,\tau_2}(\mathbf{T_2^*})$.

Lemma 5.3: $|R_{l,\tau_1}(\mathbf{T_2^*}) - R_{l,\tau_2}(\mathbf{T_2^*})| < k_o\beta(\tau_1,\tau_2)$.

Using Lemma 5.3, we have that

$$\sum_{l\in\mathcal{L}} Q_l(t-\tau_2)R_{l,\tau_1}(\mathbf{T_2^*}) \geq \sum_{l\in\mathcal{L}} Q_l(t-\tau_2) \times (R_{l,\tau_2}(\mathbf{T_2^*}) - k_o\beta(\tau_1,\tau_2)).$$

With the fact that $\mathbf{T_2^*}$ is an optimal set of thresholds for the proposed algorithm with NSI $(\mathbf{Q}[\mathbf{t}](\tau_2 : t), \mathbf{C}[\mathbf{t}](\tau_2 : t))$, we have

$$\sum_{l\in\mathcal{L}} Q_l(t-\tau_2)R_{l,\tau_1}(\mathbf{T_2^*}) \geq \sum_{l\in\mathcal{L}} Q_l(t-\tau_2) \times (R_{l,\tau_2}(\mathbf{T_1^*}) - k_o\beta(\tau_1,\tau_2)).$$

Employing Lemma 5.3 once again, we have

$$\begin{aligned}
&\sum_{l\in\mathcal{L}} Q_l(t-\tau_2)R_{l,\tau_1}(\mathbf{T_2^*}) \\
&\geq \sum_{l\in\mathcal{L}} Q_l(t-\tau_2)\big(R_{l,\tau_1}(\mathbf{T_1^*}) - 2k_o\beta(\tau_1,\tau_2)\big) \\
&\geq \sum_{l\in\mathcal{L}} Q_l(t-\tau_2)R_{l,\tau_1}(\mathbf{T_1^*}) - (LQ_{max})2k_o\beta(\tau_1,\tau_2) \\
&= \sum_{l\in\mathcal{L}} Q_l(t-\tau_2)R_{l,\tau_1}(\mathbf{T_1^*})\left(1 - \frac{(LQ_{max})2k_o\beta(\tau_1,\tau_2)}{\sum_{l\in\mathcal{L}} Q_l(t-\tau_2)R_{l,\tau_1}(\mathbf{T_1^*})}\right).
\end{aligned}$$

Using the bound

$$\sum_{l\in\mathcal{L}} Q_l(t-\tau_2)R_{l,\tau_1}(\mathbf{T_1^*}) \geq Q_{max}\sum_j c_j \min P_{ij}^{\tau_1},$$

we have

$$\sum_{l\in\mathcal{L}} Q_l(t-\tau_2)R_{l,\tau_1}(\mathbf{T_2^*})$$
$$\geq \sum_{l\in\mathcal{L}} Q_l(t-\tau_2)R_{l,\tau_1}(\mathbf{T_1^*})\left(1-\frac{2Lk_o\beta(\tau_1,\tau_2)}{\sum_j c_j \min P_{ij}^{\tau_1}}\right)$$
$$= (1-\alpha(\tau_1,\tau_2))\sum_{l\in\mathcal{L}} Q_l(t-\tau_2)R_{l,\tau_1}(\mathbf{T_1^*}).$$

Using Lemma 5.2 now yields $\Lambda_{\tau_2} \supseteq (1-\alpha(\tau_1,\tau_2))\Lambda_{\tau_1}$ as desired. ■

Next, we analyze the behavior of $\alpha(\cdot,\cdot)$ and provide corollaries that characterize the throughput region Λ_∞ in terms of Λ_τ. For the sake of simplicity, we assume $P_{ij} > 0$ for all i and j. Even if P_{ij} are not all positive, we can find an integer m_o (since the Markov chain is aperiodic, irreducible and finite) such that $P_{ij}^{m_o} > 0$ for all i and j.

Corollary 5.4:

$$a)\quad \alpha(\tau_{min},\tau_{max}) \leq \frac{4Lk_o(1-M\delta)^{\tau_{min}}}{\sum_j c_j \min_i P_{ij}^{\tau_{min}}},$$

$$b)\lim_{\tau_{max}\to\infty}\alpha(\tau_{min},\tau_{max}) \leq \frac{2Lk_o(1-M\delta)^{\tau_{min}}}{\sum_j c_j \min_i P_{ij}^{\tau_{min}}},$$

where $\delta = \min_{ij} P_{ij}$.

Proof: The proof is based on the exponential convergence property [16] of finite state Markov chains and detailed proof is presented in Appendix. ■

VI. Conclusion

In this paper, we have addressed the problem of distributed scheduling in wireless networks with Markovian channels and heterogeneously delayed NSI. We have proposed a threshold-type distributed scheduling algorithm that is provably throughput optimal. We have shown that thresholds depend only up on the critical set of NSI. We have also characterized the effect of delayed NSI on the network throughput region.

Appendix

Proof: (Lemma 4.1) Firstly, if the arrival rates belong to the throughput region (i.e., $E[\vec{A}[t]] \in \Lambda$), then from the definition of throughput region we can construct a set of channel state dependent policies (i.e., f_l's) and "time-share" over those policies to stabilize the network.

Now for the other direction, given $\vec{A}[t]$ is supportable we have that there exists a scheduling algorithm $\mathcal{F}$ which can make the network stable. Let us define the state of the network at time t as $\mathbf{Y}[t] = \{Q_l[t](0:\tau_{max}), C_l[t](0:\tau_{max})\}_{l\in\mathcal{L}}$ and denote the network state under policy $\mathcal{F}$ by $\mathbf{Y}^{\mathcal{F}}[t]$. Since the arrival rate vector is supportable under the policy $\mathcal{F}$ and the fact that $\mathbf{Y}^{\mathcal{F}}[t]$ is Markovian, we have the result [16] that the Markov chain $\mathbf{Y}^{\mathcal{F}}[t]$ to be positive recurrent. Therefore $\mathbf{Y}[t]$ exhibits a steady stationary distribution. Let us denote the scheduling decision under policy $\mathcal{F}$ as $S^{\mathcal{F}}(\mathbf{Y}[t])$ We will now construct a time-sharing scheduling policy $\mathcal{F}_s$ that depends on the steady state distribution of queue lengths and channel states (denoted as $\pi(\mathbf{y}), y = \{\mathbf{q}(0:\tau_{max}), \mathbf{c}(0:\tau_{max})\}$) under policy $\mathcal{F}$. Let $r(\mathbf{y}) = \Pr(\mathbf{q}|\mathbf{c})$, computed using $\pi(\mathbf{y})$.

At each time, when delayed channel state information $\mathbf{C}[t](0:\tau_{max}) = \mathbf{c}$, the policy $\mathcal{F}_s$ probabilistically selects the scheduling decision $S^{\mathcal{F}}(\mathbf{q},\mathbf{c})$ with probability $r(\mathbf{y}=(\mathbf{q},\mathbf{c}))$. We observe that the time-sharing policy $\mathcal{F}_s$ allocates the same amount of service to each link as $\mathcal{F}$. Since $\vec{A}[t]$ can be supported by the time sharing policy, we have that $E[\vec{A}[t]] \in \Lambda$. ■

Proof: (Theorem 4.2) We first show the following threshold property for the optimal solution to the optimization problem defined in equation (1),

$$F_l^*(\mathcal{P}_l(\mathbf{C}[t](0:\tau_{max}))) = 1_{C_l[t]\geq T_l^*(\mathcal{P}_l(\mathbf{C}[t](0:\tau_{max})))},$$

Let us assume that we partly know the optimal solution. In particular, we assume that we are given the entire $\{F_l^*(\mathcal{P}_l(\mathbf{C}[t](0:\tau_{max})))\}_{l\in\mathcal{L}}$ except $F_k^*(\mathcal{P}_k(\mathbf{C}[t](0:\tau_{max})))$ at two different values of NSI ($\mathcal{P}_k(\mathbf{C}[t](0:\tau_{max})) = \{(C_k[t]=c_i,\vec{r}),(C_k[t]=c_j,\vec{r})\}$)available at transmitter k.

To find $F_k^*(C_k[t]=c_i,\vec{r}), F_k^*(C_k[t]=c_j,\vec{r})$, we can solve the optimization (1) with other variables being fixed to the optimal solution. Consider the function that needs to be optimized:

$$\sum_l Q_l E\big[C_l[t]F_l(.)\prod_{m\in I_l}(1-F_m(.))|\mathbf{c}[t-\tau_{max}]\big].$$

Expanding this out, we can write this as

$$\sum_l Q_l \sum_{\vec{z}\in\mathcal{C}^{L\tau_{max}}}\big(Pr(\vec{z}|\mathbf{c}[t-\tau_{max}])C_l(\vec{z})F_l(\vec{z})\times \prod_{m\in I_l}(1-F_m(\vec{z}))\big).$$

Since the variables in the above optimization are only $F_k(C_k[t]=c_i,\vec{r})$ and $F_k(C_k[t]=c_j,\vec{r})$, we ignore the terms in the summation that do not involve these variables (as they are constant and do not affect the arg max). Let A_i denote the set $\{\vec{z}:\vec{z}\in\mathcal{C}^{L\tau_{max}},\mathcal{P}_k(\vec{z})=(c_i,\vec{r})\}$. The new function we now have is:

$$Q_k\sum_{\vec{z}\in A_i\cup A_j}\big(Pr(\vec{z}|\mathbf{c}[t-\tau_{max}])C_k(\vec{z})F_k(\vec{z})\prod_{m\in I_k}(1-F_m(\vec{z}))\big)$$
$$+\sum_{l:l\in I_k} Q_l \sum_{\vec{z}\in A_i\cup A_j}\big(Pr(\vec{z}|\mathbf{c}[t-\tau_{max}])C_l(\vec{z})F_l(\vec{z})\times \prod_{m\in I_l}(1-F_m(\vec{z}))\big).$$

From the above expression, we observe that the above optimization for finding two variables $F_k(c_i,\vec{r}), F_k(c_j,\vec{r})$ splits into two independent optimization problems. First, let us consider the function that needs to be optimized to get $F_k(c_i,\vec{r})$:

$$Q_kF_k(c_i,\vec{r})c_i\sum_{\vec{z}\in A_i}\big(Pr(\vec{z}|\mathbf{c}[t-\tau_{max}])\prod_{m\in I_k}(1-F_m(\vec{z}))\big)+$$
$$(1-F_k(c_i,\vec{r}))\sum_{l:l\in I_k}Q_l\sum_{\vec{z}\in A_i}\big(Pr(\vec{z}|\mathbf{c}[t-\tau_{max}])C_l(\vec{z})F_l(\vec{z}).\prod_{m\in I_l,m\neq k}(1-F_m(\vec{z}))\big).$$

From the above equation, we observe that the optimization function is *linear* in the variable $F_k(c_i, \vec{r})$. Using the fact that channels are independent across links, we have the above function of the form $Pr(C[t] = c_i|\vec{r})(ac_iF_k(c_i,\vec{r}) + b(1 - F_k(c_i,\vec{r})))$, where parameters a and b are independent of value of c_i. Similarly, we can show that the function that needs to be optimized for variable $F_k(c_j,\vec{r})$ is of form $ac_jF_k(c_j,\vec{r}) + b(1 - F_k(c_i,\vec{r}))$. Thus the optimal solution is of the form

$$F_k^*(c_i,\vec{r}) = \begin{cases} 1 & \text{if } ac_i \geq b, \\ 0 & \text{if } ac_i < b. \end{cases}$$

The above solution implies that if $c_j \geq c_i$ and $F_k^*(c_i,\vec{r}) = 1$, then $F_k^*(c_j,\vec{r}) = 1$. This proves the threshold nature of optimal solution.

The proof technique used to show the threshold nature of optimal solution can be used to show that any variable involving channel state information other than critical set NSI results in the same optimization problem. Therefore we have that the optimal solution depends only on the critical set information. ■

Proof: (Lemma 4.3) Let us define the Lyapunov function $V[t]$ as follows,

$$V[t] := \sum_{l\in\mathcal{L}} Q_l^2[t].$$

We thus have,

$$E[V[t+1] - V[t]|\mathbf{Q}[t-\tau_{max}], \mathbf{C}[t-\tau_{max}]] = E[\sum_{l\in\mathcal{L}}(\Delta Q_l[t])(Q_l[t+1] + Q_l[t])|\mathbf{Q}[t-\tau_{max}], \mathbf{C}[t-\tau_{max}]]$$

where $\Delta Q_l[t]$ is the difference $Q_l[t+1] - Q_l[t]$. Using the fact that arrivals and services are bounded in each time slot, we have

$$E\big[V[t+1] - V[t]|(\mathbf{Q}[t-\tau_{max}], \mathbf{C}[t-\tau_{max}])\big] \leq K + E\big[\sum_{l\in\mathcal{L}}(\Delta Q_l[t])(2Q_l[t-\tau_{max}])|(\mathbf{Q}[t-\tau_{max}], \mathbf{C}[t-\tau_{max}])\big].$$

Using the queue update equation, we have

$$E\big[V[t+1] - V[t]|(\mathbf{Q}[t-\tau_{max}], \mathbf{C}[t-\tau_{max}])\big] \leq K + E\big[\sum_{l\in\mathcal{L}}(\mathbf{R}_{l,\tau_{max}}(\mathbf{F}^*(.)))(2Q_l[t-\tau_{max}])| (\mathbf{Q}[t-\tau_{max}], \mathbf{C}[t-\tau_{max}])\big]. \quad (4)$$

Since $(1+\epsilon)\vec{\lambda} \in \Lambda$, there exists $\{\bar{\eta}(\mathbf{c})\}_{\mathbf{c}}$ such that

$$\sum_{\mathbf{c}\in\mathcal{C}^L} \pi(\mathbf{c})\big((1+\epsilon)\lambda_l - \bar{\eta}_l(\mathbf{c})\big) \leq 0.$$

From the scheduling algorithm optimization, we also have that

$$E\big[(\sum_{l\in\mathcal{L}}(\mathbf{R}_{l,\tau_{max}}(\mathbf{F}^*(.)))|\mathbf{C}[t-\tau_{max}] - \bar{\eta}_l(\mathbf{C}[t-\tau_{max}]))Q_l[t-\tau_{max}]\big] \leq 0.$$

Taking the expectation on both sides of inequality (4) over $\mathbf{C}[t-\tau_{max}]$, we have that

$$E\big[V[t+1] - V[t]|\mathbf{Q}[t-\tau_{max}]\big] \leq K_1 - 2\epsilon\sum_l Q_l[t-\tau_{max}]\lambda_l.$$

Now, from the Foster-Lyapunov criteria [16], we have that the network is stochastically stable. ■

Proof: (Lemma 5.3) From the equation (3), we have

$$|R_{l,\tau_1}(\mathbf{T_2^*}) - R_{l,\tau_2}(\mathbf{T_2^*})| = |(\sum_{i=1}^{M} c_i P_{.i}^{\tau_1} 1_{c_i \geq T_{2,l}^*}) \prod_{m\in I_l}(\sum_{i=1}^{M} P_{.i}^{\tau_1} 1_{c_m \geq T_{2,l}^*}) - (\sum_{i=1}^{M} c_i P_{.i}^{\tau_2} 1_{c_i \geq T_{2,l}^*}) \prod_{m\in I_l}(\sum_{i=1}^{M} P_{.i}^{\tau_2} 1_{c_m \geq T_{2,l}^*})|.$$

Let us denote the summation $\sum_{i=1}^{M} c_i P_{.i}^{\tau_1} 1_{c_i \geq T_{2,l}^*}$ by $f_l(\tau_1)$ and the summation $\sum_{i=1}^{M} P_{.i}^{\tau_1} 1_{c_m \geq T_{2,l}^*}$ by $g_m(\tau_1)$. Thus, we have

$$|R_{l,\tau_1}(\mathbf{T_2^*}) - R_{l,\tau_2}(\mathbf{T_2^*})| = |f_l(\tau_1)\prod_{m\in I_l} g_m(\tau_1) - f_l(\tau_2)\prod_{m\in I_l} g_m(\tau_2)|.$$

We now rewrite the above expression by adding and subtracting the term $f_l(\tau_2)\prod_{m\in I_l} g_m(\tau_1)$ to get

$$|R_{l,\tau_1}(\mathbf{T_2^*}) - R_{l,\tau_2}(\mathbf{T_2^*})| = |f_l(\tau_1)\prod_{m\in I_l} g_m(\tau_1) - f_l(\tau_2)\prod_{m\in I_l} g_m(\tau_1) + f_l(\tau_2)\prod_{m\in I_l} g_m(\tau_1) - f_l(\tau_2)\prod_{m\in I_l} g_m(\tau_2)|.$$

Using the triangle inequality results in

$$|R_{l,\tau_1}(\mathbf{T_2^*}) - R_{l,\tau_2}(\mathbf{T_2^*})| \leq |f_l(\tau_1) - f_l(\tau_2)(\prod_{m\in I_l} g_m(\tau_1))| + |f_l(\tau_2)\prod_{m\in I_l} g_m(\tau_1) - f_l(\tau_2)\prod_{m\in I_l} g_m(\tau_2)|.$$

Let the set I_l be expressed as $\{m_1, m_2, m_3, ..., m_l\}$. By iterating the above idea of adding and subtracting terms on the second component of the above expression and using the triangle inequality, we have

$$|R_{l,\tau_1}(\mathbf{T_2^*}) - R_{l,\tau_2}(\mathbf{T_2^*})| \leq |f_l(\tau_1) - f_l(\tau_2)(\prod_{m\in I_l} g_m(\tau_1))| + ... + |f_l(\tau_2)(g_{m_l}(\tau_1) - g_{m_l}(\tau_2))\prod_{k:m_k\in I_l, k\neq l} g_{m_k}(\tau_2)|.$$

Using the following upper bounds, $|f_l(\tau_1) - f_l(\tau_2)| \leq \sum c_i\beta(\tau_1,\tau_2)$, $|g_l(\tau_1) - g_l(\tau_2)| \leq M\beta(\tau_1,\tau_2)$ and $|f_l(\tau_1)| \leq \sum c_i$, we have

$$|R_{l,\tau_1}(\mathbf{T_2^*}) - R_{l,\tau_2}(\mathbf{T_2^*})| \leq$$
$$(\sum c_i)\beta(\tau_1,\tau_2) + (\sum c_i)|I|M\beta(\tau_1,\tau_2)$$
$$= (1+M|I|)(\sum c_i)\beta(\tau_1,\tau_2).$$

■

Proof: (Corollary 5.4) From equation (2), we have

$$\alpha(\tau_1,\tau_2) := \frac{2Lk_o\beta(\tau_1,\tau_2)}{\sum_j c_j \min_i P_{ij}^{\tau_1}}.$$

It is sufficient to prove that $\beta(\tau_1,\infty) \leq (1-M\delta)^{\tau_1}$ and $\beta(\tau_1,\tau_2) \leq 2(1-M\delta)^{\tau_1}\ \forall \tau_2 \geq \tau_1$. Consider the following difference:

$$\begin{aligned} P_{ij}^{\tau} - P_{kj}^{\tau} &= \sum_u (P_{iu} - P_{ku}) P_{uj}^{\tau-1} \\ &= \sum_{u:P_{iu}\geq P_{ku}} (P_{iu}-P_{ku})P_{uj}^{\tau-1} + \\ &\quad \sum_{u:P_{iu}<P_{ku}} (P_{iu}-P_{ku})P_{uj}^{\tau-1}. \end{aligned}$$

Let us denote $\min_u P_{uj}^{\tau}$ by m_j^{τ} and $\max_u P_{uj}^{\tau}$ by M_j^{τ}. We now bound the above difference using m_j^{τ} and M_j^{τ}, we have

$$\begin{aligned} P_{ij}^{\tau} - P_{kj}^{\tau} \leq & \sum_{u:P_{iu}\geq P_{ku}} (P_{iu}-P_{ku})M_j^{\tau-1} \\ &+ \sum_{u:P_{iu}<P_{ku}} (P_{iu}-P_{ku})m_j^{\tau-1}. \end{aligned}$$

By noticing that $\sum_{u:P_{iu}<P_{ku}}(P_{iu}-P_{ku}) + \sum_{u:P_{iu}\geq P_{ku}}(P_{iu}-P_{ku}) = 0$, we have

$$\begin{aligned} &P_{ij}^{\tau} - P_{kj}^{\tau} \leq \sum_{u:P_{iu}\geq P_{ku}} (P_{iu}-P_{ku})\big(M_j^{\tau-1} - m_j^{\tau-1}\big) \\ &= \big(M_j^{\tau-1} - m_j^{\tau-1}\big)\big(\sum_{u:P_{iu}\geq P_{ku}} P_{iu} - \sum_{u:P_{iu}\geq P_{ku}} P_{ku}\big) \\ &= \big(M_j^{\tau-1} - m_j^{\tau-1}\big)\big(1 - \sum_{u:P_{iu}<P_{ku}} P_{iu} - \sum_{u:P_{iu}\geq P_{ku}} P_{ku}\big) \\ &\leq (1-M\delta)\big(M_j^{\tau-1} - m_j^{\tau-1}\big), \end{aligned}$$

where the last inequality follows from the definition of δ. Using the definition of M_j^{τ} and m_j^{τ}, we have that

$$\begin{aligned} M_j^{\tau} - m_j^{\tau} &\leq (1-M\delta)\big(M_j^{\tau-1} - m_j^{\tau-1}\big) \\ &\leq (1-M\delta)^{\tau}. \end{aligned}$$

Using the fact that m_j^{τ} monotonically increases with τ, M_j^{τ} monotonically decreases with τ, and both have a common limit π_j, we have

$$|P_{ij}^{\tau} - \pi_j| \leq (1-M\delta)^{\tau}. \quad (5)$$

Consider the following difference:

$$\begin{aligned} |P_{ij}^{\tau_2} - P_{kj}^{\tau_1}| &= |P_{ij}^{\tau_2} - \pi_j + \pi_j - P_{kj}^{\tau_1}| \\ &\leq |P_{ij}^{\tau_2} - \pi_j| + |\pi_j - P_{kj}^{\tau_1}|. \end{aligned}$$

Using (5) in the above inequality, we have the desired corollary.

■

REFERENCES

[1] M. Andrews, K. Kumaran, K. Ramanan, A. Stolyar, R. Vijayakumar, and P. Whiting, "CDMA data QoS scheduling on the forward link with variable channel conditions," *Bell Labs Tech. Memo*, April 2000.

[2] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Automat. Contr.*, vol. 4, pp. 1936–1948, December 1992.

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

[4] S. Shakkottai and A. Stolyar, "Scheduling for multiple flows sharing a time-varying channel: The exponential rule," *Ann. Math. Statist.*, vol. 207, pp. 185–202, 2002.

[5] A. Stolyar, "Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm," *Queueing Systems*, vol. 50, pp. 401–457, August 2005.

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

[7] M. J. Neely, E. Modiano, and C. E. Rohrs, "Power and server allocation in a multi-beam satellite with time varying channels," in *Proc. IEEE Infocom.*, (New York, NY), pp. 1451–1460, June 2002.

[8] M. Neely, E. Modiano, and C. Li, "Fairness and optimal stochastic control for heterogeneous networks," in *Proc. IEEE Infocom.*, vol. 3, (Miami, FL), pp. 1723–1734, March 2005.

[9] K. Kar, X. Luo, and S. Sarkar, "Throughput-optimal scheduling in multichannel access point networks under infrequent channel measurements," in *Proceedings of IEEE Infocom*, (Anchorage, AK), May 2007.

[10] L. Ying and S. Shakkottai, "On throughput optimality with delayed network-state information," in *Technical Report*, 2008.

[11] L. Ying and S. Shakkottai, "Scheduling in mobile ad hoc networks with topology and channel-state uncertainty," in *Proc. IEEE Infocom.*, pp. 2347–2355, 2009.

[12] A. Gopalan, C. Caramanis, and S. Shakkottai, "On wireless scheduling with partial channel-state information," in *Proc. Ann. Allerton Conf. Communication, Control and Computing*, (Urbana, IL), September 2007.

[13] S. Ahmad, M. Liu, T. Javidi, Q. Zhao, and B. Krishnamachari, "Optimality of myopic sensing in multichannel opportunistic access," *Information Theory, IEEE Transactions on*, vol. 55, pp. 4040 –4050, sep. 2009.

[14] J. Liu, A. L. Stolyar, M. Chiang, and H. V. Poor, "Queue back-pressure random access in multihop wireless networks: optimality and stability," *IEEE Trans. Inf. Theor.*, vol. 55, no. 9, pp. 4087–4098, 2009.

[15] X. Qin and R. Berry, "Opportunistic splitting algorithms for wireless networks," in *Proc. IEEE Infocom.*, March 2004.

[16] P. Billingsley, *Probability and Measure*. New York: Wiley Interscience Publication, 1994.

ACKNOWLEDGEMENT

This work was partially supported by NSF grants CNS-0721380, CNS-0831756, the DARPA ITMANET program and DTRA grant HDTRA1-09-1-0055.