IEEE 802.11n MAC Frame Aggregation Mechanisms for Next-Generation High-Throughput WLANs
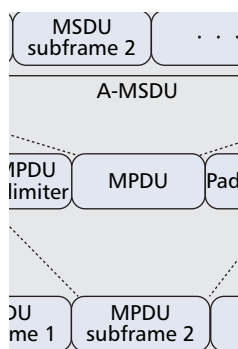
Dionysios Skordoulis and Qiang Ni, Brunel University
Hsiao-Hwa Chen, National Sun Yat-Sen University
Adrian P. Stephens, Intel, UK
Changwen Liu, Entropic Communications
Abbas Jamalipour, University of Sydney

The authors investigate the key MAC enhancements that help 802.11n achieve high throughput and high efficiency. A detailed description is given for various frame aggregation mechanisms proposed in the latest 802.11n draft standard.

Abstract

IEEE 802.11n is an ongoing next-generation wireless LAN standard that supports a very high-speed connection with more than 100 Mb/s data throughput measured at the medium access control layer. This article investigates the key MAC enhancements that help 802.11n achieve high throughput and high efficiency. A detailed description is given for various frame aggregation mechanisms proposed in the latest 802.11n draft standard. Our simulation results confirm that A-MSDU, A-MPDU, and a combination of these methods improve extensively the channel efficiency and data throughput. We analyze the performance of each frame aggregation scheme in distinct scenarios, and we conclude that overall, the two-level aggregation is the most efficacious.

Introduction

Over the last decade, the applications of wireless and cellular devices have expanded rapidly. The most popular network in the wireless domain is the IEEE 802.11 wireless local area network (WLAN) [1] mainly because of the advantages that its systems possess. Its key characteristics, such as interoperability, mobility, flexibility, and cost-effective deployment, have led to its gaining vast support across enterprises, the public sector, homes, and data service providers. However, the low efficiency of its medium access control (MAC) and physical (PHY) layer protocols restricts its applications to support high data rate multimedia services. Current WLAN systems endure difficulties with the increasing expectations of end users and with volatile bandwidth and delay-boundary demands from new higher data rate services, such as high-definition televi-

sion (HDTV), video teleconferencing, multimedia streaming, voice over IP (VoIP), file transfer, and online gaming.

In July 2002, the IEEE 802.11 standard working group established the High-Throughput Study Group (HTSG) with the aim to achieve higher data rate solutions by means of existing PHY and MAC mechanisms [2, 3]. Its first interest was to achieve a MAC data throughput over 100 Mb/s using the 802.11a standard. However, the objective proved to be infeasible as the estimated throughput bounds are well below the theoretical maximum link rate because of the existing MAC and PHY overhead [4, 5]. So, in September 2003, the HTSG set off the IEEE 802.11n ("n" represents next-generation) resolution to compose a high-throughput (HT) extension of the current WLAN standard that will increase transmission rate and reduce compulsory overhead. The main goal of the IEEE 802.11n Task Group (TGn) is to define an amendment that will have maximum data throughput of at least 100 Mb/s, as measured at the MAC data service access point (SAP), and at the same time, to allow coexistence with legacy devices. Some of their recent propositions [6] for the PHY include multiple input multiple output (MIMO) antennas with orthogonal frequency division multiplexing (OFDM) and various channel binding schemes. For the MAC, the main developments are the introduction of frame aggregation at the same time as multiple protection schemes have been designed to allow coexistence of "n" with *legacy* devices.

In this article we mainly focus on the MAC frame aggregation methods that TGn has proposed in the latest 802.11n standard draft, and how these can improve WLANs by way of higher throughput and maximized efficiency. Thus, we begin with a brief outline of the current IEEE
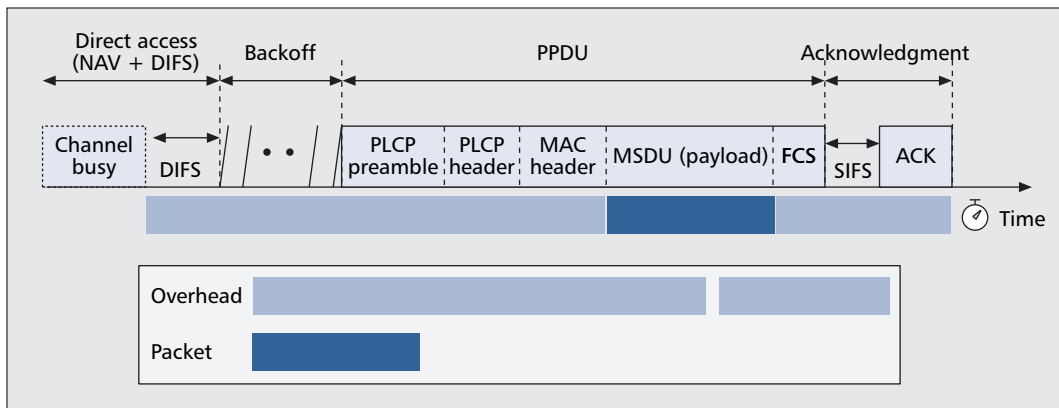
■ Figure 1. *DCF basic operation.*

802.11 WLAN standards, followed by a discussion of its maximal throughput limitations because of overhead. We then concentrate on the latest work of 802.11n on both PHY and MAC, but elaborate more on the latter. We also evaluate, via simulations, the improved performance when using frame aggregation and its efficiency over distinct scenarios. Finally, we conclude by summarizing this article's findings.

OVERVIEW OF IEEE 802.11
IEEE 802.11 PHY

The IEEE 802.11 PHY layer specification concentrates mainly on wireless transmission and concurrently, performs secondary functions, such as assessing the state of the wireless medium and reporting it back to the MAC sublayer. The original specification was first approved in 1997 [1] and includes a primitive MAC architecture and three basic over-the-air communication techniques with maximal raw data rates of 1 and 2 Mb/s. Because of their fairly low data bandwidths, further amendments have been proposed throughout the years: IEEE 802.11a [7], 802.11b, and 802.11g [8]. Both 802.11a and 802.11b were finalized in 1999 and support raw data rates up to 11 Mb/s and 54 Mb/s, respectively. In June 2003, a third PHY specification (802.11g) was introduced, with similar maximum raw data rate as 802.11a but operating in separate frequency bands. For this period, there were many amendments and countless research works for improved PHY specifications that mostly aim to provide reliable connections and higher data rates. This is mainly because there is a continuous rapid increase in user demand for faster connections.

In spite of establishing novel techniques that theoretically can be used for higher data transmission rates, the throughput outcomes at the MAC data SAP are surprisingly low and in most cases, half of what the underlying PHY rates can offer [4].

IEEE 802.11 MAC

The MAC architecture is based on the logical coordination functions that determine who and when to access the wireless medium at any time. It supports fragmentation and encryption and acts as an interface between the logical link con-

trol (LLC) sublayer and the PHY layer. In the legacy IEEE 802.11 standard, there are two types of access schemes: the mandatory distributed coordination function (DCF), which is based on the carrier sense multiple access with collision avoidance (CSMA/CA) mechanism; and the optional point coordination function (PCF), which is based on a poll-and-response mechanism. Because these MAC schemes are inadequate to resolve differentiation and prioritization between frames and multimedia applications such as VoIP and audio/video conferencing with strict performance constraints have become widely popular, a new extension was vital. In late 2005, IEEE 802.11 Working Group (WG) approved the IEEE 802.11e amendment [9] to provide an acceptable level of quality of service (QoS) for multimedia applications. The 802.11e proposes the hybrid coordination function (HCF), which uses a contention-based channel access method, known as enhanced DCF channel access (EDCA). EDCA has the ability to operate simultaneously with a polling-based HCF controlled channel access (HCCA). In addition to the differentiation and prioritization that IEEE 802.11e offers, the transmission opportunity (TXOP), an interval of time in which multiple data frames can be transferred from one STAtion (STA) to another — also known as bursting — was introduced as a way to improve MAC efficiency. Along with frame bursting, another type of acknowledgment (ACK), known as block ACK, was established so receivers can acknowledge multiple received data frames sufficiently and economically by using just a single extended ACK frame. The traffic flows are characterized through the traffic identifiers (TID) and the traffic specification (TSPEC) frames, a set of information that specifies the corresponding QoS requirements.

THROUGHPUT LIMITATIONS

To understand the inefficiency of IEEE 802.11 over higher data rates, we must briefly describe the legacy DCF. This method operates with a first-in-first-out (FIFO) transmission queue that is situated for receiving and buffering incoming data from the higher layers. The basic operation of DCF is illustrated in Fig. 1. After a frame, also known as a MAC service data unit (MSDU), arrives from the LLC at the head of the trans-

The key underlying model defined in the specification for the PHY layer operates multiple antennas for both transmitter and receiver. The use of MIMO provides many benefits, such as antenna diversity and spatial multiplexing.

mission queue, the DCF operation instructs the MAC to wait for a global defined interframe interval called DCF interframe space (DIFS) before any other actions can be taken. If the PHY reports back to the MAC that the wireless channel is busy, the STA's MAC halts until the medium becomes free. On the other hand, if the medium remains idle during DIFS deference, the STA enters a back-off procedure where a slot is selected from a random back-off counter within a contention window (CW). Next, the counter starts a decrement process while the channel remains idle for each slot interval. When the counter reaches zero, the STA obtains an affirmation to send the information through the wireless link. Now, each STA that receives a data frame utilizes an error checking processes to detect the presence of errors. If no errors are found, it sends back an ACK frame after a specified short interframe space (SIFS) to verify that the information was successfully received. If the sending STA does not receive an ACK after SIFS, it will assume that the communication was broken or interfered, and it will start a new DCF process for retransmission. If there is a case of collision, then the MAC extends its CW, selects a new slot, and repeats the previous steps. Finally, there is an optional mechanism known as request to send/clear to send (RTS/CTS) that intends to resolve the so-called hidden and exposed node scenarios that usually occur in ad hoc networks. With RTS/CTS, after a STA is granted access to transmit, it first sends an RTS frame and then holds back for the CTS response from the receiver. Obviously, this situation can be disadvantageous if actual data frame size is small because the RTS/CTS exchange produces further overhead and consequently reduces the effective throughput.

We can comprehend clearly the consequences of that hefty overhead on the system throughput from Fig. 1 (RTS/CTS has been omitted for simplicity). The figure illustrates the required procedure that each single packet traverses from the time it arrives at the MAC until it is successfully received by the receiver, with different headers and tails added on the original payload over different sublayers. Note that [1] states that all physical layer convergence protocol (PLCP) preambles and PLCP headers shall be transmitted using the basic (most of the time, the minimum) link rate, which is much less than the rate used for data transmission. A complete transmission cycle of a simple DCF consists of DIFS deferral, back off, data transmission, SIFS deferral, ACK transmission, and propagation delay. As shown in Fig. 1, to transmit a data packet, a large overhead is associated. Be aware that the overhead shown does not correspond to real time lengths, as payload varies, but it shows the additional time that is required to have a successful transmission. So, the higher the packet rate — meaning the number of packets that are injected to the MAC per second — the higher the relative overhead the system introduces. At this point the analysis in [5] has shown that the maximum ideal throughput is bounded at 50 Mb/s when the frame size is 1 KB (or 1024 B), and the link rate increases to infinite. In general, when a legacy device uses a 54 Mb/s link, it can

provide only a maximum relative MAC throughput that is less than 50 percent of the peak PHY or around 25 Mb/s. Consequently, it was confirmed that to reach the TGn target of an achievable MAC layer throughput higher than 100 Mb/s, it is essential to adopt innovative MAC techniques. Because most of these types of MAC and PHY overheads are essential, we must focus on how we can reduce the frequency of them during the transmission of multiple packets. This can be achieved by concatenating or packing them together and adding the overheads over a group of packets rather separate ones, and this is known as aggregation. Various methods of aggregation (e.g., [5, 6, 10–13]) have been proposed, but all of them follow a similar logic. In the next section, we will describe the TGn proposal for frame aggregation.

IEEE 802.11N ARCHITECTURE
IMPROVEMENTS ON PHY LAYER

The key underlying model defined in the specification for the PHY layer operates multiple antennas for both transmitter and receiver. The use of MIMO provides many benefits, such as antenna diversity and spatial multiplexing. In general, MIMO technology increases the spectral efficiency of a wireless communication system. Traditional single-input single-output (SISO) systems tend to immobilize certain propagation phenomena, such as multipath propagation. Multipath is typically perceived as interference degrading the ability of a receiver to recover intelligent information. However, by using multiple antennas, an exploitation of the multipath phenomena can turn advantageous as the data throughput and range increases, and the bit error rate decreases. MIMO technology has the ability to simultaneously resolve information from multiple signal paths using spatially separated receive antennas.

Another important proposition increases the bandwidth of the current channel from 20 MHz to 40 MHz. Using a wider channel bandwidth, we improve the theoretical capacity limits; this can easily be seen from Shannon's capacity equation $C = B \log_2(1 + SNR)$. Thus, if properly implemented, the 40 MHz channels can become more desirable than two times the usable channel bandwidth of two 802.11 legacy channels. Diversity in a multi-antenna system is achievable either through space-time coding or using the channel state information at the transmitter intelligently. Before transmission, a session exchange of sounding PHY protocol data units (PPDUs) is required to calibrate the radio channel. Based on the information elements collected from sounding and calibration sequences, beam forming can be used to boost signal quality by selecting the proper power and coding scheme for the spatial streams. Higher signal quality means that a given data rate can be available at longer range. For a given signal-to-noise ratio (SNR), a beam-formed transmission can carry a higher data rate. Finally, transmission uses different coding schemes as they perform analogously, and thus MIMO has two coding schemes, space time block coding (STBC) and low density parity check coding (LDPC).
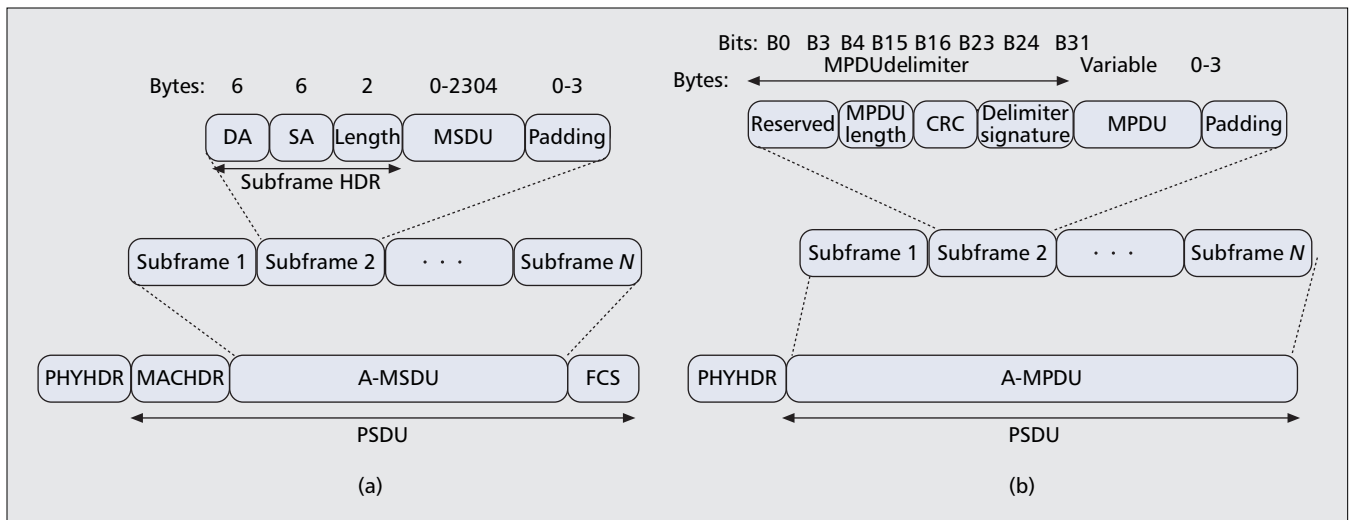
■ Figure 2. *One-level frame aggregation: a) A-MSDU; b) A-MPDU.*

Some of the improvements have been classified as optional because they cause technical hitches when non-802.11n devices are present in the same WLAN. Nonetheless, this matter does not demote their importance; on the contrary if these mechanisms are employed over a network than includes only HT nodes, the PHY rate can increase as much as 600 Mb/s [6].

MAC ENHANCEMENTS

In this subsection, we briefly mention the most important MAC enhancements that are proposed in [6] with a more detailed explanation on frame aggregation, which maximizes throughput and efficiency. There are two methods available to perform frame aggregation: aggregate MAC protocol service unit (A-MSDU) and aggregate MAC protocol data unit (A-MPDU). The main distinction between an MSDU and an MPDU is that the former corresponds to the information that is imported to or exported from the upper part of the MAC sublayer from or to the higher layers, respectively, whereas the later relates to the information that is exchanged from or to the PHY by the lower part of the MAC. Aggregate exchange sequences are made possible with a protocol that acknowledges multiple MPDUs with a single block ACK in response to a block acknowledgment request (BAR). Another key enhancement that the 802.11n specifies is the bidirectional data transfer method over a single TXOP, known as reverse direction. This feature permits the transportation of data frames, even aggregates, in both directions in one TXOP. Until now, when the sender STA is allocated with a TXOP, it informs surrounding STAs about how long the wireless medium will be engaged. However, this approximation of channel use is not always accurate, and often the transmission ends sooner. As a result, contended STAs assume that the channel is still occupied when this is not the case. With reverse direction, the initial receiver STA is allowed to send any packets available that are addressed to the sender for the remaining TXOP time. This fits especially well with TCP because it allows a TCP link to *piggyback* TCP ACK collection onto TCP data transmission. The long network allocation vector (long-NAV) is another enhancement that improves scheduling, given that a station that holds a TXOP may set a longer NAV value intended to protect multiple PPDUs. Another mandatory feature is phased coexistence operation (PCO) which protects stations using either 20 MHz or 40 MHz channel spectrum at the same time. Finally, the reduced IFS (RIFS) is proposed to allow a time interval of 2 μs between multiple PPDUs, which is much shorter than SIFS as defined in the legacy standards.

A-MSDU — The principle of the A-MSDU (or MSDU aggregation) is to allow multiple MSDUs to be sent to the same receiver concatenated in a single MPDU. This definitely improves the efficiency of the MAC layer, specifically when there are many small MSDUs, such as TCP acknowledgments. This supporting function for A-MSDU within the 802.11n is mandatory at the receiver. For an A-MSDU to be formed, a layer at the top of the MAC receives and buffers multiple packets (MSDUs). The A-MSDU is completed either when the size of the waiting packets reaches the maximal A-MSDU threshold or the maximal delay of the oldest packet reaches a pre-assigned value. Its maximum length can be either 3839 or 7935 bytes; this is 256 bytes shorter than the maximum PHY PSDU length (4095 or 8191 bytes, respectively), as predicted space is allocated for future status or control information. The size can be found in the HT capabilities element that is advertised from an HT STA in order to declare its HT status. The maximal delay can be set to an independent value for every AC but is usually set to 1 μs for all ACs. There are also certain constraints when constructing an A-MSDU:

- All MSDUs must have the same TID value
- Lifetime of the A-MSDU should correspond to the maximum lifetime of its constituent elements
- The destination address (DA) and sender address (SA) parameter values in the subframe header must match to the same receiver address (RA) and transmitter address (TA) in the MAC header.
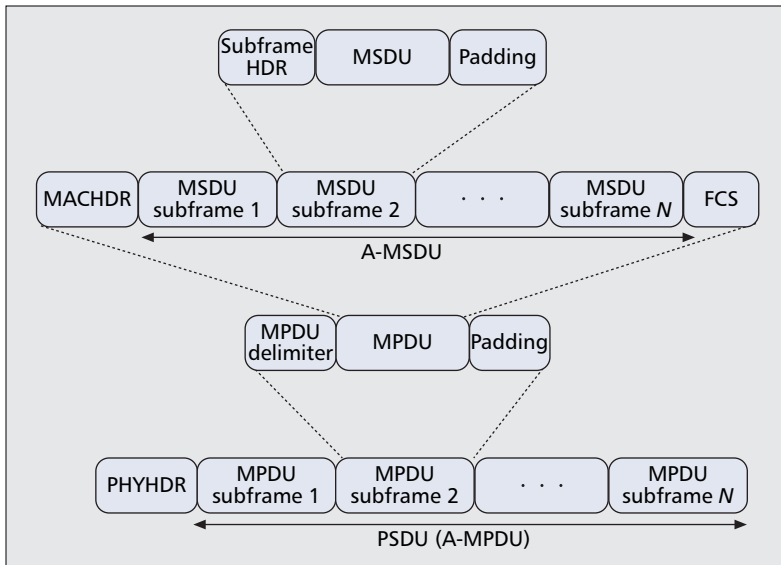
■ Figure 3. *Two-level frame aggregation.*

Thus, broadcasting or multicasting is not allowed.

Figure 2a describes a simple structure of a carrier MPDU that contains an A-MSDU. Each subframe consists of a subframe header followed by the packet that arrived from the LLC and 0 – 3 bytes of padding. The padding size depends on the rule that each subframe, except for the last one, should be a multiple of four bytes, so the end receiver can approximate the beginning of the next subframe. A major drawback of using A-MSDU is under error-prone channels. By compressing all MSDUs into a single MPDU with a single sequence number, for any subframes that are corrupted, the entire A-MSDU must be retransmitted. This situation has been addressed in [5, 13], where additional frame structures or optimum frame sizes have been proposed to improve performance under noisy channels.

A-MPDU — The concept of A-MPDU aggregation is to join multiple MPDU subframes with a single leading PHY header. A key difference from A-MSDU aggregation is that A-MPDU functions after the MAC header encapsulation process. Consequently, the A-MSDU restriction of aggregating frames with matching TIDs is not a factor with A-MPDUs. However, all the MPDUs within an A-MPDU must be addressed to the same receiver address. Also, there is no waiting/holding time to form an A-MPDU so the number of MPDUs to be aggregated totally depends on the number of packets already in the transmission queue. The maximum length that an A-MPDU can obtain — in other words the maximum length of the PSDU that may be received — is 65,535 bytes, but it can be further constrained according to the capabilities of the STA found in the HT capabilities element. The utmost number of subframes that it can hold is 64 because a block ACK bitmap field is 128 bytes in length, where each frame is mapped using two bytes. Note that these two bytes are required to acknowledge up to 16 fragments but because A-MPDU does not allow fragmentation, these extra bits are excessive. As a result, a new variant has been implemented, known as compressed block ACK with a bitmap field of eight bytes long. Finally, the size of each subframe is limited to 4095 bytes as the length of a PPDU cannot exceed the 5.46-ms time limit; this can be derived from the maximum length divided by the lowest PHY rate, which is 6 Mb/s and is the highest duration of an MPDU in 802.11a.

The basic structure is shown in Fig. 2b. A set of fields, known as delimiters are inserted before each MPDU and padding bits varied from 0 – 3 bytes are added at the tail. The basic operation of the delimiter header is to define the MPDU position and length inside the aggregated frame. It is noted that the cyclic redundancy check (CRC) field in the delimiter verifies the authenticity of the 16 preceding bits. The padding bytes are added such that each MPDU is a multiple of four bytes in length, which can assist subframe delineation at the receiver side. In other words, the MPDU delimiters and PAD bytes determine the structure of the A-MPDU. After the A-MPDU is received, a de-aggregation process initiates. First it checks the MPDU delimiter for any errors based on the CRC value. If it is correct, the MPDU is extracted, and it continues with the next subframe till it reaches the end of the PSDU. Otherwise, it checks every four bytes until it locates a valid delimiter or the end of the PSDU. The delimiter signature has a unique pattern to assist the de-aggregation process while scanning for delimiters.

Two-Level Aggregation — A two-level frame aggregation comprises a blend of A-MSDU and A-MPDU over two stages. In Fig. 3 we illustrate how this new scheme can be achieved.

The basic operation is explained as follows: In the first stage, if any MSDUs that are buffered in the A-MSDU provisional storage area justify the A-MSDU constraints explained in the previous related subsection, these data units can be compacted into a single A-MSDU. If the TIDs are different, all these aberrant frames can move to the second stage where they will be packed together with any A-MSDUs derived from the first stage or other single MSDUs by using A-MPDU aggregation. However, it must be mentioned that given that the maximum MPDU length for an A-MPDU data frame is limited to 4095 bytes, then A-MSDUs or MSDUs with lengths larger than this threshold can not be transmitted. Conjointly, any fragments from an A-MSDU or MSDUs also cannot be included in an A-MPDU. In the following section, we evaluate how this synthesis is more efficient in most of the cases than A-MPDU and A-MSDU aggregation operating alone.

PERFORMANCE EVALUATION

In this section, we compare the performance of the latest draft of the IEEE 802.11n A-MSDU and A-MPDU aggregation schemes along with the two-level aggregation. For the simulations, we used a simulation model implemented by Intel, based on the OPNET Modeler [14] with the latest 802.11n PHY and MAC enhancements.

SIMULATION SCENARIOS

The IEEE 802.11n Task Group has predefined specific usage models [15] based on various market-based use cases. The usage models intend to support the definitions of network simulations that will allow them to evaluate performance of various proposals in terms of, for example, network throughput, delay, packet loss, and other metrics. In this subsection, we test the efficiency of aggregation over two high throughput stations. This scenario also is known by the TGn as Scenario 17 (point-to-point goodput test). Here, we study the maximum throughput for each proposed aggregating scheme while increasing the offered load or when the load remains constant but the packet size is increasing. From these scenarios, we also can observe the degrading channel efficiency when aggregation is disabled, and the system uses in full its latest PHY techniques (e.g., MIMO).

The scenario is situated in an infrastructure service area that operates under the EDCA mode. It includes a fixed HT AP and a fixed HT STA, both operating in a 20-MHz channel and using the 64-QAM 3/4 modulation coding scheme (MCS) as their operational data rate with two antennas. The devices are placed over a distance of ten meters with line of sight (LOS). The first station has a synthetic data source that provides varying offered loads (in Mb/s) of User Data Protocol (UDP) traffic. These UDP sources have no timeout values specified, and they all have the same TIDs. To understand in full the remarkable improvement of frame aggregation, we assume that transmission proceeds with no interference or channel fading and that all frames are received successfully, and no retransmission is required. The simulation time is ten seconds, which is an adequate period for multiple packets to be transmitted over the two STAs, taking into consideration that packet arrival at the MAC layer sets off immediately.

SIMULATION RESULTS AND DISCUSSION

Figure 4 illustrates the throughput results (in Mb/s) obtained from the MAC SAP while the offered load (OL) from the associated HT STA accumulates gradually. For the frame aggregation, we define four different settings:
• Both A-MPDU and A-MSDU are enabled
• Only the A-MPDU algorithm is used
• Only the A-MSDU algorithm is used
• No aggregation at all

In the first simulation, the OL is incremented by just varying the packet size while keeping the constant packet generation interval, also known as the constant packet rate (CPR), which is 40 μs. The initial OL is 25 Mb/s, and it increases up to 300 Mb/s with the increase of packet size (i.e., 125/250/500/750/1000/1500 bytes). For each case, the traffic generation rate is configured high enough to saturate the air link rate that corresponds to the "PHY peak" (144 Mb/s) on each figure, and the maximal A-MSDU length is 4 KB. As shown in Fig. 4, we observe that all throughputs increase according to the load. In general, as the packet size increases, the A-MSDU stays below 75 Mb/s, while A-MPDU and the two-level aggregation achieve maximum
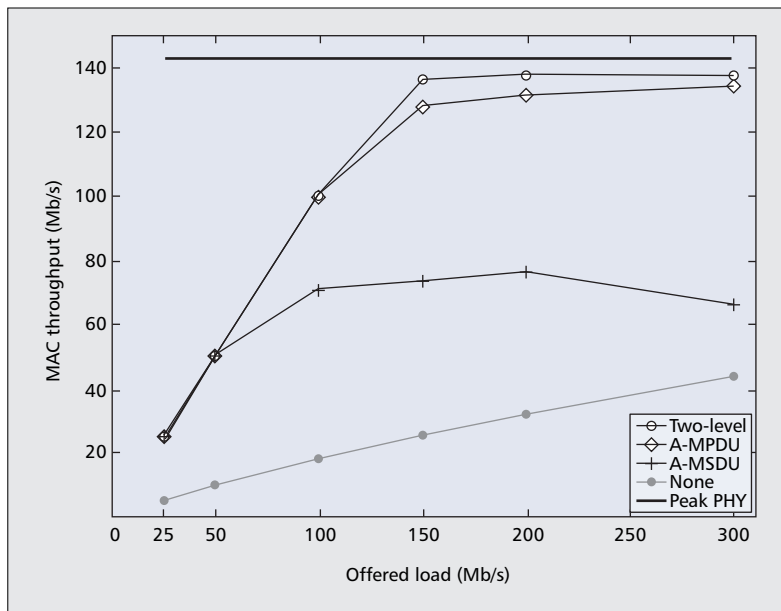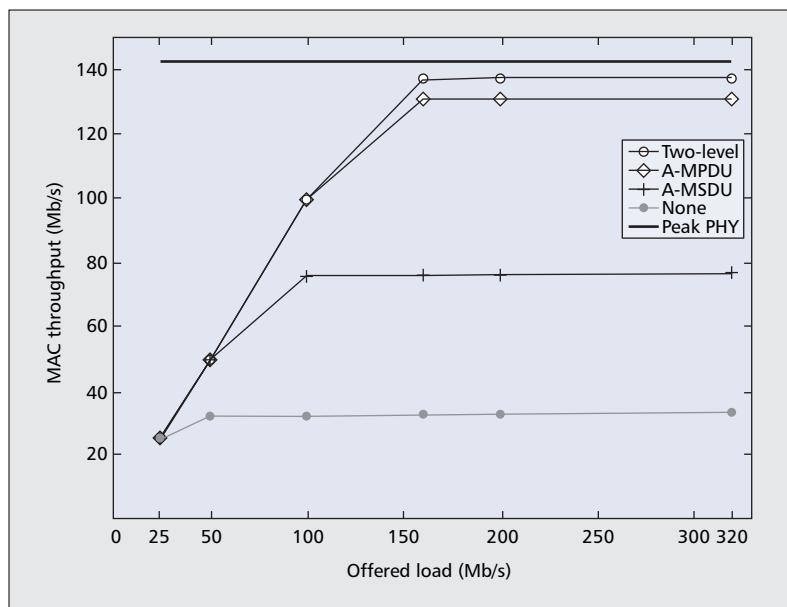

■ **Figure 4.** *Throughput vs. increased offered load by varying the packet size.*

throughputs of 136 and 134 Mb/s, respectively. When the packet size maintains values of 125 B (see OL = 25 Mb/s) and 250 B (see OL = 50 Mb/s), the corresponding throughputs for any type of aggregation are similar. This is because A-MSDU can aggregate several small packets within a single MPDU, even if the length is limited to 4 KB; the same way that A-MPDU can place multiple MPDUs in a single PSDU. Thus, for small packet sizes, we can choose any type of aggregation. On the other hand, when the packet size is larger than 250 KB, the throughput of A-MSDU distinguishes significantly from A-MPDU and the two-layer aggregation, with much lower values because the number of MSDUs that fit into a single A-MSDU is becoming less than the other cases. We can monitor this behavior even more closely when packet size increases from 1000 B (see OL = 200 Mb/s) to 1500 B (see OL = 300 Mb/s). The A-MSDU throughput drops slightly for the reason that we had four packets of 1000 B, fitting in one A-MSDU, where in the case of 1500 B packet size, only two of them can occupy the same space. The same behavior occurs with the two-level aggregation but only because of the A-MSDU stage. This is also why the A-MPDU throughput increases further when two-level aggregation remains at the same levels. The throughput for the no-aggregation case always increases with the increase of the OL by varying the packet size even after the channel is saturated. However, by only increasing the packet size up to the maximum Ethernet transmission unit (1500 bytes), without aggregation, will achieve throughput about three times lower than that of the A-MPDU and the two-level aggregation. This clearly demonstrates that small packet size is the key factor that lowers the throughput efficiency.
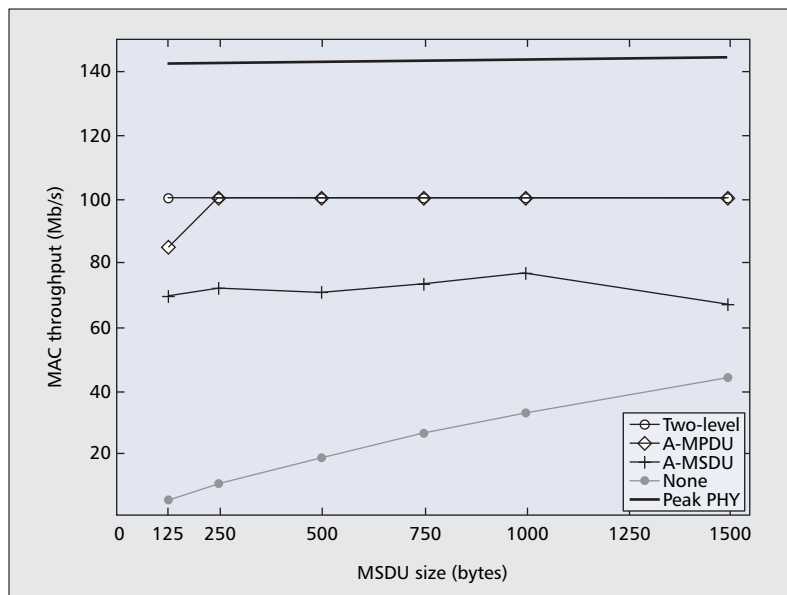
In the second simulation, we increase the OL by altering the packet interarrival rate, that is, variable packet rate (VPR), instead of increasing the packet sizes. So the packet size remains constant at 1,000 B during the simulation test. From

Fig. 5, we can observe that the MAC throughput performance of all the schemes first increases with an increase of the OL. However, after channel is saturated, the throughput of all the schemes remains constant even when the OL keeps increasing, which is different from the first simulation. This behavior is characterized as normal because of channel saturation, meaning the resources are limited to the impending demand. Furthermore, the throughput achieved by the A-MPDU and the two-level aggregation after saturation is approximately 4.5 times higher than the no aggregation scheme.

The last simulation represents a scenario for a fixed OL of 100 Mb/s with variant packet sizes and appropriate interval times. In this simulation, there is no channel saturation. Along with the throughput values for each type of aggrega-

tion, we investigate the degree of aggregation that the two-level aggregation performs, by comparing the number of indicated MSDUs and the received MPDUs. In Fig. 6, the A-MPDU and the two-level aggregation achieve the 100 Mb/s goal that TGn has set, whereas A-MSDU falls below that threshold at around 75 Mb/s. In conclusion, although the A-MSDU mechanism can achieve higher throughput than the legacy 802.11 standards, it does not utilize the channel as fully as A-MPDU and the two-level aggregation do. However, there is an exception for the A-MPDU function when packet size is 125 B and packet rate is 10 µs each. This shows that this type of aggregation cannot handle the bulk traffic effectively, and thus, a lower throughput than the two-level aggregation because the A-MSDU algorithm facilitates to overcome that issue at the first stage of aggregation. It is very important to understand how this blend of A-MSDU and A-MPDU, in most cases, is capable of improving the effectiveness of the MAC layer, specifically when there are many small MSDUs, such as TCP acknowledgments. For example, when the packet size is set for 125 B, there are approximately 999,800 MSDUs generated during the simulation period. In spite of this, the number of MPDUs received at the receiver is 34,500 MPDUs, and that suggests that a huge MAC and PHY overhead was avoided. For the largest packet size of 1500 B, there are approximate 83,300 indicated MSDUs and about 41,650 received MPDUs, a 50 percent improvement in efficiency.

CONCLUSION

This article investigates several frame aggregation schemes proposed in the latest IEEE 802.11n draft, namely, A-MSDU, A-MPDU, and the combined two-level aggregation. Our simulation results demonstrated that concatenation mechanisms performed over different sublayers can actually increase the channel efficiency and reduce the overhead of the future 802.11 MAC.

All types of aggregation schemes are highly recommended as they resolve the fundamental problem of inefficiency due to the PHY/MAC overhead. However, the IEEE 802.11n draft only identifies the basic concepts and data frame structures. In a flawless environment, it can deliver attractive results but in terms of its functionality in a real working environment, there are still many issues that require further investigation. For example, the processing time required to compute these aggregates can increase the overall delays. Actually, as the efficiency of aggregation increases, its operation becomes more complex (e.g., two-level aggregation). Another unresolved issue is how large a concatenation threshold the devices should set. Ideally, the maximum value is preferable but in a noisy environment, short frame lengths are preferred because of potential retransmissions. The A-MPDU concatenation scheme operates only over the packets that are already buffered in the transmission queue, and thus, if the CPR data rate is low, then efficiency also will be small. There are many ongoing studies on alternative queuing mechanisms different from the standard

■ Figure 5. *Throughput vs. increased offered load by varying the packet arrival interval (packet size = 1 kbyte).*

■ Figure 6. *Throughput vs. increased MSDU size.*

FIFO. A combination of frame aggregation and an enhanced queuing algorithm could increase channel efficiency further.

Future work includes examining multidestination aggregation, frame aggregation performing over error-prone channels, and the co-existence of HT stations with the legacy ones.

REFERENCES

[1] IEEE Std. 802.11 WG, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," Aug. 1999 (reaffirmed June 2003).
[2] V. Jones, R. DeVegt, and T. Jerry, "Interest for Higher Data Rates (HDR) Extension to 802.11a," IEEE 802.11n working doc. 802.11-02-081r0, Jan. 2002.
[3] J. Rosdahl, "Draft Project Authorization Request (PAR) for High Throughput Study Group," IEEE 802.11n working doc. 802.11-02/798r2, Mar. 2003.
[4] Y. Xiao and J. Rosdahl, "Throughput and Delay Limits of IEEE 802.11," IEEE Commun. Letters, vol. 6, no. 8, Aug. 2002, pp. 355–57.
[5] Q. Ni et al., "AFR partial MAC proposal for IEEE 802.11n," IEEE 802.11n working doc. 802.11-04-0950-00-000n, Aug. 2004.
[6] IEEE P802.11n, Draft 2.0, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Enhancements for Higher Throughput," Feb. 2007.
[7] IEEE Std. 802.11a, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High-Speed PHY in the 5 GHz Band," 1999 (reaffirmed June 2003).
[8] IEEE Std. 802.11g WG, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Further High Date Rate Extension in the 2.4 GHz Band," June 2003.
[9] IEEE Std. 802.11e WG, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements," Nov. 2005.
[10] Y. Kim et al., "Throughput Enhancement of IEEE 802.11 WLAN via Frame Aggregation," IEEE VTC, vol. 4, Sept. 2004, pp. 3030–34.
[11] Y. Xiao, "Packing Mechanisms for the IEEE 802.11n Wireless LANs," IEEE GLOBECOM '04, pp. 3275–79.
[12] Y. Chen, S. Emeott, and R. Choudhury, "Performance Analysis of Data Aggregation Techniques for Wireless LAN," Motorola Labs, 2005.
[13] Y. Lin and V. Wong, "Frame Aggregation and Optimal Frame Size Adaptation for IEEE 802.11n WLANs," Proc. IEEE GLOBECOM, San Francisco, CA, Nov. 2006, pp. 1–6.
[14] OPNET Technologies, Inc., OPNET Modeler: Accelerating Network R&D; http://www.opnet.com/products/modeler/home-2.html (accessed Feb. 20, 2007).
[15] A. P. Stephens et al., "IEEE P802.11 Wireless LANs: Usage Models," IEEE 802.11n working doc. 802.11-03/802r23, May 2004.

BIOGRAPHIES

DIONYSIOS SKORDOULIS (Dionysios.Skordoulis@brunel.ac.uk) received his B.Sc. in computer science with distributed systems from City University, London, in 2005 and his M.Sc. in wireless communication systems from Brunel University, London, in 2006. Currently, he is working toward his Ph.D., supervised by Dr. Qiang Ni at Brunel University. His main research interests are QoS and performance optimization for current and next-generation WLANs. He has received the IEEE Communications prize for best project and the Rohde & Schwarz award for best overall PG student.

QIANG NI [M'04] (Qiang.Ni@ieee.org) received his Ph.D. in 1999 from Huazhong University of Science and Technology (HUST), China, and subsequently spent two years as a post-doctoral fellow at the Wireless Communication Research Laboratory, HUST. He is currently a lecturer in the School of Engineering & Design, Brunel University, London. Prior to that he was a senior researcher at Hamilton Institute, National University of Ireland Maynooth. His research interests are wireless networking and mobile communications. He has published over 40 refereed papers in his research fields. He worked with INRIA France as a researcher for three years (2001–2004). He has been active as an IEEE 802.11 Wireless Standard Working Group voting member since 2002.

HSIAO-HWA CHEN [SM'00] (hshwchen@ieee.org) is currently a full professor and was the founding director of the Institute of Communications Engineering at National Sun Yat-Sen University, Taiwan. He has authored or co-authored over 200 technical papers in major international journals and conferences, five books and several book chapters in the areas of communications, including the following: Next Generation Wireless Systems and Networks and The Next Generation CDMA Technologies (Wiley, 2005 and 2007). He has been an active volunteer in various IEEE technical activities for over 20 years. Currently, he is serving as chair of the IEEE Communications Society Radio Communications Committee. He served or is serving as symposium chair/co-chair of many major IEEE conferences, including VTC, ICC, GLOBECOM, and WCNC. He has served or is serving as associate editor and/or guest editor of numerous technical journals in communications. He is serving as chief editor (Asia and Pacific) for Wiley's Wireless Communications and Mobile Computing Journal and International Journal of Communication Systems. He is editor-in-chief of Wiley's Security and Communication Networks (http://www.inter-science.wiley.com/journal/security).

ADRIAN P. STEPHENS [M] (Adrian.P.Stephens@intel.com) received his B.A. and Ph.D. from Cambridge University. He is a principal engineer in Intel's Wireless Standards and Technology group (part of the Mobile Wireless group). His focus is on developing next-generation IEEE 802.11 WLAN standards. He is Technical Editor of IEEE 802.11n and coordinated (internally and externally) Intel's MAC proposal for IEEE 802.11n. He has over 20 years of product development and research experience working for the industry and government. He is a member of the IEE and is a chartered engineer.

CHANGWEN LIU (changwen@gmail.com) received his B.S. in mathematics from Hunan Normal University, his M.S. in applied mathematics from HUST, and his Ph.D. in computer science from the University of Illinois at Urbana Champaign in 1995. He is now a staff MAC engineer with Entropic Communications Inc., San Diego, California. His research interests include network protocol design, simulation, and analysis for wireless and wired data communications. From 2002 to 2007 he was a senior research scientist in Intel's Communication Technology Lab.

ABBAS JAMALIPOUR [S'86, M'91, SM'00, F'07] (a.jamalipour@ieee.org) holds a Ph.D. from Nagoya University, Japan. He is the author of the first book on wireless IP and two other books. He has co-authored five books and over 180 journal and conference papers, and holds two patents, all in the field of wireless networks. He is an IEEE Distinguished Lecturer and a Fellow of Engineers Australia. He was chair of the Satellite and Space Communications Technical Committee (2004–2006) and is currently vice chair of the Communications Switching and Routing Technical Committee and chair of the Asia-Pacific Board, Chapters Coordinating Committee. He is a technical editor of IEEE Communications Magazine, Wiley's International Journal of Communication Systems, and several other journals. He is a voting member of the IEEE GITC and was a vice chair of IEEE WCNC '03–'06, chair of IEEE GLOBECOM '05 (Wireless Communications), and a symposium co-chair of IEEE ICC '05–'08 and IEEE GLOBECOM '06–'07, among many other conferences. He is the recipient of several international awards, most recently the Best Tutorial Paper Award and Distinguished Contribution to Satellite Communications Award, both from the IEEE Communications Society in 2006.

There are many ongoing studies on alternative queuing mechanisms different from the standard FIFO. A combination of frame aggregation and an enhanced queuing algorithm could increase channel efficiency further.