

Scheduling in Multi-Channel Wireless Networks^{*}

Vartika Bhandari^{**} and Nitin H. Vaidya

University of Illinois at Urbana-Champaign, USA
vartikab@acm.org, nhv@illinois.edu

Abstract. The availability of multiple orthogonal channels in a wireless network can lead to substantial performance improvement by alleviating contention and interference. However, this also gives rise to non-trivial channel coordination issues. The situation is exacerbated by variability in the achievable data-rates across channels and links. Thus, scheduling in such networks may require substantial information-exchange and lead to non-negligible overhead. This provides a strong motivation for the study of scheduling algorithms that can operate with *limited information* while still providing acceptable worst-case performance guarantees. In this paper, we make an effort in this direction by examining the scheduling implications of multiple channels and heterogeneity in channel-rates. We establish lower bounds on the performance of a class of *maximal* schedulers. We first demonstrate that when the underlying scheduling mechanism is “imperfect”, the presence of multiple orthogonal channels can help alleviate the detrimental impact of the imperfect scheduler, and yield a significantly better efficiency-ratio in a wide range of network topologies. We then establish performance bounds for a scheduler that can achieve a good efficiency-ratio in the presence of channels with heterogeneous rates without requiring explicit exchange of queue-information. Our results indicate that it may be possible to achieve a desirable trade-off between performance and information.

1 Introduction

Appropriate scheduling policies are of utmost importance in achieving good throughput characteristics in a wireless network. The seminal work of Tassiulas and Ephremides yielded a *throughput-optimal* scheduler, which can schedule all “feasible” traffic flows without resulting in unbounded queues [8]. However, such an optimal scheduler is difficult to implement in practice. Hence, various imperfect scheduling strategies that trade-off throughput for simplicity have been proposed in [5, 9, 10, 7], amongst others.

The availability of multiple orthogonal channels in a wireless network can potentially lead to substantial performance improvement by alleviating contention and interference. However, this also gives rise to non-trivial channel coordination issues. The situation is exacerbated by variability in the achievable data-rates across channels and links. Computing an optimal schedule, even in a single-channel network, is almost always intractable, due to the need for global information, as well as the computational

^{*} This research was supported in part by NSF grant CNS 06-27074, US Army Research Office grant W911NF-05-1-0246, and a Vodafone Graduate Fellowship.

^{**} Vartika Bhandari is now with Google Inc.

complexity. However, imperfect schedulers requiring limited *local* information can typically be designed, which provide acceptable worst-case (and typically much better average case) performance degradation compared to the optimal. In a multi-channel network, the local information exchange required by even an imperfect scheduler can be quite prohibitive as information may be needed on a per-channel basis. For instance, Lin and Rasool [4] have described a scheduling algorithm for multi-channel multi-radio wireless networks that requires information about *per-channel* queues at all interfering links.

This provides a strong motivation for the study of scheduling algorithms that can operate with limited information, while still providing acceptable worst-case performance guarantees. In this paper, we make an effort in this direction, by examining the scheduling implications of multiple channels, and heterogeneity in channel-rates. We establish lower bounds on performance of a class of *maximal* schedulers, and describe some schedulers that require limited information-exchange between nodes. Some of the bounds presented here improve on bounds developed in past work [4].

We begin by analyzing the performance of a centralized greedy maximal scheduler. A lower bound for this scheduler was established in [4]. However, in a large variety of network topologies, the lower bound can be quite loose. Thus is particularly true for multi-channel networks with single interface nodes. We establish an alternative bound that is tighter in a range of topologies. *Our results indicate that when the underlying scheduling mechanism is imperfect, the presence of multiple orthogonal channels can help alleviate the impact of the imperfect scheduler, and yield a significantly better efficiency-ratio in a wide range of scenarios.*

We then consider the possibility of achieving efficiency-ratio comparable to the centralized greedy maximal scheduler using a simpler scheduler that works with limited information. We establish results for a class of maximal schedulers coupled with local queue-loading rules that do not require queue-information from interfering nodes.

2 Preliminaries

We consider a multi-hop wireless network. For simplicity, we largely limit our discussion to nodes equipped with a single half-duplex radio-interface capable of tuning to any one available channel at any given time. All interfaces in the network have identical capabilities, and may switch between the available channels if desired. Many of the presented results can also be used to obtain results for the case when each node is equipped with multiple interfaces; we briefly discuss this issue.

The wireless network is viewed as a directed graph, with each directed link in the graph representing an available communication link. We model interference using a *conflict* relation between links. Two links are said to conflict with each other if it is only feasible to schedule one of the links on a certain channel at any given time. The conflict relation is assumed to be symmetric. The conflict-based interference model provides a tractable approximation of reality – while it does not capture the wireless channel precisely, it is more amenable to analysis. Such conflict-based interference models have been used frequently in the past work (e.g., [11, 4]).

Time is assumed to be slotted with a slot duration of 1 unit time (i.e., we use slot duration as the time unit). In each time slot, the scheduler determines which links should transmit in that time slots, as well as the channel to be used for each such transmission.

We now introduce some notation and terminology.

The network is viewed as a collection of directed links, where each link is a pair of nodes that are capable of direct communication with non-zero rate.

- \mathcal{L} denotes the set of directed links in the network.
- \mathcal{C} is the set of all available orthogonal channels. Thus, $|\mathcal{C}|$ is the number of available channels.
- We say that a scheduler schedules link-channel pair (l, c) if it schedules link l for transmission on channel c .
- r_l^c denotes the rate achievable on link l by operating link l on channel c , provided that no conflicting link is also scheduled on channel c . For simplicity, we assume that $r_l^c > 0$ for all $l \in \mathcal{L}$ and $c \in \mathcal{C}$.¹ The rates r_l^c do not vary with time. We also define the terms: $r_{max} = \max_{l \in \mathcal{L}, c \in \mathcal{C}} r_l^c$, and $r_{min} = \min_{l \in \mathcal{L}, c \in \mathcal{C}} r_l^c$. When two conflicting links are scheduled simultaneously on the same channel, both achieve rate 0.
- β_s denotes the *self-skew-ratio*, defined as the minimum ratio between rates supportable over *different* channels on a *single* link. Therefore, for any two channels c and d , and any link l , we have $\frac{r_l^d}{r_l^c} \geq \beta_s$. Note that $0 < \beta_s \leq 1$.
- β_c denotes the *cross-skew-ratio*, defined as the minimum ratio between rates supportable over the *same* channel on *different* links. Therefore, for any channel c , and any two links l and l' : $\frac{r_{l'}^c}{r_l^c} \geq \beta_c$. Note that $0 < \beta_c \leq 1$.

Let $r_l = \max_{c \in \mathcal{C}} r_l^c$. Let $\sigma_s = \min_{l \in \mathcal{L}} \frac{\sum_{c \in \mathcal{C}} r_l^c}{r_l}$. Note that $\sigma_s \geq 1 + \beta_s(|\mathcal{C}| - 1)$. Moreover, in typical scenarios, σ_s will be expected to be much larger than this worst-case bound. σ_s is largest when $\beta_s = 1$, in which case $\sigma_s = |\mathcal{C}|$.

- $b(l)$ and $e(l)$, respectively, denotes the nodes at the two endpoints of a link. In particular, link l is directed from node $b(l)$ to node $e(l)$.
- $\mathcal{E}(b(l))$ and $\mathcal{E}(e(l))$ denote the set of links incident on nodes $b(l)$ and $e(l)$, respectively. Thus, the links in $\mathcal{E}(b(l))$ and $\mathcal{E}(e(l))$ share an endpoint with link l . Since we focus on single-interface nodes, this implies that if link l is scheduled in a certain time slot, no other link in $\mathcal{E}(b(l))$ or $\mathcal{E}(e(l))$ can be scheduled at the same time. We refer to this as an *interface conflict*. Let $\mathcal{A}(l) = \mathcal{E}(b(l)) \cup \mathcal{E}(e(l))$. Note that $l \in \mathcal{A}(l)$. Links in $\mathcal{A}(l)$ are said to be *adjacent* to link l . Links that have an interface conflict with link l are those that belong to $\mathcal{E}(b(l)) \cup \mathcal{E}(e(l)) \setminus \{l\}$. Let $A_{max} = \max_l |\mathcal{A}(l)|$.
- $\mathbf{I}(l)$ denotes the set of links that conflict with link l when scheduled on the same channel. $\mathbf{I}(l)$ may include links that also have an interface-conflict with link l . By convention, l is considered included in $\mathbf{I}(l)$. The subset of $\mathbf{I}(l)$ comprising interfering links that are not adjacent to l is denoted by $\mathbf{I}'(l)$, i.e., $\mathbf{I}'(l) = \mathbf{I}(l) \setminus \mathcal{A}(l)$. Let $I_{max} = \max_l |\mathbf{I}'(l)|$.

¹ Though we assume that $r_l^c > 0$ for all l, c , the results can be generalized very easily to handle the case where $r_l^c = 0$ for some link-channel pairs.

- K_l denotes the maximum number of non-adjacent links in $\mathbf{I}'(l)$ that can be scheduled on a given channel simultaneously if l is not scheduled on that channel. $K_l(|C|)$ denotes the maximum number of non-adjacent links in $\mathbf{I}'(l)$ that can be scheduled simultaneously using any of the $|C|$ channels (without conflicts) if l is not scheduled for transmission. Note that here we exclude links that have an interface conflict with l .
- K is the largest value of K_l over all links l , i.e., $K = \max_l K_l$. $K_{|C|}$ is the largest value of $K_l(|C|)$ over all links l , i.e., $K_{|C|} = \max_l K_l(|C|)$. Let $I_{max} = \max_l |\mathbf{I}'(l)|$. It is not hard to see that for *single-interface* nodes:

$$K \leq K_{|C|} \leq \min\{K_{|C|}, I_{max}\} \quad (1)$$

We remark that the term K as used by us is similar, but not exactly the same as the term K used in [4]. In [4], K denotes the largest number of links that may be scheduled simultaneously if some link l is not scheduled, including links adjacent to l . We exclude the adjacent links in our definition of K . Throughout this text, we will refer to the quantity defined in [4] as κ instead of K .

- Let γ_l be 0 if there are no other links adjacent to l at either endpoint of l , 1 if there are other adjacent links at only one endpoint, and 2 if there are other adjacent links at both endpoints.
- γ is the largest value of γ_l over all links l , i.e., $\gamma = \max_l \gamma_l$.
- *Load vector*: We consider single-hop traffic, i.e., any traffic that originates at a node is destined for a next-hop node, and is transmitted over the link between the two nodes. Under this assumption, all the traffic that must traverse a given link can be treated as a single flow. The traffic arrival process for link l is denoted by $\{\lambda(t)\}$. The arrivals in each slot t are assumed i.i.d. with average λ_l . The average load on the network is denoted by load vector $\vec{\lambda} = [\lambda_1, \lambda_2, \dots, \lambda_{|L|}]$, where λ_l denotes the arrival rate for the flow on link l . λ_l may possibly be 0 for some links l .
- *Queues*: The packets generated by each flow are first added to a queue maintained at the source node. Depending on the algorithm, there could be a single queue for each link, or a queue for each (link, channel) pair.
- *Stability*: The system of queues in the network is said to be stable if, for all queues Q in the network, the following is true [2]:

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t E[q(\tau)] < \infty \quad (2)$$

where $q(\tau)$ denotes the backlog in queue Q at time τ

- *Feasible load vector*: In each time slot, the scheduler used in the network determines which links should transmit and on which channel (recall that each link is a directed link, with a transmitter and a receiver). In different time slots, the scheduler may schedule a different set of links for transmission. A load vector is said to be *feasible*, if there exists a scheduler that can schedule transmissions to achieve stability (as defined above), when using that load vector.

- *Link rate vector*: Depending on the schedule chosen in a given slot by the scheduler, each link l will have a certain transmission rate. For instance, using our notation above, if link l is scheduled to transmit on channel c , it will have rate r_l^c (we assume that, if the scheduler schedules link l on channel c , it does not schedule another conflicting link on that channel). Thus, the *schedule* chosen for a time-slot yields a *link rate vector* for that time slot. Note that *link rate vector* specifies rate of transmission used on each link in a certain time slot. On the other hand, *load vector* specifies the rate at which traffic is generated for each link.
- *Feasible rate region*: The set of all feasible load vectors constitutes the feasible rate-region of the network, and is denoted by Λ .
- *Throughput-optimal scheduler*: A *throughput-optimal* scheduler is one that is capable of maintaining stable queues for any load vector $\vec{\lambda}$ in the interior of Λ . For simplicity of notation, we use $\vec{\lambda} \in \Lambda$ in the rest of the text to indicate a load-vector vector λ lying in the interior of a region Λ .
From the work of [8], it is known that a scheduler that maintains a queue for each link l , and then chooses the schedule given by $\text{argmax}_{\vec{r}} \sum_l q_l r_l$, is throughput-optimal for scenarios with single-hop traffic (q_l is the backlog in link l 's queue, and the maximum is taken over all possible link rate vectors \vec{r}). Note that q_l is a function of time, and queue-backlogs at the start of a time slot are used above for computing the schedule (or link-rate vector) for that slot.
- *Imperfect scheduler*: It is usually difficult to determine the throughput-optimal link-rate allocations, since the problem is typically computationally intractable. Hence, there has been significant recent interest in *imperfect* scheduling policies that can be implemented efficiently. In [5], cross-layer rate-control was studied for an imperfect scheduler that chooses (in each time slot) link-rate vector \vec{s} such that $\sum_l q_l s_l \geq \delta \text{argmax}_{\vec{r}} \sum_l q_l r_l$, for some constant δ ($0 < \delta \leq 1$).
It was shown [5] that any scheduler with this property can stabilize any load-vector $\vec{\lambda} \in \delta\Lambda$. Note that if a rate vector $\vec{\lambda}$ is in Λ , then the rate vector $\delta\vec{\lambda}$ is in $\delta\Lambda$. $\delta\Lambda$ is also referred to as the δ -*reduced rate-region*. If a scheduler can stabilize all $\vec{\lambda} \in \delta\Lambda$, its *efficiency-ratio* is said to be δ .
- *Maximal scheduler*: Under our assumed interference model, a schedule is said to be maximal if (a) no two links in the schedule conflict with each other, and (b) it is not possible to add any link to the schedule without creating a conflict (either conflict due to interference, or an interface-conflict).

We will also utilize the Lyapunov-drift based stability criterion from Lemma 2 of [6].

3 Scheduling in Multi-channel Networks

As was discussed previously, throughput-optimal scheduling is often an intractable problem even in a single-channel network. However, imperfect schedulers that achieve a fraction of the stability-region can potentially be implemented in a reasonably efficient manner. Of particular interest is the class of imperfect schedulers known as *maximal schedulers*, which we defined in Section 2. The performance of maximal schedulers

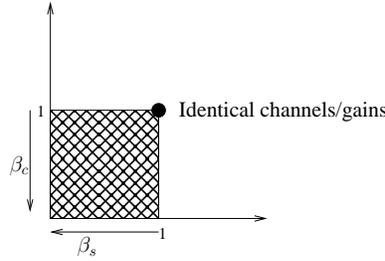


Fig. 1. 2-D visualization of channel heterogeneity

under various assumptions has been studied in much recent work, e.g., [10, 7], with the focus largely on single-channel wireless networks. The issue of designing a distributed scheduler that approximates a maximal scheduler has been addressed in [3], etc.

When there are multiple channels, but each node has one or few interfaces, an additional degree of complexity is added in terms of channel selection. In particular, when the link-channel rates r_l^c can be different for different links l , and channels c , the scheduling complexity is exacerbated by the fact that it is not enough to assign different channels to interfering links; for good performance, the channels must be assigned taking achievable rates into account, i.e., individual channel identities are important.

Scheduling in multi-channel multi-radio networks has been examined in [4], which argues that using a simple maximal scheduler is used in such a network could possibly lead to arbitrary degradation in efficiency-ratio (assuming arbitrary variability in rates) compared to the efficiency-ratio achieved with identical channels. A queue-loading algorithm was proposed, in conjunction with which, a maximal scheduler can stabilize any vector in $(\frac{1}{\kappa+2})\Lambda$, for arbitrary β_c and β_s values. This rule requires knowledge of the length of queues at all interfering links, which can incur substantial overhead.

While variable channel gains are a real-world characteristic that cannot be ignored in designing effective protocols/algorithms, it is important that the solutions not require extensive information exchange with large overhead that offsets any performance benefit. In light of this, it is crucial to consider various points of trade-off between information and performance. In this context, the quantities β_s, β_c and σ_s defined in Section 2 prove to be useful. The quantities β_s and β_c can be viewed as two orthogonal axes for worst-case channel heterogeneity (Fig. 1). The quantity σ_s provides an aggregate (and thus averaged-out) view of heterogeneity along the β_s axis. $\beta_s = 1$ corresponds to a scenario where all channels have identical characteristics, such as bandwidth, modulation/transmission-rate, noise-levels, etc., and the link-gain is a function solely of the separation between sender and receiver. $\beta_c = 1$ corresponds to a scenario where all links have the same sender-receiver separation, and the same conditions/characteristics for any given channel, but the channels may have different characteristics, e.g., an 802.11b channel with a maximum supported data-rate of 11 Mbps, and an 802.11a channel with a maximum supported data-rate of 54 Mbps.

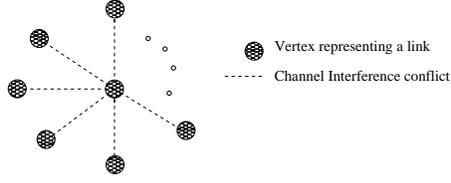


Fig. 2. Example of improved bound on efficiency ratio: link-interference topology is a star with a center link and x radial links

In this paper, we show that in a single-interface network, a simple maximal scheduler augmented with local traffic-distribution and threshold rules achieves an efficiency-ratio at least $\left(\frac{\sigma_s}{K_{|C|} + \max\{1, \gamma\}|C|}\right)$. The noteworthy features of this result are:

1. This scheduler does not require information about queues at interfering links.
2. The performance degradation (compared to the scheduler of [4]) when rates are variable, i.e., $\beta_s, \beta_c \neq 1$, is not arbitrary, and is at worst $\frac{\sigma_s}{|C|} \geq \frac{1 + \beta_s(|C| - 1)}{|C|} \geq \frac{1}{|C|}$. Thus, even with a purely local information based queue-loading rule, it is possible to avoid arbitrary performance degradation even in the worst case. Typically, the performance would be much better.
3. In many network scenarios, the provable lower bound of $\left(\frac{\sigma_s}{K_{|C|} + \max\{1, \gamma\}|C|}\right)$ may actually be better than $\frac{1}{K+2}$. This is particularly likely to happen in networks with single-interface nodes, e.g., suppose we have three channels a, b, c with $r_l^a = 1, r_l^b = 1, r_l^c = 0.5$ for all links l . Then, in the network in Fig. 2 (where the link-interference graph is a star with x radial vertices, and there are no interface-conflicts), $K_{|C|} = x, \gamma = 0, \sigma_s = 2.5$, and we obtain a bound of $\frac{1}{0.4x + 1.2}$, whereas the proved lower bound of the scheduler of [4] is $\frac{1}{x+2}$.

The multi-channel scheduling problem is further complicated if the rates r_l^c are time-varying, i.e., $r_l^c = r_l^c(t)$. However, handling such time-varying rates is beyond the scope of the results in this paper, and we address only the case where rates do not exhibit time-variation. Note that related prior work on multi-channel scheduling [4] also addresses only time-invariant rates.

4 Summary of Results

For multi-channel wireless networks with single-interface nodes, we present lower bounds on the efficiency-ratio of a class of maximal schedulers (including both centralized and distributed schedulers), which indicate that the worst-case efficiency-ratio can be higher when there are multiple channels (as compared to the single-channel case). More specifically, we show that:

- The number of links scheduled by any maximal scheduler are within at least a δ fraction of the maximum number of links activated by any feasible schedule, where:

$$\delta = \max \left\{ \frac{|C|}{K_{|C|} + \max\{1, \gamma\}|C|}, \frac{1}{\max\{1, K + \gamma\}} \right\}$$

- A centralized greedy maximal (CGM) scheduler achieves an efficiency-ratio which is at least $\max\left\{\frac{\sigma_s}{K_{|C|} + \max\{1, \gamma\}|C|}, \frac{1}{\max\{1, K + \gamma\}}\right\}$. This constitutes an improvement over the lower bound for the CGM scheduler proved in [4]. Since $K_{|C|} \leq \min\{K|C|, I_{max}\} \leq \kappa|C|$, this new bound on efficiency-ratio can often be substantially tighter.
- We show that any maximal scheduler, in conjunction with a simple local queue-loading rule, and a threshold-based link-participation rule, achieves an efficiency-ratio of at least $\left(\frac{\sigma_s}{K_{|C|} + \max\{1, \gamma\}|C|}\right)$. This scheduler is of significant interest as it does not require information about queues at all interfering links.

Due to space constraints, proofs are omitted. Please see [1] for the proofs.

Note that the text below makes the natural assumption that two links that conflict with each other (due to interference or interface-conflict) are **not** scheduled in the same timeslot by any scheduler discussed in the rest of this paper.

5 Maximal Schedulers

We begin by presenting a result about the cardinality of the set of links scheduled by any maximal scheduler.

Theorem 1. *Let S_{opt} denote the set of links scheduled by a scheduler that seeks to maximize the number of links scheduled for transmission, and let S_{max} denote the set of links activated by any maximal scheduler. Then the following is true:*

$$|S_{max}| \geq \max\left\{\frac{|C|}{K_{|C|} + \max\{1, \gamma\}|C|}, \frac{1}{\max\{1, K + \gamma\}}\right\} |S_{opt}| \quad (3)$$

The proof is omitted due to lack of space. Please see [1].

6 Centralized Greedy Maximal Scheduler

A centralized greedy maximal (CGM) scheduler operates in the manner described below.

In each timeslot:

1. Calculate link weights $w_l^c = q_l r_l^c$ for all links l and channels c .
2. Sort the link-channel pairs (l, c) in non-increasing order of w_l^c .
3. Add the first link-channel pair in the sorted list (i.e., the one with highest weight) to the schedule for the timeslot, and remove from the list all link-channel pairs that are no longer feasible (due to either interface or interference conflicts).
4. Repeat step 3 until the list is exhausted (i.e., no more links can be added to the schedule).

In [4], it was shown that this centralized greedy maximal (CGM) scheduler can achieve an approximation-ratio which is at least $\left(\frac{1}{\kappa+2}\right)$ in a multi-channel multi-radio network, where κ is the maximum number of links conflicting with a link l that may

possibly be scheduled concurrently when l is not scheduled. This bound holds for arbitrary values of β_s and β_c , and variable number of interfaces per node.

However, this bound can be quite loose in multi-channel wireless networks where each device has one or few interfaces.

In this section, we prove an improved bound on the efficiency-ratio achievable with the CGM scheduler for *single-interface* nodes. We also briefly discuss how it can be used to obtain a bound for multi-interface nodes.

Theorem 2. *Let S_{opt} denote the set of links activated by an optimal scheduler that chooses a set of link-channel pairs (l, c) for transmission such that $\sum w_l^c$ is maximized.*

Let $c^(l)$ denote the channel assigned to link $l \in S_{opt}$ by this optimal scheduler.*

Let S_g denote the set of links activated by the centralized greedy maximal (CGM) scheduler, and let $c^g(l)$ denote the channel assigned to a link $l \in S_g$.

Then:

$$\frac{\sum_{l \in S_g} w_l^{c^g(l)}}{\sum_{l \in S_{opt}} w_l^{c^*(l)}} \geq \max \left\{ \frac{\sigma_s}{K|C| + \max\{1, \gamma\}|C|}, \frac{1}{\max\{1, K + \gamma\}} \right\} \quad (4)$$

The proof is omitted due to lack of space. Please see [1].

Theorem 2 leads to the following result:

Theorem 3. *The centralized greedy maximal (CGM) scheduler can stabilize the δ -reduced rate-region, where:*

$$\delta = \max \left\{ \frac{\sigma_s}{K|C| + \max\{1, \gamma\}|C|}, \frac{1}{\max\{1, K + \gamma\}} \right\}$$

Proof. We earlier discussed a result from [5] that any scheduler, which chooses rate-allocation \vec{s} such that $\sum q_l s_l \geq \delta \operatorname{argmax} \sum q_l r_l$, can stabilize the δ -reduced rate-region. Using Theorem 2 and this result, we obtain the above result.

We remark that the above bound is independent of β_c .

6.1 Multiple Interfaces per Node

We now describe how the result can be extended to networks where each node may have more than one interface.

Given the original network *node-graph* $G = (V, E)$, construct the following transformed graph $G' = (V', E')$:

For each node $v \in V$, if v has m_v interfaces, create m_v nodes v_1, v_2, \dots, v_{m_v} in V' . For each edge $(u, v) \in E$, where u, v have m_u, m_v interfaces respectively, create edges (u_i, v_j) , $1 \leq i \leq m_u, 1 \leq j \leq m_v$, and set $q_{(u_i, v_j)} = q_{(u, v)}$. Set the achievable channel rate appropriately for each edge in E' and each channel. For example, assuming that the channel-rate is solely a function of u, v and c , then: for each channel c , set $r_{(u_i, v_j)}^c = r_{(u, v)}^c$.

The transformed graph G' comprises only single-interface links, and thus Theorem 2 applies to it. Moreover, it is not hard to see that a schedule that maximizes $\sum q_l r_l$ in G' also maximizes $\sum q_l r_l$ in G . Thus, the efficiency-ratio from Theorem 2 for network graph G' yields an efficiency-ratio for the performance of the CGM scheduler in the multi-interface network.

We briefly touch upon how one would expect the ratio to vary as the number of interfaces at each node increases. Note that the efficiency-ratio depends on $\beta_s, |\mathcal{C}|, K_{|\mathcal{C}|}, \gamma$. Of these β_s and $|\mathcal{C}|$ are always the same for both G and G' . γ is also always the same for any G' derived from a given node-graph G , as it depends only on the number of other node-links incident on either endpoint of a node-link in G (which is a property of the node topology, and not the number of interfaces each node has). However, $K_{|\mathcal{C}|}$ might potentially increase in G' as there are many more non-adjacent interfering *links* when each interface is viewed as a distinct node. Thus, for a given number of channels $|\mathcal{C}|$, one would expect the provable efficiency-ratio to initially decrease as we add more interfaces, and then become static.

While this may initially seem counter-intuitive, this is explained by the observation that multiple orthogonal channels yielded a better efficiency-ratio in the single-interface case since there was more spectral resource, but limited hardware (interfaces) to utilize it. Thus, the additional channels could be effectively used to alleviate the impact of sub-optimal scheduling. When the hardware is commensurate with the number of channels, the situation (compared to an optimal scheduler) increasingly starts to resemble a single-channel single-interface network.

6.2 Special Case: $|\mathcal{C}|$ Interfaces per Node

Let us consider the special case where each node in the network has $|\mathcal{C}|$ interfaces, and achievable rate on a link between nodes u, v and all channels $c \in \mathcal{C}$ is solely a function of u, v and c (and not of the interfaces used). In this case, it is possible to obtain a simpler transformation. Given the original network node-graph $G = (V, E)$, construct $|\mathcal{C}|$ copies of this graph, viz., $G_1, G_2, \dots, G_{|\mathcal{C}|}$, and view each node in each graph as having a single-interface, and each network as having access to a single channel. Then each network graph G_i can be viewed in isolation, and the throughput obtained in the original graph is the sum of the throughputs in each graph. From Theorem 2, in each graph we can show that the CGM scheduler is within $\left(\frac{1}{\max\{1, K+\gamma\}}\right)$ of the optimal. Thus, even in the overall network, the CGM scheduler is within $\left(\frac{1}{\max\{1, K+\gamma\}}\right)$ of the optimal.

7 A Rate-Proportional Maximal Multi-Channel (RPMMC) Scheduler

In this section, we describe a scheduler where a link does not require any information about queue-lengths at interfering links.

The set of all links is denoted by \mathcal{L} . The arrival process for link l is i.i.d. over all time-slots t , and is denoted by $\{\lambda_l(t)\}$, with $E[\lambda_l(t)] = \lambda_l$. We make no assumption about independence of arrival processes for two links l, k . However, we consider only

the class of arrival processes for which $E[\lambda_l(t)\lambda_k(t)]$ is bounded, i.e., $E[\lambda_l(t)\lambda_k(t)] \leq \eta$ for all $l \in \mathcal{L}, k \in \mathcal{L}$, where η is a suitable constant.

Consider the following scheduler:

Rate-Proportional Maximal Multi-Channel (RPMMC) Scheduler

Each link maintains a queue for each channel. The length of the queue for link l and channel c at time t is denoted by $q_l^c(t)$. In time-slot t : only those link-channel pairs with $q_l^c(t) \geq r_l^c$ participate, and the scheduler computes a maximal schedule from amongst the participating links. The new arrivals during this slot, i.e., $\lambda_l(t)$ are assigned to channel-queues in proportion to the rates, i.e., $\lambda_l^c(t) = \frac{\lambda_l(t)r_l^c}{\sum_{b \in \mathcal{C}} r_l^b}$

Theorem 4. *The RPMMC scheduler stabilizes the queues in the network for any load-vector within the δ -reduced rate-region, where:*

$$\delta = \frac{\sigma_s}{K_{|C|} + \max\{1, \gamma\}|C|}$$

The proof is omitted due to space constraints. Please see [1].

Corollary 1 *The efficiency-ratio of the RPMMC scheduler is always at least:*

$$\left(\frac{\sigma_s}{|C|}\right) \left(\frac{1}{K + \max\{1, \gamma\}}\right)$$

Proof. The proof follows from Theorem 4 and (1).

8 Discussion

The intuition behind the RPMMC scheduler is simple: by splitting the traffic across channels in proportion to the channel-rates, each link sees the average of all channel-rates as its *effective rate*. This helps avoid worst-case scenarios where the link may end up being repeatedly scheduled on a channel that yields poor rate on that link. The algorithm is made attractive by the fact that no information about queues at interfering links is required. Furthermore we showed that the efficiency-ratio of the RPMMC scheduler is always at least $\left(\frac{\sigma_s}{|C|}\right) \left(\frac{1}{K + \max\{1, \gamma\}}\right)$. Note that $1 + \beta_s(|C| - 1) \leq \sigma_s \leq |C|$. Thus, the efficiency ratio of this algorithm does not degrade indefinitely as β_s becomes smaller. Moreover, in many practical settings, one can expect σ_s to be $\Theta(|C|)$ and the performance would be much better compared to the worst-case of $\sigma_s = 1 + \beta_s(|C| - 1)$.

9 Future Directions

The RPMMC scheduler provides motivation for further study of schedulers that work with limited information. The scheduler of Lin-Rasool [4] and the RPMMC scheduler represent two extremes of a range of possibilities, since the former uses information

from all interfering links, while the latter uses no such information. Evidently, using more information can potentially allow for a better provable efficiency-ratio. However, the nature of the trade-off curve between these two extremities is not clear. For instance, an interesting question to ponder is the following: If interference extends up to M hops, but each link only has information upto $x < M$ hops, what provable bounds can be obtained? This would help quantify the extent of performance improvement achievable by increasing the information-exchange, and provide insights about suitable operating points for protocol design, since control overhead can be a concern in real-world network scenarios.

References

1. V. Bhandari. Performance of wireless networks subject to constraints and failures. Ph.D. Thesis, UIUC, 2008.
2. L. Georgiadis, M. J. Neely, and L. Tassiulas. Resource allocation and cross-layer control in wireless networks. *Found. Trends Netw.*, 1(1):1–144, 2006.
3. C. Joo and N. B. Shroff. Performance of random access scheduling schemes in multi-hop wireless networks. In *Proceedings of IEEE INFOCOM*, pages 19–27, 2007.
4. X. Lin and S. Rasool. A Distributed Joint Channel-Assignment, Scheduling and Routing Algorithm for Multi-Channel Ad-hoc Wireless Networks. In *Proceedings of IEEE INFOCOM*, pages 1118–1126, May 2007.
5. X. Lin and N. B. Shroff. The impact of imperfect scheduling on cross-layer rate control in wireless networks. In *Proceedings of IEEE INFOCOM*, pages 1804–1814, 2005.
6. M. J. Neely, E. Modiano, and C. E. Rohrs. Dynamic power allocation and routing for time varying wireless networks. In *Proceedings of IEEE INFOCOM*, 2003.
7. G. Sharma, R. R. Mazumdar, and N. B. Shroff. On the complexity of scheduling in wireless networks. In *MobiCom '06: Proceedings of the 12th annual international conference on Mobile computing and networking*, pages 227–238. ACM, 2006.
8. L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12):1936–1948, Dec. 1992.
9. X. Wu and R. Srikant. Scheduling efficiency of distributed greedy scheduling algorithms in wireless networks. In *Proceedings of IEEE INFOCOM*, 2006.
10. X. Wu, R. Srikant, and J. R. Perkins. Queue-length stability of maximal greedy schedules in wireless networks. In *Workshop on Information Theory and Applications*, 2006.
11. X. Wu, R. Srikant, and J. R. Perkins. Scheduling efficiency of distributed greedy scheduling algorithms in wireless networks. *IEEE Trans. Mob. Comput.*, 6(6):595–605, 2007.