Aggregating Packets to Improve 802.11 Wireless Network Performance

M P Morrison, G A D Punchihewa, L De Silva, F M Al-Ali
Institute of Information Sciences & Technology
Massey University, Palmerston North, New Zealand

Email: mpmorrison@iee.org, g.a.punchihewa@massey.ac.nz, l.desilva@massey.ac.nz, f.m.al-ali@massey.ac.nz

Abstract: This paper proposes that packet aggregation can be used to increase the throughput of 802.11 wireless networks. An overview of the current 802.11 standards is given, with a review of the current MAC layer and its effect on throughput. A brief look at typical wireless network traffic is given. Performance improvement through a dynamic RTS/CTS threshold is described. This is followed by a description and comparison of the two types of packet aggregation at the MAC layer – forced delay and congestion triggered. The paper looks at how queues are used within aggregation, with several types of queueing control identified. Simulations were carried out to compare queueing controls and measure the throughput improvement offered by aggregation.

1 Introduction

Currently, IEEE 802.11 wireless networks are gaining more and more acceptance. Currently there are three main extensions to the original 2 Mb/s 802.11 standard: 802.11a operating at 5.2 GHz with 54 Mb/s maximum raw data speed, 802.11b operating at 2.4 GHz with 11 Mb/s maximum raw data speed, and 802.11g operating at 2.4 GHz with a maximum of 54 Mb/s raw data speed. Other extensions have also been ratified by the IEEE, such as 802.11e, but these focus on added functionality rather than improvements in performance. Currently, the IEEE is in the process of developing a new standard designed for higher throughput – 802.11n – and will specify extensions to both the physical and MAC layers. This is scheduled for finalisation in early 2006, and aims for 100 Mb/s actual throughput (as opposed to raw data rate) [1].

However, these maximum speeds advertised are the fastest possible data speed. The medium access coordination functions and headers use lower speeds for enhanced reception and backwards compatibility, and often conditions do not allow data transfer at maximum speeds. Additionally, some of the extensions to 802.11 add overhead or reduce efficiency for extra features, e.g. security in 802.11i, or quality of service in 802.11e.

The performance of wireless networks can be evaluated using many indicators. One of the more important and often most easily recognised indicators is the throughput, or actual data rate, of the network. The technical definition of throughput is the data rate at the MAC service access point. In other words, it is the data rate that the wireless network can offer to the network layer in the OSI network model.

Currently, the throughput of 802.11 wireless networks is not enough to support high bandwidth applications, such as streaming video, or several voice data streams. One of the main influences on throughput in 802.11 is the control on medium access. This is performed by coordination functions that determine when, and for how long, a station can transmit over the wireless medium. The time taken in coordinating the stations in the network influences the network's throughput.

1.1 Coordination Functions

The most basic coordination function, the distributed coordination function (DCF), is specified in the original 802.11 standard [2]. The DCF defines two access mechanisms used in transmission, basic collision sense multiple access with collision avoidance (CSMA/CA) access and a Request-To-Send/Clear-To-Send (RTS/CTS) 'handshake'. For CSMA/CA, a station senses the medium activity - if it is idle, the station transmits. If the medium is busy, the station monitors the medium until it is idle for a 'distributed inter-frame space' (DIFS) time period. The station then waits a random backoff interval before transmitting the data frame.

CSMA/CA is extended if the collision probability is high and the packet size is larger than a set threshold. In this event, the RTS/CTS reservation scheme is used, where short RTS and CTS frames are exchanged to reserve the medium for the transmission of the frame. This scheme shortens the collision duration (collisions can only occur over the period of the short RTS and CTS frames) and copes with hidden stations.

In addition to DCF, an optional, centralised, point coordination function (PCF) was specified [2]. PCF provides contention-free channel access, controlled using polling by a central access point (therefore PCF cannot be used for ad-hoc networks), and has higher precedence than DCF. In PCF, the access point polls each station, and grants channel access if it is requested. Higher overall throughputs can be gained by using PCF by allowing demanding users more access, but this reduces or removes less demanding user's access to the network.

If the optional 802.11e standard is used, replacement coordination functions provide quality of service differentiation for stations. Quality of service is needed in applications where some data requires higher priority – such as voice over IP. Enhanced DCF (EDCF) replaces DCF, but decreases efficiency by up to 30% with no increase in throughput. Hybrid coordination function (HCF) replaces PCF, and can offer up to 20% increase in throughput. However, it cannot be used in ad-hoc networks, it is less stable than PCF, and does not guarantee fair access to all stations [3]. These coordination functions are necessary for the operation of 802.11, but have a major impact on throughput as they must run separately from data transfers. The coordination functions, combined with headers and trailers for the MAC and physical layers, mean that the actual throughput is well below the raw data rate.

1.2 Realistic Throughput

For 802.11b, the actual throughput is limited to about 5 Mb/s; with 802.11a and 802.11g limited to about 24 Mb/s. Even these rates are seen only if the packet sizes are consistently close to the maximum allowable size (usually set to 1500 or 2048 bytes payload, but can be a maximum of 2300 bytes for payload and header as defined by the standard). Packets of this size are only, in general, seen in file transfers.

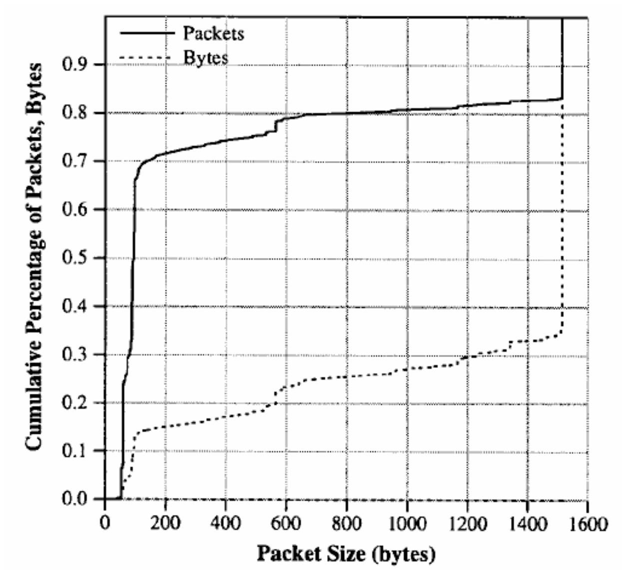

Figure 1: Packet distribution from Tang & Baker [4]

For general Internet and other TCP usage however, the average packet size is far lower than this maximum size – which means that the percentage overhead will be increased accordingly. In a semester-long study of a campus-wide wireless network, Kotz & Essein analysed over 350 million wireless frames [4]. They found that 99.7% of the frames were IP traffic, with 94.3% alone for TCP. In a similar study, over 70 million frames, by Tang & Baker, over 70% of all packets were below 200 bytes (see Figure 1) [5].

As the overheads introduced by the MAC layer are significant – especially in the case of small packets – the MAC layer efficiency needs to be improved. Indeed, simulations carried out as part of the 802.11n development process looked at the efficiency of the 802.11e MAC, combined with the proposed physical layers for 802.11n. These simulations found that for physical layer modes giving over 140 Mbps, the 802.11e MAC layer is at best 64% efficient, and in some cases only 54% efficiency, even with 1500 byte packets [6].

2 Dynamic RTS/CTS Threshold

The first improvement that was proposed in our research was controlling the RTS/CTS handshake used in medium access. For 802.11, a variable named *dot11RTSThreshold* (number of bytes) is used to decide whether a packet is large enough to warrant the use of a RTS/CTS handshake, or whether just the basic access mechanism is used. The basic access method has less overhead – especially for smaller frames – but the collision probability increases with larger frames. The RTS/CTS handshake secures the wireless medium, thus reducing the collision probability, making it more effective for transmission of larger frames that offset the extra overhead. Secondly, the RTS/CTS handshake is inefficient when only a small number of stations are in the current service set due to the reduced probability of collisions, but becomes progressively more efficient with larger numbers of stations.

Currently, *dot11RTSThreshold* is statically set when a wireless network is created. However, if a network becomes smaller, or the load on the network reduces, a higher *dot11RTSThreshold* can be used to reduce overheads and increase performance. Similarly, a growing network in terms of either stations or load will benefit from a lower *dot11RTSThreshold* to reduce the collision probability.

The effect of the value of *dot11RTSThreshold* has been studied previously [7,8,9,10]. All of these studies found that the optimum *dot11RTSThreshold* scales according to the number of stations in the network - as the number of stations increases, the optimal value for *dot11RTSThreshold* decreases.

If the number of the contending stations is relatively small (e.g. 5), the basic access achieves equivalent packet delay compared to the RTS/CTS mechanism, for all packet size values, due to the low collision probability. However, when a network's size increases (e.g. 25 stations) and the packet size exceeds a specific threshold (~250 bytes), the RTS/CTS mechanism achieves a lower packet delay value compared to the basic access. This threshold decreases to about 150 bits for large network scenarios ($n \approx 50$). In this case, the RTS/CTS mechanism is particularly beneficial due to the increase in collision probability [11].

Therefore a dynamic RTS threshold will provide an increase in the throughput of a wireless network by either reducing overheads, or decreasing the number of collisions. A dynamic scheme would not remove anything from the current standard either – meaning that it could be used with existing 802.11 wireless networks.

3 Packet Aggregation

The primary target for this research into improving MAC layer performance is aggregating several packets from the network layer into one frame at the MAC layer. This means that only one physical layer header and one MAC layer frame header are needed, as well as only a single coordination function round, for several packets. This mechanism significantly reduces the overhead per packet. Although either an increased header size, or use of several sub-headers, is needed for the receiver to split the aggregated frame up, the addition is very small. Also, the largest percentage overhead is below the MAC layer, due to the slower speed used for physical layer headers and medium access mechanisms.

This idea has been attempted previously. Due to the cut-throat competition in the 802.11 hardware market, aggregation has been used as a proprietary means to achieve higher speeds from 802.11 in two commercially available wireless chipsets, from KarlNet [12] and Atheros [13]. There is also research into packet aggregation in 802.11 networks by Jain et al [14]. Packet aggregation has also been suggested as a likely component of the 802.11n MAC layer.

The KarlNet 'TurboCell' 802.11 chipset uses a patented 'forced delay aggregation', where frames are placed in queues according to the next routing step. Transmission of the frame occurs either when it has been in a queue for a certain delay period, or a queue is full [14]. This method is not ideal however, as it forces a delay on all packets, even if the station can transmit, the TurboCell chipset will wait to check if more packets are coming for possible aggregation.

The Atheros 'Super-G' chipset uses 'bursting' and 'fast frames' [13]. Bursting allows each station to transmit several frames at once with only a single medium access – this is similar to the operation of HCF in 802.11e. Although the fast frame mode claims to use frame aggregation, in fact it uses an extra negotiation handshake between the station and access point to agree to send larger packets i.e. larger than the maximum specified at the network setup. This would give high throughput for file transfers, but not for other types of IP traffic. Also, the extra features in the TurboCell or SuperG chipsets will not work for ad-hoc wireless networks.

These two chipsets have a major drawback, irrespective of how good the actual technologies are – they are non-standard extensions to 802.11, and as such are not shared by any other manufacturers. This means that a wireless network must be made up entirely of the one set of products offered by the company. If any other product is used – either other manufacturers or an older model – the network will still operate as specified in the 802.11 standards, but will not see the benefits claimed.

3.1 Aggregation types

The research carried out by Jain et al [13] identified two types of packet aggregation – forced delay aggregation as used by the TurboCell chipset, and 'congestion triggered aggregation' developed by themselves earlier. Forced delay aggregation (figure 2a) puts packets into time bounded queues, and then 'pushes' them through to be transmitted. Congestion triggered aggregation (figure 2b) relies on the fact that in a congested network, packets will naturally queue as they wait for transmission. In this mechanism, the transmitter 'pulls' a frame from a queue when it can transmit – this frame may have one packet, or several, in its payload. This does not introduce delays as in the TurboCell implementation, as if only one packet is to be sent, it will still be sent with the same delay as standard 802.11.

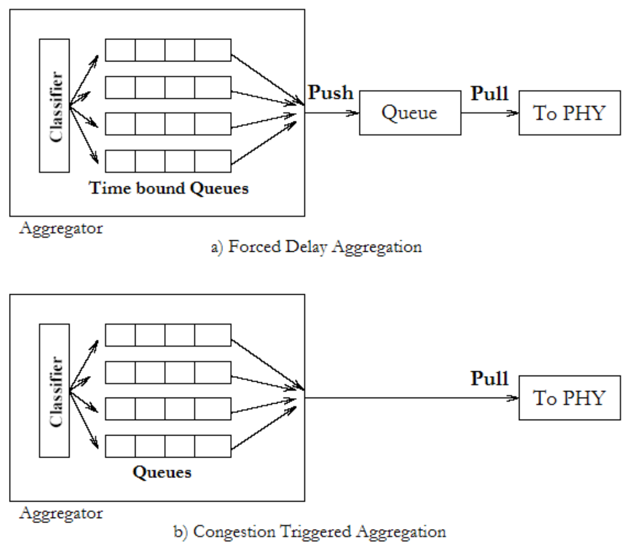

Figure 2: Aggregation types [13]

Jain et al [13] looked at aggregation implemented in their test 802.11b wireless network, and looked at the throughput and median latency. At low throughputs, current 802.11b performed the same as the congestion triggered aggregation, with forced delay aggregation giving similar throughput, but larger median delay than existing 802.11b. As the throughput increased, the aggregation showed much better results in both received throughput and median latency than 802.11b, but forced delay aggregation lagged behind congestion triggered aggregation in terms of median delay.

3.2 Queues in Aggregation

In the aggregation mechanism, each incoming packet from the network layer (either originating from the local computer, or as part of a multi-hop route) is sorted into a queue depending on its destination. At low traffic loads, the choice of queue has little effect on the throughput offered or the packet delay. However, the authors theorized that with increasing load, the decision of which queue to aggregate packets out of would produce significant differences in performance.

Our research identified six main types of queueing control in the aggregation queueing decision for comparison, as outlined in the following list.

- Round-robin
- Priority based – assign a priority to each queue, and choose the highest priority queue that is non-empty
- Weighted fair queueing – combination of priority and round-robin, i.e. every queue gets a turn, but some queues get precedence over others
- Load-based – choose queue with the largest number of bytes in it
- First-in, first-out (FIFO) - choose queue with the earliest timestamp at its head
- Maximise aggregation – choose queue giving the highest possible gain from aggregation, i.e. the queue with the largest number of smaller packets

From the paper by Jain et al [14], both the forced delay aggregation and the congestion triggered aggregation use the first in, first served queueing control. These six queueing controls were simulated, using congestion triggered aggregation, along with a model of existing 802.11 MAC. All queueing controls had the same amount of queueing memory, with the same equivalent amount used for the existing 802.11 simulation.

4 Simulations

The simulations assumed that the 802.11a physical layer was being used, with the highest data rate (54 Mb/s) used. As yet, no sub-frame format has been proposed. However, a space of 6 bytes was allocated for the purpose of simulating packet sub-headers within the frame. The coordination function used to control the stations was DCF.

The simulations were implemented as a MATLAB script, with the packet sizes and arrival times at each station calculated beforehand. Traffic was generated randomly using a packet distribution function similar to that measured by Tang & Baker (see Figure 1) [5]. The decision to use MATLAB was made as the researchers could not find any existing simulation tools that met the requirements – most importantly the handling of variable packet sizes that aggregation created.

Simultaneous transmissions by 1, 2, 5, 10 and 20 stations were simulated, using set packet rates at each station's MAC layer service access point (giving throughputs ranging between about 1.5 Mb/s (500 packets/s) and 16 Mb/s (5000 packets/s)). From the simulation, the main statistics generated were the average packet delay and the overall throughput.

4.1 Simulation Results

Figure 3 shows the results from the simulation with 5 stations transmitting. At lower packet rates (about 1.5 and 3 Mb/s per station), there is virtually no packet delay, and all queueing controls give very similar throughputs of about 8 Mb/s and 16 Mb/s. This level of traffic is well within the capability of existing 802.11a. However, the next step, at about 6 Mb/s per station (2000 packets/s), shows a significant difference between the simulations with aggregation enabled, and the existing 802.11a simulation.

With this level of traffic, the existing 802.11 MAC cannot cope. The average packet delay is high, and the overall throughput is visibly limited. The effect of the aggregation can be seen clearly, with the overall throughput and packet delay significantly better than the existing MAC for all queueing mechanisms. With this rate of data packets (and higher rates), a difference in queueing mechanism can be seen.

Consistently the best queueing control mechanism in terms of throughput is the load based control, with the FIFO control close behind. In terms of average packet delay, the 'maximise aggregation' control performs best as it sends more packets sooner. FIFO control also achieves low packet delay due to its time based nature. Interestingly, the load-based control is the second-worst performing control in terms of average packet delay. The 'priority based' control is clearly the worst performing controls in terms of throughput and average packet delay, due to the unfairness towards low priority stations.

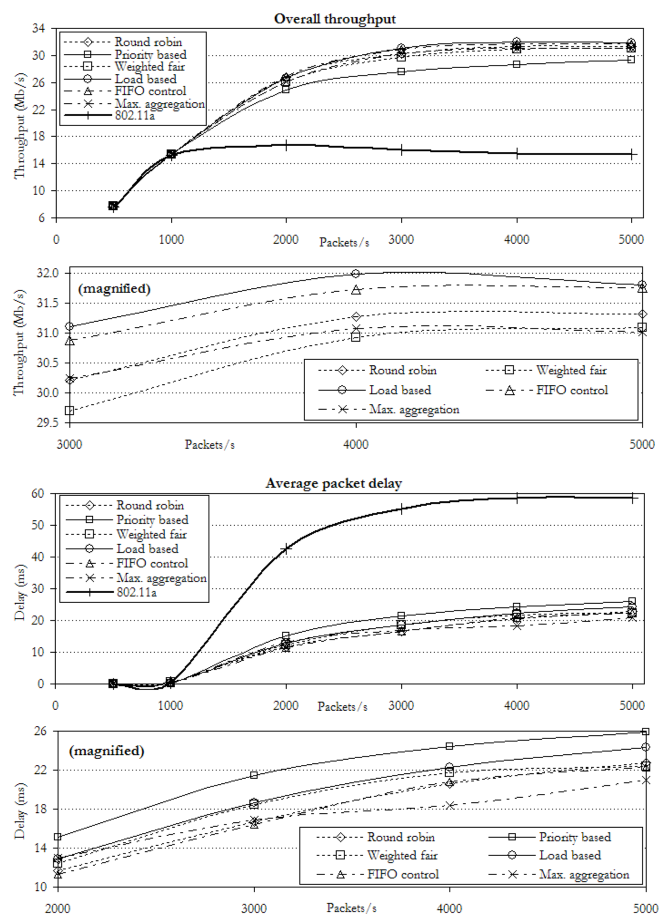

Figure 3: Simulation results for 5 stations transmitting: overall throughput and average packet delay

The simulations show that using aggregation with load-based control attains an overall throughput reaches approximately 32 Mb/s. This is much higher than is currently achievable with 802.11a/g, and this performance is not gained with continuously large packets, but with a wide range of packets, including a large percentage smaller than 200 bytes.

Figure 4 shows the results of simulations carried out with the same traffic rate per station (about 3 Mb/s), but with the increasing numbers of stations transmitting. For a single station, 2 stations and 5 stations, the traffic is within the capabilities of existing 802.11 MAC layer – with data rates of approximately 3, 6 and 15 Mb/s respectively. However, as the number of stations transmitting at once increases, the existing 802.11 MAC cannot cope - the overall throughput levels off at around 17 Mb/s, and actually begins to fall in the 20 station simulation. The average packet delay also increases significantly under the higher loads.

The simulation of the aggregation mechanism under the same load performs much better, with about 25 Mb/s throughput overall for 10 stations, and close to 30 Mb/s for 20 stations. The average packet delay increases, but not to the extent of that seen in the existing 802.11 simulation. Once again, the load-based and FIFO controls have the highest throughput, however the maximise aggregation and round-robin mechanisms are not far behind. Also, the priority based scheme performs the worst, with the weighted fair queueing (with is partly a priority based mechanism) also not performing as well. The results are similarly repeated in the average packet delay, with the lowest delay given by the maximise aggregation followed by the FIFO control.

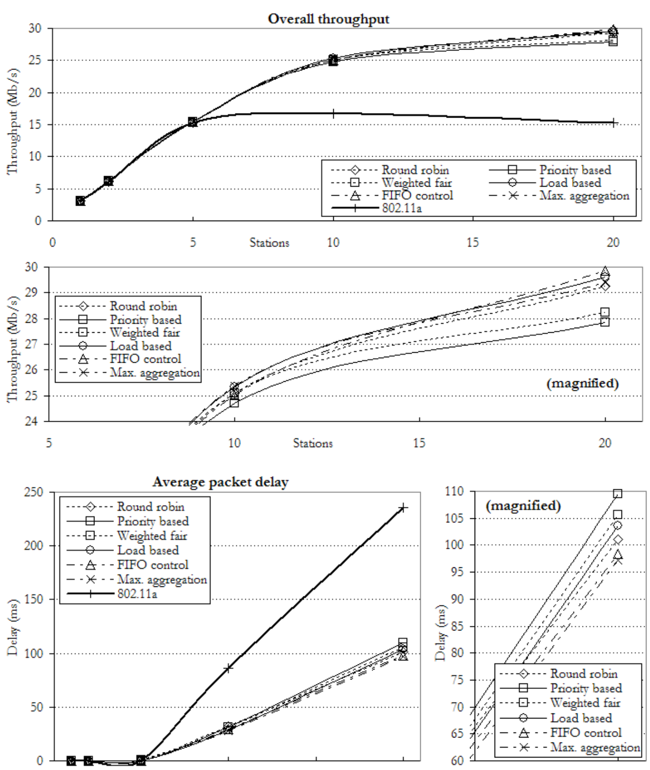

Figure 4: Simulation results for different network sizes (1000 packet/s at each MAC service access point) - average packet delay and overall throughput

From all the simulations carried out, the overall throughput appears to be limited at about 32 Mb/s. This is higher than the 24 Mb/s limit for existing 802.11a/g, and is achieved with a more reasonable packet size distribution compared with the large packet sizes needed for maximum speed under the existing MAC. This is mainly because the aggregation combines the smaller frames into larger ones, thus lowering the overhead percentage. The aggregation allows the full speed of the physical layer to be used more often – rather than simply using a higher physical layer speed to get a higher throughput.

4.2 Further Improvements

From the simulations we have conducted, the best queueing control is the load based mechanism. Currently, the load is calculated solely by the byte-size of the queue. Another method would be to look at the packet loading of each queue. As well as this, our research has identified a dynamic extension to the queueing control. Due to the often random nature of queueing, some queues can be full while other queues are comparatively empty.

The proposed scheme looks for queues that are almost full. These queues are defined to be those that cannot fit the maximum sized packet if it were to be sent down by the network layer. The dynamic scheme would give this queue temporary priority to transmit frames, thus preventing the packet from timing out because of the queueing mechanism. The default TCP timeout for a packet is 3 seconds. This means that if a packet is dropped by the queueing mechanism, it would be delayed by at least 3 seconds – which is relatively large compared with the millisecond delays experienced in normal operation.

Basic aggregation can also be improved with a backoff scheme in case of transmission failure. In all wireless networks, the longer a frame is, the more prone it is to errors, which often requires the frame to be retransmitted. The backoff proposed would lower the upper limit for the frame payload size due to aggregation. For instance, all simulation carried out used a maximum payload size of 1500 bytes. In the case of a transmission failure, the maximum size would be reduced by half to 750 bytes. This reduction in maximum frame payload size would continue until successful transmission. The maximum value would reset to 1500 bytes upon success. Some history of transmission failures could also be kept – allowing the maximum payload size to be kept at a lower amount to ensure a lower rate of failure. This could be used to counter the time-variant effects that often plague wireless networks.

5 Conclusion

From the simulations carried out, packet aggregation can be used to increase the throughput of 802.11 wireless networks. Under heavy traffic, MAC layer aggregation can achieve a much higher throughput than the existing 802.11 MAC. This increase in throughput is due to the increased payload size under aggregation compared with existing 802.11 – thus reducing the overhead percentage.

Combined with other MAC layer enhancements, packet aggregation could be used in conjunction with the next generation of 802.11 physical layer – most likely to be a form of MIMO (multiple-input, multiple output – where multiple antennae are used with advanced coding schemes to get multiple channels within a single frequency band). This combination will allow a throughput of at least 100 Mb/s at the MAC service access point, with raw data rates of at least 160 Mb/s.

6 Further Research

The first task to complete is formulating a sub-header system for aggregation. The simulations carried out allowed for 6 bytes per extra packet aggregated – therefore the initial target is to fit the packet sub-headers within this limit. Thus far, it is envisioned that only a sequence number, a length field and a packet check sequence is needed, as all of the 802.11 specific details of all packets aggregated into a single frame will be identical and therefore contained in the standard header. The maximum payload size backoff scheme for packet aggregation also needs to be developed to a testable state for simulation.

Beyond this, other MAC layer improvements will be investigated. All MAC layer improvements identified will also be considered from a standardisation point of view. This process would entail preparing formal descriptions of the improvements. In conjunction with this, hardware implementation will be investigated, and if feasible, hardware based implementations will be simulated.

7 References

1 J. Rosdahl. "Project Authorisation Request for High Throughput Study Group". IEEE Working Group Paper 802.11-02/0798r7, March 2003.

2 Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. IEEE Std. 802.11, 1999.

3 P. Garg, R. Doshi, R. Greene, M. Baker, M. Malek and X. Cheng. "Using IEEE 802.11e MAC for QoS over Wireless". Proceedings of the IEEE Performance, Computing and Communications Conference, Phoenix, Arizona, USA, pp. 537-542, 2003.

4 D. Kotz and K. Essein. "Analysis of a Campus-wide Wireless Network". Proceedings of the Eighth Annual International Conference on Mobile Computing and Networking, Atlanta, Georgia, USA, pp. 107-118, September 2002.

5 M. Baker and D. Tang. "Analysis of a Local Area Wireless Network". Proceedings of the 6th Annual International Conference on Mobile Computing and Networking. ACM Press, pp. 1–10, 2000.

6 A. R. Dias, S. Rouquette-Leveil and S. Simeons. "Multiple Antenna OFDM Solutions for Enhanced PHY". IEEE Working Group Paper 802.11-04/0229r0, March 2004.

7 W. Y. Choi and S. K. Lee. "Maximizing MAC Throughputs by Dynamic RTS-CTS Threshold". IEEE Working Group Paper 802.11-04/0312r0, March 2004.

8 S. T. Sheu, T. Chen, J. Chen, and F. Ye. "The Impact of RTS Threshold on IEEE 802.11 MAC Protocol". Proceedings of the Ninth International Conference on Parallel and Distributed Systems, Taiwan, China, December 2002.

9 S. Yan, Y. Zhou and S. Wu. "An Adaptive RTS Threshold Adjust Algorithm based on Minimum Energy Consumption in IEEE 802.11 DCF". Proceedings of International Conference on Communication Technology, Beijing, China, pp. 1210-1214, April 2003.

10 M. Natkaniec and A. R. Pach. "An Analysis of the Influence of the Threshold Parameter on the IEEE 802.11 Network Performance". Proceedings of the IEEE Wireless Communications and Networking Conference, Chicago, Illinois, USA, pp. 819-823, September 2000.

11 P. Chatzimisios, A.C. Boucouvalas and V. Vitsas. "Packet Delay Analysis of IEEE 802.11 MAC Protocol". *Electronics Letters*, Vol. 39 No. 18, 4th September 2003.

12 D. Karl, S. Monahan and A. Whitman. "Integrity TurboCell – Enhancing the Capabilities of Standard 802.11". KarlNet Inc. Available from http://www.karlnet.com/Documents/WhitePaper/TurboCellWhitePaper/TurboCell%20White%20Paper.htm. 16 July 2004.

13 Atheros Communications, Inc. "White Paper: Super-G Maximizing Wireless Performance". Atheros Communications, Inc. Available from http://www.atheros.com/pt/atheros_superg_whitepaper.pdf. 16 July 2004.

14 A. Jain, M. Gruteser, M. Neufeld and D. Grunwald. "Benefits of Packet Aggregation in Ad-Hoc Wireless Network". Technical Report CU-CS-960-03, Department of Computer Science, University of Colorado, Boulder, Colorado, USA. August 2003.