Long-term Time-share Guarantees are Necessary for Wireless LANs

Godfrey Tan and John Guttag
MIT Computer Science and Artificial Intelligence Laboratory
{godfreyt, guttag}@csail.mit.edu

Abstract

Wireless local area networks (WLANs) based on a family of 802.11 technologies are becoming ubiquitous. These technologies support multiple data transmission rates. Transmitting at a lower data rate (by using a more resilient modulation scheme) increases the frame transmission time but reduces the bit error rate. In non-cooperative environments such as public hot-spots, individual nodes attempt to maximize their achieved throughput by adjusting the data rate or frame size used, irrespective of the impact of this on overall system performance.

In a series of experiments, we demonstrate that the existing distributed MAC protocol encourages non-cooperative nodes to use globally inefficient transmission strategies that lead to degraded aggregate throughputs. We also show that by establishing independence between the allocation of the shared channel time and the strategies used by individual nodes, an improved MAC protocol can lead rational but non-cooperative nodes to make choices that increase aggregate throughputs by as much as 30% under some conditions.

1 Introduction

802.11 is the *de facto* wireless networking standard. An 802.11 network may operate in two modes. In the *infrastructure mode*, a mobile node or station equipped with an 802.11b interface communicates via an access point or base station that is connected to the Internet, typically via a wired network. In the *ad hoc mode*, mobile nodes communicate in a peer-to-peer fashion in the absence of access points.

This paper deals with issues related to the way in which 802.11 WLANs resolve contention for the channel in non-cooperative environments, such as public *hot-spots* (e.g. WLANs at airports and coffee shops) or private enterprises that are physically close to each other (e.g. neighboring office suites in a commercial building or neighbor residences). The phenomenon observed in this paper applies to 802.11 networks running in either the infrastructure mode or the *ad hoc* mode. In non-cooperative environments, multiple nodes may compete for channel access in a *rational* but *non-cooperative* manner. That is each competing node will maximize its utility regardless of what other nodes achieve. We assume that the utility of each node is its achieved throughput.

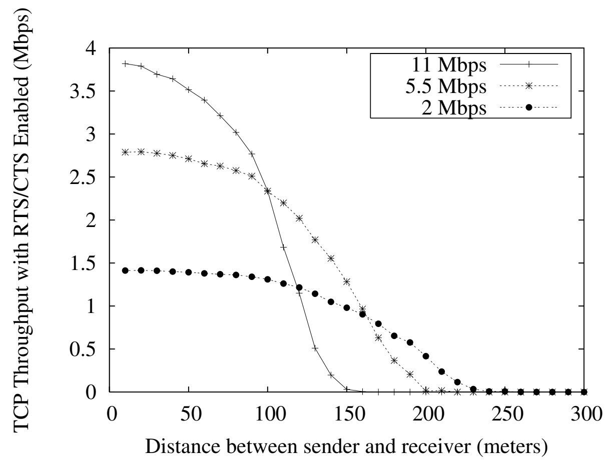

Figure 1: TCP throughputs achieved at various data rates in a simulated environment.

During congested periods, contention among nodes significantly affects the performance of indoor wireless systems. The bandwidth of the 2.4GHz band used by 802.11 is only wide enough for 3 orthogonal 802.11b or 802.11g channels, i.e., a maximum of 3 nodes can simultaneously transmit data with little interference. Contention among nodes using the same channel is resolved using a distributed medium access protocol (MAC), called the Distributed Coordination Function (DCF).

In practice, the channel loss rates of nodes vary widely due to the complex characteristics of indoor RF channels. Even receiving nodes that are equi-distant from a common sender experience differing channel conditions [3]. When the average signal strength at the receiver is lower than the threshold required for successful frame reception, the sender can unilaterally elect to use an alternative coding scheme that exploits the trade-off between data rate and BER [2]. Transmitting at a lower data rate by using a more resilient modulation scheme leads to higher frame transmission time but reduces the frame loss rate.

Figure 1 shows the achieved TCP throughput of a sender as a function of the distance between it and a receiver in a simulated environment. The channel model used in this

Data Rates (Mbps)	(11,5.5)	(11, 2)	(2,2)
$n0$'s throughput	2.955	2.009	0.992
$n1$'s throughput	0.556	0.655	0.411
Total (Mbps)	3.512	2.665	1.403

Table 1: The aggregate throughput is highest when $n0$ and $n1$ transmit at 11 and 5.5 Mbps respectively. However, at steady state, $n1$ lowers its data rate to 2 Mbps to achieve higher throughput, while significantly degrading the aggregate throughput. $n0$ cannot benefit by lowering its data rate and thus transmits at 11 Mbps at steady state.

simulation is described in detail in Section 3. For each pair of data rates, there exists a *cross-over distance* at which using a lower data rate yields higher throughput because the reduction in frame loss rate at the lower data rate is high enough to compensate for the slower transmission rate.

We also assume that each node has an implementation of DCF that complies with the 802.11 specification. Non-compliant malicious nodes are beyond the scope of this paper. In general, each 802.11 card is equipped with a proprietary implementation of DCF in the form of firmware and undergoes a certification process administered by the Wi-Fi Alliance, a nonprofit international association formed in 1999 to certify interoperability of WLAN products based on the 802.11 specification [10].

Each competing 802.11 node can use any 802.11-compliant strategy to maximize its achieved throughput. An 802.11 node can determine, for each frame transmission, the frame size and the data transmission rate. For the rest of this paper, we assume that each competing node uses maximum-sized data frames. If each node uses the most efficient transmission strategy, i.e. the strategy that yields the highest achievable throughput when the node alone occupies the channel, the resulting aggregate throughput will be optimal with respect to a particular allocation of channel capacity.

Table 1 shows the achieved TCP throughputs of two sending nodes $n0$ and $n1$, each of which sends data to a receiver. The distance between $n0$ and its receiver is 10 m whereas the distance between $n1$ and its receiver is 140 m. All nodes are within radio range of each other. As shown in Figure 1, in the absence of contention, the data rates that yield the highest achievable throughputs for $n0$ and $n1$ are 11 and 5.5 Mbps respectively. However, in the presence of competition, rational node $n1$ would lower its data rate to 2 Mbps to increase its achieved throughput by 18%. Unfortunately, this comes at the expense of reducing the aggregate throughput by 26%. Note that $n0$ would not benefit by reducing its data rate, so it transmits at 11 Mbps at steady state.

The root cause of this behavior is the mechanism used by DCF to dictate how the medium is shared. This variant of the carrier sense multiple access (CSMA) MAC protocol is designed to give an approximately equal probability of channel access (measured in number of *transmission opportunities*) to each competing node with similar loss characteristics. That is to say, over any period lasting several dozens milliseconds, each node will be able to transmit an equal number of frames, *irrespective of the amount of channel time required to transmit the frame*. Therefore, $n1$ being rational will lower its data rate if the reduction in its frame loss rate potentially results in higher throughput. However, by doing so, the throughput of $n0$ as well as the aggregate throughput will be significantly reduced because of the increased time required to transmit $n1$'s frames.

In general, if a node at a location to the left of a cross-over distance has the channel all to itself, it will never use a data rate lower than its optimal data rate, since doing so would result in reduced throughput. However, in the presence of another competing node, that node may transmit at a lower data rate, since by doing so it can use the channel longer and experience a lower loss rate. As a result, the overall efficiency of the network significantly suffers.

In this paper, we:

- Show through simulation the regions in which DCF allows competing nodes to use strategies that are inefficient in utilizing the shared medium,
- Argue our position that by guaranteeing the allocation of long-term shares of channel time to competing nodes with respect to a desired fairness constraint, the MAC protocol can force rational nodes to efficiently use the shared medium, thereby improving the achieved throughputs of all competing nodes.

The rest of this paper is organized as follows. The next section discusses related work. Section 3 examines the conditions under which DCF forces rational nodes to use inefficient strategies. Section 4 describes our proposed solution at the MAC layer and related preliminary results. We conclude in Section 5 by discussing the potential impact of our work and what we plan to do in the future.

2 Related Work

Previous studies [1, 9] have discussed some undesirable effects that DCF has on overall network performance when multiple competing nodes use different data rates. For instance, when competing nodes transmit at different data rates, the aggregate throughput, under many conditions, will be dominated by the lowest transmission rate.

Tan *et al.* [9] proposes that multi-rate WLANs use *time-based fairness*, in which each node is given an equal amount of channel time, as the fairness criterion and explains in detail why time-based fairness is more suit-

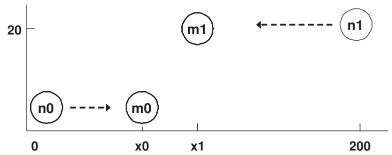

Figure 2: n0 and n1 transmit to m0 and m1 respectively.

able than *throughput-based fairness*, in which each node achieves equal throughputs.

Our work is independent of the choice of fairness criterion. We analyze the conditions under which DCF allows rational non-cooperative nodes to use inefficient strategies. The observations and results shown in later sections do not depend on a particular notion of fairness but rather concern with the existing mechanics of allocating the channel capacity. To our knowledge, we are the first to show that DCF encourages rational nodes to use undesirable strategies under certain conditions.

3 Analysis

In this section, we analyze the conditions that cause rational nodes to intentionally lower their data rates as the result of competition for the wireless channel.

3.1 Simulation Environment

We use ns [5] to conduct our simulations. RF propagation is modeled using a two-ray ground large-scale radio propagation model and a Rayleigh fast-fading model [6]. The latter models the fading phenomenon on short time-scales, which arises due to moving transmitters, receivers, or objects along transmission paths. The received power thresholds for various data rates are based on the Orinoco 802.11b Gold Card data sheet.

In practice, wireless card vendors employ proprietary auto-rate adaptation schemes that adjust the data transmission rate based on estimated channel conditions. Our results do not depend on a particular auto-rate protocol. For concreteness in our examples, we use the Receiver Based Auto Rate protocol (RBAR) [2]. RBAR relies on the existing RTS/CTS (request to send/ clear to send) mechanism, which deals with interference from hidden terminals, to assess the channel conditions experienced by each receiver. In particular, the sender sends an RTS frame and the receiver reports the received signal strength of the RTS frame in a replying CTS frame. Based on the signal strength information, the sender then chooses the highest transmission rate that has a high probability of a successful data frame transmission under the assumption that the channel conditions will remain unchanged for the transmission period. Figure 3 shows in most cases RBAR performs better than any single-rate protocol.

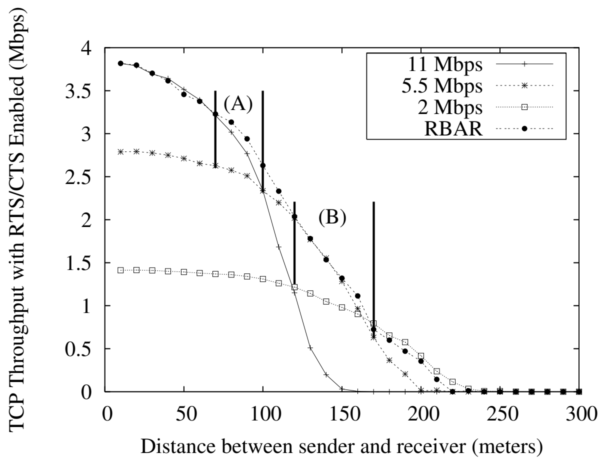

Figure 3: TCP throughputs achieved when using various fixed data rates and RBAR. A rational node whose transmission path distance is in region (A) or region (B) can benefit under DCF by intentionally using an inefficient strategy when competing against another node with a better channel condition, i.e. a competing node whose transmission path distance is shorter.

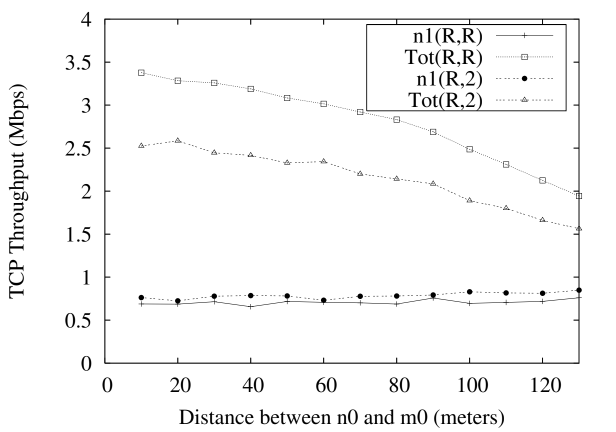

Figure 4: TCP throughputs achieved by $n1$ and the aggregate achieved throughputs under two pairs of strategies, $(R, 2)$ and (R, R), as a function of the transmission path distance of competing node $n0$. $(R, 2)$ denotes that $n0$ uses RBAR and $n1$ transmits at a fixed data rate of 2 Mbps. Tot plots the aggregate throughputs under a given strategy pair. The most efficient strategy for $n1$ is to transmit at 5.5 Mbps, which is what RBAR running at $n1$ would do. Thus, (R, R) denotes the most efficient strategy pair yet, at steady-state, $(R, 2)$ is used instead since $n1$ gains higher throughputs by transmitting at 2 Mbps.

3.2 Conditions Leading to Inefficiencies

We ran experiments using the setup shown in Figure 2. There are two TCP flows, one from $n0$ to $m0$ and the other from $n1$ to $m1$. Note that $m0$ and $m1$ also send TCP acknowledgment packets to $n0$ and $n1$ respectively. The goal is to investigate under what competition conditions $n1$ would benefit by intentionally lowering its data rate. Thus, we fixed the transmission path distance between $n1$ and $m1$ at 130 meters (i.e. $x1 = 70$ m) but varied that between $n0$ and $m0$ from 10 to 130 meters by moving $m0$. All nodes are within radio transmission range of each other.

When both transmitters used RBAR, $n1$ achieved lower throughput than $n0$ when $n1$'s transmission path distance was longer than $n0$'s. Notice that the optimal data rate for $n1$ would be 5.5 Mbps if $n1$ had the channel all to itself (see Figure 1). In fact, this was what RBAR did most of the time. However, in the presence of a competing flow, $n1$ could achieve higher throughput by transmitting at 2 Mbps. This behavior is evident in Figure 4, which shows the achieved throughputs of $n1$ and the aggregate throughputs as a function of the transmission path distance of competing node $n0$. For example, when $n0$'s transmission path distance is 10 m, node $n1$ can achieve an 11% increase in throughput by always transmitting at a lower (inefficient) data rate instead of using RBAR. However, as a result, the achieved throughput of $n0$ (not shown in the figure) and the aggregate throughput would decrease by 53% and 34% respectively.

We ran numerous experiments to determine the regions in which a rational node could benefit by transmitting at an inefficient data rate. In Figure 3, a node whose transmission path distance falls in region A or region B can achieve higher throughputs by intentionally choosing a data rate lower than the most efficient data rate whenever it competes against another node that experiences a lower loss rate. The wide ranges of regions A and B highlight the importance of incorporating mechanisms to reduce inefficiencies as a result of competition among rational but noncooperative nodes.

4 Solutions

There are two major ways to keep rational nodes from intentionally using inefficient strategies, i) Stipulate as part of the 802.11 specification that each card manufacturer implements a rate adaptation scheme that never intentionally uses an inefficient transmission strategy or ii) Modify the MAC protocol so that a node cannot gain additional throughputs by employing an inefficient transmission strategy.

The first approach requires that each card manufacturer implements a scheme like RBAR that selects a transmission strategy solely based on channel conditions and not according to the observed throughputs. When each rate adaptation protocol never selects an inefficient strategy intentionally, the resulting aggregate throughputs will be high. This solution although plausible has several drawbacks. First, since each manufacture employs its own proprietary auto-rate protocol, it will become much harder to verify whether a particular product conforms to the specification, and also to detect a malfunctioning or malicious card during the existing Wi-Fi certification process. Second, such a solution can still lead to unpredictable outcomes. This is because the channel capacity allocation among nodes can vary over time as DCF continues to allocate varying amounts of channel time according to the differing data rates used by competing nodes. Unpredictable capacity allocations may not be desirable for streaming media applications which often demand consistent allocations of the channel capacity.

An alternative approach is to develop an ideal MAC protocol that prevents a node's transmission strategy from affecting its share of channel time. Thus, the only way a node can maximize its throughput is by using its share of allocated channel time efficiently. Each card manufacturer can still employ a proprietary rate adaptation protocol to maximize the node's achieved throughput, but no such scheme can benefit by using an inefficient strategy. The rest of this section discusses how to develop such a MAC protocol.

4.1 Impact of Limiting Burst Duration

For 802.11-based wireless LANs, channel occupancy time, not the transmission opportunities, is the fundamental resource that needs to be shared among competing nodes. The channel occupancy time of a node is the total channel time used to transmit data frames. The time needed to transfer a data frame includes i) the transmission time of the data frame, ii) the transmission time of a synchronous MAC-layer acknowledgment, iii) the inter-frame idle periods as required by DCF and iv) the time required to transmit RTS and CTS frames when necessary.

It appears that if one fixes the amount of channel time a node garners per transmission opportunity as well as the probability of gaining transmission opportunities, a rational node will not intentionally lower its data rate since it cannot gain a higher amount of channel time by doing so. Unfortunately, this is not the case.

The reason is because for a given channel condition, the average number of frames transmitted per transmission opportunity or the average *burst length* varies with the data rate used since:

- The maximum number of successive frames [1] that can be successfully transmitted varies depending upon channel conditions. For instance, a sending node can almost always successfully transmit bursts of 5 consecutive frames at 11 Mbps to a receiver 10 meters away. But, when the same amount of frames is used to transmit to a receiver 100 meters away, the node will almost always experience a frame loss (within each burst).

[1] Each frame in a burst is still acknowledged and retransmitted in the usual manner.

- DCF requires all nodes to use an exponential backoff scheme. As soon as a frame loss is detected, DCF will perform a backoff by picking a random number of slots between 0 and the contention window size cw, which is doubled for every consecutive loss. A node cannot attempt any transmission during the backoff period.

When two nodes that have different average burst lengths compete against each other, the node with a smaller burst length will use a smaller share of channel time since it cannot make use of the maximum time allowed for each burst transmission. Therefore, that node may intentionally lower its data rate to obtain a higher fraction of channel time (by using up all of its allowed channel time in an inefficient manner) and to reduce its loss rate while reducing the aggregate throughput.

One may argue that DCF's backoff algorithm can be modified so that a node only backs off after sending the maximum number frames in each transmission opportunity, regardless of failures within the burst. Such a technique can potentially prevent rational nodes from intentionally lowering their data rates. However, it can also lead to lower overall network efficiency for the following reasons.

In indoor mobile environments, channel conditions are time-correlated on short time scales due to multipath and mobility [7], and thus, whenever a frame transmission fails due to channel errors, it is likely that successive frame transmissions will also fail. Thus, when multiple nodes are competing for channel access and losses are bursty (as they often do [4, 8]), a node will avoid likely failed transmissions by backing off as soon as it experiences a frame loss. Meanwhile, a competing node with a better channel condition can transmit, improving the overall efficiency. Studies have observed that the channel qualities of different transmission paths are often independent [4, 8]. Therefore, the existing backoff scheme of DCF that requires a node to back off as soon as a loss is encountered can potentially improve the overall network efficiency in the presence of competition.

We note that the 802.11e working group is currently drafting a standard containing specification of an enhanced version of DCF (EDCF) that allows nodes to transmit bursts of frames in a similar manner as we have presented here. For the reasons explained in this subsection, EDCF can force rational nodes to use inefficient strategies under some conditions.

4.2 Long-term Time-share Guarantees Work

We propose that the MAC protocol:

- Determines the desired long-term allocation of the shared medium based on channel time, instead of transmission opportunities,
- Limits the amount of channel time used per transmission opportunity (for the sole purpose of avoiding starvation), and
- Dynamically allocates the probability of transmission opportunities as a function of the observed channel time share so that the observed long-term global allocation of channel time is not affected by the transmission strategies used by nodes.

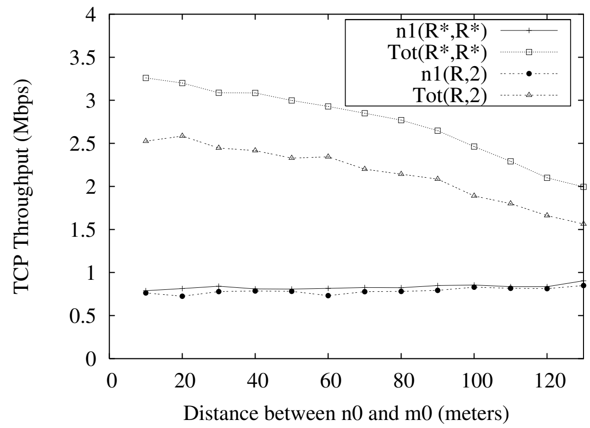

Figure 5: TCP throughputs achieved by $n1$ and the aggregate throughputs as the transmission path distance of competing node $n0$ varies. As explained in Figure 4, under DCF, $n1$ gains higher throughputs by transmitting at a less efficient rate of 2 Mbps. $n1(R, 2)$ and $Tot(R, 2)$ plot the achieved throughputs of $n1$ and the aggregate throughputs under DCF respectively. Under hypothetical MAC protocol DCF*, using (R^*, R^*) leads to the highest achieved throughputs for both nodes. $Tot(R^*, R^*)$ plots the aggregate throughputs under DCF*, which are superior to those achieved under DCF.

When the MAC protocol provides long-term channel time guarantees, rational competing nodes, knowing that they cannot gain additional channel time under any circumstances, will choose the most efficient strategies for the observed channel conditions.

DCF's distributed random backoff mechanism dictates the allocation of transmission opportunities for competing nodes. As we explained earlier, the contention window size is doubled with each loss experienced. Conversely, if the previous frame transmission is successful, the contention window is reset to a pre-determined minimum value, cw_{min}. Under DCF, cw_{min} for each node is 31. Under our proposed approach, cw_{min} of each node will be adjusted based on its observed channel time share. We note that other approaches such as directly adjusting the contention window size for each frame transmission (instead of cw_{min}) are also possible.

In practice, when competing nodes experience varying loss characteristics, DCF does not provide long-term guarantees of either transmission opportunities (because of different average backoff times) or channel time shares (because of varying per-frame channel times). In the DCF example of the previous section, the ratio of channel time shares used by $n1$ and $n0$ (at 10 m away from $m0$) was 0.6 when $n1$ was using the most efficient data rate (with RBAR). When $n1$ optimized its local throughput by transmitting at the less efficient data rate of 2 Mbps, the ratio rose to 1.7.

Under our approach, any desired allocation of long-term channel occupancy time can be ensured. To illustrate potential realizable gains under our approach, we assume that a desired ratio of channel time shares of $n1$ to $n0$ of 0.7. Clearly, this desired goal is not achievable under DCF if $n0$ uses RBAR. Under our approach, $n1$ will be assigned a larger probability of winning transmission opportunities than under DCF, so that the system achieves the desired long-term allocation of channel time.

For this particular example, we set $cw_{min} = 26$ for $n1$ and $m1$, and $cw_{min} = 36$ for $n0$ and $m0$. Figure 5 shows that under DCF*, i.e. DCF with our modifications, the achieved throughput of $n1$ using RBAR, $n1(R^*, R^*)$, is always higher than what $n1$ achieved by always transmitting at 2 Mbps under DCF, $n1(R, 2)$.

In practice, cw_{min} should be dynamically adjusted in a distributed fashion as a function of the observed and desired channel time shares. This ensures that the long-term channel time share of a node will be the same regardless of the data rate used. Therefore $n1$, being rational, will not lower its data rate under DCF*, and aggregate throughput is improved by as much as 30%.

Unlike the other approach of stipulating auto-rate protocol as part of the specification, this approach allows for easier testing and verification of a particular implementation of the MAC protocol, since the behaviors of a specification-compliant MAC can be easily understood.

5 Discussion

The proliferation of hotspots and independently manged WLANs leads to a situation in which there is increasing competition for the wireless channel by non-cooperative devices. The MAC layer protocols in use today, namely DCF and its future successor EDCF, encourage rational nodes to use inefficient strategies under some realistic conditions, degrading aggregate network performance.

We also showed that long-term channel time share guarantees (rather than transmission opportunity guarantees) can be used to ensure that rational competing nodes use the channel time allocated to them in the most efficient manner. A MAC protocol can achieve this goal by dynamically adjusting in a distributed manner the contention window size of each node as a function of its observed channel time share.

We believe that such a scheme is practical and should be built into future 802.11 standards so that the 802.11-compliant devices use efficient strategies. There are two major challenges in developing such a scheme. First, each node must observe its share of channel time. The information needed to do this is already available under DCF, since a node knows the time taken to transfer its own frames and can infer the time taken by neighboring nodes from existing physical and virtual carrier sensing mechanisms. Second, each node must periodically determine either its contention window size or cw_{min} as a function of its channel time share. This must be done in a way that ensures the convergence of the observed global channel time allocation to the desired allocation. We are currently developing such a scheme. We believe that the increased protocol complexity will be more than offset by realizable significant performance gains in non-cooperative environments, but have yet to demonstrate this experimentally.

References

[1] Martin Heusse, Franck Rousseau, Gilles Berger-Sabbatel, and Andrzej Duda. Performance anomaly of 802.11b. In *INFOCOM*, April 2003.

[2] Gavin Holland, Nitin H. Vaidya, and Paramvir Bahl. A rate-adaptive MAC protocol for multi-hop wireless networks. In *Mobile Computing and Networking*, pages 236–251, 2001.

[3] David Kotz, Calvin Newport, and Chip Elliott. The mistaken axioms of wireless-network research. Technical Report TR2003-467, Dept. of Computer Science, Dartmouth College, July 2003.

[4] Allen Miu, Godfrey Tan, Hari Balakrishnan, and John Apostolopoulos. Divert: Fine-grained Path Selection for Wireless LANs. In *Proc. of ACM Mobisys 2004*, Boston, MA, Jun 2004.

[5] ns-2 Network Simulator. `http://www.isi.edu/vint/nsnam/`.

[6] R. Punnoose, P. Nikitin, and D. Stancil. Efficient simulation for Ricean fading within a packet simulator. In *IEEE Vehicular Technology Conference*, pages 764–767, 2000.

[7] T.S. Rappaport. *Wireless Communications*. Prentice Hall, Upper Saddle River, N.J., 1996.

[8] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly. Opportunistic media access for multirate ad hoc networks. In *MOBICOM*, sept 2002.

[9] Godfrey Tan and John Guttag. Time-based Fairness Improves Performance in Multi-rate WLANs. In *Proc. of the Usenix Annual Technical Conference*, Boston, MA, June 2004.

[10] Wi-Fi Alliance Web Site. `http://www.wi-fi.org`.