E-CLIC European Collaborative Innovation Centres for Broadband Media Service	**Place for logotype**

An Experimental Comparison of Burst Packet Transmission Schemes in IEEE 802.11-based Wireless Mesh Networks

Case Study Final Report

EUROPEAN UNION
European Regional
Development Fund

Investing in the future by working together for a sustainable and competitive region

The Interreg IVB
North Sea Region
Programme

An Experimental Comparison of Burst Packet Transmission Schemes in IEEE 802.11-based Wireless Mesh Networks

Peter Dely, Andreas J. Kassler[1], Nico Bayer, Dimitry Sivchenko[2]

Abstract
Wireless mesh networks (WMNs) are wireless multi-hop networks comprised of mesh routers, which relay traffic on behalf of clients and other nodes. Using the standard IEEE 802.11 distributed coordination function (DCF) as MAC layer, a node needs to contend for the medium each time it wants to transmit a packet. This creates high overhead in particular for small packets and leads to poor performance for real-time applications such as Voice over IP (VoIP) or online gaming. Burst packet transmission can increase the efficiency. For example, using the Transmission Opportunity limit (TXOPlimit) in IEEE 802.11e, a station may transfer several packets without contending for the channel in between. Similarly, IP packet aggregation combines several IP packets together and sends them as one MAC Service Data Unit. Originally, both schemes have been developed for single-hop networks only. Thus the impact on WMNs is unclear if the packets need to contend over multiple hops. In this paper, we use measurements from a 9-node WMN testbed to compare TXOPs and IP packet aggregation for VoIP in terms of fairness, network capacity and quality of user experience. We show that for low networks loads, both TXOPs and IP packet aggregation increase the VoIP quality compared to IEEE 802.11 DCF. However, in highly loaded networks IP packet aggregation outperforms the other schemes.

Index Terms
Performance Evaluation, Testbed, Burst Transmission

1 Introduction

Wireless mesh networks (WMNs) are a promising technology for providing cost-efficient wireless Internet access to e.g. rural or urban areas. In a WMN, mesh routers relay traffic on behalf of clients or other mesh routers and thus form a wireless multi-hop network. Most WMNs are based on IEEE 802.11 versions and 802.11s is standardising mesh mode operation. With the ever increasing transmission speed of 802.11 based technology due to the introduction of MIMO, OFDM, etc, transmission time for user payload decreases rapidly. However, the performance of WMNs based on 802.11 can still be low as the time spent on overhead (such as backoff, MAC and PHY layer headers) dominates for small packets. As an example, the transmission time of a 100 byte packet sent at 54 Mbit/s consists to 95% of overhead created by the IEEE 802.11 MAC layer. This problem becomes more severe as the data rate increases because most of the headers are transmitted at a lower data rate. Also, this effect is more prominent for short frames, such as those typically used for VoIP. Furthermore, John et al. [1] report that 50% of the packets on the Internet are smaller than 700 bytes. As VoIP is an important service to be considered for mesh network operators, it is important to transmit small packets in an efficient way.
One possibility to increase transmission efficiency is to aggregate multiple smaller frames together into a larger one for transmission in one burst. This approach has multiple benefits because it reduces the PHY and MAC header overhead. Also, it reduces the total number of transmissions, which reduces contention and collision probabilities. This is especially important in a multi-hop setting under presence of hidden nodes as collisions lead to low throughput.
Burst transmission schemes have been investigated at different layers. MAC layer frame aggregation is a key mechanism to achieve higher throughput in IEEE 802.11n. Using the concept of transmission opportunity limit (TXOPlimit) in IEEE 802.11e allows a sender to send multiple packets once it has gained access to the medium. While those schemes work very efficient, a drawback is that they require a specific MAC layer. On the other hand, there have been several packet aggregation schemes developed that work with any 802.11 based MAC layer. In such approaches, a layer 2.5 module aggregates

[1] Peter Dely and Andreas J. Kassler are with Karlstad University, Universitetsgatan 2, 651 88 Karlstad, Sweden (e-mail: peter.dely@kau.se and kassler@ieee.org).
[2] Nico Bayer and Dimitry Sivchenko are with Deutsche Telekom Laboratories, Ernst-Reuter-Platz 7, 10587 Berlin, Germany (e-mail: nico.bayer@telekom.de and dmitry.sivchenko@telekom.de).

multiple IP packets together and forms an aggregated IP packet which is then transmitted at the MAC layer.
To have enough packets to aggregate, it is possible to delay packets to achieve a higher aggregation ratio. Such delay might seem counterproductive but can increase aggregation ratio, especially in low traffic scenarios. Collision probability will thus be further reduced leading to overall lower end-to-end delay as the average backoff delay will decrease. While in a single hop infrastructure network, the access point has all information to derive an optimum packet size to be used for aggregation, this is complicated in a multi-hop scenario where other factors play an important role because packets are relayed along multiple hops. Hence, it is not clear, which burst transmission scheme works best in a multi-hop setting.
In this report, we evaluate in detail the performance of IP packet aggregation in a real wireless mesh testbed under several scenarios varying network load and traffic direction. In addition, we evaluate user achievable quality of experience by analyzing MOS scores of aggregated VoIP calls. We compare the performance with other schemes, including the use of TXOPs. We also study the impact of RTS-CTS usage in a multi-hop setting. The performance evaluation is based on a real operational wireless mesh network, looking at the end-user perceived quality. Our report is organized as follows. Section 2 outlines existing burst transmission schemes. Section 3 details the experimental setup and analyzes and compares the performance of the different schemes. We present our conclusions and draft future work in Section 4.

2 Background

2.1 IEEE 802.11e RTS-CTS and transmission opportunities

There are two main operations defined for 802.11 distributed coordination function (DCF) MAC (see Figure 1). In the simple mode, a station waits until the medium is idle for DIFS and transmits after backoff. If the transmission is successful, the receiver sends back an ACK. If the packet (or the ACK) has not been received correctly, the sender will resend after timeout following a backoff procedure. In the other mode, each transmission starts with an exchange of request-to-send (RTS) and clear-to-send (CTS) handshake to virtually reserve the medium. Such control packets can improve the performance in the single-hop case as the data packets are usually larger compared to RTS-CTS control packets and more effected by collisions. However, RTS-CTS packets are transmitted without any further protection and still may collide. Especially in the multi-hop case under the presence of hidden terminals, frequent collisions among RTS or CTS packets may severely degrade the performance. Also, the usage of RTS-CTS packets should be avoided for small packets due to the increased overhead.

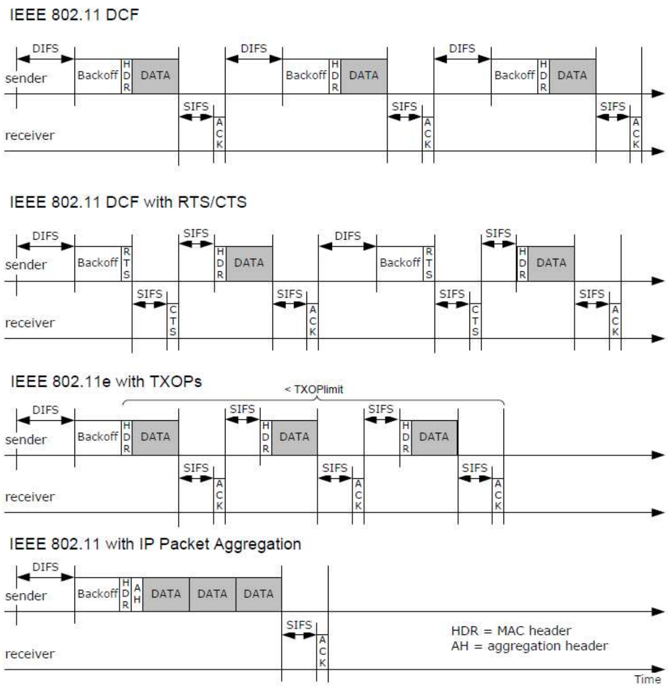

Figure 1: Burst Packet Transmission Schemes for 802.11

Increasing the performance and Quality of Service provisioning can be achieved using the concept of transmission opportunity (TXOP) in 802.11e [2]. Once a sender has successfully contended for the medium, it can send several frames separated by SIFS without contending for the medium in-between (see Figure 1). The TXOP is defined by its starting time and duration during which a station may transfer data of a particular traffic class. TXOPs can be either obtained via contention-based medium access (EDCA-TXOPs) or via controlled medium access (HCCA-TXOP or polled TXOP). The maximum duration of an EDCA-TXOP is limited by the parameter TXOPlimit, which is distributed periodically through beacon messages. The TXOPlimit allows controlling the maximum time a station can allocate the medium for the delivery of MAC Service Data Units. As different service classes can define different TXOPlimits, this mechanism enables an effective control of the delay. IEEE 802.11e allows the use of block-ACKs which enables the receiver to acknowledge the successful reception of multiple frames using a single ACK packet.

2.2 IEEE 802.11 A-MSDU/A-MPDU

IEEE 802.11n [3] introduces MAC frame aggregation, where the sender either aggregates MAC Protocol Data Units (A-MPDU) or MAC Service Data Units (A-MSDU). In the A-MSDU mode, the MAC layer aggregates multiple packets from the upper layer by adding a single MAC header and check-sum. In contrast, the A-MPDU mode concatenates multiple 802.11 MAC frames each having its own MAC header and check-sum. By introducing a MAC delimiter, a receiver is able to separate each subframe, even if some of the sub-frames are corrupted. It also supports a block ACK scheme which allows the sender to retransmit only erroneous subframes. This can improve performance for channels having high bit error rates. The standard does not specify when packets should be aggregated but normally this is done when there is more than one frame available in the sender queue. Hence, under low load, most packets will be sent unaggregated. Skordoulis et al. [3] show that frame aggregation in 802.11n can lead to performance improvements in single hop cases, if both modes are combined effectively. Kim [4] evaluates the performance of an early version of 802.11n frame aggregation as a function of payload size and data rate also in the single hop case.

2.3 IP Packet Aggregation

With IP packet aggregation packets destined for the same next-hop are concatenated before passing them to the MAC layer. An extra IP-header is added which enables the next hop to de-aggregate the packet (see Figure 1). This mechanism is transparent to the MAC layer and thus no partial MAC-layer retransmission of erroneous segments is possible. While in theory limiting the maximum burst length to a value smaller than the MTU on weak links can reduce the packet error rate, this is not very applicable in practice. The transition region from a good link (that allows to fill the whole MTU) to a bad link (which requires bursts smaller than the MTU) is about 1-2 dB in SNR [5]. Due to the coarse quantization of SNR measurements on current cards and due to small scale fading, tuning the maximum burst size not very effective. Instead, the rate adaptation scheme should select a PHY rate that supports large frames.

Since IP packet aggregation is decoupled from the MAC layer, it cannot utilize the inherent delay for access the medium in the MAC. Artificially delaying packets by the right amount of time is thus crucial. Kyungtae and Ganguly[6] propose to let ingress mesh router probe the path to the destination to determine the end-to-end latency. The aggregation delay is set so that end-to-end latency plus the buffer delay does not exceed a pre-configured threshold. Intermediate nodes are not allowed to artificially delay packets further, but can aggregate additional packets whenever available. Riggio et al. [7] use a combination of probe messages, channel monitoring and an analytical model to derive an optimum packet size for a given network condition. Packets are delayed to create packet bursts of the optimum packet size.

3 Performance Evaluation

3.1 Experimental Setup

We compared the performance of multiple burst transmission schemes in the KAUMesh testbed, which consists of 20 Cambria GW2358-4 based mesh routers deployed in the ceiling of the engineering building of Karlstad University. The nodes are equipped with Atheros 5212-based IEEE 802.11a/b/g wireless cards. A wired Ethernet card was used for time synchronization and to transfer traffic log-files. Mesh nodes run Linux 2.6.22 and MadWIFI 0.9.4, which we modified to support IEEE 802.11e in ad-hoc mode. In order to avoid CPU bottlenecks and unwanted effects of the rate adaptation scheme, we disabled Auto-Rate and fixed the PHY data rate to 6 Mbit/s. The cards are operated in the 5 GHz frequency range to avoid interference from the campus WLAN. A subset of 9 nodes was used for our evaluation topology (see Figure 2).

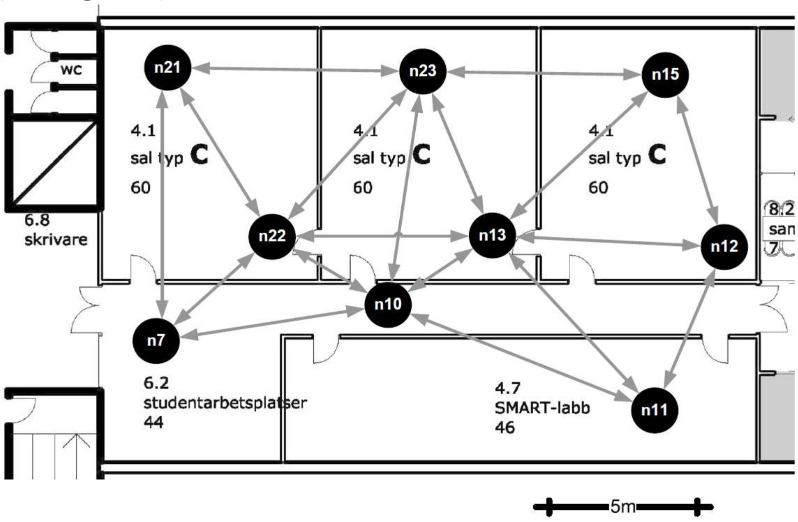

Figure 2: Evaluation topology within KAUMesh

We have implemented IP packet aggregation as a module for the Linux traffic control subsystem. The module contains a virtual FIFO-queue for each neighbor. When an IP-packet is sent from the user space or forwarded, it is marked with an expiration timestamp and enqueued in a virtual queue. After a packet is enqueued, the network card requests packets from the operating system or a timer expires, the aggregation module selects a virtual queue and concatenates all packets up to a size of MTU. Virtual queues are only dequeued, if packets have surpassed their expiration times (“aggregation delay”) or enough packets are available to fill up the whole MTU. An extra IP header indicating an aggregated packet is prepended and the aggregation packet is sent. A new timer is set to trigger a dequeue of the virtual queue when the next packet expires. The aggregation delay is configured statically. On the receiving node, aggregated incoming packets are identified by the extra IP header, de-aggregated in a netfilter-module and inserted into the normal Linux IP-stack.

For each IEEE 802.11e access category, the network card (in our case based on the Atheros 5212 chipset) has a hardware FIFO queue, in which MAC frames are stored before transmission. As soon as the station has successfully contended for the medium, it can transmit frames of one access category that are available in the MAC frame buffer for a maximum time of TXOPlimit. To the best of our knowledge, support for block-ACKs cannot be configured with the Atheros 5212 chipset. As our WLAN NICs do not support the IEEE 802.11n standard, we could not evaluate the A-MSDU scheme.

3.2 Single-Hop Performance

We compare the IEEE 802.11 DCF, IEEE 802.11e with TXOPlimits of 1, 2, 3 and 8ms and IP packet aggregation with aggregation delays of 1, 2, 3 and 8 ms in a single-hop scenario. We transmitted parallel UDP flows (200 bytes payload) at rates of 300 up to 650 packets/s (in steps of 25 packets) from nodes 10, 13, 23, 22 and 7 to node 21, using [8]. Each test was executed for 60 seconds and repeated 5 times.

3.2.1 Maximum Achievable Rate

Figures 3a) and 3 b) show average end-to-end packet loss and delay, error bars show standard deviation of individual test runs, indicating the well known behavior [9] of the IEEE 802.11 DCF. For lightly loaded networks (e.g. load < 3.0 Mbit/s), packet loss ratio and delay are low. In the transition from a non-saturated to a saturated network, packet loss ratio and delay rise quickly. While the network throughput in saturated IEEE 802.11 networks is at its or close to its peak, delay and packet loss ratio is usually unacceptable for VoIP, which mandates delay below 150 ms at a loss below 3% [6]. Thus, the optimum operation point for a network is just before the saturation. For the standard IEEE 802.11 MAC layer this is at about 3 Mbit/s, while using TXOPs increases it to about 3.4 to 3.7 Mbit/s (a higher

TXOPlimit permits a higher traffic injection rate). With IP packet aggregation the operation point is between 4.1 to 4.8 Mbit/s, whereas higher aggregation delays allow higher traffic injection rates. However, for low loads of 3 Mbit/s the end-to-end delay increases by using a higher aggregation delay. The better performance of IP aggregation under high load is due to the more effective overhead reduction.

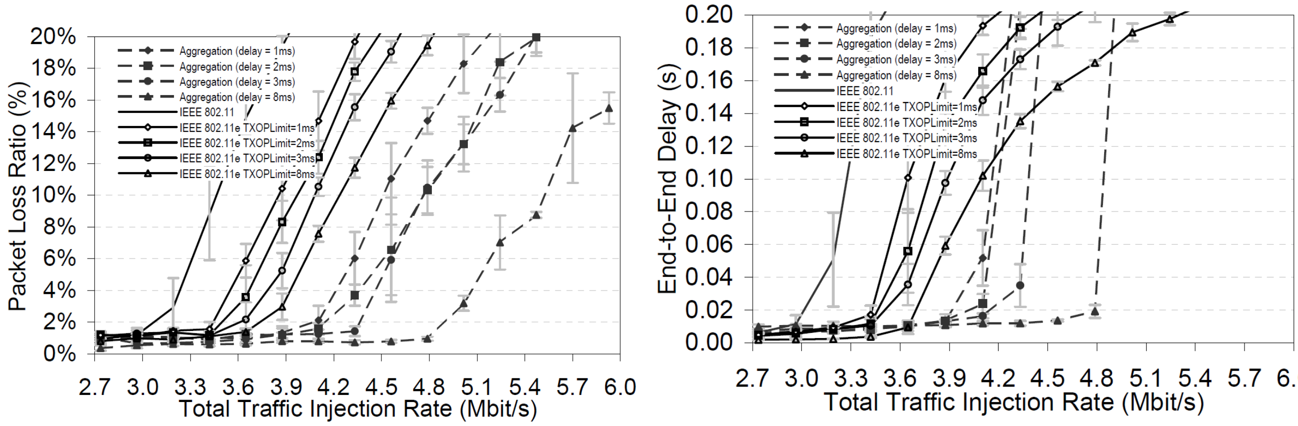

Figure 3: Packet loss (a) and delay (b) vs. offered load for IEEE 802.11 DCF and IP packet aggregation (single hop)

3.2.2 Average Burst Length

Using TXOPs, the number of channel access attempts is reduced, less time is spent in backoff-phases and fewer collisions occur. IP packet aggregation in addition reduces the number of ACKs and inter-frame waiting times. The efficiency depends on the number of packets sent within one burst. For IEEE 802.11e, we measured the average length of a burst by capturing all traffic with a wireless NIC in RF-monitor mode. This allows to determine when a packet arrived at the network card using the MAC-timestamp field in the Radiotap header. Measurements show that transferring one packet requires around 440µs (including all headers and the ACK, but not the channel access). If difference in arrival times of two subsequent packets is smaller than 460µs (440µs transfer + 20µs error margin), we conclude that both were sent within the same TXOP. For the IP packet aggregation we obtained the average burst size directly from the aggregation module statistics. Using the method described in [10] we also measured channel busy fraction (fraction of time the channel is sensed busy due to transmission or collisions). This is a good indicator for network congestion but also can be used to measure transmission efficiency.

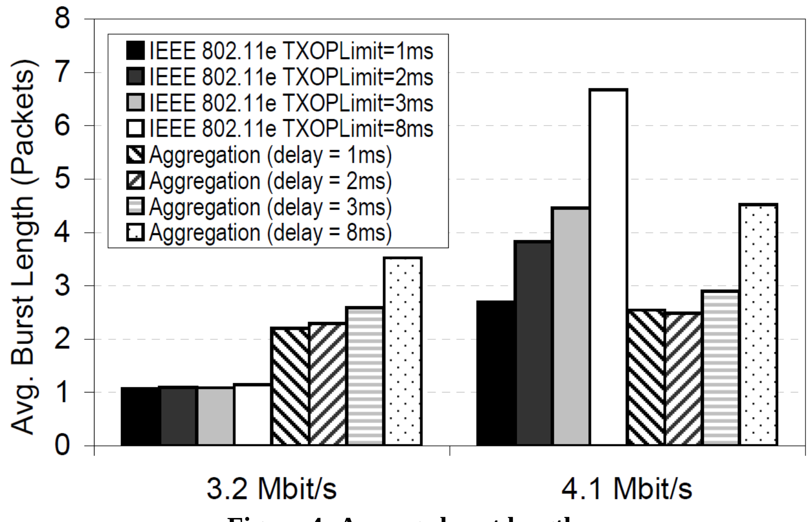

Figure 4: Average burst length

From Figure 4we observe that for low offered load (3.1 Mbit/s) IEEE 802.11e with TXOPs hardly sends more than one packet within one burst. At 3.2 Mbit/s data injection rate the channel busy fraction is 74% (TXOP case), which indicates that collisions are rare and MAC layer queues are short. Therefore, almost no packets are available in the queue for sending within one TXOP. In contrast, IP packet aggregation artificially delays packets and thereby can on average send more than two packets at once. The measured channel busy fraction here is only 60% (67%) for an aggregation delay of 8 ms (1 ms), which shows the higher efficiency of the IP packet aggregation leading also to lower MAC layer utilisation. Higher aggregation delay leads to larger burst length as more packets are available to be aggregated.

For higher rates (4.1 Mbit/s) the channel busy fraction increased to 89% (TXOP case) and we observed a considerable amount of collisions. As a consequence, the queue builds up in the MAC hardware buffers and therefore several packets are available to be sent within one TXOP. Interestingly, when setting TXOPlimit to 1 ms, there are packet bursts with 3 packets, although transmitting three packets (transmission time 1.3 ms) would exceed the TXOPlimit. Apparently the hardware does not comply with the IEEE 802.11e standard here. In contrast, using aggregation, the channel busy fraction was reduced to 78 % for an aggregation delay of 8 ms.

Queuing in the MAC layer only has a minor effect on the burst size of IP packet aggregation. The aggregation module cannot utilize packets waiting in the MAC layer queue to create longer bursts. Only if the MAC layer queue is full (maximum length 50 by default), packets queue up in the aggregation module. Otherwise only artificially delayed packets are available for aggregation. Thus the average burst size is lower for IP packet aggregation. Due to the reduction in inter-frame waiting times and ACKs, aggregation is still more efficient than TXOPs. Reducing the length of the MAC layer queue could increase the aggregation burst length. However, a short MAC layer queue requires very fast packet processing in the higher layers, since the MAC layer needs to have packets available as soon as the medium becomes idle.

3.3 Multi-Hop Performance

Next, we compare IEEE 802.11, IEEE 802.11e with TXOPlimit of 8ms (best performance under high load in single-hop) and IP packet aggregation in a multi-hop scenario with VoIP traffic. Furthermore we investigate the impact of RTS/CTS on TXOPs and IP packet aggregation.

3.3.1 Traffic Generation and Quality Evaluation

We assume that one VoIP call consists of two G.711 audio flows: one from the mesh gateway to the mesh router and one in the reverse direction. We emulate a flow using a stream of 200 byte UDP datagrams with a constant arrival rate of 50 packets per second. We evaluated 50 different scenarios with 8 concurrent calls and 50 different scenarios with 12 concurrent calls. We created one scenario by randomly selecting one node as gateway. Among the remaining mesh nodes, we then randomly chose sources/destinations (one node can be source/destination for several calls), which communicate with the gateway node. Since our evaluation mainly focuses on MAC layer issues, we used static routes.

For each scenario, end-to-end delay, packet loss ratio and jitter of each flow was measured for 60 seconds. We emulated a fixed play-out buffer, which drops packets if their jitter according to [11] is greater than 30 ms. We estimated the perceived user experience of the VoIP call by (see also [12]) first calculating the R-factor (considering the impairment by packet loss including drops by the play-out buffer and the delay) and then converting it to the Mean Opinion Score (MOS), as described by ITU-T E-model using eq. b-4 of [13]. The MOS describes the average user satisfaction, where 5 is “Excellent”, 4 is “Good”, 3 is “Fair”, 2 “Poor” and 1 is “Bad”.

3.3.2 Average Quality

In Figures 5a) and 5b) depict the cumulative distribution of the MOS for VoIP flows with 8 (left) and 12 (right) concurrent calls (combined over all scenarios and flows). For low loaded networks (8 calls), the standard IEEE 802.11 MAC layer provides good quality (MOS\ge 4) to 92% of the flows. The remaining 8% is constituted mainly by flows that need to be relayed over 3 or 4 hops. Using TXOPs or IP packet aggregation reduces the overall network load and gives good quality to all flows. For highly loaded networks (12 calls), none of the compared modes provides good quality to all flows. The best choice here is IP packet aggregation, where approx. 73% of the flows have a MOS greater than 4. Using TXOPs, the number of flows having MOS smaller than 4 increased to 46%. Interestingly, enabling RTS/CTS never improved performance.

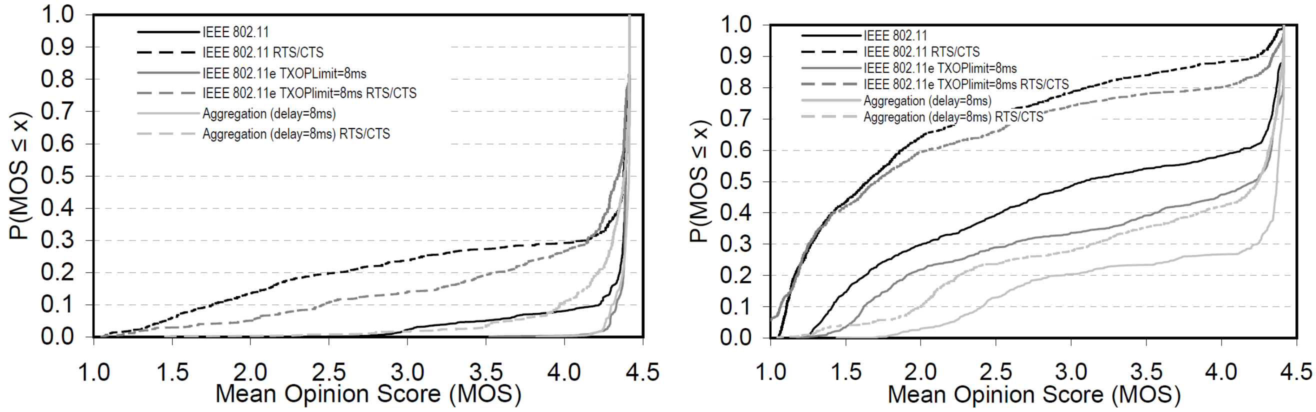

Figure 5: CDF for MOS for 8 calls (a) vs. 12 calls (b) multi-hop scenario comparing TXOPs, RTS/CTS, IP packet aggregation.

A larger hop-count creates more opportunities for packet loss and increases end-to-end delay, which negatively impacts VoIP quality. Figure 6 shows this relation where we group the results according to the hop-counts that flows traversed. The large deviations for flows of a given hop count and be explained as follows: in some scenarios there are only a few flows with high hop counts (2 or 3 hops). In this case the resulting network load is moderate and thus even the flows traversing more hops have good quality. In other scenarios however, many flows have high hop counts. The resulting network load is high and the quality is poor. Using IEEE 802.11e with a TXOPlimit of 8 ms the average burst length was 1.52 packets. In contrast, with IP packet aggregation (delay=8 ms) the average burst length was 1.96. As IP packet aggregation delays packets artificially, it can send more packets at once. In particular in multi-hop networks, where traffic is forwarded and thus not always back-logged, artificially delaying packets may lead to better performance, especially in higher loaded networks.

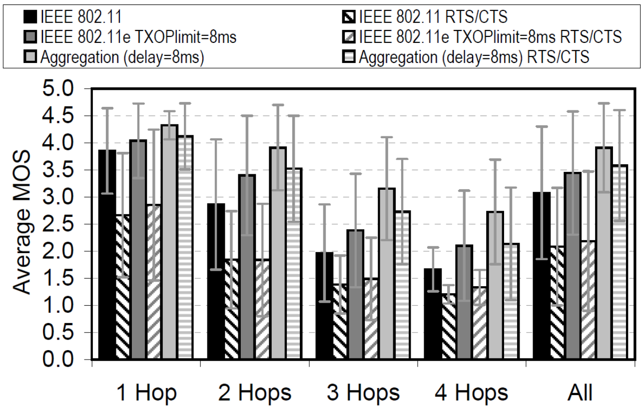

Figure 6: Avg. MOS for 12 concurrent calls split by hop count

The general trend is clear: The average MOS decreases when number of hops increases. RTS/CTS does not improve the performance in average. Due to the higher overhead of RTS/CTS the medium is saturated earlier and hence the VoIP is degraded. The potential benefit of using RTS/CTS by reserving the medium so that less colissions occur is counteracted by the lower efficiency and other problems created by RTS/CTS (such as increased number of exposed nodes when RTS/CTS collide). Aggregation outperforms the use of TXOPs, which again is superior to the standard mode. Therefore, the key to improve performance for VOIP under high load is to enable an efficient transmission mode (e.g. by using packet aggregation).

3.3.3 Fairness

To study the impact of RTS/CTS on fairness, we calculated Jain's fairness index [14] for each scenario, where a value of 1 implies perfect fairness (all flows have same MOS). Figure 7 displays the average over all scenarios. If the network load is low (8 concurrent calls) and thus no congestion occurs also flows traversing three or four hops receive good quality, which leads to a high fairness index. Compared to the standard IEEE 802.11e, the use of TXOPs or IP packet aggregation reduces overhead and consequently network load. Therefore fairness is increased by those burst transmission schemes. Even with 12 concurrent calls the overhead reduction due to burst packet transmission improves fairness.

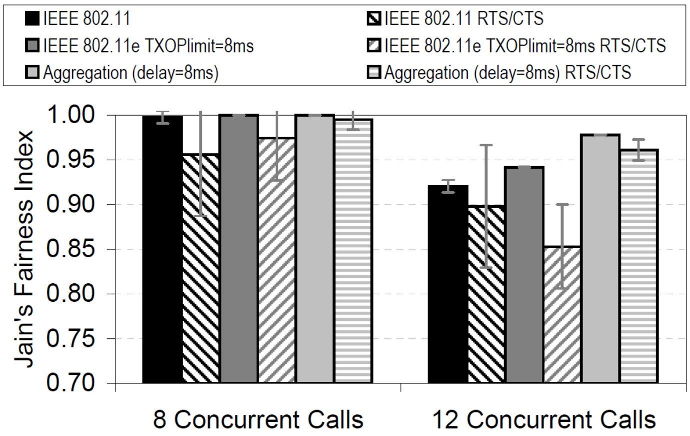

Figure 7: Fairness Comparison

4 Conclusions

We have evaluated the VoIP performance of IP packet aggregation and IEEE 802.11e TXOPs a wireless mesh network. Compared to the standard IEEE 802.11 DCF, both schemes significantly increase the average VoIP quality and fairness. RTS/CTS does not contribute to a better average quality or more fairness. In multi-hop scenarios artificially delaying packets can be beneficial, since it creates longer bursts and enhances efficiency. As future work we plan to study the frame aggregation of IEEE 802.11n in a multi-hop setting and the adaptive control of the aggregation delay for IP packet aggregation to cope with different traffic load.

Bibliography

[1] W. John and S. Tafvelin, “Analysis of internet backbone traffic and header anomalies observed,” in Proc. ACM SIGCOM IMC ’07. New York, NY, USA: ACM, 2007, pp. 111–116.
[2] “IEEE Standard for Information technology - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements,” IEEE Std 802.11e-2005 (Amendment to IEEE Std 802.11, 1999 Edition (Reaff 2003), pp. 1–189, 2005.
[3] D. Skordoulis, Q. Ni, H.-H. Chen, A. Stephens, C. Liu, and A. Jamalipour, “Ieee 802.11n mac frame aggregation mechanisms for nextgeneration high-throughput wlans,” IEEE Wireless Communications, vol. 15, no. 1, pp. 40–47, February 2008.
[4] Y. Kim, S. Choi, K. Jang, and H. Hwang, “Throughput Enhancement of IEEE 802.11 WLAN via Frame Aggregation.” in Proc. IEEE VTC 2004, September 2004.
[5] S. Mangold, S. Choi, and N. Esseling, “An error model for radio transmissions of wireless LANs at 5GHz,” in Proc. Aachen Symposium 2001, pp. 209–214.
[6] K. Kim, S. Ganguly, R. Izmailov, and S. Hong, “On packet aggregation mechanisms for improving voip quality in mesh networks,” in Proc. IEEE VTC 2006-Spring, vol. 2, May 2006, pp. 891–895.
[7] R. Riggio, D. Miorandi, F. De Pellegrini, F. Granelli, and I. Chlamtac, “A traffic aggregation and differentiation scheme for enhanced qos in ieee 802.11-based wireless mesh networks,” Comput. Commun., vol. 31, no. 7, pp. 1290–1300, 2008.
[8] “Multi-generator (mgen),” URL: http://cs.itd.nrl.navy.mil/work/mgen/.
[9] D. Malone, K. Duffy, and D. Leith, “Modeling the 802.11 distributed coordination function in nonsaturated heterogeneous conditions,” IEEE/ACM Trans. on Networking, vol. 15, no. 1, pp. 159–172, Feb. 2007.
[10] P. Dely, A. Kassler, and D. Sivchenko, “Theoretical and Experimental Analysis of the Channel Busy Fraction in IEEE 802.11,” in Proceedings of Future Network & Mobile Summit 2010, June 2010.
[11] H. Schulzrinne and S. Petrack, “Rtp payload for dtmf digits, telephony tones and telephony signals,” United States, 2000.
[12] R. G. Cole and J. H. Rosenbluth, “Voice over IP performance monitoring,” SIGCOMM Comput. Commun. Rev., vol. 31, no. 2, pp. 9–24, 2001.
[13] “ITU-T Recommendation G.107, International Telephone Connection and Circuits - General Definitions. The E model, a Computational Model for Use in Transmission Planning,” ITU-T, Tech. Rep., March 2005.
[14] R. Jain, D. Chiu, and W. Hawe, “A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Systems,” Technical Report DEC-TR-301, Tech. Rep., 1984.