Delay Guarantees for Throughput-optimal Wireless Link Scheduling

Koushik Kar, *Member, IEEE,* Saswati Sarkar, *Member, IEEE,* Abouzar Ghavami and Xiang Luo

***Abstract*—We consider the question of obtaining tight delay guarantees for throughout-optimal link scheduling in arbitrary topology wireless ad-hoc networks. Two classes of scheduling policies are considered: 1) a maximum queue-length weighted independent set scheduling policy, 2) a randomized independent set scheduling policy where the set scheduling probabilities are selected optimally. Both policies stabilize all queues for any set of feasible packet arrival rates, and are therefore throughput-optimal. For these policies, we show that the average packet delay is bounded by a constant that depends on the *chromatic number* of the interference graph, and the *arrival slack* in the system – a metric representing the overall load on the network. We prove that this upper bound is asymptotically tight in the sense that there exist classes of topologies where the expected delay attained by any scheduling policy is lower bounded by the same constant. We extend our upper bounds to the case of multi-hop sessions. Through simulations, we study how our analysis compares with actual delays computed for i.i.d., Markovian and trace-driven packet arrival processes.**

***Index Terms*—Delay analysis, maximum weight scheduling, randomized scheduling.**

I. Introduction

Recent proliferation of commercial wireless services has created large scale demands for transmission of traffic that require stringent quality-of-service (throughput, delay etc.) guarantees. Intelligent scheduling of wireless links is imperative for providing such guarantees. The main challenge in scheduling wireless links is that multiple links in a vicinity can not successfully transmit simultaneously. Efficient resolution of scheduling constraints is the main bottleneck in providing analytical performance guarantees.

Tassiulas *et al.* [19] obtained a link scheduling policy that attains the maximum possible throughput in presence of arbitrary scheduling constraints, by scheduling in each time slot an independent set (in the link interference or conflict graph) that has the maximum aggregate queue length. This policy, referred to as *Maximum Weighted Scheduling (MWS)* henceforth, schedules at any given time instant (a) the set of links that can be simultaneously scheduled while obeying necessary constraints, and (b) has the maximum sum of queue lengths among all such sets.

Obtaining delay guarantees is substantially more difficult than obtaining throughput guarantees, which is itself a challenging problem, due to the following reasons. Throughput guarantees can be obtained by any scheduling policy as long as it ensures that the expected time intervals between successive instants in which the system is empty is finite. However, obtaining delay guarantees is contingent upon ensuring that the above expected duration is low. Specifically, consider a family of variants of MWS which does not schedule *any* link in the system if the queue length of every link is below a certain threshold, say L. For finite L, any such variant ensures that the above expected duration is finite, and therefore maximizes throughput. Yet, the above expected duration, and therefore the delay, attained increases monotonically with increasing L.

Since both throughput and delay are important performance metrics, we seek to obtain provable guarantees on expected delay for policies that maximize throughput. We focus on the following two throughput-optimal policies: 1) MWS [19], and 2) a *Randomized Scheduling (RS)* policy that schedules (independent sets of) links with a fixed probability irrespective of the queue length of the links. MWS has been empirically observed to attain low delay, and does not use any information about the arrival statistics in the scheduling process. While MWS requires solving an maximum weight independent set problem, under certain scenarios MWS can be well-approximated by low-complexity *maximal* scheduling policies (e.g., [16]). Recent experimental work on maximum weight and backpressure scheduling/routing have reported successful distributed implementation (approximation) of MWS variants in the 802.11 framework, by using the MadWiFi device driver to adapt the MAC contention parameters accordingly [15], [20]; RS can also be implemented (approximately) using similar methods. Calculation of the optimum scheduling probabilities for RS requires solving a graph coloring problem, and knowledge of arrival statistics. Once these the arrival rates are estimated (possibly through measurement), however, the scheduling computation needs to be performed only once, and RS can be executed without any further computation or knowledge of global network states. Thus the per-slot (amortized) complexity of implementing RS can be substantially lower than that of MWS if the network topology and packet average arrival rates do not change significantly over time.

We prove that in any network $\mathcal{N}$ the expected delay attained by both MWS and RS is $O(C(\mathcal{N})/\beta)$, where $C(\mathcal{N})$ is the *chromatic number of the link interference graph* for network $\mathcal{N}$, and $1-\beta$ ($0<\beta<1$) is a measure of the *network load* (Section III). More precisely, $C(\mathcal{N})$ represents the minimum number of independent sets ("colors") into which the link interference graph of $\mathcal{N}$ can be partitioned, and β, henceforth referred to as the *arrival slack* in the system, is a measure of the distance between the arrival rate vector and the stability region boundary.

Subsequently, we prove that there exist classes of network topologies where the expected delay attained by *any* scheduling policy is $\Omega(C(\mathcal{N}))$ (Section IV). Thus, for constant β, the delay guarantees attained by MWS and RS are asymptotically tight. Somewhat contrary to intuition, our results show that there exist topology classes under which the static, idling RS policy attains similar delays as the dynamic, non-idling MWS policy. The dependence of the delay bounds on, $C(\mathcal{N})$, the chromatic number of the interference graph, is significant as it identifies the essential topological properties of the network the worst-case delay depends on. It is also useful

K. Kar, A. Ghavami and X. Luo are with the Department of Electrical, Computer and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180, USA (email: {kark,ghavaa,luox3}@rpi.edu).

S. Sarkar is with the Department of Electrical Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA (email: swati@seas.upenn.edu).

This is an extended version of a work that appeared in *IEEE INFOCOM 2009* [11].

from a network design/formation perspective: it implies that for low delay, networks should be formed so that $C(\mathcal{N})$ is small. This can be attained, for example, by keeping the network interference degree low, and forming smaller clusters (as $C(\mathcal{N})$ is upper-bounded by the interference degree plus one, and lower-bounded by the maximum clique size in the interference graph).

We also consider multi-hop generalizations of the MWS and RS policies, and obtain upper bounds on the expected delay as a function of the maximum number of hops of an end-to-end session (Section V).

It is worth noting that the system model and scheduling algorithms that we consider are compatible with both with TDMA and FDMA technologies, provided the interference constraints and the independent sets are defined accordingly. If different nodes in the neighborhood are using different frequencies (FDMA), then only nodal transceiver constraints (or "primary" interference constraints) need to be considered in computing the schedules. However, if nodes share the same channel using TDMA, then two transmissions in the neighborhood may interfere with one another even if their transmitter-receiver pairs are different; in that case, such "secondary interference" constraints also need to be taken into consideration.

The Lyapunov function method that we use for deriving the upper bound on the average delay for MWS has been used earlier in [12], [13] for delay analysis, but for more restrictive network topologies or scheduling policies. Leonardi *et al.* have upper bounded the delay attained by MWS but for an input-queued switching network; in contrast, the networks we consider in this paper can have arbitrary topologies. For the special case of an $N \times N$ switch with traffic rate λ, the general MWS delay upper bound that we derive reduces to $N/(1-\lambda)$, which for large N is nearly the same as the bound of $(N-\lambda)/(1-\lambda)$ obtained in [12]. Neely *et al.* [13] considered a specific scheduling policy, maximal scheduling, and showed that if the arrival traffic is in the stability region of maximal scheduling, the expected delay under maximal scheduling is $O(\log N)$ where N is the number of links in the network. Maximal scheduling can however provide poor throughput guarantees, as depending on the network topology, the stability region of maximal scheduling can become arbitrary small as compared to the optimum. Delay attained by maximal scheduling policies have also been analyzed in [12], [18] albeit in the context of combined input-output queued switches. Shah *et al.* [10] and Sarkar *et al.* [16] have shown that an $O(1)$ expected queue length per link is attainable for the special class of non-expanding graphs (which include random geometric interference graphs). Asymptotic guarantees on queue lengths do not imply similar guarantees on delay (since expected delay is expected queue length divided by the expected arrival rate), without additional assumptions on how the arrival rate scales with increase in the size of the topology. Also, the above guarantees do not apply for arbitrary network topologies. In addition, novel contributions of our work include delay analysis of the RS policy, analysis of the asymptotic tightness of the derived delay guarantees for the two policies (MWS and RS), and extending the results to multi-hop flows.

II. System Model

We consider scheduling at the Medium Access Control (MAC) layer in a wireless network. We assume that time is slotted. A wireless network can be modeled as a directed graph $G = (V, E)$, where V and E respectively denote the sets of nodes and links, and $|E| = N$. A link exists from a node u to another node v if and only if v can receive u's signals. The link set E depends on the transmission power levels of nodes and the propagation conditions in different directions.

Definition 1: A link i *interferes* with a link j if j can not successfully transmit a packet when i is transmitting. The *interference set* of a link i, S_i, is the set of links j such that either i interferes with j or j interferes with i.

Definition 2: The *interference graph* $I^{\mathcal{N}} = (V_I^{\mathcal{N}}, E_I^{\mathcal{N}})$ of a network $\mathcal{N}$ is an undirected graph in which the vertex set $V_I^{\mathcal{N}}$ corresponds to the set of links in $\mathcal{N}$ and there is an edge between two vertices i and j if $j \in S_i$.

Definition 3: An *independent set* in a graph is a subset of its vertices such that there does not exist an edge between any two vertices in the subset. Let $J_1, \ldots, J_M$ be the independent sets of $I^{\mathcal{N}}$, and let $\vec{J^i}$ be the indicator vector representing any independent set J_i. Let $\mathcal{J} = \{J_1, \ldots, J_M\}$.

Definition 4: A *coloring* of a graph is an allocation of colors to vertices of the graph such that no two vertices that have an edge between them is assigned the same color. The *chromatic number* of a graph is the minimum number of colors required for coloring the graph. Equivalently, it is the minimum number of independent sets of a graph that can partition its vertex set.

Let $\mathcal{C}(\mathcal{N}) = \{V_1, \ldots, V_{C(\mathcal{N})}\}$ represent a minimum coloring of the link interference graph $I^{\mathcal{N}}$, where $V_1, \ldots, V_{C(\mathcal{N})}$ are the subsets of the vertices of $I^{\mathcal{N}}$ that have been assigned the same color. Clearly, $\mathcal{C}(\mathcal{N}) \subseteq \mathcal{J}$, and $C(\mathcal{N}) = |\mathcal{C}(\mathcal{N})|$ represents the chromatic number of $I^{\mathcal{N}}$.

We initially consider only single-hop flows, and in Section V we generalize our results to multi-hop flows. We now describe the arrival process for the single-hop flows (links). Let $A_i(t)$ be the number of packets that link i generates in interval $(t, t+1]$, $i = 1, \ldots, N$. We assume that for each i, $\mathbf{E}(A_i(t)) = \lambda_i$, where λ_i is referred to as the *arrival rate* of link i. We also assume that $\mathbf{E}(A_i^2(t)) \leq \gamma \mathbf{E}(A_i(t))$, where γ is a constant that depends on the distribution of the arrival process. A sufficient (but not necessary) condition for this to hold is that the maximum number of packets that arrive in a slot is upper bounded by γ.

Definition 5: The *arrival rate vector* $\vec{\lambda}$ is an $N-$dimensional vector of the arrival rates.

A *scheduling policy* decides in each slot the subset of links that would transmit packets in the slot. Clearly, a scheduling policy must select an element of $\mathcal{J}$ in each slot.

Every packet has a transmission time of one slot. We assume that any packet arriving in a slot may be transmitted in the next slot. Let $D_i(t)$ be the number of packets that link i transmits in interval $(t, t+1]$, $i = 1, \ldots, N$. Clearly the transmissions depend on the scheduling policy. Let $Q_i(t)$ be the number of packets that are waiting for transmission in link i at the beginning of slot t. Let $\vec{Q}(t), \vec{A}(t), \vec{D}(t)$ be the queue length, arrival and departure vectors respectively, with components, $Q_i(t), A_i(t), D_i(t)$ respectively. In each time slot, we assume an integral number of arrivals (possibly multiple) and departures (at most one). Thus $A_i(\cdot) \in \{0, 1, \ldots\}, D_i(\cdot) \in \{0, 1\}, Q_i(\cdot) \in \{0, 1, \ldots\}$, and

$$Q_i(t+1) = Q_i(t) + A_i(t) - D_i(t). \tag{1}$$

We now describe two scheduling policies that we analyze in this paper: 1) maximum weighted scheduling (MWS), and 2) randomized scheduling (RS($\vec{p}$)). MWS considers the weight of an independent set J_i as the sum of the queue lengths of the links in J_i, and in each slot t schedules the independent set that has the maximum weight among all independent sets in $\mathcal{J}$. In each slot t, RS($\vec{p}$), schedules independent set J_i with probability p_i irrespective of the queue lengths of the links, and the schedules selected in different slots are mutually independent. Here, $\vec{p}$ is a M-dimensional probability vector such that $\sum_{i=1}^{M} p_i = 1$ and $p_i \geq 0$ for each i. Note that since RS schedules independent sets at random without considering queue length information, it can possibly schedule an independent set with all empty queues, while there are other independent sets with backlogged queues in the system. In contrast, MWS is non-idling, in the sense that it will never serve an independent set with all empty queues as long as there are backlogged queues in the system.

Definition 6: The network is said to be *stable* if $\lim_{T\to\infty} \sum_{n=1}^{T} Q_i(t)/T$ is finite.

Alternative but closely related definitions of network stability have also been considered in the literature, like finiteness of the $\limsup$ of the average queue-lengths [14], positive recurrence of the queue-length process [19], and boundedness of the expected queue-lengths [17]. The two policies we analyze in this paper result in network stability (for any "admissible" arrival rate vector) under all these stability criteria.

Definition 7: The *stability region* of a scheduling policy is the set of arrival rate vectors for which the network is stable when the policy is used. An arrival rate vector $\vec{\lambda}$ is said to be *feasible* if it is in the stability region of some scheduling policy. The *network stability region* Λ is the set of all feasible arrival rate vectors.

Define $\Lambda_\beta = \{\vec{\lambda} : \vec{\lambda} = \sum_{i=1}^{M} w_i \vec{J}^i$, for some $w_1, \ldots, w_M$, such that $w_i \geq 0, \sum_{i=1}^{M} w_i = 1 - \beta\}$. From the above definition, an arrival rate vector $\vec{\lambda}$ is in Λ_β if $(1 - \sum_{i=1}^{M} w_i) = \beta$. In other words, if an arrival rate vector $\vec{\lambda}$ is in Λ_β then its "distance" from the boundary of the network stability region is β, and β is then denoted as the *arrival slack*.

Definition 8: The *expected delay* in a network is the expected number of time slots that elapse between the arrival and departure of a packet.

III. Upper Bounds on Expected Delay

The maximum weighted scheduling policy (MWS) is known to attain the network stability region, and thereby maximizes network throughput. With an appropriate choice of scheduling probabilities $\vec{p}$, the randomized scheduling policy (RS) can be easily shown to attain maximum throughput as well. In this section, we upper bound the expected delays of MWS and RS($\vec{p}$).

We analyze the delay under the assumption of i.i.d. arrivals, which allows a simpler analysis and slightly tighter bounds; the analysis can however be extended to Markovian arrivals as well.

We will show that for any given network $\mathcal{N}$, when the arrival slack is β, the expected delays attained by these policies are $O(C(\mathcal{N})/\beta)$ (Theorems 1, 2).

Theorem 1: Consider a network $\mathcal{N}$, and a $\vec{\lambda} \in \Lambda_\beta$ where $\beta \in (0, 1)$. Then, the expected delay attained by MWS in $\mathcal{N}$ is at most $\frac{\gamma+1}{2}\frac{C(\mathcal{N})}{\beta}$.

Proof: For the Lyapunov function $U(\vec{x}) = \sum_{i=1}^{N} x_i^2$, using (1),

$$\begin{aligned} U(\vec{Q}(t+1)) - U(\vec{Q}(t)) \;\leq\; & 2\left(\vec{A}(t) - \vec{D}(t)\right)^T \vec{Q}(t) + \\ & \left(\vec{A}(t)\right)^T \left(\vec{A}(t)\right) + \left(\vec{D}(t)\right)^T \left(\vec{D}(t)\right). \end{aligned} \tag{2}$$

Now, since $\vec{\lambda} \in \Lambda_\beta$, $\beta \in (0, 1)$, $\{\vec{Q}(u), u = 0, 1, 2, \ldots\}$ constitutes a positive recurrent Markov chain [19]. We take expectations of both sides of (2) considering that the distribution of $\vec{Q}(t)$ is the stationary distribution of this Markov chain. Considering the left side,

$$\mathbf{E}\left(U(\vec{Q}(t+1)) - U(\vec{Q}(t))\right) = 0. \tag{3}$$

Now, consider the right hand side. Since $\vec{\lambda} \in \Lambda_\beta$, there exists $w_1, \ldots, w_M$ such that $\sum_{i=1}^{M} w_i = 1 - \beta$ and $\vec{\lambda} = \sum_{k=1}^{M} w_k \vec{J}^k$. Thus, since $\vec{A}(t)$ is independent of $\vec{Q}(t)$,

$$\mathbf{E}\left(\left(\vec{A}(t)\right)^T \vec{Q}(t) | \vec{Q}(t)\right) \;=\; \vec{\lambda}^T \vec{Q}(t) \;=\; \sum_{k=1}^{M} w_k (\vec{J}^k)^T \vec{Q}(t).$$

Thus, since $\sum_{i=1}^{M} w_i = 1 - \beta$, $\mathbf{E}\left(\left(\vec{A}(t)\right)^T \vec{Q}(t)\right) \leq (1 - \beta)\mathbf{E}\max_{k=1}^{M} (\vec{J}^k)^T \vec{Q}(t)$.

Using the properties of MWS scheduling, $\left(\vec{D}(t)\right)^T \vec{Q}(t) = \max_{k=1}^{M} (\vec{J}^k)^T \vec{Q}(t)$. Thus,

$$\mathbf{E}\left(\vec{A}(t) - \vec{D}(t)\right)^T \vec{Q}(t) \leq -\beta \mathbf{E} \max_{k=1}^{M} (\vec{J}^k)^T \vec{Q}(t). \tag{4}$$

Since $\mathbf{E}\left(A_i^2(t)\right) \leq \gamma \mathbf{E}\left(A_i(t)\right)$, we have

$$\mathbf{E}\left(\vec{A}(t)\right)^T \left(\vec{A}(t)\right) \leq \gamma \sum_{i=1}^{N} \lambda_i. \tag{5}$$

Since the components of the departure vector are either 0 or 1, $D_i^2(t) = D_i(t)$. Since the distribution of $\vec{Q}(t)$ is the stationary distribution of the Markov chain $\{\vec{Q}(u), u = 0, 1, \ldots\}$, $\mathbf{E}\left(\vec{D}(t)\right) = \vec{\lambda}$. Thus,

$$\mathbf{E}\left(\vec{D}(t)\right)^T \left(\vec{D}(t)\right) = \mathbf{E}\left(\sum_{i=1}^{N} D_i(t)\right) = \sum_{i=1}^{N} \lambda_i. \tag{6}$$

From (3) to (6), we have $0 \leq -2\beta \mathbf{E} \max_{k=1}^{M} \left(\vec{J}^k\right)^T \vec{Q}(t) + (\gamma+1)\sum_{i=1}^{N} \lambda_i$.

$$\text{Thus, } \mathbf{E} \sum_{k \in J} Q_k(t) \leq \frac{(\gamma+1)\sum_{i=1}^{N} \lambda_i}{2\beta} \text{ for any } J \in \mathcal{J}. \quad (7)$$

Now, $\sum_{k=1}^{N} Q_k(t) = \sum_{j=1}^{C(\mathcal{N})} \sum_{k \in V_j} Q_k(t)$. Thus, from (7) and since $V_j \in \mathcal{J}$,

$$\mathbf{E}\left(\sum_{k=1}^{N} Q_k(t)\right) \leq C(\mathcal{N})(\gamma+1)\frac{\sum_i \lambda_i}{2\beta}.$$

Since the expected delay is $\mathbf{E}\left(\sum_{k=1}^{N} Q_k(t)\right) / \left(\sum_{i=1}^{N} \lambda_i\right)$, the result follows. ∎

Theorem 2: Consider a network $\mathcal{N}$, and $\vec{\lambda} \in \Lambda_\beta$ where $\beta \in (0,1)$. When the arrival process is Bernoulli, there exists a M-dimensional probability vector $\vec{p}$, such that the expected delay attained by RS($\vec{p}$) in $\mathcal{N}$ is at most $\frac{C(\mathcal{N})}{\beta}$.

Proof: Since $\vec{\lambda} \in \Lambda_\beta$, there exists $w_1, \ldots, w_M$ such that $\sum_{i=1}^{M} w_i = 1-\beta$ and $\vec{\lambda} = \sum_{k=1}^{M} w_k \vec{J}^k$. Recall that $\mathcal{C}(\mathcal{N}) \subseteq \mathcal{J}$ is a collection of independent sets that constitutes a minimum coloring of the link interference graph. Let $p_k = w_k$ if $J_k \in \mathcal{J} \setminus \mathcal{C}(\mathcal{N})$, and $p_k = w_k + \beta/C(\mathcal{N})$, if $J_k \in \mathcal{C}(\mathcal{N})$. Note that $p_k \geq 0$, and $\sum_{k=1}^{M} p_k = \sum_{k=1}^{M} w_k + |\mathcal{C}(\mathcal{N})|\beta/C(\mathcal{N})$. Since $|\mathcal{C}(\mathcal{N})| = C(\mathcal{N})$, $\sum_{k=1}^{M} p_k = \sum_{k=1}^{M} w_k + \beta = 1$. Thus, $\vec{p}$ is a M-dimensional probability vector.

When an independent set in $\mathcal{J}$ is scheduled, all links in it are scheduled. Also, each link l is in one independent set in $\mathcal{C}(\mathcal{N})$. Thus, each link l is scheduled with probability $\sum_{i=1}^{M} w_i \vec{J}^i_l + \beta/C(\mathcal{N})$, where $\vec{J}^i_l$ is the component corresponding to link l in $\vec{J}^i$, the indicator vector representing the independent set J_i. Since $\lambda_l = \sum_{i=1}^{M} w_i \vec{J}^i_l$, l is scheduled with probability $\lambda_l + \beta/C(\mathcal{N})$. Since the arrival process is Bernoulli, it follows from eqn. (16) of [9] that the expected delay for link l is $\frac{1-\lambda_l}{\lambda_l+\beta/C(\mathcal{N})-\lambda_l}$ which is at most $C(\mathcal{N})/\beta$. ∎

Note that unlike Theorem 1, Theorem 2 relies on an additional assumption that the arrival process is Bernoulli. Specifically, for Bernoulli arrivals, $\gamma = 1$, and the upper bounds on the delay in Theorems 1 and 2 turn out to be the same in that case. Computing the probability vector $\vec{p}$ that attains the above delay guarantee requires finding the weights $\{w_k\}$ from the arrival vector $\vec{\lambda}$, and finding the chromatic number of the interference graph. The rate decomposition problem to calculate the independent set weights $\{w_k\}$ is not known to be polynomially solvable in the general case; in presence of primary interference constraints only (i.e., if two links can transmit together successfully as long as they do not have any common end node), this problem can be solved in polynomial time [8]. Finding the chromatic number requires solving a graph coloring problem, which is NP-hard. However, this complex calculation need not be done on a per slot basis – it can be done once at the very beginning, and the scheduling probabilities $\{p_k\}$ thus computed can be used thereafter, until the network topology changes. In contrast, MWS requires solving the NP-hard maximum-weighted independent set problem at each scheduling instant.

In a contemporary but independent work, Gupta and Shroff [5], [6] have recently upper bounded the expected delay for a different version of the MWS algorithm, where the the queue-length Q_i for any link i is weighted by $w_i = \frac{1}{\mu_i - \lambda_i}$ while computing the maximum weighted schedule (independent set). The delay bound is expressed in terms of the additional parameters $(\mu_i, i = 1, \ldots N)$, which must be chosen carefully so as to attain low expected delay. Typically, a good choice of the parameters μ_i would require knowledge of the arrival rates λ_i and the structure of the stability region. We however provide an upper bound on the expected delay of MWS, as defined in [19], which does not require any knowledge of arrival rates or the structure of the stability region (other than knowledge of the different independent sets).

IV. Lower Bounds on Expected Delay

We now obtain a lower bound on the expected delay of an arbitrary policy (Theorem 3). The bound is derived for a specific class of networks (as described in the proof of Theorem 3), for which we show that the delay attained by any policy must grow asymptotically with the chromatic number at the same rate as indicated in the upper bounds derived in Section III. While the obtained lower bound is not applicable for general topology networks, it shows that the delay bounds in Theorems 1 and 2 cannot be improved asymptotically across all networks.

Theorem 3: For any real numbers $\beta \in (0,1), \epsilon \in (0,1)$, and any positive integer $\tilde{C}$, there exists a network $\mathcal{N}$ with $C(\mathcal{N}) = \tilde{C}$ and an arrival rate vector in Λ_β, such that the expected delay attained in $\mathcal{N}$ by any scheduling policy is at least $\frac{(1-\epsilon)^2}{(1+\epsilon)} \frac{\tilde{C}}{2}$.

We next show that a tighter lower bound can be obtained for any randomized scheduling by exploiting the relation between the expected delay and β. This bound is also derived for a specific class of networks (as described in the proof of Theorem 4, and is different from that used in the proof of Theorem 3), but serves to show that there exist networks for which the delay bound derived in Theorem 2 is asymptotically tight.

Theorem 4: For any real number $\beta \in (0,1)$, any probability vector $\vec{p}$, and any positive integer $\tilde{C}$, there exists a network $\mathcal{N}$ with $C(\mathcal{N}) = \tilde{C}$ and an arrival rate vector in Λ_β, such that the expected delay attained in $\mathcal{N}$ by RS($\vec{p}$) is at least $\frac{\tilde{C}-1}{\beta} + 1$.

Note that for large $\tilde{C}$, the lower bound $\frac{\tilde{C}-1}{\beta} + 1 \approx \frac{\tilde{C}}{\beta}$. Proving a similar lower bound that involves the arrival slack β for an arbitrary scheduling policy remains open.

Proof of Theorem 3: We first describe a network $\mathcal{N}$ with $C(\mathcal{N}) = \tilde{C}$. The network consists of $\tilde{C}$ disjoint groups (sets) of links, $J_1, \ldots, J_{\tilde{C}}$, each of size K. (The value for K will be specified later.) Thus each J_i consists of K links, and $J_i \cap J_l = \emptyset$ if $i \neq l$. Also, links belonging to different groups interfere with each other, i.e., links i, l can be served simultaneously if and only if $i, l \in J_w$ for some w. Sets $J_1, \ldots, J_{\tilde{C}}$ and their subsets constitute the independent sets of the interference graph of $\mathcal{N}$. Clearly, the colors assigned to links in J_i and J_l must be different if $i \neq l$. It is easy to see that the chromatic number of the interference graph is equal to $\tilde{C}$.

Let $\alpha = (1-\beta)/\tilde{C} \in (0,1)$. Consider an arbitrary probability distribution F on non-negative integers with expectation α that satisfies the following technical condition: the moment generating function $Z(\tau)$ of the distribution F (i.e., $\mathbf{E}\left(\exp(\tau X)\right)$ where X is a random variable with distribution F) is finite in some neighborhood of $\tau = 0$. Note that F can be selected from a large class of probability distributions which consists of, but is not limited to, Bernoulli(α), Poisson(α), Binomial(x, y) with $xy = \alpha$, etc. Let packets arrive in each link as per temporally and mutually independent random processes with distribution F each. Note that the corresponding arrival rate vector $\vec{\lambda} = (\alpha, \ldots, \alpha)$ can be expressed as $\vec{\lambda} = \sum_i w_i \vec{J^i}$, where $w_i = \alpha$, $i = 1, \ldots, \tilde{C}$, and $w_i = 0$ for $i > \tilde{C}$. Since $\sum_i w_i = \alpha|\tilde{C}| = 1-\beta$, $\vec{\lambda} \in \Lambda_\beta$.

Consider all the packets that arrive in an arbitrary slot t, and let the delay of these packets be denoted by a random variable $\hat{D}$. In the following we show that $\mathbf{E}\hat{D} \geq (1-\epsilon)^2\tilde{C}/(2(1+\epsilon))$. for large enough K. Let X_i be the total number of packets received by links in J_i in t, for $i = 1, \ldots, \tilde{C}$. Consider an event A in which $K\alpha(1-\epsilon) \leq X_i \leq K\alpha(1+\epsilon)$ for each i. Clearly, $\mathbf{E}\hat{D} \geq Pr(A)\mathbf{E}(\hat{D}|A)$. We bound $\mathbf{E}(\hat{D}|A)$ first. Under A, it is easy to see that the total delay of all packets under consideration (i.e. packets that have arrived in slot t) must be at least $K\alpha(1-\epsilon)(1+2+\ldots+\tilde{C}) = K\alpha(1-\epsilon)\frac{\tilde{C}(\tilde{C}+1)}{2}$, since at most one of the independent sets J_i, $i = 1, \ldots, \tilde{C}$, can be scheduled in any slot. Under A, since the number of such packets is upper bounded by $K\alpha(1+\epsilon)\tilde{C}$, we have $\mathbf{E}(\hat{D}|A) \geq \left(K\alpha(1-\epsilon)\frac{\tilde{C}(\tilde{C}+1)}{2}\right) / \left(K\alpha(1+\epsilon)\tilde{C}\right) > \frac{\tilde{C}}{2}\frac{1-\epsilon}{1+\epsilon}$.

Next we bound $Pr(A)$. From large deviation results, for all large enough K, the probability that $X_i \notin [K\alpha(1-\epsilon), K\alpha(1+\epsilon)]$ is at most $3\exp(-K\nu)$ where ν is a positive constant that depends on the distribution F and α, ϵ (Section 5.11 [4])). Using union bound, for all large enough K, the probability that $X_i \notin [K\alpha(1-\epsilon), K\alpha(1+\epsilon)]$ for at least one i is at most $3\tilde{C}\exp(-K\nu)$. Thus, $Pr(A) \geq 1-\epsilon$ for all large enough K. The result follows. ∎

Proof of Theorem 4: We first describe a network $\mathcal{N}$ with $C(\mathcal{N}) = \tilde{C}$. The network consists of $\tilde{C}$ links such that any two links in the network interfere with each other. Thus, the chromatic number of the link interference graph is $\tilde{C}$. The independent sets of the interference graph consist of $J_0, J_1, \ldots, J_{\tilde{C}}$, where $J_0 = \emptyset$, and J_i consists only of link i when $i \geq 1$.

Consider a Bernoulli arrival process for which the arrival rate vector is in Λ_β of $\mathcal{N}$. Let $\alpha = (1-\beta)/\tilde{C} \in (0,1)$. Let the arrival process at each link be Bernoulli(α), independent of the arrival processes at other links. The arrival processes in different slots are also independent. Note that the corresponding arrival rate vector $\vec{\lambda} = (\alpha, \ldots, \alpha)$ can be expressed as $\vec{\lambda} = \sum_{i=0}^{\tilde{C}} w_i \vec{J^i}$, where $w_i = \alpha$, $i = 1, \ldots, \tilde{C}$, and $w_0 = 0$. Since $\sum_i w_i = \alpha|\tilde{C}| = 1-\beta$, $\vec{\lambda} \in \Lambda_\beta$.

Consider RS($\vec{p}$) for an arbitrary $(\tilde{C}+1)-$dimensional probability vector $\vec{p}$. Then, in each slot, RS($\vec{p}$) serves link i with probability p_i for $i \geq 1$, and the service opportunities are temporally independent. From eqn. (16) of [9], the expected delay in link i is $\frac{1-\alpha}{p_i-\alpha}$. Since all links have equal arrival rates, the overall expected delay is $\frac{1}{\tilde{C}}\sum_{i=1}^{\tilde{C}}\frac{1-\alpha}{p_i-\alpha}$. Thus, if $\vec{p}^*$ minimizes the expected delay for RS($\vec{p}$) among all possible choices for $\vec{p}$, then $\vec{p}^*$ must be the optimum solution of the following symmetric convex optimization problem:

$$\text{minimize} \quad \frac{1}{\tilde{C}}\sum_{i=1}^{\tilde{C}}\frac{1-\alpha}{p_i-\alpha},$$
$$\text{subject to:} \quad p_i \geq 0,\ i = 1, \ldots, \tilde{C}, \quad \text{and} \quad \sum_{i=1}^{\tilde{C}} p_i \leq 1.$$

It is easy to show that the optimum solution is given by $p_i^* = 1/\tilde{C}$. Since $\alpha = (1-\beta)/\tilde{C}$, the expected delay under RS($\vec{p}^*$) is $\frac{(1-\frac{1-\beta}{\tilde{C}})\tilde{C}}{\beta} = \frac{\tilde{C}-(1-\beta)}{\beta} = \frac{\tilde{C}-1}{\beta} + 1$. ∎

In a contemporary but independent work, Gupta and Shroff [6] have obtained lower bounds for delays for arbitrary policies in arbitrary networks. Since our goal has been to show the tightness of the upper bounds for MWS and RS, we have focused on specific classes of networks. The lower bound in [6] is expressed in terms of *local* (node-specific) *exclusive sets* (cliques), whereas our bounds (both upper and lower) are derived in terms of the *chromatic number* of the *entire network*. For the specific networks we have considered (those used in the proof of Theorem 3), the lower bound derived in [6], reduces to $1/(2\beta)$ when $\tilde{C} = 1$, and 0 otherwise. Thus, for a small β and $C = 1$, this lower bound is tighter than our bound of $\tilde{C}/2$ (approximately); Theorem 3 provides a tighter bound otherwise.

V. Multi-hop sessions

We consider a network where sessions (flows) can span multiple links, and obtain upper bounds on the delays of two throughput-optimal policies that are natural generalizations of MWS and RS. Let there be K end-to-end flows, $1, \ldots, K$ each spanning at most P links. Flow i traverses $|P_i|$ links denoted by $l_1, l_2, \ldots, l_{P_i}$, and corresponds to *flow-link* pairs $(i, l_1), \ldots, (i, l_{P_i})$. Let $\mathcal{F}$ be the set of flow-links, and l_{zi} denote the flow-link of flow i in its z-th hop. Two flow-links interfere if they correspond to the same link, or if the corresponding links interfere as per the pairwise link interference relations. The interference graph for the flow-links can be described as in Section II. Let $C^{\mathcal{M}}(\mathcal{N})$ and $\mathcal{J}^{\mathcal{M}}$ respectively denote the chromatic number and the collection of independent sets for this interference graph – there exists $C^{\mathcal{M}}(\mathcal{N})$ disjoint independent sets in $\mathcal{J}^{\mathcal{M}}$, $V_1^{\mathcal{M}}, V_2^{\mathcal{M}}, \ldots$ whose union equals $\mathcal{F}$.

A flow-link (i, l) has a packet arrival if (i) a new packet is generated for flow i and $l = l_{1i}$, or (ii) the previous hop for i transmits a packet. We assume that the extraneous packet arrival process for the flows satisfy the same assumptions as in Section II, with rates λ_i, $i = 1, \ldots, K$. A scheduling policy can schedule any independent set in $\mathcal{J}^{\mathcal{M}}$. Now, let the vectors $\vec{Q}(t), \vec{A}(t), \vec{D}(t)$ respectively denote the packet queue lengths (packets waiting for transmission), arrivals and departures of the flow-links. The notions of stability, network stability region Λ^M, arrival slack β, and $\Lambda_\beta^{\mathcal{M}}$ can be defined similar to Section II, but with flow-links instead of links. Finally, let $B_{i,l}(t)$ be the difference in backlog (queue length difference) across flow-link (i, l), i.e. $B_{i,l}(t) = Q_{i,l}(t) - Q_{i,l'}(t)$ and l' is the hop next to l for flow i (if l is not the last hop for i), and $B_{i,l}(t) = Q_{i,l}(t)$ otherwise. $\vec{B}(t)$ is the vector of these backlog differences.

Next, we describe the throughput-optimal policies we consider. RS($\vec{p}$) schedules independent set $J \in \mathcal{J}^{\mathcal{M}}$ w.p. p_J in any slot, where p_Js are chosen appropriately. With multi-hop sessions in wireless networks, the throughput-optimality of the MWS policy – that schedules the independent set with the maximum sum of backlogs (queue lengths) at any time – remains an open question. It is worth noting here that in certain classes of networks, like a network of switches as considered in [2], the MWS scheduling policy may not be throughput-optimal for multi-hop flows. Since our focus is on throughput-optimal policies, we consider the *back-pressure* MWS policy described in [19], denoted here by *MWS-BP*, which is known to be throughput-optimal. MWS-BP is the same as MWS except that it considers the weight of an independent set J as the sum of the backlog differences across the flow-links in J. Both MWS-BP and RS($\vec{p}$) (for an appropriate choice of $\vec{p}$) attain the network stability region. We now upper bound their delays.

Theorem 5: Let $\vec{\lambda} \in \Lambda_\beta^{\mathcal{M}}$ where $\beta \in (0,1)$. Then, the expected delay attained by MWS-BP in $\mathcal{N}$ is at most $\frac{\gamma+1}{2} P^2 \frac{C^{\mathcal{M}}(\mathcal{N})}{\beta}$.

Proof: Consider the quadratic Lyapunov $U(\vec{x}) = \sum_{i=1}^N x_i^2$, and note that (2), (3) hold since $\vec{\lambda} \in \Lambda_\beta$, $\beta \in (0,1)$. Using an analysis similar to [19], and the properties of back-pressure policy, and considering the expectations w.r.t the stationary distribution of $(\vec{Q}(t), \vec{A}(t), \vec{D}(t))$

$$\mathbf{E}\left(\vec{A}(t) - \vec{D}(t)\right)^T \vec{Q}(t) \leq -\beta \mathbf{E} \max_{J \in J^{\mathcal{M}}} (\vec{J})^T \vec{B}(t). \quad (8)$$

Note that $\mathbf{E}\left(A_{i,l}^2(t)\right) \leq \gamma \mathbf{E}\left(A_{i,l}(t)\right)$ for any flow-link pair (i,l). This follows from the statistical assumptions on the extraneous arrivals of each flow, and also because each intermediate hop of a flow receives at most 1 arrival in each slot. Also, $\mathbf{E}\left(A_{i,l}(t)\right) = \lambda_i$ since the system is stable as $\vec{\lambda} \in \Lambda_\beta \subset \Lambda$.

$$\text{Thus, } \mathbf{E}\left(\vec{A}(t)\right)^T \left(\vec{A}(t)\right) \leq \gamma \sum_{i=1}^K \lambda_i |P_i|. \quad (9)$$

$$\text{Similar to (6), } \mathbf{E}\left(\vec{D}(t)\right)^T \left(\vec{D}(t)\right) = \sum_{i=1}^K \lambda_i |P_i|. \quad (10)$$

From (2), (3), (8), (9), (10), similar to (7), and since $|P_i| \leq P$ for all i,

$$\mathbf{E} \sum_{f \in J} B_f(t) \leq \frac{(\gamma+1) P \sum_{i=1}^K \lambda_i}{2\beta} \quad \text{for any } J \in \mathcal{J}^{\mathcal{M}}. \quad (11)$$

Let $\mathcal{F}^{(z)}$ constitute the flow-links that correspond to the z-th hops (l_{zi}) of the respective flows, and $\tilde{\mathcal{F}}^{(z)}$ constitute the flow links in $\mathcal{F}^{(z)} \cup \mathcal{F}^{(z+1)} \cup \ldots$. Now, $\sum_{i=1}^K Q_{i,l_{zi}}(t) = \sum_{f \in \tilde{\mathcal{F}}^{(z)}} B_f(t) = \sum_{j=1}^{C^{\mathcal{M}}(\mathcal{N})} \sum_{k \in (V_j^{\mathcal{M}} \cap \tilde{\mathcal{F}}^{(z)})} B_k(t)$. (If flow i does not have z hops, $Q_{i,l_{zi}}(t) = 0$.) Thus, from (11) and since $V_j^{\mathcal{M}} \cap \tilde{\mathcal{F}}^{(z)} \in \mathcal{J}^{\mathcal{M}}$, for each $z = 1, \ldots P$,

$$\begin{aligned} \mathbf{E}\left(\sum_{i=1}^K Q_{i,l_{zi}}(t)\right) &\leq PC^{\mathcal{M}}(\mathcal{N})(\gamma+1)\frac{\sum_i \lambda_i}{2\beta} \\ \mathbf{E}\left(\sum_{i=1}^K \sum_{z=1}^P Q_{i,l_{zi}}(t)\right) &\leq P^2 C^{\mathcal{M}}(\mathcal{N})(\gamma+1)\frac{\sum_i \lambda_i}{2\beta} \end{aligned}$$

Since the expected delay is $\mathbf{E}\left(\sum_{i=1}^K \sum_{z=1}^P Q_{i,l_{zi}}(t)\right) / \left(\sum_{i=1}^K \lambda_i\right)$, the result follows. ■

In contemporary but independent work, Bui *et al.* [3] and Gupta and Shroff [7] have obtained bounds on the average delay (or queue-lengths) of MWS-BP or its variants, for multi-hop flows. The authors in [3] analyze the back-pressure scheduling policy and also obtain a (P^2) dependence of the delay upper bound on the maximum hop-count P, but without link interference constraints. Therefore the model and results in [3] are more applicable to wired networks, whereas we consider wireless link interference constraints as well. Even if link interference constraints do not exist, Theorem 5 will in general provide a tighter characterization of the delay in terms of the network topology and traffic load, than the corresponding delay bound in [3] which depends linearly with the number of flows, and inversely with the traffic load of the flow, and uses a stronger notion (smaller value) of the “arrival slack” than ours. The authors in [7] do consider general interference models, but focus of deriving lower bounds on the delay across all scheduling policies, and study through simulations how MWS-BP and a proposed variant of it, perform with respect to the lower bound. Unlike our work, [7] does not derive any upper bound on the delay of MWS-BP or its variants for networks with general topologies.

Theorem 6: Let $\vec{\lambda} \in \Lambda_\beta^{\mathcal{M}}$ where $\beta \in (0,1)$. When the extraneous arrival process is Bernoulli, there exists a $|\mathcal{J}^{\mathcal{M}}|$-dimensional probability vector $\vec{p}$, such that the expected delay attained by RS($\vec{p}$) in $\mathcal{N}$ is at most $P\frac{C^{\mathcal{M}}(\mathcal{N})}{\beta}$.

Proof: As in the proof of Theorem 2 for single-link sessions, each flow link l_{zk} is scheduled in each slot with probability ϑ_{zk} independent of the scheduling events in other slots, where ϑ_{zk} is the sum of the probabilities of selections of the independent sets containing l_{zk}. Now, the arrival process (which is extraneous) to the flow link l_{1k} at the first hop of the kth flow is Bernoulli(λ_k) for each z. Thus, the departure process of flow-link l_{1k}, which is also the arrival process of the flow-link l_{2k}, is Bernoulli(λ_k) as well (last paragraph of Section II, p. 359, [9]). Thus, applying the same argument recursively, the arrival process of the flow-links l_{zk} is Bernoulli(λ_k) for each k. Thus, using a line of analysis similar to proof of Theorem 2, the average single-hop delay of packets in the network, $\mathbf{E}(\sum_{f \in \mathcal{F}} Q_f)/(\sum_{f \in \mathcal{F}} \lambda_f)$ is upper bounded by $\frac{C^{\mathcal{M}}(\mathcal{N})}{\beta}$. Here λ_f is the packet arrival rate on flow-link $f = (i,l)$, which equals λ_i, the extraneous packet arrival rate of flow i, due to system stability. Since each flow traverses at most P hops, $\sum_{f \in \mathcal{F}} \lambda_f \leq P \sum_{i \in K} \lambda_i$. Thus the average end-to-end delay, $\mathbf{E}(\sum_{f \in \mathcal{F}} Q_f)/(\sum_{i \in K} \lambda_i)$ is upper bounded by $P\frac{C^{\mathcal{M}}(\mathcal{N})}{\beta}$. ■

From Theorems 5 and 6, we observe that the upper bound on

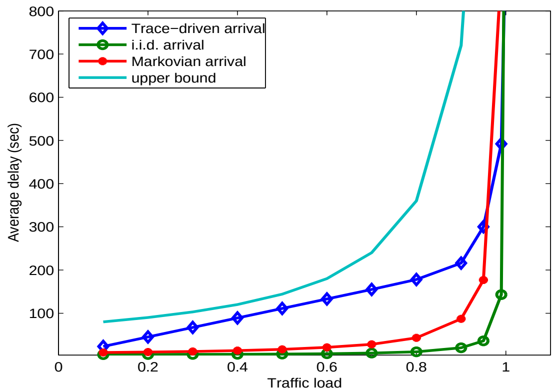

Fig. 1. Delay vs. utilization for MWS-BP with two flows sharing a path with $P = 3$ links for i.i.d., Markovian and trace-driven packet arrival processes.

average delay scales faster with P for MWS-BP, as compared to RS. For single-hop flows, RS is in general non work-conserving, unlike MWS-BP (MWS). However, for multi-hop flows, MWS-BP can be non work-conserving as well, as packets of a flow are forwarded to the next hop only if the backlog difference is positive (or at least non-negative); a temporary overload in a link due to one flow may therefore shut down service of other flows on other links in the network. For instance, consider a link $l = (v, w)$ that is being shared by a "short" (single-hop) flow and a "long" (multi-hop) flow, where link l is not the first hop of the long flow. If the short flow generates packets at a high rate in a given interval, it may not allow the long flow to receive service on link l during that period. The queue buildup of the long flow at the start node v of l will render the difference in backlog for that flow across its previous link say $l' = (u, v)$ to be negative, which will in turn prevent service to the long flow on l'. This effect will propagate upstream and may prevent the long flow from being served (temporarily) on the links from its source to l. This effect is likely to have a more pronounced effect on the delay performance in the overall network when per-flow path-lengths are longer. Since RS allocates pre-determined fractions of slots to different flows in different links, which are calculated based on their average arrival rates, it renders each flow in each link immune to congestion arising from temporary overloads of other flows and other links. It is worth noting however, that such detrimental "chain-effect" arises rarely for MWS-BP and our experiments over a wide range of network topologies and arrival rate processes suggest that MWS-BP generally out-performs RS, even for multi-hop flows, as the simulation results presented next exemplifies (also see [11]).

To study how the average delay compares with the derived delay bound for different arrival processes, we consider a simple topology where two flows share a path with P links, and scheduling must obey primary interference constraints. Other than i.i.d. arrival process, we also consider a more bursty arrival process, modeled as a two-state on-off Markov chain, and adjust the state transition probabilities so that the average length of the on time is 4 times of that in the i.i.d. case, for the same average load. We also consider a representative real data trace measured at an access point from [1], and generate arrival processes at different average loads with similar statistical properties. From Figure 1, which shows the average delays for MWS-BP along with the upper bound (computed according to Theorem 5), we observe that the nature of variation of the average delays with increasing (normalized) traffic load – for the three arrival processes and the delay upper-bound – are largely similar. As expected, the Markovian arrival process leads to higher delay as compared to i.i.d. arrivals due to higher burstiness. Realistic traffic will typically have higher burstiness than i.i.d. arrival process, and we observe this for the traces in [1] as well. In Figure 1, the average delay for trace-driven data is naturally observed to be higher than that for the i.i.d. case. Thus, while our upper bound (derived assuming i.i.d. arrivals) is in general loose as compared to the average delay for i.i.d. arrivals, it may provide a better (closer/tighter) characterization of the average delay for more realistic traffic that is likely to exhibit a higher degree of burstiness.

Acknowledgements

This work was supported by the National Science Foundation through grants CNS-0435141, CNS-0448316, NCR-0238340, CNS-0721308 and ECS-0622176.

References

[1] The CRAWDAD Wireless Database. http://crawdad.cs.dartmouth.edu/. Dartmouth campus, spring02 dataset.

[2] M. Andrews and L. Zhang. Achieving stability in networks of input-queued switches. *IEEE/ACM Tranactions on Networking*, 11(5):848–857, October 2003.

[3] L. X. Bui, R. Srikant, and A. Stolyar. A novel architecture for reduction of delay and queueing structure complexity in the back-pressure algorithm. *IEEE/ACM Transactions on Networking*, 19(6):1597 – 1609, December 2011.

[4] G. Grimmett and D. Stirzaker. *Probability and Random Processes*. Oxford University Press, 3rd edition, 2001.

[5] G. R. Gupta and N. B. Shroff. Scheduling with queue length guarantees for shared resource systems. In *ACM Sigmetrics (Poster Paper)*, Annapolis, MD, June 2008.

[6] G. R. Gupta and N. B. Shroff. Delay analysis for wireless networks with single hop traffic and general interference constraints. *IEEE/ACM Tranactions on Networking*, 18(2):393–405, April 2010.

[7] G. R. Gupta and N. B. Shroff. Delay analysis and optimality of scheduling policies for multihop wireless networks. *IEEE/ACM Transactions on Networking*, 19(1):129–141, February 2011.

[8] B. Hajek and G. Sasaki. Link scheduling in polynomial time. *IEEE Trans. Information Theory*, 34(5):910–917, Sep 1988.

[9] J. Hsu and P. Burke. Behavior of tandem buffers with geometric input and markovian output. *IEEE/ACM Transactions on Communications*, 24(3):358–361, March 1976.

[10] K. Jung and D. Shah. Low delay scheduling in wireless networks. In *Proc. IEEE ISIT*, France, 2007.

[11] K. Kar, X. Luo, and S. Sarkar. Delay guarantees for throughput-optimal wireless link scheduling. In *IEEE INFOCOM 2009*, pages 2331 –2339, April 2009.

[12] E. Leonardi, M. Mellia, F. Neri, and M.A. Marsan. Bounds on delays and queue lengths in input-queued cell switches. *Journal of the ACM*, 50(4):520–550, July 2003.

[13] M. Neely. Delay analysis for maximal scheduling in wireless networks with bursty traffic. In *Proc. IEEE INFOCOM*, Phoenix, AZ, April 2008.

[14] M.J. Neely, E. Modiano, and C-P. Li. Fairness and optimal stochastic control for heterogeneous networks. *IEEE/ACM Transactions on Networking*, 16(2):396 –409, April 2008.

[15] J. Ryu, V. Bhargava, N. Paine, and S. Shakkottai. Back-pressure routing and rate control for icns. In *Proc. ACM MOBICOM*, Chicago, IL, September 2010.

[16] S. Sarkar and S. Ray. Arbitrary throughput versus complexity tradeoffs in wireless networks using graph partitioning. *IEEE Transaction on Automatic Control*, November 2008.

[17] S. Sarkar and L. Tassiulas. A framework for routing and congestion control for multicast information flows. *IEEE Transactions on Information Theory*, 48(10):2690 – 2708, October 2002.

[18] D. Shah. Maximal matching scheduling is good enough. In *Proceedings of Globecom*, December 2003.

[19] L. Tassiulas and A. Ephremidis. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12):1936–1948, Dec 1992.

[20] A. Warrier, S. Janakiraman, S. Ha, and I. Rhee. Diffq: Practical differential backlog congestion control for wireless networks. In *Proc. IEEE INFOCOM*, Rio de Janeiro, Brazil, April 2009.