Modelling throughput and starvation in 802.11 wireless networks with multiple flows

A. Margolis, R. Vijayakumar, Sumit Roy
Dept. of Electrical Engineering
University of Washington, Box 352500
Seattle, WA 98195-2500
{amargoli, rajivv, sroy}@u.washington.edu

***Abstract*—Carrier sense multiple access (CSMA) protocols require stations that wish to transmit to first check the shared medium for ongoing transmissions. However, in wireless networks stations may not be able to sense transmissions from all other stations; hence those at different locations can experience different amounts of contention, resulting in unfairness and possible "starvation" (inability to acquire channel access for long periods). In this paper we model how the 802.11 MAC protocol allocates bandwidth among multiple saturated flows in a linear topology when not all senders are within sensing range. We consider a specific class of topologies consisting of n independent flows, where each sender can sense k neighbors on either side. Our work uncovers global interactions among flows leading to startling sensitivities in node throughput. A new model to predict the long-term throughput of each flow under saturation is presented and our model results validated via OPNET simulations.**

I. Introduction

The IEEE 802.11 MAC protocol uses physical carrier sensing (PCS) and optional RTS/CTS handshaking or virtual carrier sensing (VCS) to prevent wireless nodes from transmitting simultaneously. If all senders in a network are within carrier sense range—or if they all communicate with a single access point and RTS/CTS is used—they should share the channel fairly in the long run. In this paper, we consider a class of topologies in which the 802.11 MAC allocates bandwidth unfairly among the flows. This long-term unfairness is a result of the unsynchronized CSMA nature of the 802.11 MAC, and it has been pointed out before in the literature. In the most basic case, called (here and elsewhere) the "three flows problem," a sender is within sensing range of two other senders that can't sense each other (one on either side), and it is prevented from accessing the channel for long periods of time due to its neighbors' overlapping transmissions. The scenarios that we consider here are extensions of this case, consisting of n senders, arranged linearly such that each can sense at most k senders on either side. We present a model, based on the "contention graph" of the network, to predict throughput for each flow under saturation. While the $n = 3$, $k = 1$ case (the three-flows problem) has been modeled before, our model provides simple yet accurate approximation for any n, k. We compare our model's predictions to OPNET simulation results. We additionally observe, based on both model and simulation results, that the allocation of bandwidth among flows can change dramatically if a single flow is added or removed from the network.

II. The Three-Flows Problem

In Figure 1, the three sender-receiver pairs support separate, saturated flows, i.e., we assume each sender has an infinite queue of MAC-layer waiting packets. The distances between the senders are such that node 2 can sense the transmissions of both 1 and 3, but 1 and 3 are out of CS range of each other. Therefore sources 1 and 3 can be active simultaneously and we assume that the signal-to-interference ratio at their destinations is high enough to allow both transmissions to be successful.

The 802.11 MAC protocol (DCF) uses a random binary exponential backoff procedure to resolve contention: a sender must wait a "DCF Interframe Space" (DIFS) and then a random number of "slot times" before it can send each packet (see [1]). If the count-down is interrupted by a transmission, the sender freezes the count-down timer until the next transmission. The DIFS, slot time (σ), and minimum contention windows (CW_{min}) are fixed values that differ among the different PHY-layer standards.

In this scenario, flow 2 will get a much lower throughput than the other two flows, regardless of the specific parameter values. We consider each sender's view of the channel as being composed of *busy periods*, *contention periods* and transmissions; during busy periods there are ongoing transmissions from neighbors and the sender must wait, and during contention periods the sender is waiting out the DIFS or decrementing its timer. During contention periods seen by sender 2, it is at a disadvantage due to the fact that it has to contend with both senders 1 and 3, whereas each of those senders is contending with only one other sender (2). Furthermore, busy periods at sender 2 tend to be longer than busy periods at the other two senders because of overlapping transmissions by 1 and 3, as illustrated in Figure 2. As is discussed elsewhere, the problem is in essence caused by the fact that senders 1 and 3 are not synchronized in their transmissions. Sender 2 can fail to see the channel idle for the necessary minimum time (DIFS), and thus not get to decrement its timer, for many consecutive transmissions of its neighbors. If neighbors are not within "communication range" an additional problem may occur: the protocol standard mandates that if a sender receives

1930-529X/07/$25.00 © 2007 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2007 proceedings.

a neighbor's packet with errors, it must wait an "Extended Interframe Space" (EIFS) after the channel is idle, rather than the usual DIFS. Our simulation results ([2]), as well as results in [3] show that a longer packet transmission time (due to low data rate and/or large packet sizes), longer DIFS, or use of the EIFS (due to being out of communication range of neighbors) all tend to increase the length of time and number of consecutive transmissions of neighbors composing the busy periods seen by sender 2.

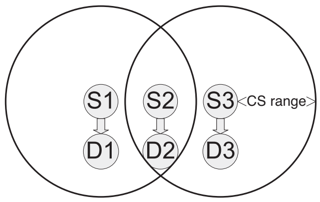

Fig. 1. The "three-flows" problem.

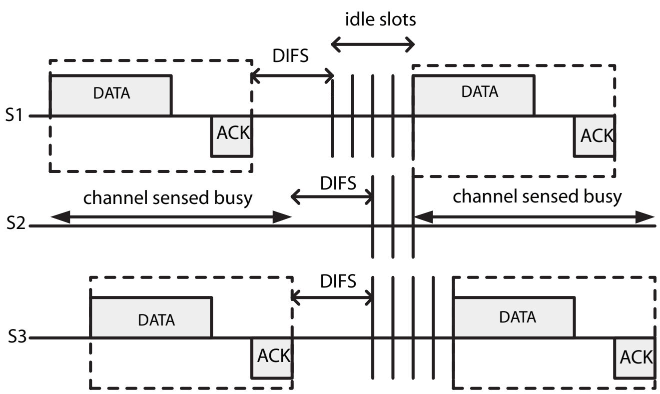

Fig. 2. Sender 2 has to wait until both senders 1 and 3 have been idle for at least a DIFS before it can decrement its backoff counter. In the specification of the 802.11 MAC, the time between a packet arrival at its receiver and transmission of its ACK is less than a DIFS, so sender 2 will not get to decrement its timer during this interval.

III. Related Work

The three-flows problem and variations have been considered by other researchers under several names, e.g., the"exposed terminal problem"([4],[5]). The issue of asynchronization as a cause of unfairness under 802.11 has been noted specifically in [6] and [7]. In [7], the author describes the problem in greater detail with the idea of "neighborhood capture," in which a set of senders can essentially seize the channel for long periods of time. The author notes that in chain network topologies where the CS range is smaller than the length of the chain, nodes in the middle are disadvantaged for two reasons: they have more contenders than the end nodes, and some of their contenders can transmit simultaneously. The lack of synchronization between the contenders who can't hear each other implies that in certain topologies these middle nodes get almost no throughput. The author discusses the implications of this problem in different topologies, noting, for example, that in a ring of four senders in which two opposite senders can transmit simultaneously but adjacent senders cannot, once one pair of senders are active, the other two can be squeezed out for relatively long periods of time. Thus the lack of synchronization can cause periods of starvation for one or the other pairs of flows, even though the topology is not biased to favor any particular set of flows.

Our work, inspired by [7], advances the state-of-art by presenting a model for a specific set of topologies to estimate the channel share of each flow with any number of senders. A handful of other researchers have presented models for similar scenarios. In [8] a detailed analysis of the three-flows problem is presented, which aims to predict throughput of each flow under saturation. Their model uses a Markov chain that keeps track of backoff timers of all (three) senders, an approach that is intractable as the number of sender-receiver pairs scales. Boorstyn et al ([9]) presented a general method for modelling CSMA networks based on a continuous-time Markov chain whose state description consists of the set of active nodes; this approach has been adopted by others, such as [10]. The significant advantage of the Markov description in [9] and [10] — which we adopt — is the effectiveness of their state representation, which only grows at a modest rate with network size. Our work applies the state representation in [9] to the class of linear 1-hop topologies, but our approach differs from theirs in several ways. Their model uses a Poisson packet arrival rate and exponential transmission rate, whereas we use a saturated traffic model. Their approach uses queueing theory to calculate the probability of being in each state, and uses an iterative method to compute the maximum possible throughput for each node. However, our approach is based mainly on combinatorial techniques to analyze a model for the n linear flow case; we develop this model based upon the observation that the middle sender in the three-flows problem is essentially starved. In [10], the starvation of the middle node is noted for a limiting case, but no further analysis is presented for extended topologies.

Garetto el al present a modelling approach in [11], [12], that in theory allows for the calculation of throughput for any node in any 802.11 multiflow network where not all nodes can sense each other. However, obtaining results in our model is significantly simpler as it is based on observations about node interactions in linear arrangements of 1-hop flows. Our model also makes apparent the effect of adding or removing a flow in these topologies, which is not as immediately apparent from their model.

Graph models have often been used to incorporate constraints on nodes or links that can transmit simultaneously, as we do here. For instance, [13] introduces a "flow contention graph," in which each flow is represented by a vertex and an edge is placed between two vertices if the corresponding nodes cannot be active at the same time (due to interference or contention in the CSMA algorithm). Only one flow in each clique [1] can be active at a time, and further, it must be the

[1]In graph theory, a *clique* is set of vertices for which there are edges between every pair in the set.

1930-529X/07/$25.00 © 2007 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2007 proceedings.

only one transmitting in all of the cliques it belongs to. In [14], the problem of finding the maximum possible number of simultaneous transmissions in a network graph is shown to be equivalent to the problem of finding the number of edges in a graph that aren't connected by another edge, a problem which we solve here for the particular class of graphs under consideration. In [15], a conflict graph (based on links rather than flows) is presented; they define a "usage vector" that gives the fraction of time that each link in the network can be active, noting that the usage for each link can be computed by summing the fraction of time given to each of the independent sets that the link belongs to in the conflict graph. They note that finding the optimal throughput of a network is equivalent to finding the size of the largest independent set (called the independence number) of a graph.

Our approach shares some of the techniques in [13], [14], and especially in [15] in its use of the link conflict graph and calculation of throughput by summing over fraction of time allocated to maximal independent sets. However, our method for calculating the fraction of time allocated to each maximal independent set is based on our observations about the three-flows problem. While similar observations have been noted before (see for example, [10]), we believe we are the first to use these observations to derive actual throughput results and validate them via credible simulations.

IV. Throughput Estimation in Linear Topologies

We consider a general extension of the three flow case: *there are n independent saturated flows and each sender is within sensing range of k flows on either side*. Our approach makes use of the *conflict graph* for the network, as discussed in [13], [15] and elsewhere. Two nodes can transmit simultaneously (without collision or interference) if and only if they are not connected in the conflict graph.[2] Two examples ($n = 5, k = 1$ and $n = 5, k = 2$) and their corresponding conflict graphs are shown in Figure 3. A set of nodes that can transmit simultaneously forms an "independent set" ([9], [12], [10], [15]) in the conflict graph; an independent set which can't have any more nodes added to it is *maximal.* In Figure 3a, the maximal independent sets are (1,3,5), (1,4), (2,5), and (2,4). In Figure 3b, the maximal independent sets are (1,4), (1,5), (2,5) and (3). In the three-flows problem ($n = 3$, $k = 1$), the maximal independent sets are (1,3) and (2).

The basis for our model is the description of the system at any time by a *state* of active nodes which form a maximal independent set. To make this precise, we say that a node is "active" starting from the point at which it begins a successful transmission until another sender within sensing range begins a successful transmission; it is also considered active when all its neighbors within sensing range are being silenced by their neighbors. An illustration is as follows: in the three-flows case, if sender 2 is currently transmitting, the system is considered to be in state (2) until either sender 1 or 3 begins transmitting, at which point it is considered to be in state (1,3). This makes sense because, although senders 1 and 3 do not begin transmitting exactly at the same time, it is guaranteed that when one of them takes the channel, the other only has to wait until its countdown timer expires to begin transmitting successfully.

[2]In the specific cases that we consider, in which there is one flow and one sender on each link, we use the terms "node", "link", "sender", and "flow" interchangeably. Each corresponds to one vertex in the conflict graph.

As discussed above, the idea of representing the system as a sequence of states of active or transmitting nodes, where each possible state is an independent set in the conflict graph, is not a new concept. The analysis in [9], which has been adapted in [12] and [10], uses a Markov chain over sets of active nodes. Also, the method of summing over fraction of time given to independent sets in a link conflict graph to calculate throughput was used in [15].

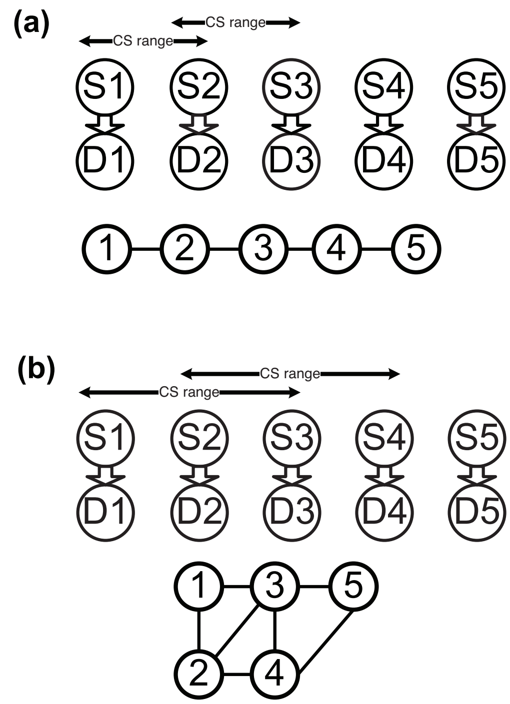

Fig. 3. Examples of the considered topologies and their conflict graphs: (a) for the $n = 5$, $k = 1$ topology, and (b) for the $n = 5$, $k = 2$ topology.

The key approximation of our model is that the system spends *all* of its time in the largest-size states and that all states with lower cardinality get zero time. This is more accurate when the protocol parameters are such that a "middle sender" will tend to see longer busy periods due to the simultaneous, non-synchronized transmissions of its two neighbors. This behavior is noted in [10] and illustrated with three-flows problem, where, using exponential (μ) transmission rates and Poisson arrivals (λ), they show that as $\frac{\lambda}{\mu} \rightarrow \infty$, sender 2 (the node not in a largest-size independent set) gets throughput approaching 0.

In our extended topologies, there are many "middle" nodes which can sense two senders who can't sense each other. Our approximation is based on the following argument: for the system to transition from a largest-size state to a smaller one would require two active nodes who can't hear each other to be idle simultaneously for a sufficient length of time for a "middle" node to decrement its timer to zero. On the other hand, it is not uncommon for a sender to become active if it has only one active neighbor. So for instance, in the $n = 5$, $k = 1$ topology, the transition from state $(1, 3, 5)$ to state $(2, 5)$ would be a low-probability event—it requires that sender 2 win the channel from its two neighbors who can't hear each other. However, once in state (2,5) it is easy to transition back to state (1,3,5), since this requires only that either sender 1 or 3 wins the channel back from sender 2. Therefore, we can expect the time spent in state (2,5), and more generally, in any smaller-size state, to be small. Therefore we expect the throughput to be very low for any node which is not part of one of the largest states.[3]

We also make the approximation that the system spends *equal* amounts of time in each of the largest-size states. A detailed argument, which explicitly models the sequence of states as a Markov chain, is presented in [2]. The argument is based on the idea that sender i is equally likely to yield the channel to any sender that is within sensing range and that does not sense any other active senders. For instance, in the $n = 5$, $k = 2$ topology, state (1,5) is equally likely to transition into states (1,4) and (2,5). In the $n = 6$, $k = 2$ topology, state (1,6) is equally likely to transition into states (1,5), (1,4), (2,6), or (3,6).

Estimation of Node Throughput: Based on the above, we derive the following results. Let $L = \lfloor \frac{n-1}{k+1} \rfloor + 1$ be the size of the largest maximal independent set. Clearly the number of such largest sets is the number of ways to choose L items from a list of n such that the chosen items are all at least $k+1$ places apart. We denote this number $b_{n,k}$. In the case $k = 1$ and n even, there are $\frac{n}{2}+1$ (=$L+1$) sets. In the general case, we derived an expression for $b_{n,k}$ by first considering the case where each of the L items is selected as close as possible to the top of the list, i.e., the first item is at position 1, the second at position $k+2$, etc. This leaves $U = n - ((L-1)(k+1)+1)$ empty positions at the bottom of the list. Thus there are $U + 1$ possible positions for the last item; considering then the possible positions for the next-to-last item leads to the recursive expression $b_{n,k} = \sum_{i=0}^{U} b_{((L-2)(k+1)+1+i),k}$ (see [2]). In the case that $k + 1$ divides $n - 1$, then $U = 0$, so there is only one largest state, and the nodes in this state are active almost all of the time. In the case where $k+1$ does not divide $n-1$, there are more than one largest state. For example, if $n = 8$ and $k = 2$, the size $L = 3$ sets are $(1, 4, 7)$, $(2, 5, 8)$, $(1, 4, 8)$, and $(1, 5, 8)$. Since we made the argument that all size-L states occur an equal fraction of the time, the long-run fraction of time that a node i gets the channel is given by the fraction of such states (size L maximal independent sets) that i belongs to: $\frac{N_i}{b_{n,k}}$ where N_i is the total number of states that i belongs to. Now N_i is the total number of ways to select a total of $L - 1$ items from above and below i on the list at least $k + 1$ places apart, if this is possible; therefore $N_i = \left(b_{i-(k+1),k}\right) \cdot \left(b_{n-(i+k),k}\right)$ if $b_{i-(k+1),k} + b_{n-(i+k),k} = L-1$, or $N_i = 0$ if $b_{i-(k+1),k} + b_{n-(i+k),k} < L-1$ (in the case that it is not possible to select $L-1$ items after selecting i.)

[3]Note that in some topologies, such as $n = 4, k = 1$, there are no smaller-size states; all maximal independent sets are the maximum size.

Table I lists predicted and actual (simulated) results for a few networks, in terms of *normalized throughput*. This represents the fraction of the theoretical maximum throughput of a single flow transmitting continuously. We computed the normalized throughput from OPNET simulation results as

$$normalized\ throughput = throughput \cdot \frac{T + \frac{CW_{min}-1}{2}\sigma}{packet\ size} \tag{1}$$

where $throughput$ is measured in bps, $packet\ size$ is measured in bits, and

$$\begin{aligned} T =& mean\ data\ packet\ trans.\ time \\ &+ SIFS + ACK\ trans.\ time + DIFS. \end{aligned} \tag{2}$$

The simulations used packet sizes of 1500 bytes and the 802.11a PHY settings. The model does not take data rate into account, but for comparison, simulation results are given for both a low data rate (6Mbps) and a high data rate (12Mbps). Due to the expected (and observed) symmetry, only results for flows $\lceil \frac{n}{2} \rceil$ are given. In the simulations, the EIFS is triggered by neighboring sender's transmissions. The simulation throughputs are averaged over one 2-minute run.

TABLE I
COMPARISON OF SIMULATION AND MODEL PREDICTIONS FOR A FEW EXAMPLE TOPOLOGIES. SIMULATION RESULTS ARE FROM OPNET, NORMALIZED AS EXPLAINED ABOVE.

	$n=3, k=1$		$n=5, k=1$		
	flow 1	flow 2	flow 1	flow 2	flow 3
model	1	0	1	0	1
sim, 6Mbps	0.98	0.01	0.97	0.01	0.97
sim, 12Mbps	0.96	0.02	0.96	0.03	0.94
	$n=4, k=1$		$n=6, k=1$		
	flow 1	flow 2	flow 1	flow 2	flow 3
model	0.6667	0.3333	0.75	0.25	0.50
sim, 6Mbps	0.71	0.27	0.79	0.20	0.50
sim, 12Mbps	0.69	0.30	0.76	0.23	0.49
	$n=4, k=2$		$n=5, k=2$		
	flow 1	flow 2	flow 1	flow 2	flow 3
model	1	0	0.6667	0.3333	0
sim, 6Mbps	0.97	0.00	0.65	0.30	0.00
sim, 12Mbps	0.94	0.02	0.66	0.30	0.01

An important observation that we make, and that is implied by our model, is that the fairness of channel allocation among the nodes can change dramatically if a single flow is added or removed from the end of the network. Figure 4 illustrates this effect. The figure shows the results of a simulation in OPNET consisting of 10 sender-receiver pairs that can sense $k = 1$ pair on either side. Initially, all but the first is silent; every 10 seconds a new saturated flow is switched on. The throughput at the receiver for each flow is color-coded by flow number (due to symmetry, only flows 1-5 are plotted), and the flat plots in the same color are the model predictions. The predicted throughput in bps was computed from the predicted normalized throughput using Eq. (1).

1930-529X/07/$25.00 © 2007 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2007 proceedings.

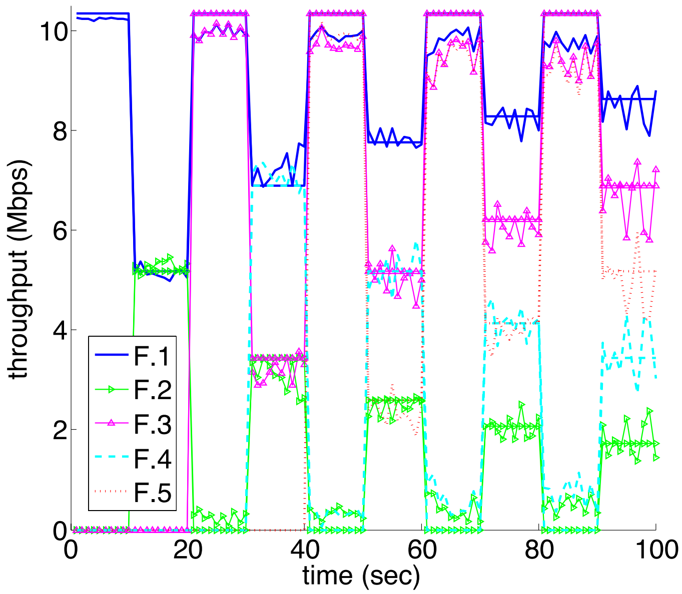

Fig. 4. Predicted and simulated throughput at receiver for 10 sender-receiver pairs, where each pair can sense $k = 1$ pair on either side. Initially only flow 1 is active, and every 10 seconds the next flow is switched on. (Due to symmetry, only flows 1-5 are shown). The "flat" plots are the predicted long-term throughput for each flow.

The results in Table I and Figure 4, as well as further results in [2] for 802.11b, showed that the model is less accurate when either:

(i) the data rate is high and the conflict graph has some smaller-size maximal independent sets;

(ii) there are many "nodes in the middle" (i.e., k is large), particularly when the data rate is high;

(iii) $\frac{n-1}{k+1}$ is large, so that both largest and smaller-size maximal independent sets contain many nodes

The explanation for (i) is that if the packet transmission time is short, it is less rare for a middle sender to win the channel from its two neighbors. Thus the system transitions more frequently into (and spends more time in) smaller-size states. This effect can be seen in Table I for the $n = 3, k = 1$ topology, the $n = 4, k = 2$ topology, and the $n = 5, k = 1$ topology, for which the model is less accurate at the higher data rate. Note the the same effect occurs with shorter packet sizes, smaller DIFS, or when the EIFS isn't used (due to neighbors being within communication range). The explanation for (ii) is that we would expect higher cumulative throughput for the "starved" nodes in the middle if there are more of them—both the $n = 4$, $k = 2$ topology and the $n = 3$, $k = 1$ topology have one largest state, but the former will spend more cumulative time in its smaller states (2 and 3) than the latter will in its smaller state (2). The explanation for (iii) is that if a "starved" node belongs to a smaller-size maximal independent set that contains many other nodes, its chances of becoming active are much higher since its neighbors can be silenced by one of the other nodes in its set. For example, the $n = 9, k = 1$ topology has only one largest state (containing all of the odd nodes) but when it does enter the smaller state containing the even nodes, it is likely to remain there for some time. This is evident in Figure 4: when $n = 9$ flows are active, the "starved" flows (2,4,6,8) get higher thoughput than the "starved" flow 2 gets when only $n = 3$ flows are active.

V. Conclusion

In this paper, we considered chain-like networks of multiple saturated 1-hop flows. We have argued that the bandwidth allocated to each flow in such a topology is well approximated, under certain conditions, by the fraction of largest-size maximal independent sets to which it belongs. An implication of our analysis is that certain combinations of n and k lead to very unfair behavior, while others are more fair. Adding a single sender/receiver pair to the end of the topology can dramatically change the way the channel is allocated among all the flows; hence there is no "limiting behavior" as $n \to \infty$.

References

[1] "IEEE 802.11-1999," 2003. [Online]. Available: http://standards.ieee.org/getieee802/802.11.html

[2] A. Margolis, "Modelling throughput and starvation in 802.11 wireless networks with multiple flows," Master's thesis, Department of Electrical Engineering, University of Washington, Seattle, WA, Aug 2006.

[3] C. Chaudet, D. Dhoutaut, and I. Lassous, "Experiments of some performance issues with IEEE 802.11b in ad hoc networks," in *WONS 2005*, Jan. 2005, pp. 158–163.

[4] J. Stine, G. de Veciana, K. Grace, and R. Durst, "Orchestrating spatial reuse in wireless ad hoc networks using synchronous collision resolution (SCR)," *World Scientific Journal of Interconnection Networks*, vol. 3, no. 3–4, pp. 167–195, 2003.

[5] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang, "MACAW: a media access protocol for wireless lan's," in *SIGCOMM '94*. New York, NY, USA: ACM Press, 1994, pp. 212–225.

[6] J. He and H. K. Pung, "Fairness of medium access control protocols for multi-hop ad hoc wireless networks," *Computer Networks*, vol. 48, pp. 867–890, Aug. 2005.

[7] M. Benveniste, "Wireless LANs and 'neighborhood capture'," in *The 13th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, 2002*, vol. 5, 2002, pp. 2148–2154.

[8] C. Chaudet, I. G. Lassous, E. Thierry, and B. Gaujal, "Study of the impact of asymmetry and carrier sense mechanism in IEEE 802.11 multi-hops networks through a basic case," in *PE-WASUN '04*. New York, NY, USA: ACM Press, 2004, pp. 1–7.

[9] R. R. Boorstyn, A. Kershenbaum, B. Maglaris, and V. Sahin, "Throughput analysis in multihop CSMA wireless packet radio networks," *IEEE Trans. on Communications*, vol. COM-35, no. 3, pp. 267–274, Mar 1987.

[10] X. Wang and K. Kar, "Throughput modelling and fairness issues in CSMA/CA based ad-hoc networks," in *Proceedings of IEEE INFOCOM*, 2005.

[11] M. Garetto, J. Shi, and E. W. Knightly, "Modeling media access in embedded two-flow topologies of multi-hop wireless networks," in *MobiCom '05*. New York, NY: ACM Press, 2005, pp. 200–214.

[12] M. Garetto, T. Salonidas, and E. Knightly, "Modeling per-flow throughput and capturing starvation in CSMA multi-hop wireless networks," in *IEEE INFOCOM 2006, Barcelona, Spain*, Apr. 2006.

[13] T. Nandagopal, T.-E. Kim, X. Gao, and V. Bharghavan, "Achieving MAC layer fairness in wireless packet networks," in *Mobile Computing and Networking*, 2000, pp. 87–98.

[14] H. Balakrishnan, C. L. Barrett, V. S. A. Kumar, M. V. Marathe, and S. Thite, "The distance-2 matching problem and its relationship to the MAC-layer capacity of ad hoc wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 22, pp. 1069–1079, Aug. 2004.

[15] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu, "Impact of interference on multi-hop wireless network performance," *Wireless Networks*, vol. 11, pp. 471–487, 2005.

[16] X. Yang and N. H. Vaidya, "Priority scheduling in wireless ad hoc networks," *Wireless Networks*, vol. 12, pp. 273–286, 2006.

1930-529X/07/$25.00 © 2007 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2007 proceedings.