DYNAMIC CHANNEL RATE ASSIGNMENT FOR MULTI-RADIO WLANS

Qian Wu Carey Williamson
Department of Computer Science, University of Calgary
Calgary, AB, Canada T2N 1N4
Email: {qianwu,carey}@cpsc.ucalgary.ca

ABSTRACT
The dynamic transmission rate selection feature of IEEE 802.11b WLANs can cause throughput degradation when stations with different transmission rates share the same physical channel. One solution to this problem is to use a Multi-Rate Multi-Channel (MRMC) MAC protocol. However, static channel rate assignment in MRMC can lead to load imbalance across channels. In this paper, we present a dynamic channel rate assignment scheme for the MRMC MAC protocol. The ns-2 network simulator is used to evaluate the approach in an MRMC WLAN with up to 20 mobile hosts. Our simulation results show that dynamic channel rate assignment can improve aggregate TCP throughput by 30-54% compared to static channel rate assignment.

KEY WORDS
IEEE 802.11b, Wireless LANs, MAC protocol design, Multi-channel MAC, Load Balancing

1 Introduction

In recent years, IEEE 802.11 [2] wireless local area networks (WLANs) have been widely deployed in universities, offices, hotels, airports, and other public places. These WLANs, which rely on an Access Point (AP) to relay traffic to and from the wired Internet, offer physical-layer transmission rates of 11 Mbps for IEEE 802.11b [3] and 54 Mbps for IEEE 802.11a [4].

IEEE 802.11b WLANs can use dynamic transmission rate selection on a per-frame basis to overcome adverse wireless channel conditions. However, overall system throughput may suffer when stations with different transmission rates share the same physical channel [7, 10]. Basically, low-rate stations monopolize the channel, forcing high-rate stations to wait.

In prior work, we proposed a Multi-Rate Multi-Channel (MRMC) MAC-layer protocol [13] to solve the throughput degradation problem. This protocol exploits the capabilities of multi-channel or multi-radio WLANs [1, 5, 9, 21]. In the MRMC protocol, *multiple* channels are used in an AP *simultaneously*. A mobile station associates with a particular channel based on its received signal strength from the AP. The MRMC protocol keeps low-rate and high-rate stations on different physical channels, so that the channel access of low-rate stations does not affect high-rate stations.

In our prior work [13], we assumed 4 channels, with transmission rates of 1, 2, 5.5, and 11 Mbps. To simplify the protocol, the bit rate was statically assigned to each channel, However, when mobile stations are unevenly distributed around an AP, static channel rate assignment can cause load imbalance (e.g., one crowded channel, and several underutilized channels).

In this paper, we propose and evaluate a dynamic channel rate assignment scheme for the MRMC protocol. The key idea is to balance load across channels. By avoiding excessive load on a crowded channel, we can improve the aggregate throughput performance of the MRMC protocol. We evaluate the approach using ns-2 network simulation, finding that dynamic channel rate assignment can improve the TCP throughput of MRMC by 30-54%.

The rest of the paper is organized as follows. Section 2 provides a brief discussion of background and related work. Section 3 describes our proposed algorithm for dynamic channel rate assignment. Section 4 describes the simulation methodology for performance evaluation of the protocol, and Section 5 presents the simulation results. Section 6 discusses issues related to the dynamic rate assignment scheme. Section 7 concludes the paper.

2 Background and Related Work

2.1 IEEE 802.11 WLANs and MRMC MAC

An IEEE 802.11 WLAN can be configured in either *infrastructure* mode or *ad hoc* mode. In infrastructure mode, an AP acts as a central point to relay traffic to and from the Internet. In ad hoc mode, no AP is required; stations communicate among themselves in a peer-to-peer fashion.

IEEE 802.11b WLANs support dynamic transmission rate selection to combat adverse wireless channel conditions. Four transmission rates are allowed: 1, 2, 5.5, and 11 Mbps. The sender of a frame can decide the transmission rate to use based on recent observations of wireless channel conditions (e.g., successful transmissions, missing ACKs, excessive retransmissions). While the multi-rate scheme offers performance advantages, it can suffer from throughput degradation when several mobile stations share the same physical channel [7, 10].

The MRMC protocol [13] requires a multi-channel or multi-radio WLAN operating in *infrastructure* mode. That is, multiple channels are used by the AP simultaneously.

The same protocol could be used in a multi-radio wireless mesh network, though we have not yet explored this idea.

The MRMC protocol uses different transmission channels to isolate high-rate stations from low-rate stations. The AP broadcasts a beacon frame every 100 ms. A mobile station receiving the beacon frames uses the received signal strength to determine the appropriate channel and transmission rate to use. In our prior work [13], we used static channel rate assignment.

In this paper, we discuss a dynamic channel rate assignment scheme for MRMC. To the best of our knowledge, there has been no study of dynamic channel rate assignment for multi-channel WLANs. M. El-Kadi *et al* [8], A. Malla *et al* [14], and B. Wolfinger *et al* [20] have exploited dynamic resource allocation for QoS provisioning in cellular networks and wireline networks. Their proposed schemes were not applied to wireless LANs. Some research has studied dynamic load balancing in Wireless LANs between APs [6, 15, 17]. However, our problem is slightly different. In particular, dynamic channel rate assignment must consider not only load (i.e., number of associated stations), but also transmission rates and wireless channel conditions. These factors significantly affect WLAN performance, as discussed next.

2.2 Wireless Channel Modeling

In digital communication theory [16], the bit error rate (BER) of a modulation scheme depends on the received *signal-to-noise ratio* (SNR). In general, the higher the received SNR, the lower the BER. On the other hand, for a given SNR, simpler modulation schemes tend to have lower BER. That is, since simpler modulation schemes generally represent lower bit rates, a frame transmitted with a lower bit rate is less likely to experience errors than a frame transmitted with a higher bit rate at the same SNR.

The received SNR is largely determined by the propagation environment. In a wireless channel, the large-scale path loss and small-scale (multipath) fading are the two main factors that affect the SNR [18]. Path loss determines the *mean* signal strength at a certain receiver distance. Multipath fading is caused by the superposition of multiple in-phase and out-of-phase copies of the original transmitted signal. It can cause rapid fluctuations in received signal strength over very short time scales. The received signal strength depends on the transmitted power, the path loss, and the multipath fading characteristics.

There are well-established mathematical models for path loss and multipath fading [18]. In a non-line-of-sight wireless propagation environment, multipath fading is usually modeled by the Rayleigh channel model, which we use in our simulations. The path loss (PL) can be calculated using the log-distance path loss model [18]:

$$\overline{PL(dB)} = \overline{PL(d_0)} + 10\, n\, log(\frac{d}{d_0}) \qquad (1)$$

where n is the *path loss exponent* (typically 2-6 for indoor

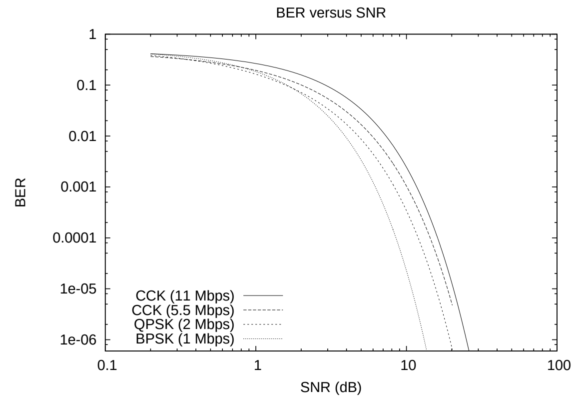

Figure 1. BER for IEEE 802.11b Modulation Schemes

propagation environments, and set to $n = 3$ in our simulations), d is the distance between the transmitter and the receiver, and d_0 is the close-in reference distance (typically 1 meter for indoor propagation environments). $\overline{PL(d_0)}$ is the mean received power (in dB) at the close-in reference distance d_0. The mean received power (in Watts) at d_0 can be estimated using the Friss free space propagation model:

$$\overline{P_r(d_0)} = \frac{P_t G_t G_r \lambda^2}{(4\pi)^2 d_0^2 L} \qquad (2)$$

where P_t and P_r are the transmit and receive powers, G_t and G_r are the transmit and receive antenna gains (typically 1), λ is the carrier wavelength, and L is the system loss factor (typically 1). $\overline{PL(dB)}$ is then equal to $10\, log_{10}(\overline{P_r(d_0)})$.

The BER for a given modulation scheme can be calculated from the received SNR. For well-known modulation schemes, theoretical work expresses the relationship between BER and SNR. For example, the BER of Binary Phase Shift Keying (BPSK), used for the 1 Mbps transmission rate in IEEE 802.11b, is given by:

$$P_b = Q(\sqrt{\frac{2 \cdot SNR \cdot B}{R}}) \qquad (3)$$

where B is the bandwidth (in Hz) of the modulated signal and R is the bit rate of the modulation scheme [16].

Figure 1 shows the BER versus SNR performance for all modulation schemes used in IEEE 802.11b. The 2 Mbps transmission rate uses Quadrature Phase Shift Keying (QPSK), while the 5.5 and 11 Mbps transmission rates use Complementary Code Keying (CCK) [3]. The results in Figure 1 were calculated using equation 5.2-34 in [16]. The other relevant BER formula can also be found in [16].

From Figure 1, we can determine the required SNR for a specific BER and modulation scheme. This knowledge can be used to set the threshold for transmission rate selection (i.e., modulation scheme) based on the received SNR. For example, if the maximum BER is set to 10^{-5} and the SNR falls below the required value for the current modulation scheme, the station needs to adjust its rate.

3 Dynamic Channel Rate Assignment

The motivation for dynamic channel rate assignment is to improve load balancing, and thus the effective utilization of the resources available in a multi-channel WLAN. Stations should share channel resources fairly and efficiently, based on location, load, wireless channel conditions, and desired transmission rates.

An important design decision concerns the locus of control. That is, should channel rate assignments be made individually by the mobile stations or centrally by the AP? We evaluated both approaches in preliminary experiments, finding that central assignment by the AP is superior.

There are two reasons why centralized decision-making is preferable. The first reason, which is intuitively obvious, is that the AP has more information with which to make decisions about load balancing. The second reason, which is less obvious, relates to system stability. If the AP periodically broadcasts the current list of channels (with their corresponding data rates and load levels) in the WLAN, then mobile stations tend to undergo synchronized channel switching (i.e., a ping-pong-like shifting of load, which does not resolve load imbalance). The synchronization effect reduces system efficiency, mitigating the advantages of the MRMC protocol. One solution to the synchronization problem is the randomization of channel switching decisions. Another solution (our approach) is to have the AP assign channels, in response to probe requests generated asynchronously by the mobile stations.

This design provides the framework for our dynamic channel rate assignment scheme. Rate assignment is basically an extended version of the channel association procedure used in most WLANs.

Figure 2 shows our algorithm for dynamic channel rate assignment. At boot time, an AP assigns a default rate (e.g., 11 Mbps) to each channel that it supports. This transmission rate is only temporary; it can be changed later as needed. The AP also maintains a channel association table, to keep track of the stations associated with each channel, and the channel's current transmission rate.

Each mobile station periodically sends a probe frame to its associated AP. Upon receiving a probe frame, the AP calculates a moving average SNR for that station using:

$$SNR_{avg} = \alpha SNR_{avg} + (1-\alpha) SNR_{new}$$

where SNR_{new} is the SNR of the incoming probe frame, and $\alpha = 0.9$ [13].

Based on the SNR result, the AP uses a threshold-based algorithm to determine the *candidate transmission rate* R for the mobile station. Given a set of thresholds T_i ($i = 1, 2, \ldots, N-1$), determined from a graph like Figure 1 for the target BER, the rate R is:

R_1 (highest) if $SNR_{avg} > T_1$

R_i (medium) if $T_i < SNR_{avg} \leq T_{i-1}$

R_N (lowest) otherwise

Given the candidate rate R, the AP must find a *candidate channel* C for the station to associate with. If there is an empty channel with rate R, then this channel is assigned. Otherwise, the AP searches for *any* empty channel. If there is one, then its rate is changed to R. Otherwise, the AP must find the most suitable channel that is currently in use by other stations. The preference is for the least loaded channel with the proper rate R. If there is no such channel, then channels with "close" rates are considered.

We consider two different heuristics for the latter search: "higher rates first" and "lower rates first". If the transmission rate of the (occupied) candidate channel C differs from the desired candidate rate R, then the candidate rate R is altered to match that of the candidate channel C. This approach minimizes the impact of a new station joining an occupied channel.

The last step of the association process switches the mobile station to the candidate channel C (if different from the station's current channel). The AP sends to the mobile station a channel association frame with the channel id and rate. The mobile station sends back a channel association ack frame on the old channel. The AP updates the channel association table as required. The mobile station switches to the new channel C, modifies its transmission rate to be R, and resumes normal communication with the AP.

4 Simulation Methodology

We evaluate our dynamic channel rate assignment scheme using the `ns-2` network simulator [19]. The primary performance metric is aggregate TCP throughput. We compare the dynamic channel rate assignment scheme with static rate assignment in both stationary and mobile scenarios, using the same wireless channel error model as in [13].

Two main assumptions are made for the evaluation. First, we assume that mobile stations are active most of time. This assumption ensures that the number of stations on each channel is a good indicator for channel load. Second, we assume that 4 channels and 4 transmission rates are available. This assumption is consistent with our previous work for static channel rate assignment [13].

The network model simulated is shown in Figure 3. An FTP (TCP) server on a wired network is connected to an AP via 100 Mbps Ethernet. The transmission power of the AP covers (on average) a circular area with a radius of 45 meters. Multiple stations within this range communicate with the AP via the wireless channels. They act as TCP clients, receiving fixed-size 1500-byte packets (frames) from the FTP server. This is the same network model and workload as in [13]. We use FTP traffic in the experiments in order to assess steady-state TCP throughput performance. Our dynamic rate assignment scheme should benefit other traffic types as well.

The `ns-2` TCP model is used to simulate the FTP server and clients. All simulations use the TCP NewReno model. The application layer has infinite data to send (i.e., FTP-like bulk transfers).

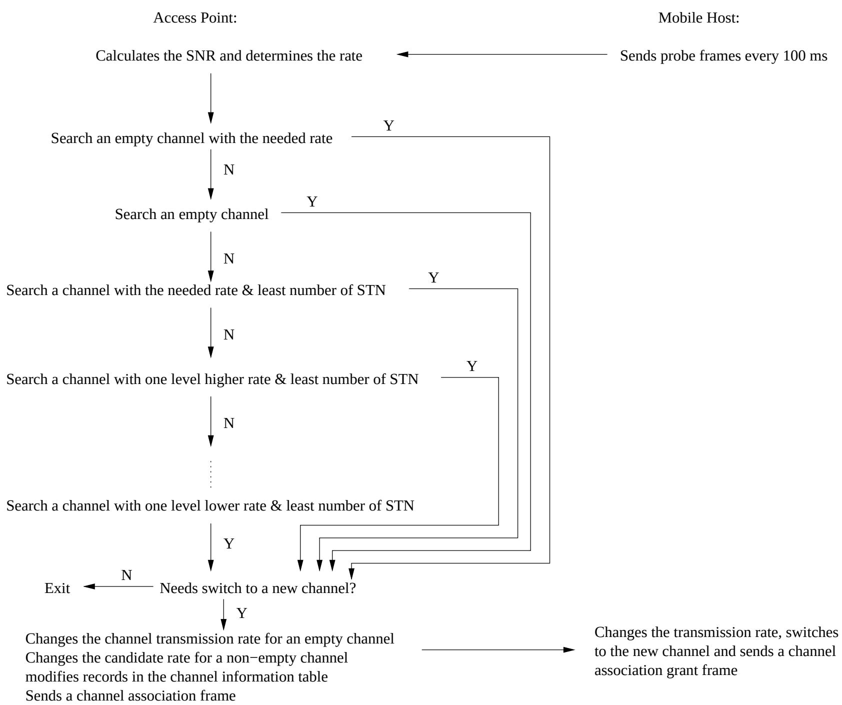

Figure 2. Flowchart of the Dynamic Rate Assignment Scheme

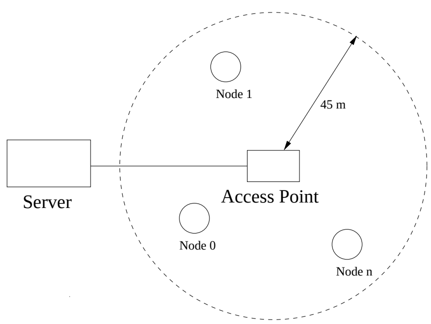

Figure 3. Simulation Model

Two sets of experiments were conducted: one for stationary hosts and one for mobile hosts. Each experiment runs for 300 seconds of simulated time.

The first set of experiments evaluates dynamic channel rate assignment scheme in a scenario with stationary hosts. These stations were placed randomly around the AP, and activated at random start times. The number of stations was varied from 2 to 20. Each station initiated a TCP transfer from the server. We repeated each experiment 30 times, with different seeds for the random placement of stations, the station activation times, and the FTP start times. We calculate total throughput and 99% confidence intervals.

The second set of experiments considers host mobility. Mobile stations move randomly in the circular area around the AP at the average speed of 2 m/s. The number of mobile stations was varied from 2 to 20. Again, we repeat each experiment 30 times, with different seeds for the station activation times and the FTP start times. We report aggregate throughput and 99% confidence intervals.

5 Simulation Results

5.1 Stationary Scenario

Figure 4 shows the simulation results for the scenario with stationary hosts. The horizontal axis shows the number of hosts, while the vertical axis shows the aggregate TCP throughput achieved by these hosts. Recall that a single host on an 11 Mbps WLAN can typically achieve a throughput of 5 Mbps.

There are four different lines represented on the graph in Figure 4. The top line (with boxes) is for the “lower rates first” version of the algorithm in Figure 2. The second line (with ’X’) is for the “higher rates first” variation. The third line (solid) is for static channel rate assignment, as in the original MRMC protocol [13]. The lowest line is for a conventional single-channel IEEE 802.11b WLAN, using the WaveLAN-II algorithm for dynamic rate adaptation [12]. The latter line is presented for reference purposes.

Three main observations are evident from Figure 4. First, the MRMC MAC protocol provides 200-500% improvement compared to a conventional IEEE 802.11b WLAN. These results are consistent with our prior work [13]. Second, dynamic channel rate assignment improves upon the static channel rate assignment in MRMC. The magnitude of this improvement depends on the num-

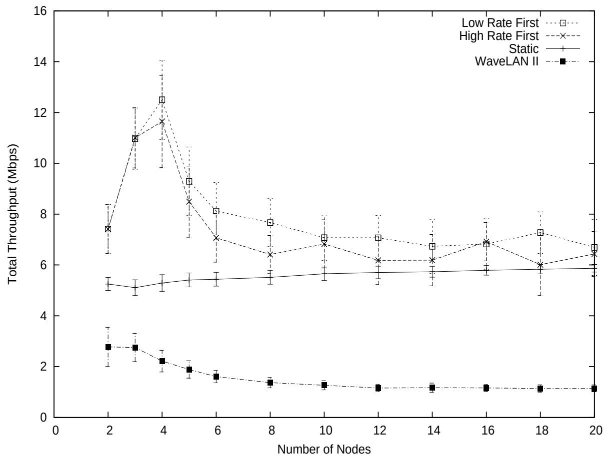

Figure 4. Results for Stationary Scenario

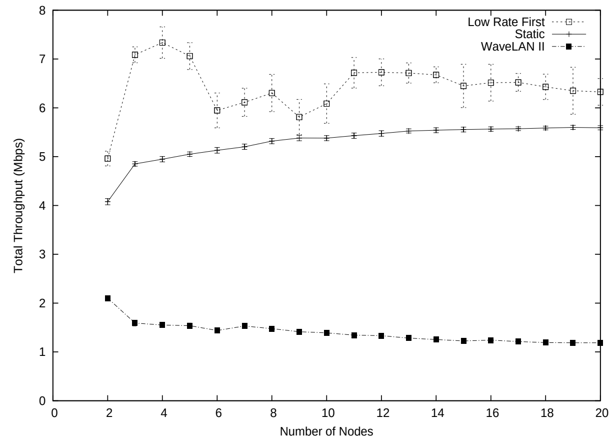

Figure 5. Results for Mobile Scenario

ber of hosts: there is a large advantage (100-250%) with 2-5 hosts, and a small advantage (10-20%) with many hosts. Third, the mean throughput for "lower rates first" is slightly higher than that for "higher rates first".

The performance advantages of dynamic channel rate assignment for 2-5 hosts are easily explained. With few hosts, empty channels are almost always available, so load balancing amongst occupied channels is trivial. Furthermore, *all* of the channels can potentially operate at 11 Mbps, depending on the host location and the received SNR. In static channel rate assignment, the channel rates are always 1, 2, 5.5, and 11 Mbps.

When the number of stations exceeds the number of channels, channel contention is inevitable, and the total throughput decreases. However, there is still an advantage over static channel rate assignment. The advantage comes from the traffic isolation property of the MRMC protocol, and the improved load balancing.

Dynamic channel rate assignment outperforms static rate assignment in all cases studied. The average performance gain is 54% for "lower rates first", and 43% for "higher rates first".

"Lower rates first" outperforms "higher rates first" for two reasons. First, since higher transmission rates require more complicated modulation schemes, assigning a station with marginal SNR to such a channel means that their frames are vulnerable to transmission errors [11, 18]. Second, an error-prone station can degrade the throughput of other users on the high-rate channel (i.e., retransmissions). Retransmitted frames, when required, use the same transmission rate selection algorithm as the original transmission attempt.

Figure 4 also shows that dynamic channel rate assignment is quite sensitive to host location. That is, the 99% confidence intervals show higher variance in the aggregate throughput, based on the random host placement in the 30 simulation runs. The variance is much higher than for the other MAC protocols.

This increased variance arises because of the higher potential throughput available using dynamic channel rate assignment. For example, with 4 nodes (all on separate channels), the aggregate throughput could range from 3 Mbps (all far from the AP, using 1 Mbps) to 20 Mbps (all close to the AP, using 11 Mbps). Static channel rate assignment would produce results ranging from 0.7 Mbps to 5 Mbps for these same host placement scenarios. Thus the higher variance in throughput should not be construed as a disadvantage of the protocol.

5.2 Mobile Scenario

Figure 5 presents the simulation results from the mobile host scenario. For clarity, we remove the results for "higher rates first", since "lower rates first" is always better.

Figure 5 shows that dynamic channel rate assignment consistently outperforms static rate assignment in the mobile scenario. While the overall throughput results are lower than in Figure 4 because of host mobility, there is still a consistent performance improvement of 20-30%. The largest advantage occurs when the number of mobile stations matches the number of channels.

6 Discussion

One possible concern about dynamic rate assignment is the overhead caused by channel switching and channel rate changes. This overhead depends on the channel load and on station mobility. For wireless LANs, mobile stations usually move at walking speed, so the overheads caused by channel switching and channel rate changes are small compared to the performance gains from balancing load.

In our study, we use the number of mobile stations as the load estimator, assuming that mobile stations are active most of time. For this workload, dynamic rate assignment is most effective when the mobile stations are unevenly distributed in the WLAN. If the traffic from mobile stations is bursty, or if stations sleep occasionally, then other criteria are needed to estimate the traffic load accurately.

7 Conclusions

In this paper, we proposed a dynamic channel rate assignment scheme for the MRMC protocol [13]. The main purpose of the proposed scheme is to balance load across the channels used by an AP. Our dynamic channel rate assignment scheme tries to find empty or lightly loaded channels.

We evaluated the performance of the proposed dynamic rate assignment scheme using simulation, with the ns-2 network simulator. We considered a 4-channel WLAN environment, with both stationary and mobile host scenarios. Our simulation results show that the dynamic channel rate assignment scheme improves aggregate TCP throughput by 54% compared to static channel rate assignment in MRMC. The performance gain is 30% in the scenario with mobile hosts.

Future work will explore the applicability of our protocol to wireless mesh networks.

References

[1] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou, A Multi-Radio Unification Protocol for IEEE 802.11 Wireless Networks, *Proc. IEEE BROADNETS'04*, San Jose, USA, 2004, 344-354.

[2] ANSI/IEEE Standard 802.11, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 1999.

[3] ANSI/IEEE Standard 802.11b, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Higher-Speed Physical Layer Extension in the 2.4 GHz Band, 1999.

[4] ANSI/IEEE Standard 802.11a, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Higher Speed Physical Layer in the 5 GHz Band, 1999.

[5] P. Bahl, A. Adya, J. Padhye, and A. Wolman, Reconsidering Wireless Systems with Multiple Radios, *ACM SIGCOMM Computer Communication Review*, 34(5), 2004, 39-46.

[6] Y. Bejerano, S. Han, and L. Li, Fairness and Load Balancing in Wireless LANs Using Association Control, *Proc. ACM/IEEE MOBICOM'04'*, Philadelphia, USA, 2004, 315-329.

[7] G. Cantieni, Q. Ni, C. Barakat, and T. Turletti, Performance Analysis of Finite Load Sources in 802.11b Multirate Environments, to appear in *Computer Communications Journal, special issue on Performance Issues of Wireless LANs, PANs, and Ad Hoc Networks*.

[8] E. El-Kadi, S. Olariu, and H. Abdel-Wahab, A Rate-Based Borrowing Scheme for QoS Provisioning in Multimedia Wireless Networks, *IEEE Trans. on Parallel and Distributed Systems*, 13(2), 2002, 156-166.

[9] Engim, Inc., EN-3001 Intelligent Wideband WLAN Technology, http://www.engim.com/products_technology.html.

[10] M. Heusse, F. Rousseau, G. Berge-Dabbatel, and A. Duda, Performance Anomaly of 802.11b, *Proc. IEEE INFOCOM'03*, San Francisco, USA, 2003, 836-843.

[11] G. Holland, N. Vaidya, and P. Bahl, A Rate-Adaptive MAC Protocol for Multi-Hop Wireless Networks, *Proc. ACM/IEEE MOBICOM'01*, Rome, Italy, 2001, 236-251.

[12] A. Kamerman and L. Monteban, WaveLAN-II: A High-Performance Wireless LAN for the Unlicensed Band, Bell Labs Technical Journal, 1997.

[13] T. Kuang, Q. Wu, and C. Williamson, MRMC: A Multi-Rate Multi-Channel MAC Protocol for Multi-Radio Wireless LANs, submitted for publication, 2005.

[14] A. Malla, M. El-Kadi, S. Olariu, and P. Todorova, A Fair Resource Allocation Protocol for Multimedia Wireless Networks, *IEEE Trans. on Parallel and Distributed Systems*, 14(1), 2003, 63-71.

[15] I. Papanikos and M. Logothetis, A Study on Dynamic Load Balance for IEEE 802.11b Wireless LAN, *Proc. COMCON 2001*, Crete, Greece, 2001.

[16] J. Proakis, *Digital Communications*, Fourth Edition, (New York: McGraw-Hill, 2001).

[17] S. Sheu and C. Wu, Dynamic Load Balance Algorithm (DLBA) for IEEE 802.11 Wireless LAN, *Tamkang Journal of Science and Engineering*, 2(1), 1999, 45-52.

[18] R. Theodore, *Wireless Communications: Principles and Practice*, Second Edition, (Upper Saddle River: Prentice Hall PTR, 2002).

[19] VINT Group, Network Simulator ns-2 (version 2.26), available at: http://www.isi.edu/nsnam/ns.

[20] B. Wolfinger, J. Wolf, and G. Le Grand, Improving Node Behavior in a QoS Control Environment for Local Broadcast Networks, *Proc. SPECTS 2003*, Montreal, Canada, 2003, 361-371.

[21] S. Wu, Y. Tseng, C. Liu, and J. Sheu, A Multi-Channel MAC Protocol with Power Control for Multi-Hop Mobile Ad Hoc Networks, *The Computer Journal*, 45(1), 2002, 101-110.