A Simulation Study of IEEE 802.11 with Optimal Rate Control in Multi-hop Topologies

Apoorva Jindal, Konstantinos Psounis

◆

1 INTRODUCTION

The most commonly used scheduling scheme to build wireless multi-hop networks is CSMA-CA as off-the shelf IEEE 802.11 radios can be used to easily and cheaply deploy a multi-hop network [5], [6], [24]. However, several researchers have observed that using CSMA-CA with either TCP or without any rate control and pumping data as fast as possible leads to unfair and inefficient throughouts, and even starvation for certain flows [9], [11], [16], [20], [22], [27]–[29].

Researchers have explored two different directions to allocate fair and efficient throughputs in wireless multi-hop networks. (i) New rate control protocols while retaining the scheduling protocol to be CSMA-CA [22], [27], [29]. (ii) New scheduling alternatives to CSMA-CA [7], [15], [18], [25]. An important question to ask here is: Which amongst the two is a better direction? And the answer to this question is dictated by the answer to the following question. Is it possible to allocate fair and efficient rates over CSMA-CA? If the answer to this question is yes, then the former approach is a better direction because of economic reasons as off-the-shelf IEEE 802.11 radios are easily available at low-cost. If the answer to this question is no, then obviously, the latter approach is the better direction.

To answer the proposed question, we evaluate CSMA-CA with optimal rate control and compare it to optimal scheduling with optimal rate control for a number of carefully constructed multi-hop topologies. Through this exercise, we make the following contributions.

(i) We propose a new, general methodology to evaluate CSMA-CA with optimal rate control. This methodology yields an important simulation benchmark to evaluate any new rate control and scheduling protocol for multi-hop networks. The importance of this benchmark lies in the following observation. Both TCP and no rate control yield extremely unfair and inefficient throughputs for most multi-hop topologies. Thus, any new rate control protocol which does much better than TCP/no rate control over CSMA-CA may still be orders of magnitude away from the best or optimal throughput achievable with CSMA-CA. Similarly, a new scheduling protocol may yield better rates with TCP or with no rate control than CSMA-CA, but its capacity region may be much smaller than that of CSMA-CA. Hence, the performance of any new rate control and scheduling protocol should be compared to the performance of CSMA-CA with optimal rate control using the methodology proposed in this work.

(ii) We carefully construct a number of topologies, each with different interference characteristics. This rich set of topologies will also serve as simulation benchmarks to evaluate new protocols for multi-hop networks. For the convenience of researchers, we also select a smaller set of topologies which capture all the different facets of multi-hop networks, and hence, make perfect candidates to evaluate a new protocol.

(iii) *Our most important contribution is to establish that CSMA-CA performs reasonably well in multi-hop networks. Specifically, CSMA-CA is always within* 40% *of the optimal for all topologies we study, and, for most topologies, the throughput loss is less than* 30%*, that is, CSMA-CA is within* 70% *of the optimal. Finally, for all topologies which are expected to be constructed in practice, CSMA-CA yields throughputs within* 75% *of the optimal.*

Our final contribution has the following two implications. (i) For topologies which will occur in practice, our results establish that CSMA-CA with a well-designed rate control scheme will achieve close-to-optimal throughputs, and its not clear if alternative scheduling schemes can do better than CSMA-CA. (ii) Our results also construct topologies for which CSMA-CA cannot allocate efficient throughputs, even with optimal rate control. Thus, network designers can use these results to determine what to avoid when building a mesh network. Also, these topologies will serve as benchmark topologies to evaluate alternative scheduling protocols.

2 METHODOLOGY

To understand whether fair and efficient rates are achievable with CSMA-CA or not, we compare the following four results for a number of carefully constructed

---

- *Apoorva Jindal is with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor.*
  E-mail: apoorvaj@umich.edu
- *Konstantinos Psounis is with the Department of Electrical Engineering, University of Southern California, Los Angeles.*
  E-mail: kpsounis@usc.edu

topologies. (i) The achievable rate region with optimal scheduling, (ii) the achievable rate region with CSMA-CA, (iii) the rates achieved by TCP over CSMA-CA, and (iv) throughputs achieved assuming that all sources always have a packet to send (no rate control) with CSMA-CA. This is referred to as the saturation throughput [4] and has been studied in detail for CSMA-CA in multi-hop networks [9], [11], [17].

2.1 Deriving Each Result

We first describe the methodology adopted to derive each result. Note that a transport protocol like TCP also guarantees reliability. The destination sends back an ACK packet for each received DATA packet. This results in a flow of transport layer ACK packets from the destination to the source. Since the ACK packet flow will interfere with the flow of DATA packets, it will reduce the achievable rate region. For a fair comparison, we will incorporate the effect of the flow of ACK packets while deriving the achievable rate regions and saturation throughput. The size of the ACK packets is set to the size of ACK packets in TCP.

(i) Achievable rate region with optimal scheduling: We use the methodology proposed by Jain *et al.* [12]. For a fair comparison, we assume that the overhead imposed by the control message exchange and protocol headers is equal for IEEE 802.11 (the protocol which implements CSMA-CA scheduling) and optimal scheduling[1]. The control messages exchanged for IEEE 802.11 scheduling are: RTS, CTS and 802.11-ACK[2]. And the protocol headers included in the DATA packet are: PHY, MAC, IP and UDP headers. We measure the bandwidth consumed by these messages and headers, and factor it in the calculations for optimal scheduling. This ensures that the loss in throughput with CSMA-CA is entirely due to inefficiency in scheduling and random backoffs.

(ii) Achievable rate region with CSMA-CA: We simulate sources generating DATA packets and destinations generating ACK packets at a constant bit rate. UDP is used at the transport layer. To derive the entire region, we simulate all possible combinations of flow rates with each flow rate varying from 0 to 1 Mbps in steps of 10Kbps and checking if the input rates were achievable[3]. Note that for a given rate to be achievable, the network should be able to support both the DATA and ACK flows.

(iii) Performance of TCP over CSMA-CA: We simulate sources generating FTP traffic and use TCP at the transport layer to deliver packets. We use TCP SACK but with Nagle's algorithm [19] and the delayed ACK mechanism [8] turned off in our simulations.

(iv) Saturation throughput: We simulate two CBR flows (DATA from source to the destination and ACK from the destination to the source) for each source-destination pair. Rates of both the flows is set to 100 Mbps so that the queues at the source are always full.

Packet Payload	1024 bytes
Transport layer ACKs	40 bytes
MAC Header	34 bytes
PHY Header	16 bytes
ACK	14 bytes + PHY header
RTS	20 bytes + PHY header
CTS	14 bytes + PHY header
Channel Bit Rate	1 and 11 Mbps
Propagation Delay	1 μs
Slot Time	20 μs
SIFS	10 μs
DIFS	50 μs
W_0	31
m	6

TABLE 1
Simulation parameters.

2.2 Simulation Setup

To generate simulation results with CSMA-CA scheduling, we use Qualnet 4.0 [2] as the simulation platform. All our simulations are conducted using an unmodified IEEE 802.11(b) MAC (DCF) with RTS/CTS. We use the default parameters of IEEE 802.11(b) (summarized in Table 1) in Qualnet unless otherwise stated. Auto-rate adaptation at the MAC layer is turned off. Simulations are conducted using for the following two data rates - 1 Mbps and 11 Mbps, which are the minimum and the maximum data rates allowed by IEEE 802.11(b). All our simulations are conducted with zero channel losses. We set the buffer size and maximum retry limit in IEEE 802.11 (the number of retransmission attempts after which the packet is dropped) to a very large value to avoid packet losses. This allows us to generate the achievable rate region without having to worry about transport layer retransmissions to recover from these losses. The packet size is fixed to be 1024 bytes. We use AODV [21] to set up the routes. Strict priority is given to control packets exchanged to set up the routes over DATA packets. FIFO is used to serve packets of the same priority.

2.3 Proportionally Fair and Max Sum Throughput Rate Allocations

Using the achievable rate region, it is easy to determine both the proportionally fair rate allocation and the rate allocation which maximizes the sum throughput. We move along the boundary of the achievable rate region and determine the allocation which maximizes

---

[1] We expect the actual overhead required to achieve optimal scheduling to be much higher. Hence, the comparison is geared to favor optimal scheduling which makes the obtained good results for CSMA-CA even stronger.

[2] We call the ACK messages exchanged by the IEEE 802.11 protocol as 802.11-ACK to distinguish it from the transport layer ACKs.

[3] Note that simulating every possible combination of flow rates to determine the achievable rate region does not scale. Our prior work [13] presents an analytical method to determine the achievable rate region of any given topology for CSMA-CA-scheduled networks. However, since most topologies used in this paper are small, we use simulations to determine the achievable rate region.

$\sum_{f \in F} log\left(x_f\right)$ and $\sum_{f \in F} x_f$ respectively, where F denotes the set of end-to-end flows and x_f denotes the rate of flow $f \in F$.

2.4 Max-Min Fair Rate Allocation

Unless explicitly stated, if there are more than 3 flows in the network, we do not calculate the entire achievable rate region for CSMA-CA because simulating every possible combination of flow rates to determine the achievable rate region does not scale. Instead, we will compare the max-min fair rate allocation with CSMA-CA and optimal scheduling because it can be derived through simulations without having to find the entire rate region.

First, we define the max-min fair rate allocation and then present the methodology used to derive it.

A feasible allocation of rates $\vec{x}$ is max-min fair if and only if an increase of any rate within the domain of feasible allocations must be at the cost of a decrease of some already smaller rate. Formally, for any other feasible allocation $\vec{y}$, if $y_f > x_f$, $f \in F$, then there must exist some $f' \in F$ such that $x_{f'} \leq x_f$ and $y_{f'} < x_{f'}$ [3].

We use the following methodology to determine the max-min fair rate allocation. (a) Initialize the rate of all flows to be 0 Kbps. Let S denote a set of flows. Place all flows in the topology in S. (b) Keep increasing the rate of all flows in S in steps of 10Kbps till one of the queue becomes fully utilized (arrival rate = service rate). Label this queue Q. (c) Assign the current rate to the flows which pass through the neighborhood of Q (see Section 2.6 for a precise definition of neighborhood) and remove these flows from S. (d) If there are any flows remaining in S, go to step (b).

2.5 Comparing CSMA-CA and Optimal Scheduling

To compare CSMA-CA and optimal scheduling, we compare the sum throughput at the max-min allocation with CSMA-CA and optimal scheduling. Specifically, if $x_f^{CSMA-CA}$ and x_f^{OPT} denote the rate of flow $f \in F$ at the max-min rate allocation with CSMA-CA and optimal scheduling respectively, then we say that CSMA-CA is within $x\%$ of the optimal if $x = \frac{100 \sum_{f \in F} x_f^{CSMA-CA}}{\sum_{f \in F} x_f^{OPT}}$.

Amongst the commonly used rate allocation points, like proportionally fair rate allocation, maximum sum throughput rate allocation and max-min rate allocation, we choose the max-min rate point to compare because of the following two reasons. (i) For topologies where CSMA-CA suffers a throughput loss of more than 40%, amongst the commonly used rate allocation points, the worst throughput ratio is observed at the max-min rate allocation. (ii) The max-min fair rate point can be determined without having to find the entire achievable rate region. This allows us to compare CSMA-CA and optimal scheduling for larger topologies with several flows.

2.6 Congested Neighborhoods

We formally defined congested neighborhoods in our prior work [22]. For completeness, we reproduce the definition here. We then describe the methodology used to determine the congested neighborhoods in a topology and then comment on their significance.

We first define when two edges are said to interfere with each other, then define the term neighborhood and finally explain the term congested neighborhood. Let T_e and R_e denote the transmitter and the receiver of an edge e. Two edges e_1, e_2 are said to interfere with each other if either T_{e_1} interferes with either T_{e_2} or R_{e_2}, or R_{e_1} interferes with either T_{e_2} or R_{e_2}. A neighborhood of an edge e is defined to be the set of edges which interfere with e. At the max-min allocation, the edges whose queues are fully utilized (arrival rate = service rate) are defined to be congested edges, and the neighborhood of congested edges is defined to be congested neighborhoods.

To determine the congested neighborhood, we first determine the max-min allocation using the methodology described in Section 2.4. We then simulate the topology with CBR flows and fixing the rate of each end-to-end flow to be equal to its corresponding rate in the max-min allocation. To account for simulation errors, any edge whose queue utilization is greater than 0.95 is labelled a congested edge and its neighborhood is marked as a congested neighborhood. For each topology, the congested edge will be denoted by a symbol depicting a queue.

We will use this information in the intuitive discussion about the characteristics of each topology as congested neighborhoods form the basic building blocks in any topology. How two flows interfere with each other [4] can be directly studied by studying how edges interfere in the congested neighborhood shared by both flows.

3 CSMA-CA WITH OPTIMAL RATE CONTROL

In this section, we compare the achievable rate region with CSMA-CA and optimal scheduling for several different multi-hop topologies. We divide the topologies into six groups. For each group, we first describe the characteristics of the topologies belonging to that group, and then present results for specific topologies

Before studying multi-hop topologies, we first present the throughput results on an one-edge topology with one transmitter and one receiver and no interfering edge in Table 2. This serves as a baseline as it allows us to determine the throughput loss which can be attributed to different overheads. For multi-hop topologies, this baseline case allows us to determine how much is the additional loss due to scheduling inefficiencies and collisions. We make the following observations from these results. (i) At 1 Mbps, CSMA-CA is within 96% of the optimal, while at 11 Mbps, it is within 85%. We first discuss why CSMA-CA achieves a lower throughput

---

[4] Two flows are said to interfere with each other if they flow through edges which interfere with each other.

Experiment Description	Throughput with 1 Mbps data rate	Throughput with 11 Mbps data rate
Achivable rate with optimal scheduling	0.695 Mbps	3.23
Achievable rate with CSMA-CA	0.67 Mbps	2.76 Mbps
TCP over CSMA-CA	0.66 Mbps	2.62 Mbps
Saturation Throughput	0.67 Mbps	2.76 Mbps

TABLE 2
Results for a Single-edge topology

with 11 Mbps data rate. Before each packet transmission, the transmitter back-offs (waits) for a randomly selected duration. The expected value of this random duration is a constant irrespective of the data rate. Thus, a fraction of available throughput is lost due to back-offs. We refer to this loss in throughput as the random access overhead. Now, lower the packet transmission time, more is the random access overhead, which explains the extra loss in throughput at the higher data rate. (ii) Even though the data rate is 11 Mbps, even with optimal scheduling, one can achieve a throughput of only 3.23 Mbps. Recall that the control overhead of optimal scheduling is assumed to be the same as that of CSMA-CA. Thus, more than 70% of the available throughput is consumed by control overhead (which includes MAC control packets and protocol headers). Which aspects of the protocol introduce this significant control overhead is discussed in detail in Section 3.6.

3.1 Two-Edge Topologies

The four two-edge topologies [11], [14] represent the fundamental ways in which edges interact with each other.

Coordinated Stations: In this topology, the two transmitters interfere with each other. Figure 1(a) shows an example of this topology. This is similar to a single-hop WLAN topology for which CSMA-CA was originally designed. Hence, we expect the performance of CSMA-CA to be close to the optimal. Figures 1(b) and 1(c) compare the achievable rate region of CSMA-CA and optimal scheduling at 1 and 11 Mbps data rates respectively. We make the following observations. (i) At the max-min allocation, CSMA-CA is within 97.1% and 88.8% of the optimal respectively. (ii) For this topology, the max-min rate allocation is also proportionally fair and maximizes the sum throughput. (iii) TCP and saturation yield fair and efficient throughputs.

Near Hidden Edges: Both transmitters do not interfere with each other, but they interfere with each other's receiver. Figure 1(d) depicts the topology. The RTS/CTS control messages were introduced to avoid significant throughput losses due to collisions in this topology. Hence, we expect CSMA-CA to yield rates close to the optimal for this topology too. Figures 1(e) and 1(f) compare the achievable rate regions at 1 and 11 Mbps data rates respectively. We make the following observations. (i) At the max-min allocation, CSMA-CA is within 94.2% and 81.4% of the optimal respectively. (ii) The max-min rate allocation is also proportionally fair. (iii) Starving one of the flows maximizes the sum throughput. (iv) TCP allocates unfair rates. (v) Saturation yields fair but inefficient rates.

Asymmetric Topology: The transmitter of the second edge interferes with the receiver of the first edge, while the transmitter of the first edge does not interfere with either the transmitter or the receiver of the second edge. Figure 1(g) depicts the topology. Since node 3 can overhear the CTS from node 2, its aware of the transmission on edge 1 → 2, however, node 1 will never be aware of the transmission on edge 3 → 4. Thus, the topology is asymmetric. Figures 1(h) and 1(i) compare the achievable rate regions at 1 and 11 Mbps data rates respectively. We make the following observations. (i) At the max-min allocation, CSMA-CA is within 80.7% and 74% of the optimal respectively. *Amongst the four two-edge topologies, CSMA-CA suffers the maximum throughput loss in this topology.* (ii) The max-min rate allocation is also proportionally fair. (iii) Starving one of the flows maximizes the sum throughput. (iv) TCP and saturation starve the flow 1 → 2.

Far Hidden Edges: Only the receivers interfere in this topology. Figure 1(j) depicts the topology. Figures 1(k) and 1(l) compare the achievable rate regions at 1 and 11 Mbps data rates respectively. We make the following observations. (i) At the max-min allocation, CSMA-CA is within 86.8% and 78.7% of the optimal respectively. (ii) The max-min rate allocation is also proportionally fair. (iii) Starving one of the flows maximizes the sum throughput. (iv) Both TCP and saturation yield fair but inefficient rates.

3.2 Flow in the Middle and Variants

In topologies belonging to this category, each flow experiences a different level of interference. For example, in the topology presented in Figure 2(a), the middle flow interferes with the two outer flows while the outer flows do not interfere with each other; hence the middle flow experiences more interference than the outer two flows. With TCP and with no rate control (saturation) over CSMA-CA, the flow which experiences a higher level of interference experiences unfair throughputs, and even starvation. Hence, these topologies have been extensively studied in the literature by works which focus on the unfairness and starvation issues with TCP/saturation over CSMA-CA [11], [22], [28]. Hence, to avoid repetition, we will not discuss the throughput results with TCP and saturation for these set of topologies. Instead, we focus on the throughputs achievable with optimal rate control.

Flow in the Middle: Figure 2(a) depicts the topology. There is only one congested edge - 3 → 4. All three flows contribute to congestion on this edge. The middle flow 3 → 4 interferes with both the outer flows, while the

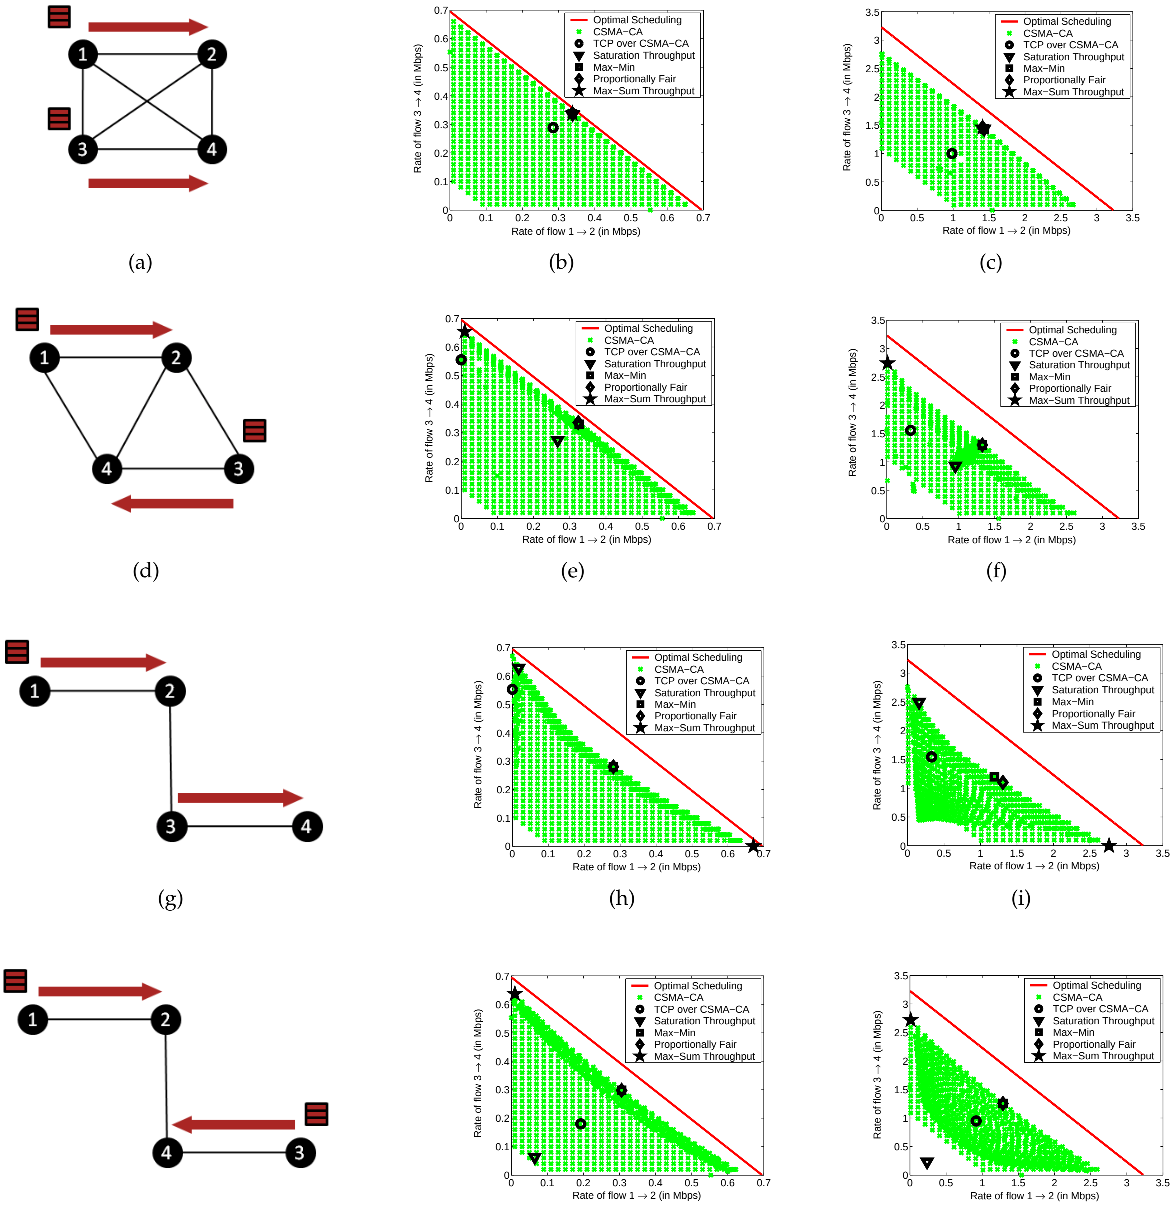

(a) (b) (c)

(d) (e) (f)

(g) (h) (i)

(j) (k) (l)

Fig. 1. Coordinated Stations. (a) The topology. (b) Achievable rate region with data rate $= 1$ Mbps. (c) Achievable rate region with data rate $= 11$ Mbps. Near Hidden. (d) The topology. (e) Achievable rate region with data rate $= 1$ Mbps. (f) Achievable rate region with data rate $= 11$ Mbps. Asymmetric. (g) The topology. (h) Achievable rate region with data rate $= 1$ Mbps. (i) Achievable rate region with data rate $= 11$ Mbps. Far Hidden. (j) The topology. (k) Achievable rate region with data rate $= 1$ Mbps. (l) Achievable rate region with data rate $= 11$ Mbps.

outer flows do not interfere with each other. Figures 2(b) and 2(c) compare the achievable rate regions at 1 and 11 Mbps data rates respectively. (By symmetry, the outer two flows will achieve approximately the same rate for any scheme.) We make the following observations. (i) At the max-min allocation, CSMA-CA is within 72.6% and 62.5% of the optimal respectively. Thus, CSMA-CA with a well-designed rate control protocol will yield fair and efficient throughputs for this topology. (ii) Proportional fairness allocates a lower rate to the middle flow than the outer two flows. Also, starving the middle flow maximizes the sum throughput. These two are common trends which we will observe in all the topologies belonging to this category. The reason is as follows. Allocating a non-zero rate to the middle flow will reduce the throughput of both the outer flows. Hence, proportional fairness allocates a lower rate to the middle flow and allocating a zero rate to the middle flow maximizes the sum throughput. To avoid repetition, we will not repeat these two observations for other topologies of this category.

Stack: Figure 2(d) depicts the topology. It is similar to the flow in the middle topology, except now each flow goes through two hops instead of one. Also, only the

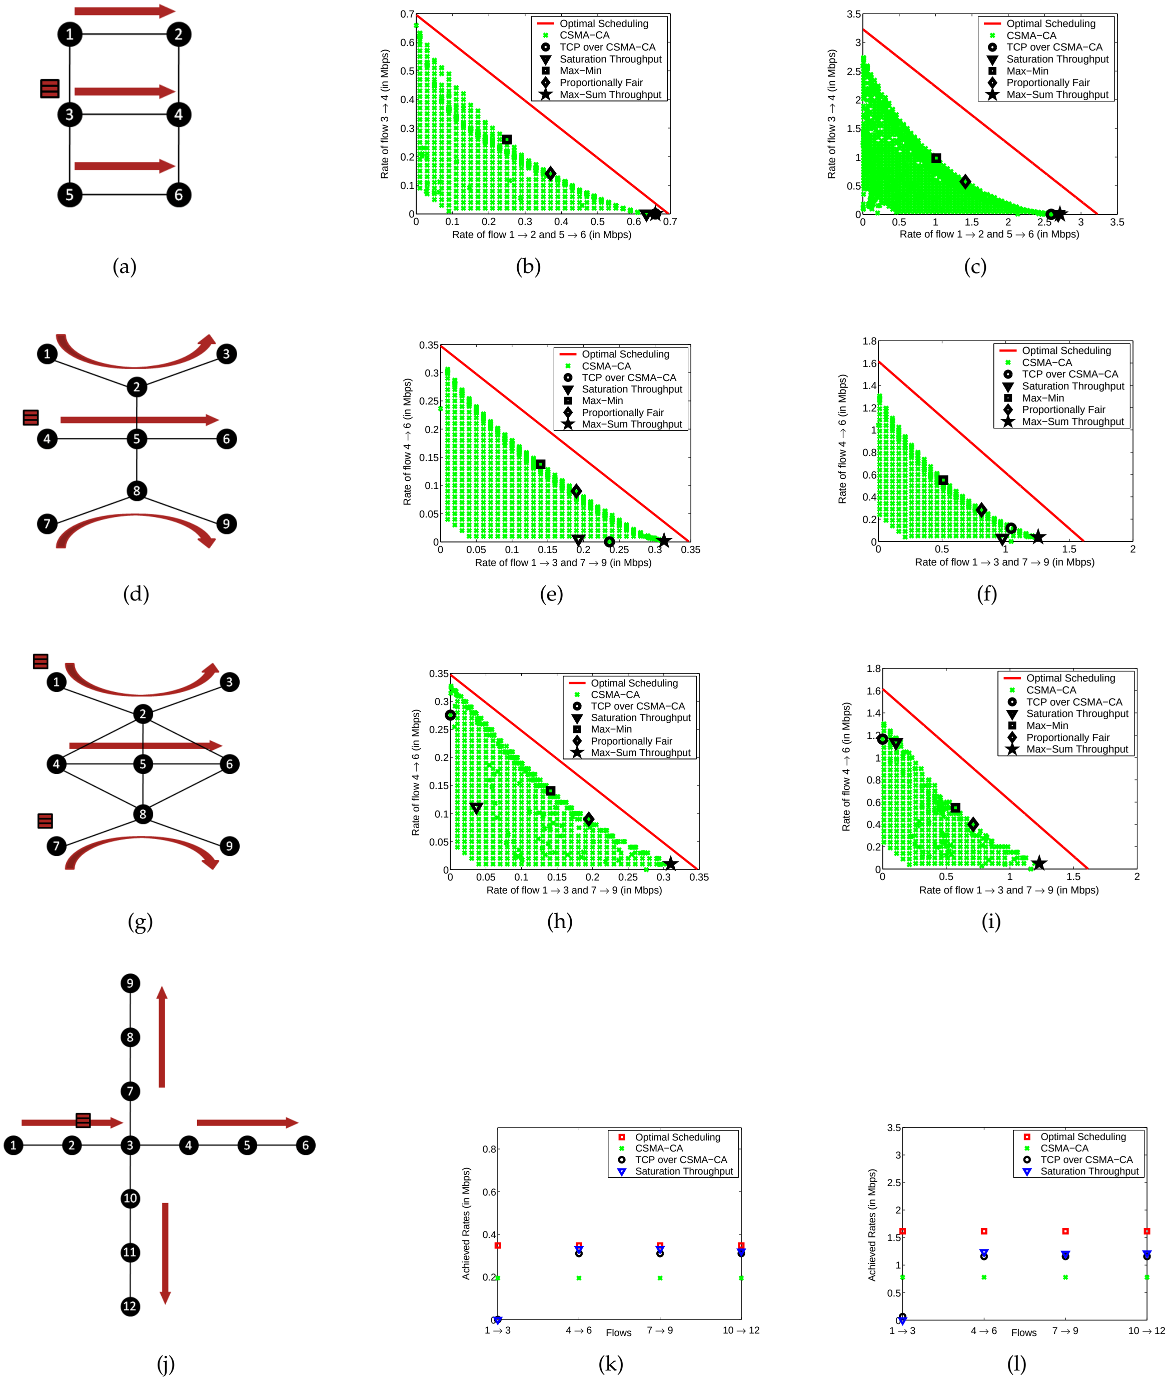

Fig. 2. Flow in the Middle. (a) The topology. (b) Achievable rate region with data rate $= 1$ Mbps. (c) Achievable rate region with data rate $= 11$ Mbps. Stack. (d) The topology. (e) Achievable rate region with data rate $= 1$ Mbps. (f) Achievable rate region with data rate $= 11$ Mbps. Diamond. (g) The topology. (h) Achievable rate region with data rate $= 1$ Mbps. (i) Achievable rate region with data rate $= 11$ Mbps. Fork. (j) The topology. (k) Max-min rate allocation with data rate $= 1$ Mbps. (l) Max-min rate allocation with data rate $= 11$ Mbps.

first-hop nodes, 2 and 5, and 5 and 8, interfere with each other. There is still only one congested edge $4 \rightarrow 5$. Figures 2(e) and 2(f) compare the achievable rate regions at 1 and 11 Mbps data rates respectively. (By symmetry, the outer two flows will achieve approximately the same rate for any scheme.) At the max-min allocation, CSMA-CA is within 79.9% and 64.7% of the optimal respectively. *As expected, reducing interference on nodes* $1, 3, 4, 6, 7$ *and* 9 *has increased the rate region with CSMA-CA as compared to the flow in the middle topology.*

Diamond: Figure 2(g) depicts the topology. It is similar to the stack topology, except now nodes 4 and 6 also interfere with nodes 2 and 8. This changes the congested edges. Now, there are two congested edges $1 \rightarrow 2$ and $7 \rightarrow 8$ which lie on the outer two flows. Figures 2(h) and 2(i) compare the achievable rate regions at 1 and 11 Mbps data rates respectively. (By symmetry, the outer two flows will achieve approximately the same rate for any scheme.) We make the following observations. (i) At the max-min allocation, CSMA-CA is within 80.9% and 69.8% of the optimal respectively. Thus, making nodes 4 and 6 more aware of the interference in the topology has increased CSMA-CA's rate region. (ii) Now, TCP and saturation starve the outer two flows and not the flow in the middle. This occurs because the outer two flows go over the congested edge and suffer from losses instead of the middle flow [22].

Fork: Figure 2(j) depicts the topology. It is similar to the flow in the middle topology except that now the middle flow interferes with three non-interfering flows instead of just two. Figures 2(k) and 2(l) compare the max-min rate allocation [5] at 1 and 11 Mbps data rates respectively. At the max-min allocation, CSMA-CA is within 56.1% and 48.3% of the optimal respectively. Thus, there is an additional loss in throughput over the flow in the middle topology. *Hence, we conclude that more the number of non-interfering flows which interfere with the middle flow, worse is the performance of CSMA-CA.*

3.3 Chain and Variants

In topologies belonging to this category, there is at least one flow which goes over multiple hops. With TCP and with no rate control over CSMA-CA, the throughput achieved for the flow which goes over multiple hops is very inefficient. Hence, these topologies have been extensively studied in the literature which focus on the inefficiencies in throughput with TCP/saturation over CSMA-CA [16], [20], [27]. However, with optimal rate control, we observe that CSMA-CA allocates rates within 50% of the optimal for the topologies belonging to this category.

Chain: Figure 3(a) depicts the topology. It has two long flows, $1 \rightarrow n$ and $n \rightarrow 1$. Figures 3(b) and 3(c) compare the achievable rate regions at 1 and 11 Mbps data rates respectively for $n = 7$. We make the following observations. (i) At the max-min allocation, CSMA-CA is within 58.3% and 50.4% of the optimal respectively. (ii) The max-min rate allocation is also proportionally fair. (iii) Sum throughput is maximized at the max-min rate allocation at 11 Mbps. At 1 Mbps, allocating a higher throughput to the forward flow $1 \rightarrow n$ maximizes the sum throughput. (iii) TCP and saturation allocate inefficient rates over CSMA-CA. (iv) *We performed additional simulations for different values of* $4 \leq n \leq 15$, *and observed that changing the number of hops does not change the performance of CSMA-CA with optimal rate control.*

Chain with Two Interfering Short Flows: Figure 3(d) depicts the topology. It has one long flow and two short flows $2 \rightarrow 8$ and $2 \rightarrow 9$, both of which interfere with each other as well as the long flow. Figures 3(e) and 3(f) compare the achievable rate regions at 1 and 11 Mbps data rates respectively. (By symmetry, the two short flows will achieve approximately the same rate for any scheme.) We make the following observations. (i) At the max-min allocation, CSMA-CA is within 82.5% and 63.9% of the optimal respectively. *Thus, contrary to intuition, increasing the interference around the congested edge improves the throughput ratio.* (ii) The proportional fair rate allocation allocates a higher rate to the two smaller flows. (iii) Allocation a zero rate to the long flow maximizes the sum throughput. The reason is the same as discussed for the flow in the middle topology. (iv) TCP and saturation starve the long flow.

Smaller Chain with Two Interfering Short Flows: Figure 3(g) depicts the topology. It is similar to the previous topology, except that now the long flow goes over only 4 hops instead of 6. Figures 3(h) and 3(i) compare the achievable rate regions at 1 and 11 Mbps data rates respectively. (By symmetry, the two short flows will achieve approximately the same rate for any scheme.) We observe that the rate region for CSMA-CA has increased, and now, at the max-min allocation, it is within 87.1% and 82.6% of the optimal respectively. *Thus, for a topology with shorter flows interfering with the long flow, reducing the number of hops improves the performance of CSMA-CA.* However, a reduction in the number of hops does not change any other observations.

Chain with Three Interfering Short Flows: Figure 3(j) depicts the topology. It is similar to the topology depicted in Figure 3(d) except for the additional flow $6 \rightarrow 7$. Now there are two congested neighborhoods in the topology, around edges $1 \rightarrow 2$ and $4 \rightarrow 5$, and the long flow passes through both. Figures 3(h) and 3(i) compare the achievable rate regions at 1 and 11 Mbps data rates respectively. (By symmetry, the two short flows $2 \rightarrow 8$ and $2 \rightarrow 9$ will achieve approximately the same rate for any scheme.) We observe that the rate region for CSMA-CA decreases, and now, at the max-min allocation, it is within 57.2% and 54.5% of the optimal respectively. *Hence, we conclude that more the number of congested neighborhoods the long flow goes through, the smaller is CSMA-CA's achievable rate region.*

---

[5] Recall that we compare only the max-min rate allocation for topologies with more than 3 flows.

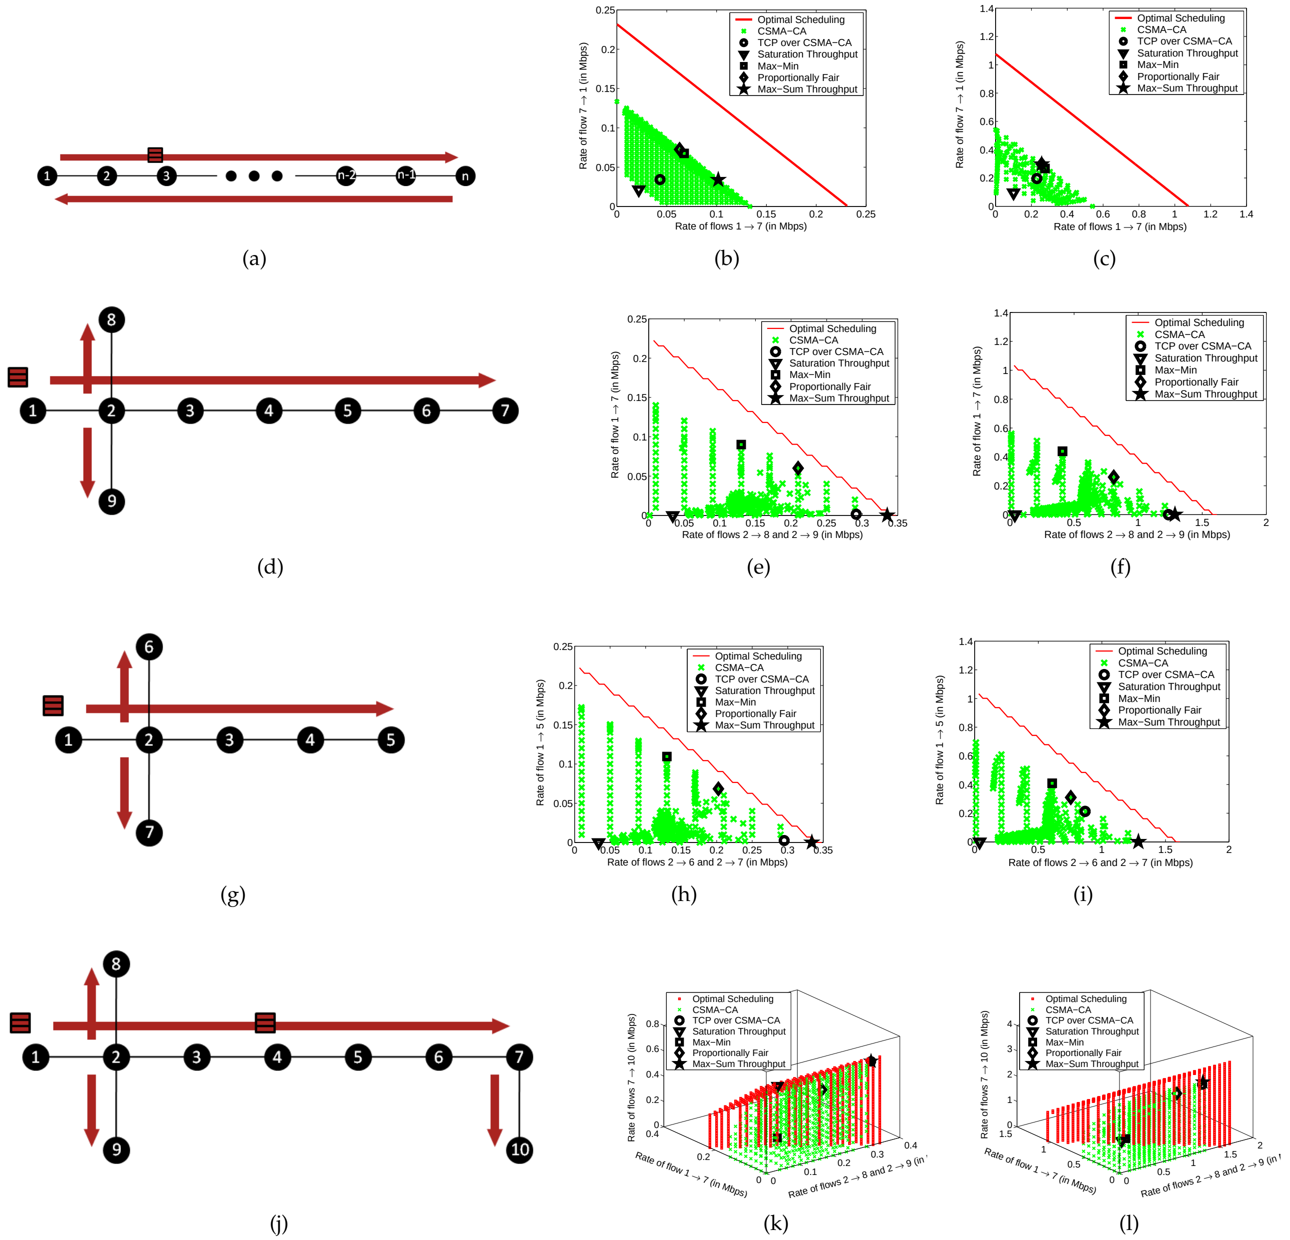

Fig. 3. Chain. (a) The topology. (b) Achievable rate region with data rate $= 1$ Mbps. (c) Achievable rate region with data rate $= 11$ Mbps. Chain with one Interfering Flow. (d) The topology. (e) Achievable rate region with data rate $= 1$ Mbps. (f) Achievable rate region with data rate $= 11$ Mbps. Smaller Chain with One Interfering Flow. (g) The topology. (h) Achievable rate region with data rate $= 1$ Mbps. (i) Achievable rate region with data rate $= 11$ Mbps. Chain with Two Interfering Flows. (j) The topology. (k) Achievable rate region with data rate $= 1$ Mbps. (l) Achievable rate region with data rate $= 11$ Mbps.

3.4 Combining Flow in the Middle and Chain

In this category of topologies, there exists at least one flow which goes over multiple hops as well as experiences a higher level of interference than other flows. This paper is the first to study this category of topologies. The importance of these topologies is that both the starvation and inefficiency issues observed with TCP and with no rate control over CSMA-CA can be observed in a single topology. Also, for both data rates, with optimal rate control, CSMA-CA has the lowest throughput for this category of topologies amongst all the categories we study. Hence, these topologies will serve as important benchmarks to evaluate any new rate control and scheduling protocol.

Chain-Cross: Figure 4(a) depicts the topology. It is similar to the topology depicted in Figure 3(j) except for the additional flow $1 \rightarrow 2$ and the short flows around node 2 ($8 \rightarrow 9$ and $10 \rightarrow 11$) do not interfere with each other. Thus, flow $1 \rightarrow 2$ is a flow in the middle, and the long flow $1 \rightarrow 7$ is not only a flow in the middle but also goes over 7 hops and two congested neighborhoods. Figures 4(b) and 4(c) compare the max-min

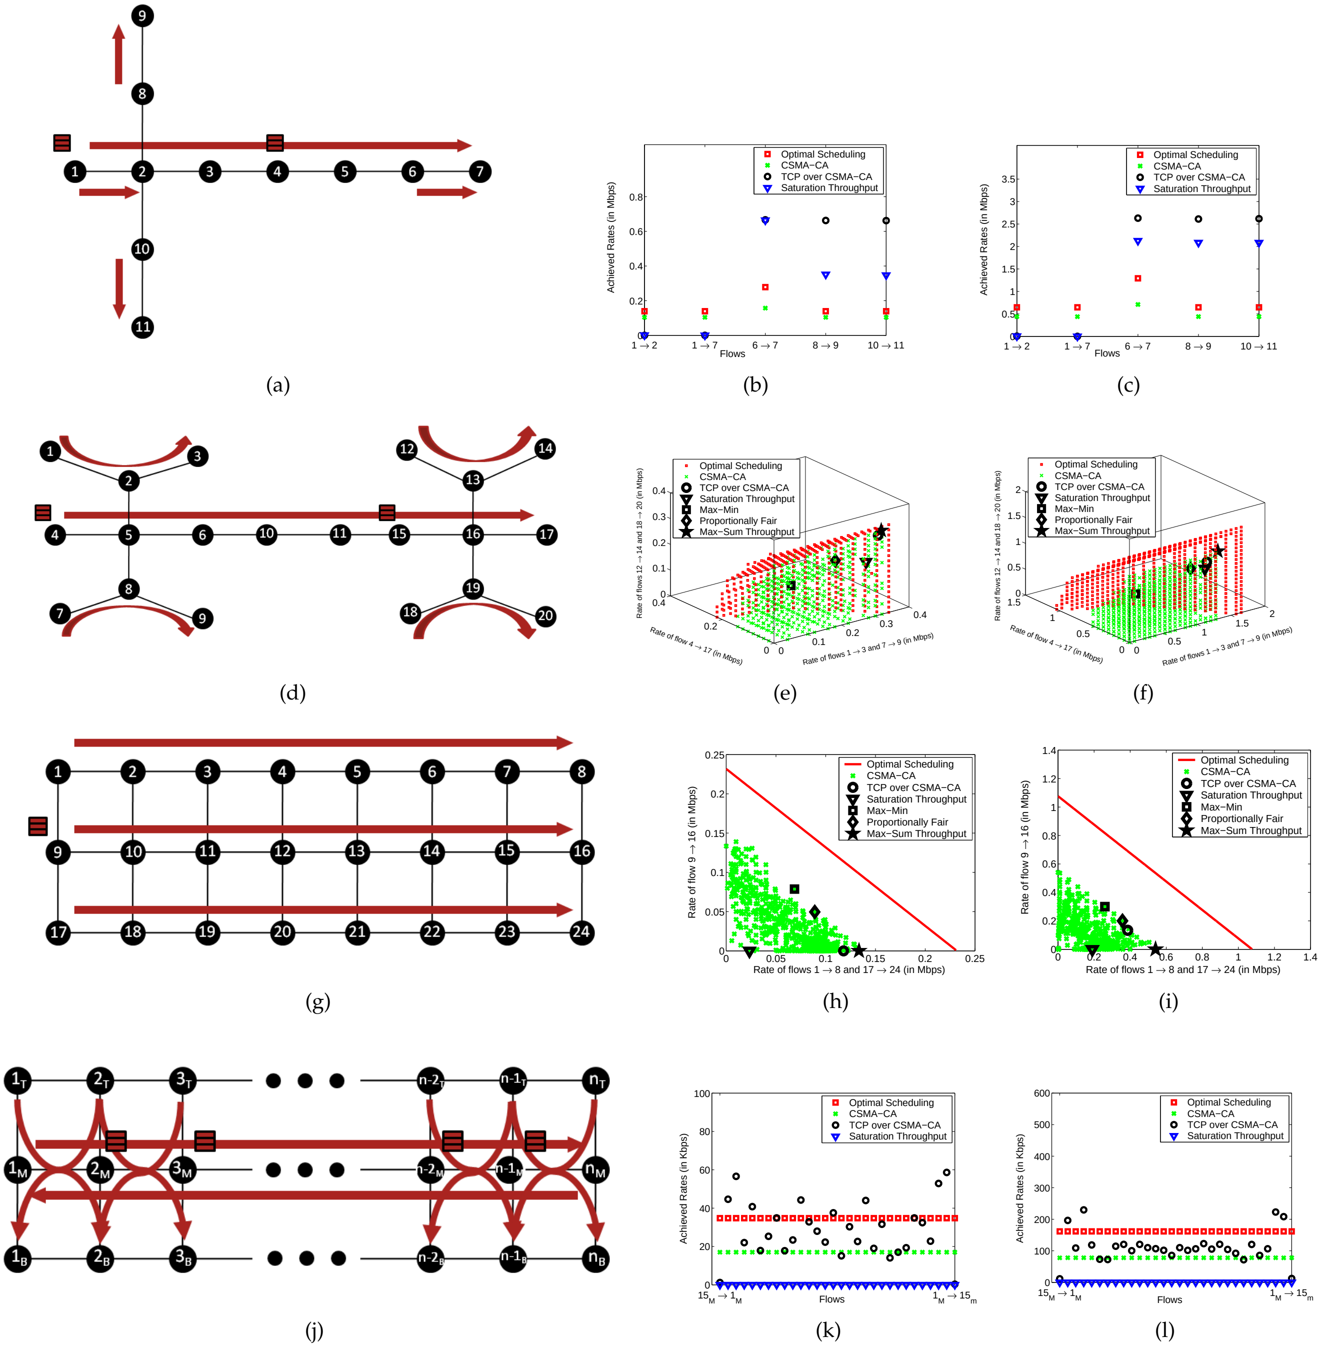

(a) (b) (c)

(d) (e) (f)

(g) (h) (i)

(j) (k) (l)

Fig. 4. Chain-Cross. (a) The topology. (b) Max-Min Fair Rate Allocation with data rate $= 1$ Mbps. (c) Max-Min Fair Rate Allocation with data rate $= 11$ Mbps. Stack-Stack. (d) The topology. (e) Achievable rate region with data rate $= 1$ Mbps. (f) Achievable rate region with data rate $= 11$ Mbps. Long Flow in the Middle. (g) The topology. (h) Achievable rate region with data rate $= 1$ Mbps. (i) Achievable rate region with data rate $= 11$ Mbps. Parking-Lot. (j) The topology. (k) Max-Min Fair Rate Allocation with data rate $= 1$ Mbps. (l) Max-Min Fair Rate Allocation with data rate $= 11$ Mbps.

rate allocation at 1 and 11 Mbps data rates respectively. We make the following observations. (i) At the max-min allocation, CSMA-CA is within 68.7% and 63.7% of the optimal respectively. Thus, with flows around node 2 ($8 \rightarrow 9$ and $10 \rightarrow 11$) not interfering with each other, the performance of CSMA-CA deteriorates (as compared to the chain topology with three interfering short flows). *Thus, we conclude that the presence of flows which interfere with a common flow but do not interfere with each other degrades CSMA-CA's achievable rate region.* (ii) Like optimal scheduling, CSMA-CA allocates a higher rate to flow $6 \rightarrow 7$ which interferes with only one flow $1 \rightarrow 7$. (iii) TCP and saturation starve flows $1 \rightarrow 2$ and $1 \rightarrow 7$.

Stack-Stack: Figure 4(d) depicts the topology. It has two stack topologies in serial, and one long flow $4 \rightarrow 17$ which is also the flow in the middle for both these stack topologies. Figures 4(e) and 4(f) compare the achievable rate regions at 1 and 11 Mbps data rates respectively. (By symmetry, the outer two flows in both the stack topologies will achieve approximately the same rate for any scheme.) We make the following observations. (i) At the max-min allocation, CSMA-CA is within 72.3% and 61.2% of the optimal respectively. (ii) Similar to the flow in the middle topologies, $4 \rightarrow 17$ is allocated a lower throughput at the proportionally fair rate allocation and is allocated a zero rate to maximize sum throughput. (iii) Both TCP and saturation starve the long flow $4 \rightarrow 17$.

We also evaluate topologies with two diamond topologies in serial as well as a stack topology followed by a diamond topology and vice versa, and all our observations qualitatively remain the same. Also, amongst all these topologies, CSMA-CA has the smallest rate region for the stack-stack topology.

Long Flow in the Middle: Figure 4(g) depicts the topology. It is similar to the flow in the middle topology depicted in Figure 2(a), except that all the flows go through 7 hops instead of 1. Figures 4(h) and 4(i) compare the achievable rate regions at 1 and 11 Mbps data rates respectively. (By symmetry, the outer two flows will achieve approximately the same rate for any scheme.) At the max-min allocation, CSMA-CA is within 63.6% and 50.6% of the optimal respectively. *Thus, the flow in the middle as well as the outer flows going over multiple hops degrades CSMA-CA's achievable rate region.* Rest of the observations remain the same as the stack-stack topology.

Parking-Lot: Figure 4(j) depicts the topology. It has two long flows, one in each direction similar to the chain topology depicted in Figure 3(a). In addition, it has a number of smaller flows, most of which do not interfere with each other, interfering with the two long flows. Finally, both long flows traverse multiple congested neighborhoods in the topology. Figures 4(k) and 4(l) compare the max-min rate allocations at 1 and 11 Mbps data rates respectively for $n = 15$. CSMA-CA is within 48.9% and 48.3% of the optimal respectively at the max-min allocation. *Not surprisingly, amongst all the topologies studied, this topology is one of the two for which CSMA-CA has the worst performance.* Also, note that TCP not only completely starves the two long flows but is also unfair to the intermediate small flows. Finally, saturation yields extremely inefficient rates for this topology.

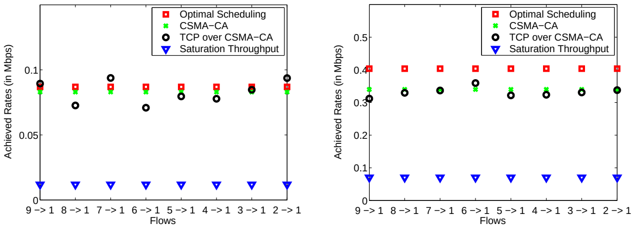

(a) (b)

Fig. 5. Clique. (a) Max-Min Fair Rate Allocation with data rate $= 1$ Mbps. (b) Max-Min Fair Rate Allocation with data rate $= 11$ Mbps.

3.5 Real Topologies

Topologies belonging to this category are the ones which we expect to commonly occur in real multi-hop wireless mesh networks. These topologies may be a result of careful placement of nodes or the result of carefully building a mesh-tree towards an access point. (However, there may be multiple interfering access points.) A study of these topologies not only allows us to comment on how CSMA-CA is expected to perform in real mesh topologies, but also yields an understanding of what kind of topologies should we avoided when building a mesh.

Clique: This topology consists of n edges and $2.n$ nodes. Every node interferes with each other, that is, the nodes form a complete graph. Its interference characteristics are similar to that of a single-hop WLAN topology. Since CSMA-CA was originally designed to support such a topology, we expect its performance to be very close to the optimal. Figures 5(a) and 5(b) compare the max-min rate allocations at 1 and 11 Mbps data rates respectively for $n = 9$. CSMA-CA is within 95.5% and 84.2% of the optimal. TCP allocates fair and efficient rates while saturation leads to inefficiencies.

Tree: Figure 6(a) depicts the topology. Each node sends a flow towards node 1 which is assumed to be the base station or the access point connected to the wired Internet. This is the most common topology one expects to construct. Figures 6(b) and 6(c) compare the max-min rate allocations at 1 and 11 Mbps data rates respectively. CSMA-CA is within 75.2% and 70.3% of the optimal respectively. Hence, CSMA-CA achieves rates close to the optimal for this real topology. TCP allocates efficient, but unfair rates while saturation yields extremely inefficient rates.

Star: Figure 6(d) depicts the topology. Each node sends a flow towards node 1 which is assumed to be the base station or the access point. And the other nodes are the users connecting to the access point. This is the single-hop topology one will build around any access

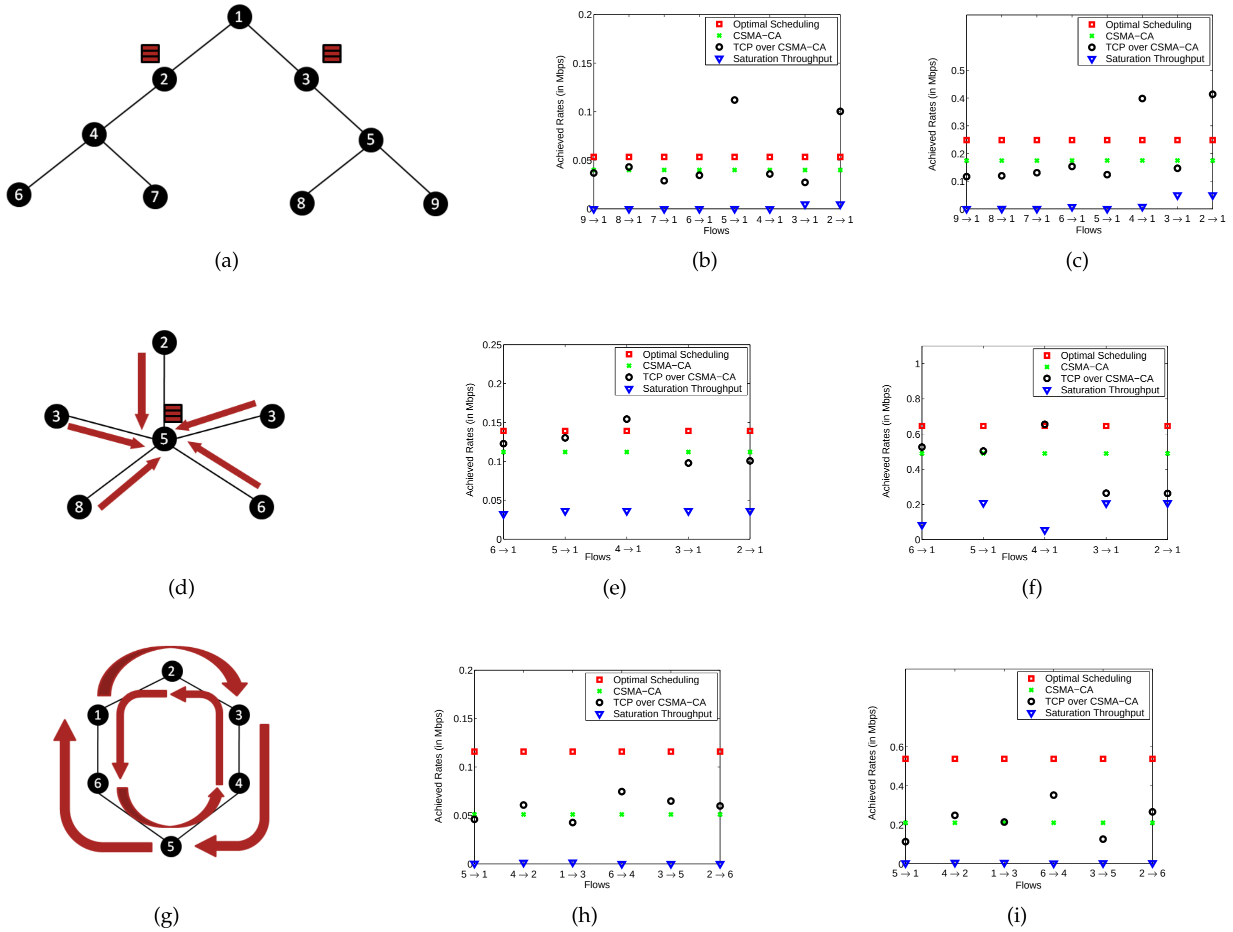

(a) (b) (c)

(d) (e) (f)

(g) (h) (i)

Fig. 6. Tree. (a) The topology. (b) Max-Min Fair Rate Allocation with data rate $= 1$ Mbps. (c) Max-Min Fair Rate Allocation with data rate $= 11$ Mbps. Star. (d) The topology. (e) Max-Min Fair Rate Allocation with data rate $= 1$ Mbps. (f) Max-Min Fair Rate Allocation with data rate $= 11$ Mbps. Ring. (g) The topology. (h) Max-Min Fair Rate Allocation with data rate $= 1$ Mbps. (i) Max-Min Fair Rate Allocation with data rate $= 11$ Mbps.

point. Note that as we studied in the clique topology, if these nodes interfere with each other, then CSMA-CA can allocate rates very close to the optimal. Hence, we construct this topology assuming that the users do not interfere with each other. Figures 6(e) and 6(f) compare the max-min rate allocations at 1 and 11 Mbps data rates respectively. CSMA-CA is within 80.5% and 75.6% of the optimal respectively. Again, CSMA-CA achieves rates close to the optimal for this real topology while TCP allocates efficient, but unfair rates and saturation yields extremely inefficient rates.

Ring: Figure 6(g) depicts the topology. The nodes are arranged as a ring, and there are 6 flows routed around the topology. Figures 6(h) and 6(i) compare the max-min rate allocations at 1 and 11 Mbps data rates respectively. CSMA-CA is within 44% and 39% of the optimal respectively. *Thus, amongst all the topologies studied, the ring topology is one of the two for which CSMA-CA has the worst performance.* Hence, a designer should avoid building a topology which has a ring around which flows are being routed.

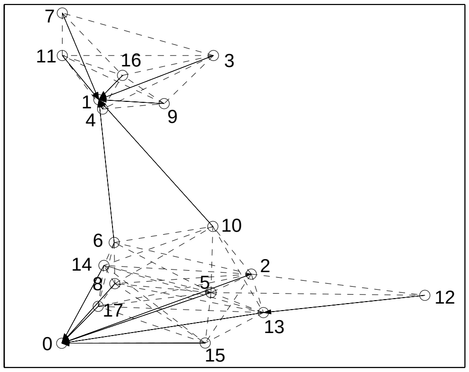

Fig. 7. Deployment at Houston: The topology.

Deployment at Houston: The final topology we study is the real topology of an outdoor residential deployment in a Houston neighborhood [6]. The node locations (shown in Figure 3.5) are derived from the deployment and fed into the Qualnet simulator. The physical channel that we use in the simulator is a two-ray path loss model with Log-normal shadowing [23]. AODV is used to set up the routes. Nodes 0 and 1 are connected to the wired world and serve as gateways for this deployment. All other nodes route their packets towards one of these nodes (whichever is closer in terms of hop count). The resulting topology as well as the routing tree is also shown in Figure 3.5. There are 16 flows in this topology. Figures 8(a) and 8(b) compare the max-min rate allocations at 1 and 11 Mbps data rates respectively. CSMA-CA is within 81.3% and 68.9% of the optimal respectively. Again, CSMA-CA achieves rates close to the optimal for this real topology while TCP and saturation starve a number of flows.

3.6 Overheads

Control Overhead of CSMA-CA: As discussed at the start of this section, at 11 Mbps data rate, more than 70% of the available throughput is consumed by control overhead (which includes MAC control packets and protocol headers). Here, we discuss which control packet can most of this loss in throughput be attributed to.

Prior works have observed that RTS/CTS control messages [10], [26] incur a significant overhead. To indirectly quantify their overhead, we determine the throughput achieved with optimal scheduling for a single-edge topology (with one transmitter and one receiver and no interfering edge) with the same overhead as IEEE 802.11 without RTS/CTS to be equal to 4.89 Mbps at 11 Mbps data rate and 0.785 Mbps at 1 Mbps data rate.

We observe that RTS/CTS control messages incur a overhead which is more significant at higher data rates. RTS and CTS are messages of 20 bytes and 14 bytes respectively. So, why does exchanging such small messages at 11 Mbps incur so much overhead. Also, there is still a throughput loss of more than 6 Mbps at 11 Mbps data rate which is still unaccounted for. Is this high overhead an artifact of CSMA-CA scheduling or a result of protocol inefficiencies in IEEE 802.11? The following observation answers this question.

An IEEE 802.11 transmitter can transmit at one of the four available basic data rates [1]. The actual data rate employed depends on the condition of the channel and can change as the channel conditions change. This is called auto-rate adaptation. The PHY layer header contains information used to determine the data rate of the incoming transmission, and hence is always transmitted at 1 Mbps [1]. And the PHY layer header is exchanged for both control (RTS, CTS and 802.11-ACK) and DATA packets. For a data rate of 11 Mbps, the transmission time of the 1024 byte DATA packet is comparable to the transmission time of the PHY layer header which is transmitted at 1 Mbps. Note that a similar overhead is incurred for the much smaller 40 byte transport layer ACK packets. Hence, the control overhead to allow for auto-rate adaptation accounts for the large overhead incurred by RTS/CTS control messages as well as the additional throughput loss at 11 Mbps data rate. Thus, the large control overhead is an artifact of the auto-rate adaptation implemented at the PHY layer in IEEE 802.11 and has nothing to do with the scheduling protocol CSMA-CA.

The Random Access Overhead: Before each packet transmission, the transmitter back-offs (waits) for a randomly selected duration. Thus, a fraction of available throughput is lost due to back-offs. Recall that we refer to this loss in throughput as the random access overhead. Now, lower the packet transmission time, more is the random access overhead as the expected value of this random duration is a constant irrespective of the data rate.

All topologies we study achieves a lower throughput ratio at 11 Mbps data rate than at 1 Mbps data rate. The reason is the larger random access overhead at 11 Mbps. Choosing a smaller value of W_0 when data rate is equal to 11 Mbps will reduce this random access overhead. For example, choosing $W_0 = 8$ and retaining the default values for the rest of the IEEE 802.11 parameters compensates the extra loss in throughput at 11 Mbps data rate. The flow in the middle, chain and stack-stack (a representative topology for each category) achieve throughputs within 70.6%, 58.7% and 74.3% of the optimal respectively; thus the performance at 11 Mbps becomes equivalent to the performance at 1 Mbps. Note that the observation of choosing a smaller W_0 for higher data rate/smaller packet sizes is not new and prior works have made similar observations too [17].

To summarize, the extra throughput loss with 11 Mbps data rate can be compensated by using a smaller value of W_0. Hence, the throughput ratio comparison at 1 Mbps yields a better and more fair evaluation of the performance of CSMA-CA for any topology with default 802.11 parameters. So, our summary of this study will be constructed based on the throughput ratios evaluated at 1 Mbps.

3.7 Summary

In this section, we summarize the intuitions derived regarding which topology characteristics deteriorates or improves CSMA-CA's performance with optimal rate control. (i) Presence of asymmetric edges in a topology leads to a throughput loss larger than the loss caused by any other two-edge topology. (ii) Even with a flow in the middle, CSMA-CA allocates fair and efficient rates with optimal rate control. (iii) Presence of flows which interfere with a common flow but do not interfere with each other deteriorates CSMA-CA's performance. This is a counter-intuitive observation as one may expect that increasing interference in a topology should degrade CSMA-CA's throughput performance. But, recall that our metric is the throughput ratio of CSMA-CA's and optimal scheduling's throughput performance. Non-interfering flows which interfere with a common flow yield better absolute throughputs, but a worse throughput ratio as compared to a topology with interfering flows. (iv) For a chain topology with more than 4 hops, changing the number of hops does not change CSMA-CA's throughput performance. However, in the presence of short flows interfering with the long flows, reducing

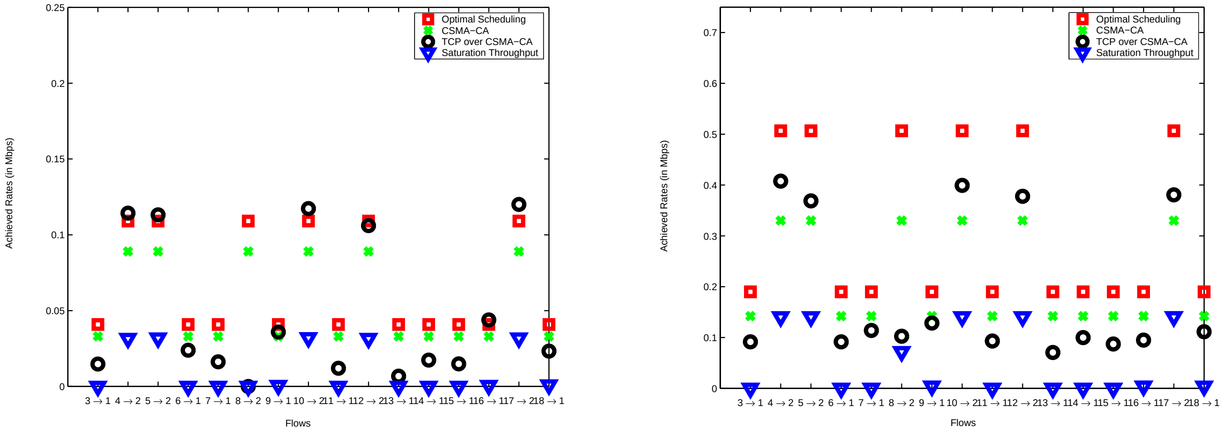

(a) (b)

Fig. 8. Deployment at Houston. (a) Max-Min Fair Rate Allocation with data rate $= 1$ Mbps. (b) Max-Min Fair Rate Allocation with data rate $= 11$ Mbps.

the number of hops improves CSMA-CA's performance. (v) More the number of congested neighborhoods the long flow in a chain topology passes through, the smaller is CSMA-CA's achievable rate region. (vi) Amongst all the topologies we study, the parking lot and the ring topology have the worst performance for CSMA-CA.

4 CONCLUSIONS

To conclude, we summarize the important observations made during this study.

(i) CSMA-CA is always within 40% of the optimal for all topologies we study. For most topologies, the throughput loss is less than 30%, that is, CSMA-CA is within 70% of the optimal. Topologies which combine flow in the middle and chain, like the parking lot topology, and the ring topology see the largest drop in throughput with CSMA-CA. Hence, these two topologies make ideal candidates to evaluate any new scheduling protocol. *For topologies which are expected to be constructed in practice, like the tree topology, the star topology and the topology derived from the deployment at Houston, CSMA-CA yields fair and efficient throughputs (within 75% of the optimal) with optimal rate control.*

(ii) Both TCP and saturation yield extremely unfair and inefficient rates over CSMA-CA in multi-hop networks. Hence, comparing the performance of any new rate control or scheduling protocol with TCP or saturation over CSMA-CA can be misleading. For example, a new rate control protocol may be doing much better than TCP/saturation over CSMA-CA, but still maybe orders of magnitude away from the best (both in terms of fairness and efficiency) throughput we can achieve with CSMA-CA. Similarly, a new scheduling protocol may achieve a better throughput with TCP/saturation as compared to CSMA-CA, but its achievable rate region can be much smaller than CSMA-CA. Hence, for a fair comparison, the performance of any new rate control and scheduling protocol should be compared to the performance of CSMA-CA with optimal rate control using the methodology proposed in this work.

(iii) Topologies which combine the flow in the middle and the chain topology (Section 3.4) display both starvation and inefficient throughputs with TCP/saturation over CSMA-CA. Hence, they make perfect candidate topologies to evaluate any new rate control protocol.

REFERENCES

[1] Part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications - higher-speed physical layer extension in the 2.4 ghz band. *IEEE Std 802.11b-1999*, nov 2002.
[2] Qualnet, network simulator, 2007.
[3] D. P. Bertsekas and R. G. Gallagher. *Data Networks*. Prentice Hall, 2 edition.
[4] G. Bianchi. Performance analysis of the ieee 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18:535–547, Mar. 2000.
[5] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *Proceedings of ACM MOBICOM*, 2005.
[6] J. Camp, J. Robinson, C. Steger, and E. Knightly. Measurement driven deployment of a two-tier urban mesh access network. In *Proceedings of ACM MOBISYS*, 2006.
[7] P. Chaporkar, K. Kar, X. Luo, and S. Sarkar. Throughput and Fairness Guarantees Through Maximal Scheduling in Wireless Networks. *IEEE Transactions on Information Theory*, 2008.
[8] D. Clark. Window and acknowledgement strategy in TCP. *RFC 813 Internet Engineering Task Force*, 1982.
[9] M. Durvy, O. Dousse, and P. Thiran. Border effects, fairness, and phase transition in large wireless networks. In *Proceedings of IEEE INFOCOM*, 2008.
[10] J. Eriksson, S. Agarwal, P. Bahl, and J. Padhye. Feasibility study of mesh networks for all-wireless offices. In *Proceedings of ACM MOBISYS*, 2006.
[11] M. Garetto, T. Salonidis, and E. Knightly. Modeling per-flow throughput and capturing starvation in csma multi-hop wireless networks. In *Proceedings of IEEE INFOCOM*, 2006.
[12] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. In *Proceedings of ACM MOBICOM*, 2003.
[13] A. Jindal and K. Psounis. Achievable rate region and optimality of multi-hop wireless 802.11-scheduled networks. In *Proceedings of ITA*, 2008.
[14] A. Jindal and K. Psounis. The Achievable Rate Region of 802.11-Scheduled Multihop Networks. *IEEE/ACM Transactions on Networking*, 17(4):1118–1131, 2009.

[15] C. Joo, X. Lin, and N. Shroff. Understanding the capacity region of the greedy maximal scheduling algorithm in multi-hop wireless networks. In *Proceedings of IEEE INFOCOM*, 2008.

[16] J. Li, C. Blake, D. D. Couto, H. Lee, and R. Morris. Capacity of ad-hoc wireless networks. In *Proceedings of ACM MOBICOM*, 2001.

[17] K. Medepalli and F. A. Tobagi. Towards performance modeling of ieee 802.11 based wireless networks: A unified framework and its applications. In *Proceedings of IEEE INFOCOM*, 2006.

[18] E. Modiano, D. Shah, and G. Zussman. Maximizing throughput in wireless networks via gossip. In *Proc. ACM SIGMETRICS*, 2006.

[19] J. Nagle. Congestion control in IP/TCP internetworks. *RFC 896 Internet Engineering Task Force*, 1984.

[20] K. Nahm, A. Helmy, and C. J. Kuo. TCP over multihop 802.11 networks: issues and performance enhancements. In *Proceedings of ACM MOBIHOC*, 2005.

[21] C. E. Perkins and E. M. Royer. Ad hoc on-demand distance vector routing. In *Proceedings of WMCSA*, 1999.

[22] S. Rangwala, A. Jindal, K. Jang, K. Psounis, and R. Govindan. Understanding congestion control in multi-hop wireless mesh networks. In *Proceedings of ACM MOBICOM*, 2008.

[23] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, 2 edition, 1996.

[24] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh. Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols. In *Proceedings of IEEE WCNC*, 2005.

[25] S. Sanghavi, L. Bui, and R. Srikant. Distributed link scheduling with constant overhead. In *Proc. ACM SIGMETRICS*, 2007.

[26] Y. Sun, I. Sheriff, E. Belding-Royer, and K. Almeroth. An experimental study of multimedia traffic performance in mesh networks. In *Proceedings of WitMeMo*, 2005.

[27] K. Tan, F. Jiang, Q. Zhang, and X. Shen. Congestion Control in Multihop Wireless Networks. *IEEE Transactions on Vehicular Technology*, 2006.

[28] X. Wang and K. Kar. Throughput modeling and fairness issues in csma/ca based ad-hoc networks. In *Proceedings of IEEE INFOCOM*, 2005.

[29] K. Xu, M. Gerla, L. Qi, and Y. Shu. TCP unfairness in ad-hoc wireless networks and neighborhood RED solution. In *Proceedings of ACM MOBICOM*, 2003.