

High-Performance MAC for High-Capacity Wireless LANs

Yuan Yuan, Daqing Gu, William Arbaugh, Jinyun Zhang

TR2005-080 October 2004

Abstract

The next-generation wireless technologies, e.g., 802.11n and 802.15.3a, offer a physical-layer speed at least an-order-of-magnitude higher than the current standards. However, direct application of current MACs leads to high protocol overhead and significant throughput degradation. In this paper, we propose ADCA, a high-performance MAC that works with high-capacity physical layer. ADCA exploits two ideas of adaptive batch transmission and opportunistic selection of high-rate hosts to simultaneously reduce overhead and improve aggregate throughput. It seeks to provide high-rate hosts temporal fair share of the channel similar to the single-rate IEEE 802.11, and low-rate hosts proportional temporal fairness in long term. Simulations show that the ADCA design increases the throughput by 112% and reduces average delay by 54% compared with the legacy DCF. It delivers more than 100 Mbps MAC-layer throughput, compared with 35Mbps by the legacy MAC.

International Conference on Computer Communications and networks (ICCCN)

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

High-Performance MAC for High-Capacity Wireless LANs

Yuan Yuan*, Daqing Gu[†], William Arbaugh*, Jinyun Zhang[†]

* Computer Science Department, University of Maryland, College Park, MD, 20740. {yuanyuan, waa}@cs.umd.edu

[†]Mitsubishi Electric Research Laboratories, Cambridge, MA, 02139. {dgu,jzhang}@merl.com

Abstract—The next-generation wireless technologies, e.g., 802.11n and 802.15.3a, offer a physical-layer speed at least an-order-of-magnitude higher than the current standards. However, direct application of current MACs leads to high protocol overhead and significant throughput degradation. In this paper, we propose ADCA¹, a high-performance MAC that works with high-capacity physical layer. ADCA exploits two ideas of adaptive batch transmission and opportunistic selection of high-rate hosts to simultaneously reduce overhead and improve aggregate throughput. It seeks to provide high-rate hosts temporal fair share of the channel similar to the single-rate IEEE 802.11, and low-rate hosts proportional temporal fairness in long term. Simulations show that the ADCA design increases the throughput by 112% and reduces average delay by 54% compared with the legacy DCF. It delivers more than 100Mbps MAC-layer throughput, compared with 35Mbps by the legacy MAC.

I. INTRODUCTION

Recent advances in wireless communications, smart antennas, and digital signal processing have made it feasible to provide very high-capacity wireless links at the physical layer. These emerging physical-layer technologies offer at least an-order-of-magnitude larger bandwidth than the current generation standards. IEEE 802.11n [6], for example, is standardizing specifications that offer up to 100 Mbps throughput at the MAC layer. IEEE 802.15.3a [4], based on ultra-wideband (UWB) communications, aims at data rates of 110 Mbps or higher in personal area networks.

In this paper, we design a high-performance MAC for such high-capacity physical-layer technologies. The proposed MAC will follow the CSMA/CA paradigm, the dominant MAC approach in wireless data networks. Our goal is to deliver much higher throughput at the MAC layer than what the current MAC solutions can offer with a high-capacity physical layer in place. The proposed MAC should significantly increase the throughput gain in terms of the ratio of MAC-layer throughput to physical-layer bandwidth. This needs to be done by simultaneously reducing the MAC overhead and increasing aggregate channel throughput. While the direct application of current 802.11 MAC can only deliver less than 50Mb/s throughput at the MAC layer with a 216Mb/s physical-layer rate, we seek to reach about 100Mb/s in the same setting. Moreover, the MAC should support diverse applications, including both loss-sensitive data and latency-sensitive multimedia applications, with QoS assurance.

There are two main challenges for this high-performance MAC design. First, how to minimize the protocol overhead is a nontrivial issue to address. The current 802.11 DCF [1] uses control messages of RTS, CTS, and ACK, contention backoff and various inter-frame spacing parameters, in order for the CSMA/CA-based MAC to function properly. However, these parameters incur high protocol overhead. When the physical-layer rate further increases, the high overhead becomes even more significant since the data-carrying time shrinks as the overhead part remain constant. Second, how to improve the overall channel throughput by leveraging the good channel quality of hosts poses another challenge. Wireless channel condition, as well as the perceived transmission rate of a host at any given time, is location dependent and time varying. If we can exploit this channel dynamics and opportunistically select hosts with best channel conditions, we can achieve significant improvement in system throughput. In addition, the MAC solution needs to handle the issue of variable packet size, which is common for both data and multimedia applications.

The state-of-the-art MAC solutions are not designed for the high-capacity physical layer. They do not address issues of minimized overhead and maximized channel throughput simultaneously. The 802.11 MAC [1] incurs considerable protocol overhead. 802.11e [2] focuses on achieving MAC QoS but does little to improve channel efficiency. Dynamic TDMA-based design, e.g., the 802.15.3 MAC [3], works well for constant-bit-rate multimedia applications, but is not very efficient for bursty data applications.

In this paper, we propose Adaptive Distributed Channel Access (ADCA), a highly efficient MAC for high-capacity physical layer. ADCA primarily targets at the infrastructure mode, the dominant operation mode in practice. ADCA uses two main ideas, adaptive batch transmission and opportunistic selection of high-rate hosts, to improve channel efficiency. Instead of transmitting a single packet for a host after a successful control handshake, ADCA allows for each host to transmit multiple back-to-back packets and reply with an ACK message for a block of packet transmissions. At a given time, ADCA also adaptively favors hosts in good channel conditions, thus with high transmission rates, to contend for the channel. For high-rate hosts, ADCA achieves long term throughput proportional fairness in that such hosts transmit data in proportional to their current transmission rates. For low-rate hosts, ADCA achieves proportional temporal fairness in that their channel access time is proportional to their rates.

¹This design is US patent held by Mitsubishi Electric Research Laboratories, USA

This way, each host receives a minimum fair share, the overall throughput is additionally improved by opportunistically favoring hosts with higher rates at any time. Moreover, ADCA provides service differentiation via differential backoffs for hosts in multiple service categories, and handles variable packet size for each application.

Extensive simulations show that ADCA can deliver more than 100Mb/s MAC-layer throughput for a 216Mb/s physical layer and improves the current 802.11MAC (without RTS/CTS turned on) by up to 112%. ADCA is also able to support both data applications and multimedia streams in a single framework. It reduces the average delay by up to 55% in our simulations. The performance remains stable with a larger number of hosts, thus scalable to large user population, compared with the current 802.11 MAC.

The rest of the paper is organized as follows. Section 2 illustrates the limitations of the current 802.11 MAC. Section 3 describes the design of ADCA. Section 4 evaluates ADCA via extensive simulations, and Section 5 compares it with the related work. Section 6 concludes this paper.

II. LIMITATIONS OF CURRENT IEEE CSMA/CA MAC

We focus on high-capacity, packet-switched wireless LANs operating at the infrastructure mode, and the design can also work within wireless PANs. Within wireless cell, an Access Point (AP) coordinates packet transmissions for all residing hosts. The wireless channel is shared for both uplink (from a host to an AP) and downlink (from an AP to a host) flows, and for both data and signaling.

The popular IEEE 802.11 Distributed Coordination Function (DCF) MAC applies in both infrastructure and ad-hoc modes and follows the CSMA/CA paradigm. Upon a packet transmission, a host uses carrier sensing and waits until the channel becomes idle, then defers for DIFS time interval. The host then backs off for a value randomly chosen between zero and Contention Window (CW). Once the backoff timer expires, an optional RTS/CTS handshake is initiated between the two hosts, followed by DATA packet and ACK transmissions. In the infrastructure mode, when the packet size is larger than the RTS threshold, RTS/CTS is recommended to be turned on to reduce the damage of collisions [1].

The current 802.11 MAC is fairly inefficient for the high-capacity physical layer. Figure 1 illustrates the MAC-layer throughput achieved by DCF when the physical layer is 216Mb/s. The figure shows that the current DCF MAC can only deliver about 48Mb/s throughput without RTS/CTS, and merely 30Mb/s when the RTS/CTS option is on. The DCF MAC has three main limitations that cause such channel inefficiency:

- The IEEE 802.11 MAC incurs high protocol overhead, which comes from RTS/CTS, packet preambles, acknowledgments, contention windows and various interframe-spacing parameters. The overhead becomes more significant as the physical-layer rate increases substantially.
- The current MAC does not effectively exploit the multirate capability to increase the overall channel utilization. DCF




Fig. 1. Throughput of Legacy MAC

ensures roughly the same long-term access probability, hence throughput fairness for each host, no matter at what speed each host can transmit at a given time. As a result, the channel is unnecessarily monopolied by low-rate hosts in terms of access time. The high-rate hosts only receive a disproportionately lower amount of access time. Both aspects lead to reduced overall throughput.

- Various multimedia applications, such as video conferencing and HDTV, made possible by the new physical layer. These applications typically require QoS assurances in terms of minimum delay and/or minimum bandwidth in order for them to function properly. Current DCF MAC solutions, however, do not supply any QoS provision.

III. ADCA DESIGN

ADCA uses a combination of adaptive batch transmission and opportunistic selections of high-rate hosts, to both reduce MAC overhead and improve overall channel throughput. It offers multiple Access Categories (ACs) to provide service differentiation. ADCA uses a Reference Parameter Set (RPS) and a Credit Counter Set (CCS) for each AC to regulate channel access within CSMA/CA. Each host periodically receives RPS, comprised of Reference Rate (R_f), Reference Packet Size (S_f), and Reference Batch Size (B_f), announced and adjusted by the AP in its BEACON frames. A host also maintains CCS to record its current channel access credit and handle variable packet sizes. With RPS and CCS, each host will independently determine whether it is eligible for accessing channel to transmit packet after *winning channel contention*, defined as the host senses channel idle after its backoff timer expires. When a host's transmission rate is larger than a reference rate R_f specified in the RPS, it is eligible for channel access, and it may transmit several back-to-back packets in a batch once succeeding in channel contention. The receiver also replies with a combined ACK upon receiving a block of packets. When a host's rate is smaller than R_f , it has to wait to accumulate access credit for accessing the channel to transmit packets. Once its credit exceeds a threshold, it can also transmit a batch of packets on winning channel contention.

In summary, ADCA uses batch packet transmission and block ACK to reduce protocol overhead. It opportunistically favors high-rate hosts that are in good channel condition and provides temporal fair access for them as in single-rate legacy MAC. ADCA also ensures a adaptive minimum fair share for

low-rate hosts in that it provides proportional temporal fairness for these hosts in long term. The pseudo code for ADCA is shown in Figure 2.

```

ON JOINING NETWORK
1  Initiate values for  $S_f$ ,  $R_f$ ,  $B_f$  and  $A_f$  according to received Beacon Frame;
2   $credit_h = 0$ ;
3   $credit_l = 0$ ;
ON WINNING CHANNEL CONTENTION
1  /*  $R, S$  are stations's data rate and packet size */
2  if  $R \geq R_f$ 
3      /* keep proportional throughput fairness among high rate stations */
4      int  $B = ((S_f/R_f) * B_f + credit_h)/(S/R)$ ;
5      if  $B < B_f$ 
6           $credit_h + = (S_f/R_f) * B_f$ ;
7           $credit_l + = 1$ ;
8          resume to backoff procedure;
9      else
10         Transmit up to  $B$  number of packets with  $A_f$  parameter;
11          $credit_h = 0$ ;
12          $credit_l = 0$ ;
13  else /* for low rate stations  $R < R_f$  */
14     /* keep proportional temporal fairness for them */
15     if  $credit_l < (R_f/R)$ 
16          $credit_l + = 1$ ;
17          $credit_h + = (S_f/R_f) * B_f$ ;
18         resume to backoff procedure;
19     else
20          $B = ((S_f/R_f) * B_f)/(S/R)$ ;
21         Transmit up to  $B$  number of packet with parameter  $A_f$ 
22          $credit_h = 0$ ;
23          $credit_l = 0$ ;

```

Fig. 2. Pseudo Code of ADCA MAC

A. Adaptive batch transmission and block acknowledgment

ADCA allows for multiple back-to-back packet transmissions between the AP and a given host, to reduce protocol overhead. To achieve this goal, the AP advertises three parameters R_f , S_f , and B_f to each host. These three parameters state that, given a given time, for a host transmitting at rate R_f , it can potentially transmit a batch of B_f consecutive packets, each of which is S_f bytes. During each batch transmission for a given host, the consecutive B_f packets are separated only by the smallest time interval SIFS. This is in sharp contrast to the legacy 802.11 MAC, in which two consecutive packet transmissions are separated by the time of DIFS plus random contention backoff and possibly the optional RTS/CTS exchange. This way, batch transmission significantly reduces the MAC overhead. The choice of B_f reflects the average channel coherence time [9] for hosts in the local cell.

The above batch transmission is adaptive to each host's current rate. When a host is not operating at the reference rate R_f , its batch size or access probability is adjusted in proportional to its current transmission rate.

In order to further reduce the MAC overhead, we also use block ACK via the parameter A_f , which is negotiated between two communicating hosts according to their perceived channel conditions. In ADCA, a single ACK signal is sent back to the sender for a block of A_f number of back-to-back transmitted packets, instead of per-packet ACK in the current 802.11 MAC. If some packets in the block are not received correctly, the sender retransmits corrupted packets indicated the block ACK packet. This further reduces the protocol overhead of the current MAC.

B. Opportunistic selection of high-rate hosts

ADCA preferentially grant those hosts under best channel conditions for channel contention, while restraining others when they perceive poor channel.

Specifically, for high-rate hosts with rates $R \geq R_f$, ADCA allows for them to access channel immediately after they win channel contention, as long as they can transmit a batch of B_f packets for transmission, it will resume to channel contention but retain channel time credit via its credit counter $credit_h$. When its accumulated credit is sufficient for a batch of B_f packets, it can transmit immediately after succeeding in channel contention.

For low-rate hosts with rates $R < R_f$, we will defer their channel access and control their access probability, and only allow them to contend for the channel approximately every R_f/R interval. This is realized by its credit counter $credit_l$. This credit counter is incremented each time low-rate host wins the channel contention, but not access channel to transmit packets. When its accumulated credit reaches R_f/R and it succeeds in channel contention, the low-rate host can transmit a batch of $(R/R_f) * B_f$ packets, each of which has a size S_f . At the same time we still keep $credit_h$ for them in case they get channel time compensation as soon as they perceive good channel conditions during this process.

The above design balances between maximizing overall channel throughput (by providing high-rate hosts currently in good conditions with higher access probability) and providing minimum share for the channel to avoid starvation for low-rate hosts. It, therefore, provides minimum fair share to each host and additionally maximizes channel throughput. The provided fairness for a contending host is as follows. For high-rate hosts, we provide them temporal fairness in terms of identical channel access time, which similar to legacy MAC in single-rate. Equivalently, this implies proportional throughput fairness in that the throughput is proportional to its current transmission rate R . For low-rate hosts, each host is provided proportional temporal fairness in that its access time is roughly proportional to its current transmit rate. Its throughput, accordingly, is in square proportion to its rate.

In the above design, we also handle the issue of variable packet size. When the packet size is different from S_f , the batch size and the credit counters will also be re-calibrated by S/S_f . This is also obvious from the pseudo code.

C. Achieving service differentiation

ADCA also achieves service differentiation through differential backoffs in multiple service categories. It prioritizes traffic with different QoS requirements in terms of throughput and latency via several Access Categories (ACs). Each AC has a separate backoff value. A higher-priority AC has smaller backoff values, whereas lower-priority AC has larger backoff ones. This way, the higher-priority AC always has preference over channel access. This mechanism is similar to EDCA in 802.11e.

D. Implementation Issues

In ADCA, each AC has its own RPS settings that can be adjusted depending on the current system performance and channel condition. The RPS settings are managed by the AP, and their configurations are included in the periodic Beacon frame.

In ADCA, each host may adapt its current transmission rate depending on its perceived SNR. ADCA can work with any rate adaptation mechanism, e.g., Auto Rate Fallback (ARF) or BRAF. Similar to BRAF, ADCA can use the RTS/CTS handshake to let the receiver choose the best rate for data transmissions. The parameter A_f can also be negotiated and adjusted via RTS/CTS according to the current channel condition; we expand fields in RTS and CTS to carry parameter A_f . ADCA recommends turning on the RTS/CTS for these purposes. Also we need to accommodate ACK packet format to notify sequence numbers of corrupted packets for retransmission.

IV. SIMULATIONS

We use ns-2 simulations to evaluate the performance of ADCA, and compare it with the IEEE 802.11 MAC scheme, which is also enhanced with EDCA functionalities for service differentiation. Table I lists the physical-layer parameters used in our simulations; the values are based on the 802.11n physical-layer design proposed by MERL². Most of these parameters are compatible with the 802.11a specification. Each host is also allowed to transmit and receive messages at different rates depending on the channel condition.

TABLE I
PHY/MAC PARAMETERS USED IN SIMULATION

SIFS	16 μ s	AIFS[AC0,1]	54 μ s
DIFS	34 μ s	CWmin[AC0,1]	31
Slot Time	9 μ s	CWmax[AC0,1]	1023
ACK size	14 bytes	AIFS[AC2]	43 μ s
MAC Header	28 bytes	CWmin[AC2]	15
Peak DataRate	216Mb/s	CWmax[AC2]	500
Basic DataRate	24Mb/s	AIFS[AC3]	34 μ s
PLCP Preamble Length	20 μ s	CWmin[AC3]	7
PLCP Header Length	4 μ s	CWmax[AC3]	100

A. MAC Throughput Gain At Various Physical-Layer Rates

We first evaluate the throughput gain of ADCA at different physical-layer rates, compared with the legacy IEEE 802.11 MAC. The scenario has five hosts, each carrying a UDP traffic source in AC0 at 40Mb/s and 1280B per packet. In ADCA, we set $(S_f/R_f) * B_f$ as 3ms, with R_f as 216Mb/s and S_f as 1280B, in order to enable batch transmission. A_f is set to 3 to have an acknowledgment every three-packet transmission.

Figure 3 depicts the throughput for both ADCA and the legacy MAC as the physical-layer transmission rate varies





Fig. 3. Throughput as function of transmission rate


from 24Mb/s to 216Mb/s. As shown, the MAC-layer throughput increases for both ADCA and legacy MAC, as the physical-layer rate increases. The MAC-layer throughput in the legacy MAC only increases by 204% as the physical layer increases from 24Mb/s to 216Mb/s. ADCA, in contrast, improves its MAC-layer throughput by 450% in the same setting. Moreover, the throughput of ADCA almost obtain linear throughput improvement with the physical-layer rate, whereas the throughput increases sublinearly in IEEE 802.11 MAC. ADCA, therefore, is more efficient with the high-capacity physical layer. Our study further shows that, the overhead incurred by the physical-layer preamble and header are the limiting factor for the throughput of ADCA.

B. Exploiting Multirate Transmissions by Different Hosts

We next demonstrate that ADCA can adaptively exploit the channel conditions experienced by different hosts to further improve the aggregate throughput. We consider the simulation setting with ten hosts, each carrying a UDP flow. Five hosts are transmitting at 216Mb/s, and each UDP source rate is 20Mb/s. The remaining five hosts are transmitting at a lower rate of 54Mb/s, and each UDP source rate is 5Mb/s. The reference packet size is 1280KB, the batch transmission interval $(S_f/R_f) * B_f$ is 3ms, and A_f is set as 1.



(a) Throughput vs R_f



(b) Throughput gain vs background traffic

Fig. 4. Throughput vs R_f and as a function of background traffic

Figure 4(a) shows the aggregate throughput performance for high-rate hosts, low-rate hosts, and all hosts, respectively. The high-rate hosts can increasingly exploit their good channel conditions and high transmission rates in ADCA by accessing the channel longer than the low-rate hosts to improve the overall channel throughput. This is realized by increasing a single parameter R_f . When R_f is set as 54Mb/s, temporal fairness is achieved among all hosts, which brings 83%

²Mitsubishi Electric Research Laboratories, Cambridge, MA, USA

throughput gain compared with the 802.11 MAC that penalizes high-rate hosts to achieve long-term throughput fairness. When R_f is further increased to 108Mb/s and 216Mb/s, respectively, ADCA acquires 121% and 154% overall throughput gains compared with the legacy MAC. In these cases, the channel access probabilities of low-rate hosts are reduced in proportion to the ratio of R/R_f . The access time by low-rate hosts, consequently, decreases proportionally. By limiting the access probability by low-rate hosts that are nevertheless not in good channel conditions, ADCA provides more transmission opportunities for high-rate hosts perceiving good channels. This feature, enabled by tuning one single parameter R_f , is important for high-speed wireless LANs to mitigate the severe throughput degradation incurred by low-rate hosts and to take advantage of channel dynamics.

C. Service Differentiation

In this set of experiment, we evaluate effects of ADCA design on service differentiation. We consider a multimedia scenario, which includes five high-priority (AC3) CBR on-off audio sources with a sending rate 64kbps and 160B per packet, five medium-priority (AC2) VBR video sources with an average rate of 200kbps, and a number of low-priority (AC1) CBR video sources of 3.2Mb/s rate and 800B packet size. The VBR traffic is generated based on the videoconferencing tool *vic*.

Figure 4(b) plots the throughput gain achieved by ADCA over the legacy MAC, as the number of CBR video sources increases from 10 to 30. Observe that ADCA ensures throughput gain up to 90%, as the number of CBR traffics grows. This enables ADCA to support more bandwidth-demanding applications through its improved channel efficiency. At the same time, this causes minor increased delay for the traffic.

TABLE II
MEAN DELAY (MS) AS FUNCTION OF BACKGROUND FLOWS

Num. of flows	10	15	20	25	30
Audio (ADCA)	0.270	0.801	1.320	1.595	2.335
Audio (802.11)	0.307	0.316	0.335	0.374	0.336
Video (ADCA)	0.694	1.443	3.112	6.252	6.071
Video (802.11)	0.685	0.851	0.883	0.921	1.028

Table II lists the average delay of Audio and VBR video traffics when the number of CBR video sources varies from 10 to 30. The mean delay increases faster in ADCA than in the legacy MAC as more sources join the arena. This is because the high-priority applications have to wait $(S_f/R_f) * B_f$ to resume channel contention, rather than one packet transmission time in the 802.11 MAC. The absolute values of mean delay, however, remain small and tolerable to most multimedia applications. Thus service differentiation is still effective to provide lower delay bounds for high-priority traffics.

D. The Impact of Parameters

In the final experiments, we study the impact of two RPS parameters on the performance of ADCA. We consider five

hosts, each carrying one UDP traffic source in AC0 and transmitting at their peak rate 216Mb/s. The UDP sources generate packets of 1280B at a rate of 40Mb/s. As the channel will be saturated by these five flows, we evaluate the impact of parameter B_f and A_f in RPS on channel efficiency while keeping R_f at 216Mb/s and S_f as 1280B.

1) *Batch Size B_f* : Figure 5(a)(b) depicts the throughput and the mean delay for both ADCA and the legacy MAC as the parameter B_f varies. In this case, A_f is set to 1. We observe that without RTS/CTS handshake, ADCA results in significant throughput gain compared with the 802.11 MAC in the range of 14.5% to 68% as the batch transmission interval $(S_f/R_f) * B_f$ increases from 0.2ms to 6ms. The batch transmission begins to take major impact when the batch transmission interval grows larger than 1ms. The average delay is simultaneously reduced by about 50% as B_f increases. ADCA achieves such gains primarily by reducing the overhead incurred by deferral and backoff procedures of the MAC.

Furthermore as we turn on RTS/CTS handshake, the recommended practice to handle large packet size, we achieve even greater throughput gain by as high as 128% and reduce the mean delay about 55% compared with the legacy MAC. On the other hand, turning on RTS/CTS in the legacy MAC results in 27.4% throughput degradation and around 39% mean delay increase, compared with the case of no RTS/CTS. However turning on/off RTS/CTS only has minor effect on ADCA.

2) *Block Acknowledgment Frequency A_f* : We now study the impact of A_f on performance of our ADCA design. Here, we keep $(S_f/R_f) * B_f$ as 3ms. The throughput and mean delay results for ADCA and the legacy MAC as the function of parameter A_f are shown in Figure 5(c)(d). The results indicate that, increasing parameter A_f can further increase the MAC-layer throughput to over 100Mb/s, the goal set by the current IEEE 802.11n working group. The throughput gain obtained by increasing A_f can be 112% over 802.11 DCF. In the meantime, the average delay has been reduced by approximately 40%.

In addition, we see that significant throughput gain can be achieved while A_f is set to be 2 or 3. Therefore, even hosts adopting small values for A_f according to their channel conditions can utilize this parameter to greatly improve their throughput and delay performance.

V. RELATED WORK

Two most popular approaches to wireless MAC are CSMA/CA based and TDMA based schemes. IEEE 802.11 DCF, unarguable the most widely deployed MAC, uses CSMA/CA. Its downsides include inefficient channel utilization due to large MAC overhead, and lack of QoS support. Recent effort on 802.11e provides service differentiation to meet requirements by various applications. The 802.11e MAC also improves channel efficiency using the BlockAck technique, but this mechanism is quite complex in 802.11e by requiring an explicit setup and tear-down procedure. Moreover, 802.11 and 802.11e can not opportunistically use high-rate hosts to improve overall channel throughput. Our ADCA

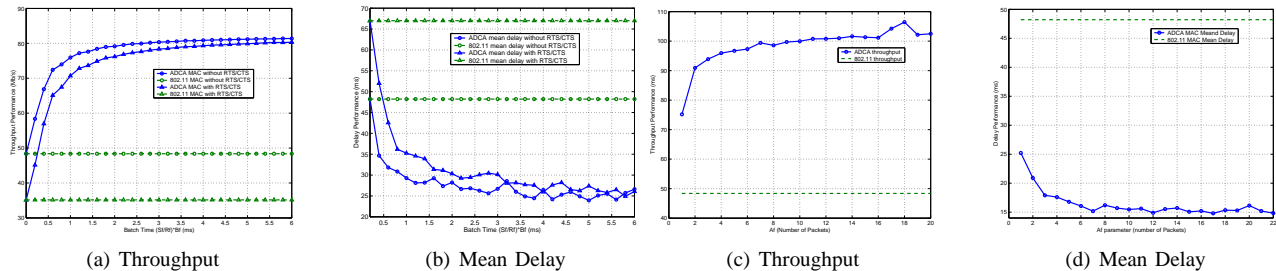


Fig. 5. Throughput and delay performance as functions of B_f and A_f

scheme addresses all these limitations. It is efficient, enabling opportunistic use of the high-rate hosts, and scaling to large user population. All these features are achieved within the CSMA/CA framework.

Dynamic TDMA is used in the HiperLAN/2 (H/2) [7]. It requires explicit resource reservation and centralized channel time allocation, which is good for constant-bit-rate multimedia but not very efficient for bursty data applications. It is also prone to the hidden terminal problem. Note that even in the infrastructure mode, hidden terminals may still exist due to insufficient numbers of physical channels used by neighboring APs. IEEE 802.15.3 [3] [4] also uses dynamic TDMA in its MAC for WPANs. It aims at high capacity physical layer. However, the scenarios are mainly in the home environment, which only accommodates very few devices and users within each AP. It also requires fine-grained, accurate time synchronization at the slot level.

There are numerous papers addressing issues of MAC fairness, energy efficiency, directional antennas, and security (see recent conference proceedings of mobicom, mobihoc, and infocom, and various wireless networking journals for a good sample). Recent studies [9] [8] also examine the issue of inefficiency of the 802.11 MAC. These proposals improve channel efficiency by more accurate channel estimation, or by opportunistically exploiting good channel conditions of a few hosts. However, none of them addresses the issue of high-capacity physical layer. They do not provide QoS, and the MAC performance is not scalable to large node population.

VI. CONCLUSION

Emerging high-speed wireless LAN and PAN technologies seek to provide high-capacity physical layer at least an order of magnitude higher than the current-generation standards. However, direct extension of the legacy MAC to such scenarios will significantly compromise the MAC efficiency. The fundamental problem is that, the current MAC incurs a significant amount of MAC-layer overhead and cannot leverage the very high transmission rates exploited by hosts to improve channel throughput. As a result, the legacy 802.11 MAC can only offer about 48Mb/s throughput at the MAC layer when the RTS/CTS option is turned off, and about 35Mb/s when the RTS/CTS is turned on, with a 216Mb/s physical layer.

In this paper, we propose ADCA, a high-performance MAC that works in concert with high-capacity physical layer in

wireless LANs and WPANs. ADCA minimizes the MAC overhead via adaptive batch transmission and block ACK. Each host transmits multiple back-to-back packets, governed by its channel coherence time, once it succeeds in channel contention. The receiving host only sends back a single ACK upon receiving multiple packets. ADCA also allows high-rate hosts (i.e., their transmission rate is higher than R_f) to contend the channel with higher probability, while limiting the access probability of low-rate hosts that are in bad channel conditions. ADCA ensures temporal fair share of the channel among high-rate hosts as in single-rate IEEE 802.11, and provides adaptively proportional temporal access to low-rate hosts in proportion to their current transmission rates in long term. As a result, each host receives a minimum share of the channel, while the overall channel throughput is improved significantly. In addition, ADCA achieves service differentiation via differential backoffs for various access categories (ACs).

Our extensive simulations show that ADCA achieves up to 128% throughput gain, and reduces the average delay by about 54%, compared with the legacy 802.11 MAC. ADCA can offer 106Mb/s MAC-layer throughput with RTS/CTS turned on, when the physical-layer rate is 216Mb/s.

REFERENCES

- [1] B. O'Hara and A. Petrick. IEEE 802.11 Handbook, A Designer's Companion. IEEE Press, 1999.
- [2] IEEE Std 802.11e/D8.0 Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY).
- [3] IEEE Std 802.15.3TM Part 15.3: Wireless MAC and PHY Specifications for High Rate WPANs.
- [4] IEEE P802.15.SG3a PAR: Higher Speed Physical Layer Extension for the High Rate Wireless Personal Area Networks (WPAN).
- [5] IEEE Std 802.11a-1999 Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY).
- [6] IEEE 802.11n PAR: Draft Amendment to STANDARD for Information Technology-Telecommunications and information exchange between systems-Local and Metropolitan networks-Specific requirements-Part 11: Wireless LAN MAC and PHY specifications: Enhancements for Higher Throughput.
- [7] EN 300 652 V1.2.1(1998-07), Broadband Radio Access Networks (BRAN); High Performance Radio Local Area Network (HIPERLAN) Type 1; Function specification, ETSI, 1998.
- [8] G. Holland, N. Vaidya, and P. Bahl. A rate-adaptive MAC protocol for multi-hop wireless networks. In MOBICOM 2001.
- [9] Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, Opportunistic media access for multirate ad hoc networks, in MOBI-COM 2002.