

Cross-Layer Exploitation of MAC Layer Diversity in Wireless Networks

Min Cao, Vivek Raghunathan and P. R. Kumar

Abstract—The conventional function of the medium access control (MAC) layer in wireless networks is interference management. We show how the MAC can also be used to mitigate the effect of fading. We begin by providing experimental data to demonstrate that multipath fading effects are seen at the MAC layer. These effects appear at timescales on the same order of the IEEE 802.11 protocol and therefore, interact negatively with the RTS-CTS-DATA-ACK handshake. We identify two types of MAC diversities to jointly combat fading and interference, called *multi-receiver diversity* and *multi-channel diversity* respectively, through canonical scenarios. In order to harness these MAC layer diversities, we propose a simple *dynamic-binding multi-channel MAC* (DB-MCMAC) protocol that is backward compatible with IEEE 802.11. DB-MCMAC exploits MAC diversities by opportunistically acquiring the floor for the best receiver on each channel, and dynamically binding data transmissions after the floor has been acquired. We employ a simple continuous time Markov chain model to analyze the expected performance of the DB-MCMAC protocol. A comprehensive performance evaluation of DB-MCMAC using ns-2 shows that it can successfully harness multi-receiver and multi-channel fading and interference diversities to provide considerable improvements over a baseline multi-channel MAC in several situations.

Index Terms—protocol design, MAC diversity, multi-channel MAC, fading, measurements, performance analysis

I. INTRODUCTION

In wireless networks, the communication channel is time-varying and shared among geographically distributed users. This leads to two fundamental challenges in protocol design: 1) *Fading*: Fading is the time variation of the wireless channel due to two effects. First, large-scale path loss and shadowing effects cause the signal to attenuate with distance. Second, multipath scattering effects result in delayed copies of the signal adding up constructively or destructively at the receiver, causing the received signal to vary on the order of carrier wavelength λ . This is the small-scale effect of multi-path fading. Fading is traditionally mitigated in the physical layer using diversity coding and modulation techniques to harness the different available degrees of freedom in the time, frequency, code and antenna domains. The goal is to provide a reliable point-to-point communication link at the physical layer. 2) *Interference*: The shared nature of wireless communication may result in a point-to-point transmission being corrupted due to interference from other transmissions originating in the vicinity of the receiver (termed a *collision*). Conventionally, collision resolution and interference management are viewed

as the main function of the medium access control (MAC) layer, and are typically handled by scheduling or random access protocols. For example, in IEEE 802.11, a four-way RTS-CTS-DATA-ACK handshake is used to acquire the floor and a *binary-exponential-backoff* (BEB) algorithm is employed to adapt to the interference levels.

This separation of point-to-point link reliability and multiple access functionality between the physical and MAC layers relies on the assumption that physical layer diversity techniques work perfectly and hide fading effects from the MAC. In this paper, we provide experimental data to show that this is not the case, and that fading can, in fact, be seen at the MAC layer. Depending on the environment, this fading can be fast or slow relative to the timescales of IEEE 802.11. This coupling between the timescales of fading and IEEE 802.11 results in the increased occurrence of hidden and exposed terminal problems in the MAC [1].

We propose a unified MAC/PHY design to mitigate the effects of fading and interference by exploiting *MAC diversity*. MAC diversities arise from the fact that links to different intended receivers, or over different frequency channels experience independent time-varying fading and interference conditions. We respectively term these as *multi-receiver diversity* and *multi-channel diversity*. The key to exploring MAC diversity in a distributed manner is to allow the MAC to adapt to channel fading and interference conditions, and assign the transmission opportunity to the best link. This is similar in nature to the notion of opportunistic communication in cellular wireless systems [2]. Our MAC diversity schemes operate at the MAC frame level, and not at the symbol level, as with physical layer diversity techniques. We eschew the design of optimal adaptation schemes. Such schemes can have high overhead and be counterproductive in practice. Instead, we start off with canonical scenarios where MAC diversity is available, and use these to motivate the design of a simple diversity-exploiting MAC protocol. This protocol, which we call *dynamic-binding multi-channel MAC* (DB-MCMAC), requires very simple modifications to the transmitter side of IEEE 802.11. DB-MCMAC does not add any new control frames, or modify any IEEE 802.11 control frames. Thus, it requires *no modifications to the IEEE 802.11 air interface*. It can be completely implemented in the MAC firmware at the sender side without any receiver side modifications. A DB-MCMAC transmitter node can inter-operate with an IEEE 802.11 receiver node and thus, DB-MCMAC is backward compatible with the original IEEE 802.11 standard. We will show that even this simple protocol leads to significant performance improvements because it utilizes the additional degrees of freedom efficiently.

DB-MCMAC is based on three key ideas: firstly, the

This material is based upon work partially supported by AFOSR under Contract No. F49620-02-1-0217, NSF under Contract Nos. NSF ANI 02-21357, CCR-0325716 and CNS 05-19535, USARO under Contract No. DAAD19-01010-465, DARPA/AFOSR under Contract No. F49620-02-1-0325, DARPA under Contract Nos. N00014-0-1-1-0576. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the above agencies.

transmitter queues packets (we use *packet* and *DATA frame* interchangeably) on a per-receiver basis to prevent head-of-line blocking. Secondly, a per-channel per-receiver contention window is used to track the fading and interference conditions to every receiver on each channel. This is the analog of transmitter-side channel state information (CSI) in cellular wireless systems [2]. Thirdly, the binding of DATA to a particular channel is dynamic. A set of backoff counters, one for each receiver on each channel, start counting down from a random value chosen according to the corresponding contention window. When one of the backoff timers fires, a DATA frame is pulled from the corresponding receiver-queue and “bound” to the appropriate channel. The transmitter then tries to acquire the floor on that channel by sending an RTS to the intended receiver. Depending on whether the attempt was successful or not, the contention window is decreased or increased respectively. Different rules could be used to drive the contention window adaptation, e.g., multiplicative increase multiplicative decrease (MIMD). This exploits diversity in two ways: when a channel becomes good, the contention window for the channel decreases, and allows us to preferentially use that channel as compared to the other channels. When a channel goes bad and a transmit attempt is unsuccessful, dynamic binding allows us to switch the packet to another channel, thus preventing data loss. It also decreases the backoff delays that arise from multiple attempts on the same channel.

An approximate Markov chain model is used to carry out performance analysis of DB-MCMAC. An exhaustive performance evaluation of DB-MCMAC in ns-2 indicates that it is quite successful in harnessing MAC diversities. In environments with multi-receiver diversity, DB-MCMAC provides order of magnitude improvements over IEEE 802.11. In environments with multi-channel diversity, DB-MCMAC provides upto 150% throughput improvements over a baseline multi-channel multi-interface MAC. Further, it mitigates inter-flow and intra-flow interference whenever possible. Dynamic binding also reduces packet loss rates in fast fading environments, and thereby boosts TCP performance by upto 110%.

The remainder of the paper is organized as follows: In Section II, we present experimental data to illustrate the fading effects at the MAC layer and motivate the exploration of MAC diversities. We describe canonical scenarios to identify *multi-receiver diversity* and *multi-channel diversity* in Section III. The DB-MCMAC scheme is proposed in Section IV to harness these diversities. An approximate Markov chain model is used to carry out performance analysis of DB-MCMAC in Section V. An exhaustive ns-2 simulation study of the protocol is carried out in Section VI to demonstrate the effectiveness of our scheme in typical scenarios. Finally, we discuss related work in Section VII and conclude in Section VIII.

II. MEASUREMENT OF FADING EFFECTS AT MAC LAYER

Reliability and multiple access have traditionally been separated across the PHY/MAC layer boundary. This assumes that the PHY layer can perfectly hide fading effects from the MAC. We will experimentally show that this is not always the case, and that fading can, in fact, be seen from the MAC layer. After taking the effect of physical layer coding schemes into account, we observe that the wireless channel is in one

of two states: in the good state, packets get through without loss, whereas in the bad state, all packets are dropped. This can be modeled using a two state Markov fading process [3], as in Fig.1. The key parameter of interest is the average time spent in each of the states. We carried out measurements to obtain typical values of these parameters in IEEE 802.11 environments.




Fig. 1. Two state Markov fading model.

We use an experimental setup with a node sending ICMP ECHO REQUEST broadcast probes to a collocated receiver on channel 10. Since longer probes reduce the granularity of measurement, and inflate the probability of packet loss for a given bit error rate, we use one byte probes to measure the channel state. Broadcasting enables us to measure the real underlying channel condition, since IEEE 802.11 broadcasts do not use the RTS-CTS-DATA-ACK handshake and are not re-transmitted. ICMP ECHO response is turned off at the receiver to kill all reverse direction traffic. We use 3M 1181 EMI copper shielding tape to reduce the effective transmit range. Each experiment is run five times, and we post-process the tcpdump logs to obtain the cdf of the average time spent in each of the two states, as well as the *ETX* of the link [4], defined as $ETX \triangleq \frac{1}{1 - P_{loss}}$.

In the absence of fading and interference, IEEE 802.11 takes (approximately) 12 ms to deliver a 1500 byte packet at a data rate of 1 Mbps, and around 2 ms at 11 Mbps. In environments with interference or fading, binary exponential backoff of IEEE 802.11 increases this figure even more. In our test setup, we have observed both slow and fast fading relative to the IEEE 802.11 timescale. In Fig.2, the performance with two nodes over an intermediate quality link ($ETX = 1.93$) is shown. Here, the mass of the cdf is concentrated in the 1 to 10 ms timescale. In other words, this is fast fading relative to the IEEE 802.11 MAC timescale. The performance with two nodes over a better quality link ($ETX = 1.27$) is captured in Fig.3. In this experiment, the fading timescale of interest is on the order of 10 to 100 ms, i.e., this is a slow fading scenario relative to the timescale of IEEE 802.11.

These observations are similar in nature to the extensive outdoor link layer measurements reported in [4]. This indicates that fading at the IEEE 802.11 timescale is a problem in both indoor and outdoor wireless settings and needs to be taken into account in protocol design.

III. MAC LAYER DIVERSITIES

A traditional way to mitigate fading effects in the physical layer is through exploiting the notion of diversity. We instead use a MAC layer approach to mitigate fading and interference. The presence of multiple receivers and channels provides *MAC diversities* that we can harness to improve performance. We start off by first identifying canonical scenarios where these diversities can be exploited.




Fig. 2. cdf of MAC fading timescale: fast fading.




Fig. 3. cdf of MAC fading timescale: slow fading.

We digress to recapitulate IEEE 802.11 operation. Traditionally, IEEE 802.11 MAC and PHY layer functions are implemented in firmware. A FIFO *interface queue* (IFQ) is maintained by the device driver. When the interface is ready, the device driver dequeues a packet and transfers it to the wireless card. The MAC then initiates a RTS-CTS-DATA-ACK handshake to the intended receiver. If the handshake fails due to fading or interference, the MAC retransmits the packet till it succeeds, or the retry limit is exceeded and the packet is discarded. The wireless interface then goes back to the ready state.

The first form of MAC diversity comes from the existence of multiple intended receivers. For example, in Fig.4, node A has two neighbors: B and C. Since small scale fading effects vary on the order of the carrier wavelength λ , link l_{AB} typically fades independent of the link l_{AC} . Assume that in slot 1, l_{AB} is in “good” state and l_{AC} is in “bad” state; while in slot 2, l_{AB} is in “bad” state and l_{AC} is in “good” state. An ideal scheme would transmit to B in slot 1 and to C in slot 2. Such a diversity arises from the different link conditions from a transmitter to different receivers. We call this *multi-receiver diversity*. We can also obtain *multi-receiver diversity* from varying interference conditions at different receivers. For example, consider Fig.4 and suppose that in slot 1, C is silenced by neighbor E, while in slot 2, B is silenced by D. Then, a good diversity exploiting scheme should transmit to node B in slot 1 and to node C in slot 2. To explore this diversity, the wireless device driver for the interface should maintain a per neighbor queue and the MAC layer should be able to track the channel fading and interference conditions on a per neighbor basis.

In many wireless systems, there are multiple available non-overlapping channels, e.g., three with IEEE 802.11b/g in the 2.4 GHz band and 13 with IEEE 802.11a in the 5 GHz band.

(Due to signal power leakage and other considerations, these numbers may be lower in practice [5].) Assuming that it is possible to transmit/receive simultaneously in these frequency bands, this provides us another form of MAC diversity, which we term *multi-channel diversity*. For instance, consider the scenario in Fig.5, where there are four nodes in a straight line and there are two channels available for each node. Let $l_{AB}^{Ch_i}$ denote channel i for link l . Suppose that $l_{AB}^{Ch_1}$ is “good” and $l_{AB}^{Ch_2}$ is “bad” in slot 1; while the channel conditions are reversed in slot 2. Then, in a diversity-exploiting scheme, node A should transmit over channel 1 in slot 1 and over channel 2 in slot 2. As with multi-receiver diversity, *multi-channel diversity* can also be obtained when the interference conditions are varying across channels. Suppose that in slot 1, node B is silenced on channel 2 by a transmission on $l_{CD}^{Ch_2}$, while in slot 2, it is silenced on channel 1 by a transmission in $l_{CD}^{Ch_1}$. Then, node A should transmit over channel 1 in slot 1, and over channel 2 in slot 2.

In the case where multiple receivers and multiple channels are both available (for example, in Fig.6), in addition to maintaining per neighbor queues, we should also allow each neighbor queue to use all possible channels to achieve an increased multiplexing gain. An ideal protocol should select the receiver with the best link condition for each channel, and the channel with the best condition for each receiver.

IV. PROPOSED PROTOCOL FOR EXPLOITING DIVERSITY

In this section, we propose a *dynamic binding multi-channel MAC* (DB-MCMAC) to harness multi-receiver and multi-channel MAC diversities by extending the classic IEEE 802.11 protocol. DB-MCMAC works by using a *contention window (CW)* adaptation scheme to track link conditions and a *dynamic binding* scheme to achieve opportunistic (*receiver, channel*) pair selection.

The following assumptions drive our protocol design:

- 1) It is possible to transmit and receive simultaneously on multiple non-overlapping frequency channels. We assume that a radio includes multiple parallel RF front-end chips and baseband processing modules to support several simultaneous channels, and there is only one MAC layer to coordinate the functions of multiple channels [6]. Such capability may soon be available in hardware.
- 2) For backward compatibility with IEEE 802.11, we do not add any new control frames, or modify any existing frames. In other words, *we do not modify the IEEE 802.11 air interface*.
- 3) It is possible to modify IEEE 802.11 sender side functionality, including the contention window adaptation and the interface queue.

We eschew the design of “optimal” diversity-exploiting protocols that try to extract every ounce of performance from the wireless network. Instead, our emphasis in protocol design is on simple schemes that can exploit MAC diversities whenever available. As demonstrated later in Section VI, even our simple protocol harnesses diversity “well enough” to obtain considerable benefits.

A. Contention Window Adaptation

A necessary condition for exploiting diversity gains available in fading/interference environments is the ability of the




Fig. 4. A multi-receiver diversity scenario.



Fig. 5. A multi-channel diversity scenario.




Fig. 6. A multi-receiver and multi-channel diversity scenario.

transmitter to track the link conditions on each channel to each receiver. This is akin to the well known notion of transmitter side channel state information (CSI) in 3G cellular systems like IS-856 [2]. In the IEEE 802.11 MAC, the contention window (CW) is used to adapt to the congestion level due to interference at the receiver. CW is initialized to $CW_{min} = 31$. Whenever the transmitter times out due to (i) failure to get a CTS after sending a RTS, or (ii) failure to get an ACK after sending a DATA packet, the CW is doubled till it reaches a maximum of $CW_{max} = 1023$. On successfully transmitting a packet, or failing to do so after *RetryLimit* successive retries, the firmware discards the packet and resets CW to CW_{min} . We observe that if we maintain per receiver CW's, then the CW can be used as an indication of the interference condition of the receiver.

One approach to estimating link fading conditions is to use a feedback loop from the receiver to the transmitter, e.g., in MOAR [7] or MAD [8], the receiver measures the RTS signal strength and piggybacks the measured value in the CTS. However, this only works for slow fading environments because it relies on feedback. It is also possible that the channel is in a “deep fade” (i.e., in bad state) and as a result, the RTS and CTS packets do not get through. Instead, the approach we use is to treat the RTS as a probing packet, and use the success or failure of this RTS probe to learn the channel fading conditions.

One particularly simple and unified scheme is to use contention window adaptation to track both interference and fading conditions. In other words, we do not differentiate between failures due to fading or collisions. This has the advantage of simplifying the estimation protocol, and does a pretty good job, as seen in Section VI. To accomplish this, at each transmitter, we maintain a CW parameter on a per channel, per intended receiver basis. For example, if a transmitter node has two intended receivers A and B, and there are two channels, then it maintains four contention windows: $CW_A^{Ch_1}$, $CW_A^{Ch_2}$, $CW_B^{Ch_1}$ and $CW_B^{Ch_2}$. The contention window $CW_R^{Ch_i}$ is adapted according to the failure or success of handshaking messages for the intended receiver R on channel i . A *multiplicative increase, multiplicative decrease* (MIMD) rule is used to adapt the CW's efficiently in response to link conditions:

$$CW = \begin{cases} \min(CW \cdot u, CW_{max}), & \text{on RTS/DATA failure,} \\ \max(CW/d, CW_{min}), & \text{on tx success,} \end{cases}$$

where $u > 1$ and $d < 1$ are the increase and decrease factors respectively and are adjustable parameters that can be tuned to the speed of variation of link conditions. Note that if we set $u = 2$ and $d = \infty$, this becomes the BEB algorithm employed in IEEE 802.11, which we call the *multiplicative increase reset decrease* (MIRD) adaptation rule.

There have been works on estimating channel fading using Markov chains and performing channel dependent scheduling in cellular networks [9]. However, the problem becomes much more sophisticated when interference is taken into account in general multi-hop wireless networks, and when there are multiple channels and intended receivers. In stead of finding the “optimal” strategy, our emphasis in this paper is on the exploration of diversity through easily implementable schemes and thus, we choose to use the simple, but possibly sub-optimal MIMD adaptation rule.

B. The DB-MCMAC Protocol

We now describe the proposed DB-MCMAC protocol to harness MAC layer multi-receiver and multi-channel diversities. In IEEE 802.11, once a packet is passed to the MAC layer, the MAC is bound to that particular packet and keeps trying a RTS-CTS-DATA-ACK handshake with the intended receiver till it succeeds, or the retry limit is exceeded. If the link condition to the bound receiver is bad, then there is increased delay due to the time spent in exponentially backing off and retrying to the bad receiver. On the other hand, it is possible that there is another receiver with good link condition to which successful transmission can be completed.

Observe that it is the static binding of the packet to a channel that commits the channel to a particular receiver and prevents the successful exploration of MAC diversity. If the binding was flexible, then the MAC could respond to a RTS failure by “unbinding” the packet from the channel, and retrying on some other channel. Meanwhile, the released channel can be used to transmit a packet for some other receiver. The MAC could do this in an intelligent manner, so that receivers with better link conditions get preferential access to the channel. The DB-MCMAC protocol uses such *dynamic binding* to enable the MAC diversities to be efficiently exploited.

The architecture of DB-MCMAC is shown in Fig.7. Assume that there are m channels, and n intended receivers denoted by $\{R_j, j = 1, \dots, n\}$. In the MAC layer, packets are queued on a per-neighbor basis. The capacity of each of these neighbor queues is m .

In the neighbor queue to R_j , we associate a backoff timer and contention window for each channel k . We denote them by (R_j, k) and $CW_{R_j}^{Ch_k}$ respectively. Thus, in each channel k , there are n backoff timers $\{(R_j, k), j = 1, \dots, n\}$, each competing to win channel access for the neighbor queue it is associated with. When channel k becomes idle, all the backoff timers on that channel $\{(R_j, k), j = 1, \dots, n\}$ start counting down after a DIFS period (as in IEEE 802.11). When one of the backoff counters, say (R_i, k) , fires off first, it freezes all the other backoff timers on this channel, i.e., $\{(R_j, k), j =$




Fig. 7. Architecture of the DB-MCMAC protocol.

$1, \dots, n, j \neq i$. It then finds the first packet in the MAC neighbor queue R_i that has not been attached to any other channel, attaches this packet to channel k , and starts a RTS-CTS-DATA-ACK handshake with the intended receiver R_i to send out the packet over this channel.

If the packet is successfully delivered, the corresponding contention window, $CW_{R_i}^{Ch_k}$, is decreased using the MIMD rule. On the other hand, if the RTS or DATA packet fails (i.e., we get no CTS or ACK after a timeout period), the corresponding contention window, $CW_{R_i}^{Ch_k}$, is increased according to the MIMD rule. In addition, the unsuccessful packet is detached from channel k and replaced back in the neighbor queue R_i . In either case, we release the channel k by unfreezing all the other backoff timers on this channel, i.e., $\{(R_j, k), j = 1, \dots, n, j \neq i\}$, and restart $\{(R_i, k)\}$ with the updated $CW_{R_i}^{Ch_k}$.

It is important to note that in DB-MCMAC, the backoff timer does not correspond to an individual packet. Instead, each backoff timer (R_j, k) serves as an agent to compete for access on channel k on behalf of its MAC neighbor queue R_j . Thus, with DB-MCMAC, if a packet transmission attempt fails on one of the channels (say k), the packet is allowed to switch to another channel, and channel k can be taken over by a packet intended for a different neighbor. This mechanism is called *dynamic binding*, and allows us to fully explore the multi-receiver and multi-channel diversities in the MAC layer. As in IEEE 802.11, we also associate a *STA short retry count* (ssrc) and a *STA long retry count* (slrc) with each packet in the MAC neighbor queue to record the number of RTS and DATA failures it encounters respectively. If these retry limits are exceeded, the packet is dropped to keep the MAC from being blocked by packets that cannot get through. This completes the description of the transmit state machine of the DB-MCMAC. The receive state machine of DB-MCMAC is exactly the same as that of IEEE 802.11.

Since DB-MCMAC uses the same handshaking mechanism as IEEE 802.11, it is backward compatible with the IEEE 802.11 standard. Thus, a node running IEEE 802.11 and a node running DB-MCMAC can still exchange packets and communicate with each other, since the protocol embeds all of its intelligence on the transmit side and does not change the receive state machine.

C. Fairness

It is well known that there is a fundamental tradeoff between the system throughput and fairness [10]. Since DB-MCMAC is always trying to opportunistically serve packets with the best link conditions, it is possible that links experiencing

severe interference or fading have less chance to grab the channel. This will lead to a fairness problem between flows with different intended receivers. DB-MCMAC uses a layered design that decouples fairness issues from MAC layer functionalities. A per-neighbor queue is implemented in the interface queue (IFQ) (in addition to the one in the MAC). The scheduler (see Fig.7) is implemented in the interface queue and provides a *scheduling mechanism* to decide the service order of packets from the per-neighbor IFQ queues. Different scheduling policies can be used to provide applications with different tradeoff points. The design of scheduling policies is not the focus of this paper; we simply implement a FCFS policy that passes packets down to the MAC as soon as possible.

V. MODELING AND ANALYSIS OF DB-MCMAC

We now return to the modeling and performance analysis of DB-MCMAC in the presence of fading effects. We introduce an approximate continuous-time Markov chain model that captures the interaction between fading and MAC in the simple scenario shown in Fig.8. To our knowledge, this is the first attempt to model the interaction between fading timescales and MAC protocol operation (although the IEEE 802.11 MAC has been modeled in great detail in the literature [11], [12]).

In this scenario, we have node A transmitting to receiver B across two independent fading channels. Let us suppose that the two channels have the same data rate. As in [12], we consider the *saturation* condition where node A always has packets available for transmission on both channels, i.e., the number of outstanding packets in A 's neighbor queue to B is always 2. We model the wireless channel as a two state continuous time Markov chain (see Fig.1), as in [3]. Denote by "G" the good state and "B" the bad state. The durations of the good and bad states for channel i are assumed to be exponential random variables with means $1/\lambda_{g,i}$ and $1/\lambda_{b,i}$ respectively for $i = 1, 2$. Then, the continuous-time Markov chain model for the two independent fading channels is shown in Fig.9, where $S_1 S_2$ denotes the joint state with S_1 and S_2 representing the states of channel 1 and 2 respectively, $S_1, S_2 \in \{G, B\}$. Let the packet error probability be $p_{i,G}$ when channel i is in good state, and $p_{i,B}$ when it is in bad state; and assume that these probabilities remain constant during each state. (In general, packet error probabilities also depend on the length of the packets, and any of the RTS-CTS-DATA-ACK packets can be destroyed due to fading. To simplify the analysis, we consider the case where fading only corrupts the RTS packets.)

In the scenario shown in Fig.8, DB-MCMAC maintains two backoff counters $(B, 1)$ and $(B, 2)$ to track the fading conditions of channel 1 and 2 respectively. When the RTS-CTS handshake is successful, DATA-ACK packets are exchanged over the channel, and we say that the channel is in "sending" state, denoted by "s". When the RTS fails, the CW is adjusted and the backoff counter starts again. For simplicity, we use MIRD contention window adaptation with $u = 2$.

In the analysis, we adopt the convention used in [12]. Define $W_0 = CW_{min}$. Let m (called the "maximum backoff stage") be the value such that $CW_{max} = 2^m W_0$. The contention window can take values on $2^i W_0$, where $i \in (0, m)$ is

called the ‘‘backoff stage’’. Let $c_j(t)$ be the stochastic process representing the state of the transmitting station on channel j at time t , where $c_j \in T \triangleq \{s, 0, \dots, m\}$, $j = 1, 2$. Let σ denote the slot time in IEEE 802.11, τ_{Sifs} denote the SIFS spacing time, τ_{Difs} denote the DIFS spacing time, and τ_{Rts} , τ_{Cts} , τ_{Data} , τ_{Ack} denote the length of RTS, CTS, DATA, ACK packets respectively. Denote the backoff time at stage i by $\tau_{BO(i)}$, which is uniformly distributed on $(0, (CW_i - 1)\sigma)$. Neglecting timeouts and propagation delays, the holding time that c_j stays in state i can be approximated as

$t_i \approx \tau_{Difs} + \tau_{BO(i)} + \tau_{Rts} + \tau_{Cts} + 2\tau_{Sifs}$, $i = 0, \dots, m$, and the time that the station stays in the sending state is around

$$t_s \approx \tau_{Data} + \tau_{Ack} + \tau_{Difs} + \tau_{Sifs},$$

where the parameters are the same as in the IEEE 802.11 standard [13]. In order to employ a continuous-time Markov chain model for the analysis, we make an approximation that the holding times t_i , $i \in T$ are exponentially distributed with mean $E[t_i]$, as in [11]. For convenience, we denote by:

$$g \triangleq E[t_s] \approx \tau_{Data} + \tau_{Ack} + \tau_{Difs} + \tau_{Sifs},$$

$$f(i) \triangleq E[t_i] \approx \tau_{Difs} + \tau_{Rts} + \tau_{Cts} + 2\tau_{Sifs} + 2^{i-1}W_0\sigma.$$

With this assumption, we can model the bi-dimensional process $\{c_1(t), c_2(t)\}$ using the continuous-time Markov chain depicted in Fig.10. In this model, there are four ‘‘superstates’’: $S_1S_2 \in \{GG, GB, BG, BB\}$, which correspond to the channel fading states in Fig.9. Inside each superstate S_1S_2 , the state transition rates are

$$\begin{cases} q_{(s,j) \rightarrow (0,j)} = 1/g, & j \in T \\ q_{(i,s) \rightarrow (i,0)} = 1/g, & i \in T \\ q_{(i,j) \rightarrow (i+1,j)} = p_{1,S_1}/f(i), & i \in T \setminus \{s, m\}; j \in T \\ q_{(i,j) \rightarrow (0,j)} = (1 - p_{1,S_1})/f(i), & i \in T \setminus \{s\}; j \in T \\ q_{(i,j) \rightarrow (i,j+1)} = p_{2,S_2}/f(j), & i \in T; j \in T \setminus \{s, m\} \\ q_{(i,j) \rightarrow (i,0)} = (1 - p_{2,S_2})/f(j), & i \in T; j \in T \setminus \{s\}. \end{cases} \quad (1)$$

The first two equations in (1) account for the time in which the transmitting station finishes a sending stage. The third equation accounts for the rate at which a RTS failure occurs when channel 1 is in the backoff stage i ; while the fourth equation is the rate at which RTS succeeds when channel 1 is in the backoff stage i . The last two equations correspond to the rates of RTS failure and success on channel 2 respectively. Besides the transitions between states inside each superstate, we also add transitions between the corresponding (i, j) states in each superstate to model the channel fading effects. These transition rates are decided by the corresponding transition rates between the superstates, as labeled in the Markov chain model for two independent fading channels shown in Fig.9. Note that in Fig.10, we have only shown the intra-superstate transitions for the ‘‘GG’’ superstate. The intra-superstate transitions for the other superstates are similar.

Let $\pi_{(i,j)}^{s_1s_2} = \lim_{t \rightarrow \infty} P\{c_1(t) = i, c_2(t) = j, S_1(t) = s_1, S_2(t) = s_2\}$, $i, j \in T$, $s_1, s_2 \in \{G, B\}$ be the invariant probabilities of the Markov chain. Since we have specified all the transitions in the Markov chain, we can compute the invariant probabilities by numerical methods and various performance metrics can be estimated based on the invariant probabilities of the states. For example, the aggregate goodput η of the two channels can be computed from the datarate η_0

TABLE I
SYSTEM PARAMETERS TO OBTAIN NUMERICAL RESULTS

slot time (σ)	20 μ s	datarate	1 Mbit/s	RTS, CTS	320 bits
τ_{Sifs}	10 μ s	CW_{min}	32	DATA	4088 bits
τ_{Difs}	50 μ s	CW_{max}	1024	ACK	320 bits

TABLE II
ESTIMATED GOODPUT UNDER DIFFERENT FADING TIMESCALES

fading timescale	slow	medium	fast
$\lambda_{g,1}$ (/s)	10.0	100.0	1000.0
$\lambda_{b,1}$ (/s)	10.0	100.0	1000.0
$\lambda_{g,2}$ (/s)	10.0	100.0	1000.0
$\lambda_{b,2}$ (/s)	10.0	100.0	1000.0
goodput (Mbit/s)	0.7534	0.7599	0.9248

as follows:

$$\eta = \eta_0 \cdot \frac{\tau_{Data}}{g} \cdot \sum_{s_1, s_2 \in \{G, B\}} \left(\sum_{j \in T} \pi_{(s,j)}^{s_1s_2} + \sum_{i \in T} \pi_{(i,s)}^{s_1s_2} \right). \quad (2)$$

We use the system parameters from the IEEE 802.11 standard [13] shown in Table I to compute numerical results. We assume that the packet error probability for both channels is the same in both the good state, $p_{1,G} = p_{2,G} = 0.1$, and the bad state, $p_{1,B} = p_{2,B} = 0.9$. The goodput obtained under different channel fading scenarios (specified by the transition rates between good and bad states) is shown in Table II. The results show that as the rate at which the channels switch between good state and bad state becomes faster, the goodput is increased by the DB-MCMAC protocol. This is because the transmitter can explore the multi-channel diversity by dynamically selecting the good channel for data transmissions.

We note that this is an approximate model for a simple scenario with two independent fading channels with many simplifying assumptions. It serves our purpose of capturing the interaction between fading and MAC protocol behaviors, and demonstrates the effectiveness of DB-MCMAC’s multi-channel diversity exploration.

VI. PERFORMANCE EVALUATION

We have implemented DB-MCMAC in ns-2 and carried out an extensive performance evaluation of the protocol.

A. Simulation Methodology

We have implemented two new propagation modules in ns-2 to model the effect of small-scale fading. In both these modules, each (node, neighbor, channel) 3-tuple can independently transition between good and bad states. In the good state, the large scale two ray ground fading model decides received signal strength. In the bad state, received signal strength is 0.

- 1) *TwoRayGround/TwoStateFading*: The transition between the good and bad states is modeled by a two state Markov fading process, as in Fig. 1.
- 2) *TwoRayGround/TwoStateConfigurableFading*: The transitions can be manually configured through *oTcl*.

We have also implemented a static routing agent (Agent/WirelessManual) to configure routes manually through *oTcl*. This eliminates routing overhead and any negative interactions between routing protocols and the MAC.

We compare DB-MCMAC against two baseline schemes:

- 1) For single channel scenarios, we use IEEE 802.11 DCF.
- 2) For multi channel scenarios, we use a static binding multi-interface multi-channel MAC (SB-MCMAC). In SB-MCMAC,



Fig. 8. Simple scenario for modeling and analysis.




Fig. 9. Markov chain model for two independent fading channels.




Fig. 10. Markov chain model of DB-MCMAC for the scenario shown in Fig. 8.




Fig. 11. Multi-receiver diversity scenario: 1 sender, k receivers, 1 channel.

each interface is tuned to a different channel and runs an independent instance of IEEE 802.11 DCF. Packets for all receivers are queued in a global FIFO queue. When an interface finishes a transmission, a new packet is dequeued from the global queue and transferred to the interface. Note that SB-MCMAC can also exploit multi-channel diversity at a coarse level of granularity, since it delivers more packets to good channels than to bad channels.

Transport layer goodput is the primary metric we will use to characterize performance. In all the simulations, we use a transmit range of 250m. Carrier sensing range is set to the transmission range. The IFQ length is set to 50 packets. The data rate is set to 1 Mbps and the MIRD CW adaptation rule is used. We use CBR sources with a size of 210 bytes on top of UDP. Simulations are run for 100 seconds.

B. Does contention window adaptation track link conditions?

In DB-MCMAC, we use contention window adaptation to estimate the link interference and fading conditions. To verify its effectiveness, we simulated a sender sending to two collocated receivers on a single channel, as in Figure 11 with $k = 2$. We manually configured a “negatively correlated” fading scenario such that the link from 0 to 1 was good whenever the link from 0 to 2 was bad and vice versa. The fading timescale is 10s. The channel condition and contention window adaptation are plotted in Fig.12. It can be seen in the figure that the contention window converges quickly to $CW_{max} = 1024$ in bad state and to $CW_{min} = 32$ in good state. An analysis of the traces shows that the convergence speed is on the order of milliseconds.




Fig. 12. Tracking the fading condition: CW adaptation with NAV overriding.

C. Impact of multi-receiver diversity

1) *Fading*: We next proceed to evaluate the performance of DB-MCMAC in environments with multi-receiver MAC diversity. We use a collocated topology with one sender, k receivers and one channel, as shown in Fig.11. The two state Markov model is used with fading timescales varying from 1ms to 1s. Each link fades independently and $ETX = 2$. With three receivers, DB-MCMAC achieves throughput improvements over IEEE 802.11 ranging from 200% at timescales of 1 ms to 350% at timescales of 100 ms, as seen in Fig.13. In Fig.14, we plot the percentage improvement in throughput against the number of receivers for different fading timescales. We observe that the throughput gain due to multi-receiver diversity increases with the number of receivers at all fading timescales.

DB-MCMAC obtains these improvements because it can efficiently track link conditions using CW adaptation on a per-receiver basis. When a receiver is good, its CW adapts accordingly and its backoff timer fires earlier. Thus, DB-MCMAC opportunistically prefers good receivers and removes the head-of-line blocking effect in IEEE 802.11.

2) *Interference*: Multi-receiver diversity can also arise from multiple receivers with heterogeneous interference conditions. The scenario shown in Fig.15 provides a good example of




Fig. 13. Multi-receiver diversity gain with three receivers.




Fig. 14. Multi-receiver diversity gain against number of receivers.

such diversity. In this scenario, there are three one-hop flows: $1 \rightarrow 0$, $1 \rightarrow 2$ and $3 \rightarrow 4$. With IEEE 802.11, the flow $1 \rightarrow 2$ gets starved due to contention from the flow $3 \rightarrow 4$. Head-of-line blocking at node 1 also kills the performance of the flow $1 \rightarrow 0$, even though it is sending to an uncongested receiver. On the other hand, with DB-MCMAC, flow $1 \rightarrow 0$ can be boosted without sacrificing the throughput of the other two flows, as seen in Figure 16. This boost in performance arises from the fact that $CW_{1 \rightarrow 0}$ converges to CW_{min} , while $CW_{1 \rightarrow 2}$ converges to CW_{max} and thus, node 1 sends preferentially to receiver 0. In addition, DB-MCMAC removes head-of-line blocking.

D. Impact of Multi-Channel Diversity

1) *Fading*: We now proceed to study the MAC diversity provided by multiple channels. As before, we start with a simple scenario: a single flow from node 0 to node 1 over k channels, as shown in Fig.17. In this scenario, each of the channels is fading independently and has an ETX of 4. The performance improvement of DB-MCMAC grows monotonically with fading timescale and ranges from 15-150%, as seen in Fig.18 where we plot the goodput at different timescales for a scenario with three channels. As seen in the Markov chain analysis of DB-MCMAC in Section V, the fading timescale is crucial in determining the performance:

- 1) In slow multipath fading environments, there are few degrees of freedom to explore diversity. In such scenarios, the coarse diversity harnessed by SB-MCMAC is good enough.
- 2) In environments where fast fading interacts with MAC handshaking, DB-MCMAC's dynamic binding opportunistically exploits diversity by using good channels for transmissions. Thus, it outperforms SB-MCMAC in this regime.



Fig. 15. Multi-receiver interference diversity scenario.




Fig. 16. Multi-receiver diversity gain: interference.

In Fig.19, we plot the percentage improvement in goodput against the number of channels. From the figure, it appears that two channels are enough to harness the available multi-channel diversity gain.




Fig. 17. Multi-channel diversity scenario: 2 nodes, k channels, 1 flow.




Fig. 18. Multi-channel diversity gains in fast fading (three channels).

2) *Inter-flow interference mitigation*: Since DB-MCMAC dynamically attempts transmissions across multiple channels depending on the underlying network conditions, it can be used for inter-flow interference mitigation. For example, consider the classic hidden terminal scenario shown in Fig.20 where there are two available channels. In Fig.21, we observe that the goodput of each flow is close to that of a single one-hop IEEE 802.11 flow. In other words, DB-MCMAC mitigates inter-flow interference by reducing the RTS failure rate and orthogonalizing channel accesses.

3) *Intra-flow interference mitigation*: When IEEE 802.11 is used for multi-hop routing, a flow's transmissions at a node interfere with its transmissions at upstream and downstream neighbors. This intra-flow interference kills multi-hop performance in single channel networks. We investigate whether DB-MCMAC can mitigate intra-flow interference by using a single multi-hop UDP flow shown in Fig.22. In Fig.23, we show the goodput obtained by DB-MCMAC and SB-MCMAC against the number of hops. DB-MCMAC provides performance improvements over SB-MCMAC that increase from 19.09% for 3 hops to 53.73% for 7 hops.




Fig. 19. Multi-channel diversity gains against fading timescale as the number of channels increases.



Fig. 20. Inter-flow interference scenario: 4 nodes, 2 channels, 2 flows.

E. Impact on TCP performance

We now consider the impact of MAC diversity on TCP performance. Consider the two node, k channel, independent fading setting shown in Fig. 17 with a single TCP SACK flow and $k = 4$. As the fading speed varies, DB-MCMAC does 61.31% to 111.14% better than SB-MCMAC in terms of TCP goodput, as shown in Fig.24.

This improvement arises from two facts. As with UDP, DB-MCMAC increases MAC layer goodput by preferentially using good channels to reduce the average cost of backoff. In addition, DB-MCMAC reduces the probability of packet drop. To see why this is true, observe that with SB-MCMAC, if a packet is assigned to a channel and the channel then goes into outage, the packet is lost (after *RetryLimit* tries). Thus, mistakes made in selecting a bad channel to transmit on are very costly. On the other hand, DB-MCMAC can dynamically unbind the packet from the channel after a failed attempt and retry on some other channel. Since TCP throughput goes inversely with the square root of dropping probability [14], TCP goodput of DB-MCMAC is much better than SB-MCMAC.




Fig. 21. Inter-flow interference mitigation.




Fig. 22. Intra-flow interference scenario: k -hop flow, 2 channels, 11 Mbps data rate




Fig. 23. Intra-flow interference mitigation against the number of hops.




Fig. 24. TCP performance against fading timescale.

F. Impact of MAC Diversity on Random Topologies

Finally, we verify that MAC multi-channel and multi-receiver diversities is available in a large range of wireless scenarios, and that DB-MCMAC is able to efficiently harness these diversities to provide improvements. We generate 32 random topologies and traffic flow patterns over a 1000m \times 1000m grid with two channels. Each scenario has 32 nodes and 12 UDP flows. Routing is done using DSDV, and fading is introduced at different timescales with $ETX = 4$. The goodput is averaged over all scenarios and plotted in Figure 25, along with error bars at one standard deviation on both sides. DB-MCMAC does consistently better than SB-MCMAC and the throughput improvements increase monotonically from 34.34% in the absence of fading to 131.60% in fast fading. It should be noted that this improvement in throughput is due to diversity gains and does not come at the expense of fairness; see Table III for a comparison of the average value of Jain's fairness index for DB-MCMAC and SB-MCMAC.




Fig. 25. MAC diversity gains with random scenarios: 11 Mbps data rate

TABLE III
FAIRNESS INDEX AGAINST FADING TIMESCALE

Fading timescale	Fairness Index	
	SB-MCMAC	DB-MCMAC
No fading	0.599	0.685
100 ms	0.446	0.641
10 ms	0.482	0.709
1 ms	0.557	0.721

VII. RELATED WORK

Physical layer diversity techniques have been used to enhance the effectiveness of collision resolution in the MAC layer [15], [16] using a physical layer capable of achieving *multi-packet reception* (MPR) capabilities. However, these schemes require sophisticated signal processing and typically assume a cellular network model. Our approach identifies another form of diversity that can be completely explored at the MAC layer without relying on physical layer capabilities.

The multi-channel opportunistic auto rate (MOAR) protocol [7] exploits frequency diversity to transmit packets at a higher rate on higher quality channels. MOAR's rate adaptation mechanism relies on the CTS to provide link condition feedback. MOAR is designed for "slow fading", where the channel stays good on the order of 100 ms and multiple packets can be transmitted on good channels. Our measurement data indicates that fading timescales can vary from 1ms to 100ms. Thus, the "fast fading" regime, where fading interacts with IEEE 802.11, may also be experimentally significant and we can no longer rely on CTS feedback to track the channel conditions. It is known in cellular systems that opportunistic scheduling can even take advantage of fast fading [2]. The dynamic binding mechanism we describe is its counterpart in multi-hop random access networks. There has been a lot of research in the design of multi-channel MAC protocols [17], [18]. These can be differentiated into schemes that require multiple transceivers/antennas, and those that require only one transceiver/antenna pair. These protocols typically use one channel as a control channel to resolve the collisions and reserve other channels for data transmissions. However, as far as we are aware, none of these schemes consider channel fading effects. For example, if the control channel that is critical to the effectiveness of these schemes is in a "deep fade", the performance of these protocols will degrade dramatically. On the other hand, the diversity-exploiting DB-MCMAC protocol works well even in the presence of adverse link fading and interference conditions.

VIII. CONCLUDING REMARKS

The traditional approach to wireless system design decouples physical layer link reliability from MAC layer multiple access and implicitly assumes that the physical layer can perfectly hide fading from the MAC. We have demonstrated that fading can, in fact, be observed at the MAC layer at 1 to 100ms timescales that interact unfavorably with MAC protocol operation. To mitigate fading and interference, we have identified multi-receiver and multi-channel MAC diversities through canonical scenarios. The multi-channel dynamic binding DB-MCMAC protocol has been proposed to harness these MAC diversities. In DB-MCMAC, link conditions are

tracked through CW adaptation and links with good conditions are dynamically selected for data transmission. The protocol design has been deliberately simplified in order to ensure backward compatibility with IEEE 802.11.

We have analyzed the performance of DB-MCMAC in a simple two channel scenario using a simplified continuous-time Markov chain to capture the interaction of fading with MAC protocol behavior. Through ns-2 simulations, we have shown that DB-MCMAC can successfully harness MAC diversities, providing order of magnitude throughput improvements in single channel environments, and upto 150% throughput improvement in multi-channel environments. It automatically suppress intra and inter-flow interference when there are multiple channels and boosts TCP goodput by reducing packet drop rates. In random scenarios, DB-MCMAC consistently improves performance by 30 – 130% depending on the fading timescale. It is possible to design more sophisticated schemes that harness the diversity efficiently. This is the subject of future work.

REFERENCES

- [1] M. Cao, V. Raghunathan, and P. R. Kumar, "Robust floor acquisition in the presence of multiple fading channels," in *Proceedings of CISS 2006*, 2006.
- [2] D. Tse and P. Viswanath, *Fundamentals of Wireless Communications*. Cambridge University Press, 2005.
- [3] H. M. Chaskar, T. V. Lakshman, and U. Madhoo, "TCP over wireless with link level error control: analysis and design methodology," *IEEE/ACM Transactions on Networking*, vol. 7, no. 5, 1999.
- [4] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *Proceedings of ACM SIGCOMM*, 2004.
- [5] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou, "A multi-radio unification protocol for IEEE 802.11 wireless networks," in *Proceedings of IEEE BROADNETS*, 2004.
- [6] I. Akyildiz, X. Wang, and W. Wang, "Wireless mesh networks: A survey," *Computer Networks Journal*, vol. 47, no. 3, 2005.
- [7] V. Kanodia, A. Sabharwal, and E. Knightly, "MOAR: A multi-channel opportunistic auto-rate media access protocol for ad hoc networks," in *Proceedings of IEEE BROADNETS*, 2004.
- [8] Z. Ji, Y. Yang, J. Zhou, M. Takai, and R. Bagrodia, "Exploiting medium access diversity in rate adaptive wireless lans," in *Proceedings of ACM MOBICOM*, 2004.
- [9] M. Hassan, M. M. Krunz, and I. Matta, "Markov-based channel characterization for tractable performance analysis in wireless packet networks," *IEEE Transactions on Wireless Communications*, vol. 3, no. 3, 2004.
- [10] H. Luo, S. Lu, and V. Bharghavan, "A new model for packet scheduling in multihop wireless networks," in *Proceedings of ACM MOBICOM*, 2000.
- [11] F. Cali, M. Conti, and E. Gregori, "IEEE 802.11 Wireless LAN: Capacity analysis and protocol enhancement," in *Proceedings of IEEE INFOCOM*, 1998.
- [12] G. Bianchi, "Performance analysis of the IEEE 802.11 Distributed Co-ordination Function," *IEEE Journal of Selected Areas in Communications*, 2000.
- [13] IEEE 802.11 Wireless Local Area Network Working Group, *IEEE 802.11 Standard*. IEEE Standards for Information Technology, 1999.
- [14] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP throughput: a simple model and its empirical validation," in *Proceedings of ACM SIGCOMM '98*, 1998.
- [15] M. K. Tsatsanis, R. Zhang, and S. Banerjee, "Network-assisted diversity for random access wireless networks," *IEEE Transactions on Signal Processing*, 2000.
- [16] L. Tong, Q. Zhao, and G. Mergen, "Multipacket reception in random access wireless networks: From signal processing to optimal medium access control," *IEEE Communications Magazine*, 2001.
- [17] N. Jain, S. Das, and A. Nasipuri, "A multichannel CSMA MAC protocol with receiver-based channel selection for multihop wireless networks," in *Proceedings of IEEE ICCCN*, 2001.
- [18] J. So and N. Vaidya, "Multi-channel MAC for ad hoc networks: Handling multi-channel hidden terminals using a single transceiver," in *Proceedings of ACM MOBIHOC*, 2004.