

Opportunistic Media Access for Multirate Ad Hoc Networks *

B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly
Department of Electrical and Computer Engineering
Rice University
Houston, TX 77005
{bahar,kanodia,ashu,knightly}@rice.edu

ABSTRACT

The IEEE 802.11 wireless media access standard supports multiple data rates at the physical layer. Moreover, various auto rate adaptation mechanisms at the medium access layer have been proposed to utilize this multi-rate capability by automatically adapting the transmission rate to best match the channel conditions. In this paper, we introduce the Opportunistic Auto Rate (OAR) protocol to better exploit durations of high-quality channels conditions. The key mechanism of the OAR protocol is to opportunistically send multiple back-to-back data packets whenever the channel quality is good. As channel coherence times typically exceed multiple packet transmission times for both mobile and non-mobile users, OAR achieves significant throughput gains as compared to state-of-the-art auto-rate adaptation mechanisms. Moreover, over longer time scales, OAR ensures that all nodes are granted channel access for the same time-shares as achieved by single-rate IEEE 802.11. We describe mechanisms to implement OAR on top of any existing auto-rate adaptation scheme in a nearly IEEE 802.11 compliant manner. We also analytically study OAR and characterize the gains in throughput as a function of the channel conditions. Finally, we perform an extensive set of *ns-2* simulations to study the impact of such factors as node velocity, channel conditions, and topology on the throughput of OAR.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless communication*

General Terms

Design, Performance

*This research is supported by NSF Grants ANI-9979465 and ANI-0085842, a Texas Advanced Technology Program Grant, a Sloan Fellowship, and by a gift from Texas Instruments.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MOBICOM'02, September 23–26, 2002, Atlanta, Georgia, USA.
Copyright 2002 ACM 1-58113-486-X/02/0009 ...\$5.00.

Keywords

IEEE 802.11, Scheduling, Medium Access, Wireless Channels, *Ad Hoc* Networks

1. INTRODUCTION

The IEEE 802.11a and 802.11b media access protocols provide a physical-layer multi-rate capability [13]. With the original IEEE 802.11 protocol, all transmission takes place at a single base rate, typically 2 Mbps. With the multi-rate enhancement, transmission can take place at a number of rates according to channel conditions. Higher data rates than the base rate are possible when the signal-to-noise ratio (SNR) is sufficiently high such that channel-resiliency demands of error correcting codes and modulation schemes can be relaxed. Consequently, with IEEE 802.11a the set of possible data rates is 6, 9, 12, 18, ..., 54 Mbps whereas for IEEE 802.11b the set of possible data rates is 1, 2, 5.5, and 11 Mbps.

As the multi-rate IEEE 802.11 enhancements are physical layer protocols, MAC mechanisms are required to exploit this capability. The Auto Rate Fallback (ARF) protocol [6] was the first commercial implementation of a MAC that utilizes this feature. With ARF, senders attempt to use higher transmission rates after consecutive transmission successes (which indicate high channel quality) and revert to lower rates after failures. Under most channel conditions, ARF provides a performance gain over pure single-rate IEEE 802.11.

In [4], a protocol termed Receiver Based Auto Rate (RBAR) is proposed. The core idea of RBAR is for *receivers* to measure the channel quality using physical-layer analysis of the request-to-send (RTS) message. Receivers then set the transmission rate for each packet according to the highest feasible value allowed by the channel conditions. As the RTS message is sent shortly before data transmission, the estimation of the channel condition is quite accurate, so that RBAR yields significant throughput gains as compared to ARF (as well as compared to single-rate IEEE 802.11). Moreover, as request- and clear-to-send messages are necessarily sent at the base rate so that all nodes can overhear them, overhearing nodes are informed of the modified data-transmission times so that they can set their backoff timers accordingly.

In practice, the channel quality of wireless hosts can vary significantly, both for mobile and stationary nodes [16]. In particular, a node's received signal is composed of both a

line-of-sight component as well as numerous delayed signals reflected off of surrounding objects. All such signals (plus random noise) combine and must be decoded by the receiver. In favorable cases, these signals add *coherently* in a way that enhances the channel quality (increases received power, increases signal-to-noise ratio, etc.). However, in unfavorable cases, these signals will tend to cancel each other out and lead to a poor-quality channel. Even for stationary nodes, any change in the line-of-sight path or any reflected path (e.g., due to an intermediary pedestrian) will change the quality of the channel and hence, change the data rate that is feasible with multi-rate IEEE 802.11.

In this paper, we introduce Opportunistic Auto Rate (OAR), an enhanced protocol for multi-rate IEEE 802.11 in wireless ad hoc networks. The key idea of OAR is to opportunistically exploit high quality channels when they occur via transmission of multiple back-to-back packets. In particular, when the multi-rate MAC indicates that the channel quality allows transmission above the base rate, OAR grants channel access for multiple packet transmissions in proportion to the ratio of the achievable data rate over the base rate. Consequently, OAR nodes transmit more packets under high quality channels than under low quality channels. However, OAR cannot arbitrarily favor flows with the best channel quality, as access for flows with perhaps perpetually bad channels must also be ensured. Consequently, we ensure that all flows are granted the same *temporal-share* of channel access under OAR as under single-rate IEEE 802.11.

This assurance of achieving the same time shares as single-rate IEEE 802.11 implies that OAR inherits the long-time-scale temporal fairness properties of IEEE 802.11. Thus, for scenarios in which all nodes are within radio range of each other, OAR can provide flows with dramatically different *throughputs* as dictated by their channel conditions, but all flows will achieve approximately identical time shares. For example, with OAR, a flow with an average channel condition providing a data rate of 11 Mbps will obtain approximately 5 times the throughput of a flow with an average channel condition providing a data rate of 2 Mbps, while the two flows will both access the channel approximately 50% of the time. The distinction between time-share fairness and throughput fairness is critical in multi-rate networks. Throughput fairness has received a great deal of attention in both base-station and ad hoc networks (e.g., [9, 12] and [10, 11, 17, 18] respectively) and indeed, throughput and temporal fairness are equivalent in single-rate networks. However, normalizing flow throughputs in a multi-rate network would result in significant inefficiency and mitigate the gains of the multi-rate physical layer, as poor-channel flows would consume disproportionately more time and system resources. In more complex topologies, difficulties in achieving *equal* flow time shares are well established [10, 11, 17, 18]. With OAR, we do not attempt to equalize temporal shares but instead extract an opportunistic throughput gain while maintaining the same temporal shares provided by single-rate IEEE 802.11.

We utilize two mechanisms to realize the OAR protocol. First, OAR requires a multi-rate MAC protocol such as RBAR or ARF to access the medium at rates above the base rate. While OAR can be applied to both sender- and receiver-based protocols, we consider RBAR as it results in significantly higher throughput gains than ARF as described above. Second, OAR requires a mechanism to

hold the channel for an extended number of packet transmissions when a high-rate channel is provided by RBAR. We show how use of the IEEE 802.11 mandated fragmentation field provides a simple mechanism for a sender to maintain channel access for the extended opportunistic duration dictated by OAR. While alternate mechanisms are also possible for multi-packet transmission, we find that RBAR together with the fragmentation mechanism yields the highest performance solution.

To study the performance of OAR, we develop an analytical model that characterizes the throughput gains of OAR and RBAR as compared to IEEE 802.11 as a function of the physical-layer channel conditions. The key technique is to relate the contention times in IEEE 802.11 to those of OAR and RBAR and to relate physical layer parameters such as path loss, distance, and background noise, to the data-rate achievable by multi-rate IEEE 802.11. While necessarily highly simplified, the model provides a key link between media access performance and physical layer conditions, in contrast to previous work, which abstracts the channel as a constant-rate medium (e.g., [2]).

Finally, we perform an extensive simulation study to evaluate OAR in realistic scenarios and to isolate the performance factors that determine OAR's throughput gains. In particular, we study the role of node density, node velocity, channel conditions, node location distributions, network topology, and interactions with TCP. Example findings are as follows. (1) In most cases, we find that OAR achieves throughput gains of 40% to 50% over RBAR, on top of the substantial gains that RBAR achieves as compared to ARF and single-rate IEEE 802.11. (2) The throughput gain of OAR as compared to RBAR increases with the number of nodes as OAR also reduces contention. (3) The total throughput gain of OAR increases nearly linearly with the line-of-sight component of the Ricean fading channel, with a slope greater than that achieved by RBAR. (4) User mobility affects gains primarily via the user location distribution with mobility speeds in the range expected by pedestrians. (5) Small bursts due to OAR's consecutive packet transmissions do not adversely affect TCP with even moderate node buffer sizes.

In cellular networks, opportunistic scheduling refers to a base station's selection of the best-channel-node for packet transmission [3, 8]. OAR is quite different from such techniques as only the centralized control of a base station can allow selection among users using measurements of all channels. In contrast, OAR applies to ad hoc networks as well as infrastructure-based networks with access points as no centralization is required. Instead, nodes transmit multiple packets when channels are good and a single packet when only the base rate is available. Regardless, if additional information is available about other nodes' channel conditions, OAR could be extended to also temporarily defer access when other nodes have better channels as advocated in [3, 8].

The remainder of this paper is organized as follows. First, in Section 2, we describe the state-of-the-art wireless channel model, and review IEEE 802.11 and multi-rate RBAR and ARF. Next, in Section 3, we present the OAR protocol. We present an analytical model of OAR in Section 4 and present the results of simulation experiments in Section 5, and conclude in Section 6.

2. BACKGROUND AND MOTIVATION

2.1 Channel Model

The transmitted radio frequency signal is reflected by both natural and man-made objects. Thus, the signal at the receiver is a superposition of different reflections of the same signal, received with varying delays and attenuations. Based on the relative phases of different reflections at the receiver, the different copies of the same signal may add coherently or tend to cancel out. Coherent addition of the copies can result in large received signal powers and cancellation eventually leads to zero received signal power. The signal power strength is heavily dependent on the spatial location of the transmitter, receiver, the reflecting objects and the material of the reflecting objects.

Any motion, be that of the surrounding environment, the sender, or the receiver, causes the strength of the received signal to vary with time. The speed of variation is directly governed by the speed of change in the communication medium (which consists of all intermediate objects in the direct and reflected space). Physical layer design and analysis typically consider detailed propagation models that characterize all reflections and their time-variations [14]. An accurate and widely utilized model which considers time-varying multi-path propagation [16] is¹

$$y(t) = \sum_{i=1}^{p(t)} A_i(t)x(t - \tau_i(t)) + z(t), \quad (1)$$

where $x(t)$ is the transmitted signal and $y(t)$ is the received signal. The time-varying multi-path propagation is captured by the attenuation of each path $A_i(t)$, the time delays $\tau_i(t)$ and the number of paths $p(t)$. The additive term $z(t)$ is generally labeled as the background noise and represents the thermal noise of the receiver. Note that the loss suffered by the signal during its propagation along different paths is captured in $A_i(t)$, and depends on the distance between the sender and the receiver.

Typically, physical layer algorithms (error correcting codes, channel modulation, demodulation and decoding) use the elaborate models in Equation (1). The performance of any physical layer implementation is well captured by observing its packet loss rate as a function of the received signal to noise ratio (SNR). Received signal to noise ratio measures the extent of the received signal power over the channel background noise. Typically, the larger the SNR, the better the chance of any packet being received error free. Actual performance (packet loss rate as a function of SNR) is dependent on a particular implementation.²

Recognizing that the received SNR can be used to capture the packet level performance of *any* physical layer implementation, we use the following model for the received signal to noise ratio for transmitter power P at packet transmission time t_p

$$SNR(t_p) = Pd(t_p)^{-\beta} \frac{\rho(t_p)}{\sigma^2}, \quad (2)$$

where $d(t_p)$ is the distance between the sender and the receiver at time t_p , β is the path loss exponent, $\rho(t_p)$ is the

¹This is a baseband model which assumes perfect carrier demodulation at the receiver radio frequency front-end.

²For example, we have found in controlled laboratory tests with hardware-emulated channel conditions, that 802.11b compliant cards from different manufacturers perform differently under identical channel conditions.

average channel gain for the packet at time t_p , and σ^2 is the variance of the background noise $z(t)$.

The short time-scale variation in the received SNR is captured by the time-varying parameter $\rho(t_p)$, known as the *fast fading* component of the fading process. The time-variation of $\rho(t_p)$ is typically modeled by a probability distribution and its rate of change [16]. An accurate and commonly used distribution for $\rho(\cdot)$ is the Ricean distribution,

$$p(\rho) = \frac{\rho}{\sigma^2} e^{-\left(\frac{\rho}{2\sigma^2} + K\right)} I_0(2K\rho), \quad (3)$$

where K is the distribution parameter representing the strength of the line of the sight component of the received signal and $I_0(\cdot)$ is the modified Bessel function of the first kind and zero-order [16]. For $K = 0$, the Ricean distribution reduces to the Rayleigh distribution, in which there is no-line-of-sight component.

The rate of change of $\rho(t_p)$ depends on a mobile host's relative speed with respect to its surroundings. Among the several models available in the literature we use the Clarke and Gans model [16]. (Also see [16] for a survey.) The motion of nodes causes a Doppler shift in the frequency of the received signal, and the extent of the Doppler shift depends on the relative velocity of the sender and the receiver. Let f_m denote the maximum Doppler frequency during the communication between the two nodes. Then according to the Clarke-Gans model, the received signal is modulated in the frequency domain by the following spectrum

$$S(f) = \frac{1.5}{\pi f_m b \sqrt{1 - \left(\frac{f-f_c}{f_m}\right)^2}}. \quad (4)$$

In Equation (4), f_c represents the carrier frequency of the transmitted signal.³ The spectral shape of the Doppler spectrum in Equation (4) determines the time domain fading waveform and hence the temporal correlation. The inverse of the maximum Doppler frequency of f_m , $T_c = \frac{1}{f_m}$, is known as the *coherence interval*, and represents the average time of decorrelation. In essence, the channel SNR values $\rho(\cdot)$ separated by more than T_c , are approximately independent. At mobile speeds of 1 m/s (3.6 km/hr), the coherence interval is approximately 122.88 ms for a center frequency of 2.4 GHz. The coherence interval reduces to 24.57 ms, 12.28 ms and 6.14 ms for mobile speeds of 5 m/s (18 km/hr), 10 m/s (36 km/hr) and 20 m/s (72 km/hr). In engineering design [16], a more conservative estimate of coherence interval is used which is around 43% of the above numbers: 51.98 ms, 10.39 ms, 5.20 ms and 2.59 ms for speeds of 1, 5, 10 and 20 m/s. At 2, 5.5 and 11 Mbps, a 1000 byte packet takes 4 ms, 1.45 ms and 0.73 ms. *The fact that coherence intervals are on the order of multiple packet transmission times provides a key motivating factor for designing opportunistic scheduling policies.*

Finally, we note that in practice, the channels between any sender and receiver pair are independent of other sender-receiver pairs. Namely, received signals are correlated only if nodes are extremely close to each other: for a 2.4 GHz central frequency, the critical distance is less than 12 cm [5].

2.2 Review of IEEE 802.11

Here, we briefly review the IEEE 802.11 Distributed Coordinated Function (DCF) which serves as a starting point

³For 802.11b, carrier frequency is in 2.4-2.481 GHz range.

both for development of the RBAR [4] and OAR protocols.

As described in [13], IEEE 802.11 media access is based on the RTS/CTS mechanism [1]. In particular, a transmitting node must first sense an idle channel for a time period of Distributed InterFrame Spacing (DIFS) after which it generates a random backoff timer chosen uniformly from the range $[0, w - 1]$, where w is referred to as the contention window. At the first transmission attempt, w is set to CW_{min} (minimum contention window). After the backoff timer reaches 0, the node transmits a short request-to-send (RTS) message. If successfully received, the receiving node responds with a clear-to-send (CTS) message. Any other node which hears either the RTS or CTS packet uses the data packet length information to update its Network Allocation Vector (NAV) containing the information of the period for which the channel will remain busy. Thus, all nodes including hidden node can defer transmission suitably to avoid collision. Finally, a binary exponential backoff scheme is used such that after each unsuccessful transmission, the value of w is doubled, up to the maximum value $CW_{max} = 2^m CW_{min}$, where m is the number of unsuccessful transmission attempts.

2.3 Multi-rate IEEE 802.11

The IEEE 802.11a and IEEE 802.11b protocols are *multi-rate* in that they provide physical-layer mechanisms to transmit at higher rates than the base rate if channel conditions permit.⁴ Figure 1 shows a sample channel variation with time for a mobile speed of 2.5 m/s. The received power shows wide fluctuations such that the supported data rate varies between 5.5 and 2 Mbps with almost equal probability. In practice, depending on the line-of-sight factor K in (3) and the distance from the transmitter d in (2), the channel rates can vary within the entire range of the lowest to highest possible data rate.⁵ As noted earlier, the highest available rate in IEEE 802.11a is 54 Mbps and 11 Mbps for IEEE 802.11b.




Figure 1: Illustration of channel condition variation

The first commercial implementation that exploits this multi-rate capability is termed Auto Rate Fallback (ARF)

⁴IEEE 802.11a is based on Orthogonal Frequency Division Multiplexing in the 5 GHz band whereas IEEE 802.11b is based on Direct Sequence Spread Spectrum in the 2.4 GHz band.

⁵Achievable data rates as a function of distance for 802.11a are available in a white paper from <http://www.atheros.com>. For 802.11b, we use the specifications for the OrinocoTM wireless NIC which can be found at <http://www.orinocowireless.com>

[6]. With ARF, senders use the history of previous transmission error rates to adaptively select future (attempted) transmission rates. That is, after a number of consecutive successful transmissions, the sender changes its modulation scheme to attempt transmission at a higher rate, and vice versa after consecutive losses. Consequently, if a mobile user has (for example) a perpetually high-quality channel, the user will eventually transmit at higher data rates while accessing the medium according to the same IEEE 802.11 MAC as described above.

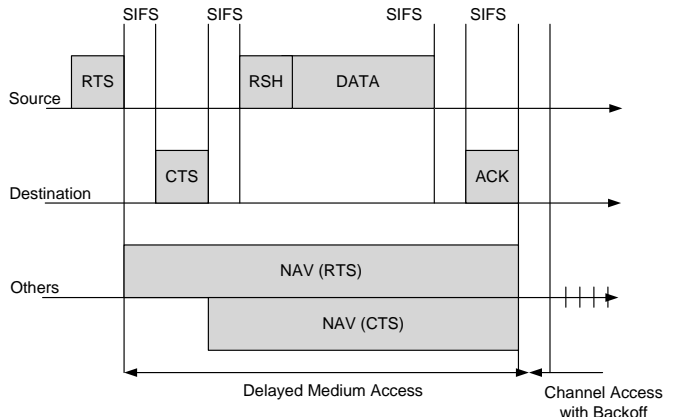


Figure 2: IEEE 802.11 with RBAR enhancement

An enhanced protocol to exploit the multi-rate capabilities of IEEE 802.11a termed Receiver Based Auto Rate (RBAR) was proposed in [4]. The key idea of RBAR is for receivers to control the sender's transmission rate. In IEEE 802.11a, all RTS/CTS messages must be sent at the base rate to ensure that all stations are able to receive these messages error free. RBAR uses physical-layer analysis of the received RTS message to determine the maximum possible transmission rate for a particular bit error rate. The receiver inserts this rate into a special field of the CTS message to inform the sender and other overhearing nodes of the potentially modified rate. This message is termed reservation-sub-header (RSH) and is inserted preceding data transfer as illustrated in Figure 2. With the RSH message, overhearing nodes can modify their NAV values to the new potentially decreased transmission time. In this way, RBAR quickly adapts to channel variations and extracts significant throughput gains as compared to ARF.

3. THE OAR PROTOCOL

In this section we present the Opportunistic Auto Rate (OAR) MAC protocol. The key observation is that channel coherence times (durations for which mobile stations have better-than-average channels) are typically at least multiple packet transmission times. Consequently, when a mobile user is granted channel access while encountering a high-quality channel, OAR grants the user a channel access time that allows multiple packet transmissions. As the subsequent packet transmissions are also highly likely to be successful at the higher data rate, OAR obtains a throughput gain as compared to RBAR and ARF. However, OAR must also limit the extent to which it is opportunistic in order to ensure that users with perpetually bad channels obtain their fair share of time accessing the channel. Hence,

OAR is opportunistic to the largest extent possible while maintaining the temporal fairness properties of the base-rate IEEE 802.11 protocol.

The OAR protocol is designed as an enhancement to any automatic rate adaptation protocol for IEEE 802.11. We next describe OAR in the context of RBAR, and below discuss alternate realizations.

3.1 Protocol

The key OAR mechanism is for a flow to keep the channel for an extended number of packets (instead of for a single packet) once the channel is measured to be of sufficient quality to allow transmission at rates higher than the base rate.

With RBAR, the receiver measures the channel quality and transmits the channel transmission rate in the CTS and RSH messages. As all overhearing nodes can decode the CTS and RSH messages sent at the base rates, they can appropriately set their NAV values to the potentially shorter transmission times due to sending a fixed amount of data at a higher rate.

With OAR, the same *time* is granted to a sender as if the sender is transmitting at the base rate. For example, if the base rate is 2 Mbps and the channel condition is measured such that transmission at 11 Mbps is feasible, the sender is granted a channel access time to send $\lfloor 11/2 \rfloor = 5$ packets. Thus, OAR inherits the same temporal fairness properties as the IEEE 802.11 base-rate protocol.⁶

In a fully connected topology in which all nodes are in radio range of each other, base-rate IEEE 802.11 indeed provides long-term fairness. With OAR, if all nodes are in radio range, they will still obtain identical long-term time shares but will obtain quite different throughputs. For example, suppose there are two flows, one with a low SNR such that it can only transmit at the base rate of 2 Mbps and the other with a high SNR so that it can transmit at rate 54 Mbps. If the channel conditions remain identical, OAR will provide equal time shares to the two flows but the latter flow will obtain 26 times the throughput of the former flow. Thus, in contrast to the focus on *throughput* fairness of [9, 12] which attempt to normalize flow throughputs, temporal fairness is more suitable for multi-rate networks as normalizing flow throughputs would cancel the throughput gains available due to a multi-rate physical layer.

More formally, to make the distinction between temporal and throughput fairness, we define $\alpha_i(t_1, t_2)$ as the service in *time* that flow i receives during (t_1, t_2) . The measure of fairness for OAR with equal flow weights is therefore

$$|\alpha_i(t_1, t_2) - \alpha_j(t_1, t_2)|. \quad (5)$$

Moreover, denoting $\alpha_i^B(t_1, t_2)$ as the time share for flow i using base-rate IEEE 802.11 and $\alpha_i^{\text{OAR}}(t_1, t_2)$ as the time share under OAR, we also consider the fairness index

$$|\alpha_i^B(t_1, t_2) - \alpha_i^{\text{OAR}}(t_1, t_2)| \quad (6)$$

as the deviation of OAR time shares from IEEE 802.11 base-rate time shares. As ensuring that Equation (5) is close to

⁶As described in the Introduction, it is well established that IEEE 802.11 ad hoc networks will not in general attain global fairness and equal flow time shares. As solving the general fairness MAC problem is beyond the scope of this work, we match the IEEE 802.11 time shares. Enhancements to OAR to provide global fairness along the lines of [10, 11, 17, 18] is an important topic for future work.

0 requires sophisticated mechanisms such as [10, 11, 17, 18], we target ensuring that the fairness measure of Equation (6) is close to 0 to ensure that OAR provides near base-rate time shares.

We next describe a mechanism for *how* to hold the channel for these extended times using the fragmentation field in IEEE 802.11.

3.2 Fragmentation in IEEE 802.11

Figure 3 illustrates the IEEE 802.11 time-line for transmission of a fragmented data packet. Each frame contains information that defines the duration of the next transmission. The duration information from the RTS/CTS frame is used to update the NAV to indicate that the channel is busy until the end of ACK 0. Both Fragment 0 and ACK 0 also contain information to update the NAV to indicate a busy channel until the end of ACK 1. This continues until the last fragment which carries the duration of one ACK time plus one SIFS time in its duration field. The ACK for the last fragment has the duration field set to zero. Each fragment/ACK acts as a virtual RTS/CTS so that no further RTS/CTS frames are used. Also each fragment (except the last one) has the *more_fragments_indicator* in the MAC header set to 1. The MAC header of each fragment also carries the fragment number which is used by the destination to reassemble the data packet. For an un-fragmented data packet this is set to 0. Also note that each subsequent fragment is sent after a time interval SIFS after receiving the ACK (instead of the DIFS period for data).

3.3 Integration of OAR with IEEE 802.11 via the Fragmentation Field

The fragmentation mechanism defined as part of the IEEE 802.11 standard and mandated for implementation provides a simple and practical way for nodes to hold the channel for multiple packets when high data rates are measured. In particular, as in RBAR, the receiver indicates the available physical-layer rate via the RTS/CTS messages. If the data rate is above the base rate, the *more_fragments* flag in *frame control* Field of MAC header is set by the sender until $\lfloor \text{transmission_rate} / \text{base_rate} \rfloor$ packets are transmitted. The duration field carried by the data and the subsequent ACK is also updated to indicate that the medium will be busy until the end of next data packet. Thus, as described above, each data packet and ACK serve as a virtual RTS/CTS so that no additional RTS/CTS frames need to be generated after the initial RTS/CTS handshake.

The sender must also set the *fragment number* subfield in the *sequence control* field of the MAC header to be 0. This prevents the receiver from treating the data packet as a part of an actual fragmented packet.

3.4 Traffic and Channel Conditions

Two issues arise while using the OAR protocol. The first issue revolves around the behavior of the OAR protocol in case there are no data packets available in the interface queue⁷ to exploit the high-rate channel. In this case, the protocol reverts back to the default RBAR protocol and sends as many packets as are available subject to the upper bound imposed by the above time-share criterion. Since each data packet carries the “more fragments” indicator for

⁷The queue between the link layer and MAC layer.




Figure 3: Illustration of fragmentation in IEEE 802.11

only the next data packet, the sender can revert back to the default RBAR at any time without any loss in throughput resulting from incorrect settings of other nodes' NAVs.

Second, while coherence times are on *average* sufficiently long to allow multiple packet transmissions at the same rate, it is possible that the channel condition will significantly change during the opportunistic multi-packet-transmission sequence. If transmission at the original rate is maintained, error rates may become unacceptably large if the channel quality worsens, or further throughput gains are foregone if the channel quality improves.

To address this issue, OAR receivers continually monitor the channel quality, and if a significant change is detected, additional RSH messages are used to notify the receiver and adapt the rate. The subsequent transmission rate can be increased or decreased provided that the flow has not exceeded its base-rate time share.

3.5 Example

A representative example of the protocol time-lines for OAR and RBAR is shown in Figure 4. In the example, node 1 has a good channel resulting in a 11 Mbps rate. Node 2 has a poor channel which can only support the base-rate of 5.5 Mbps. Node 1 gains access for both OAR and RBAR, which results in a packet transmission at the highest rate. In OAR, node 1 retains channel access and sends four more packets without any channel contention. On the other hand, RBAR goes into contention immediately after the transmission of node 1 is completed. The backoff counter for node 2 was frozen while node 1 was transmitting, and node 1 picks a new backoff counter after finishing its transmission. Thus, with high probability, node 2 gets access to the channel and sends its data at 2 Mbps. Thus, RBAR loses in throughput due to extra contention after every packet in addition to not capitalizing fully when good channels are encountered.

If node 2 had gained access first (instead of node 1), then again OAR would have gained over RBAR. In this case, whenever node 1 gets access (possibly after node 2 is done transmitting), it will transmit 5 packets since the channel will still be in a good state with high probability. But RBAR would have again continued to contend for each packet, losing time in both contention and transmission at low rates.

3.6 Alternate Implementation Mechanisms

In the discussion above, OAR is described using RBAR rate adaptation and fragmentation to hold the channel.

OAR can also be applied to sender-based protocols. For example, with an ARF-like protocol, senders control the transmission rate. In such cases, senders can also hold the channel for extended durations as defined by OAR. However, the throughput gains of OAR for sender-based protocols would be reduced since, as demonstrated in [4], receiver-based control of transmission rates results in a significant increase in throughput as compared to sender-based control.

Second, alternate mechanisms other than the fragmentation field could be employed for senders to hold the channel for the extended durations defined by OAR. For example, after a successful transmission above the base-rate, a sender can set its contention window to 0 making the sender highly likely to re-capture the channel. This process can be repeated for the number of packet transmissions defined by OAR. Compared to use of fragmentation, such an approach would result in higher overhead due to the additional RTS/CTS messages, and potential collisions if other senders have low values of the contention window (recall that contention windows are generated uniformly random between 0 and 64 slots). The other alternate mechanism is the use of *packet bursting*, which is a measure introduced in IEEE 802.11e⁸ to better utilize the medium and enhance the performance [7]. With packet bursting, a station is allowed to transmit as many as frames it wishes without contending for the medium again after accessing the channel, as long as the total access time does not exceed a certain limit.

4. A PERFORMANCE MODEL FOR OPPORTUNISTIC MEDIA ACCESS

In this section, we analyze the throughput of OAR and RBAR for fully connected topologies. Without channel variations, nodes can hear transmissions from other nodes if they are within a critical distance from each other.⁹ With channel variations, the path losses are time-varying so that nodes are not guaranteed to hear all transmissions of other nodes, even if they are within the critical distance. Regardless, in the discussion below, we consider a (deterministically) fully connected topology in which all nodes can hear one another.

4.1 Contention

By sending more packets when the channel has high SNR, OAR reduces the average time spent in contention per packet

⁸An extension to IEEE 802.11 with goal of enhancing the access mechanisms and providing service differentiation

⁹250 meters is a value typically used in *ns-2*




Figure 4: Illustration of OAR and RBAR Time-lines for a two node system. Node 1 is in a better channel state than node 2 (11 versus 5.5 Mbps).

as compared to RBAR. We denote this gain as the *opportunistic contention gain*. The following observation quantifies this gain.

OBSERVATION 1 (CONTENTION TIME). *Let t_{co} represent the average time (in seconds) per packet spent in contention for the single-rate IEEE 802.11 protocol (operating at 2 Mbps). Denote the total time spent in contention (in seconds) by single-rate IEEE 802.11 for an experiment spanning T seconds by T_{co} . Then for a system with constantly backlogged queues and equal-size packets, the following statements hold:*

1. *The time spent in contention per packet by RBAR is exactly equal to t_{co} .*
2. *The total time in contention by OAR is approximately equal to T_{co} .*

PROOF: Both observations regarding the contention time rely on the fact that both RBAR and OAR use the same contention resolution mechanism as single-rate IEEE 802.11 at transmission rate 2 Mbps.

1. In RBAR, the system state during contention is exactly the same as that in base-rate IEEE 802.11. That is, when nodes are contending to gain access to the channel, the system undergoes the same states as base-rate IEEE 802.11. Since the contention mechanism also operates at the same data rate in RBAR and single-rate IEEE 802.11, the time spent in contention for any packet is identical for the two protocols. The only difference between the two protocols is the time used to send the actual data. Since the local timers of all nodes are frozen during data transmission (since all nodes can hear each other), the system state during contention is unaffected by the packet transmission rates.
2. In OAR, the total packet transmission times for the three data rates are approximately equal. In other words, the times taken to send 5 packets at 11 Mbps, 3 packets at 5.5 Mbps, and 1 packet at 2 Mbps are approximately the same. Again, the contention resolution operates at the base-rate of 2 Mbps as in base-rate IEEE 802.11. Thus, if we were to operate OAR and base-rate IEEE 802.11 in parallel, we would observe that the times at which nodes contend for the channel are nearly identical. Thus, the total time spent in contention in OAR is the same as that for base-rate IEEE 802.11. ■

In RBAR, the total time spent in transmitting any packet is the same as for single-rate IEEE 802.11. Thus, the gain in throughput can be completely attributed to receiver-based rate adaptation based on channel conditions. Thus, the

gain in RBAR throughput is due to reducing data transmission time and not due to reducing contention time. On the other hand, in OAR gains are also extracted from reducing contention time since OAR does not contend for every packet in good channel conditions. Moreover, by reducing contention and transmitting multiple packets at instants of high-quality channels, OAR obtains further throughput gains. The claims made in Observation 1 are verified using *ns-2* simulations, and a sample output is summarized in Table 1.

Table 1: Verification of Observation 1 using $n/2$ single hop flows for a simulation time of 25 seconds.

n	t_{co}		T_{co}	
	RBAR	IEEE 802.11	OAR	IEEE 802.11
8	8.05×10^{-4}	8.20×10^{-4}	3.45	3.82
20	8.30×10^{-4}	8.50×10^{-4}	3.50	3.94
40	8.45×10^{-4}	8.66×10^{-4}	3.63	4.00

4.2 Throughput

Given the channel model of Equation (2) and its distribution of Equation (3), the distribution of available data rates can be calculated as follows. Let SNR_2 , $\text{SNR}_{5.5}$ and SNR_{11} denote the minimum required SNR to support 2, 5.5 and 11 Mbps transmission rates, respectively. Then the probability that rate R is feasible is calculated as follows:

$$\begin{aligned}
 p(R = 0) &= p(\text{SNR} < \text{SNR}_2) \\
 p(R = 2) &= p(\text{SNR}_2 \leq \text{SNR} < \text{SNR}_{5.5}) \\
 p(R = 5.5) &= p(\text{SNR}_{5.5} \leq \text{SNR} < \text{SNR}_{11}) \\
 p(R = 11) &= p(\text{SNR}_{5.5} \leq \text{SNR})
 \end{aligned}$$

where $p(\text{SNR}) = p(\rho\sigma^2 d(t_p)^\beta P^{-1})$ is the distribution of the received SNR, $\text{SNR}(t_p)$. Using Observation 1, the throughput of OAR and RBAR can be characterized in terms of t_{co} and T_{co} as follows. For the rest of this section, *RTS*, *CTS*, *DIFS*, *SIFS*, *ACK*, *PHY-HDR* and *MAC-HDR* represent the respective time required (in seconds) for RTS, CTS, DIFS, SIFS, ACK, physical-layer header transmission, and MAC-layer header transmission.

RESULT 1 (FULLY CONNECTED TOPOLOGIES). *Let $p(r)$ denote the probability distribution of the instantaneous transmission rate r obtained when a node gains access to the channel. Further assuming that all sender-receiver pairs have statistically identical and independent channels, denote the average transmission rate of any node by $r_{avg} = \sum_i r_i p(r_i)$. Let t_{co} and T_{co} respectively denote the contention time per packet and the total contention time for n nodes over T seconds. Also let T_o represent the overhead time spent per packet equal to $(\text{DIFS} + 3\text{SIFS} + 2\text{PHY-HDR} + 2\text{MAC-HDR} + \text{ACK})$. Then the number of packets*

transmitted by RBAR and OAR over T seconds is approximately:

$$\eta_{\text{RBAR}} \approx \frac{r_{\text{avg}}T}{L + r_{\text{avg}}t_{\text{co}} + r_{\text{avg}}T_o}, \quad (7)$$

$$\eta_{\text{OAR}} \approx \frac{r_{\text{avg}}(T - T_{\text{co}})}{L + r_{\text{avg}}T_o}, \quad (8)$$

where L is the length of the packet.

DERIVATION: To derive η_{RBAR} and η_{OAR} , we use the following fact in addition to Observation 1. For fully connected topologies in base-rate IEEE 802.11, the probability of channel access for any node is $\frac{1}{n}$. For a large number of nodes, the probability that a node can access the channel for k consecutive transmissions is approximately $\frac{1}{n^k}$. Thus, for large node densities, we can assume that any node can capture channel for multiple consecutive transmissions with very low probability. Therefore, with very high probability, subsequent packet transmissions by the same node see approximately independent channels; this is especially true if n times average packet transmission time is on the order of T_c . As we also verified the above observation by extensive simulations, we do not account for the time correlation in the channel in the following derivation.

For RBAR, since the contention time per packet is the same as for single-rate IEEE 802.11, the total throughput in bits, $L\eta_{\text{RBAR}}$, is given by

$$L\eta_{\text{RBAR}} \approx r_{\text{avg}}(T - t_{\text{co}}\eta_{\text{RBAR}} - T_o\eta_{\text{RBAR}}). \quad (9)$$

Solving (9) for η_{RBAR} gives Equation (7).

For OAR, since the total contention time is approximately T_{co} , the amount of useful data time is $(T - T_{\text{co}} - \eta_{\text{OAR}}T_o)$. Thus, Equation (8) follows. ■

To compute t_{co} and T_{co} for base-rate IEEE 802.11, we use the following results from [2],

RESULT 2 (802.11 CONTENTION TIME [2]). *Let CW_{\min} denote the minimum contention window size (in number of slots) and assume that the binary exponential back-off doubles its contention window size a maximum of m times, i.e., $CW_{\max} = 2^m CW_{\min}$. Then, for a fully connected topology with n nodes,*

$$\begin{aligned} T_{\text{co}} &= (T - \eta T_p), \\ t_{\text{co}} &= \frac{T_{\text{co}}}{\eta}. \end{aligned}$$

The utilization η for an average packet length of L bits is calculated as follows

$$\begin{aligned} \eta &= \frac{P_s P_{tr} L}{(1 - P_{tr})S + P_{tr} P_s T_s + P_{tr}(1 - P_s)T_c} \\ P_{tr} &= 1 - (1 - \tau)^n \\ P_s &= \frac{n\tau(1 - \tau)^{n-1}}{P_{tr}} \\ T_c &= \text{RTS} + \text{DIFS} + \delta \\ T_s &= T_c + 3 \text{SIFS} + 3\delta + \text{CTS} + \frac{L}{2 \times 10^6} + \text{ACK} \end{aligned}$$

where τ is obtained by solving the two equations $\tau = \frac{2(1-2p)}{(1-2p)(CW_{\min}+1)+pCW_{\min}(1-(2p)^m)}$ and $p = 1 - (1 - \tau)^{n-1}$. The parameter δ represents the propagation time and S is the slot duration. The average packet time T_p is calculated using

$$T_p = \text{PHY-HDR} + \text{MAC-HDR} + \text{ACK} + \frac{L}{2 \times 10^6}.$$

Table 2: Comparison of simulation and analysis for OAR ($T=25$ seconds).

n	Simulation		Analysis	
	T_{co} (sec)	Rate (Mbps)	T_{co}	Rate
8	3.45	6.02	3.53	6.00
20	3.50	6.00	3.51	6.00
40	3.63	5.96	3.59	5.98

Table 3: Comparison of simulation and analysis for RBAR ($T=25$ seconds).

n	Simulation		Analysis	
	t_{co} (sec)	Rate (Mbps)	t_{co}	Rate
8	8.05×10^{-4}	4.26	7.47×10^{-4}	4.09
20	8.30×10^{-4}	4.24	7.42×10^{-4}	4.03
40	8.45×10^{-4}	4.20	7.61×10^{-4}	4.01

Contention times and throughputs of OAR and RBAR calculated using ns-2 simulations and Result 2 are shown in Tables 2 and 3. In this simulation, the distance between sender-receiver pairs is 50 m, and thus all channels are statistically identical but independent. In this example, the results from theory and simulation matched well for the case of OAR. For RBAR, the predicted throughputs were less than those obtained in the simulations. The primary reason for that difference stems from ignoring the coherence time in the analysis, and for RBAR, its impact is more significant compared to OAR. For a small number of nodes and low mobility, nodes can gain access to the channel more than once in the *same* coherence interval, a fact which is ignored in our analysis. In our analysis, the nodes had a good channel (11 or 5.5 Mbps) with high probability, thus RBAR simulation performance is consistently better than prediction. We conjecture that as the number of nodes becomes larger, our analysis will be more accurate for RBAR since the approximations used in Result 1 become more exact.

4.3 Discussion

Typically, the analysis of CSMA/CA protocols for wireless channels abstracts the physical link as a constant-rate channel [2, and references therein]. In contrast, we have included the channel variations in the analysis of the MAC protocol. The difficulty in modeling such systems stems from the fact that both the MAC protocol and the channel are random processes with memory. In the IEEE 802.11 MAC, the contention resolution protocol has memory (nodes resume their counters while waiting for other nodes to finish transmitting, instead of resetting them). That results in requiring a sophisticated mathematical analysis (see [2] for example). In our case, in addition to the MAC memory, the channel is slowly varying and has strong correlation in time. To complicate the analysis further, OAR uses the channel correlation to increase throughput, but in turn, changes the memory structure of the MAC protocol (especially true for arbitrary multi-hop topologies).

In light of such complexities of the throughput analysis, we note the importance of Observation 1. Though deceptively simple, Observation 1 clearly identifies the source of gain in throughput for OAR over RBAR. In addition, we show how both RBAR and OAR perform in relation to the base-rate IEEE 802.11 and show how they obtain their respective gains. This comparison to IEEE 802.11 allows

us to use any available analysis or simulation for base-rate IEEE 802.11 to compute RBAR and OAR throughputs. We also note that the result regarding the total contention time for OAR in Observation 1 holds for arbitrary topologies, as well as for fully connected topologies, since the packet times in OAR have been normalized to those of base-rate IEEE 802.11. However, the actual throughputs of RBAR and OAR are topology dependent, and are tractable only in fully connected topologies that we considered above.

5. SIMULATION EXPERIMENTS

In this section, we use *ns-2* simulations to evaluate the performance of OAR as compared to RBAR and base-rate IEEE 802.11. Our methodology is to isolate the impact of each performance factor to the largest extent possible and then consider more complex scenarios to study the joint effects of numerous factors. We begin with a fully connected topology and study the effects of node density, channel conditions, mobility, and node location. We then consider more general topologies consisting of a simple asymmetric topology and more complex random topologies. Finally, we consider interactions between a multi-rate MAC-layer and TCP for multi-hop flows.

All experiments use the fast fading model of Equation (2). In particular, we use the Ricean probability density (3) implemented in the *ns-2* extension [15]. In [15], a packet level simulation is used to model the short time-scale fading phenomenon using the procedure suggested in [16]. A pre-computed lookup table containing the components of a time-sequenced fading envelope are modulated in frequency domain using the Doppler spectrum in Equation (4). Although the *ns-2* extensions implemented in [15] result in an accurate simulation of the wireless channel for each individual flow, the fading components of channels for different flows are *identical*, a scenario not encountered in practice. This arises due to the fact that the index into the pre-computed channel table is chosen based on the simulator's time instant, which is identical for all flows. Thus, to more realistically model the wireless channel for multiple users in a manner consistent with [16], we modified the extensions of [15] such that channel lookup indexes are a function of the flow and time. As in [15], background noise is modeled with $\sigma = 1$.

The available rates for both RBAR and OAR, based on IEEE 802.11b, are set to 2 Mbps, 5.5 Mbps, and 11 Mbps, so that with OAR, nodes can respectively transmit 1, 3, or 5 consecutive packets depending on their channel condition. The values for received power thresholds for different data rates were chosen based on the distance ranges specified in the OrinocoTM802.11b card data sheet.¹⁰ As specified by the IEEE 802.11 standard, we set the rate for sending physical-layer headers to 1 Mbps for all packets. Unless otherwise noted, each transmitter generates constant-rate traffic such that all nodes are continuously backlogged. Moreover, packet sizes are set to 1000 bytes and all reported results are averages over multiple 50-second simulations.

5.1 Fully Connected Topologies

Here, we study the performance factors for opportunistic scheduling in fully connected topologies in which all nodes are within radio range of each other.

¹⁰For the path loss component of the received power, the distance thresholds for 11 Mbps, 5.5 Mbps, and 2 Mbps are 100 m, 200 m, and 250 m respectively.

5.1.1 Number of Flows

In this set of experiments, we vary the number of flows which also varies the number of nodes, as each flow is between a unique source-destination pair of nodes. Figure 5 shows the throughput gain of OAR over RBAR for different number of flows for a Ricean parameter $K = 5$.

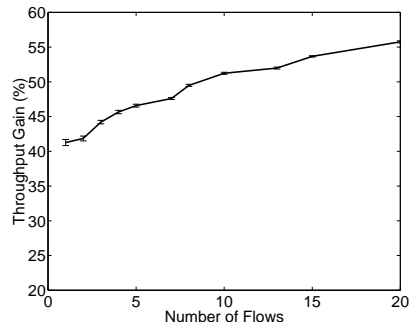


Figure 5: OAR throughput gain as a function of the number of flows.

Observe that in all cases, OAR results in significant throughput gains as compared to RBAR in the range of 42% to 56%. OAR extracts this gain by holding the channel when it is good to the longest extent possible subject to maintaining the base-rate time shares. Moreover, observe that as the number of flows increases, the throughput gain of OAR as compared to RBAR also increases. This increase is due to two factors. First, with a higher number of contending flows, the fraction of time that a flow with a good channel is accessing the medium is higher, resulting in increasing gains. That is, since flows send a single packet at the base rate, but up to 5 consecutive packets at higher rates, additional users provide more opportunities for a flow to be in a good-channel state thereby extracting further gains from opportunistic scheduling. Second, OAR reduces the contention overhead since nodes with a good channel exploit the channel for transmission of additional packets. Since IEEE 802.11 as well as RBAR have increasing contention times for an increasing number of nodes, OAR extracts increased throughput gains with a larger number of nodes by decreasing contention time.

Finally, we note that for both RBAR and OAR, the throughput gains as compared to base-rate IEEE 802.11 are significant. For example, for 10 flows, RBAR obtains a gain of 230% throughput as compared to IEEE 802.11, and OAR obtains an additional throughput gain of 51% as compared to RBAR, or 398% above IEEE 802.11.

5.1.2 Line-of-Sight Component K

Here, we explore the effect of the Ricean parameter K on the performance of OAR and RBAR. For $K = 0$, the channel has no line-of-sight component such that only reflected signals are received and hence, overall channel quality is poor. With increasing K , the line-of-sight component is stronger such that the overall channel SNR increases as described by Equation (3), and a higher transmission rate is feasible more often.

Figure 6 depicts the throughput for RBAR and OAR as a function of the Ricean parameter K for a four-flow scenario. Observe that both RBAR and OAR exploit the improved channel conditions represented with increasing K and obtain

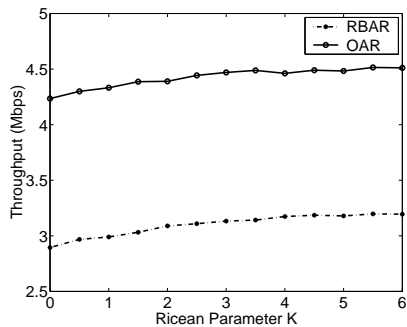


Figure 6: Throughput as a function of the Ricean parameter K .

correspondingly greater system-wide throughputs. Moreover, note that OAR achieves a higher throughput compared to RBAR over the simulated range of K due to its enhanced exploitation of high-quality channel conditions when they occur, as well as OAR's corresponding reductions in the total contention time.

5.1.3 Mobility

A node's mobility affects its channel in two ways. First, it changes the nodes' location which affects a pair of nodes' line-of-sight Ricean parameter K . Second, it affects the average channel coherence time as a node with higher velocity has a lower average coherence time hindering the ability to exploit opportunistic scheduling.

To study the effect of mobility on OAR and RBAR, we consider a single flow, where the receiver node travels to and from the transmitter node with different velocities in an oscillatory fashion as described in *Configuration 1* of [4] with a small difference: the simulation time is long enough for any speed of the node so that it has enough time to finish the path at least once. This ensures that for all velocities, the *average* time spent at each distance is independent of the node velocity, and thus only speed of channel variations is the determining factor.

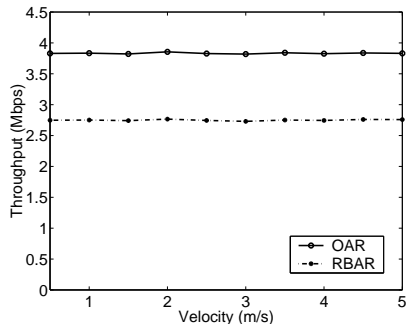


Figure 7: Throughput as a function of mobile speed.

The throughputs for OAR and RBAR are depicted in Figure 7 for speeds up to 5 m/s. As described in Section 2, this corresponds to an average coherence time of approximately 10 msec, which corresponds to slightly larger than 2 packet transmission times at the base rate of 2 Mbps and 13 packet transmission times at 11 Mbps. As shown, OAR and RBAR have a throughput that is nearly independent of velocity with OAR having approximately 35% greater throughput than RBAR. The key reason is that within this range of

velocities, the coherence time is sufficiently large to extract the full performance gains of both RBAR and OAR. While [4] did show more pronounced effects of velocity, the experiments of [4] jointly consider the effects of location and mobility, as nodes did not spend the same fraction of time in each location independent of velocity.

5.1.4 Location Distribution

For the final experiments in a fully connected topology, we study the effect of the node location distribution on the throughput of OAR and RBAR. We consider a scenario representing a room such that there are N flows altogether and a subset $N' \leq N$ located in the room which has area $1/100^{th}$ of the total area of approximately 0.2 km². Within such a small room, the N' flows are highly likely to have high data-rate channels.




Figure 8: Throughput as a function of node clustering density.

Figure 8 depicts the results for $N = 10$. The figure depicts the OAR and RBAR throughputs as a function of the density of flows located in the room N'/N . As shown as the nodes become increasingly clustered into a small location, the channel conditions improve and both OAR and RBAR obtain increased throughputs, with OAR's remaining approximately 50% above RBAR.

Finally, we note that in all experiments, the long-term temporal shares of all flows are nearly identical for the OAR protocol, since IEEE 802.11 base-rate media access is temporally fair at long-time-scales when all nodes are within radio range and are equally likely to obtain channel errors. We next turn to more complex and asymmetric topologies.

5.2 Asymmetric Topologies

In systems with topologies that are not fully connected, i.e., all nodes are *not* within range of each other, nodes can have different probability of channel capture due to one node hearing an RTS or CTS that another node does not hear. This unequal channel access probability can result in large differences in the time shares of channel access among nodes. Such behaviors are due to asymmetry in information available to each flow and are well documented in the context of IEEE 802.11 [1].

An illustrative example of asymmetric information among nodes is depicted in Figure 9, in which the receiver of Flow A (node 2) is in direct radio range of Flow B, whereas the sender (node 1) has no knowledge of Flow B. In the scenario of Figure 9, Flow B obtains a significantly higher share of the channel access time as compared to Flow A, namely 80% vs. 20% when using IEEE 802.11. This disparity in total share is attributed to the fact that Flow B can hear packets from




Figure 9: Asymmetric topology

the receiver of Flow A, and hence knows exactly when to contend for the channel. On the other hand, the transmitter of Flow A does not hear any packets from Flow B, and thus has to discover an available time-slot randomly; hence Flow A continually attempts to gain access to the channel via repeated RTS requests which in most cases result in doubling of Flow A’s contention window. As a result, the probability of Flow A capturing the channel is significantly less than that of Flow B.

In this section we show that in general topologies, even with asymmetric information, OAR will still have throughput gain over RBAR and at the same time complies with pure IEEE 802.11 in sense of channel access time shares.

In the experiments for Figure 10, we consider four scenarios for the topology of Figure 9. The scenarios from the left to right pairs of bars depict four combinations of high and low-quality channels denoted by H and L. A high quality channel is most often 11 Mbps whereas a low quality channel is most often 2 Mbps due to small and large distances among nodes. In particular, the four cases for channel conditions for Flow A and Flow B are given by [L,L], [H,L], [L,H], and [H,H] respectively. The y-axis of Figure 10(a) depicts the normalized deviation from the single-rate IEEE 802.11 temporal shares as in Equation (6) for the total duration of the simulation. The y-axis of Figure 10(b) depicts the throughput gain as compared to single-rate IEEE 802.11. Thus, the left-most bar illustrates that when both channel conditions are low quality denoted by [L,L], both OAR and RBAR provide the same temporal shares as single-rate IEEE 802.11 (20% and 80% respectively) while achieving throughput gains of approximately 60% and 35% respectively. However, note that in other scenarios, RBAR deviates significantly from the IEEE 802.11 time shares whereas in OAR does not. Yet, in all cases, OAR is able to extract significant throughput gains as compared to both RBAR and IEEE 802.11 while maintaining the IEEE 802.11 time shares.

5.3 Random Topologies

Here, we consider random topologies which combine the previous effects. In particular, we consider a scenario in which nodes are uniformly distributed in a circular arena with diameter of 1250 meters. Moreover, the sender and receiver of one flow are moving across the arena during the simulation. Figure 11 shows the throughput of both OAR and RBAR for different numbers of nodes in the arena. Observe that the gains of OAR as compared to RBAR remain in the range of 40% to 50% as in previous experiments, despite the more complex scenario.

5.4 Interactions with TCP

Our final experiments address the issue of whether the throughput gains due to opportunistic scheduling at the MAC layer can be exploited by TCP at the transport layer. On one hand, OAR provides increased throughput which

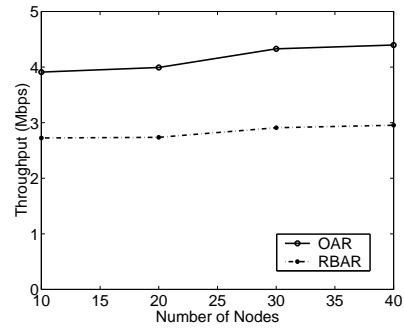


Figure 11: OAR and RBAR throughputs for random topologies.

a backlogged TCP flow should be able to exploit. On the other hand OAR’s transmission of multiple packets back-to-back in cases of good channel conditions could also lead to increased packet drops at downstream nodes if insufficient buffer space is available.




Figure 12: Topology for multi-hop TCP experiments

In these experiments, we explore the net effect of these two factors using the topology of Figure 12. In this scenario, there are 15 2-hop flows between nodes 1 through 15 and node B. Each flow represents a large file transfer using TCP Reno. Nodes A and B are sufficiently far apart so that the path loss coupled with channel variation results in a data rate of 2 Mbps most of the time. In contrast, the distances between nodes 1 through 15 and node A are sufficiently small such that a data transmission rate of 11 Mbps is feasible most of the time.

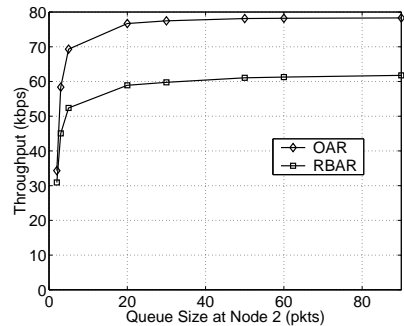

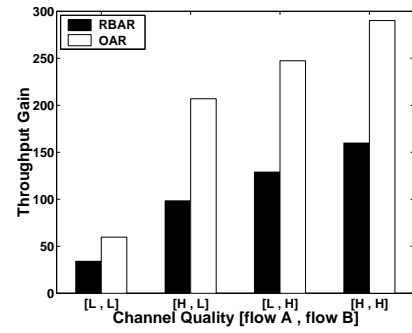


Figure 13: Throughput as a function of queue size.

Figure 13 shows the average throughput obtained for both RBAR and OAR as a function of the queue size in packets of the intermediate node A. Observe that for queue sizes of approximately 5 packets, TCP/OAR and TCP/RBAR obtain throughputs nearly identical to that achieved with a very large buffer size. Thus, as OAR can transmit up to 5



(a) Time Share



(b) Throughput Gain

Figure 10: OAR and RBAR time share deviations and throughput gains from IEEE 802.11.

packets back-to-back, it is critical for the intermediate node to be able to transmit at least this many packets. However, as RBAR also requires a similar buffer size to obtain its full throughput gain, it is clear that the buffer requirement is primarily due to the mis-match in average channel rates as opposed to burstiness introduced by the MAC layer. Finally, we observe that with buffer sizes above 5 packets, OAR obtains a throughput gain of approximately 30% as compared to RBAR.

6. CONCLUSIONS

In this paper, we introduced OAR, an opportunistic media access protocol for multi-rate ad hoc networks. With OAR, nodes with good channel conditions are granted access to the channel for a duration that allows multiple packet transmissions vs. a single packet when nodes access the channel at the base rate. Consequently, by exploiting inherent variations in channel conditions, nodes will transmit more data during epochs of high-quality channels thereby increasing the total throughput of all users. Moreover, OAR ensures that all nodes, regardless of their channel condition, access the channel for a time-share equal to that achieved with single-rate IEEE 802.11. To study the performance gains of OAR, we developed an analytical model that characterizes the impact of channel conditions on the throughput obtained by both OAR and RBAR. Finally, we performed an extensive set of *ns-2* simulations with realistic channel conditions and found that OAR obtains throughput gains of up to 50% as compared to state-of-the-art auto-rate protocols.

7. REFERENCES

- [1] V. Bharghavan, S. Demers, S. Shenker, and L. Zhang. MACAW: A media access protocol for wireless LANs. In *Proceedings of ACM SIGCOMM '94*, London, UK, 1994.
- [2] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, Mar. 2000.
- [3] S. Borst and P. Whiting. Dynamic rate control algorithms for HDR throughput optimization. In *Proceedings of IEEE INFOCOM '01*, Anchorage, Alaska, Apr. 2001.
- [4] G. Holland, N. Vaidya, and P. Bahl. A rate-adaptive MAC protocol for multi-hop wireless networks. In *Proceedings of ACM MOBICOM'01*, Rome, Italy, 2001.
- [5] W. C. Jakes. *Microwave Mobile Communications*. Wiley, 1974.
- [6] A. Kamerman and L. Monteban. WaveLAN II: A high-performance wireless LAN for the unlicensed band. *Bell Labs Technical Journal*, pages 118–133, Summer 1997.
- [7] A. Lindgren, A. Almquist, and O. Schelén. Evaluation of Quality of Service schemes for IEEE 802.11 wireless LANs. In *Proceedings of the 26th Annual IEEE Conference on Local Computer Networks (LCN 2001)*, November 2001.
- [8] X. Liu, E. Chong, and N. Shroff. Transmission scheduling for efficient wireless network utilization. In *Proceedings of IEEE INFOCOM '01*, Anchorage, Alaska, Apr. 2001.
- [9] S. Lu, V. Bharghavan, and R. Srikant. Fair scheduling in wireless packet networks. *IEEE/ACM Transactions on Networking*, 7(4):473–489, Aug. 1999.
- [10] H. Luo, S. Lu, and V. Bharghavan. A new model for packet scheduling in multihop wireless networks. In *Proceedings of ACM MOBICOM'00*, Boston, MA, Aug. 2000.
- [11] H. Luo, P. Medvedev, J. Cheng, and S. Lu. A self-coordinating approach to distributed fair queueing in Ad Hoc wireless networks. In *Proceedings of IEEE INFOCOM'01*, Apr. 2001.
- [12] T. Ng, I. Stoica, and H. Zhang. Packet fair queueing algorithms for wireless networks with location dependent errors. In *Proceedings of IEEE INFOCOM'98*, San Francisco, CA, May 1998.
- [13] B. O'Hara and A. Petrick. *IEEE 802.11 Handbook, A Designer's Companion*. IEEE Press, 1999.
- [14] J. G. Proakis. *Digital Communications*. McGraw Hill, 1995.
- [15] R. Punnoose, P. Nikitin, and D. Stancil. Efficient simulation of Ricean fading within a packet simulator. In *IEEE Vehicular Technology Conference*, pages 764–767, 2000.
- [16] T. S. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, 1999.
- [17] N. Vaidya and P. Bahl. Fair scheduling in broadcast environments, Aug. 1999. Microsoft Research Tech. Rep. MSR-TR-99-61.
- [18] N. Vaidya, P. Bahl, and S. Gupta. Distributed fair scheduling in a wireless LAN. In *Proceedings of ACM MOBICOM'00*, Boston, MA, Aug. 2000.