Throughput Guarantees Through Maximal Scheduling in Wireless Networks

Prasanna Chaporkar, Koushik Kar, Saswati Sarkar

Abstract

We address the question of providing throughput guarantees through distributed scheduling, which has remained an open problem for some time. We consider a simple distributed scheduling strategy, *maximal scheduling*, and prove that it attains a guaranteed fraction of the maximum throughput region in arbitrary wireless networks. The guaranteed fraction depends on "interference degree" of the network which is the maximum number of sessions that interfere with any given session in the network and do not interfere with each other. Depending on the nature of communication, the transmission powers and the propagation models, the guaranteed fraction can be lower bounded by the maximum link degrees in the underlying topology, or even by constants that are independent of the topology. The guarantees also hold in networks with arbitrary number of frequencies. We prove that the guarantees are tight in that they can not be improved any further with maximal scheduling.

I. Introduction

Maximizing the network throughput by appropriately scheduling sessions is a key design goal in wireless networks. Tassiulas *et al.* characterized the maximum attainable throughput region and also provided a scheduling strategy that attains this throughput region in any given wireless network [8]. The policy, however, is centralized and can have exponential complexity depending on the network topology considered. Later, Tassiulas [7] and Shah *et al.* [6] provided linear complexity randomized scheduling schemes that attain the maximum achievable throughput region; both scheduling strategies however require centralized control.

Designing a distributed scheduling policy that attains the throughput region in wireless networks has remained elusive. Recently, Lin *et al.* [4] proved that a distributed maximal matching scheduling strategy is guaranteed to attain at least half of this region for the node-exclusive spectrum sharing model. In the node-exclusive spectrum sharing model, the only scheduling constraint is that a node cannot communicate with multiple nodes simultaneously. This specific interference model holds only when every node has a unique frequency in its two-hop neighborhood.

Different wireless networks have significantly different interference constraints. Bluetooth networks satisfy the node-exclusive spectrum sharing model. On the other hand, IEEE 802.11 networks have limited number of frequencies that may not permit the allocation of unique frequencies in a two-hop neighborhood. Furthermore, the interference regions of nodes involved in transmissions may vary widely depending on the signal propagation conditions, and may be different for different transmitter-receiver pairs. A basic question that remains open is whether a distributed scheduling strategy can attain a guaranteed fraction of the maximum achievable throughput region for arbitrary interference models. Our investigation takes a step forward in solving this open problem.

Our contribution is to characterize the maximum throughput region attained by a distributed scheduling strategy under arbitrary topologies and interference models. The simple scheduling policy we consider, referred to as *maximal scheduling*, only ensures that if a transmitter u has a packet to transmit to a receiver v, either (u, v) or a transmitter-receiver pair that can not simultaneously transmit with (u, v) is scheduled for transmission; the scheduling is otherwise arbitrary. Our investigation of this maximal scheduling policy has been motivated by the following observations. In the specific node-exclusive spectrum sharing model,

P. Chaporkar and S. Sarkar are with the department of Electrical and Systems Engineering at the University of Pennsylvania. K. Kar is with the department of Electrical, Computer and Systems engineering at Renessealer Polytechnic Institute. Their emails are chaporkar@seas.upenn.edu, koushik@ecse.rpi.edu and swati@seas.upenn.edu.

the maximal scheduling policy becomes the maximal matching policy considered by Lin *et. al.*, and is therefore guaranteed to attain at least half of the maximum throughput region [4]. Dai *et. al.* [3] has also obtained a similar guarantee for the maximal matching policy in input-queued switches where the scheduling constraints are similar to that in the node-exclusive spectrum sharing model. Last but not the least, the simplicity and localized nature of maximal scheduling imply that it can be readily implemented in a distributed manner with low overhead and computation cost. It is therefore interesting and important to examine whether maximal scheduling can provide any throughput guarantee under arbitrary interference models and topologies.

We prove that the maximum throughput region attained by maximal scheduling is significantly different for different interference models. We first consider a "bidirectional equal power" interference model in which the network has a single frequency, and all communications use the same power and involve bidirectional message exchanges (e.g., RTS, CTS, data, ACK exchanges in IEEE 802.11). Using a combination of Lyapunov theory and geometric packing, we prove that in this interference model, maximal scheduling is guaranteed to attain at least $1/8$th of the maximum throughput region. This result therefore guarantees that as in the node-exclusive spectrum sharing model, a distributed scheduling can attain a constant fraction of the maximum throughput region in this case as well. Furthermore, we show that the guarantee can not be improved any further in this case as there exists topologies for which maximal scheduling will attain at most $1/8$th of the maximum throughput region.

We next consider a "unidirectional equal power" interference model in which all communications involve unidirectional message exchanges. The network still has a single frequency and all communications use the same power. In this case, however, the performance of maximal scheduling can become arbitrarily bad. More precisely, given any constant Z, there exists topologies in which maximal scheduling will attain less than $1/Z$ of the maximum throughput region. On the other extreme, as discussed before, in the node-exclusive spectrum sharing model, maximal scheduling is guaranteed to attain at least half of the maximum throughput region [4]. We also demonstrate that in this case there exists topologies in which maximal scheduling, and hence maximal matching, will attain at most $1/2$ of the maximum throughput region.

We conclude that a slight variation in the interference constraints may significantly alter the throughput region attained by maximal scheduling (and possibly by other distributed scheduling strategies as well). We can not therefore draw conclusions about the performance of maximal scheduling for arbitrary interference constraints from the results in a few representative scenarios. Also, given that large number of interference models arise, case by case investigations may not be feasible. We therefore proceed to design a framework for characterizing the throughput region of maximal scheduling in arbitrary wireless networks.

We characterize the fraction of the maximum throughput region attained by maximal scheduling in any given topology and interference model. Let $K(\mathcal{N})$ be the maximum interference degree in an arbitrary wireless network $\mathcal{N}$, where the "interference degree" of any transmitter-receiver pair (u, v) is the maximum number of transmitter-receiver pairs that interfere with (u, v) but do not interfere with each other. We prove that maximal scheduling is guaranteed to attain at least $1/K(\mathcal{N})$ of the maximum throughput region in the given network $\mathcal{N}$. Also, there exists an arrival process in the given network $\mathcal{N}$ for which maximal scheduling will attain at most $1/K(\mathcal{N})$ of the maximum throughput region. Given a network, the maximum interference degree may be computed using geometric or graph-theoretic techniques. These results therefore allow us to obtain performance guarantees for maximal scheduling for arbitrary node locations, propagation conditions, interference models and channel allocations.

The paper is organized as follows. We describe the system model and the maximal scheduling policy in Section II. We describe some example communication and interference models in Section III. We characterize the throughput regions of maximal scheduling for some representative interference models in Section IV, and for arbitrary wireless networks in Section V. We conclude in Section VI. We present the proofs in [2].

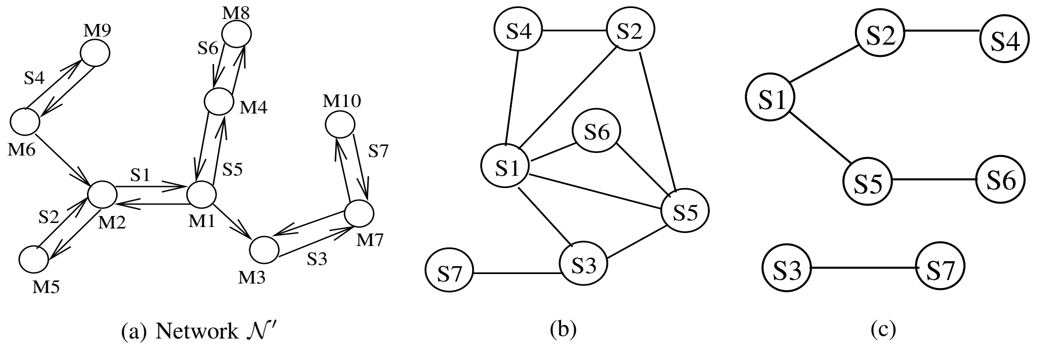

Fig. 1. Panel (a) shows a directed graph with $V = \{M1, \ldots, M10\}$. The arrows between the nodes indicate the directed links. There are 7 sessions: $S1, \ldots, S7$. Nodes $M2, M5, M3, M6, M1, M8$ and $M10$ are the transmitters of sessions $S1, S2, S3, S4, S5, S6$ and $S7$, respectively. Node $M2$ has 3 neighbors: $M1$, $M5$, $M6$. Nodes $M1$ and $M2$ have degree 5; hence the degrees of edges $(M1, M2)$ and $(M2, M1)$ are 10. Here, $\delta_G = 10$. Both the out-degree of $M1$ and in-degree of $M2$ are 3. Thus, the directed degree of $(M1, M2)$ is 6. Here, $\Delta_G = 6$. Sessions $S5$ and $S6$ interfere with each other, as $M4$ has a single transceiver.
Panels (b) and (c) show the interference graphs for the network shown in (a) under bidirectional and unidirectional communication models, respectively. As panels (b) and (c) show, the interference sets of $S6$ are $\{S1, S5\}$ and $\{S5\}$ under the bidirectional and unidirectional communication models, respectively.

II. SYSTEM MODEL

We consider scheduling at the MAC layer in a wireless network. We assume that time is slotted. The topology in 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. 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.

We now introduce terminologies that we use throughout the paper. Some of these are well-known in graph theory; we mention these for completeness.

Definition 1: A node i is a *neighbor* of a node j, if there exists a link from i to j, i.e., $(i, j) \in E$.

The *degree of a node* u is the number of links in E originating from or ending at u. The *degree of a link* $e = (u, v)$ is defined as the sum of the degrees of u and v. The *maximum link degree* in G, δ_G, is the maximum degree of any link in E.

The *out-degree* of a node u is the number of links in E originating from u. The *in-degree* of a node u is the number of links in E ending at u. The *directed degree of a link* $e = (u, v)$ is defined as the sum of the out-degree of u and in-degree of v. The *maximum directed link degree* in G, Δ_G, is the maximum directed degree of any link in E.

At the MAC layer, each session traverses only one link. If a session i traverses link (u, v) then u and v are i's transmitter and receiver respectively, and the session is completely specified by the $3-$tuple, (i, u, v). Multiple sessions may traverse the same link. Without loss of generality, we assume that every node in V is either the transmitter or the receiver of at least one session. If this assumption does not hold, we can consider G to be a subgraph obtained from the original topology by removing the nodes that are not the end points of sessions.

Definition 2: A session i *interferes* with session j if j can not successfully transmit a packet when i is transmitting.

In Section III, we will describe broad classes of communication and interference models and how to obtain the pairwise interference relations in each case.

A wireless network $\mathcal{N}$ can be described by the topology $G = (V, E)$, the $3-$tuple specifications of the sessions and the pair-wise interference relations between the sessions. We consider a network with N sessions.

Definition 3: The *interference set* of a session i, S_i, is the set of sessions j such that either i interferes with j or j interferes with i.
Note that if $j \in S_i$, then $i \in S_j$.

Definition 4: 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 sessions in $\mathcal{N}$ and there is an edge between two vertices i and j if $j \in S_i$.

We elucidate these definitions through examples in Fig. 1.

We now describe the arrival process. Let $A_i(n)$ be the number of packets that session i generates in interval $(0, n)$, $i = 1, \ldots, N$. We assume that any packet arriving in a slot arrives at the beginning of the slot, and may be transmitted in the slot. The arrival process $\{A_i(.), i = 1, \ldots, N\}$ satisfies a strong law of large numbers (SLLN). Thus, there exists non-negative real numbers $\lambda_i, i = 1, \ldots, N$ such that with probability 1,

$$\lim_{n \to \infty} A_i(n)/n = \lambda_i, \;\; i = 1, \ldots, N. \tag{1}$$

The condition (1) on the arrival processes is mild. Several arrival processes including all jointly stationary and ergodic arrival processes satisfy (1).

Definition 5: The *arrival rate* of session i is λ_i, $i = 1, \ldots, N$. The *arrival rate vector* $\vec{\lambda}$ is an $N-$dimensional vector whose components are the arrival rates.

Definition 6: A *scheduling policy* is an algorithm that decides in each slot the subset of sessions that would transmit packets in the slot.
Clearly, a subset S of sessions can transmit packets in any slot if no two sessions in S interfere with each other and every session in S has a packet to transmit. Every packet has length 1 slot. Thus, if a session is scheduled in a slot, it transmits a packet in the slot.

Let $D_i(n)$ be the number of packets that session i transmits in interval $(0, n)$, $i = 1, \ldots, N$. Clearly the transmissions depend on the scheduling policy.

Definition 7: The network is said to be *stable* if with probability 1,

$$\lim_{n \to \infty} D_i(n)/n = \lambda_i, \;\; i = 1, \ldots, N. \tag{2}$$

Thus, a network is stable if the arrival and departures rates are equal for each session.

Definition 8: The *throughput region* of a scheduling policy is the set of arrival rate vectors $\vec{\lambda}$ such that the network is stable under the policy for any arrival process that satisfies (1) and has arrival rate vector $\vec{\lambda}$.

Definition 9: An arrival rate vector $\vec{\lambda}$ is said to be *feasible* if it is in the throughput region of some scheduling policy.

Definition 10: The *maximum throughput region* Λ is the set of feasible arrival rate vectors.
Note, Λ depends on the network $\mathcal{N}$.

Example 1: Consider the network shown in Fig. 2(a). Consider a scheduling policy π_1, that serves session $t \bmod 9 + 1$ in slot t, where "mod" is a modulo operator. Under π_1, each session $i \in \{1, \ldots, 9\}$ can transmit at the rate of at most $1/9$. Thus, the throughput region of π_1, Λ^{π_1}, is characterized as follows:

$$\Lambda^{\pi_1} = \{(\lambda_1, \ldots, \lambda_9) \; : \; \lambda_i \leq 1/9 \; \forall \, i\}.$$

In this case, the maximum throughput region Λ is given by

$$\Lambda = \left\{(\lambda_1, \ldots, \lambda_9) \; : \; \lambda_1 + \max_{2 \leq i \leq 9}\{\lambda_i\} \leq 1\right\}.$$

Therefore, in this example, scheduling policy π_1 achieves only a small fraction of the maximum throughput region.

We now describe the "maximal scheduling" policy we consider. This policy schedules a subset S of sessions such that (i) every session in S has a packet to transmit, (ii) no session in S interferes with any other session in S, (iii) if a session i has a packet to transmit, then either i or a session in

S_i, is included in S. Clearly, many subsets of sessions satisfy the above criteria in each slot, e.g., in Fig. 1(b), $\{S1, S7\}, \{S2, S3, S6\}$ satisfy the above criteria in any slot in which all sessions have packets to transmit. Maximal scheduling can select any such subset. If each session knows its interference set, maximal scheduling can be implemented in distributed manner using standard algorithms [5]. In most cases of practical interest, sessions can determine their interference sets using local message exchange.

III. Interference Models

The pairwise interference relations between the sessions depend on topology $G = (V, E)$ and the nature of communication. The topology G is determined by the transmission powers, propagation conditions and node locations. Communication can either be bidirectional or unidirectional. In the former, when a session is scheduled, both the transmitter and the receiver transmit sequentially. For example, the transmitter may transmit data and control messages while the receiver may transmit control messages. Such bidirectional communications occur in IEEE 802.11. Thus, there must be links in both directions between a session's transmitter and receiver. In unidirectional communication, when a session is scheduled, it transmits packets from only the transmitter to the receiver. For example, unidirectional communication occurs in IEEE 802.11 when control messages are disabled (e.g., in broadcast mode). Different combinations of these conditions lead to different interference relations. We next characterize the pair-wise interference relations for some of these combinations.

We assume that each node has a single transceiver. Thus a node can be involved in at most one transmission. In other words, sessions that have a node in common interfere with each other. We initially assume that all transmissions use the same frequency. Thus, node j can not receive any packet successfully if more than one of its neighbors are transmitting simultaneously (we do not assume any capture). Thus, a transmission on link $(i, j) \in E$ is successful in a slot if and only if no neighbor of j other than i transmits in the slot. For example, in Fig. 1(a), transmission along $(M5, M2)$ is successful if $M1$ and $M6$ do not transmit. For bidirectional communication, when a session (i, u, v) is scheduled, transmissions proceed along both (u, v) and (v, u). For unidirectional communication, when a session (i, u, v) is scheduled, transmissions proceed only along (u, v). The above constraints provide the interference relations for both the bidirectional and unidirectional communication models.

In the *bidirectional communication model*, a session i interferes with session j if i and j have a common end point, or one end point (transmitter or receiver) of j is a neighbor of an end point of i. For example, in Fig. 1(a), $S1, S5, S7$ interfere with $S3$. This is also clearly evident from Fig. 1(b). In the *unidirectional communication model*, session i interferes with session j if i and j have a common end point, or j's receiver is a neighbor of i's transmitter. For example, in Fig. 1(a), only $S7$ interferes with $S3$. Observe that the interference relations may be asymmetric, i.e., i may interfere with j but j may not interfere with i. For example, under bidirectional communication model, in Fig. 1(a), $S1$ interferes with $S3$ but $S3$ does not interfere with $S1$.

We now describe several important special cases. First assume that the propagation conditions are identical in all directions. Each node transmits at a fixed power level which can be different for different nodes. The power level of a node u determines its transmission range, and all nodes within u's transmission range receive u's signal. Thus, the link set E has the following structure: a link exists from u to v if and only if the distance between u and v is less than or equal to u's transmission range. In the bidirectional communication model, session i interferes with session j if one end point of j is within the transmission range of an end point of i. In the unidirectional communication model, session i interferes with session j if j's receiver is within the transmission range of i's transmitter.

Let us further assume that all nodes transmit at the same power. Thus, all nodes have the same transmission range d which is determined by the transmission power. Now, the link set E has the following structure: a link exists from u to v if and only if the distance between u and v is less than d. Now, in the bidirectional communication model, a session i interferes with session j if one end point of j is within distance d from an end point of i (*bidirectional equal power model*). In the unidirectional interference model, a session i interferes with session j if j's receiver is within distance d from i's transmitter (*unidirectional equal power model*). Refer to Fig. 2(a) and (b) for examples of both cases. Note that now the interference relation is symmetric in the bidirectional communication model, i.e., if node i interferes with node j, then node j also interferes with node i. However, interference relationships could still be asymmetric in the unidirectional communication model.

We also consider multi-channel networks. We assume that the network has a large number of frequencies such that every node has a unique frequency in its two-hop neighborhood. Now, for both bidirectional and unidirectional communications, only the sessions that have common end point interfere. This model arises in Bluetooth communications, and is commonly referred to as the *node-exclusive spectrum sharing model* (Fig. 3). The framework and the analytical results for arbitrary interference models however extend to the more general case where the network has an arbitrary number of frequencies [2].

We observe that the pairwise interference relations are significantly different in each of the cases discussed above. There is however one important similarity. If session i interferes with another session j, the distance between the transmitters of i and j is at most three hops. Thus, a session can use local message exchange to determine its interference set. Hence, maximal scheduling can be implemented in distributed manner in each of these cases. But, given the significant difference between the interference relations, it is not clear how similar the performance of maximal scheduling will be in these different cases. In the next section, we assess this difference by characterizing the throughput region of maximal scheduling in a few representative scenarios.

IV. Performance of Maximal Scheduling for Specific Interference Models

We characterize the throughput regions of maximal scheduling Λ^{MS} for some representative interference models. We focus on the bidirectional equal power (Subsection IV-A) and unidirectional equal power models (Subsection IV-B). We subsequently compare the throughput regions obtained in these cases with that in the well-investigated node-exclusive spectrum sharing model (Subsection IV-C). We conclude that the throughput regions are significantly different in different cases.

A. Throughput region of maximal scheduling for bidirectional equal power model

Lemmas 1 and 2 show that in the bidirectional equal power model, Λ^{MS}, is $1/8$th of the maximum throughput region Λ.

Lemma 1: For the bidirectional equal power model, if $\vec{\lambda} \in \Lambda$, $\vec{\lambda}/8 \in \Lambda^{MS}$.

We describe the intuition behind the result. Let arrival rate vector $\vec{\lambda} \in \Lambda$. Then, from (2), under some scheduling policy the packet arrival rate λ_j for each session j equals j's departure rate. Thus, for each session i, the sum of its arrival rate and the arrival rates of the sessions in its interference set S_i must equal the sum of the corresponding departure rates.

In each slot, either i or one or more sessions in S_i may transmit packets, but i can not simultaneously transmit with any session in S_i. We prove using geometry that for any i at most 8 sessions in S_i can simultaneously transmit packets [2]. Thus in any slot at most 8 packets can be transmitted by sessions in $\{i\} \cup S_i$. Thus, the sum of the departure rates of sessions in $\{i\} \cup S_i$, and hence the sum of the corresponding arrival rates, is at most 8. Thus, when the arrival rate vector is $\vec{\lambda}/8$ instead of $\vec{\lambda}$, the sum of the arrival rates of sessions in $\{i\} \cup S_i$ is at most 1.

Let the arrival rate vector be $\vec{\lambda}/8$, and let maximal scheduling be used. For any session i, maximal scheduling always serves 1 packet from either i or a session in S_i in any slot in which i has a packet to transmit. Thus, whenever i has a packet to transmit, the sum of the departure rates for these sessions is 1, which is greater than or equal to the sum of the arrival rates of these sessions. Now, since the departure rate of any session j cannot exceed j's arrival rate, for all i, the sum of the departure rates from the sessions in $\{i\} \cup S_i$ equals the sum of the corresponding arrival rates. It follows that the departure rate of each session i equals i's arrival rate. Thus, the system is stable. Hence $\vec{\lambda}/8 \in \Lambda^{\text{MS}}$.

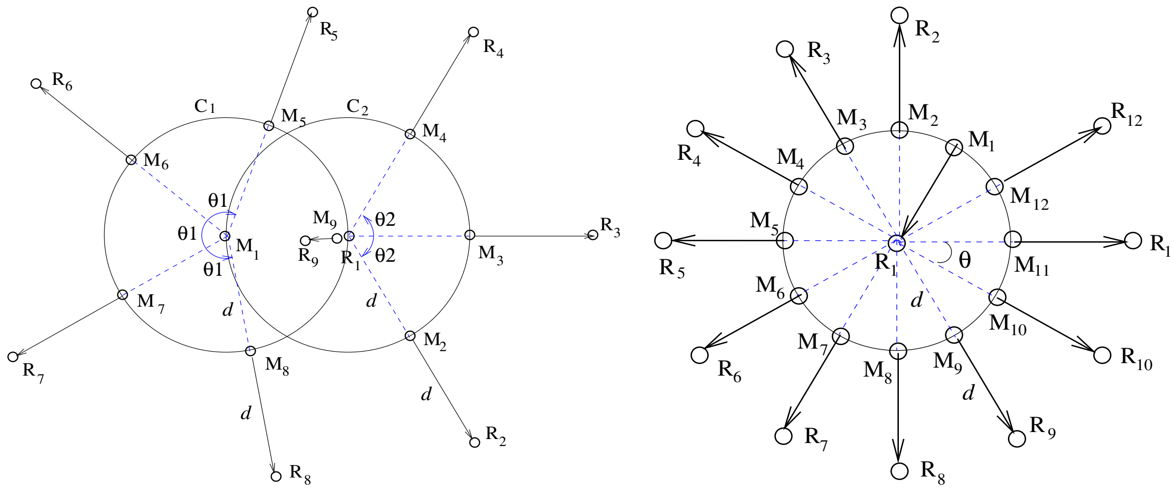

(a) (b)

Fig. 2. Fig. (a) shows a network with interference constraints given by the bidirectional equal power model and transmission range d. There are 9 sessions: $1, \ldots, 9$. Session i has transmitter M_i and receiver R_i. The interference area of session 1 is the union of circles C_1 and C_2. Here, $\theta_1 = 70$ deg, and $\theta_2 = 61$ deg. Distance between (i) M_i and R_i is d for every $i = 1, \ldots, 8$, (ii) M_9 and R_9 is $\epsilon > 0$, where ϵ is a small positive number, (iii) M_1 and M_i is d for every $i = 2, \ldots, 9$, (ii) M_j and M_k is greater than d for every $j, k \in \{2, \ldots, 9\}, j \neq k$ and (iv) M_9 and R_1 is ϵ. *Thus, session 1 interferes with all the other* 8 *sessions, but none of the other sessions interfere with each other.* Fig. (b) shows a network with interference constraints given by the unidirectional equal power model and transmission range d. There are 12 sessions: $1, \ldots, 12$. Session i has transmitter M_i and receiver R_i. The distance between M_i and R_i, and R_1 and M_i is d for every i. Thus, session 1 interferes will all the other 11 sessions, but none of the other sessions interfere with each other. We refer to sessions $2, \ldots, 12$ as non-interfering sessions. Here, θ is $\pi/6$. Note that $2\pi/\theta - 1$ non-interfering sessions can be accommodated. Thus, for any given Z, $Z+1$ non-interfering sessions can be accommodated by choosing $\theta = 2\pi/(Z+2)$.

We now describe why for any i at most 8 sessions in S_i can simultaneously transmit packets. From the interference constraints, at least one end point of each session in S_i must be within a distance d from either i's transmitter or i's receiver. Also, the distance between i's transmitter and receiver is at most d. Thus, at least one end point of each session in S_i must be in the union of two circles of radius d and centered around i's transmitter and receiver respectively (Fig. 2(a)). We refer to the area in this union as i's *interference area*. We prove using geometric arguments that at most 8 points can be present in this interference area such that the distance between any two points exceeds d. Clearly, if sessions j and k need to simultaneously transmit packets, the distance between an end point of j and an end point of k must exceed d. The result follows.

Lemma 2: Consider an arbitrary positive constant Z such that $Z < 8$. For the bidirectional equal power model, there exists a network $\mathcal{N}$ and an arrival rate vector $\vec{\lambda}$, such that $\vec{\lambda} \in \Lambda$ in $\mathcal{N}$, but $\vec{\lambda}/Z \notin \Lambda^{MS}$ in $\mathcal{N}$.

We present the intuition behind the result. Using geometry, we first demonstrate that it is possible to obtain a network with 9 sessions where one session (session 1) interferes with all other sessions and none of the other sessions interfere with each other (Fig. 2(a)). In such a network, consider an arrival rate vector $\vec{\lambda}$ in which session 1 generates packets at the rate λ_1 and session i generates packets at the rate λ_2 for all $i \in \{2, \ldots 9\}$. Consider λ_1, λ_2 such that $\lambda_1 + \lambda_2 = 1$, $\lambda_2 = Z/8$. The system can be stabilized by scheduling session 1 in λ_1 fraction of slots and the other sessions in the remaining slots. Thus, $\lambda \in \Lambda$. Now consider arrival rate vector $\vec{\lambda}/Z$. Let session 1 generate packets at the rate λ_1/Z and session i generate packets at the rate λ_2/Z for all $i \in \{2, \ldots 9\}$. Furthermore, sessions $2, 3, \ldots 9$ generate packets in non-overlapping slots (Fig. 4). Thus, since $\lambda_2/Z = 1/8$, in each slot at least one session in $\{2, \ldots 9\}$ generates packets, and therefore has a packet to transmit. Let maximal scheduling schedule session 1 only when other sessions do not have packets. But then session 1 is never scheduled. Since session 1's arrival rate is positive, the system is not stable under maximal scheduling. Thus, $\vec{\lambda}/Z \notin \Lambda^{\mathrm{MS}}$.

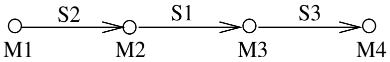

Fig. 3. Figure shows a network with 4 nodes $M1, \ldots, M4$ and 3 sessions $S1, S2$ and $S3$. Under node exclusive spectrum sharing model, $S1$ interferes with both $S2, S3$, but $S2$ and $S3$ do not interfere with each other.

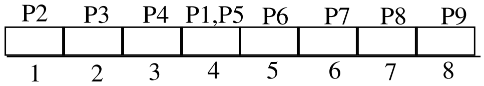

Fig. 4. Figure shows a periodic arrival process for the network in Fig. 2(a). The period is 8 slots. Session S_i generates a packet in the slot marked Pi, for each i. Here, $\lambda_i = 1/8$, for each i.

B. Throughput region of maximal scheduling for unidirectional equal power model

We now consider the unidirectional equal power model. We prove that maximal scheduling can not attain a constant fraction of the maximum throughput region.

Lemma 3: Consider an arbitrary positive constant Z. For the unidirectional equal power model, there exists a network $\mathcal{N}$, an arrival rate vector $\vec{\lambda}$, such that $\vec{\lambda} \in \Lambda$ in $\mathcal{N}$, but $\vec{\lambda}/Z \notin \Lambda^{MS}$ in $\mathcal{N}$.

We present the intuition behind the result. We could obtain the throughput guarantee of $1/8$ in the bidirectional equal power model irrespective of the network primarily because for this model in any network the interference set of any session consists of at most 8 sessions that can transmit simultaneously. This no longer holds for the unidirectional equal power model. In fact, for unidirectional equal power model given any constant Z we can construct a network where the interference set of a session consists of $\lceil Z + 1 \rceil$ sessions that can transmit simultaneously (Fig. 2(b)). We can prove that in such a network there exists an arrival rate vector $\vec{\lambda}$, such that $\vec{\lambda} \in \Lambda$, but $\vec{\lambda}/Z \notin \Lambda^{\text{MS}}$. The intuition behind the proof of this part is similar to that for Lemma 2.

C. Throughput region of maximal scheduling for node-exclusive spectrum sharing model

The throughput regions of maximal scheduling are significantly different for the bidirectional and unidirectional power models. We next mention the guarantees for a third interference model, the node-exclusive spectrum sharing model, and then contrast the guarantees in the three cases. We need the following concepts that are well-known in graph theory. Consider a graph $G' = (V, E')$ where E' consists of only those links in E that are traversed by sessions

Definition 11: A *matching* in G' is a set of links such that no two links have a common node.

Definition 12: A *maximal matching* is a matching in G' such that for any link e that is traversed by a session that has a packet to transmit, either e is in the matching or a link that has a common node with e is in the matching.

In the node-exclusive spectrum sharing model, maximal scheduling always selects sessions that constitute a maximal matching in G'. This follows from the definition of maximal scheduling and because of the pair-wise interference relations in the node-exclusive spectrum sharing model. Lin *et. al.* [4] has proved that maximal matching attains at least $1/2$ the maximum throughput region in the node-exclusive spectrum sharing model. Thus, maximal scheduling also attains at least $1/2$ the maximum throughput region. We would like to remark that in this model in any network the interference set of any session consists of at most 2 sessions; therefore, in this case too, the throughput guarantee seems to be related to this quantity.

We next prove that there exists networks where maximal scheduling attains at most $1/2$ the maximum throughput region.

Lemma 4: Consider an arbitrary positive constant Z such that $Z < 2$. For the node-exclusive spectrum sharing model, there exists a network and an arrival rate vector $\vec{\lambda}$, such that $\vec{\lambda} \in \Lambda$ in $\mathcal{N}$, but $\vec{\lambda}/Z \notin \Lambda^{MS}$ in $\mathcal{N}$.

We present the intuition behind this result. We construct a network with 3 sessions where one session (session 1) interferes with all other sessions and none of the other sessions interfere with each other (Fig. 3). Like for Lemma 2, we can prove that in such a network there exists an arrival rate vector $\vec{\lambda}$, such that $\vec{\lambda} \in \Lambda$, but $\vec{\lambda}/Z \notin \Lambda^{\text{MS}}$.

The throughput regions for the maximal scheduling in the node-exclusive spectrum sharing model are again significantly different from those in the bidirectional and unidirectional equal power models. We conclude that these guarantees will critically depend on the interference relations. Furthermore, the differences between the characterizations obtained for the bidirectional and the unidirectional interference models demonstrate that slight changes in interference conditions can significantly alter the guarantees. We can not therefore draw conclusions about the performance under different models from the results in a few representative scenarios. Also, given that large number of interference relations exist, case by case investigations may not be feasible. We therefore need a framework for characterizing the throughput region of maximal scheduling in arbitrary wireless networks.

V. Performance guarantees of maximal scheduling in arbitrary networks

We design a framework for characterizing the throughput region of maximal scheduling Λ^{MS} for an arbitrary wireless network.

We first introduce the notion of "interference degree" for sessions.

Definition 13: The *interference degree of a session* i is (i) the maximum number of sessions in its interference set S_i that can simultaneously transmit, if S_i is non-empty and (ii) 1 if S_i is empty.

The interference degrees depend on the links traversed by the sessions and the topology $G = (V, E)$ which in turn depends on the node locations, propagation conditions and interference models. For example, in Fig. 1(b), $S_{S1} = \{S2, S3, S4, S5, S6\}$, and the largest set of sessions in S_{S1} that can simultaneously transmit is $\{S3, S4, S6\}$. Thus, the interference degree of $S1$ is 3.

The characterizations of Λ^{MS} obtained so far for specific interference models are closely related to the maximum interference degrees of sessions under these models. For example, for both the bidirectional equal power and the node-exclusive spectrum sharing models maximal scheduling attains exactly $1/P$ fraction of the maximum throughput region Λ, where P is the maximum interference degree of any session. In the unidirectional equal power model, we observed that maximal scheduling can not attain a constant fraction of Λ and also that a session can have arbitrarily large interference degree. We now prove that this relation is not a coincidence but reflects a general phenomenon that holds for arbitrary networks.

We first define the interference degree of a network.

Definition 14: The *interference degree of a network* $\mathcal{N}$, $K(\mathcal{N})$, is the maximum interference degree of sessions in the network.

In Fig. 1(b) and (c), the interference degrees of the network are 3 and 2 respectively. Session $S1$ has these interference degrees in both cases.

Theorem 1: In any wireless network $\mathcal{N}$, if $\vec{\lambda} \in \Lambda$ in $\mathcal{N}$, $\vec{\lambda}/K(\mathcal{N}) \in \Lambda^{MS}$ in $\mathcal{N}$.

Theorem 2: Consider an arbitrary wireless network $\mathcal{N}$ and a constant Z such that $Z < K(\mathcal{N})$. There exists an arrival rate vector $\vec{\lambda}$ such that $\vec{\lambda} \in \Lambda$ in $\mathcal{N}$, but $\vec{\lambda}/Z \notin \Lambda^{MS}$ in $\mathcal{N}$.

The intuition behind Theorems 1 and 2 are similar to that for Lemmas 1 and 2 respectively. The generalization here is that we obtain the characterizations in terms of $K(\mathcal{N})$ because for any session i at most $K(\mathcal{N})$ sessions in S_i can transmit simultaneously.

Theorems 1 and 2 allow us to characterize Λ^{MS} under arbitrary interference relations, node locations, edge sets, session configurations and propagation models, provided we can compute $K(\mathcal{N})$ in these cases. We now obtain upper bounds for $K(\mathcal{N})$ for arbitrary bidirectional and unidirectional communications models, in terms of the maximum link degrees δ_G and Δ_G in the underlying topology G. We also prove that the bounds are tight as there exists networks $\mathcal{N}$ where $K(\mathcal{N})$ equals these bounds. These bounds and the resulting characterizations of Λ^{MS} hold even when different nodes use different transmission powers and propagation conditions in different directions are different.

Lemma 5: In a wireless network $\mathcal{N}$ with bidirectional communication and underlying topology $G = (V, E)$, $K(\mathcal{N}) \leq \max(\delta_G - 4, 1)$. Moreover, there exists a wireless network $\mathcal{N}_1$ with bidirectional communication and underlying topology $G = (V, E)$, such that $K(\mathcal{N}_1) = \max(\delta_G - 4, 1)$.

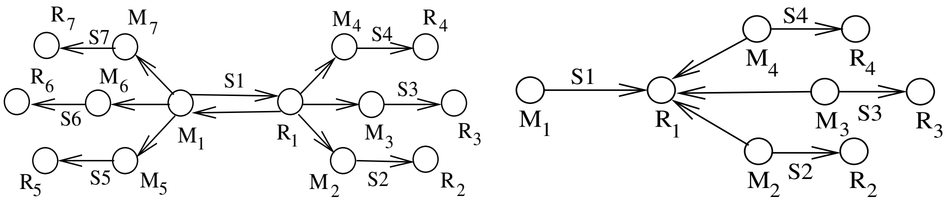

(a) Network $\mathcal{N}_1'$ (b) Network $\mathcal{N}_2'$

Fig. 5. Fig. (a) shows a network $\mathcal{N}_1'$ with bidirectional communication model and 7 sessions: $(S1, M_1, R_1), \ldots, (S7, M_7, R_7)$. Session $S1$ interferes with all the remaining sessions, and none of the remaining sessions interferes with each other. Thus, $K(\mathcal{N}_1') = 6$. The degree of (M_1, R_1) is 10, which is also equal to δ_G. Thus, $K(\mathcal{N}_1') = \delta_G - 4 = \max(\delta_G - 4, 1)$.
Fig. (b) shows a network $\mathcal{N}_2'$ with unidirectional communication model and four sessions: $(S1, M_1, R_1), \ldots, (S4, M_4, R_4)$. Sessions $S2, S3$ and $S4$ interfere with $S1$, but not with each other. Thus, $K(\mathcal{N}_2') = 3$. The directed degree of (M_1, R_1) is 5, which is also equal to Δ_G. Thus, $K(\mathcal{N}_2') = \Delta_G - 2 = \max(\Delta_G - 2, 1)$. In both figures, arrows indicate directed links between the nodes.

The upper-bound in Lemma 5 follows because for bidirectional communication the interference-degree of each session (i, u, v) is at most the degree of link (u, v) minus 4, and the degree of any link in G is at most δ_G. Fig. 5(a) shows an example network $\mathcal{N}_1'$ where $K(\mathcal{N}_1') = \delta_G - 4$; the bound is therefore tight.

Lemma 6: In a wireless network $\mathcal{N}$ with unidirectional communication and underlying topology $G = (V, E)$, $K(\mathcal{N}) \leq \max(\Delta_G - 2, 1)$. Moreover, there exists a wireless network $\mathcal{N}_1$ with unidirectional communication and underlying topology $G = (V, E)$, such that $K(\mathcal{N}_1) = \max(\Delta_G - 2, 1)$.

Lemmas 5 and 6 are similar. Due to unidirectional communication, the bound in Lemma 6 however depends on Δ_G instead of δ_G. Fig. 5(b) provides an example to illustrate the tightness.

Theorems 1 and 2 explain the characterizations of Λ^{MS} for all the specific interference models considered so far. For the bidirectional equal power model, for any network $\mathcal{N}$, $K(\mathcal{N}) \leq 8$ [2], and there exists a network $\mathcal{N}_1$ where $K(\mathcal{N}_1) = 8$ (Fig. 2(a)). Thus, Lemmas 1 and 2 follow as special cases of Theorems 1 and 2. For the unidirectional equal power model, given any Z a network can be constructed so as to attain the interference degree $Z + 1$ (Fig. 2(b)). Theorem 2 now explains Lemma 3. Theorem 1 also explains the throughput characterization for the maximal matching policy in the node-exclusive spectrum sharing model obtained by Lin *et. al* [4]. In this model, for any network $\mathcal{N}$, $K(\mathcal{N}) \leq 2$. Also, there exists a network $\mathcal{N}_1$ with $K(\mathcal{N}_1) = 2$ (Fig. 3). Thus, the throughput guarantee of $1/2$ obtained in this case follows as a special case of Theorem 1, and Lemma 4 follows as a special case of Theorem 2.

The characterizations of Λ^{MS} for specific interference models are often obtained for the worst network under the interference model. This observation applies to all results obtained in Section IV and also the guarantees obtained by Lin *et. al* [4]. Theorems 1 and 2 allow the guarantees to cater to specific networks, and therefore often provide better guarantees. For example, Lemma 3 states that for the unidirectional equal power model, given a constant, there exists topologies where the throughput region of maximal scheduling is less than that constant fraction of the maximum throughput region. But, Lemma 6* shows that even in this model maximal scheduling attains a guaranteed fraction of the maximum throughput region; Lemma 6 shows that the guarantee however depends on the degrees in the underlying topology G. Although in the worst case, these degrees can be arbitrarily large (and therefore the guaranteed fraction can not be lower bounded by a constant in the worst case), these degrees are usually small. Thus, for several topologies Lemma 6 guarantees an acceptable performance even for the unidirectional equal power model. Similarly, for the bidirectional equal power model, whenever $\delta_G < 12$, Lemma 5 guarantees that $K(\mathcal{N}) \leq 7$, and Theorem 1 provides a throughput guarantee that is better than the lower bound of $1/8$ in Lemma 1.

*Note that Lemma 6 holds for all unidirectional communication models and hence for the unidirectional equal power model.

VI. Discussion and Conclusion

In this paper, we have addressed the long-standing open question of attaining throughput guarantees with distributed scheduling in wireless networks. We have studied the performance of a simple distributed scheduling policy, maximal scheduling, which had earlier been investigated in context of node-exclusive spectrum sharing model and input-queued switches. We have obtained tight performance guarantees for maximal scheduling under arbitrary interference models and topologies, and have characterized the throughput region attained by maximal scheduling in terms of the interference degree of the network. The characterizations demonstrate that the performance bounds depend heavily on the nature of communication and interference models. We prove that maximal scheduling is guaranteed to attain a constant fraction of the maximum throughput region for certain communication and interference models, while it is also guaranteed to not attain a constant fraction in the worst case for some other models.

Concurrently with our work, Wu *et. al.* have obtained bounds for the throughput region of maximal scheduling [9]. Specifically, they proved that in the bidirectional and unidirectional interference models, maximal scheduling is guaranteed to attain at least $1/N_{\mathcal{E}}$ of the maximum throughput region, where $1/N_{\mathcal{E}}$ is the maximum number of links interfering with a given link. They also proved that in the bidirectional interference model there exists an arrival rate vector and a network such that maximal scheduling will attain at most $2/N_{\mathcal{E}}$ of the maximum throughput region. The upper bound is clearly interesting when $N_{\mathcal{E}} > 2$. Note that for any network $\mathcal{N}$, $K(\mathcal{N}) \leq \mathcal{N}_{\mathcal{E}}$, and in several cases $K(\mathcal{N}) << \mathcal{N}_{\mathcal{E}}$. Thus, the lower bound we obtained in Theorem 1 is tighter than that obtained by Wu *et. al.* [9]. Similarly, given a $N_{\mathcal{E}}$, one can construct a network with the same $N_{\mathcal{E}}$ and $K(\mathcal{N}) = N_{\mathcal{E}} - 1$, and when $N_{\mathcal{E}} \leq 2$, $N_{\mathcal{E}} - 1 > N_{\mathcal{E}}/2$. Thus, again the upper bound we obtained in Theorem 2 is tighter than that obtained by Wu *et. al.* [9]. Nevertheless, the proof techniques used by Wu *et. al.* [9] are certainly illuminating, and may be useful in chracaterizing the delay under maximal scheduling.

The class of maximal scheduling policies is quite broad, and our performance bounds apply to all policies in this class. However, it remains to be seen whether certain policies in this class can attain better performance bounds, while still being amenable to low-complexity distributed implementation. Similar questions remain open for distributed scheduling policies outside this class as well. Finally, recently, Wu *et. al.* [9], [10] and Bui *et. al.* [1] have shown that under specific interference models a combination of maximal scheduling with either regulators or congestion control techniques attains guaranteed fraction of the maximum throughput region when sessions traverse multiple hops. The question whether maximal scheduling alone attains similar performance bounds for multihop sessions remains open.

References

[1] L. Bui, A. Eryilmaz, R. Srikant, and X. Wu. Joint asynchronous congestion control and distributed scheduling for multihop wireless networks. In *Preprint*, 2005.

[2] P. Chaporkar, K. Kar, and S. Sarkar. Throughput guarantees through maximal scheduling in multihop wireless networks. Technical report, Univ. of Pennsylvania, Philadelphia, PA, Jul 2005. http://www.seas.upenn.edu/∼swati/publication.htm.

[3] J. Dai and B. Prabhakar. The throughput of data switches with and without speedup. In *IEEE INFOCOM'00*, pages 556–564, Tel Aviv, Israel, Mar 2000.

[4] X. Lin and N. Shroff. The impact of imperfect scheduling on cross-layer rate control in multihop wireless networks. In *IEEE INFOCOM'05*, Miami, FL, Mar 2005.

[5] D. Peleg. *Distributed Computing: A Locality-sensitive Approach.* Society of Industrial and Applied Mathematics, Philadelphia, PA, 2000.

[6] D. Shah, P. Giaccone, and B. Prabhakar. An efficient randomized algorithm for input-queued switch scheduling. *IEEE Micro*, 22(1):19–25, Jan-Feb 2002.

[7] L. Tassiulas. Linear complexity algorithms for maximum throughput in radio networks and input queued switches. In *INFOCOM'98*, pages 533–539, 1998.

[8] 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.

[9] X. Wu and R. Srikant. Bounds on the capacity region of multihop wireless networks under distributed greedy scheduling. In *Preprint*, 2005.

[10] X. Wu and R. Srikant. Regulated maximal matching: a distributed scheduling algorithm for multihop wireless networks with node-exclusive spectrum sharing. In *Proceedings of IEEE CDC-ECC'05*, Seville, Spain, Dec 2005.