

Multipacket Reception Enabled Aggregation for Very High-Speed WLANs

Tianji Li, David Malone, Douglas Leith

Hamilton Institute, National Univ. of Ireland at Maynooth, Ireland

Email: {tianji.li, david.malone, doug.leith}@nuim.ie

Abstract—In order to improve MAC layer efficiency for future very high-speed WLAN such as IEEE 802.11n, one-to-many aggregation has been proposed where multiple packets destined to multiple receivers are aggregated into one large frame which is then transmitted. On receiving successfully a frame, each receiver in this scheme sends back an ACK in sequence. The overhead caused by the multiple ACK transmissions however greatly restricts the effectiveness of aggregation. Fortunately, due to the recent development of signal processing and antenna array techniques, it is now possible to achieve multipacket reception (MPR) where, even though there are multiple simultaneous transmissions, the physical layer can still separate signals from different users. In this paper, we exploit the use of this technique for decoding multiple ACKs in one-to-many aggregation. In particular, we show by theoretical analysis that while one-to-many aggregation alone achieves reasonable improvements, MPR enabled one-to-many aggregation achieves fundamental improvements over the non-MPR version one. In fact, the effectiveness of such one-to-many aggregation is comparable to that of one-to-one aggregation which is a key technique for achieving high throughput efficiency in future WLANs.

Index Terms—Multipacket reception (MPR), Medium access control (MAC), Cross-layer design, Wireless LAN (WLAN), IEEE 802.11, IEEE 802.11n.

I. INTRODUCTION

Wireless LANs based on 802.11 technology are becoming increasingly ubiquitous. With the aim of supporting rich multimedia applications such as high-definition television (HDTV, 20Mbps) and DVD (9.8Mbps), the technology trend is towards increasingly greater bandwidths. Some recent 802.11n proposals seek to support PHY rates of up to 600 Mbps ([2], [3], [4]). However, higher PHY rates do not necessarily translate into corresponding increases in MAC layer throughput. Indeed, it is well known that the MAC efficiency of 802.11 typically decreases with increasing PHY rate [5], [12]. The reason is that while increasing PHY rates lead to faster transmission of the MAC frame payload, overhead such as PHY headers and contention time typically do not decrease at the same rate and thus begin to dominate frame transmission times. To mitigate the impact of overhead, we have developed a novel scheme called Aggregation with Fragment Retransmission (AFR) in [10] [8] and [9] (Fig. 1). In the AFR scheme, multiple packets are aggregated into and transmitted as a single large frame¹. If errors occur during transmission, only the corrupted fragments of the frame are retransmitted.

¹We define a *packet* as what MAC receives from the upper layer, a *frame* as what MAC transfers to the PHY layer, and a *fragment* as part(s) of a frame.

One concern for the one-to-one aggregation is that there may not be enough traffic between any single pair of stations for aggregation into large frames to be feasible. In this case, one-to-many aggregation may be an option. A simple such scheme called MMP (Multiple Receiver Aggregate Multi-Poll MPDU) (Fig. 2) is mentioned in the TGn Sync proposal to the IEEE 802.11n [2] working group. In this scheme, multiple sequential ACKs are used as feedback for the frame transmission. We identify fundamental properties that must be satisfied by any CSMA/CA based aggregation, and show that MMP violates the basic scaling requirement of aggregation and hence inefficiency can be expected with increasing PHY rate. The reason is simply due to the sequential ACKs transmissions, each of which carries a significant overhead that does not scale with the PHY rate.

Fortunately, the multipacket reception (MPR) ability, which is enabled by recent developments in signal processing and antenna array techniques, can be used to mitigate the ACK overhead [7] [11]. In [14], the authors enhance the throughput efficiency of upload traffic in WLANs by exploiting the MPR capability. More specifically, orthogonal training sequences are assigned to each STA by the AP in a WLAN, where the AP is equipped with multiple antennas and each STA has only one antenna. STAs then transmit data to the AP simultaneously but each with a distinct training sequence which makes it possible for the AP to successfully decode all transmissions. In order to assign the special training sequences, the authors propose to use RTS/CTS before data transmissions. For very high-speed WLANs, however, the use of RTS/CTS may create an excessive overhead and thus should be avoided if possible.

In this paper, we propose the use of MPR for ACK transmissions in a one-to-many aggregation scheme in order to avoid the large overhead associated with multiple sequential ACKs. In particular, in a WLAN the AP manages all of the STAs and is able to assign training sequences without RTS/CTS. This can be achieved by, for example, adding an extra field in the MAC header of each frame. STAs that decode correctly the information directed to themselves may initiate ACK transmissions at the same time due to the MPR ability of the AP, so that the MPR-enabled one-to-many aggregation behaves just like one-to-one aggregation, and thus satisfies the scaling requirements. We develop, based on our previous one-to-one aggregation scheme called AFR, a one-to-many AFR and enhance it with MPR capability (Fig. 3). An analytical model is derived to evaluate the throughput and delay performance




Fig. 1. One-to-one aggregation.




Fig. 2. One-to-many aggregation.




Fig. 3. MPR enabled one-to-many aggregation.

of the one-to-many AFR scheme with and without MPR. Results show that AFR with MPR significantly outperforms that without MPR.

In the upcoming IEEE 802.11n standard, the use of both multiple antennas and aggregation is proposed to be obligatory for high efficiency. Our proposed MPR-enabled aggregation scheme combines both of these techniques and thus fits naturally within the proposed 802.11 framework.

The remainder of the paper is organized as follows. In Section II we identify the fundamental properties that must be satisfied by all CSMA/CA based aggregation schemes. Section III describes the one-to-many AFR scheme with and without MPR. A theoretical analysis is given in Section IV. Finally we summarise our conclusions in Section V.

II. BASIC CONSIDERATIONS OF AGGREGATION

For any aggregation (no matter one-to-one or one-to-many aggregation), the basic requirement for high efficiency is to aggregate packets into large frames so as to spread the cost of fixed overheads across multiple packets. To reduce the overhead associated with transmission errors, each frame can be sub-divided into fragments, with packets that exceed the fragment size being divided. Fragments are the unit used in the retransmission logic, i.e., damaged fragments are retransmitted rather than the entire frame.

The time to transmit a packet is $T_p = L_p/R$, where L_p is the packet size and R is the PHY rate. Hence, the per packet throughput efficiency is

$$\eta_p = \frac{T_p}{T_p + T_{oh}^p} = \frac{L_p/R}{L_p/R + T_{oh}^p} \quad (1)$$

We can see that $T_p = L_p/R$ scales with $1/R$. In order to maintain throughput efficiency η_p , we require that the per packet overhead T_{oh}^p also scales with $1/R$. Considering T_{oh}^p in more detail, we can typically decompose it into the following elements (where r denotes the number of transmissions before all fragments from this packet are transmitted successfully, and other notation is listed in Table I):

$$T_{oh}^p = \frac{(T_{hdr}^{phy} + T_{hdr}^{mac} + T_{hdr}^{frag} + T_{CW} + T_{ack}) \cdot r}{M} \quad (2)$$

To ensure that T_{oh}^p scales with $1/R$, we require that:

- The number of packets M in a frame should be proportional to R , that is $M = bR$ for some constant b . This ensures that the overhead T_{hdr}^{phy} , T_{hdr}^{mac} , T_{ack} and T_{CW} translate into a per packet overhead that scales with R .

- Since there is only one MAC header, when M is proportional to R there is no fundamental constraint on the rate at which MAC headers are transmitted. The same is not true for fragment headers.
- For a given fragment size L_{frag} , the number of fragments in a frame m increases with the number of packets M in a frame, i.e., $m = m'M$ where m' is the number of fragments per packet, we thus have $m = m'bR$ when $M = bR$. Hence, for T_{hdr}^{frag}/M to scale with $1/R$ the rate at which fragment headers are transmitted must be chosen proportional to R , in which case $T_{hdr}^{frag}/M = mL_1/R = m'L_1/R$.

For one-to-one aggregation schemes such as AFR, all of the above requirements can be satisfied if the scheme properly designed. For one-to-many aggregation without MPR, however, the first rule that M scales with R is violated since when packets in one frame go to M distinct receivers then $T_{ack} = MT'_{ack}$ where T'_{ack} denotes the time taken to transmit one ACK frame, i.e., even though M is proportional to R , the ACK overhead T_{ack} does not scale with R , hence a low efficiency η_p is expected.

Fortunately, the multipacket reception (MPR) ability can be used to mitigate the overhead caused by this sequential ACK

n	Number of STAs
M	Number of packets in a frame
m	Number of fragments in a frame
m'	Number of fragments in a packet
T_{CW}	Contention time
T_{SIFS}	Time duration of SIFS
T_{DIFS}	Time duration of DIFS
T_{ack}	Overhead for transmitting an ACK frame ^a
T_{EIFS}	Time duration of EIFS ^b
T_{hdr}^{phy}	Time duration to transmit the PHY headers of a frame
T_p	Time duration to transmit one packet
T_f	Time duration to transmit payload of one frame
T_{oh}^p	Overhead for transmitting one packet
σ	PHY layer time slot
L_f	Payload size in one frame (bytes)
L_p	Packet size (bytes)
L_1	Fragment header size (bytes)

TABLE I

NOTATION USED IN THIS PAPER.

^a $T_{ack} = T_{SIFS} + T_{hdr}^{phy} + T_{ack}^{pld}$, where T_{ack}^{pld} denotes the time duration to transmit an ACK frame. Note that we define T_{ack} in this way for notation brevity.

^b $T_{EIFS} = T_{ack}$

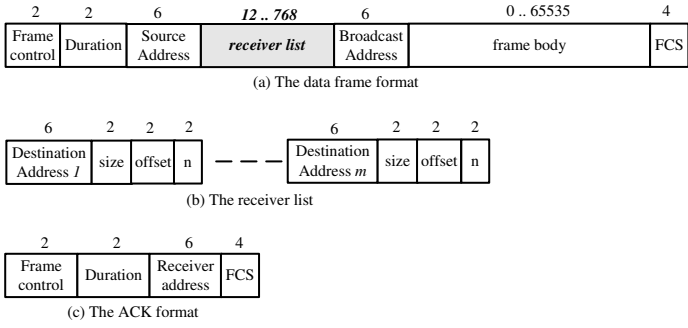


Fig. 4. The data format of one-to-many AFR.

transmission. A MPR-enabled STA can separate signals from different users even if there are multiple transmissions. There are a few techniques can provide MPR ability: separating signals at modulation level (e.g., CDMA), spatial multiplexing enabled by antenna array, and geographical locations of users [13] [11].

For the future very high-speed WLANs, multiple antennas has been proposed to be obligatory for high efficiency. Meanwhile, aggregation is also proposed. The combination of both techniques however is missing. If properly designed, the multiple ACKs for a one-to-many aggregation can be send and received simultaneously. As a result, such MPR-enabled one-to-many aggregation would satisfy all the scaling requirements above.

III. THE ONE-TO-MANY AFR SCHEME

Based on the insight provided by the foregoing analysis, we describe in detail the one-to-many AFR scheme with and without MPR ability, which is a naturally extension of our previous one-to-one aggregation scheme called AFR [8] [9].

As in the legacy DCF scheme of 802.11, there are four types of frame in the one-to-many AFR scheme: RTS, CTS, data frame, and ACK (See Fig. 4 for the format of the latter two. RTS/CTS is not shown because they are not preferred in very high-speed WLANs.).

The data frame consists of *MAC header*, *receiver-list*, *frame body* and *FCS* fields. In the MAC header, the *broadcast address* is used to fill the original 802.11's destination address so that all the STAs are able to decode it. The STAs to which the AP is sending information are recorded in the *receiver list*. In the *receiver-list*, the *destination address*, *size*, *offset*, and *n* fields records respectively the destinations of this frame, the size of each packet, the start position of each packet in the frame, and the number of same-sized packets destined to the same destination. The last field *n* is designed for the case where there are more than one packets destined for a STA. The *frame body* is the aggregate of packets that will be transmitted in this frame. Note that since the MAC header and the *receiver list* are gathered together, robust FEC can easily be used to protect them.

For one-to-many AFR without MPR, binary exponential backoff of the MAC contention window may happen in two

T_{SIFS} (μs)	16
Idle slot duration (σ) (μs)	9
T_{DIFS} (μs)	34
$T_{PHY_{hdr}}$ (μs)	20
CW_{min}	16
Propagation delay (μs)	1
Symbol delay (μs)	4
Retry limit	4
Packet size (bytes)	1024
PHY data rate (Mbps)	216
PHY basic rate (Mbps)	216

TABLE II
MAC AND PHY PARAMETERS USED IN THIS PAPER.

cases. One is when more than one node starts transmissions at the same time so that none of the senders receive any ACKs from receivers. In this case, the sender STAs double their contention window sizes as in the legacy DCF [1]. Doubling window sizes is also possible if there is only one sender but none of the receivers are able to decode the information correctly. For the MPR-enabled one-to-many AFR, however, the retransmission logic is exactly the same as the legacy DCF scheme.

IV. PERFORMANCE GAIN

To evaluate the performance of the new scheme, we extend our previous analysis [9] of the one-to-one AFR scheme, which has been verified against NS simulations, to model the one-to-many AFR scheme with and without MPR.

For simplicity, with one-to-many traffic we assume that in each frame there is one packet for each n destinations. This assumption can be readily relaxed to include more general situations, at the expense of more complex notation in the model. We also assume that senders are saturated i.e. always have enough packets to fill a frame. Finally, we assume that the number of antennas at the sender is equal to the number of receivers – again, this assumption can be relaxed, but simplifies exposition of the model.

A. Throughput Performance

The throughput is defined as the expected payload size of a successfully transmitted frame in an expected slot duration, i.e., $S = E[L_f]/E[T]$. We first compute the expected slot duration $E[T]$. There are three kinds of duration in a WLAN if we assume the channel is error free as in [6]:

- Let n be the number of nodes in the network. If none of them transmits any frame, they all wait for an idle duration T_I , the length of which is the default slot duration.
- Let T_C denote the duration during which at least one node transmits. In this case, the channel is kept busy for the time taken to transmit a frame and the corresponding ACKs. For the one-to-many AFR scheme without MPR, there are altogether n' packets are aggregated in one frame and there are n' ACK transmit durations associated




Fig. 5. The x-axis represents the number of receivers. The y-axis represents the MAC throughput. Parameters used are listed in Table II.

with each frame transmission. For MPR-enabled one-to-many AFR, each frame contains n' packets and only one ACK transmit duration is incurred.

- In the case of successful transmissions, for simplicity the duration is taken to be the same as in the collision case.

Therefore, the slot durations for one-to-many AFR without MPR are (the notation is listed in Table I):

$$\begin{aligned} T_I &= \sigma \\ T_C &= T_{DIFS} + T_{hdr}^{phy} + n'T_p + n'T_{ack} \\ T_S &= T_C \end{aligned} \quad (3)$$

The durations for MPR-enabled one-to-many AFR are:

$$\begin{aligned} T'_I &= \sigma \\ T'_C &= T_{DIFS} + T_{hdr}^{phy} + n'T_p + T_{ack} \\ T'_S &= T'_C \end{aligned} \quad (4)$$

Let τ denote a STA's transmission probability in a slot, the corresponding possibilities for these durations are (The method to solve τ is the same as in [6] and [9]):

$$\begin{aligned} P_I &= (1 - \tau)^n \\ P_S &= n \cdot (\tau(1 - \tau)^{n-1}) \\ P_C &= 1 - P_I - P_S \end{aligned} \quad (5)$$

Therefore, the throughput of one-to-many AFR without MPR S , and MPR-enabled one-to-many AFR S_{mpr} are:

$$S = \frac{P_S \cdot n' \cdot L_p}{P_I T_I + P_S T_S + P_C T_C} \quad (6)$$

$$S_{mpr} = \frac{P_S \cdot n' \cdot L_p}{P_I T'_I + P_S T'_S + P_C T'_C} \quad (7)$$

In Fig. 5, the throughputs of one-to-many AFR with and without MPR are plotted against the number of receivers is varied. As expected, the MPR-enabled version achieves fundamental improvement, e.g., around 60% with 8 receivers.

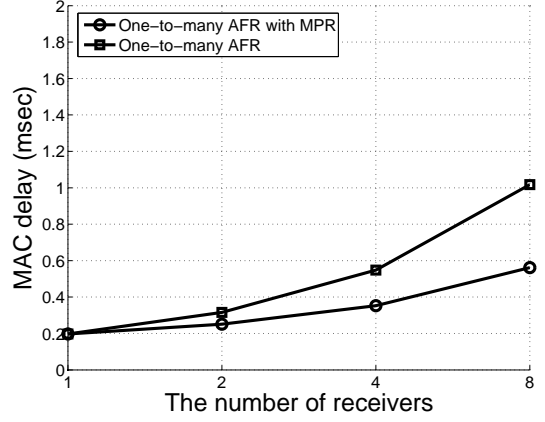


Fig. 6. The x-axis represents the number of receivers. The y-axis represents the MAC throughput. Parameters used are listed in Table II.

B. Delay Analysis

Our model can be extended to estimate the MAC layer delay, i.e., the mean time between a packet reaching the head of the MAC interface queue and being successfully transmitted. Let S^{frame} be the system throughput in frames-per-second rather than bits-per-second. That is, the MAC layer can transport S^{frame} frames in one second, thus the delay to successfully transmit one frame is $1/S^{frame}$, where

$$S^{frame} = \frac{E[\text{number of frames}]}{E[T]}. \quad (8)$$

In the AFR scheme, a packet is fragmented and may be only partially transmitted in one transmission. Thus, we need to know the mean delay before all fragments of a packet are successfully transmitted. Each fragment will be successfully transmitted in $\leq r'$ successful frame transmissions with probability (where p_e^{frag} denoted the error probability of a fragment.)

$$\begin{aligned} &(1 - p_e^{frag}) + (p_e^{frag})(1 - p_e^{frag}) + \dots + (p_e^{frag})^{r'-1}(1 - p_e^{frag}) \\ &= 1 - (p_e^{frag})^{r'}. \end{aligned} \quad (9)$$

Suppose that a packet arrives and is divided into m' fragments. The probability of successfully transmitting m' fragments in $\leq r'$ attempts is $(1 - (p_e^{frag})^{r'})^{m'}$. Further, assuming that errors are independent, the probability of transmitting a packet in exactly r' attempts is $(1 - (p_e^{frag})^{r'})^{m'} - (1 - (p_e^{frag})^{r'-1})^{m'}$. So the expected number of retransmission attempts can be written as

$$r = \sum_{r'=1}^{\infty} r' \left[(1 - (p_e^{frag})^{r'})^{m'} - (1 - (p_e^{frag})^{r'-1})^{m'} \right]. \quad (10)$$

Here, the sum may be truncated to account for the finite number of retransmission attempts. Therefore we have that the per packet MAC delay for one-to-many AFR without MPR D

and that of MPR-enabled version D_{mpr} are

$$D = r \cdot \frac{P_I T_I + P_3 T_3 + P_C T_C}{P_3}. \quad (11)$$

$$D_{mpr} = r \cdot \frac{P_I T'_I + P_3 T'_3 + P_C T'_C}{P_3}. \quad (12)$$

The delay performance of both schemes when the bit error rate is 0 is illustrated in Fig. 6 where we can see that the MPR-enabled version enjoys low delay with increased number of receivers.

V. CONCLUSION

In order to improve the MAC layer efficiency in future very high-speed WLAN such as IEEE 802.11n, one-to-many aggregation has previously been proposed. In this approach, multiple packets destined to multiple receivers are aggregated into a single large frame which is then transmitted. On successfully receiving a frame, receivers in this scheme each send back an ACK sequentially. The overhead caused by the multiple ACK transmissions however greatly restricts the effectiveness of aggregation. In this paper we exploit recent developments in signal processing and antenna array techniques which mean that it is now possible to achieve multipacket reception (MPR) where, even though there are multiple simultaneous transmissions, the physical layer can still separate signals from different users. A one-to-many MPR-enabled AFR scheme is proposed. We show by theoretical analysis that MPR enabled one-to-many aggregation achieves fundamental improvements over non-MPR schemes.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the support of Science Foundation Ireland grant IN3/03/I346.

REFERENCES

- [1] IEEE std 802.11-1999, "Part 11: wireless LAN MAC and physical layer specifications", reference number ISO/IEC 8802-11:1999(E), IEEE Std 802.11, 1999.
- [2] S. A. Mujtaba, et. al., "TGn Sync Proposal Technical Specification", www.tgnsync.org, IEEE 802.11-04/889r6, May. 2005.
- [3] J. Ketchum, et. al., "System Description and Operating Principles for High Throughput Enhancements to 802.11", IEEE 802.11-04/0870r0, Aug. 2004.
- [4] M. Singh, B. Edwards, et. al., "System Description and Operating Principles for High Throughput Enhancements to 802.11", IEEE 802.11-04-0886-00-000n, Aug. 2004.
- [5] Magis Networks White Paper, "IEEE 802.11 e/a Throughput Analysis", 2004, www.magisnetworks.com.
- [6] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function", IEEE Journal on Selected Areas in Communications, Vol. 18, Number 3, March 2000.
- [7] S. Ghez, S. Verdu, and S. C. Schwartz, "Stability Properties of Slotted Aloha with Multipacket Reception Capability", IEEE Transactions on Automatic Control, vol. 33, no. 7, pp. 640-649, July 1988.
- [8] T. Li, Q. Ni, D. Malone, D. Leith, Y. Xiao, and T. Turletti, "A New MAC Scheme for Very High-Speed WLANs". IEEE WOWMOM 2006.
- [9] T. Li, Q. Ni, D. Malone, D. Leith, Y. Xiao, and T. Turletti, "Aggregation with Fragment Retransmission for Very High-Speed WLANs". submitted to IEEE/ACM Transactions on Networking.
- [10] Q. Ni, T. Li, T. Turletti and Y. Xiao, "AFR partial MAC proposal for IEEE 802.11n", IEEE 802.11-04-0950-00-000n, Aug. 2004.
- [11] 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, Nov. 2001.
- [12] Y. Xiao and J. Rosdahl, "Performance analysis and enhancement for the current and future IEEE 802.11 MAC protocols", ACM SIGMOBILE Mobile Computing and Communications Review (MC2R), special issue on Wireless Home Networks, Vol. 7, No. 2, Apr. 2003, pp. 6-19.
- [13] Q. Zhao, L. Tong, "Semi-Blind Collision Resolution in Random Access Wireless Ad Hoc Networks", IEEE Transactions on Signal Processing, VOL. 48, NO. 10, Oct. 2000.
- [14] P. X. Zheng, Y. J. Zhang and S. C. Liew, "Multipacket Reception in Wireless Local Area Networks", IEEE ICC 2006.