JOINT RATE ADAPTATION AND CHANNEL-ADAPTIVE RELAYING IN 802.11 AD HOC NETWORKS

Michael R. Souryal* and Nader Moayeri
Wireless Communication Technologies Group
National Institute of Standards and Technology
Gaithersburg, Maryland, USA

ABSTRACT

Channel-adaptive relaying has recently been proposed as a means to exploit spatial and temporal diversity in multihop ad hoc networks with fading. In conjunction with appropriate routing protocols, adaptive relaying enables each forwarding node in a multihop path to dynamically select the next-hop relay as a function of the measured (time-varying) channel state, providing a form of selection diversity at each hop. Based on the notion that links to diversity-selected relays have higher information capacity and therefore can support higher data rates than links obtained with traditional routing, this paper proposes marrying channel-adaptive relaying with rate adaptation (or adaptive modulation-coding). In particular, we specify a protocol for performing joint rate and relay adaptation in 802.11 ad hoc networks with geographic routing. Using both analytical and simulation tools, synergistic gains are observed in throughput, capacity and delay. Performance results are given for individual links as well as for multihop networks, in time-varying, correlated Rayleigh and Ricean fading channels over a range of channel speeds. Of particular interest in this study is the robustness of the adaptation to increasing channel Doppler. As a by-product of this work, we propose a new, SNR-based rate adaptation scheme for use in 802.11 systems that requires no modification to the standard 802.11 frame structure.

I. INTRODUCTION

Link quality in ad hoc networks varies due to mobility and bursty interference, and can change dramatically from one transmission to the next. Most ad hoc network routing protocols rely on the consistent and stable performance of individual links, and intermittent links can result in high packet loss rates and control overhead [1]. Adaptive techniques can help mitigate the effects of varying channel state as well as exploit favorable channel conditions when they exist. For example, with adaptive modulation and/or coding, a more robust signaling scheme is used on lower quality links while a more spectrally efficient scheme is used on higher quality links. Adaptive power, or power control, adjusts the transmission power based on the channel quality, and is used when the objective is to conserve energy or limit interference. Furthermore, these adaptive techniques can be used jointly, as has been analyzed for the two-user [2] and cellular [3] cases.

In the context of wireless LANs and ad hoc networks, link adaptation has been proposed for IEEE 802.11 systems, as well. Since versions of 802.11 already specify multiple data rates that use different modulation-coding schemes, rate adaptation aims to choose the most appropriate data rate at any given time. It exploits periods of time during which channel quality is good (e.g., favorable fading, low interference) to increase throughput. The AutoRate Fallback (ARF) protocol uses information available at the medium access control (MAC) layer in the form of received and missing acknowledgments along with a heuristic algorithm to select the data rate [4]. The authors of [5] improve upon this algorithm by increasing its responsiveness to channel variation. Other approaches select the data rate using information reported by the physical layer, such as the received signal strength or signal-to-noise ratio (SNR) [6–8]. Of these analyses, only [6] includes multihop performance, with the others focusing on single-hop performance.

While rate adaptation exploits temporal changes in link quality to improve overall throughput, it does not take advantage of the spatial diversity in multihop networks which could yield even greater efficiencies. Because multipath fading and possibly shadow fading are typically uncorrelated from a given trasmitting node to different receiving nodes, some forwarding relays will offer better performance than others on any given hop. Extremely Opportunistic Routing (ExOR) [9] takes advantage of this spatial diversity by multicasting a packet to a set of candidate relays, after which one relay is chosen to forward the packet. However, it is not clear if and how ExOR can be utilized in conjunction with fast rate adaptation since the forwarding node is selected only after the transmission is complete. In previous work, we pro-

*This research was performed while the author held a National Research Council Research Associateship Award at the National Institute of Standards and Technology.

posed a technique for selection of the forwarding relay on a per-hop and per-packet basis using timely channel measurements, termed "channel-adaptive relaying," but it was formulated for a single-rate system [10]. In this paper, we describe an approach for joint rate adaptation and relay selection in a multirate system. We specify protocols and metrics appropriate for 802.11b multihop networks employing geographic routing as an example.[1] The analysis shows that combined relay selection and rate adaptation provides synergistic gains, for relay selection tends to select links with higher capacity while rate adaptation makes efficient use of those links.

After reviewing channel-adaptive relaying, we describe our proposed rate adaptation scheme and explain how it is combined with channel-adaptive relaying. An analytical link performance evaluation indicates the potential gains, and a more detailed evaluation of network performance by simulation examines end-to-end measures for a range of mobility levels.

II. CHANNEL-ADAPTIVE RELAY SELECTION

Channel-adaptive relaying was described in [10] as a cross-layer means for achieving spatial diversity in multihop networks to combat fading. It operates in conjunction with a routing protocol that provides each forwarding node with multiple candidate next-hop relays to which to forward the packet. We briefly review the approach here for clarity and completeness. For each packet transmission, adaptive selection of the next relay is performed at two levels. At the routing layer, the forwarding node selects up to L relays based on its measurement of long-term channel characteristics (e.g., average SNR of the link). The routing layer passes this list of next-hop relays to the MAC layer which then performs a rapid poll of each relay and transmits the packet to the "best" relay based on a current channel measurement of each link.[2] Thus, channel-adaptive relaying provides a form of selection diversity against both shadow and multipath channel fading.

The metric by which a relay is evaluated depends on the type of routing protocol and the level of information available. For example, if SNR is available, the forwarding node may select the relay link with maximum SNR. In the context of geographic (or position-based) routing, where position information of each relay and of the final destination is available, the forwarding node may seek to maximize the expected progress, defined as the product of the probability of successful transmission on the link and the progress (in units of distance) offered by that relay towards the final destination. That is, the forwarding node selects the next-hop relay as

$$l^{*} = \arg\max_{1\leq l\leq L} \left\{ Z_l P_s\left(\gamma_l\right) \right\} \tag{1}$$

where Z_l is the progress offered by relay l (i.e, the distance from the forwarding node to the final destination, less the distance from relay l to the final destination), and $P_s\left(\gamma_l\right)$ is the probability of successful transmission on the link to relay l estimated from the measured SNR on that link, γ_l. By taking into account progress as well as link quality in selecting the next-hop relay, a balance is drawn between taking many short, reliable hops versus few long, unreliable hops.

The 802.11-based protocol described in [10] for performing channel-adaptive relay selection at the link layer is based on the concept of MAC-layer anycasting [12, 13]. The concept is illustrated in Fig. 1(a). A new frame is introduced, a *multicast* request-to-send (MRTS), which is identical to the 802.11 RTS but expanded to hold L destination addresses. Each relay in the address list that successfully receives the MRTS, and that is available to receive the data packet, replies with a clear-to-send (CTS). To avoid collisions, the relays reply in the order specified by the MRTS address list. A new CTS frame is introduced which contains three new fields: the address of the relay sending the CTS, the signal-to-interference-and-noise ratio (SINR) measured by the relay upon receipt of the MRTS, and the current geographic coordinates of the relay. The forwarding node records the information contained in each CTS it receives as well as the channel measurement upon receipt of each CTS. With this information, the forwarding node is able to evaluate (1), and its evaluation of P_s can account for the local interference conditions at both ends of the link.

Fig. 1(b) illustrates the timeline of a sample exchange where the MRTS is multicast to three nodes, each one replies with a CTS, and the forwarding node selects the second relay. The timeline also shows the network allocation vector (NAV) states of nearby nodes as a result of the exchange. The overhead of the new frame fields and of the wider channel reservation due to multiple CTS transmissions is discussed in detail in [10].

III. RATE ADAPTATION

This section describes a rate adaptation protocol operating separately from relay selection. Prior to transmission of a data frame, the receiver measures the link state, selects a data rate using a metric defined below, and indicates that selection to the sender. The sender then transmits the data frame at the selected data rate. If the data frame is successfully received, the receiver responds with an acknowledgment (ACK) frame

[1] The approach can be extended to other systems that offer multiple data rates and forwarding options.

[2] Experimental results with commercially available 802.11 cards confirm the usefulness of physical layer channel measurements for link quality assessment [11].

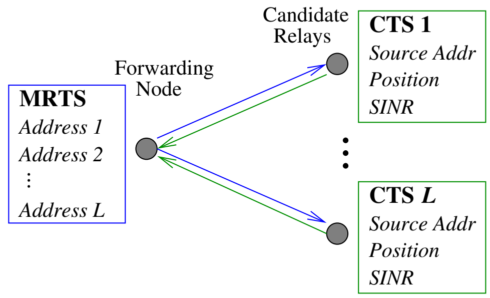

(a) Polling current position and channel state information

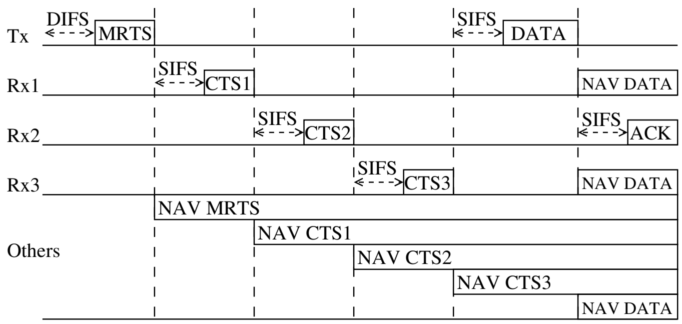

(b) Timeline of polling protocol for three next-hop nodes

Fig. 1. MAC-layer channel-adaptive relay selection

transmitted at the lowest data rate (for added reliability in time-varying channels). The data rate selection is made independently for each data frame transmission or retransmission.

Of the previous work on rate adaptation referenced in Section I, the proposed scheme is most similar to the Receiver-Based AutoRate (RBAR) protocol [6]. In both, the data rate selection is made by the receiver rather than the sender. Furthermore, both schemes measure the received SNR prior to the data transmission through use of the RTS/CTS mechanism. For this reason, the use of these schemes is restricted to systems that employ RTS/CTS, the original intent of which was to reserve the channel at both the sender and receiver for an upcoming data exchange. The tradeoff between the overhead of RTS/CTS and the benefit of channel measurement and reservation is investigated in the performance analysis.

The proposed scheme differs from RBAR, however, in two respects. Whereas a new MAC data frame is introduced to support RBAR, the proposed rate adaptation scheme, when used separately from channel-adaptive relay selection, requires no new frame structure and, therefore, can coexist with existing 802.11 systems. The two also differ in the data rate selection criterion. RBAR selects the data rate using a set of SNR thresholds, while our scheme maximizes a throughput metric.

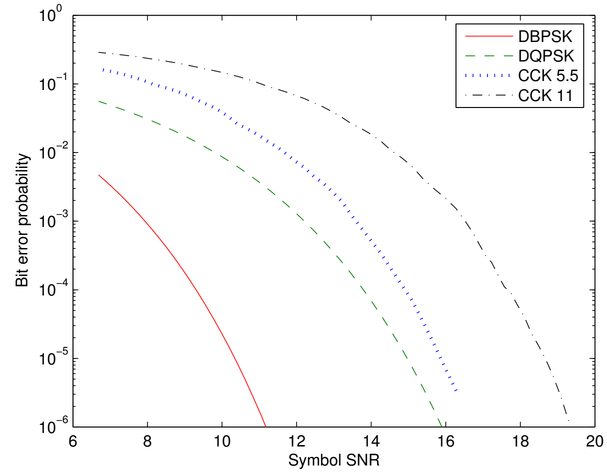

Fig. 2. Bit error rate vs. SNR of 802.11b data rate modes in AWGN

A. Rate Selection

The 802.11b physical layer, for example, makes available four data rate modes: 1 and 2 Mbps using differential binary and quaternary phase shift keying (DBPSK and DQPSK), respectively, and 5.5 and 11 Mbps using two forms of complementary code keying (CCK). Naturally, the higher the data rate is, the higher an SNR is required to achieve a given bit error rate (BER). Fig. 2 illustrates the BER versus the received symbol SNR for these modulation schemes in additive white Gaussian noise (AWGN). The DBPSK and DQPSK curves are from expressions in [14], and CCK performance data can be found in [15].

In the proposed scheme, the sender transmits an RTS at the broadcast data rate (2 Mbps). As specified by 802.11, the sender advertises the duration of the rest of the data exchange, including the ACK, in the *duration* field of the RTS frame. This duration is calculated based on the lowest available data rate and is used by other nodes within listening range of the sender to set their network allocation vectors (NAVs), thereby reserving the channel for the sender's upcoming exchange.

Upon receipt of the RTS, the receiver records the SNR reported by the physical layer. From the duration field, and prior knowledge of the lowest data rate, the receiver calculates the length of the upcoming data frame. With the SNR measurement and frame length known, the receiver evaluates, for each data rate, the expected throughput in packets per second, defined as

$$\lambda_i(\gamma) = \frac{P_{s,i}(\gamma)}{D_i} \tag{2}$$

where γ is the measured SNR, $P_{s,i}(\gamma)$ is the probability of successful transmission at SNR γ and data rate i, and D_i is the duration (in seconds) of a successful data frame and ACK transmission at data rate i. The receiver then selects the data rate that maximizes (2):

$$i^* = \arg\max_i \lambda_i(\gamma). \tag{3}$$

In the case of 802.11b, which uses no forward error correction, the probability of successful transmission is simply the probability that no channel bit errors occur. Moreover, for a data frame to be successfully transmitted after the RTS/CTS exchange has completed, both the data and ACK frames must be received error-free. The probability of successful transmission, then, is approximated by

$$P_{s,i}(\gamma) \cong [1 - P_{b,i}(\gamma)]^N \tag{4}$$

where $P_{b,i}(\gamma)$ is the bit error probability at SNR γ and data rate i, and N is the total number of bits to be transmitted, including the ACK. The righthand side of (4) is an approximation rather than exact for two reasons. First, it assumes bit errors are independent, which may not be the case in practice in the presence of bursty interference. Second, the SNR upon receipt of the data frame may not be the same as that upon receipt of the ACK, due to time-varying channel gain as well as potentially different local interference conditions at the sender and receiver. Nevertheless, (4) provides a first-order approximation of the true success probability that is accurate when the packet duration is less than the coherence time of the channel.

B. Indication of Data Rate Selection to Sender

After the receiver has selected the data rate using (3) and the SNR measured on receipt of the RTS, it needs to indicate that selection to the sender. It does so by advertising the duration of the data/ACK exchange at the selected data rate, D_{i^*}, in the *duration* field of the CTS frame, which is also transmitted at the broadcast data rate. The sender, knowing the data frame length, infers the selected data rate from the value of the advertised duration and proceeds to transmit the data frame at that rate. If it is successfully received, the receiver replies with an ACK at the lowest data rate.

Except when the receiver selects the lowest data rate, the duration advertised by the CTS will be less than that advertised by the RTS. To prevent unnecessarily long channel reservations, listening nodes update their NAVs with the new, shorter value advertised by either the CTS frame or the subsequent data frame of that exchange. It is possible, though not very likely, that a listening node hears only the RTS and not any of the subsequent frames of that exchange, thereby resulting in its NAV being longer than necessary.

IV. **JOINT RELAY SELECTION AND RATE ADAPTATION**

Since channel-adaptive relay selection effectively selects links with favorable channel conditions, we propose marrying it with rate adaptation to achieve more spectrally efficient use of those links. In joint relay selection and rate adaptation, the forwarding node transmits an MRTS to a set of candidate relays as in basic channel-adaptive relaying, at the broadcast data rate. Each relay that receives the MRTS measures the SNR on receipt of the MRTS, determines the data rate that maximizes (2), and advertises the associated duration in its CTS. After waiting for all the possible CTS replies, the forwarding node transmits the data frame to the relay that maximizes the expected progress–throughput product. In terms of the notation above, the sender selects the relay as

$$l^* = \arg\max_{1\le l\le L} \left\{ Z_l \lambda_{i_l^*}(\gamma_l) \right\} \tag{5}$$

where i_l^* is the data rate selected by relay l using (3), and $\lambda_{i_l^*}(\gamma_l)$ is the expected throughput to relay l at that data rate. The joint use of channel-adaptive relaying and SNR-based rate adaptation tends to choose links with favorable channel conditions and progress and then exploits those links with spectrally efficient modulation schemes.

In time-varying channels, the rate adaptation scheme is susceptible to making incorrect rate selections when the channel state changes between the time the SNR measurement is made and the time the data and ACK frames are transmitted. An incorrect decision can be particularly harmful when the SNR decreases sufficiently during that time that a higher data rate transmission is unsuccessful, requiring a retransmission. Joint relay selection and rate adaptation mitigates this risk by using diversity to yield link SNRs with a statistical distribution shifted towards higher values.

V. **LINK PERFORMANCE ANALYSIS**

In this section, we evaluate the link performance of ideal rate adaptation in a quasi-static Rayleigh fading channel. In particular, we evaluate the average MAC-layer throughput on a link, averaged over the fading channel. This analysis indicates the potential gains associated with rate adaptation both with and without the selection diversity made available by adaptive relay selection.

We assume the channel amplitude is fixed during the exchange of the data and ACK frames, and that it is independent and identically Rayleigh distributed from one data transmission to the next. As a result, the SNR, which is proportional to the square of the channel amplitude, has an exponential distribution. Averaging (2) over this distribution, the average throughput using data rate i is

$$\lambda_i = \frac{1}{D_i} \int_0^\infty P_{s,i}(\gamma) f(\gamma)\, d\gamma \tag{6}$$

where $f(\gamma) = \exp(-\gamma/\overline{\gamma})/\overline{\gamma}$ is the density function of the SNR and $\overline{\gamma}$ is the average SNR.

When ideal rate adaptation is used, the sender transmits at the data rate that maximizes (2) for the current SNR. Conditioned on a given SNR, the conditional throughput with rate adaptation is

$$\lambda(\gamma) = \max_i \left\{ \frac{P_{s,i}(\gamma)}{D_i} \right\}$$

and the average throughput with rate adaptation is

$$\lambda = \int_0^\infty \max_i \left\{ \frac{P_{s,i}(\gamma)}{D_i} \right\} f(\gamma)\, d\gamma\,. \tag{7}$$

When relay selection is used, the statistics of the channel improve. Assuming that the channel gain on L available links is independent and identically distributed, and that the link with the maximum channel gain is selected, the distribution function of the selected link gain is

$$\begin{aligned} F(\gamma) &\triangleq \Pr\left[\max\{\Gamma_1, \Gamma_2, \ldots, \Gamma_L\} \le \gamma\right] \\ &= \left(1 - e^{-\gamma/\overline{\gamma}}\right)^L. \end{aligned} \tag{8}$$

Differentiating (8) with respect to γ yields the density function of the channel gain with Lth-order selection diversity:

$$f(\gamma) = \frac{L}{\overline{\gamma}} e^{-\gamma/\overline{\gamma}} \left(1 - e^{-\gamma/\overline{\gamma}}\right)^{L-1}. \tag{9}$$

The average link throughput with relay selection is obtained by using (9) in (6) and (7).

Numerical evaluations of (6) and (7) are shown in Fig. 3 using the 802.11b data rates as an example, and using (4) with the BERs shown in Fig. 2 and $N = 5360$ bits to evaluate $P_{s,i}(\gamma)$. Fig. 3 plots the throughput in packets per second versus the average SNR of a relay link. The throughput of each fixed rate scheme plateaus beyond a certain SNR, with higher rates achieving higher throughput, as expected, but also requiring higher SNR to reach that throughput. Ideal rate adaptation without diversity ($L = 1$) achieves throughput that is at least as high as that of any fixed rate scheme over the entire range of average SNR. Except for extremely low and high SNR, the throughput with rate adaptation is higher than that of the non-adaptive schemes. The reason is that the throughput at a given average SNR represents the average over an ensemble of SNRs, and rate adaptation maximizes the throughput for every SNR in that ensemble.

The results for increasing orders of selection diversity ($L > 1$) show the potential increase in throughput when rate adaptation is combined with adaptive relay selection. They also demonstrate the diminishing marginal returns with the increasing choice of relays, which is characteristic of diversity schemes in general.

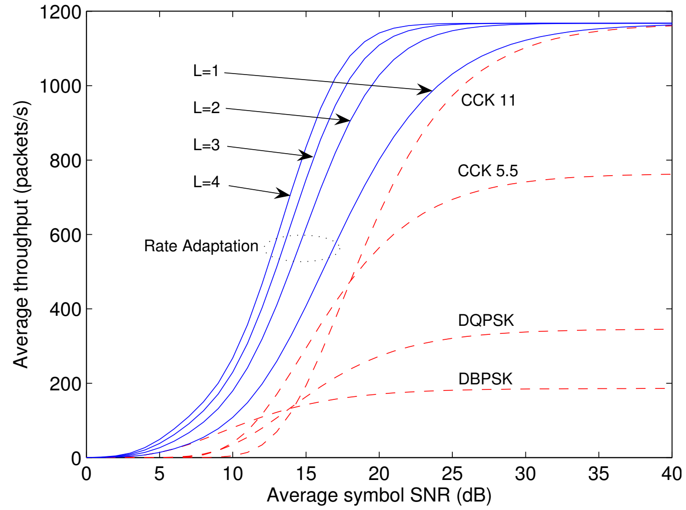

Fig. 3. Link throughput vs. average SNR in quasi-static Rayleigh fading

VI. NETWORK PERFORMANCE ANALYSIS

While the previous section used analytical methods to illustrate the potential benefits in link performance of rate adaptation, with and without link diversity, this section uses simulation results to characterize the end-to-end performance of these adaptive techniques in multihop networks.

The protocols for channel-adaptive relay selection and rate adaptation described in Sections II and III were implemented and evaluated in the QualNet 3.7 simulation environment. Routing is based on the description of Greedy Perimeter Stateless Routing (GPSR) given in [16], with an average beacon interval of 1.5 s. Medium access is based on the IEEE 802.11 DCF with extensions to support channel-adaptive relaying. Relay selection at the routing and MAC layers utilizes the selection rule in (5).

The physical layer is that of 802.11b with all four data rates available. Antennas are omnidirectional, and channel propagation is modeled as a combination of two-ray path loss and time-varying correlated fading. Transmission power is 4.145 dBm, and the receiver sensitivity is -93 dBm, resulting in a transmission range of approximately 300 m at the 1 Mbps data rate. Results for two scenarios are discussed below—FTP traffic on a grid topology and CBR traffic on random, mobile topologies.

A. Grid Topology

A single file transfer protocol (FTP) connection transmits 512-byte packets from a source to a destination on a fixed

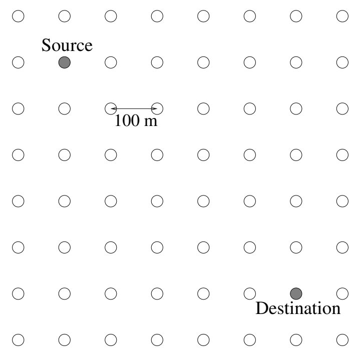

Fig. 4. 8×8 grid topology

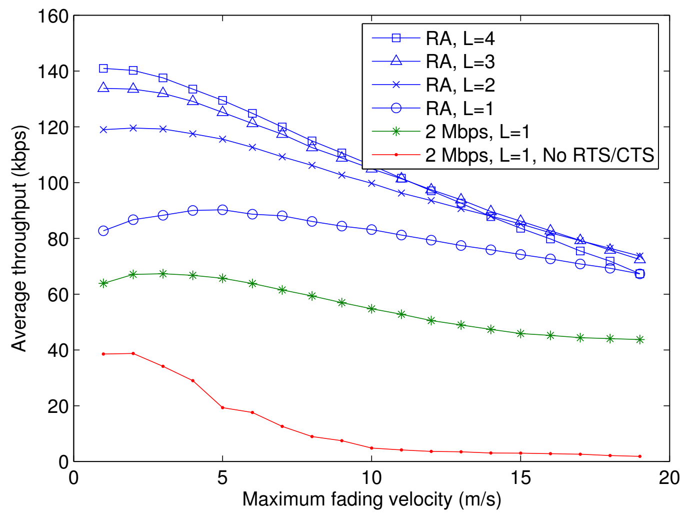

Fig. 5. FTP throughput on 8×8 grid with Rayleigh fading (RA = rate adaptation)

8×8 grid of nodes spaced 100 m apart as illustrated in Fig. 4. Although the nodes are fixed, fading channels for a range of mobility levels are simulated to test the responsiveness of the channel-adaptive protocols to time-varying channels. Each data point is averaged over six independent channel realizations, with each realization lasting 900 s. File transfer starts at $t = 6$ s and ends at $t = 880$ s.

Fig. 5 plots the average end-to-end throughput as a function of channel speed for a Rayleigh fading channel. Channel speed is quantified in terms of the equivalent maximum velocity in meters per second. In all cases, relay selection is employed at the routing layer for adaptivity to slow channel variations. Relay selection at the MAC layer for faster channel adaptivity is indicated by the value of L, the maximum number of relays polled.

We observe that rate adaptation alone (i.e., result "RA, $L = 1$") increases throughput relative to a 2 Mbps fixed rate scheme by 30% or more, depending on the channel speed. When rate adaptation is combined with adaptive relay selection at the MAC layer (the results with $L > 1$), additional gains in throughput are realized, increasing 80% for

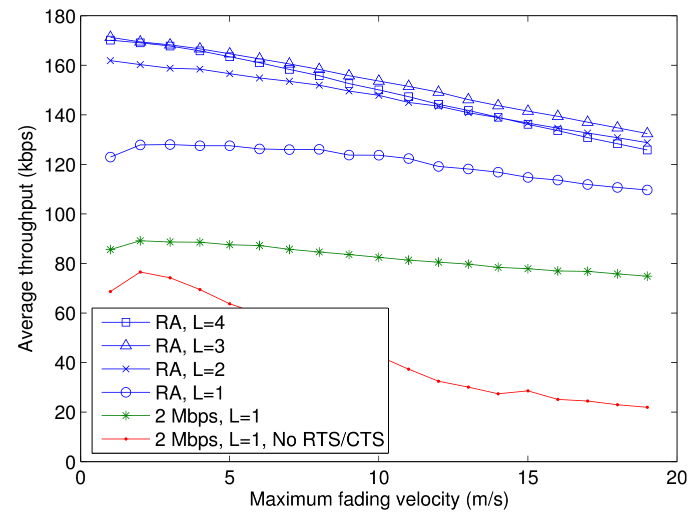

Fig. 6. FTP throughput on 8×8 grid with Ricean fading ($K = 5$ dB)

$L = 2$ and doubling for $L = 3$, relative to the fixed rate scheme. As observed in Section V, most of the additional gain is achieved with 2nd-order diversity, with diminishing marginal returns for $L = 3$ and $L = 4$. At higher channel speeds, breakpoints are observed at which the benefit of additional selectivity is outweighed by the increased overhead and measurement delay incurred by larger L. However, a key difference relative to the fixed rate channel-adaptive relaying results presented in [10] is that, with joint rate adaptation and relay selection, these breakpoints occur at three times the channel speed, meaning that the jointly adaptive system is significantly more robust to time variability of the channel.

Also shown in Fig. 5 are results for standard 802.11b without the use of RTS/CTS for channel reservation. The poor performance of this scheme relative to those which do require RTS/CTS for rapid channel measurement indicates that the gains of channel reservation and adaptivity more than compensate for the associated overhead of RTS/CTS.

Fig. 6 shows analogous results for Ricean fading, which models multipath fading with a specular, line-of-sight component. The Rice factor (ratio of specular power to scattered power) is $K = 5$ dB, here. Because the distribution of the channel gain is more concentrated, the diversity gain with increasing L is less than it is in Rayleigh fading. In fact, we observe no additional gain in going from $L = 3$ to $L = 4$, even at slow speeds, indicating that the incremental diversity gain does not offset the added overhead of polling a fourth relay. As in the case of Rayleigh fading, most of the gain is achieved with 2nd-order link layer diversity.

B. Random, Mobile Topologies

The following results were obtained for mobile networks of 200 nodes with random topology. The initial locations

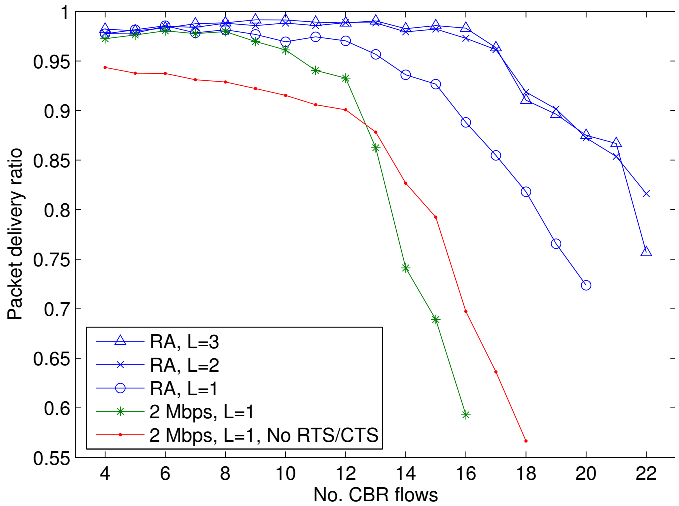

(a) Packet delivery ratio

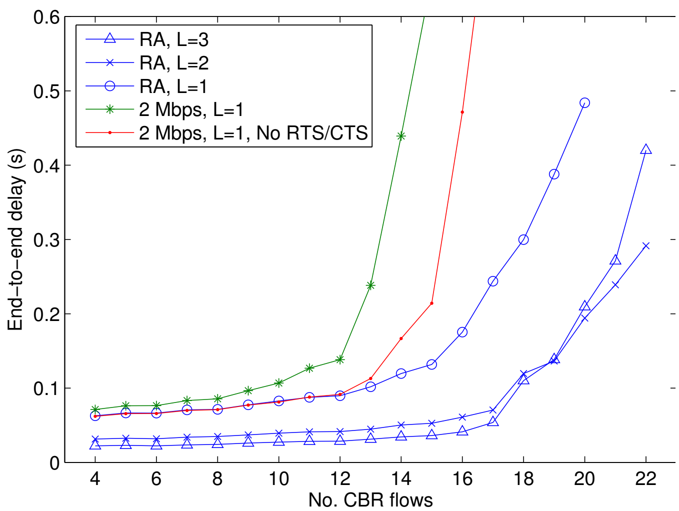

(b) End-to-end delay

Fig. 7. CBR performance as a function of traffic load for random, mobile topologies with Rayleigh fading, 2 m/s maximum velocity

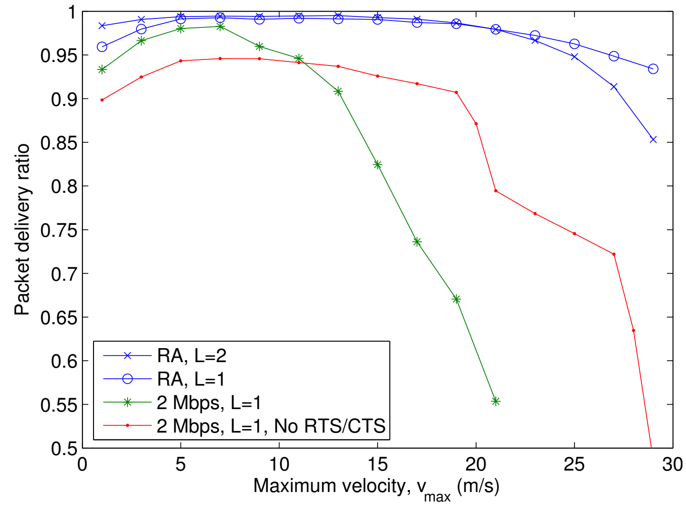

(a) Packet delivery ratio

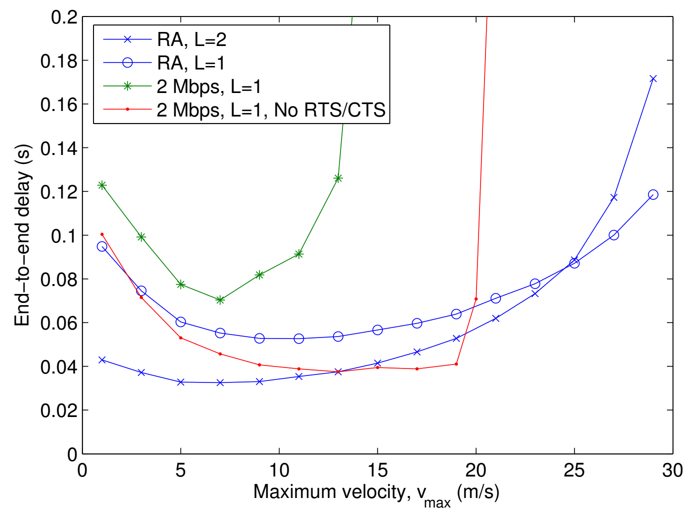

(b) End-to-end delay

Fig. 8. CBR performance as a function of node velocity for random, mobile topologies with Rayleigh fading, 10 CBR flows

of the nodes are randomly, uniformly distributed in a rectangular region of size 3000 m $\times$ 600 m. Node mobility follows the random waypoint model with zero pause time and speeds selected randomly and uniformly in the interval $(0, v_{\max})$. Traffic is generated from multiple constant bit rate (CBR) flows, each generating 64-byte packets at 4 packets/s. Source-destination pairs are mutually exclusive, so that a node is the source or destination of no more than one CBR flow. Each data point is averaged over six independent realizations of a topology/mobility pattern, with each realization lasting 900 s. The start times of the CBR flows are randomly, uniformly distributed on the interval $[6, 180)$ s, and they all end at 880 s.

Fig. 7 illustrates end-to-end performance in terms of the packet delivery ratio and the average end-to-end delay, as a function of the number of CBR flows in the network, for channels with Rayleigh fading and a maximum node velocity of $v_{\max} = 2$ m/s. In this traffic scenario, with shorter 64-byte packets, the overhead associated with channel measurement is more significant relative to the data payload. Nevertheless, the gain realized by rate adaptation (RA, $L = 1$) more than compensates for this overhead, yielding a 25% increase in sustainable CBR traffic at a packet delivery rate of 90%, relative to the fixed rate scheme with or without RTS/CTS. Furthermore, joint rate adaptation and relay selection with only 2nd-order link layer diversity provides a corresponding capacity increase of 58% (19 flows versus 12 flows). The added benefit of 3rd-order diversity ($L = 3$) is negligible in terms of packet delivery rate, but noticeable in terms of delay, reducing delay by 30% at traffic levels up to 16 flows.

Fig. 8 shows results when the traffic load is fixed at 10 CBR flows and the maximum node velocity is varied. These results reveal that while the fixed rate scheme breaks down around 10-15 m/s and the standard non-RTS/CTS scheme around 20 m/s, the rate adaptive schemes are robust for a wider range of node velocity, up to 30 m/s (67 mph). Furthermore, joint rate adaptation and 2nd-order relay selection reduces end-to-end delay at speeds up to 25 m/s (55 mph). We observe, therefore, that channel-adaptive relaying based on measured SNR can be useful even at vehicular speeds.

VII. CONCLUSION

Motivated by the theoretical increase in information capacity afforded by selection diversity in multihop relaying systems, a protocol was specified for joint relay adaptation and relay selection in a multihop 802.11 ad hoc network using geographic routing and RTS/CTS at the MAC layer. At each hop and for each packet transmission, a channel measurement is made and the data rate and next-hop relay are selected to maximize a progress–throughput metric for that transmission. Results of network simulations demonstrate that the proposed rate adaptation improves performance relative to a fixed rate scheme which does not use RTS/CTS, and that joint rate and relay adaptation provides synergistic gains in throughput of delay-insensitive traffic and in capacity and delay of CBR traffic. The jointly adaptive scheme was also robust to time-varying channels corresponding to vehicular speeds.

While results were given here for a system using geographic routing, channel-adaptive relaying can be applied to non-geographic routing protocols that also provide multiple forwarding options, such as Ad Hoc On-Demand Multipath Distance Vector (AOMDV) routing [17].

ACKNOWLEDGMENT

The authors would like to thank Marc Heissenbüttel of the Institute of Computer Science and Applied Mathematics, University of Bern, Switzerland, for providing the basic GPSR source code which we extended for this work.

REFERENCES

[1] H. Lundgren, E. Nordström, and C. Tschudin, "Coping with communication gray zones in IEEE 802.11b based ad hoc networks," in *Proc. 5th ACM Int. Workshop on Wireless Mobile Multimedia*, 2002, pp. 49–55.

[2] A. J. Goldsmith and S.-G. Chua, "Variable-rate variable-power MQAM for fading channels," *IEEE Trans. Commun.*, vol. 45, no. 10, pp. 1218–1230, Oct. 1997.

[3] X. Qiu and K. Chawla, "On the performance of adaptive modulation in cellular systems," *IEEE Trans. Commun.*, vol. 47, no. 6, pp. 884–895, June 1999.

[4] A. Kamerman and L. Monteban, "WaveLAN-II: A high performance wireless LAN for the unlicensed band," *Bell Labs Technical J.*, pp. 118–133, Summer 1997.

[5] D. Qiao and S. Choi, "Fast-responsive link adaptation for IEEE 802.11 WLANs," in *Proc. ICC*, vol. 5, May 2005, pp. 3583–3588.

[6] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive MAC protocol for multi-hop wireless networks," in *Proc. MobiCom*, 2001, pp. 236–251.

[7] D. Qiao, S. Choi, and K. G. Shin, "Goodput analysis and link adaptation or IEEE 802.11a wireless LANs," *IEEE Trans. Mobile Computing*, vol. 1, no. 4, pp. 278–292, Oct.-Dec. 2002.

[8] J. del Prado Pavon and S. Choi, "Link adaptation strategy for IEEE 802.11 WLAN via received signal strength measurement," in *Proc. ICC*, vol. 2, May 2003, pp. 1108–1113.

[9] S. Biswas and R. Morris, "ExOR: Opportunistic multi-hop routing for wireless networks," in *Proc. ACM Special Interest Group on Data Communication (SIGCOMM)*, 2005, pp. 133–143.

[10] M. R. Souryal and N. Moayeri, "Channel-adaptive relaying in mobile ad hoc networks with fading," in *Proc. Second Annual IEEE Commun. Society Conf. on Sensor and Ad Hoc Commun. and Networks (SECON)*, Sept. 2005.

[11] M. R. Souryal, L. Klein-Berndt, L. E. Miller, and N. Moayeri, "Link assessment in an indoor 802.11 network," in *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, Apr. 2006.

[12] R. R. Choudhury and N. H. Vaidya, "MAC-layer anycasting in ad hoc networks," in *Proc. 2nd workshop on Hot Topics in Networks (HotNets II)*, 2003.

[13] S. Jain and S. R. Das, "Exploiting path diversity in the link layer in wireless ad hoc networks," in *Proc. 6th IEEE Int. Symp. on a World of Wireless Mobile and Multimedia Networks (WoWMoM)*, June 2005, pp. 22–30.

[14] J. G. Proakis, *Digital Communications*. 4th ed., New York: McGraw-Hill, 2001.

[15] C. Andren and M. Webster, "CCK modulation delivers 11Mbps for high rate IEEE 802.11 extension," in *Proc. Wireless Symposium/Portable by Design Conference*, Spring 1999.

[16] B. Karp and H. T. Kung, "GPSR: Greedy perimeter stateless routing for wireless networks," in *Proc. MobiCom*, 2000, pp. 243–254.

[17] M. K. Marina and S. R. Das, "On-demand multipath distance vector routing in ad hoc networks," in *Proc. IEEE Int. Conf. for Network Protocols*, 2001, pp. 14–23.