Throughput and Fairness 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. Our results can be generalized to networks with multicast communication, arbitrary number of frequencies and end-to-end sessions. Finally, we enhance maximal scheduling to guarantee fairness of rate allocation.

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 [16]. The policy, however, is centralized and can have exponential complexity depending on the network topology considered. Later, Tassiulas [15] and Shah *et al.* [14] 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.* [7] 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,

P. Chaporkar is with INRIA, Paris, France. K. Kar is with the department of Electrical, Computer and Systems engineering at Renessealer Polytechnic Institute. S. Sarkar is with the department of Electrical and Systems Engineering at the University of Pennsylvania. Their emails are Prasanna.Chaporkar@ens.fr, koushik@ecse.rpi.edu, swati@seas.upenn.edu. Parts of this paper appeared in Allerton 2005 [3], ITA 2006 [8] and Wiopt 2006 [11]. This work was supported by the National Science Foundation under grants NCR-0238340, CNS-0435306, NCR-0448316 and CNS-0435141.

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, 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 [7]. Dai *et. al.* [4] 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 [7]. 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 comparisons between the throughput region of maximal scheduling with the maximum possible throughput region of the network characterizes the penalty due to the use of only local information in the scheduling. The characterizations of the throughput region of maximal scheduling obtained so far bounds the performance of the network in terms of that of the worst session. However, depending on the interference in individual neighborhoods, different sessions may be able to accommodate different arrival rates. The natural next question now is whether it is possible to obtain better non-uniform bounds by considering the constraints of individual sessions. We prove that under maximal scheduling the performance of each session can be characterized by the interference degree of only the links in its path, and the interference degrees of the neighbors of these links. Thus the performance penalty for each session, due to the use of local information based scheduling, depends only on the neighborhoods of the links in its path. The result is somewhat counterintuitive as the overall performances of sessions may depend on each other even when they are separated by several hops. Furthermore, we prove that the performance penalties under maximal scheduling can not be localized any further. Specifically, the interference degrees of the links of a session alone can not determine its throughput guarantee.

Maximal scheduling is really a class of policies, and some policies in this class could allocate bandwidth very unfairly. Recently, Lin *et al.* [7] and Bui *et. al.* [2] have shown that in the node exclusive spectrum sharing model, maximal scheduling can be used for maximizing the network utility and congestion control. We obtain global fairness guarantees in wireless networks with arbitrary interference models using maximal scheduling. First, using the characterizations for the throughput region for maximal scheduling, we characterize the feasible set of service rate allocations for maximal scheduling, and prove that a combination of a token generation scheme together with maximal scheduling attains maxmin fairness in this feasible set. We next show that the rate vector attained by the above combination is fairer than the overall maxmin fair rate vector times the reciprocal of the maximum interference degree in the network. The token generation scheme allows each session to estimate its maxmin fair rate in a distributed manner. Sessions contend for channel access in accordance with this estimate, and the contention is resolved using maximal scheduling. The token generation and the contention resolution can be executed in parallel. The maxmin fair rates need not be computed explicitly, and no knowledge of the statistics of the packet arrival process is necessary for executing the algorithm. The computation need not restart when the topology or the arrival rates change. The scheme is therefore robust.

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. In Section VI, we generalize the analytical results and the framework so as to include multicast and multichannel networks, different throughput guarantees for different sessions, stronger notions of stability and end-to-end performance guarantees. We describe how maximal scheduling can be enhanced so as to guarantee fairness in Section VII. We conclude in Section VIII. We present the proofs in appendix.

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

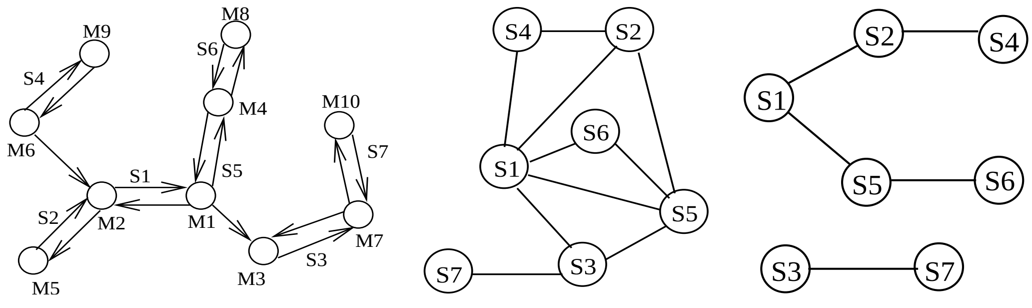

(a) Network $\mathcal{N}'$ (b) (c)

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.

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. We assume that at most $\alpha_{\max} > 1$ packets arrive for any session in any slot. 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). For simplicity, we will sometimes consider special cases of the above general model (Sections VI-D, VI-E, VII), and explicitly state whenever we do so.

Definition 5: The *arrival rate* of session i is $\lambda_i,\ i = 1, \ldots, N$. 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 [9]. 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 (Section VI-B).

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 (Appendix E). 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 $\vec{\lambda}/8 \in \Lambda^{\text{MS}}$.

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

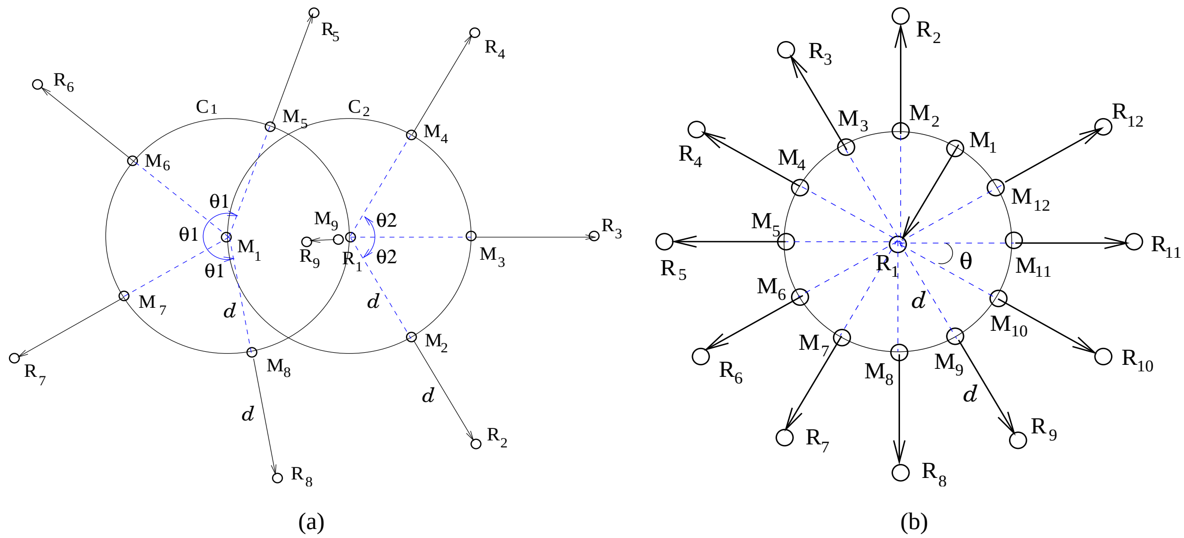

Fig. 2. Fig. (a) shows a network with interference constraints given by the bidirectional equal power model and transmission range d. There are 9 sessions: $1, \ldots, 9$. Session i has transmitter M_i and receiver R_i. The interference area of session 1 is the union of circles C_1 and C_2. Here, $\theta_1 = 70$ 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 $\theta = 2\pi/(Z + 2)$.

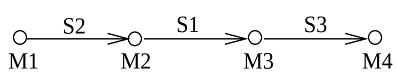

Fig. 3. Figure shows a network with 4 nodes $M1, \ldots, M4$ and 3 sessions $S1, S2$ and $S3$. Under node exclusive spectrum sharing model, $S1$ interferes with both $S2, S3$, but $S2$ and $S3$ do not interfere with each other.

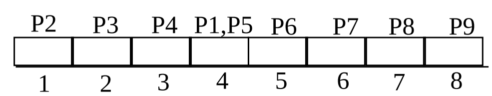

Fig. 4. Figure shows a periodic arrival process for the network in Fig. 2(a). The period is 8 slots. Session S_i generates a packet in the slot marked Pi, for each i. Here, $\lambda_i = 1/8$, for each i.

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 $2, 3, ...9$ 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, $\vec{\lambda}/Z \notin \Lambda^{\text{MS}}$.

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.* [7] 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.

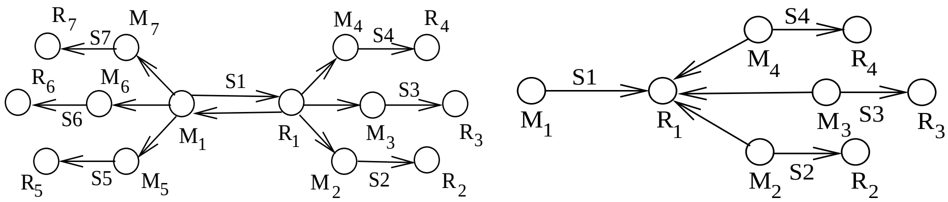

(a) Network $\mathcal{N}_1'$ (b) Network $\mathcal{N}_2'$

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.

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 $K(\mathcal{N}_1) = \max(\delta_G - 4, 1)$.

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 $\mathcal{N}$, $K(\mathcal{N}) \leq 8$ (Appendix E), 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* [7]. 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* [7]. 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.

VI. Generalizations of Throughput Guarantees

We first generalize the framework to characterize Λ^{MS} for some additional scenarios of practical interest. In subsection VI-A we consider a network with multicast sessions. In subsection VI-B, we consider a network with multiple (M) frequencies. Here, M may not be so large that every node can be allocated a frequency that is unique in its 2-hop neighborhood, and thus the node exclusive spectrum sharing model may not apply. We demonstrate that the overall framework may easily be extended to consider both cases, and Theorems 1 and2 hold.

Next, the characterizations of Λ^{MS} obtained so far demonstrate that maximal scheduling does not attain the maximum throughput region of a network. This is clearly expected as maximal scheduling uses only local information and the maximum throughput region has so far only been obtained by centralized scheduling policies [16], [15]. The contribution of these results is to characterize the penalty due to the use of such limited information, and provide tight “uniform” bounds on the penalty in the arbitrary networks. The bounds are “uniform” because they uniformly apply to all sessions. In subsection VI-C, we generalize Theorems 1 and 2 to obtain better throughput guarantees for specific sessions by allowing different bounds for different sessions (Lemma 9).

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

We have so far considered the notion of stability which guarantees that arrival rates of sessions equal their departure rates. This does not however provide guarantees on the expected queue lengths of the sessions. In subsection VI-D, we characterize the performance of maximal scheduling under a stronger notion of stability which guarantees that the expected queue lengths of all sessions are finite (Lemma 11).

Finally, in subsection VI-E, we relax the assumption that each sessions traverses only one hop, and provide throughput guarantees for maximal scheduling when sessions traverse arbitrary number of hops (Lemmas 12,13,14).

A. Multicast Networks

We now generalize the framework to support multicast (one-to-many) communications. Each multicast session has one sender and one or more receivers, and therefore has two or more end points. Thus, unicast sessions (which we considered so far) are special cases of multicast.

A session i has transmitter u, G_i receivers $(v_1, \ldots, v_{G_i})$ and is completely specified by $(i, u, v_1, \ldots, v_{G_i})$. For the bidirectional communication model, the description of the pairwise interference relations remain the same as in the unicast case. For the unidirectional communication models, the description must be generalized as follows: session i interferes with session j if i and j have a common end point, or one or more of j's receivers are neighbors of i's transmitter.

Given that the interference relations are still between two sessions, maximal scheduling can be used to schedule sessions. All the definitions introduced in context of arbitrary wireless networks again remain valid in this case. We now characterize Λ^{MS} in arbitrary wireless networks with multicast sessions. Theorems 1 and 2 also hold for multicast networks.

We now introduce some additional notations to generalize the results for specific interference models. The *multicast degree of a session* $(i, u, v_1, \ldots, v_{G_i})$ is the sum of the degrees of $u, v_1, \ldots, v_{G_i}$ and $-4G_i$. Let $\gamma(\mathcal{N})$ be the maximum multicast degree of a session in a network. The *multicast directional degree of a session* $(i, u, v_1, \ldots, v_{G_i})$ is the sum of the out-degree of u, and in-degrees of $v_1, \ldots, v_{G_i}$ and $-2G_i$. Let $\Gamma(\mathcal{N})$ be the maximum multicast directional degree of a session in a network. Let $G(\mathcal{N})$ be the maximum number of receivers in a multicast session in a network.

We first upper bound $K(\mathcal{N})$ for specific interference models, which would in turn provide lower bounds for Λ^{MS} using Theorem 1.

Lemma 7: Consider a wireless network $\mathcal{N}$ with multicast sessions.

1) *In the bidirectional communication model,* $K(\mathcal{N}) \leq \max(\gamma(\mathcal{N}), 1)$.
2) *In the unidirectional communication model,* $K(\mathcal{N}) \leq \max(\Gamma(\mathcal{N}), 1)$.
3) *In the bidirectional equal power model,* $K(\mathcal{N}) \leq 25$.
4) *In the node exclusive spectrum sharing model,* $K(\mathcal{N}) \leq G(\mathcal{N}) + 1$.

We now lower bound $K(\mathcal{N})$ for specific interference models, which would in turn provide upper bounds for Λ^{MS} using Theorem 2.

Lemma 8: 1) *In the bidirectional communication model, there exists a wireless network $\mathcal{N}$ such that* $K(\mathcal{N}) = \max(\gamma(\mathcal{N}), 1)$.
2) *In the unidirectional communication model, there exists a wireless network $\mathcal{N}$ such that* $K(\mathcal{N}) = \max(\Gamma(\mathcal{N}), 1)$.
3) *In the bidirectional equal power model, there exists a wireless network $\mathcal{N}$ such that* $K(\mathcal{N}) \geq 19$.
4) *In the unidirectional equal power model, given any constant Z there exists a wireless network $\mathcal{N}$ such that* $K(\mathcal{N}) > Z$.
5) *In the node exclusive spectrum sharing model, there exists a wireless network $\mathcal{N}$ such that* $K(\mathcal{N}) = G(\mathcal{N}) + 1$.

Using Lemmas 7 and 8 and Theorems 1 and 2, Λ^{MS} can now be characterized for specific interference models.

The generalizations in Lemmas 7 and 8 for the bidirectional and unidirectional communication models have been obtained by substituting $\max(\delta_G - 2, 1)$ and $\max(\Delta_G - 2, 1)$ in Lemmas 5 and 6 with

$\max(\gamma(\mathcal{N}), 1)$ and $\max(\Gamma(\mathcal{N}), 1)$ in Lemmas 7 and 8 respectively. Note that when all sessions are unicast, $\gamma(\mathcal{N}) = \delta_G - 4$ and $\Gamma(\mathcal{N}) = \Delta_G - 2$. Thus, Lemma 5 can be obtained as a special case of Lemmas 7 and 8.

Lemma 8 shows that in the unidirectional equal power model, maximal scheduling may not in general attain a constant fraction of the maximum throughput region. This is expected as a similar negative result holds for unicast networks (Lemma 3) and unicast is a special case of multicast.

When all sessions are unicast, $G(\mathcal{N}) = 1$. Then, Lemmas 7 and 8 and Theorems 1 and 2 guarantee that in the node exclusive spectrum sharing model, maximal scheduling attains at least $1/2$, and in some topologies no more than $1/2$ the maximum throughput region. This is consistent with the result obtained by Lin *et. al.* [7] and Lemma 4.

B. Multichannel Wireless Networks

We consider a wireless network with M channels. We assume that the transmissions from a session always use the same frequency which is pre-determined. We characterize Λ^{MS} for arbitrary frequency allocation strategies, but do not investigate the design of such strategies. A session i that traverses link (u, v) and transmits in channel k is now completely specified by the $4-$tuple (i, u, v, k). We first describe the transmission constraints. Now, node j can not receive any packet successfully in channel k if more than one of its neighbors are transmitting simultaneously in channel k. Thus, a transmission on edge $(i, j) \in E$ using channel k is successful in a slot if and only if no neighbor of j other than i transmits in channel k in the slot.

We now obtain the pairwise interference relations for both the bidirectional and unidirectional communication models using the above constraints. In the bidirectional communication model, a session i interferes with session j if they have a common end point (transmitter or receiver), or if they have the same frequency and one end point of j is a neighbor of an end point of i. In the unidirectional communication model, session i interferes with session j if they have a common end point, or if they have the same frequency and j's receiver is a neighbor of i's transmitter.

Given the above pairwise interference relations, all the definitions introduced in context of arbitrary wireless networks remain valid in this case. Theorems 1 and 2 also hold for arbitrary multichannel wireless networks. Both Lemmas 5 and 6 can be generalized to obtain specific results for bidirectional and unidirectional communication models.

C. Nonuniform Bounds

We now describe how we obtain different performance guarantees for different sessions. In Theorems 1 and 2, we have proved that in an arbitrary network $\mathcal{N}$, due to the use of maximal scheduling, the arrival rate that can be accommodated for each session reduces by at most $K(\mathcal{N})$, and the arrival rate that can be accommodated for at least one session reduces by at least $K(\mathcal{N})$. This uniform bound of a factor of $1/K(\mathcal{N})$ is obtained considering the worst session, and it is possible that for most sessions the penalty is less. We now prove that it is possible to obtain better non-uniform bounds by considering the constraints of individual sessions. Specifically, we show that the performance of each session i can be characterized by its *two-hop interference degree*, $\beta_i(\mathcal{N})$, which is the maximum of the interference degrees in its neighborhood (i.e., $\beta_i(\mathcal{N}) = \max_{j \in S_i \cup \{i\}} K_i(\mathcal{N})$), but not by its interference degree alone.

Lemma 9: If $(\lambda_1, \ldots, \lambda_N) \in \Lambda$, *then* $(\lambda_1/\beta_1(\mathcal{N}), \ldots, \lambda_N/\beta_N(\mathcal{N})) \in \Lambda^{MS}$.

Thus, due to the use of local information based scheduling, the performance of each session i decreases by a factor of $\beta_i(\mathcal{N})$; the penalty for each session therefore depends only on its two-hop neighborhood. Note that in many networks $\beta_i(\mathcal{N})$ may be significantly less than $K(\mathcal{N})$ for most sessions i (Figure 6(b)). The following result shows that a similar characterization in terms of the single-hop neighborhood does not hold in general.

Lemma 10: There exists a wireless network $\mathcal{N}$ *and an arrival rate vector* $(\lambda_1, \ldots, \lambda_N)$ *such that* $(\lambda_1, \ldots, \lambda_N) \in \Lambda$ *in* $\mathcal{N}$, *but* $(\lambda_1/K_1(\mathcal{N}), \ldots, \lambda_N/K_N(\mathcal{N})) \notin \Lambda^{MS}$.

D. Stronger Notion of Stability

In this subsection, we consider a stronger notion of stability, *queue length stability*, which guarantees that the expected queue lengths of sessions are finite in stable systems. We provide guarantees on the stability region of maximal scheduling under this notion and under some stronger assumptions on the arrival process. We first mention the additional assumptions on the arrival process and formally define the notion of queue-length-stability.

Now, $\alpha_j(t)$ and $\tilde{D}_j(t)$ denote the number of arrivals and departures respectively for session j in slot t. We assume that the arrival process $(\alpha_1(.), \ldots, \alpha_N(.))$ constitute an irreducible, aperiodic markov chain with a finite number of states. We refer to this assumption as the *jointly markovian* assumption. Note that such an arrival process satisfies (1).

Let $Q_i(n)$ be the number of packets waiting for transmission at the source of session i at the beginning of slot n.

Definition 15: The network is said to be *queue-length-stable* if there exists non-negative real numbers q_i, $i = 1, \ldots, N$, such that with probability 1,

$$\lim_{n \to \infty} Q_i(n)/n = q_i, \ i = 1, \ldots, N. \tag{3}$$

The *queue-length-stability 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 the jointly markovian assumption and has arrival rate vector $\vec{\lambda}$. The *maximum queue-length-stability region* Λ_{Q} is the union of the queue-length-stability region of all scheduling policies.
Note that if a network is queue-length-stable it is also stable, but the converse is not true. Thus, queue-length-stability is a stronger notion of stability.

We now obtain a lower-bound[†] on the queue-length-stability region of maximal scheduling $\Lambda_{\text{Q}}^{\text{MS}}$.

Lemma 11: Consider a jointly markovian arrival process with the arrival rate vector $(\lambda_1', \ldots, \lambda_N')$ such that $\lambda_1' < \lambda_1/\beta_1(\mathcal{N}), \ldots, \lambda_N' < \lambda_N/\beta_N(\mathcal{N})$, where $(\lambda_1, \ldots, \lambda_N) \in \Lambda_Q$. Then, $(\lambda_1', \ldots, \lambda_N') \in \Lambda_Q^{MS}$.

E. Multi-hop sessions

We now obtain performance guarantees for maximal scheduling when sessions traverse arbitrary number of links. We first mention the differences from the model in Section II. The network has N end-to-end sessions, each of which can be viewed as a collection of several hop-by-hop connections, one for each link it traverses; each of these hop-by-hop connections is called a *session-link* of the session considered. Each session-link is of the form (u, v), where u and v represent the transmitter and the receiver, respectively, of the corresponding session-links. For any session i, let P_i denote the set of its session-links. Let $q(j)$ denote the session of session-link j, i.e., $q(j) = \{i : j \in P_i\}$. We assume that there are a total of M session-links in the network (over all sessions), and these are indexed by $1, \ldots, M$.

The notions of interference, interference-set and interference-degrees are now defined for session-links instead of sessions. Specifically, a session-link j *interferes* with session-link k if k can not successfully transmit a packet when j is transmitting. The *interference set* of session-link j, S_j, denotes the set of session-links k such that either k interferes with j or j interferes with k (Fig. 6(a)). The *interference degree* of a session-link j in network $\mathcal{N}$, $K_j(\mathcal{N})$ is (i) the maximum number of session-links in its interference set S_j that can simultaneously transmit, if S_j is non-empty, and (ii) 1, if S_j is empty. The *two-hop interference degree of session-link* j, is defined as $\beta_j(\mathcal{N}) = \max_{m \in S_j \cup \{j\}} K_m(\mathcal{N})$. The *two-hop interference degree of session* i $\tilde{\beta}_i(\mathcal{N})$ denote the maximum two-hop interference degree of all session-links of session i, i.e., $\tilde{\beta}_i(\mathcal{N}) = \max_{j \in P_i} \beta_j(\mathcal{N})$. The *interference degree of a network* $\mathcal{N}$, $K(\mathcal{N})$, is the maximum interference degree of session-links in the network.

[†]We presented this result in ITA workshop [8]. Wu *et al.* [20] also obtained this result independently, and presented it in the same workshop.

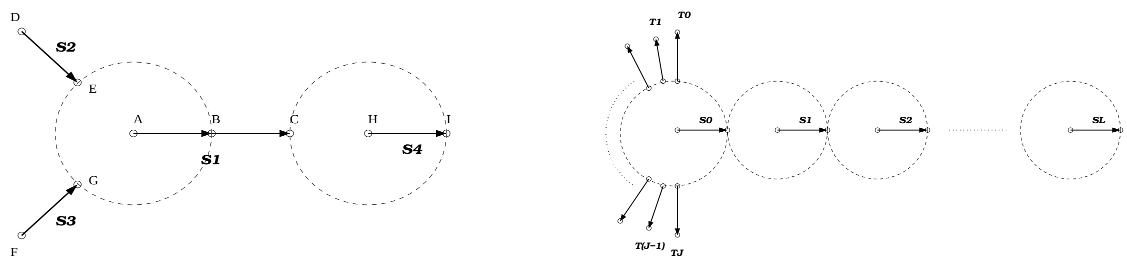

(a) (b)

Fig. 6. In both figures, all sessions and session-links are unidirectional, and the arrows show the direction of data transfer. The circles indicate the interference regions of session-links AB and HI (Fig. (a)) and S0, S1, ..., SL (Fig. (b)).
In Fig. (a), session S1 consists of two session-links, AB and BC, whereas sessions S2, S3, S4 are single-hop sessions. Session-link AB interferes with session-links DE (session S2) and FG (session S3) and session-link HI (session S4) interferes with session-link BC. Now, $S_{\text{AB}} = \{$BC, DE, FG$\}, S_{\text{BC}} = \{$AB, HI$\}, S_{\text{DE}} = S_{\text{FG}} = \{AB\}, S_{\text{HI}} = \{BC\}$. Thus, token-buckets at nodes A, B, D, F, H consist of token-queues corresponding to session-links $\{$AB, BC, DE, FG$\}$, $\{$ AB, BC, HI$\}$, $\{$AB, DE$\}$, $\{$AB, FG$\}$, and $\{$BC, HI$\}$. Thus, token-buckets associated with session-link AB (BC) are at nodes A, B, D, F (A, B, H); these are denoted buckets $1, \ldots 4$ of AB ($1, 2$ of BC). The token generation for AB at bucket 4 depends on that for AB at bucket 3 and BC at bucket 1 of BC.
In Fig. (b), network consists of single-hop sessions only. Session S0 interferes with sessions T0, ..., TJ, whereas session Si interferes with session S(i-1), for i=1,2, ..., L. Thus, $K_i(\mathcal{N}) = 1$ for $i \in \{\text{T0}, \ldots \text{TJ}, \text{SL}\}$, $K_i(\mathcal{N}) = 2$ for $i \in \{\text{S1}, \ldots \text{S(L-1)}\}$, $K_{\text{S0}}(\mathcal{N}) = J + 2$, $\beta_i(\mathcal{N}) = J + 2$ for $i \in \{\text{T0}, \ldots \text{TJ}, \text{S0, S1}\}$, and $\beta_i(\mathcal{N}) = 2$ for $i \in$ S2, ..., SL, $K(\mathcal{N}) = (J + 2)$. If J and L are large, but L $\gg$ J, then K_i, β_i for most sessions are substantially smaller than $K(\mathcal{N})$.

The packet arrival and departure processes now need to be defined for session-links. Now, $A_j(n)$ denotes the number of arrivals for session-link j in the time interval $(0, n]$, $j = 1, \ldots, M$. The arrival process at the first session-link of any session consists only of exogenous packets, and satisfies the SLLN as described in (1). Thus, if F_i denotes session-link corresponding to the first link for session i, then there exists non-negative real numbers $\lambda_i, i = 1, \ldots, N$ such that with probability 1,

$$\lim_{n \to \infty} A_{F_i}(n)/n = \lambda_i, \; i = 1, \ldots, N. \tag{4}$$

Now, $D_j(n)$ denotes the number of packets that session-link j transmits in interval $(0, n]$, $j = 1, \ldots, M$. Note that if j and $j + 1$ are consecutive session-links of a session, then $A_{j+1}(n) = D_j(n)$. Now, let L_i be the session-link corresponding to the last hop of session i. If for some constant d_i, the limit $\lim_{n \to \infty} D_{L_i}(n)/n = d_i$ with probability 1, then d_i is denoted as the departure rate of session i.

Definition 16: The network is said to be *stable* if there exists a departure rate vector $\vec{d} = (d_1, \ldots, d_N)$ such that with probability 1, for each session i

$$\lim_{n \to \infty} D_{L_i}(n)/n = d_i = \lambda_i, \; i = 1, \ldots, N. \tag{5}$$

Thus, again a network is stable if the arrival and departures rates are equal for each session. Now, using the above definition for stability, the maximum throughput region, Λ, and the throughput region for maximal scheduling, Λ^{MS}, can be defined as in Section II. Note that maximal-scheduling can be described similar to that in Section II; the only difference is that session-links must now be used instead of sessions in the description.

We first provide an upper-bound on Λ^{MS}.

Lemma 12: Given any constant Z, there exists a network $\mathcal{N}$, an arrival rate vector $\vec{\lambda}$ such that $K(\mathcal{N}) = \mathcal{Z}$, $\vec{\lambda} \in \Lambda$ in $\mathcal{N}$, but $\vec{\lambda}/\kappa \notin \Lambda^{MS}$ in $\mathcal{N}$ for any $\kappa < K(\mathcal{N})$.

We now provide lower-bounds on Λ^{MS}. under an enhancement of maximal scheduling that has been proposed by Wu *et. al.* [18], [19]. Under this enhancement, every session-link that does not originate from the source of the session has a regulator that in each slot generates a token with a probability that equals the arrival-rate of the session. Every such session-link also maintains two-queues, a *waiting-queue* and a *release-queue*. Packets arriving at such a session-link are initially stored in its waiting-queue. Whenever the regulator generates a new token, if the waiting-queue is non-empty, a packet is transferred from the waiting-queue to the release-queue. A session-link that originates from the source of the session maintains only the release-queue, and all exogenous packets waiting for transmission are stored there. Maximal scheduling only considers the release-queues of session-links for service and contention resolution. We refer to this enhancement as *regulator-enhancement.*

Lemma 13: If $\vec{\lambda} \in \Lambda$, then $(\lambda_1/\tilde{\beta}_1(\mathcal{N}), \ldots, \lambda_N/\tilde{\beta}_N(\mathcal{N})) \in \Lambda^{\text{MS}}$ in $\mathcal{N}$ under the regulator-enhancement.

Note that from Lemma 13 and since $K(\mathcal{N}) \geq \tilde{\beta}_i(\mathcal{N})$, $i = 1, \ldots, N$, if $\vec{\lambda} \in \Lambda$, then $\vec{\lambda}/K(\mathcal{N}) \in \Lambda^{\text{MS}}$ in $\mathcal{N}$ under the regulator-enhancement.

The use of regulators requires that the arrival rate for each session must be knows an each session-link. We now investigate whether performance guarantees can be provided for maximal scheduling without using regulators. We consider a special case of the general arrival process described in (4). We refer to this special case as *exponentially-convergent arrival processes.* We assume that there exists a constant $\hat{\alpha} > 1$ such that the empirical average of the exogenous arrivals in the system in T slots converges to $\vec{\lambda}$ at a rate faster than $\frac{1}{T^{\hat{\alpha}}}$. Mathematically, there exists $\widehat{t_\delta}$ such that for every $i \in \{1, \ldots, m\}$, $T \geq \widehat{t_\delta}$, and $\delta > 0$,

$$\mathbf{P}\left\{\left|\frac{\sum_{t=1}^{T} A_{F_i}(t)}{T} - \lambda_i\right| > \delta\right\} < \frac{1}{T^{\hat{\alpha}}}. \tag{6}$$

Again, a large class of arrival processes, e.g., periodic, i.i.d., and positive recurrent Markovian arrival processes with finite state space, satisfy the above assumption. We show that, without any enhancements[‡], for exponentially-convergent arrival processes, maximal scheduling attains the following weaker notion of stability. We define a random variable $B_{j,t}$ as follows. If session-link j has a packet to transmit at time t, then $B_{j,t}$ is the length of its remaining busy period, otherwise $B_{j,t} = 0$.

Lemma 14: Consider exponentially-convergent arrival processes. Let the arrival rate vector $(\lambda_1', \ldots, \lambda_N')$ be such that $\lambda_1' < \lambda_1/\tilde{\beta}_1(\mathcal{N}), \ldots,$ $\lambda_N' < \lambda_N/\tilde{\beta}_N(\mathcal{N})$, where $(\lambda_1, \ldots, \lambda_N) \in \Lambda$. Then under maximal scheduling, the packet queue of every session-link will almost surely become empty infinitely often. Furthermore, for every session-link j and time t, $\mathbf{E}[B_{j,t}] < \infty$.

The above result implies that almost surely $\limsup_{n\to\infty} \frac{D_j(n)-A_j(n)}{n} = 0 \;\; \forall\ j = 1, \ldots, M$. Thus, if the arrival rate vector satisfies the condition in Lemma 14, and for each session link the limits of the departure and the arrival rates exist almost surely, then almost surely $\lim_{n\to\infty} D_{L_i}(n)/n = \lambda_i \;\; \forall\ i = 1, \ldots N$, and the system is stable under maximal scheduling. But, there is no guarantee that these limits exist. Thus, this is a weaker notion of stability than that in Definition 16. Whether the stronger notion of stability, holds in this case or not, remains an open question.

VII. Max-Min Fairness under Maximal Scheduling

We have so far characterized the throughput region for maximal scheduling Λ^{MS} under different system assumptions. We now describe the issues involved when the arrival rate vector is not in Λ^{MS}. Then maximal scheduling can not serve all sessions at their arrival rates, and therefore it is necessary to fairly allocate the service rates or departure rates of sessions. We describe how to enhance maximal scheduling so as to ensure maxmin fair allocation of rates in the feasible set for maximal scheduling. We also prove that the rate vector attained by this enhancement is fairer than the reciprocal of the network interference degree times the maxmin fair rate vector in the overall network feasible set. We first consider networks with single-hop sessions (Subsection VII-A) and subsequently networks with multi-hop sessions (Subsection VII-B).

A. Single-hop Sessions

We assume that every session spans one link. Thus, the framework presented in Section II applies. We introduce our fairness notions and additional assumptions in Section VII-A.1, and subsequently describe the enhancement used for attaining max-min fairness and the performance guarantees in Section VII-A.2.

[‡]Each session-link therefore has only one queue for storing the packets waiting for transmission.

1) Fairness notion and terminologies: We first present a lemma that is useful in describing the feasible set under maximal scheduling.

Lemma 15:

$$\Lambda^{\text{MS}} = \{\vec{\lambda} = (\lambda_1, \ldots, \lambda_N) : \text{ if } \lambda_i > 0, \sum_{j \in S_i \cup \{i\}} \lambda_j \leq 1, \quad \forall\, i = 1, \ldots, N.\} \tag{7}$$

We can now describe the *feasible set* Δ^{MS} *of departure rate vectors* $\vec{d} = (d_1, \ldots, d_N)$ under maximal scheduling as follows:

$$\text{if } \lambda_i > 0, \sum_{j \in S_i \cup \{i\}} d_j \;\leq\; 1, \quad \forall\, i = 1, \ldots, N, \tag{8}$$

(interference constraints)

$$d_i \;\leq\; \lambda_i \quad \forall\, i = 1, \ldots, N. \tag{9}$$

The "interference constraints" (8) capture the interference relations and are analogous to constraints (7) for the stability region. The constraints (9) follow since the departure rates can not exceed the arrival rates.

Note that $\Delta^{\text{MS}} \subseteq \Lambda^{\text{MS}}$. When $\vec{\lambda} \in \Lambda^{\text{MS}}$, the departure rate vector satisfies $d_i = \lambda_i$ for each i and hence both (8) and (9) hold. When $\vec{\lambda} \notin \Lambda^{\text{MS}}$, depending on the maximal scheduling policy used, the departure rate vector can be any element of Δ^{MS}, and hence can be unfair for some sessions. For example, if maximal scheduling provides absolute priority to a session i, and $\lambda_i > 1$, then $d_i = 1$ and the departure rates of sessions in S_i are 0. This motivates our goal of ensuring fairness using maximal scheduling.

We now define the notion of maxmin fairness that we seek to attain. For any N-dimensional vector a, let $\mathcal{I}(a)$ denote a non-decreasing ordering of the components of a. Therefore, if $a = (a_1, a_2, \ldots, a_N)$ and $\mathcal{I}(a) = (\hat{a}_1, \hat{a}_2, \ldots, \hat{a}_N)$,then $(\hat{a}_1, \hat{a}_2, \ldots, \hat{a}_N)$ is a permutation of $(a_1, a_2, \ldots, a_N)$, satisfying $\hat{a}_1 \leq \hat{a}_2 \leq \ldots \leq \hat{a}_N$. A departure rate vector $\vec{d}^*$ is said to be maxmin fair if $\vec{d}^* \in \Delta^{\text{MS}}$, and for any other departure rate vector $\vec{d} \in \Delta^{\text{MS}}$, the first non-zero component in $\mathcal{I}(\vec{d}^*) - \mathcal{I}(\vec{d})$ is positive. Intuitively, a departure rate vector is maxmin fair if it is not possible to increase any of its components without decreasing any other component of equal or lesser value [1]. Note that $\vec{d}^* \in \Lambda^{\text{MS}}$ as $\Delta^{\text{MS}} \subseteq \Lambda^{\text{MS}}$. Finally, if $\vec{\lambda} \in \Lambda^{\text{MS}}$, then $\vec{d}^* = \vec{\lambda}$.

Next, we present a condition that is both necessary and sufficient for any departure rate vector to be maxmin fair. We first introduce the notion of a bottleneck constraint.

Definition 17: For any departure rate vector $\vec{d}$, an interference constraint is a *bottleneck constraint* for a session i if (a) i is involved in the constraint, (b) $d_i \geq d_k$ for all other sessions k whose sessions are associated with the constraint and (c) the inequality in the constraint is an equality.

Lemma 16: A departure rate vector $\vec{d} \in \Delta^{MS}$ is maxmin fair if and only if the following holds: for every session i, either $d_i = \lambda_i$, or the session has a bottleneck constraint.

We omit the proof for the above lemma as the proof is similar to that for the well-known bottleneck condition for maxmin fairness in wireline networks [1].

Finally, although for notational simplicity we refer to $\vec{d}^*$ as the maxmin fair departure rate vector, it is maxmin fair only in the feasible set of maximal scheduling Δ^{MS}. The feasible set for the network Δ is the union of the feasible sets of all policies, and may be a strict superset of Δ^{MS}. Thus, the maxmin fair departure rate vector in the network ($\vec{m}^*$), which we refer to as the *globally maxmin fair departure rate vector*, is the rate vector which is maxmin fair in Δ. We now describe the relation between $\vec{d}^*$ and $\vec{m}^*$. We first describe the notion of "relative fairness" introduced in [12]. A departure rate vector $\vec{a}$ is fairer than another departure rate vector $\vec{b}$ if the first non-zero component in $\mathcal{I}(\vec{a}) - \mathcal{I}(\vec{b})$ is positive. Note that by this definition a departure rate vector is maxmin fair in any feasible set if it is fairer than any other departure rate vector in the same feasible set. Now, since $\vec{m}^* \in \Delta$, $\vec{m}^*/K(\mathcal{N}) \in \Delta^{\text{MS}}$. Thus, from the definition of $\vec{d}^*$, $\vec{d}^*$ is either fairer than $\vec{m}^*/K(\mathcal{N})$ or $\vec{d}^* = \vec{m}^*/K(\mathcal{N})$.

We will consider a special case of the general arrival model presented in (1). Specifically, we will consider the *bounded-burstiness* arrival model where (a) $\lambda_i > 0 \ i = 1, \ldots, N$[§] and there exists a burstiness vector $\vec{\sigma} = (\sigma_1, \ldots, \sigma_N)$ such that

$$|A_i(t) - \lambda_i t| \leq \sigma_i \ \forall \ t. \tag{10}$$

2) Maxmin fair rate allocation algorithm: We propose a modular approach for attaining maxmin fairness using maximal scheduling (Figure 7). The first module estimates the maxmin fair bandwidth share of each session in each node in the session's path, and releases packets for transmission in accordance with these estimates. The second module schedules the transmission of the released packets so as to attain the estimates. Note that the modules operate in parallel.

Fair bandwidth is estimated by a token generation process. The source node for each active session i maintains a token bucket for i (Fig. 6(a)). The token bucket consists of a token-queue for each session in $S_i \cup \{i\}$. Every token bucket generates tokens for all token-queues in it. The token generation process is so designed that each token-queue receives tokens at a rate that equals the maxmin fair departure rate of the corresponding session (we shortly describe how this can be done). Whenever a new token is generated for a session i at the token bucket for i at i's source, i's source releases a new packet for transmission. Thus, the packet release rates are maxmin fair and hence belong to Λ^{MS}. Only the released packets are eligible for transmission. Thus, maximal scheduling transmits the released packets at the rates at which they are released. Hence, the rate allocations are maxmin fair.

We now describe the token generation process for each token-bucket. An session i is associated with $b_i = |S_i| + 1$ token-buckets, one for each of the sessions it interferes with, and itself. Let us denote these token-buckets as $1, \ldots, b_i$. Each token-bucket samples all sessions in the bucket in a round robin order. Let $C_{i,k}(t)$ be the number of tokens generated for session i at bucket k in the interval $(0, t]$. Let token-bucket k' ($1 < k' < b_i$) associated with i be sampled in slot t. Let k' not be at the source of i. Then, k' generates a token for session i in slot t if and only if $C_{i,k'}(t) < W + \min\left(C_{i,k'-1}(t), C_{i,k'+1}(t)\right)$. Thus, i receives a token at bucket k' unless the number of tokens for i at k' substantially exceeds that at the adjacent buckets; this prohibitive difference is the window parameter, W. If k' is at i's source, k' generates a token to i in slot t if and only if the number of packets generated for i at i's source in $(0, t]$ exceeds $C_{i,k'}(t)$ and $C_{i,k'}(t) < W + \min\left(C_{i,k'-1}(t), C_{i,k'+1}(t)\right)$. In slot t, k' samples the next session in the bucket in a round robin order if and only if k' does not generate a token for i. Note that token-bucket 1 and b_i have only one adjacent token-bucket for session i, and thus decide whether to generate a token based on the number of tokens at only one adjacent token-bucket. Tokens are never removed from a bucket.

We now explain why the token generation rate for each session at each token-bucket associated with the session equals the session's maxmin fair rate. For this explanation, we assume that $\lambda_i > 1$ for each i; all performance guarantees in this section however hold for arbitrary $\vec{\lambda}$. Since $\lambda_i > 1$ for each i, constraints (8) subsume constraints (9). Also, the number of packets generated for i at i's source in $(0, t]$ exceeds the number of tokens generated in $(0, t]$ at the bucket at its source for any $t \geq \sigma_i$. Thus, the token generation process at a bucket for i at i's source does not differ from that at a node that is not i's source.

Note that each token-bucket corresponds to constraint (8) for some $j \in \{1, \ldots, M\}$. Since the goal is to allocate maxmin-fair rates, each constraint should try to allocate equal rates to all sessions in the constraint. This motivates the round robin sampling of the sessions at each token-bucket. Again, all constraints involving a session must offer the same rate to the session. This is attained by relating the token generation process for a given session at a given token-bucket to that at the adjacent token-buckets for the same session. The number of tokens for a session at two adjacent buckets associated with the session differ by at most W at any time t, and the difference is at most $b_i W$ for that at any two buckets associated with the session. Thus, the rates of token generation for a session are nearly the same at any two buckets associated with the session.

[§]This assumption requires that the arrival rate for each active session is positive. Note that if a session i is not active we do not need to consider it at all. Thus, we assume that there are N active sessions denoted $1, \ldots, N$. In this section, a session will always refer to an active session, though for brevity we omit the adjective "active".

Procedure Token Generation (node m)
begin

For all t and session i, let $C_{i,0}(t) = C_{i,b_i+1}(t) = \infty$.
Let $A_i^{\text{NR}}(t)$ be the number of packets of session i at slot t that have been generated at its source but not been released.
Let $\Theta_{i,k}(t) = A_i^{\text{NR}}(t)$ if the kth bucket of session i is at i's source-node, and $\Theta_{i,k}(t) = \infty$ otherwise.
Each bucket samples the sessions associated with it in round robin order.
When session i is sampled at its kth bucket in slot t:
if $\Theta_{i,k}(t) > 0$ and $C_{i,k}(t) < C_{i,k+1}(t) + W$ and $C_{i,k}(t) < C_{i,k-1}(t) + W$, **then**
generate a token for session i at its kth bucket ($C_{i,k}(t+1) = C_{i,k}(t) + 1$);
else
do not generate token for session i at its kth bucket ($C_{i,k}(t+1) = C_{i,k}(t)$), and
sample the next session at the kth bucket in the round robin order.

end

Procedure Packet Release (source i)
begin

Release a new session i packet for transmission at session i source node when a token is generated for the session at the bucket at its source.

end

Procedure Packet Scheduling For Transmission
begin

Transmit the released packets using maximal scheduling.

end

Fig. 7. Pseudo code of the fair departure rate allocation algorithm when each session traverses one hop

Since $\lambda_i > 1$ for each i, every session has a bottleneck constraint under the maxmin fair rate allocation. Now, the maxmin fair rate of a session is determined by the bandwidth offered by the bottleneck constraint which offers the least bandwidth to the session. The bucket corresponding to the bottleneck constraint of a session is denoted as the *bottleneck bucket* for the session. Now, a session's token generation rate at any token-bucket equals that at its bottleneck bucket, which turns out to be the session's maxmin fair rate. If a session has a low maxmin fair rate, then its bottleneck constraint offers it a low rate, and it does not receive tokens several times it is sampled at other buckets; other sessions with less severe constraints receive these tokens. Thus, the following performance guarantee holds.

Lemma 17: Consider token-bucket k of session i. For the bounded-burstiness arrival model and arbitrary $\vec{\lambda}$, there exists constants ϱ, W_0, such that if $W \geq W_0$, then for any interval $(n_1, n_2]$, $|\frac{C_{i,k}(n_2) - C_{i,k}(n_1)}{n_2 - n_1} - d_i^*| \leq \frac{\varrho}{n_2 - n_1}$.

The token generation scheme here is based on the same design principle as that for an existing centralized fair bandwidth allocation algorithm [13], [17]. However, the constraints characterizing the feasibility set for maximal scheduling are significantly different from those characterizing the feasibility set in [13], [17]; therefore, the scheme differs significantly in the two cases.

We now describe the packet scheduling policy. Whenever the source node of a session i generates a new token for i at i's token-bucket at the source (the one associated with sessions in $S_i \cup \{i\}$), i releases a new packet. Only the sessions that have released packets waiting for transmission contend for scheduling, and are scheduled as per maximal scheduling. When these sessions are scheduled, they transmit only released packets.

Packets that contend for scheduling and are transmitted by maximal scheduling arrive as per the release process. The release rate vector is maxmin fair (Lemma 17) and is therefore in Λ^{MS}. Maximal scheduling therefore provides departure rates equal to the packet release rates. Thus, as the following result states, a combination of token generation and maximal scheduling attains the maxmin fair departure rates for every session.

Theorem 3: For the bounded-burstiness arrival model and arbitrary $\vec{\lambda}$, there exists a constant W_0, such that when $W \geq W_0$, $\lim_{n \to \infty} D_{L_i}(n)/n = d_i^*,\ i = 1, \ldots N$.

Procedure Token Generation (node m)
begin

For session-link i, let l and m respectively be the previous and next session-links of the same session.
For each slot t and session-link i,
if i is the first-session-link of its session, **then**
$C_{i,0}(t) = \infty, C_{i,b_i+1}(t) = C_{m,0}(t)$
else
if i is the last session-link of its session, **then**
$C_{i,0}(t) = C_{l,b_l+1}(t), C_{i,b_i+1}(t) = \infty$
else
$C_{i,0}(t) = C_{l,b_l}(t)$ and $C_{i,b_i+1}(t) = C_{m,0}(t)$.
Let $A_i^{\text{NR}}(t)$ be the number of packets of session-link i at slot t that are in its waiting-queue.
Let $\Theta_{i,k}(t) = A_i^{\text{NR}}(t)$ if the kth bucket of session-link i is at the source-node of session of i, and $\Theta_{i,k}(t) = \infty$ otherwise.
Each bucket samples the session-links associated with it in round robin order.
When session-link i is sampled at its kth bucket in slot t:
if $\Theta_{i,k}(t) > 0$ and $C_{i,k}(t) < C_{i,k+1}(t) + W$ and $C_{i,k}(t) < C_{i,k-1}(t) + W$, **then**
generate a token for session-link i at its kth bucket ($C_{i,k}(t+1) = C_{i,k}(t) + 1$);
else
do not generate token for session i at its kth bucket ($C_{i,k}(t+1) = C_{i,k}(t)$), and
sample the next session-link at the kth bucket in the round robin order.

end

Procedure Queue Management (session-link i)
begin

When a new packet is generated for session-link i or a new packet arrives at the source of session-link i from a previous session-link, add the new-packet in the waiting-queue for session-link i.
Transfer a session-link i packet from its waiting-queue to its release-queue at its source node when a token is generated for it at the bucket at its source.

end

Procedure Packet Scheduling For Transmission
begin

Transmit the packets in the release-queues of the session-links using maximal scheduling.

end

Fig. 8. Pseudo code of the fair departure rate allocation algorithm when sessions traverse multiple hops

B. Multi-hop Sessions

We next allow sessions to traverse multiple hops. Thus, the framework in Section VI-E applies. The *feasible set* Δ^{MS} *of departure rate vectors* $\vec{d} = (d_1, \ldots, d_N)$ can be described by (9) and

$$\text{if } \lambda_{q(j)} > 0, \sum_{k \in S_j \cup \{j\}} d_{q(k)} \leq 1, \quad \forall\, j = 1, \ldots, M. \tag{11}$$

Using the above description for Δ^{MS}, the maxmin fair departure rate vector can now be defined as in Section VII-A.

Definition 18: For any departure rate vector $\vec{d}$, an interference constraint is a *bottleneck constraint* for a session i if (a) a session-link j of i is involved in the constraint, (b) $d_{q(j)} \geq d_{q(k)}$ for all other session-links k whose sessions are associated with the constraint and (c) the inequality in the constraint is an equality. Again, with the above definition for a bottleneck constraint, Lemma 16 provides a necessary and sufficient condition for a departure rate vector to be maxmin fair.

We now describe the modifications required in the algorithm presented in Figure 7 for attaining maxmin fairness in this general case. We first describe the modifications in the token-generation procedure. Now, session-links, rather than sessions, are associated with token-buckets, and the source of each session-link j maintains the bucket consisting of session-links in $S_j \cup \{j\}$. Again, token-buckets sample session-links rather than sessions. The token generation process for the session-links are now similar to that for single-hop sessions. The main difference is that the token-generation process for a session-link j at the first (last) token-bucket of j must also depend on the number of tokens generated at the last (first) token-bucket for the previous (next) session-link k of the same session (Fig. 6(a)). We now describe the packet scheduling policy. The source of each session-link maintains two packets queues: a *waiting* packet queue, and a *released* packet queue. On arrival, a packet is queued at the waiting packet queue. A packet is forwarded from the waiting to the released queue when a new token is generated at the token-bucket for the session-link at the session-link's source. Only session-links with non-empty released queues contend for scheduling. The rest of the scheduling remains the same as that for the case of single-hop sessions. Refer to Figure 8 for a pseudo-code.

Both Lemma 17 and Theorem 1 hold; the term 'session' must now be replaced with 'session-link' in the statement of Lemma 17.

We now make a few concluding remarks on our maxmin fair packet scheduling algorithm. Note that the token-buckets associated with a session-link i need to know the number of tokens generated for i at other token-buckets associated with i. Also note that a token bucket associated with i is either at i's source or at j's source, where $j \in S_i$. Thus, a token bucket at the source of a session-link k need only know the number of tokens generated at a token-bucket at the source of a session-link l if and only if both k and l interfere with each other or with a common session-link. Since only session-links in close proximity interfere with each other in a wireless network, the token-generation process requires communication among nodes in proximity as well. Finally, the analytical guarantees hold even when nodes know the number of tokens generated at other nodes after some delay, as long as the delay is upper-bounded by a constant.

VIII. 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. Our results can be generalized to networks with multicast communication, arbitrary number of frequencies and end-to-end sessions. Finally, we enhance maximal scheduling to guarantee fairness of rate allocation.

Concurrently[¶] with our work, Wu *et. al.* have obtained bounds for the throughput region of maximal scheduling [19]. 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.* [19]. 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.* [19]. Nevertheless, the proof techniques used by Wu *et. al.* [19] 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

[¶]Our major results were presented at Allerton conference, September $28-30$, 2005, and Wiopt conference, April, $3-7$, 2006. Wu *et. al.* [19] reported their results at INFOCOM conference, April, $23-29$, 2006.

questions remain open for distributed scheduling policies outside this class as well. Recently, Sharma *et. al.* [6] have lower bounded the complexity of policies that attain the maximum stability region, or approximate the maximum stability region within constant factor, in arbitrary topologies. These results may help answer some of the above open questions.

REFERENCES

[1] D. Bertsekas and R. Gallager. *Data Networks (2nd ed.)*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1992.

[2] L. Bui, A. Eryilmaz, R. Srikant, and X. Wu. Joint asynchronous congestion control and distributed scheduling for multihop wireless networks. In *Proceedings of INFOCOM*, Barcelona, Spain, April 2006.

[3] P. Chaporkar, K. Kar, and S. Sarkar. Throughput guarantees through maximal scheduling in multihop wireless networks. In *Proceedings of 43d Annual Allerton Conference on Communication, Control and Computing*, Allerton, Monticello, Illinois, September 28-30 2005.

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

[5] G. Fayolle, V. A. Malyshev, and M. V. Menshikov. *Topics in the Constructive Theory of Countable Markov Chains*. Cambridge University Press, 1995.

[6] N. Shroff G. Sharma and R. Mazumdar. On the complexity of scheduling in multi-hop wireless systems. In *Proceedings of IEEE International Workshop on Foundations and Algorithms for Wireless Networking (FAWN)*, Pisa, Italy, Mar 2006.

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

[8] K. Kar P. Chaporkar and S. Sarkar. Achieving queue length stability through maximal scheduling in wireless networks. In *Proceedings of Information Theory and Applications Inaugural Workshop*, University of California, San Diego, February 6-10 2006.

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

[10] W. Rudin. *Real and Complex Analysis*. McGraw Hill, 1987.

[11] S. Sarkar, P. Chaporkar, and K. Kar. Fairness and throughput guarantees with maximal scheduling in multihop wireless networks. In *Proceedings of the fourth Workshop On Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, Boston, MA, April 2006.

[12] S. Sarkar and K. N. Sivarajan. Fairness in wireless mobile networks. *IEEE Transactions of Information Theory*, 48(8):2412–2426, Aug 2002.

[13] S. Sarkar and L. Tassiulas. End-to-end bandwidth guarantees through fair local spectrum share in wireless ad-hoc networks. *IEEE Transactions on Automatic Control*, September 2005.

[14] 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.

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

[16] 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.

[17] L. Tassiulas and S. Sarkar. Maxmin fair scheduling in wireless ad hoc networks. *IEEE Journal for Selected Areas in Communications, special issue on Ad Hoc Networks, Part I*, 23(1), January 2005.

[18] 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.

[19] X. Wu and R. Srikant. Bounds on the capacity region of multihop wireless networks under distributed greedy scheduling. In *Proceedings of INFOCOM*, Barcelona, Spain, April 2006.

[20] R. Srikant X. Wu and James R. Perkins. Queue length stability of maximal greedy schedules in wireless networks. In *Proceedings of Information Theory and Applications Inaugural Workshop*, University of California, San Diego, February 6-10 2006.

APPENDIX

PROOFS OF ANALYTICAL RESULTS IN SECTION V (THEOREMS 1 AND 2 AND LEMMAS 5 AND 6)

Appendix A: Proof of Theorem 1

We prove Theorem 1 using the following supporting lemmas.

Lemma 18: Let $\vec{\lambda} \in \Lambda$. Then, $\sum_{j \in S_i \cup \{i\}} \lambda_j \leq K(\mathcal{N})$.

Lemma 19: Let $\vec{\lambda} \in \{\vec{\lambda} : \text{ if } \lambda_i > 0, \ \sum_{j \in S_i \cup \{i\}} \lambda_j \leq 1, \ i = 1, \ldots, N\}$. Then $\vec{\lambda} \in \Lambda^{\text{MS}}$.

Theorem 1 follows from Lemmas 18 and 19. ■

We now prove Lemmas 18 and 19.

Appendix A.1: Proof of Lemma 18:

We prove Lemma 18 using a supporting lemma, Lemma 20, which we state and prove first. Lemma 20 is stated and proved for sessions with arbitrary number of hops.

Lemma 20: If $\vec{\lambda} \in \Lambda$, then (a) for each session-link j, $j = 1, 2, \ldots, M$, $\sum_{j \in S_k \cup \{k\}} \lambda_{q(j)}/\tilde{\beta}_{q(j)}(\mathcal{N}) \leq 1$, and (b) if each session spans one hop, for each session i, $i = 1, \ldots, N$, $\sum_{i \in S_j \cup \{j\}} \lambda_i/\beta_i(\mathcal{N}) \leq 1$.

Proof: We first prove (a). Let there exists a session-link i such that

$$\sum_{j \in S_i \cup \{i\}} \frac{\lambda_{q(j)}}{\tilde{\beta}_{q(j)}(\mathcal{N})} > 1.$$

We will show that $\vec{\lambda} \notin \Lambda$.

$$\text{Now, since } \beta_j \leq \tilde{\beta}_{q(j)}, \sum_{j \in S_i \cup \{i\}} \frac{\lambda_{q(j)}}{\beta_j(\mathcal{N})} > 1.$$

Now, note that $K_i(\mathcal{N}) \leq \beta_j(\mathcal{N})$ for every session-link $j \in S_i \cup \{i\}$. This is because if $j \in S_i$, then $i \in S_j$. Thus,

$$\begin{aligned} \sum_{j \in S_i \cup \{i\}} \frac{\lambda_{q(j)}}{K_i(\mathcal{N})} &> 1. \\ \Rightarrow \sum_{j \in S_i \cup \{i\}} \lambda_{q(j)} &> K_i(\mathcal{N}). \end{aligned} \tag{12}$$

Now consider an arbitrary scheduling policy π. Under π, $\sum_{j \in S_i \cup \{i\}} D_j(n) \leq nK_i(\mathcal{N})$ for every $n \geq 0$ as at most $K_i(\mathcal{N})$ nodes among $S_i \cup \{i\}$ can be scheduled concurrently.

$$\begin{aligned} \text{Thus, } \liminf_{n \to \infty} \sum_{j \in S_i \cup \{i\}} \frac{D_j(n)}{n} &\leq K_i(\mathcal{N}) \\ \Rightarrow \sum_{j \in S_i \cup \{i\}} \liminf_{n \to \infty} \frac{D_j(n)}{n} &\leq K_i(\mathcal{N}) \\ &< \sum_{j \in S_i \cup \{i\}} \lambda_{q(j)} \text{ (from (12))}. \\ \Rightarrow \liminf_{n \to \infty} \frac{D_j(n)}{n} &< \lambda_{q(j)} \text{ for some } j \in S_i \cup \{i\} \\ \Rightarrow \liminf_{n \to \infty} \frac{D_{L_j}(n)}{n} &< \lambda_{q(j)}. \end{aligned}$$

The last inequality follows since $D_{L_j}(n) \leq D_j(n)$ for all j, n. Thus, if $\lim_{n \to \infty} \frac{D_{L_j}(n)}{n}$ exists, then its value is less than $\lambda_{q(j)}$. Thus, the network is not stable under π. Alternatively, if the limit does not exist, then also the network is not stable under π. Thus, $\vec{\lambda} \notin \Lambda$. The result follows.

When each session spans one link, sessions and session-links are identical, $M = N$, $q(j) = j$, $\tilde{\beta}_{q(j)}(\mathcal{N}) = \beta_j(\mathcal{N})$. Thus, (b) follows from (a). ■

Lemma 18 follows from part (b) of Lemma 20 since $K(\mathcal{N}) \geq \beta_i(\mathcal{N})$ for all i.

Appendix A.2: Proof of Lemma 19:

Recall that $Q_i(n)$ denotes the queue length of session i in the beginning of the n^{th} slot. Then, for any scheduling policy,

$$Q_i(n+1) = Q_i(0) + A_i(n) - D_i(n) \; \forall \, n \geq 1 \text{ and } i = 1, \ldots, N. \tag{13}$$

We first define fluid limits. The definitions are similar to those used by Dai *et. al.* [4].

Appendix A.2.a: Definition of Fluid Limits: We denote by $\mathbb{N}$ and $\mathbb{R}$ the set of non-negative integers and reals respectively. For a random process $\{f(t)\}_{t\geq 0}$, we denote its value at time t along a sample path ω by $f(t,\omega)$.

Note that the domain of the functions $A(\cdot)$, $D(\cdot)$ and $Q(\cdot)$ is $\mathbb{N}$. Now, we define these functions for arbitrary $t \in \mathbb{R}$ by using a piecewise linear interpolation. The piecewise linear interpolation of a function $f : \mathbb{N} \to \mathbb{R}$ is defined as follows. For $t \in (n, n+1]$

$$f(t) = f(n) + (t-n)\left(f(n+1) - f(n)\right).$$

Note that $f(t)$ defined as above is a continuous function.

Consider any scheduling policy. From any sender i, at most one packet can be served in a slot. Also, the maximum number of packets arriving in a slot at i is bounded by $\alpha_{\max}$. Thus, for every i, ω, $t \geq 0$ and $\delta > 0$

$$A_i(t+\delta,\omega) - A_i(t,\omega) \leq \delta\alpha_{\max}, \tag{14}$$
$$D_i(t+\delta,\omega) - D_i(t,\omega) \leq \delta, \tag{15}$$
$$Q_i(t+\delta,\omega) - Q_i(t,\omega) \leq \delta\alpha_{\max}. \tag{16}$$

Now, let us define a family of functions for any given function $f(\cdot)$ as follows.

$$f^r(t,\omega) \stackrel{\text{def}}{=} \frac{f(rt,\omega)}{r} \text{ for every } \; r > 0.$$

It follows from (14), (15) and (16), that for every $r > 0$,

$$A_i^r(t+\delta,\omega) - A_i^r(t,\omega) \leq \delta\alpha_{\max}, \tag{17}$$
$$D_i^r(t+\delta,\omega) - D_i^r(t,\omega) \leq \delta, \tag{18}$$
$$Q_i^r(t+\delta,\omega) - Q_i^r(t,\omega) \leq \delta\alpha_{\max}. \tag{19}$$

Thus, all the above functions are Lipschitz continuous, and hence uniformly continuous on any compact interval. Clearly, the above functions are also bounded on any compact interval. Fix a compact interval $[0,t]$. Now, consider any sequence r_n such that $r_n \to \infty$ as $n \to \infty$. Then, by Arzela-Ascoli Theorem [10], there exists a subsequence r_{n_k} and continuous functions $\overline{A}_i(\cdot), \overline{D}_i(\cdot)$ and $\overline{Q}_i(\cdot)$ such that for every i, ω,

$$\lim_{k\to\infty} \sup_{\hat{t}\in[0,t]} \left|A_i^{r_{n_k}}(\hat{t},\omega) - \overline{A}_i(\hat{t},\omega)\right| = 0, \tag{20}$$
$$\lim_{k\to\infty} \sup_{\hat{t}\in[0,t]} \left|D_i^{r_{n_k}}(\hat{t},\omega) - \overline{D}_i(\hat{t},\omega)\right| = 0, \tag{21}$$
$$\lim_{k\to\infty} \sup_{\hat{t}\in[0,t]} \left|Q_i^{r_{n_k}}(\hat{t},\omega) - \overline{Q}_i(\hat{t},\omega)\right| = 0. \tag{22}$$

We now define fluid limits.

Definition 19: Any $(\overline{A}_i, \overline{D}_i, \overline{Q}_i)$ is called a fluid limit for $\mathcal{N}$ if there exists r_{n_k} such that all the relations (20) to (22) are satisfied.

Now, we state some important properties of the fluid limits which we use to prove Lemma 19.

Lemma 21: Every fluid limit satisfies, $\overline{A}_i(t) = \lambda_i t$ w.p. 1 for every session i and $t \geq 0$.

Lemma 22: Any fluid limit $(\overline{A}_i, \overline{D}_i, \overline{Q}_i)$ for $\mathcal{N}$ satisfies the following equality for every i and $t \geq 0$ with probability (w.p.) 1:

$$\overline{Q}_i(t) = \overline{Q}_i(0) + \lambda_i t - \overline{D}_i(t). \tag{23}$$

Lemma 23: Let $\overline{Q}_i(0) = 0$ for every i. Also, let $\sum_{j\in S_i\cup\{i\}} \lambda_j \leq 1$ if $\lambda_i > 0, \; i = 1,\ldots,N$. Then, under maximal scheduling, every fluid limit satisfies, $\overline{Q}_i(t) = 0$ for every $t \geq 0$ w.p. 1 for every i.

The proofs of Lemmas 21, 22, 23 are provided later, after the proof Lemma 19. We now prove Lemma 19.

Proof: First, we show that $\lim_{r\to\infty} D_i^r(t) = \lambda_i t$ w.p. 1 for every t. Then, the result follows by choosing $t = 1$.

Under maximal scheduling, if $\overline{Q}_i(0) = 0$ and $\sum_{j\in S_i\cup\{i\}} \lambda_i \le 1$ for every i for which $\lambda_i > 0$, then $\overline{Q}_i(t) = 0$ w.p. 1 for every i and $t \ge 0$ (Lemma 23). Thus, by Lemma 22, $\overline{D}_i(t) = \lambda_i t$ w.p. 1 for every $t \ge 0$. Since $\overline{D}_i(\cdot)$ is a fluid limit, there exists a subsequence r_{n_k} such that $\lim_{k\to\infty} r_{n_k} = \infty$ and $\lim_{k\to\infty} D_i^{r_{n_k}}(t) = \overline{D}_i(t) = \lambda_i t$ w.p. 1 (Section A.2.a). Thus, $\liminf_{r\to\infty} D_i^r(t) \le \lambda_i t$ w.p. 1. Now, we argue that $\liminf_{r\to\infty} D_i^r(t) = \lambda_i t$ w.p. 1.

Suppose, $\liminf_{r\to\infty} D_i^r(t) < \lambda_i t$ w.p. 1. Then, there exists a subsequence $\widehat{r}_{n_k}$ such that $\lim_{k\to\infty} \widehat{r}_{n_k} = \infty$ and $\lim_{k\to\infty} D_i^{\widehat{r}_{n_k}}(t) = \lambda_i t - \epsilon$ w.p. 1 for some $\epsilon > 0$. Now, note that

$$Q^{\widehat{r}_{n_k}}(t) = Q^{\widehat{r}_{n_k}}(0) + A^{\widehat{r}_{n_k}}(t) - D^{\widehat{r}_{n_k}}(t) \text{ (from (13))}.$$

Now, by taking limit as $k \to \infty$ on both sides of the above equation we obtain

$$\begin{aligned}\overline{Q}_i^1(t) &= \overline{Q}_i^1(0) + \lambda_i t - \overline{D}_i^1(t) \quad \text{w.p. 1 (from Lemma 21)}\\ &= \epsilon, \text{ (since } \overline{D}_i^1(t) = \lim_{k\to\infty} D_i^{\widehat{r}_{n_k}}(t) = \lambda_i t - \epsilon).\end{aligned}$$

Since, $\overline{Q}_i^1(t)$ is also a fluid limit under maximal scheduling, the above equation contradicts Lemma 23. Thus,

$$\liminf_{r\to\infty} D_i^r(t) = \lambda_i t \quad \text{w.p. } 1.$$

Now, for every $r > 0$, $D_i^r(t) \le A_i^r(t)$ as the number of departures from i can at most be equal to the arrivals for i till time rt. Thus, clearly,

$$\limsup_{r\to\infty} D_i^r(t) \le \lambda_i t \quad \text{w.p. } 1.$$

This shows that

$$\lim_{r\to\infty} D_i^r(t) = \lambda_i t \quad \text{w.p. } 1.$$

Now, select $t = 1$, and consider subsequence r_n such that $r_n = n$. Here, for every i

$$\begin{aligned}\lim_{n\to\infty} D_i^{r_n}(1) &= \lambda_i \quad \text{w.p. 1}\\ \lim_{n\to\infty} \frac{D_i(n)}{n} &= \lambda_i \quad \text{w.p. } 1.\end{aligned}$$

■

We now prove the supporting lemmas used to prove Lemma 19.

Appendix A.2.b: Proof of Lemma 21:

Proof: Since $\overline{A}_i(t)$ is a fluid limit, by Definition 19, there exists a sequence r_{n_k} such that $\lim_{k\to\infty} r_{n_k} = \infty$ and

$$\begin{aligned}\overline{A}_i(t) &= \lim_{k\to\infty} A_i^{r_{n_k}}(t) \quad \text{(from (20))}\\ &= \lim_{k\to\infty} \frac{A_i(r_{n_k} t)}{r_{n_k}}\\ &= \lim_{k\to\infty} \frac{A_i(r_{n_k} t)}{r_{n_k} t} t\\ &= \lambda_i t \quad \text{w.p. 1 (since } A_i(\cdot) \text{ satisfy SLLN).}\end{aligned}$$

The result follows. ■

Appendix A.2.c: Proof of Lemma 22:

Proof: Since $\overline{Q}_i(\cdot)$, $\overline{A}_i(\cdot)$ and $\overline{D}_i(\cdot)$ are fluid limits, there exists a sequence r_{n_k} such that $\lim_{k\to\infty r_{n_k}} = \infty$ and they are obtained as a uniform limits of functions $Q_i^{r_{n_k}}(\cdot)$, $A_i^{r_{n_k}}(\cdot)$ and $D_i^{r_{n_k}}(\cdot)$ respectively. Now, from (13) it follows that for every r_{n_k} and $t \geq 0$,

$$Q_i^{r_{n_k}}(t) = Q_i^{r_{n_k}}(0) + A_i^{r_{n_k}}(t) - D_i^{r_{n_k}}(t).$$

The result follows from Lemma 21 after taking the limit $k \to \infty$ on both sides of the above equality. ■

Appendix A.2.d: Proof of Lemma 23:

Proof: We prove the required by contradiction. Let $\overline{Q}_i(t) \neq 0$ for every t and i. Then, there exists a session i, $\widehat{t}$, $y_1 > 0$ and $x_1 > 0$ such that

$$\sum_{j\in S_i\cup\{i\}} \overline{Q}_j(\widehat{t}) \;=\; y_1, \tag{24}$$

$$\sum_{j\in S_i\cup\{i\}} \overline{Q}_j(t) \;<\; y_1 \text{ for every } t \in [0,\widehat{t}), \tag{25}$$

$$\overline{Q}_i(\widehat{t}) \;=\; x_1. \tag{26}$$

We justify (24) to (26) by constructing $x_1, y_1, \hat{t}$ that satisfy (24) to (26). Let $t' = \inf\{t : t \geq 0, \max_k \bar{Q}_k(t) > 0\}$. Since $\bar{Q}_k(t) \neq 0$ for some t and some k, t' is well-defined. From the definition of t' there exists an i such that $t' = \inf\{t : t \geq 0, \bar{Q}_i(t) > 0\}$. From the continuity of $\bar{Q}_k(t)$ for all t, k, the definition of t', and since $\bar{Q}_k(0) = 0$ for all k, $\bar{Q}_k(t_1) = 0$ for all $t_1 \leq t'$ and k. From the continuity of $\bar{Q}_i(t)$ for all t, there exists an $\epsilon > 0$ s.t. $\sum_{j\in S_i\cup\{i\}} \bar{Q}_j(t) \geq \bar{Q}_i(t) > 0$ for all $t \in (t', t'+\epsilon]$. Let $y_1 = \max_{t:t\in[0,t'+\epsilon]} \sum_{j\in S_i\cup\{i\}} \bar{Q}_j(t)$. Let $\hat{t}$ be the first time at which $\sum_{j\in S_i\cup\{i\}} \bar{Q}_j(t) = y_1$. Now, $\hat{t} \in (t', t'+\epsilon]$, since $\bar{Q}_k(t_1) = 0$ for all k and all $t_1 \leq t'$, and $\sum_{j\in S_i\cup\{i\}} \bar{Q}_j(t) \geq \bar{Q}_i(t) > 0$ for all $t \in (t', t'+\epsilon]$. Let $x_1 = \bar{Q}_i(\hat{t})$. Clearly, $x_1 > 0$.

Let $\lambda_i \leq 0$. From Lemma 22, since $\bar{Q}_i(0) = 0$, $\bar{Q}_i(\hat{t}) \leq -\bar{D}_i(\hat{t})$. Since $\bar{D}_i(.)$ is the fluid limit of $D_i(.)$, and $D_i(t) \geq 0$ at all t, $\bar{D}_i(\hat{t}) \geq 0$. Thus, $x_1 \leq 0$, which is a contradiction. Thus, $\lambda_i > 0$, and hence, $\sum_{j\in S_i\cup\{i\}} \lambda_j \leq 1$.

Clearly, $x_1 \leq y_1$ as $\overline{Q}_j(\cdot) \geq 0$ for every j. Since $\overline{Q}_i(\cdot)$ is a continuous function, there exists $t' \in [0,\widehat{t})$ such that

$$\overline{Q}_i(t) \geq \frac{x_1}{2} \text{ for every } t \in [t', \widehat{t}]. \tag{27}$$

Now, since $\overline{Q}_j(\cdot)$ is a fluid limit, by Definition 19, there exists a sequence r_{n_k} such that $\lim_{k\to\infty} r_{n_k} = \infty$ and $\lim_{k\to\infty} Q_j^{r_{n_k}}(t) = \overline{Q}_j(t)$ for every j and t in an interval $[0,\widehat{t}]$. Thus, we can draw two conclusions. First, for sufficiently large r_{n_k}, $Q_i^{r_{n_k}}(t) > x_1/4$ for every $t \in [t', \widehat{t}]$. Thus, $Q_i(r_{n_k}t) > r_{n_k}x_1/4$. This implies that for every $r_{n_k} > 4/x_1$,

$$Q_i(r_{n_k}t) > 1 \text{ for every } t \in [t', \widehat{t}]. \tag{28}$$

The second conclusion is that for every sufficiently large r_{n_k}, there exists $\epsilon > 0$ such that

$$\begin{aligned} \sum_{j\in S_i\cup\{i\}} Q_j^{r_{n_k}}(\widehat{t}) - \sum_{j\in S_i\cup\{i\}} Q_j^{r_{n_k}}(t') \;&>\; \epsilon, \\ \Rightarrow \lim_{k\to\infty}\left[\sum_{j\in S_i\cup\{i\}} Q_j^{r_{n_k}}(\widehat{t}) - \sum_{j\in S_i\cup\{i\}} Q_j^{r_{n_k}}(t')\right] \;&\geq\; \epsilon. \end{aligned} \tag{29}$$

Relation (29) follows from (24), (25), $t' < \widehat{t}$ and the definition of fluid limits. Select r_{n_k} large enough such that (28) holds. For all such r_{n_k},

$$\begin{aligned} &\sum_{j \in S_i \cup \{i\}} Q_j^{r_{n_k}}(\widehat{t}) - \sum_{j \in S_i \cup \{i\}} Q_j^{r_{n_k}}(t') \\ &= \sum_{j \in S_i \cup \{i\}} \left[A_j^{r_{n_k}}(\widehat{t}) - A_j^{r_{n_k}}(t')\right] - \left[\sum_{j \in S_i \cup \{i\}} D_j^{r_{n_k}}(\widehat{t}) - \sum_{j \in S_i \cup \{i\}} D_j^{r_{n_k}}(t')\right] \text{ (from (13)).} \qquad (30) \end{aligned}$$

Since maximal scheduling is used and (28) holds, at least one packet from some session in $S_i \cup \{i\}$ departs in every slot. Thus, $\sum_{j \in S_i \cup \{i\}} D_j^{r_{n_k}}(\widehat{t}) - \sum_{j \in S_i \cup \{i\}} D_j^{r_{n_k}}(t') \geq (\widehat{t} - t')$. Now, from (30),

$$\begin{aligned} \sum_{j \in S_i \cup \{i\}} Q_j^{r_{n_k}}(\widehat{t}) - \sum_{j \in S_i \cup \{i\}} Q_j^{r_{n_k}}(t') &\leq \sum_{j \in S_i \cup \{i\}} \left[A_j^{r_{n_k}}(\widehat{t}) - A_j^{r_{n_k}}(t')\right] - (\widehat{t} - t') \\ \Rightarrow \lim_{k \to \infty} \left[\sum_{j \in S_i \cup \{i\}} Q_j^{r_{n_k}}(\widehat{t}) - \sum_{j \in S_i \cup \{i\}} Q_j^{r_{n_k}}(t')\right] &\leq \lim_{k \to \infty} \sum_{j \in S_i \cup \{i\}} \left[A_j^{r_{n_k}}(\widehat{t}) - A_j^{r_{n_k}}(t')\right] - (\widehat{t} - t') \\ &= \left(\sum_{j \in S_i \cup \{i\}} \lambda_j - 1\right)(\widehat{t} - t') \text{ w.p. 1 (from Lemma 21)} \\ &\leq 0. \qquad (31) \end{aligned}$$

Note that (31) contradicts (29). Thus, the result follows. ■

Appendix B: Proof of Theorem 2

Proof: Consider an arbitrary network $\mathcal{N}$ with interference degree $K(\mathcal{N})$. By Definition 13, there exists i such that the interference degree of session i is $K(\mathcal{N})$. Consider sessions $j_1, \ldots, j_{K(\mathcal{N})} \in S_i$ such that they are pair-wise non-interfering. Now, consider the following arrival rate vector $\vec{\lambda}$: $\lambda_j = Z/K(\mathcal{N})$ if $j \in \{j_1, \ldots, j_{K(\mathcal{N})}\}$, and $\lambda_j = (K(\mathcal{N}) - Z)/K(\mathcal{N})$ if $j = i$, and $\lambda_j = 0$ otherwise. Thus, effectively the network consists only of sessions i and $j_1, \ldots, j_{K(\mathcal{N})}$. Note that since $1 \leq Z < K(\mathcal{N})$, $\lambda_j > 0$ for every $j \in \{i, j_1, \ldots, j_{K(\mathcal{N})}\}$. Now, consider a scheduling policy π that schedules i w.p. $(K(\mathcal{N}) - Z)/K(\mathcal{N})$ and sessions $j_1, \ldots, j_{K(\mathcal{N})}$ concurrently in the remaining slots. Clearly, π is rate stable. Thus, $\vec{\lambda} \in \Lambda$.

Now, consider arrival rate vector $\vec{\lambda}/Z$ and the following arrival pattern. A packet corresponding to session j_u arrives in slots t if $u = t \bmod K(\mathcal{N}) + 1$, where "mod" is the modulo operator. In every slot a packet arrives w.p. $(K(\mathcal{N}) - Z)/K(\mathcal{N})$. Clearly, the arrivals are in accordance with $\vec{\lambda}/Z$. Let maximal scheduling schedule i only when none of the sessions in S_i have a packet to transmit. Note that under maximal scheduling and the described arrival pattern, j_u is scheduled in slot t such that $u = t \bmod K(\mathcal{N}) + 1$, and thus i is never scheduled. Since $\lambda_i/Z > 0$, i is not stable. Thus, $\vec{\lambda}/Z \notin \Lambda^{MS}$. ■

Appendix C: Proof of Lemma 5

Proof: Consider a network $\mathcal{N}$ that has bidirectional communication and underlying topology $G = (V, E)$. Select a session i from u to v. Since we are considering bidirectional communication, $(u, v) \in E$ and $(v, u) \in E$. Note that at most one session along every link from u and v, and every link to u and v can be scheduled concurrently in the interference region of i without interfering with each other. Let $d_{(u,v)}$

denote the degree of link (u, v). Now, i's interference degree $k_i(\mathcal{N})$ satisfies the following inequality.

$$\begin{aligned} k_i(\mathcal{N}) &\leq \sum_{\substack{j\in V\\ j\neq v}} \left[\mathbf{1}_{\{(j,u)\in E\}} + \mathbf{1}_{\{(u,j)\in E\}}\right] + \sum_{\substack{j\in V\\ j\neq u}} \left[\mathbf{1}_{\{(j,v)\in E\}} + \mathbf{1}_{\{(v,j)\in E\}}\right] \\ &= \sum_{j\in V} \left[\mathbf{1}_{\{(j,u)\in E\}} + \mathbf{1}_{\{(u,j)\in E\}}\right] + \sum_{j\in V} \left[\mathbf{1}_{\{(j,v)\in E\}} + \mathbf{1}_{\{(v,j)\in E\}}\right] - 4 \\ &= d_{(u,v)} - 4 \\ \Rightarrow \quad \max_i\{k_i(\mathcal{N})\} &\leq \max_{(u,v)\in E}\{d_{(u,v)}\} - 4 \\ \Rightarrow \quad K(\mathcal{N}) &\leq \delta_G - 4. \end{aligned} \tag{32}$$

Now, Fig. 5(a) shows an example of a network that achieves the equality in (32). ■

Appendix D: Proof of Lemma 6

Proof: Consider a network $\mathcal{N}$ and with unidirectional communication on underlying topology $G = (V, E)$. Fix a session i from u to v. Since we are considering unidirectional communication, $(u, v) \in E$. Let $\widehat{d}_{(u,v)}$ denote the directional degree of link (u, v). Now, i's interference degree $k_i(\mathcal{N})$ satisfies the following inequality.

$$\begin{aligned} k_i(\mathcal{N}) &\leq \sum_{\substack{j\in V\\ j\neq v}} \mathbf{1}_{\{(u,j)\in E\}} + \sum_{\substack{j\in V\\ j\neq u}} \mathbf{1}_{\{(j,v)\in E\}} \\ &= \sum_{j\in V} \mathbf{1}_{\{(u,j)\in E\}} + \sum_{j\in V} \mathbf{1}_{\{(j,v)\in E\}} - 2 \\ &= \widehat{d}_{(u,v)} - 2 \\ \Rightarrow \quad \max_i\{k_i(\mathcal{N})\} &\leq \max_{(u,v)\in E}\{\widehat{d}_{(u,v)}\} - 2 \\ \Rightarrow \quad K(\mathcal{N}) &\leq \Delta_G - 2. \end{aligned} \tag{33}$$

Now, Fig. 5(b) shows an example of a network that achieves the equality in (33). ■

PROOFS OF ANALYTICAL RESULTS IN SECTION IV (LEMMAS 1, 2, 3 AND 4)

Appendix E: Proof of Lemma 1

We prove Lemma 1 by considering an arbitrary session (S_0, T_0, R_0) and showing that K_0, the maximum number of sessions that interfere with S_0 but do not interfere with each other, must satisfy $K_0 \leq 8$. Thus, Lemma 1 follows from Theorem 1.

We assume that the nodes are deployed on a two-dimensional Euclidean plane. Let the distance between the transmitting node T_0 and receiving node R_0 be $\rho \leq r$, where r is the transmission range of any node.

Without loss of generality let us assume that the line joining T_0 and R_0 is aligned along the x-axis. Let D_{T_0} and D_{R_0} represent disks of radius r around T_0 and R_0, respectively. Then the interference area of session S_0 is $D_{T_0} \cup D_{R_0}$.

In the following, a node is said to be the *transceiver node* of a session if it is either the transmitting node or the receiving node of that session; thus each session has two transceiver nodes. Note that if a session interferes with S_0, at least one of its transceiver nodes must lie in $D_{T_0} \cup D_{R_0}$. Now for each of the sessions that interfere with S_0 but do not interfere with each other, choose any one transceiver node of that session that lies in $D_{T_0} \cup D_{R_0}$; let $\mathcal{U}_0$ denote the set of the transceiver nodes thus chosen. We will show $K_0 \leq 8$ by showing $U_0 = |\mathcal{U}_0| \leq 8$.

The proof of $K_0 \leq 8$ is quite involved; therefore, we will first show that $K_0 \leq 9$, the proof of which is considerably simpler. We will then extend our arguments to show that $K_0 \leq 8$.

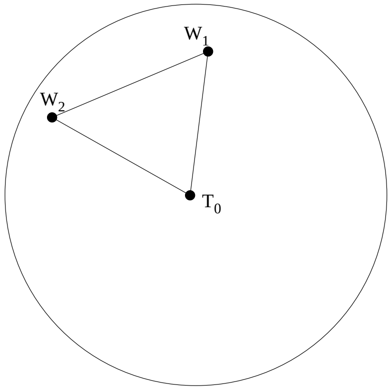

Fig. 9. Diagram used in proof of Lemma 24

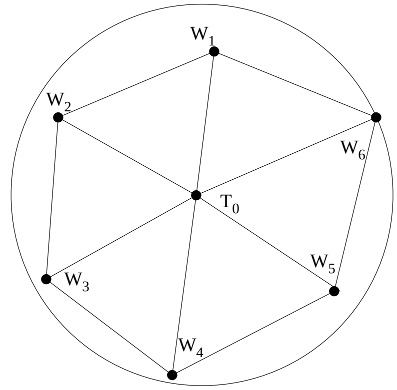

Fig. 10. Diagram used in proof of Lemma 25

Lemma 24: Let $W_1, W_2 \in \mathcal{U}_0$. If $W_1, W_2 \in D_{T_0}$ ($W_1, W_2 \in D_{R_0}$), and none of them coincide with T_0 (R_0), then the line segment joining W_1 and W_2 subtends an angle greater than $\frac{\pi}{3}$ at T_0 (R_0).

Proof: We only consider the case of $W_1, W_2 \in D_{T_0}$ (the $W_1, W_2 \in D_{R_0}$ case is similar). Let $|w_1 w_2|$ denote the length of the line segment joining w_1 and w_2, for any two points w_1, w_2. Refer to Fig. 9. Since $W_1, W_2 \in D_{T_0}$, we have $|W_1 T_0| \leq r$, $|W_2 T_0| \leq r$. Also, since sessions of transceiver nodes W_1 and W_2 do not interfere with each other, we have $|W_1 W_2| > r$. Thus, in triangle $W_1 W_2 T_0$, $W_1 W_2$ is longer than each of the other sides, and its opposite angle $\angle W_1 T_0 W_2$ must be greater than $\frac{\pi}{3}$, from elementary geometry. ■

Fig. 11 shows the area $D_{T_0} \cup D_{R_0}$. Note that $|T_0 R_0| = \rho \leq r$. Let A_1 (A_2) be a point on the circumference of D_{T_0} (D_{R_0}), such that $\angle A_1 T_0 R_0$ $\frac{\pi}{3}$ ($\angle A_2 R_0 T_0$ $\frac{\pi}{3}$). Similarly, let B_1 (B_2) be a point on the circumference of D_{T_0} (D_{R_0}), such that $\angle B_1 T_0 R_0$ $\frac{\pi}{3}$ ($\angle B_2 R_0 T_0$ $\frac{\pi}{3}$). Let line segments $T_0 A_1$ and $R_0 A_2$ intersect at A, and line segments $T_0 B_1$ and $R_0 B_2$ intersect at B.

Recall that $T_0 R_0$ is aligned along the x-axis. Let points C, D, E be points on the circumference of D_{T_0} such that CT_0, DT_0 and ET_0 subtend angles of $\frac{2\pi}{3}$, π and $\frac{4\pi}{3}$ with the x-axis, respectively. Also, let points F, G, H be points on the circumference of D_{R_0} such that FR_0, GR_0 and HR_0 subtend angles of $\frac{\pi}{3}$, 0 and $\frac{5\pi}{3}$ with the x-axis, respectively. Let P_1 (P_2) denote the points at which the line $T_0 R_0$ extended ($R_0 T_0$ extended) intersects the circumference of D_{T_0} (D_{R_0}). Thus, line segments $A_1 T_0$, CT_0, DT_0, ET_0, $B_1 T_0$ and $P_1 T_0$ divide D_{T_0} into six $\frac{\pi}{3}$ sectors. Similarly, line segments $A_2 R_0$, FR_0, GR_0, HR_0, $B_2 R_0$ and $P_2 R_0$ divide D_{R_0} into six $\frac{\pi}{3}$ sectors. From Lemma 24, it follows that each of these sectors can contain at most one node in $\mathcal{U}_0$.

Lemma 25: The number of nodes in $\mathcal{U}_0$ that lie in D_{T_0} (D_{R_0}) can be no greater than 5.

Proof: We only consider the case of D_{T_0} (the case of D_{R_0} is similar). Let $\hat{U}$ denote the number of nodes in $\mathcal{U}_0$ that lie in D_{T_0}. Since D_{T_0} is contained in six $\frac{\pi}{3}$ sectors, $\hat{U} \leq 6$.

For the sake of contradiction let us assume that $\hat{U} = 6$, and let $W_i, i = 1, \ldots, 6$ denote the six nodes in $\mathcal{U}_0$ that lie in D_{T_0}, as shown in Fig. 10. Note that none of these nodes can lie at the center of D_{T_0}, i.e., at T_0. Then, from Lemma 24, the angle subtended at T_0 by each of the line segments $W_i W_j$, $j = (i+1) \bmod 6$, is greater than $\frac{\pi}{3}$. Since the total angle subtended at T_0 cannot exceed 2π, we have a contradiction, thereby proving the lemma. ■

From Fig. 11, note that $D_{R_0} \setminus D_{T_0}$ is contained in four $\frac{\pi}{3}$ sectors. Therefore, at most 4 nodes in $\mathcal{U}_0$ can lie in $D_{R_0} \setminus D_{T_0}$. Since at most 5 nodes in $\mathcal{U}_0$ can lie in D_{T_0}, we have the following result:

Corollary 1: U_0, the number of nodes in $\mathcal{U}_0$ can be no greater than 9.

The above result implies that $K_0 \leq 9$. Now we proceed to tighten this upper bound by showing $K_0 \leq 8$.

Now let us assume, for the sake of contradiction, that $K_0 = 9$; this implies $U_0 = 9$.

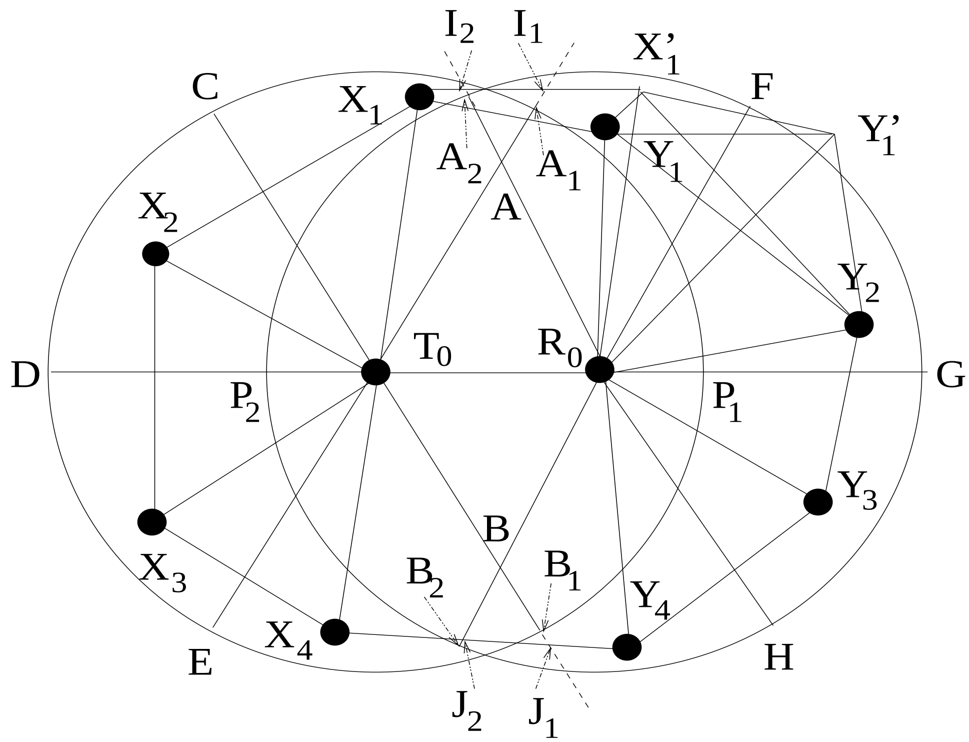

Fig. 11. Diagram used in proof of Lemma 1

Corollary 2: If $U_0 = 9$, then the number of nodes in $\mathcal{U}_0$ that lie in $D_{T_0} \setminus D_{R_0}$, $D_{T_0} \cap D_{R_0}$, and $D_{R_0} \setminus D_{T_0}$ are 4, 1 and 4, respectively.

Proof: Let U_1, U_2 and U_3 respectively denote the nodes in $\mathcal{U}_0$ that lie in $D_{T_0} \setminus D_{R_0}$, $D_{T_0} \cap D_{R_0}$, and $D_{R_0} \setminus D_{T_0}$. Then, $U_1 + U_2 + U_3 = 9$. Without loss of generality, assume $U_1 \geq U_3$.

We first argue that $U_2 \neq 0$. Note that if $U_2 = 0$, then $U_1 + U_3 = 9$, implying $U_1 \geq 5$, which is impossible since $D_{T_0} \setminus D_{R_0}$ is contained in four $\frac{\pi}{3}$ sectors. This implies that $U_2 > 0$.

Now we argue that $U_2 \leq 1$. Let us assume, for the sake of contradiction, that $U_2 \geq 2$. Then, $U_1 + U_3 = 9 - U_2 \leq 7$. Thus $U_3 \leq 3$. Therefore, $U_1 + U_2 = 9 - U_3 \geq 6$, which is impossible (from Lemma 25). Therefore, $U_2 \leq 1$. Since $U_2 > 0$ (as shown previously), we have $U_2 = 1$.

Therefore, $U_1 + U_3 = 8$. Since $U_1 \leq 4, U_3 \leq 4$ (each of $D_{T_0} \setminus D_{R_0}$ and $D_{R_0} \setminus D_{T_0}$ are contained in four $\frac{\pi}{3}$ sectors), we must have $U_1 = U_3 = 4$. ■

From the above lemma, we see that if $K_0 = 9$, then $D_{T_0} \setminus D_{R_0}$ and $D_{R_0} \setminus D_{T_0}$ must each contain 4 nodes in $\mathcal{U}_0$. For the sake of contradiction, let us assume that this is true. Note that none of these 8 nodes can lie at the centers of the two disks, i.e., at T_0 or R_0. Also, exactly one of these 8 points must lie in each of the $\frac{\pi}{3}$ sectors of $D_{T_0} \setminus D_{R_0}$ and $D_{R_0} \setminus D_{T_0}$. Let X_1, X_2, X_3 and X_4 respectively denote the nodes in $\mathcal{U}_0$ that lie in sectors A_1T_0C, CT_0D, DT_0E and ET_0B_1. Let Y_1, Y_2, Y_3 and Y_4 respectively denote the nodes in $\mathcal{U}_0$ that lie in sectors A_2R_0F, FR_0G, GR_0H and HR_0B_2. Join X_1, X_2, X_3, X_4 with T_0, and Y_1, Y_2, Y_3, Y_4 with R_0 (refer to Fig. 11). Now, construct the octagon by joining $X_1X_2, X_2X_3, X_3X_4, Y_1Y_2, Y_2Y_3, Y_3Y_4$, and X_1Y_1, X_4Y_4. Note that the length of each side of this octagon must be greater than r. Let line segment X_1Y_1 intersect line segments T_0A and R_0A (possibly extended) at points I_1 and I_2, respectively. Let line segment X_4Y_4 intersect line segments T_0B and R_0B (possibly extended) at points J_1 and J_2, respectively.

Note that the angle subtended at T_0 by $I_1X_1X_2X_3X_4J_1$ (which is a collection of the line segments $I_1X_1, X_1X_2, \ldots, X_4J_1$), is equal to $\frac{4\pi}{3}$. Similarly, the angle subtended at R_0 by $I_2Y_1Y_2Y_3Y_4J_2$ (which is a collection of the line segments $I_2Y_1, Y_1Y_2, \ldots, Y_4J_2$), is equal to $\frac{4\pi}{3}$. In the following, we show however that the angle subtended at T_0 by $I_1X_1X_2X_3X_4J_1$ plus the angle subtended at R_0 by $I_2Y_1Y_2Y_3Y_4J_2$ must be greater than $\frac{8\pi}{3}$, thus arriving at a contradiction.

We will show that the angle subtended by $X_2X_1I_1$ at T_0 plus the angle subtended by $I_2Y_1Y_2$ at R_0 is greater than π. Without loss of generality, assume that X_1 has a higher y-coordinate than Y_1 (recall that T_0R_0 is aligned along the x-axis). As shown in Fig. 11, choose $X_1^{'}$ such that $X_1T_0R_0X_1^{'}$ is a parallelogram. Join $X_1^{'}$ with Y_1 and Y_2. Note, $\angle X_1T_0I_1 + \angle I_2R_0X_1^{'} = \pi - \angle I_1T_0R_0 - \angle I_2R_0T_0 = \pi - \frac{\pi}{3} - \frac{\pi}{3} = \frac{\pi}{3}$.

We consider the following two cases separately: (i) Y_1 lies within parallelogram $X_1T_0R_0X_1^{'}$, and (ii) Y_1 lies outside parallelogram $X_1T_0R_0X_1^{'}$. Let us consider case (i) first (Fig. 11 shows this case). In this case, we claim that $\angle X_1^{'}R_0Y_2 > \frac{\pi}{3}$. To see this, choose $Y_1^{'}$ such that $X_1Y_1Y_1^{'}X_1^{'}$ is a parallelogram. Join $Y_1^{'}$ with R_0 and Y_2. Note, $|X_1^{'}Y_1^{'}| = |X_1Y_1| > r$. Note that Y_2 must lie “below” $Y_1Y_1^{'}$, since it is easy to see that there is no point in sector FR_0G that is “above” $Y_1Y_1^{'}$ and whose distance from Y_1 is greater than r.

Note that $|Y_1Y_1^{'}| = |X_1X_1^{'}| = |T_0R_0| = \rho$ (by construction). Therefore, it is easy to see that $Y_1^{'}$ must lie outside D_{R_0}. Thus, line segment $X_1^{'}Y_2$ must intersect line segment $Y_1Y_1^{'}$. In triangle $Y_1Y_2Y_1^{'}$, $|Y_1Y_2| > r$ and $|Y_1Y_1^{'}| = \rho \leq r$. Therefore, $\angle Y_1Y_1^{'}Y_2 > \angle Y_1Y_2Y_1^{'}$. Thus, $\angle X_1^{'}Y_1^{'}Y_2 \geq Y_1Y_1^{'}Y_2 > \angle Y_1Y_2Y_1^{'} \geq \angle X_1^{'}Y_2Y_1^{'}$. Thus, comparing angles in triangle $X_1^{'}Y_2Y_1^{'}$, we get $|X_1^{'}Y_2| > |X_1^{'}Y_1^{'}| > r$.

Note that since X_1 lies in sector CT_0A_1, it follows that $X_1^{'}$ must lie in sector A_2R_0F. Therefore, $X_1^{'}$ lies in D_{R_0}. In triangle $X_1^{'}R_0Y_2$, therefore, we have $|X_1^{'}R_0| \leq r, |Y_2R_0| \leq r$, and $|X_1^{'}Y_2| > r$. Therefore, $\angle X_1^{'}R_0Y_2 > \frac{\pi}{3}$.

Thus, if Y_1 lies in the parallelogram $X_1T_0R_0X_1^{'}$, we have $\angle X_1T_0I_1 + \angle I_2R_0X_1^{'} + \angle X_1^{'}R_0Y_2 > \frac{\pi}{3} + \frac{\pi}{3} = \frac{2\pi}{3}$. Moreover, since $\angle I_2R_0X_1^{'} + \angle X_1^{'}R_0Y_2 = \angle I_2R_0Y_1 + \angle Y_1R_0Y_2$, we have $\angle X_1T_0I_1 + \angle I_2R_0Y_1 + \angle Y_1R_0Y_2 > \frac{2\pi}{3}$. From Lemma 24, $\angle X_2T_0X_1 > \frac{\pi}{3}$. Therefore, $\angle X_2T_0X_1 + \angle X_1T_0I_1 + \angle I_2R_0Y_1 + \angle Y_1R_0Y_2 > \pi$. In other words, the angle subtended by $X_2X_1I_1$ at T_0 plus the angle subtended by $I_2Y_1Y_2$ at R_0 is greater than π.

Now let us consider the case where Y_1 does not lie inside parallelogram $X_1T_0R_0X_1^{'}$. Since Y_1 has a lower y-coordinate than X_1, it follows that Y_1 must lie below the line $X_1X_1^{'}$. Thus Y_1 must lie to the “right” of line $R_0X_1^{'}$. Thus, $\angle X_1T_0I_1 + \angle I_2R_0Y_1 > \angle X_1T_0I_1 + I_2R_0X_1^{'} = \frac{\pi}{3}$. From Lemma 24, we get $\angle X_2T_0X_1 > \frac{\pi}{3}, \angle Y_1R_0Y_2 > \frac{\pi}{3}$. Therefore, we obtain $\angle X_2T_0X_1 + \angle X_1T_0I_1 + \angle I_2R_0Y_1 + \angle Y_1R_0Y_2 > \pi$, implying that the angle subtended by $X_2X_1I_1$ at T_0 plus the angle subtended by $I_2Y_1Y_2$ at R_0 is greater than π.

Using similar arguments as above, it follows that the angle subtended by $X_3X_4J_1$ at T_0 plus the angle subtended by $J_2Y_4Y_3$ is greater than π. From Lemma 24, we obtain $\angle X_2T_0X_3 \geq \frac{\pi}{3}, \angle Y_2R_0Y_3 \geq \frac{\pi}{3}$. Combining all of the above results, we see that the angle subtended at T_0 by $I_1X_1X_2X_3X_4J_1$ plus the angle subtended at R_0 by $I_2Y_1Y_2Y_3Y_4J_2$ must be greater than $\pi + \pi + \frac{\pi}{3} + \frac{\pi}{3} = \frac{8\pi}{3}$. Thus we arrive at a contradiction showing that our assumption that $K_0 = 9$ was incorrect. Therefore $K_0 \leq 8$. ■

Appendix F: Proof of Lemma 2

Proof: Figure 2(a) shows a network $\mathcal{N}$ with bidirectional equal power model such that $K(\mathcal{N}) = 8$. Thus, the lemma follows immediately from Theorem 2. ■

Appendix G: Proof of Lemma 3

Proof: Consider any constant Z. In the network $\mathcal{N}$ of Fig. 2(b), for $\theta < 2\pi/(Z+2)$, $K(\mathcal{N}) > Z$ under unidirectional equal power model. Thus, the lemma follows immediately from Theorem 2. ■

Appendix H: Proof of Lemma 4

Proof: Fig. 3 shows an example of a network $\mathcal{N}$ under node exclusive spectrum sharing model with $K(\mathcal{N}) = 2$. Thus, the lemma follows immediately from Theorem 2. ■

PROOFS FOR ANALYTICAL RESULTS IN SECTION VI-A (LEMMAS 7 AND 8)

Let $\mathcal{G}_i$ denote the set of receivers for session i and let u denote the receiver for session i. Also, let $k_i(\mathcal{N})$ denote the interference degree of session i.

Appendix I: Proof for Lemma 7

Proof: 1) Consider network $\mathcal{N}$ with multicast sessions and bidirectional communication model. Since only one session can be scheduled along any link in a slot, the multicast degree of any session i (denoted by M_i) satisfies the following relation.

$$\begin{aligned} k_i(\mathcal{N}) &\leq M_i \\ \Rightarrow \quad \max_i\{k_i(\mathcal{N})\} &\leq \max_i\{M_i\} \\ \Rightarrow \quad K(\mathcal{N}) &\leq \gamma(\mathcal{N}). \end{aligned}$$

2) Consider network $\mathcal{N}$ with multicast sessions and unidirectional communication model. Note that i can interfere only with the sessions whose receiver is neighbor of u or whose sender has at least one $j \in \mathcal{G}_i$ as its neighbor. Moreover, each node can be involved in only one transmission in a slot. Thus, the directional multicast degree (denoted by $\tilde{M}_i$) satisfies the following relation.

$$\begin{aligned} k_i(\mathcal{N}) &\leq \tilde{M}_i \\ \Rightarrow \quad \max_i\{k_i(\mathcal{N})\} &\leq \max_i\{\tilde{M}_i\} \\ \Rightarrow \quad K(\mathcal{N}) &\leq \Gamma(\mathcal{N}). \end{aligned}$$

3) Let interference area for session i ($\mathcal{A}_i$) denote the area such that if an end-point of session j lies in $\mathcal{A}_i$ then session i and j interfere with each other. Note that $\mathcal{A}_i$ for any i is a subset of the area covered by a disk of radius $2d$ centered at the sender of i under bidirectional equal power model with the transmission radius of a node being d. Now, for two sessions j and k to belong to a interference set of i at least one of their end-points should belong to $\mathcal{A}_i$. Moreover, for these sessions to be mutually non-interfering, the distance between their end-points should be greater than d. Thus, if we place a disk of radius $d/2$ around one end-point of the interfering sessions j and k, then these disks do not intersect. Thus, the maximum number interfering sessions that are pair-wise non-interfering for any session i is less than or equal to the maximum number of non-overlapping disks of radius $d/2$ whose center lies in $\mathcal{A}_i$. Thus, $K(\mathcal{N})$ is less than or equal to the maximum number of non-overlapping disks of radius $d/2$ such that the disks lie completely in the area covered by a disk of radius $5d/2$. Thus, $K(\mathcal{N}) \leq 25$.
4) Note that under node exclusive spectrum sharing model, a session j interferes with session i only if i and j has common end-point. Since the number of end-points for the session i is $G_i + 1$. Thus,

$$\begin{aligned} k_i(\mathcal{N}) &\leq G_i + 1 \\ \Rightarrow \quad \max_i\{k_i(\mathcal{N})\} &\leq \max_i\{G_i + 1\} \\ \Rightarrow \quad K(\mathcal{N}) &\leq G(\mathcal{N}) + 1. \end{aligned}$$

■

Appendix J: Proof of Lemma 8

Proof: Note that unicast is a special case multicast and hence 1) and 2) follows immediately from Figures 5(a) and 5(b) respectively. Also, 4) and 5) follows from Figures 2(b) and 3 respectively. Now, we prove 3) by constructing a network $\mathcal{N}$ with a multicast session such that $K(\mathcal{N}) = 19$. We show such an network in Figure 12. Hence, the result follows. ■

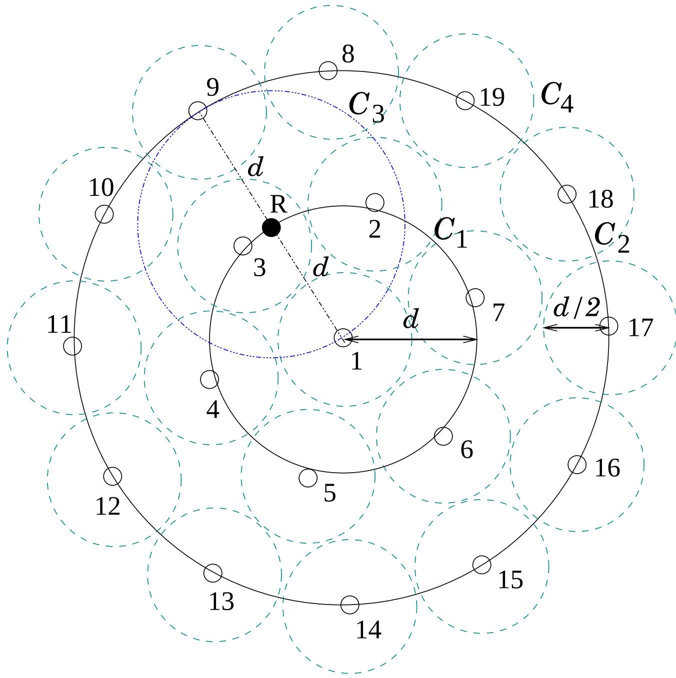

Fig. 12. Figure shows a construction to prove that a network $\mathcal{N}$ with a multicast session can be constructed such that $K(\mathcal{N}) = 19$. Now, we describe the construction. Figure shows 19 disks of radius $d/2$ whose center lies in a disk C_2 of radius $2d$. C_1 is another disk of radius d such that C_1 and C_2 are concentric. Now, consider a multicast session with 12 receivers such that the sender of the multicast session is at the center of C_1. Each of the receivers lie at the intersection of the line segment from the center of C_1 to nodes 8 to 19 and the boundary of C_1. One such receiver R that lies at the intersection of the line from the center of C_1 to node 9 and the boundary of C_1 is shown with the filled circle. This completes the construction of the multicast session. Now, we construct 19 unicast sessions that interfere with the multicast session but are pair-wise non-interfering. We place the senders of these unicast sessions at the center of each of the disks with radius $d/2$, i.e., at the location shown by the small circles numbered from 1 to 19. Note that since the disks with radius $d/2$ do not intersect, the distance between any $i \in \{1, \dots, 19\}$ to any $j \in \{1, \dots, 19\}$ is greater than $d + \epsilon$ for some $\epsilon > 0$. Now, we place receivers for each unicast session at the distance $\epsilon/4$ from its respective sender. Thus, note that the distance between any end-point of session i and any end-point of session j is strictly greater than d for every $i, j \in \{1, \dots, 19\}$. In other words, the 19 unicast sessions are pair-wise non-interfering. But, clearly, each of the unicast session interferes with the multicast session. Thus, $K(\mathcal{N}) = 19$.

Proofs of analytical results in Section VI-C (Lemmas 9 and 10)

Appendix K: Proof of Lemma 9

Lemma 9 follows from Lemma 19 and part (b) of Lemma 20. ■

Appendix L: Proof of Lemma 10

Consider a network $\mathcal{N}$ with three single-hop sessions i_1, i_2 and i_3 such that $S_{i_1} = \{i_2, i_3\}$ and $S_{i_2} = S_{i_3} = \{i_1\}$. Thus, $K_{i_1}(\mathcal{N}) = 2$ and $K_{i_2}(\mathcal{N}) = K_{i_3}(\mathcal{N}) = 1$. Let $\lambda_{i_1} = \lambda_{i_2} = \lambda_{i_3} = 1/2$. Note that a policy that schedules session i_1 in odd slots and i_2 and i_3 in the even slots stabilizes the system. Hence, $\vec{\lambda} \in \Lambda$.

Now, consider the arrival rate vector $(\lambda_{i_1}/K_{i_1}(\mathcal{N}), \lambda_{i_2}/K_{i_2}(\mathcal{N}), \lambda_{i_3}/K_{i_3}(\mathcal{N})) = (1/4, 1/2, 1/2)$, which corresponds to the following arrival process: i_2 (i_3, resp.) generates a packet every even (odd, resp.) slot, and i_1 generates a packet in slots $1, 5, 9, \dots$. Note that a maximal scheduling policy that schedules i_1 only when i_2 and i_3 do not have a packet to transmit, never schedules i_1 and is therefore unstable. Thus, $(\lambda_{i_1}/K_{i_1}(\mathcal{N}), \lambda_{i_2}/K_{i_2}(\mathcal{N}), \lambda_{i_3}/K_{i_3}(\mathcal{N})) \notin \Lambda^{\text{MS}}$. ■

Proof of analytical results in Section VI-D (Lemma 11)

Appendix M: Proof of Lemma 11

Proof: Let $\vec{\lambda} \in \Lambda_{\text{Q}}$. Then, under $\vec{\lambda}$, for some scheduling policy π, there exists a non-negative real vector $(q_1, \dots, q_N)$ such that for all i, $\lim_{n\to\infty} \sum_n Q_i(n)/n = q_i$ w.p. 1. Now, since $Q_i(n) = Q_i(0) + A_i(n-$

$1) - D_i(n-1)$, $\sum_n Q_i(n)/n = Q_i(0) + \sum_n \frac{A_i(n-1)-D_i(n-1)}{n}$. Thus, for all i, $\lim_{n\to\infty} \frac{A_i(n-1)-D_i(n-1)}{n} = 0$ w.p. 1. Since for all i, $\lim_{n\to\infty} A_i(n-1)/n = \lim_{n\to\infty} A_i(n)/n = \lambda_i$ w.p. 1, for all i, $\lim_{n\to\infty} D_i(n)/n = \lim_{n\to\infty} D_i(n-1)/n = \lambda_i$ w.p. 1. Thus, $\vec{\lambda} \in \Lambda$. Thus, from part (b) of Lemma 20, for all i, $\sum_{j\in S_i\cup\{i\}} \lambda_j/\beta_j(\mathcal{N}) \leq 1$. Thus,

$$\sum_{j\in S_i\cup\{i\}} \lambda'_j < 1 \ \forall \ i. \tag{34}$$

Let the arrival rate vector be $(\lambda'_1, \ldots, \lambda'_N)$. Consider a maximal scheduling policy. Let the state of the arrival process in the end of slot n be $\vec{B}(n)$. Clearly, $(\vec{Q}(n), \vec{B}(n))$ constitutes an irreducible aperiodic markov chain.

Consider the lyapunov function $f(t)$, where

$$f(t) = \sum_i \sum_{j\in S_i\cup\{i\}} Q_i(t)Q_j(t).$$

Clearly, $f(t) > 0$ if $Q_i(t) > 0$ for some i.

$$\begin{aligned}
& \mathbb{E}[f(n+1) - f(n)|\vec{Q}(n), \vec{B}(n)] \\
= & \sum_i \sum_{j\in S_i\cup\{i\}} \mathbb{E}[Q_i(n+1)Q_j(n+1) - Q_i(n)Q_j(n)|\vec{Q}(n), \vec{B}(n)] \\
= & \sum_i \sum_{j\in S_i\cup\{i\}} \mathbb{E}\left[\left(Q_i(n) + \alpha_i(n) - \tilde{D}_i(n)\right)\left(Q_j(n) + \alpha_j(n) - \tilde{D}_j(n)\right) - Q_i(n)Q_j(n)|\vec{Q}(n), \vec{B}(n)\right] \\
= & \sum_i \sum_{j\in S_i\cup\{i\}} \mathbb{E}\left[\left(Q_i(n) + \alpha_i(n) - \tilde{D}_i(n)\right)\left(Q_j(n) + \alpha_j(n) - \tilde{D}_j(n)\right) - Q_i(n)Q_j(n)|\vec{Q}(n), \vec{B}(n)\right] \\
\leq & \sum_i \sum_{j\in S_i\cup\{i\}} \mathbb{E}[Q_i(n)\alpha_j(n) - Q_i(n)\tilde{D}_j(n) + Q_j(n)\alpha_i(n) - Q_j(n)\tilde{D}_i(n)|\vec{Q}(n), \vec{B}(n)] \\
& +(N+1)N(\alpha^2_{\max} + 1).
\end{aligned}$$

Now,

$$\begin{aligned}
\sum_i \sum_{j\in S_i\cup\{i\}} Q_i(n)\alpha_j(n) &= \sum_i \sum_{j\in S_i\cup\{i\}} Q_j(n)\alpha_i(n), \\
\text{and } \sum_i \sum_{j\in S_i\cup\{i\}} Q_i(n)\tilde{D}_j(n) &= \sum_i \sum_{j\in S_i\cup\{i\}} Q_j(n)\tilde{D}_i(n).
\end{aligned}$$

Thus,

$$\begin{aligned}
& \mathbb{E}[f(n+1) - f(n)|\vec{Q}(n), \vec{B}(n)] \\
\leq & \ 2\sum_i Q_i(n) \sum_{j\in S_i\cup\{i\}} \mathbb{E}[\alpha_j(n) - \tilde{D}_j(n)|\vec{Q}(n), \vec{B}(n)] + (N+1)N(\alpha^2_{\max} + 1) \\
& \mathbb{E}[f(n+\tau) - f(n)|\vec{Q}(n), \vec{B}(n)] \\
\leq & \ 2\sum_i Q_i(n) \left[\sum_{j\in S_i\cup\{i\}} \sum_{k=0}^{\tau-1} \alpha_j(n+k) - \mathbb{E}[\sum_{j\in S_i\cup\{i\}} \sum_{k=0}^{\tau-1} \tilde{D}_j(n+k)|\vec{Q}(n), \vec{B}(n)\right] + (N+1)N(\alpha^2_{\max} + 1)\tau.
\end{aligned}$$

Under maximal scheduling, if $Q_i(n) > \tau + 1$, $\sum_{j\in S_i\cup\{i\}} \tilde{D}_j(l) = 1$ for each $l \in [n, n+\tau-1]$. Thus, if $Q_i(n) > \tau + 1$, $\sum_{j\in S_i\cup\{i\}} \sum_{k=0}^{\tau-1} \tilde{D}_j(n+k) = \tau$. Next, let $\delta = 1 - \max_i \sum_{j\in S_i\cup\{i\}} \lambda'_j$. From (34), $\delta > 0$.

Now, clearly, the arrival process is a positive recurrent markov chain. Hence, for any $\vec{Q}(n), \vec{B}(n)$ there exists τ_0 such that for all $\tau \geq \tau_0$, $\sum_{k=0}^{\tau-1} \alpha_j(n+k) \leq \tau(\lambda'_j + \delta/2N)$. Thus, for all $\vec{Q}(n)$, and for $\tau \geq \tau_0$,

$$\mathbb{E}[f(n+\tau) - f(n)|\vec{Q}(n) = \vec{Q}, \vec{B}(n) = \vec{B}] \leq -\delta\tau \sum_{i:Q_i(n)>\tau+1} Q_i(n) + (N+1)N(\alpha_{\max}^2 + \alpha_{\max} + 1)(\tau+1).$$

Thus, for $\tau \geq \tau_0$, $\mathbb{E}[f(n+\tau) - f(n)|\vec{Q}(n) = \vec{Q}, \vec{B}(n) = \vec{B}] < \infty$ for all $\vec{Q}, \vec{B}$, and $\mathbb{E}[f(n+\tau) - f(n)|\vec{Q}(n) = \vec{Q}, \vec{B}(n) = \vec{B}] < -1$ for all $\vec{Q}, \vec{B}$ such that $\max_i Q_i > \max(\tau+1, \frac{(N+1)N(\alpha_{\max}^2+\alpha_{\max}+1)(\tau+1)}{\delta\tau})$.

Hence, by Foster's theorem (Theorem 2.2.3 in [5]), for each $\tau \geq \tau_0$, $t \in (0, \tau-1)$, $\left(\vec{Q}(t), \vec{B}(t)\right)$, $\left(\vec{Q}(t+\tau), \vec{B}(t+\tau)\right)$, $\left(\vec{Q}(t+2\tau), \vec{B}(t+2\tau)\right)$, $\ldots$, is a positive recurrent markov chain. Also, all these markov chains have the same set of states, and same transition probabilities. Thus, under maximal scheduling, there exists a non-negative real vector $(q_1, \ldots, q_N)$ such that for all i, $\lim_{n\to\infty} \sum_n Q_i(n)/n = q_i$ w.p. 1. Thus, $(\lambda'_1, \ldots, \lambda'_N) \in \Lambda_{\mathrm{Q}}^{\mathrm{MS}}$. ■

Proofs of analytical results in Section VI-E (Lemmas 12, 13 and 14)

Appendix N: Proof of Lemma 12

Note that a network where each session traverses one link is a special case of a network where each session spans arbitrary link. In Section G, we have shown that under the unidirectional equal power model given any constant Z there exists a network $\mathcal{N}$ such that $K(\mathcal{N}) > Z$. Lemma 12 now follows from theorem 2. ■

Appendix O: Proof of Lemma 13

We prove Lemma 13 using the following supporting lemma.

Lemma 26: Let $\vec{\lambda} \in \{\vec{\lambda}: \text{ if } \lambda_{q(k)} > 0, \ \sum_{k\in S_j\cup\{j\}} \lambda_{q(k)} \leq 1, \ j = 1, \ldots, M\}$. Then $\vec{\lambda} \in \Lambda^{\mathrm{MS}}$.

Lemma 13 follows from part (a) of Lemma 20 and Lemma 26. ■

Appendix O.1: Proof of Lemma 26: We outline this proof as it is similar to that for Lemma 19. With regulators, the source of each session-link has two queues: waiting-queue and release-queue. Now, $A_j(n)$ and $D_j(n)$ denote the arrivals in and departures from the release-queue of session-link j in $(0, n]$, and $Q_j(n)$ denotes the queue length at the release-queue of session-link j at the beginning of the nth slot. For each j, $j = 1, \ldots, M$, the fluid limits of $A_j(.), D_j(.), Q_j(.)$ are defined as in Section A.2.a.

Now, we state and prove some important properties of the fluid limits which we use to prove Lemma 26.

Lemma 27: Every fluid limit satisfies, $\overline{A}_j(t) \leq \lambda_j t$ w.p. 1 for every session-link $j = 1, \ldots, M$ and $t \geq 0$.

Proof: The proof is similar to that for Lemma 21 when j is the first session-link of its session. When j is not the first session-link of its session, the proof follows because due to the regulator the release-queue of j receives packet w.p. at most $\lambda_{q(j)}$ in any slot n. ■

Lemma 28: Any fluid limit $(\overline{A}_i, \overline{D}_i, \overline{Q}_i)$ for $\mathcal{N}$ satisfies the following equality for every i and $t \geq 0$ with probability (w.p.) 1.

$$\overline{Q}_i(t) = \overline{Q}_i(0) + \overline{A}_i t - \overline{D}_i(t). \tag{35}$$

The proof is similar to that for Lemma 22.

Lemma 29: Let $\overline{Q}_i(0) = 0$ for every i. Also, let $\sum_{k\in S_j\cup\{j\}} \lambda_{q(k)} \leq 1$ if $\lambda_{q(j)} > 0$, $j = 1, \ldots, M$. Then, under maximal scheduling, every fluid limit satisfies that $\overline{Q}_i(t) = 0$ for every $t \geq 0$ w.p. 1 for every i.

The lemma follows from Lemma 27. The arguments are similar to that in the proof of Lemma 23.

We now prove Lemma 26.

Proof: We prove the following for each session-link $j = 1, \ldots, M$.

1) Every fluid limit satisfies, $\overline{A}_j(t) = \lambda_j t$ w.p. 1 for every session-link $j = 1, \ldots, M$ and $t \geq 0$.
2) $\bar{D}_j(t) = \lambda_{q(j)} t$ w.p. 1 for every t.
3) $\lim_{t\to\infty} D_j(t)/t = \lambda_{q(j)} t$ w.p. 1.

We prove using induction on the position of the session-links in the paths of their sessions.

First, let j be the first session-link of some session (i.e., the session-link originating at the source of the session). The arrivals in the release-queue of the first session-link are the exogenous arrivals. Now, (1) follows from (4). From Lemmas 28 and 23, $\overline{D}_j(t) = \overline{A}_j t$ w.p. 1 for every $t \geq 0$. Now, (2) follows from (1). Finally, using arguments similar to those in the proof Lemma 19, $\lim_{r\to\infty} D_j^r(t) = \lambda_{q(j)} t$ w.p. 1 for every t follows from (1) and (2). Now, (3) follows by choosing $t = 1$.

Now, let (1) and (2) hold for all session-links that are $1, \ldots, p$ in the paths of their sessions. We now prove (1) and (2) for a session-link j that is the $p+1$th in the path of its session. Let session-link k be the session-link of session $q(j)$ that terminate at the source of session-link j. Let $\hat{Q}_j(n)$ be the queue length at the waiting-queue of session-link j at the beginning of the nth slot. Now,

$$\hat{Q}_j(n+1) = \hat{Q}_j(0) + D_k(n) - A_j(n).$$

From (3) of induction hypothesis, $\lim_{t\to\infty} D_k(t)/t = \lambda_{q(j)} t$ w.p. 1. Note that $A_j(n) = 1$ w.p. $\lambda_{q(j)}$ if $Q_j(n) > 0$. Thus, the waiting-queue of session-link j is a queue which receives packets as per an arrival process that satisfies SLLN with rate $\lambda_{q(j)}$ and is served w.p. $\lambda_{q(j)}$ whenever it is non-empty. It follows that the departure-process of this queue $A_j(\cdots)$ satisfies SLLN with rate $\lambda_{q(j)}$. Now, (1) follows. Now, (2) and (3) follows as in the base case.

The lemma follows from (3). ■

Appendix P: Proof of Lemma 14

We prove Lemma 14 using Lemma 20 and another supporting lemma, Lemma 30, which we state and prove next.

Lemma 30: Consider an arrival rate vector $\vec{\lambda}'$ such that $\sum_{j\in S_k\cup\{k\}} \lambda'_{q(j)} < 1$. Then the packet queue of every session-link will almost surely become empty infinitely often. Furthermore, for every session-link j and time t, $\mathbb{E}[B_{j,t}] < \infty$.

Proof: Now, $\alpha_j(t)$ and $\tilde{D}_j(t)$ denote the number of arrivals and departures respectively for session-link j in slot t. Let $Q_j(t)$ be the number of packets for the session of session-link j waiting for transmission at the source of session-link j at the end of slot t. Let $S_j \cup \{j\} = \mathcal{X}_j$, and $\widehat{n} = |\mathcal{X}_j|$. First, we obtain relations among these parameters. If session-link j satisfy $Q_j(\nu) > 0$ for every $\nu \in [t, t+\tau]$, then for every $\nu \in [t, t+\tau]$,

$$\sum_{k\in\mathcal{X}_j} \tilde{D}_k(\nu) \geq 1. \tag{36}$$

$$\begin{aligned} Q_j(t) + \sum_{\nu=t+1}^{t+\tau} \alpha_j(\nu) &\leq \sum_{\nu=1}^{t+\tau} A_{q(j)}(\nu) \\ &\leq t\alpha_{\max} + \sum_{\nu=t+1}^{t+\tau} A_{q(j)}(\nu). \end{aligned} \tag{37}$$

Now we have,

$$\begin{aligned} &\mathbb{P}\left\{B_{j,t} > \tau\right\} \\ &\leq \mathbb{P}\left\{\bigcap_{v=t}^{t+\tau}\left\{\left[\sum_{k\in\mathcal{X}_j} Q_k(t) + \sum_{\nu=t+1}^{v}\sum_{k\in\mathcal{X}_j}\alpha_k(\nu) - \sum_{\nu=t+1}^{v}\sum_{k\in\mathcal{X}_j}\tilde{D}_k(\nu) > 0\right]\right\}\right\} \\ &\leq \mathbb{P}\left\{\bigcap_{v=t+1}^{t+\tau}\left\{\sum_{k\in\mathcal{X}_j} Q_k(t) + \sum_{\nu=t+1}^{v}\left(\sum_{k\in\mathcal{X}_j}\alpha_k(\nu) - 1\right) > 0\right\}\right\} \text{ (from (36))} \end{aligned}$$

$$
\begin{aligned}
&\le \ \mathbb{P}\left\{\sum_{k\in\mathcal{X}_j} Q_k(t) + \sum_{\nu=t+1}^{t+\tau}\sum_{k\in\mathcal{X}_j}\alpha_k(\nu) - \tau > 0\right\}\\
&\le \ \mathbb{P}\left\{\frac{t\widehat{n}\alpha_{\max}}{\tau} + \frac{1}{\tau}\sum_{\nu=t+1}^{t+\tau}\sum_{k\in\mathcal{X}_j} A_{F_{q(k)}}(\nu) - 1 > 0\right\} \text{ (from (37))}\\
&= \ \mathbb{P}\left\{\frac{t\widehat{n}\alpha_{\max}}{\tau} + \sum_{k\in\mathcal{X}_j}\left(\frac{1}{\tau}\sum_{\nu=t+1}^{t+\tau} A_{F_{q(k)}}(\nu) - \lambda'_{q(k)}\right) > 1 - \sum_{k\in\mathcal{X}_j}\lambda'_{q(k)}\right\}.
\end{aligned}
$$

Let $\delta = 1 - \sum_{k\in\mathcal{X}_j}\lambda'_{q(k)}$. Clearly, $\delta > 0$. Thus,

$$
\begin{aligned}
&\mathbb{P}\left\{B_{j,t} > \tau\right\}\\
&\le \ \mathbb{P}\left\{\left\{\frac{t\widehat{n}\alpha_{\max}}{\tau} > \frac{\delta}{\widehat{n}+1}\right\}\bigcup_{k\in\mathcal{X}_j}\left\{\frac{1}{\tau}\sum_{\nu=t+1}^{t+\tau} A_{F_{q(k)}}(\nu) - \lambda'_{q(k)} > \frac{\delta}{\widehat{n}+1}\right\}\right\}\\
&\le \ \mathbb{P}\left\{\frac{t\widehat{n}\alpha_{\max}}{\tau} > \frac{\delta}{\widehat{n}+1}\right\} + \sum_{k\in\mathcal{X}_j}\mathbb{P}\left\{\frac{1}{\tau}\sum_{\nu=t+1}^{t+\tau} A_{F_{q(k)}}(\nu) - \lambda'_{q(k)} > \frac{\delta}{\widehat{n}+1}\right\}\\
&= \ \sum_{k\in\mathcal{X}_j}\mathbb{P}\left\{\frac{1}{\tau}\sum_{\nu=t+1}^{t+\tau} A_{F_{q(k)}}(\nu) - \lambda'_{q(k)} > \frac{\delta}{\widehat{n}+1}\right\} \quad \text{if } \tau > \frac{\widehat{n}(\widehat{n}+1)t\alpha_{\max}}{\delta}.
\end{aligned}
$$

Now, from (6), the packet queue of every session-link will almost surely become empty infinitely often. Also,

$$
\mathbb{E}[B_{j,t}] = \sum_{\tau=1}^{\infty}\mathbb{P}\left\{B_{j,t} > \tau\right\} < \infty.
$$

■

Lemma 14 follows from part (a) of Lemma 20 and Lemma 30.

Proofs of analytical results in Section VII (Lemmas 15 and 17 and Theorem 3)

Appendix Q: Proof of Lemma 15

Let $X = \{\vec{\lambda} = (\lambda_1, \ldots, \lambda_N) : \text{ if } \lambda_i > 0, \ \sum_{j\in S_i\cup\{i\}}\lambda_j \le 1, \quad \forall\ i = 1, \ldots, N.\}$ From Lemma 19, if $\vec{\lambda} \in X$, $\vec{\lambda} \in \Lambda^{\text{MS}}$.

Now, let $\vec{\lambda} \notin X$. Then there exists a session i such that $\lambda_i > 0$ and $\sum_{j\in S_i\cup\{i\}}\lambda_j > 1$. Let $\lambda_j > 0$ for m sessions in S_i, where $m \le |S_i|$. Let these sessions be $j_1, \ldots, j_m$. For simplicity, we assume that λ_k, $k \in \{j_1, \ldots, j_m\}$ are rational numbers. Let Z be an integer such that $Z\lambda_k$ is an integer for all $k \in \{j_1, \ldots, j_m\}$. Consider an arrival process in which the arrivals for $j_1, \ldots, j_m$ are periodic with period Z, and j_l generates packets in $\left(Z\sum_{p=1}^{l-1}\lambda_p \mod Z\right)$th to $\left(Z(\sum_{p=1}^{l}\lambda_p) - 1 \mod Z\right)$th slots of the period, $l = 1, \ldots, m$. Note that for this arrival process at least one session in $\{j_1, \ldots, j_m\}$ generates a packet in every slot. Consider a maximal scheduling policy that resolves contention among sessions that have packets to transmit as follows. If j_1 has a packet to transmit, it transmits. For $2 \le k \le m$, if j_k has a packet to transmit, and none of the sessions in $\{j_1, \ldots, j_{k-1}\}$ that interfere with j_k are transmitting, j_k transmits. Note that this policy schedules one session in $\{j_1, \ldots, j_m\}$ every slot, and thus never schedules i. Thus, $d_i = 0 < \lambda_i$. Thus, $\vec{\lambda} \notin \Lambda^{\text{MS}}$. Thus, $\Lambda^{\text{MS}} = X$. The result follows. ■

Appendix R: Proof of Lemma 17

We prove Lemma 17 when each session spans one link. First, we show that if a session generates packets at rate r or higher, and if it is sampled at rate r or higher at every bucket associated with it, then it receives tokens at rate r or higher from each of its buckets (Lemma 31). We next show that a session's sampling rate at any of its buckets equals its maxmin fair rate (Lemma 32). Now, the result follows, as by definition, a session's maxmin fair rate is less than or equal to its packet generation rate. We prove Lemmas 31 and 32 in sections R.1 and R.2. Thus, like in the current section, throughout sections R.1 and R.2, we will assume that every session spans one link.

We introduce some terminologies and subsequently state Lemmas 31 and 32. Let $S_{i,n}(t)$ be the number of times session i is sampled at token-bucket n in the interval $(0,t]$, $L = \max_i b_i$, $\sigma = \max_i \sigma_i$, and β, γ are constants that are specified later.

Lemma 31: Consider an arbitrary K and a sequence of K disjoint intervals, $(t_l, w_l]$, $l = 1, \ldots, K$, that satisfies the following property for session i, for every positive integer M' and every sequence of sub-intervals $(x_m, y_m]$, $m = 1, \ldots, M'$, $(x_m, y_m] \subset (t_l, w_l]$, for some l: At every bucket n associated with i,

$$\sum_{m=1}^{M'} (S_{i,n}(y_m) - S_{i,n}(x_m)) \geq r \sum_{m=1}^{M'} (y_m - x_m) - e - M'f, \tag{38}$$

where e and f are constants that do not depend on M' and the sub intervals $(x_m, y_m]$, $m = 1, \ldots, M'$. Let $\lambda_i \geq r$ and $W \geq 3^{b_i - 1}(f + \sigma_i)/2$. Then, at every bucket n associated with i,

$$\sum_{l=1}^{K} (C_{i,n}(w_l) - C_{i,n}(t_l)) \;\; \geq \;\; r \sum_{l=1}^{K} (w_l - t_l) - 2^{b_i - 1} e - K 3^{b_i - 1}(f + \sigma_i). \tag{39}$$

Lemma 32: Consider any positive integer K, and an arbitrary non-decreasing sequence of times $x_1, y_1, \ldots, x_K, y_K$. Let $W \geq 3^{L-1}(\varepsilon_1(F) + \sigma)/2$, where $\varepsilon_1(F)$ is defined in (43) to (48). For every bucket n associated with session i,

$$\begin{aligned} \sum_{k=1}^{K} (S_{i,n}(y_k) - S_{i,n}(x_k)) \;\; &\geq \;\; d_i^* \sum_{k=1}^{K} (y_k - x_k) - \beta \\ &\quad -K\gamma, \end{aligned} \tag{40}$$

$$\begin{aligned} \sum_{k=1}^{K} (C_{i,n}(y_k) - C_{i,n}(x_k)) \;\; &\geq \;\; d_i^* \sum_{k=1}^{K} (y_k - x_k) - \beta \\ &\quad -K\gamma, \end{aligned} \tag{41}$$

$$\begin{aligned} \sum_{k=1}^{K} (C_{i,n}(y_k) - C_{i,n}(x_k)) \;\; &\leq \;\; d_i^* \sum_{k=1}^{K} (y_k - x_k) + \beta \\ &\quad +K\gamma. \end{aligned} \tag{42}$$

Here, β and γ are constants that do not depend on $x_1, y_1, \ldots, x_K, y_K$.

We introduce the notion of "rank" of a session for defining β and γ. A session has rank p if its maxmin fair rate is $\hat{d}_p$, the pth lowest among the maxmin fair rates of different sessions. Let F be the number of distinct ranks, $F \leq N$.

$$\varsigma_1(1) \;\; = \;\; 0. \tag{43}$$

$$\varepsilon_1(1) \;\; = \;\; 1. \tag{44}$$

$$\varsigma_2(p) \;\; = \;\; 2^{L-1}\varsigma_1(p). \tag{45}$$

$$\varepsilon_2(p) \;\; = \;\; 3^{L-1}(\varepsilon_1(p) + \sigma). \tag{46}$$

$$\varsigma_3(p) \;\; = \;\; 2\sigma + \max(L, 2)\,(\varsigma_2(p) + \varepsilon_2(p))$$

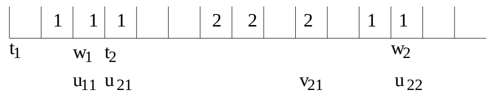

Fig. 13. We show two intervals $(t_1, w_1]$ and $(t_2, w_2]$, and some type 1 and 2 slots. We also show the corresponding u and v slots. Here $(t_1, u_{11}]$, $(t_2, u_{21}]$, $(v_{21}, u_{22}]$ are example sub-intervals that end in $u-$slots and start from the nearest $v-$slot or t_i-slot.

$$+2LW. \tag{47}$$

$$\varepsilon_3(p) = \varepsilon_2(p). \tag{48}$$

$$\varsigma_1(p+1) = (L-1)\varsigma_3(p). \tag{49}$$

$$\varepsilon_1(p+1) = (L-1)\varepsilon_3(p) + 1. \tag{50}$$

Now, $\beta = \varsigma_3(F)$ and $\gamma = \varepsilon_3(F)$.

Now, for any given $\vec{\lambda}$, Lemma 17 follows from (41) and (42) of Lemma 32 with $\varrho = \beta + \gamma$ and $W_0 = 3^{L-1}(\varepsilon_1(F) + \sigma)/2$. ■

Appendix R.1: Proof of Lemma 31: We first present the intuition behind the proof. The proof is by induction on the number of buckets associated with a session. The sessions with one bucket form the base case. Note that any such session receives a token at its bucket every time it is sampled at its bucket and has a packet that has not been released, since no adjacent bucket applies back-pressure. Now, the lemma follows for the base case from the lower bounds on the sampling and packet generation rates. We next assume that the lemma holds for all sessions with p buckets, and then prove the lemma for sessions with $p+1$ buckets. Consider a session with $p+1$ buckets and adjacent buckets n and $n+1$ associated with it. Bucket $n+1$ does not prevent the generation of any token at n unless the number of tokens at n is W more than that at $n+1$. If the number of tokens at n is W more than that at $n+1$, n does not prevent any token generation at $n+1$, and the buckets $n+1, n+2, \ldots$ generate tokens oblivious to the presence of the buckets $1, \ldots, n$, as though they constitute a session with fewer buckets. By induction hypothesis, and from the sampling and packet generation rates, the session receives tokens at rate r or higher at $n+1$ in these intervals. In all these slots, the number of tokens at n exceeds that at $n+1$ by W. Thus, n's token generation rate is lower bounded by $n+1$'s token generation rate which is at least r. In other slots, $n+1$ does not prevent the generation of any token at n. Thus, the token generation at the buckets $1, \ldots, n$ resembles that for a session with fewer buckets. Thus, by induction hypothesis and the assumption on the sampling rate, in all slots, n generates tokens at rate r or higher for the session.

Proof: We prove by induction on the number of buckets p associated with a session.

First consider a session i with one bucket n. Let n not be at the source node of i. The lemma holds from the assumption on the sampling rate (condition (38)). Now, let n be at the source node of i. Let $A_i^{\text{NR}}(t)$ be the number of packets of session i at its source at time t that have not been released. We now define a slot z_l. If $A_i^{\text{NR}}(t) > 0$ for all $t \in (t_l, w_l]$, $z_l = t_l$, else $z_l = \max_{t \in (t_l, w_l], A_i^{\text{NR}}(t)=0} t$. If $z_l > t_l$,

$$\begin{aligned} C_{i,n}(z_l) - C_{i,n}(t_l) &= A_i(z_l) - A_i(t_l) + A_i^{\text{NR}}(t_l) \\ &\geq A_i(z_l) - A_i(t_l) \\ &\geq r(z_l - t_l) - \sigma_i. \end{aligned} \tag{51}$$

The last inequality follows from (10) and since $r \leq \lambda_i$. Clearly, (51) also holds if $z_l = t_l$. Bucket n generates a token for session i every time it samples i in $(z_l, w_l]$, $\forall\ l$.

$$\begin{aligned} & \sum_{l=1}^{K} (C_{i,n}(w_l) - C_{i,n}(z_l)) \\ = & \sum_{l=1}^{K} (S_{i,n}(w_l) - S_{i,n}(z_l)) \end{aligned}$$

$$\geq \quad r\sum_{l=1}^{K}(w_l - z_l) - e - Kf \text{ (from (38)).} \tag{52}$$

$$\begin{aligned}
&\sum_{l=1}^{K}(C_{i,n}(w_l) - C_{i,n}(t_l)) \\
= \quad & \sum_{l=1}^{K}(C_{i,n}(w_l) - C_{i,n}(z_l)) + \sum_{l=1}^{K}(C_{i,n}(z_l) - C_{i,n}(t_l)) \\
\geq \quad & r\sum_{l=1}^{K}(w_l - t_l) - e - K(f + \sigma_i) \text{ (from (51) and (52)).}
\end{aligned}$$

Thus, (39) holds in the base case.

We now assume that (39) holds for all sessions with p or fewer buckets, and prove (39) for an arbitrary session i with $p+1$ buckets. Consider an arbitrary bucket n associated with i. If the number of tokens of i at n does not exceed that at buckets adjacent to n by W or more in the intervals (t_l, w_l), $l = 1, \ldots, K$, then the token generation process for i at n is not affected by back-pressure, and the proof is similar to the base case. Thus, we assume that there exists a bucket B that is adjacent to n, and $C_{i,n}(t) = C_{i,B}(t) + W$ at some time t in these intervals. Clearly $B \in \{n-1, n+1\}$. We consider the case that $B = n+1$. The proof when $B = n-1$ is similar.

Let a slot t where $C_{i,n}(t)$ exceeds $C_{i,n+1}(t)$ by W be a type 1 slot, and a slot t where $C_{i,n+1}(t)$ exceeds $C_{i,n}(t)$ by W be a type 2 slot; a slot may neither be type 1 nor type 2. Consider each $(t_l, w_l]$ interval separately. Consider the sequences of type 1 and 2 slots that are obtained after removing the slots without numbers. The last slot in such a sequence of type-1 (2) slots is denoted a “u” (“v”) slot. The mth “u-slot” (“v-slot”) of the lth interval is u_{lm} (v_{lm}) (Figure 13). Note that

$$C_{i,n}(u_{lm}) \;=\; C_{i,n+1}(u_{lm}) + W \;\forall\; l, m. \tag{53}$$

$$C_{i,n+1}(v_{lm}) \;=\; C_{i,n}(v_{lm}) + W \;\forall\; l, m. \tag{54}$$

$$C_{i,n}(t) \;\leq\; C_{i,n+1}(t) + W, \;\forall\; t. \tag{55}$$

Consider a sub-interval that ends at a u slot and starts from a t_j (not inclusive) or a $v-$slot (not inclusive), whichever is the nearest to the $u-$slot (Figure 13). Let there be J_l such sub-intervals in $(t_l, w_l]$, and $\sum_{l=1}^{K} J_l = I_1$. These sub-intervals do not consist of any type 2 slot. Thus, n does not prevent any session i token generation at $n+1$ in these sub-intervals. Hence, in these sub-intervals, the token generation for i in buckets $n+1, \ldots, p+1$ resembles that in the buckets of a session with $p+1-n$ buckets, where $n > 0$. Condition (38) holds for i in each of these buckets for every set of sub-intervals of these I_1 sub-intervals, since any such sub-interval is in $(t_l, w_l]$ for some l. Thus, the number of tokens generated for i in these I_1 sub-intervals in each of these buckets can be lower bounded using the induction hypothesis. The sub-intervals in $(t_l, w_l]$ are $(t_l, u_{l1}]$ and $(v_{lm-1}, u_{lm}]$, $m > 1$, if $v_{l1} > u_{l1}$ as in Figure 13; the sub-intervals are $(v_{lm}, u_{lm}]$, $m \geq 1$, otherwise. We assume that $v_{l1} > u_{l1}$ for all l; the argument is similar if $v_{l1} < u_{l1}$ for some or all l. From induction hypothesis,

$$\begin{aligned}
&\sum_{l=1}^{K}\left((C_{i,n+1}(u_{l1}) - C_{i,n+1}(t_l)) + \sum_{m=2}^{J_l}(C_{i,n+1}(u_{lm}) - C_{i,n+1}(v_{lm-1}))\right) \\
\geq \quad & r\sum_{l=1}^{K}\left((u_{l1} - t_l) + \sum_{m=2}^{J_l}(u_{lm} - v_{lm-1})\right) - 2^{p-1}e - I_1 3^{p-1}(f + \sigma_i).
\end{aligned} \tag{56}$$

$$\begin{aligned}
&C_{i,n}(u_{l1}) - C_{i,n}(t_l) \\
\geq \quad & C_{i,n+1}(u_{l1}) + W - C_{i,n+1}(t_l) - W \text{ (from (53) and (55))} \\
= \quad & C_{i,n+1}(u_{l1}) - C_{i,n+1}(t_l).
\end{aligned} \tag{57}$$

From (53) and (54),

$$C_{i,n}(u_{lm}) - C_{i,n}(v_{lm-1}) \;=\; C_{i,n+1}(u_{lm}) - C_{i,n+1}(v_{lm-1}) + 2W. \tag{58}$$

$$\begin{aligned}
& \sum_{l=1}^{K}((C_{i,n}(u_{l1}) - C_{i,n}(t_l)) + \sum_{m=2}^{J_l}(C_{i,n}(u_{lm}) - C_{i,n}(v_{lm-1}))) \\
\geq\; & \sum_{l=1}^{K}((C_{i,n+1}(u_{l1}) - C_{i,n+1}(t_l))) + \sum_{m=2}^{J_l}(C_{i,n+1}(u_{lm}) - C_{i,n+1}(v_{lm-1}))) \\
& +2W(I_1 - K) \text{ (from (57) and (58))} \\
\geq\; & r\sum_{l=1}^{K}\left((u_{l1} - t_l) + \sum_{m=2}^{J_l}(u_{lm} - v_{lm-1})\right) - 2^{p-1}e - K3^{p-1}(f+\sigma_i) \\
& +(I_1 - K)(2W - 3^{p-1}f - 3^{p-1}\sigma_i) \text{ (from (56))}.
\end{aligned} \tag{59}$$

Now, consider the sub-intervals obtained after removing these I_1 sub-intervals from $\cup_{l=1}^{K}(t_l, w_l]$. These new sub-intervals do not contain any type 1 slot. Thus, $n+1$ does not prevent any session i token generation at n. Hence, the session i token generation in buckets $1, \ldots, n$ resembles that of a session with n buckets, where $n \leq p$. The number of session i tokens generated at n in these sub-intervals can be lower bounded from the induction hypothesis. There are at most $I_1 + K$ such sub-intervals, which are of the form $(u_{lm}, v_{lm}]$ and $(u_{J_l}, w_l]$, since we assume that $v_{l1} > u_{l1}\ \forall\ l$.

$$\begin{aligned}
\text{Thus, } & \sum_{l=1}^{K}\left((C_{i,n}(w_l) - C_{i,n}(u_{J_l})) + \sum_{m=1}^{J_l-1}(C_{i,n}(v_{lm}) - C_{i,n}(u_{lm}))\right) \\
\geq\; & r\sum_{l=1}^{K}\left((w_l - u_{J_l}) + \sum_{m=1}^{J_l-1}(v_{lm} - u_{lm})\right) - 2^{p-1}e - (I_1 - K)3^{p-1}(f+\sigma_i) - 2K3^{p-1}(f+\sigma_i).
\end{aligned} \tag{60}$$

Adding (59) and (60),

$$\begin{aligned}
& \sum_{l=1}^{K}(C_{i,n}(w_l) - C_{i,n}(t_l)) \\
\geq\; & r\sum_{l=1}^{K}(w_l - t_l) - 2^{p}e - K3^{p}(f+\sigma_i) + (I_1 - K)(2W - 3^{p}(f+\sigma_i)).
\end{aligned} \tag{61}$$

Note that $p+1 \leq b_i$ and thus, $W \geq 3^{p}(f+\sigma_i)/2$. We have implicitly assumed that at least one type-1 slot exists in each interval $(t_l, w_l]$; this justifies the summation from $l = 1$ to K in (56). Under this assumption, $I_1 \geq K$. Hence, (39) holds for session i at bucket n. If there is no type-1 slot in $(t_l, w_l]$ for some l, then the summation in (56) must be over the intervals $(t_l, w_l]$ that have at least one type-1 slot. Let K_1 be the number of such intervals. Now, $(I_1 - K)$ must be replaced with $(I_1 - K_1)$. Since $I_1 \geq K_1$, (39) holds at all buckets associated with i. ■

Appendix R.2: Proof of Lemma 32: We outline the proof for the special case that all sessions always have packets to transmit, i.e., $\lambda_i > 1$ for all i. We use induction on the rank p of a session. For the base case ($p = 1$), using a property of the round robin sampling, we show that all sessions are sampled at a rate $\hat{d}_1$ or higher at every bucket. Now, (41), the lower bound on the token generation rate follows from Lemma 31. Next, we show (42), i.e., the token generation rates are upper bounded by $\hat{d}_1$ for all sessions with rank 1. This follows because the sampling and hence the token generation rate is upper bounded by $\hat{d}_1$ at the bottleneck bucket, and due to back-pressure the token generation rates for a session are equal at different buckets in the session's path. Now, consider the induction case, i.e., arbitrary p. The token generation rates of sessions with rank lower than p are upper bounded by their respective maxmin fair rates which are upper bounded by $\hat{d}_p$. Sessions of rank p or higher are sampled in a certain minimum fraction of the slots in which the sessions with rank lower than p do not receive tokens. Therefore, the lower bound on the sampling rate of sessions with rank p or higher follows. Again, the lower bound on the token generation rate follows from Lemma 31. We prove, as in the base case, the upper bound on the token generation rate for sessions with rank p.

In the formal proof, we relax the assumption that all sessions always have packets to transmit, i.e., we consider arbitrary $\vec{\lambda}$. We would like to clarify the usage of a particular notation before proceeding further. We have so far numbered token-buckets based on the sessions traversing them. In this terminology, bucket n of session i is i's nth bucket, and $C_{i,n}(t), S_{i,n}(t)$ are the number of tokens generated for session i at and the number of times session i is sampled at its nth bucket respectively. In the following proof, we number token-buckets separately. Thus, for example, we consider token-bucket n and all sessions associated with n. Now, $n(i)$ will denote the number for the bucket n among i's buckets. Thus, we need to use $C_{i,n(i)}(t), S_{i,n(i)}(t)$ instead of $C_{i,n}(t), S_{i,n}(t)$. For simplicity, we still use $C_{i,n}(t), S_{i,n}(t)$. Thus, in the following proof, $C_{i,n}(t), S_{i,n}(t)$ really stand for $C_{i,n(i)}(t), S_{i,n(i)}(t)$ respectively. Note that this inconsistency is limited to the following proof only, and does not lead to any error, because none of the analytical guarantees in other lemmas (including those that are used in the following proof and those whose proof use Lemma 32) depend on the token-bucket number.

Proof: We prove the following for ranks $p = 1, \ldots, F$, by induction on p.

For each bucket n, for each session i that is associated with n and has rank greater than or equal to p, for any positive integer K, and for any nondecreasing sequence of times $x_1, y_1, \ldots, x_K, y_K$,

$$\sum_{k=1}^{K} (S_{i,n}(y_k) - S_{i,n}(x_k)) \geq \hat{d}_p \sum_{k=1}^{K} (y_k - x_k) - \varsigma_1(p) - K\varepsilon_1(p). \tag{62}$$

For each bucket n, for each session i that is associated with n and has rank greater than or equal to p, for any positive integer K, and for any nondecreasing sequence of times $x_1, y_1, \ldots, x_K, y_K$,

$$\sum_{k=1}^{K} (C_{i,n}(y_k) - C_{i,n}(x_k)) \geq \hat{d}_p \sum_{k=1}^{K} (y_k - x_k) - \varsigma_2(p) - K\varepsilon_2(p). \tag{63}$$

If a session i has rank p, and $d_i^* = \lambda_i$,

$$A_i^{\text{NR}}(t) \leq \sigma_i + \varsigma_2(p) + \varepsilon_2(p) \ \forall \ t. \tag{64}$$

For each bucket n, for each session i that is associated with n and has rank p, for any positive integer K, and for any nondecreasing sequence of times $x_1, y_1, \ldots, x_K, y_K$,

$$\sum_{k=1}^{K} (C_{i,n}(y_k) - C_{i,n}(x_k)) \leq \hat{d}_p \sum_{k=1}^{K} (y_k - x_k) + \varsigma_3(p) + K\varepsilon_3(p). \tag{65}$$

We first prove (62) to (65) for $p = 1$. Note that $\hat{d}_1 = \min(1/L, \min_i \lambda_i)$. Consider a bucket n. Let $\mathcal{X}$ be the set of sessions associated with n. Since at least one session is sampled at n in a slot, in any interval $(x_k, y_k]$,

$$\sum_{j \in \mathcal{X}} (S_{j,n}(y_k) - S_{j,n}(x_k)) \quad \geq \quad y_k - x_k.$$

Since sessions are sampled in round robin order, $S_{i,n}(y_k) - S_{i,n}(x_k) \geq S_{j,n}(y_k) - S_{j,n}(x_k) - 1$ for any two sessions i, j associated with n. Thus, for any session i associated with n,

$$\begin{aligned} |\mathcal{X}| \left(S_{i,n}(y_k) - S_{i,n}(x_k) + 1\right) \quad &\geq \quad y_k - x_k, \\ S_{i,n}(y_k) - S_{i,n}(x_k) \quad &\geq \quad \frac{y_k - x_k}{|\mathcal{X}|} - 1. \end{aligned}$$

Thus, every session associated with bucket n is sampled at least $\sum_{k=1}^{Q}(y_k - x_k)/|\mathcal{X}| - Q$ times for any arbitrary sequence of nondecreasing times $x_1, y_1, \ldots, x_Q, y_Q$, and any arbitrary Q. Since $|\mathcal{X}| \leq L$, $\hat{d}_1 \leq 1/|\mathcal{X}|$. Thus, (62) holds with $\varsigma_1(1) = 0, \varepsilon_1(1) = 1$.

Since $\varepsilon_F(1) \geq \varepsilon_1(1)$, $W \geq 3^{L-1}(\varepsilon_1(1) + \sigma)/2$. Hence, (63) follows from Lemma 31 with $\varsigma_2(1) = 2^{L-1}\varsigma_1(1)$ and $\varepsilon_2(1) = 3^{L-1}(\varepsilon_1(1) + \sigma)$.

Now, we prove (64) for $p = 1$. Consider a session i with rank 1 and $d_i^* = \lambda_i$. Thus, $\hat{d}_1 = \lambda_i$. Let n be the bucket at the source node of i.

$$\begin{aligned} A_i^{\text{NR}}(t) &= A_i(t) - C_{i,n}(t) \\ &\leq (\lambda_i - \hat{d}_1)t + \sigma_i + \varsigma_2(1) + \varepsilon_2(1) \ \ \text{(from (10) and (63) for } p = 1) \\ &= \sigma_i + \varsigma_2(1) + \varepsilon_2(1) \ \text{(since } \hat{d}_1 = \lambda_i). \end{aligned}$$

Thus, (64) follows for $p = 1$.

Now, we prove (65) for $p = 1$. Consider a session i with rank 1. Let n be a bucket associated with i. Consider a sequence of non-decreasing times $x_1, y_1, \ldots, x_K, y_K$.

$$\begin{aligned} &\sum_{k=1}^{K}(C_{i,n}(y_k) - C_{i,n}(x_k)) \\ &= C_{i,n}(y_K) - C_{i,n}(x_1) - \sum_{k=1}^{K-1}(C_{i,n}(x_{k+1}) - C_{i,n}(y_k)) \\ &\leq C_{i,n}(y_K) - C_{i,n}(x_1) - \hat{d}_1\sum_{k=1}^{K-1}(x_{k+1} - y_k) + \varsigma_2(1) + (K-1)\varepsilon_2(1) \ \text{(from (63) for } p = 1). \quad (66) \end{aligned}$$

Since $\hat{d}_1 = d_i^*$ and $d_i^* \leq \lambda_i$, $\hat{d}_1 \leq \lambda_i$. First, let $\hat{d}_1 < \lambda_i$. Thus, from Lemma 16, i has a bottleneck constraint and hence a bottleneck bucket, B. Let $\mathcal{X}$ be the set of sessions associated with B. Since i has rank 1, $|\mathcal{X}| = L$, rank$(j) = 1 \ \forall \ j \in \mathcal{X}$, and $\hat{d}_1 = 1/L$.

$$\begin{aligned} &C_{i,B}(y_K) - C_{i,B}(x_1) \\ &\leq y_K - x_1 - \sum_{m \in \mathcal{X}\setminus\{i\}}(C_{m,B}(y_K) - C_{m,B}(x_1)) \\ &\leq y_K - x_1 - (L-1)\left(\hat{d}_1(y_K - x_1) - \varsigma_2(1) - \varepsilon_2(1)\right) \ \text{(from (63) since rank}(j) = 1, \ \forall \ j \ \in \mathcal{X}) \\ &= \hat{d}_1(y_K - x_1) + (L-1)(\varsigma_2(1) + \varepsilon_2(1)) \ \text{(since } \hat{d}_1 = 1/L). \quad (67) \end{aligned}$$

Now, let $\hat{d}_1 = \lambda_i$. Let B be the bucket at the source of i.

$$\begin{aligned} &C_{i,B}(y_K) - C_{i,B}(x_1) \\ &\leq A_i^{\text{NR}}(x_1) + A_i(y_K) - A_i(x_1) \\ &\leq \sigma_i + \varsigma_2(1) + \varepsilon_2(1) + \lambda_i(y_K - x_1) + \sigma_i \ \text{(from (64) and (10))} \\ &= \hat{d}_1(y_K - x_1) + 2\sigma_i + \varsigma_2(1) + \varepsilon_2(1) \ \text{(since } \hat{d}_1 = \lambda_i). \quad (68) \end{aligned}$$

From (67) and (68), there exists a bucket B associated with i such that

$$\begin{aligned} &C_{i,B}(y_K) - C_{i,B}(x_1) \\ &\leq \hat{d}_1(y_K - x_1) + 2\sigma_i + \max(L-1, 1)(\varsigma_2(1) + \varepsilon_2(1)). \quad (69) \end{aligned}$$

$$\text{Now, } |C_{i,n}(t) - C_{i,B}(t)| \leq b_i W \ \forall \ t. \quad (70)$$

$$\begin{aligned}
& C_{i,n}(y_K) - C_{i,n}(x_1) \\
\leq\ & C_{i,B}(y_K) - C_{i,B}(x_1) + 2b_i W \text{ (from (70))} \\
\leq\ & \hat{d}_1(y_K - x_1) + \max(L-1,1)\left(\varsigma_2(1) + \varepsilon_2(1)\right) + 2b_i W + 2\sigma_i \text{ (from (69))}. \qquad (71)
\end{aligned}$$

From (66) and (71),

$$\begin{aligned}
& \sum_{k=1}^{K} \left(C_{i,n}(y_k) - C_{i,n}(x_k)\right) \\
\leq\ & \hat{d}_1 \sum_{k=1}^{K}(y_k - x_k) + \max(L,2)\left(\varsigma_2(1) + \varepsilon_2(1)\right) + 2b_i W + 2\sigma_i + K\varepsilon_2(1). \qquad (72)
\end{aligned}$$

Thus, for $p = 1$, (65) follows from (72) with $\varsigma_3(1) = \max(L,2)\left(\varsigma_2(1) + \varepsilon_2(1)\right) + 2LW + 2\sigma$ and $\varepsilon_3(1) = \varepsilon_2(1)$.

Now, we assume (62) to (65) for $1, \ldots, p$, and show that (62) to (65) hold for $p+1$.

We first prove (62). Consider a session i with rank greater than or equal to $p+1$. Consider a bucket n associated with i. Let $\mathcal{Y} = \{w : w$ is associated with $n, \text{rank}(w) \leq p\}$ and $\mathcal{Z} = \{w : w$ is associated with n, $\text{rank}(w) \geq p+1\}$. In any interval $(x_k, y_k]$,

$$\begin{aligned}
& \sum_{j \in \mathcal{Z}} \left(S_{j,n}(y_k) - S_{j,n}(x_k)\right) + \sum_{j \in \mathcal{Y}} \left(C_{j,n}(y_k) - C_{j,n}(x_k)\right) \\
& \geq y_k - x_k.
\end{aligned}$$

Since sessions are sampled in round robin order, $S_{i,n}(y_k) - S_{i,n}(x_k) \geq S_{j,n}(y_k) - S_{j,n}(x_k) - 1$ for any two sessions i, j associated with n. Thus,

$$\begin{aligned}
& |\mathcal{Z}|\left(S_{i,n}(y_k) - S_{i,n}(x_k) + 1\right) \\
\geq\ & y_k - x_k - \sum_{j \in \mathcal{Y}} \left(C_{j,n}(y_k) - C_{j,n}(x_k)\right).
\end{aligned}$$

Thus,

$$\begin{aligned}
& \sum_{k=1}^{K} \left(S_{i,n}(y_k) - S_{i,n}(x_k)\right) \\
\geq\ & \frac{1}{|\mathcal{Z}|}\left(\sum_{k=1}^{K}(y_k - x_k) - K|\mathcal{Z}| - \sum_{j \in \mathcal{Y}} \sum_{k=1}^{K} \left(C_{j,n}(y_k) - C_{j,n}(x_k)\right)\right) \\
\geq\ & \frac{\left(1 - \sum_{j \in \mathcal{Y}} d_j^*\right) \sum_{k=1}^{K}(y_k - x_k)}{|\mathcal{Z}|} - \frac{|\mathcal{Y}|}{|\mathcal{Z}|}\varsigma_3(p) - K\frac{|\mathcal{Z}| + |\mathcal{Y}|\varepsilon_3(p)}{|\mathcal{Z}|}.
\end{aligned}$$

The last inequality follows since $\text{rank}(w) \leq p$, and $d_w^* = \hat{d}_{\text{rank}(w)}$, $\forall\ w \in \mathcal{Y}$. Also, $\varsigma_3(j) \geq \varsigma_3(j-1)$, $\varepsilon_3(j) \geq \varepsilon_3(j-1)$, $\forall j$. Thus, induction hypothesis (inequality (65)) applies. Now,

$$\begin{aligned}
& \sum_{k=1}^{K} \left(S_{i,n}(y_k) - S_{i,n}(x_k)\right) \\
\geq\ & \frac{\sum_{j \in \mathcal{Z}} d_j^* \sum_{k=1}^{K}(y_k - x_k)}{|\mathcal{Z}|} - \frac{|\mathcal{Y}|}{|\mathcal{Z}|}\varsigma_3(p) - K\frac{|\mathcal{Z}| + |\mathcal{Y}|\varepsilon_3(p)}{|\mathcal{Z}|} \text{ (since } \sum_{w \in \mathcal{Z}} d_w^* + \sum_{w \in \mathcal{Y}} d_w^* \leq 1) \\
\geq\ & \hat{d}_{p+1} \sum_{k=1}^{K}(y_k - x_k) - \frac{|\mathcal{Y}|}{|\mathcal{Z}|}\varsigma_3(p) - K\frac{|\mathcal{Z}| + |\mathcal{Y}|)\varepsilon_3(p)}{|\mathcal{Z}|}. \qquad (73)
\end{aligned}$$

The last step follows since rank(w) $\geq p+1$, and hence $d^*_w \geq \hat{d}_{p+1}$, $\forall\ w \in \mathcal{Z}$. Thus, from (73), (62) holds for $p+1$, with $\varsigma_1(p+1) = (L-1)\varsigma_3(p)$, and $\varepsilon_1(p+1) = (L-1)\varepsilon_3(p)+1$.

Consider a session i with rank greater than or equal to $p+1$. Note that $\lambda_i \geq \hat{d}_{p+1}$, and $W \geq 3^{L-1}(\varepsilon_1(p+1)+\sigma)/2$. Thus, (63) follows from Lemma 31, with $\varsigma_2(p+1) = 2^{L-1}\varsigma_1(p+1)$ and $\varepsilon_2(p+1) = 3^{L-1}(\varepsilon_1(p+1)+\sigma)$.

The proof for (64) is similar to that in the base case.

Now, we prove (65) for $p+1$. The argument is similar to that for the base case. We point out the differences. Consider a session i with rank $p+1$. Let n be a bucket associated with i. Consider any sequence of non-decreasing times $x_1, y_1, \ldots, x_K, y_K$.

$$\begin{aligned}
&\sum_{k=1}^{K}(C_{i,n}(y_k) - C_{i,n}(x_k)) \\
= \;& C_{i,n}(y_K) - C_{i,n}(x_1) - \sum_{k=1}^{K-1}(C_{i,n}(x_{k+1}) - C_{i,n}(y_k)) \\
\leq \;& C_{i,n}(y_K) - C_{i,n}(x_1) - \hat{d}_{p+1}\sum_{k=1}^{K-1}(x_{k+1} - y_k) + \varsigma_2(p+1) + (K-1)\varepsilon_2(p+1). \qquad (74)
\end{aligned}$$

The last inequality follows from (63) for $p+1$.

Since $\hat{d}_{p+1} = d^*_i$ and $d^*_i \leq \lambda_i$, $\hat{d}_{p+1} \leq \lambda_i$. Now, first let $\hat{d}_{p+1} < \lambda_i$. Since $d^*_i = \hat{d}_{p+1}$, $d^*_i < \lambda_i$. Thus, from Lemma 16, i is associated with a bottleneck constraint, and hence a bottleneck bucket, B. Let $\mathcal{X}$ be the set of sessions associated with B. Since i has rank $p+1$, ranks of all sessions associated with B are less than or equal to $p+1$.

$$\begin{aligned}
&C_{i,B}(y_K) - C_{i,B}(x_1) \\
\leq \;& y_K - x_1 - \sum_{m\in\mathcal{X}\setminus\{i\}}(C_{m,B}(y_K) - C_{m,B}(x_1)) \\
\leq \;& y_K - x_1 - \sum_{m\in\mathcal{X}\setminus\{i\}}(d^*_m(y_K - x_1) - \varsigma_2(p+1) - \varepsilon_2(p+1)) \ \text{(from (63))} \\
= \;& \hat{d}_{p+1}(y_K - x_1) + (|\mathcal{X}|-1)(\varsigma_2(p+1) + \varepsilon_2(p+1)). \qquad (75)
\end{aligned}$$

The last step follows since $\hat{d}_{p+1} + \sum_{m\in\mathcal{X}\setminus\{i\}} d^*_m = 1$.

Now, let $\hat{d}_{p+1} = \lambda_i$. Let B be the bucket at the source node of i. Like in the base case, using (63) and (10), we can prove that

$$C_{i,B}(y_K) - C_{i,B}(x_1) \;\leq\; \hat{d}_{p+1}(y_K - x_1) + 2\sigma_i + \varsigma_2(p+1) + \varepsilon_2(p+1). \qquad (76)$$

From (75) and (76), there exists a bucket B associated with i such that,

$$\begin{aligned}
&C_{i,B}(y_K) - C_{i,B}(x_1) \\
\leq \;& \hat{d}_{p+1}(y_K - x_1) + 2\sigma_i + \max(L-1,1)\left(\varsigma_2(p+1) + \varepsilon_2(p+1)\right). \qquad (77)
\end{aligned}$$

From (77), like in the base case,

$$\begin{aligned}
&C_{i,n}(y_K) - C_{i,n}(x_1) \\
\leq \;& \hat{d}_{p+1}(y_K - x_1) + 2\sigma_i + 2b_iW + \max(L-1,1)\left(\varsigma_2(p+1) + \varepsilon_2(p+1)\right). \qquad (78)
\end{aligned}$$

From (74) and (78),

$$\begin{aligned}
&\sum_{k=1}^{K}(C_{i,n}(y_k) - C_{i,n}(x_k)) \\
\leq \;& \hat{d}_{p+1}\sum_{k=1}^{K}(y_k - x_k) + 2b_iW + 2\sigma_i + K\varepsilon_2(p+1) + \max(L,2)\left(\varsigma_2(p+1) + \varepsilon_2(p+1)\right). \qquad (79)
\end{aligned}$$

Thus, (65) follows from (79) with $\varsigma_3(p+1) = \max(L,2)\,(\varsigma_2(p+1)+\varepsilon_2(p+1))+2LW+2\sigma$ and $\varepsilon_3(p+1) = \varepsilon_2(p+1)$. Thus, (62) to (65) hold in the induction case.

Note that $\varsigma_i(x), \varepsilon_i(x)$ are increasing in both i and x. Thus, from (62), (63) and (65), Lemma 32 holds with $\beta = \varsigma_3(F)$ and $\gamma = \varepsilon_3(F)$. ■

Appendix S: Proof of Theorem 3

We present the proof when each session traverses one link. Let $A_i^{\mathrm{R}}(t)$ be the number of packets of session i that have been released at its source node in $(0,t]$. Note that a packet is released for session i at its source if and only if a new token is generated for session i at the bucket at its source. Thus, $\forall\ t, A_i^{\mathrm{R}}(t) = C_{i,n}(t)$ where n is the bucket at i's source. Now, from Lemma 17, there exists constants ϱ, W_0, such that when $W \geq W_0,\ \forall\ t,\ |\frac{A_i^{\mathrm{R}}(t)}{t} - d_i^*| \leq \frac{\varrho}{t}$. Thus, the packet release rate vector is $\vec{d^*} \in \Lambda^{\mathrm{MS}}$. Since only the released packets are available for scheduling and the release rate vector is in Λ^{MS}, the departure rate vector exists and equals the release rate vector. The result follows. ■