Using A Backlogged Queue Approach For Adaptive MAC Frame Aggregation

Gautam Bhanage, Dipankar Raychaudhuri, Ivan Seskar
WINLAB, Rutgers University, RT 1 South, Technology Center, NJ 08902, USA
Email: {gautamb, ray, seskar}@winlab.rutgers.edu

***Abstract*—Frame aggregation is a wireless link optimization mechanism that aims to reduce transmission overheads by sending multiple frames as the payload of a single MAC frame. Static assignment of frame aggregation parameters can result in delay penalties due to variations in traffic type or load levels. Another possible side effect is an increase in packet error rate in noisy environments due to large aggregated frame size. Adaptive aggregation methods have previously been proposed to deal with the above problems independently, but there is still a need for a unified adaptation algorithm that addresses both aspects in an integrated manner. In this study, a backlogged queue (BQ) aggregation approach is proposed that considers both these aspects, and also ensures inter-operability with other WLAN devices that are not capable of frame aggregation. Performance evaluation of the proposed algorithm on the ORBIT testbed shows throughput improvements of up to 56% in the presence of channel noise and 25% in scenarios with high contention over using a simple *txop*. An experimental case study shows improvement in FTP file transfer times of up to 11% while preserving performance for real time traffic.**

***Index Terms*—MAC aggregation, 802.11e, adaptive framing, BQ framing**

I. Introduction

Frame aggregation is widely studied in 802.11 networks as a solution for reducing MAC layer overheads caused by packet headers and inter-frame spacing. Frame aggregation works by combining multiple data packets in the transmit buffer into a single MAC frame. However, aggregated frames may result in undesirable side effects, such as an increase in access delay due to wait time for frame arrivals in the transmit buffer, or degradation in packet error rate (PER) because of increased frame length. Adaptive frame aggregation which considers both these factors is a possible solution with the potential of providing improved network throughput without affecting end-user performance parameters significantly. In this study, we propose, analyze, and experimentally validate a backlogged queue (BQ) based aggregation scheme which achieves high channel efficiency in the presence of bounded delay requirements and channel noise.

Making the adaptive aggregation compatible with widespread commercial off the shelf (COTS) hardware is another important design constraint. To make our solution compliant with other 802.11 WLAN devices, we identify *txop* as a useful control mechanism in 802.11e systems. A transmission opportunity (*txop*) as described in the *802.11e* standard defines the amount of time for which a transmitter is able to reserve the medium for sending information. Typically, this *txop* determines the number of frames that the station will be able to send back to back without having to contend for channel access. Hence, within the allocated *txop*, the sender can transmit multiple consecutive frames within the SIFS time (no aggregation) or multiple frames aggregated as a single frame in the same *txop*. Our proposed BQ algorithm, leverages this standardized *txop* metric as a means of controlling aggregation.

A *delay adaptive scheme* waits only for a certain time duration for frames to be available for aggregation. Instead of making the approach explicitly *delay adaptive*, using observations made from VOIP call aggregation mechanisms presented in [21], we make the algorithm load adaptive to achieve good throughput-delay performance. Further evaluation of our model reveals that limitation on aggregated frame sizes imposed due to *txop*, and or the wireless MTU of 2272 bytes [16], results in an optimal performance with aggregation always turned on. This performance is superior than that could be achieved by making aggregation explicitly *noise adaptive*, where it would be dependent on link PER. Using these design insights, our backlogged queue (BQ) algorithm provides guidelines for making the aggregation decision on a per frame basis.

Specifically, the contributions of this paper can be enumerated as-

- Modeling: This study discusses a comprehensive analytical model for a frame aggregation mechanism that is compliant with the *802.11e* standard. Performance comparison accounts for gains due to presence of relevant transmission opportunities where frame aggregation is not used.
- BQ Frame Aggregation Algorithm: Using insights obtained from initial evaluations, we propose the Backlogged queue (BQ) frame aggregation algorithm that provides a solution which is compatible with commodity hardware, while also taking both delay and noise related Packet Error Rate (PER) into account.
- Experimental Evaluation: By doing a controlled evaluation of a static 2-Frame MAC aggregation scheme in the MadWiFi [4] driver, we emulate extreme case performances that could be achieved by our BQ algorithm in a standard WLAN deployment. Results show that aggregation based on our BQ algorithm, allows for superior performance in common WLAN conditions.
- Qualitative Discussion: Discussion based on current in-

978-1-4244-7265-9/10/$26.00 ©2010 IEEE

ternet trends and WLAN operating conditions are used to highlight salient features of our proposed algorithm and possible implications.

The rest of the paper is organized as follows. Section II provides a description of related studies that discuss link optimization using frame aggregation mechanisms. Section III describes our analytical model and shows some basic performance comparisons. These results motivate the need for our BQ frame aggregation scheme, that allows for better performance in common WLAN scenarios, while allowing a standards compliant approach to enable operation on COTS WLAN hardware. Section VI helps emulate the performance that could be achieved by our BQ aggregation algorithm by leveraging a static 2-Frame aggregation scheme in the driver. Finally, Section VII provides a few concluding points and direction to future work.

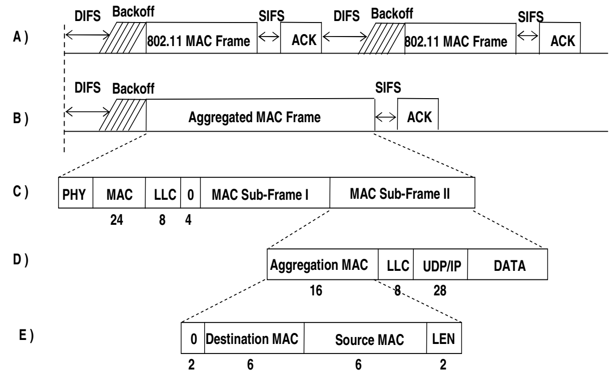

Fig. 1. Frame aggregation for the creation of a lumped MAC frame which has the two aggregated frames as the payload. This approach generates savings on *difs*, *backoff* times and physical layer headers. Subparts (A) and (B) of the figure compare timelines for transmission with and without aggregation while (C), (D), and (E) describe the payload of an aggregated frame.

II. Related Work

In the 802.11 context, frame aggregation has been proposed by studies [12] [20] as a mechanism for optimization by running it above the MAC service access point. Such an optimization also results in energy conservation for participating stations [10]. Aggregation has also been proposed as a means for optimization of real time traffic delivery [11]. However, all of these studies determine the performance of the aggregation mechanism under ideal link conditions.

Making frame aggregation adaptive has been proposed by studies for different reasons. Some studies [13] argue that increasing the number of frames aggregated will increase the amount of delay, thereby introducing a tradeoff. A key observation in the context of improving VOIP call capacity [21] showed that delays are largely determined by the backlogging of VOIP frames. We build on this approach by making generic layer 2 frame aggregation load adaptive, while also considering other features like channel noise, and compliance with the *802.11e* standard.

One of the earliest studies [14] to propose adaptive framing suggests that reduction of frame size in the presence of noise results in superior performance. However, this study does not consider the possible performance benefits due to a recent range of sophisticated transmit rate adaptation algorithms. Another study [7] discusses the option of joint rate and frame size calculation for optimum performance. However, the design is not validated on an operating prototype and is hard to implement on currently available commercial off the shelf hardware due to the use of a multi-level carrier sense mechanism. In $802.11n$, a study [15] proposes making frame aggregation adaptive for controlling link packet error rate in the presence of errors introduced due to large frames. However, this solution is not available on commercial WLAN devices. A similar scheme is proposed in [9], where partial packet recovery is made possible in 802.11 to allow for better link efficiency. However, we will show from our analysis and measurements that typically up to the 802.11 MTU, even at BERs of 10^{-4}, aggregation provides superior performance. Our measurements show that even though the individual PER increases with aggregation, savings in channel time are significantly higher even at the lowest of transmit rate to always justify aggregation.

Also, among all of the studies discussed here, only [9], [14] provide evaluations over actual hardware, where [9] uses the Click router [1] above the MAC layer and [14] relies on pre-IEEE802.11 WaveLAN cards. Our study presents results based on aggregation done at Layer2 in the network stack, while working with the commercial ATHEROS 5212 chipset cards. Though this evaluation is presented by leveraging a static $2-Frame$ aggregation in the driver , it helps to quantify the performance limits that may be achieved by an implementation of our BQ algorithm in the wireless driver.

III. Analytical Model For Aggregation

In this section, we compare the advantages of aggregating multiple frames before transmission (aggregation on) with the transmission of those frames within SIFS of each other (aggregation off). Fulfillment of this condition makes our model different from all previous approaches, and is fundamental to ensure compliance with $802.11e$ systems where a set of flows may have an equal *txop*, but the decision to aggregate could be made per flow.

A. Aggregation Model

Frame aggregation combines two or more MAC frames as shown in Figure 1 and sends them as the payload of a single frame with a custom Aggregation MAC header. To determine performance gains we estimate transmission time which includes the channel access time and transmission time for a MAC frame. The total channel time required for the transmission of a MAC frame is given as: $T_{total} = T_{backoff} + T_{difs} + T_{phy} + T_{mac} + T_{llc} + T_{data} + T_{ack} + T_{sifs}$. T_{difs}, T_{sifs}, and $T_{backoff}$ are taken in accordance with the 802.11a

[1]*Fast Framing* is a technique employed by Atheros chip manufacturers to achieve higher throughput values as an optimization for improving 802.11a/b/g system performance by utilizing MAC frame aggregation (allowing frame sizes up to 3000Bytes).

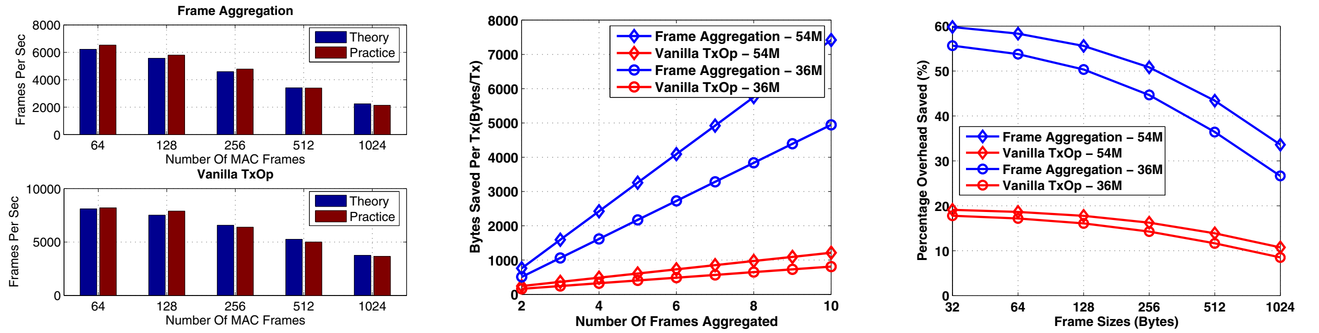

(a) Model Validation (Comparison in fps)

(b) Equivalent amount of air time saved in Bytes per transmission

(c) Percentage of overheads saved in channel time per transmission as a function of frame size.

Fig. 2. Performance comparison of frame aggregation and a vanilla transmission opportunity based setup as seen at the receiver.

standard [2]. Time taken to transmit PHY header/preamble T_{phy} and a MAC ACK T_{ack} are fixed irrespective of transmission rates. Transmission time for the MAC header T_{mac}, Link-Layer control Header T_{llc} and UDP/IP payload T_{data} depends on the transmission rate used. The total channel time for an aggregated frame will change to:

$$T_{total\_aggr\_N} = T_{backoff} + T_{difs} + T_{phy} + T_{mac} + T_{llc} + N \times (T_{aggr\_mac} + T_{llc} + T_{data}) + T_{ack} + T_{sifs} \quad (1)$$

where T_{aggr_mac} is the time spent to transmit an Aggregation header which is smaller than the standard MAC header because it will only have information of offset of actual frames within the payload.

The amount of overhead reduced due to frame aggregation of *N* frames for the same amount of UDP/IP payload can be calculated as:

$$Savings_{with\_aggr\_N} = [N-1] \times [T_{backoff} + T_{difs} + T_{phy} + T_{mac} + T_{ack} + T_{sifs}] - N \times (T_{aggr\_mac} + T_{llc}) \quad (2)$$

Since the T_{aggr_mac} and T_{llc} overheads added due to aggregation (as shown in Figure 1) are always lesser than the times spent in sending multiple ACKs, PHY, MAC headers, and channel contention, aggregation always yields bounded savings in channel time.

B. Vanilla TxOP Model

The *txop* feature denotes the transmission opportunity for the device. If the *txop* is set to the transmission time for *N* frames, the device is allowed to send *N* consecutive MAC frames within *SIFS* time of each other. While comparing performance benefits due to aggregation with that without aggregation, it is essential to determine the impact of *txop*. In the event of frame aggregation being disabled, and the transmit queue being backlogged, an $802.11e$ compliant transmitter should send multiple frames within $SIFS$ of each other as long as they use less time than the *txoplimit*. To determine savings with a vanilla *txop* setup (aggregation disabled), we calculate channel time consumed for sending *N* consecutive frames within *sifs* as:

$$T_{total\_txop\_N} = T_{backoff} + T_{difs} + N \times [T_{phy} + T_{mac} + T_{llc} + T_{data} + 2 \times T_{sifs} + T_{ack}] - T_{sifs} \quad (3)$$

Hence, savings in channel time due to the use of a vanilla *txop* scheme over conventional DCF operation are evaluated as-

$$Savings_{txop\_N} = N[T_{backoff} + (T_{difs} - T_{sifs})] \quad (4)$$

Using these calculations in our model, we present a theoretical performance comparison of using frame aggregation versus a vanilla txop scheme for different WLAN scenarios.

IV. Theoretical Performance

Using the analysis derived for a $802.11e$ system from the previous section, we present a preliminary comparison of performances with frame aggregation and the vanilla *txop* setup. Through our theoretical modeling we show that in common WLAN scenarios, aggregation up to the wireless MTU always provides superior performance, as compared to the vanilla *txop* scheme. We begin with a baseline performance estimation by evaluation under ideal link conditions. This is followed by analysis in noisy and interference dominated settings.

A. Ideal Link

We define a wireless link as ideal if we do not see any frame loss at the receiver due to noise or interference (packet error rate PER is 0). We assume such an ideal link between an access point (AP) which acts as the sender and a client acting as the receiver operating in $802.11a$ mode. The transmitted data is backlogged UDP traffic, and physical layer rate is fixed at 54Mbps.

Figure 2(a) shows the number of frames that are seen on the channel under saturation conditions with different frame sizes. Only for the sake of validation of our analytical model, we compare results with those obtained from an actual experiment under similar "ideal link" settings. In this evaluation, we measure the performance such that aggregation sends two frames combined as one, and the vanilla txop scheme sends two consecutive frames in SIFS time of each other. It can be

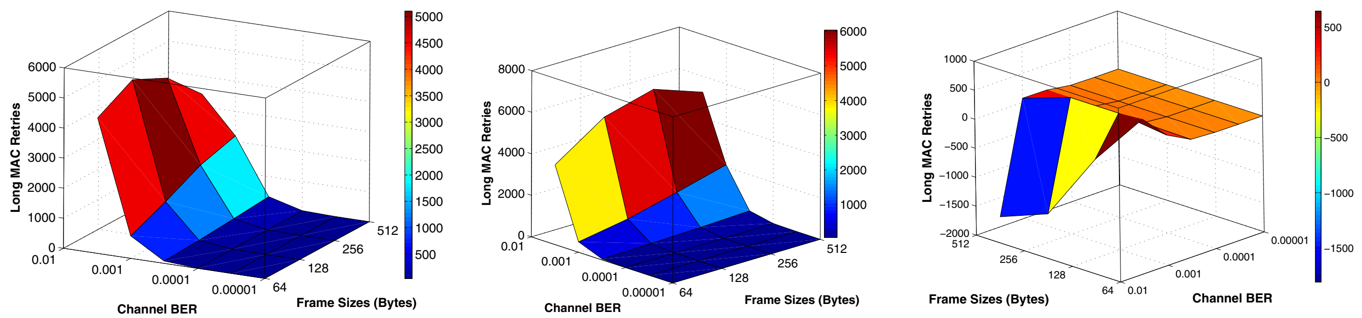

(a) Aggregation

(b) Vanilla *txop*

(c) Difference

Fig. 3. Number of long MAC retries in a single noisy link with and without aggregation.

observed that results from the experiments (practice) conform with those obtained from the model (theory), thereby helping validate the model. We also observe that the number of frames transmitted over the link are significantly reduced by aggregation, thus reducing overheads and improving performance. The rest of this section will present a discussion only based on evaluation from our analytical model.

Figure 2(b) shows the savings in number of bytes per transmission over a transmission with no aggregation or *txop* as a function of the number of 1024Byte frames aggregated. This result is determined by calculating the number of extra bytes that can be packed on the channel due to savings in channel time per transmission. These savings are determined for both aggregation, and the vanilla *txop* scheme. We observe that as the number of frames aggregated together increase, there is a corresponding increase in the number of extra bytes saved per transmission. These savings are higher for faster transmission rates because we can send more bits in the same channel time savings. In all cases we find that aggregation provides a significantly superior performance.

Figure 2(c) shows the savings in terms of percentage of channel time as a function of frame size. All results are with aggregation and *txop* are for two frames, either being aggregated or sent within SIFS of each other. It is seen that in all cases there is approximately 200% improvement in savings with aggregation over a vanilla *txop* scheme. Link optimization benefits are relatively higher for small frame sizes due to higher PHY/MAC overheads for transmitting the same amount of data.

Hence, purely based on modeling, we observe that performance benefits due to frame aggregation are maximum for higher rates and small frame sizes under our ideal channel evaluation. We will now investigate the impact of channel noise on performance with frame aggregation.

B. Noisy Environments

We simulate noise by varying BER at the receiver. To determine performance with noise, the number of long MAC retries for every frame size and BER combination are evaluated. The long MAC retries are defined as the number of frame transmission retries for long (data) frames. This metric is indicative of the link efficiency and overall performance.

The results with and without 2-frame aggregation are as shown in Figures 3(a) and 3(b). In both cases, for small BERs, the number of MAC retries per second are fairly less and the change is not noticeable across increasing frame sizes. As the level of noise on the channel increases, we see an almost identical increase in the number of retransmissions in both cases. We observe that even though the packet error rate (PER) increases with increasing frame length (doubling due to aggregation in this case), the actual number of frames transmitted are reduced to half which results in fewer retransmissions for large frames. With the highest BER, the PER reaches close to unity resulting in almost all frames being in error. In such a case since we transmit fewer frames with aggregation, we see fewer retransmissions.

These results are made clear in Figure 3(c) which shows the difference in frame retransmissions from those in Figures 3(a) and 3(b). We observe that retransmissions are more for the vanilla txop scheme in most noisy environments due to the near doubling of transmitted frames without aggregation resulting in large number of channel accesses and higher overheads. These results show that even in the presence of increased noise, performance with aggregation is superior for reasonable frame sizes close to the MTU (2048Bytes in this case) .

C. Interfering Links

Another network scenario commonly seen with multi-user WLAN deployments is interference. Previous studies [6] have shown that aggregation will consistently have superior performance with increased contention. This is also clear from Equations (2) and (4) which show that savings will scale with increasing $T_{backoff}$. Even while using RTS/CTS, use of aggregation always produces savings in overheads, since the same amount of transmitted data now requires lesser number of RTS/CTSs.

Thus evaluation based on our analytical model shows that in the presence of modest packet sizes (up to the MTU), aggregation provides significant benefits over no aggregation. In the evaluation with noise, we consider aggregation of only two frames as an estimate of worst case performance, since

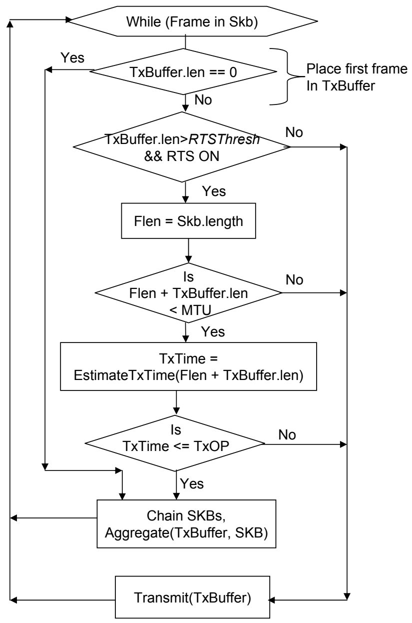

Fig. 4. Proposed algorithm for frame aggregation. Note that the *Aggregate()* routine will remove the frame from the transmit queue while combining frames and the *Transmit()* routine will initialize the transmit buffer to empty after successful transmission. This algorithm does not detail the necessary fragmentation algorithm in the event that TxOP requirements are not met even for a single non-aggregated frame.

aggregation of more number of (possibly smaller) frames together, produces a greater reduction in overheads as observed in Equation 2. Using insights obtained from these results, we formalize our BQ algorithm for frame aggregation.

V. Backlogged Queue (BQ) Frame Aggregation

We begin with a description of the algorithm which is followed by a brief description of its impact on aggregation decision over different frame sizes and link conditions.

A. BQ Algorithm

The key principle in the algorithm is to *aggregate whenever there are backlogged frames which can be transmitted within a single txop*. This algorithm should be used by the transmitting station if the receiving station supports the aggregation algorithm and is capable of extracting original frames from the received aggregated MAC frame. The algorithm is as shown in Figure 4. For every frame received from the higher layer, the algorithm tries to concatenate it to the transmit buffer as long as the necessary conditions are met.

When the transmit buffer is empty, the current frame to be transmitted is directly included in the transmit buffer. This condition is only fulfilled if the transmit buffer was recently emptied due to a transmission. Every following frame can be considered for aggregation only if: the size of the frame in the TxBuffer is larger than the RTS Threshold, the total length of the aggregated frame is smaller than the MTU, and if the total transmission time[2] to reach the destination is less than the *txop*. If either of these checks fail, the algorithm decides to just transmit the current contents of the TxBuffer[3]. On the other hand if these checks are passed, contents of the kernel socket buffer are merged with the existing transmit queue. Typically, such a merging also involves inserting proper headers within the aggregated frame itself for delineation of frames. This loop of concatenating frames continues as long as none of the checks are failed. Whenever a check fails, the TxBuffer contents are transmitted and the buffer is made empty for the next round of aggregation.

Features of our BQ algorithm that are important from a deployment perspective are:

1) It works independent of the rate control algorithm. This allows the use of state of the art approaches to rate/power control to limit losses due to channel errors.
2) Does not require a special mechanism (such as multi-level carrier sense) for making aggregation decisions.
3) It ensures that the station does not take more than *txop* of channel time and allows for co-existence with other non-aggregating stations.

We will now discuss the performance impact of the BQ algorithm on different aspects of the setup.

B. Performance With RTS/CTS

The hidden node problem problem has been cited among top ten problems for wireless data [18]. One approach to alleviate this problem is through the use of the virtual carrier sense mechanism [2]. Typically, a *RTS threshold* can be set in the wireless drivers which controls the size of the smallest MAC frame for which the driver will use the RTS/CTS handshake. Assuming that the *RTS threshold* is optimally chosen by the system administrator, the link benefits are maximum when we send the largest possible payload while using RTS/CTS and send the frame unmodified otherwise. Based on these observations, our BQ algorithm (Fig. 4) does not aggregate if the frame being transmitted is smaller than the *RTS threshold*. On the other hand, if the frame is larger than the *RTS threshold*, it tries to aggregate as per the enumerated conditions. Hence, even in a case with hidden nodes, using our algorithm to aggregate frames will result in a superior performance.

C. Impact On MAC Frames

Using modeling with a single sender - receiver link, we present results for describing the impact of our algorithm on aggregated frame sizes as shown in Figure 5.

Results in Figure 5(a) shows the amount of delay that would be experienced purely due to frame transmission with and without aggregation for two frame sizes, and transmission rates of 6 and 54Mbps. Extreme values are chosen for frame sizes and transmission rates to show range of operation.

[2]The rate information used in such calculations is either static if set explicitly, or readily available through the table maintained by the auto rate control algorithm within the driver.

[3]In some cases the un-aggregated contents of the transmit buffer might need fragmentation. For the sake of brevity we do not deal with those cases.

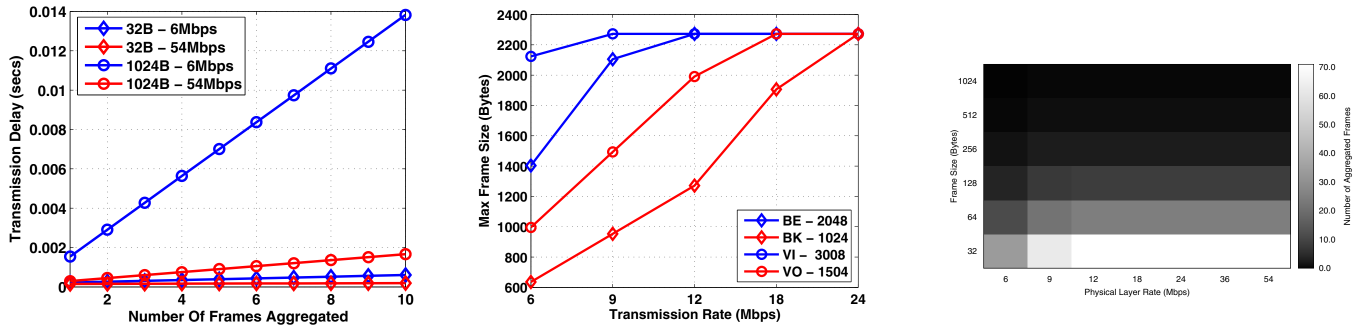

(a) TxDelay - Aggregation

(b) Max Frame Length - TOS

(c) Number Of Aggregated Frames

Fig. 5. Results from modeling show the impact of frame aggregation on the transmission delay while not accounting for the queuing delay seen by individual frames. It is assumed that in all cases, aggregation results in a reduced queuing delay at every terminal resulting in better performance. Due to TxOP constraints imposed by our opportunistic aggregation algorithm, we present results obtained experimentally by testing performance on a single link with different TOS bytes for service classifications. Finally, we show the implication of our algorithms constraints on the number of backlogged frames that may be aggregated for different rate - frame size combinations.

The delay values presented in this result are similar to the ones that will be estimated by the driver when running our aggregation algorithm. The estimated delay for 1024Bytes at 6Mbps is the highest due to the requirement of sending a large amount of data at the slowest possible speed, and vice versa for 32Bytes at 54Mbps. If the estimated transmission time of the aggregated frame falls in the *txoplimit*, aggregation is permitted. For example, with decisions based purely on txoplimits of 2048μsecs, we would not allow any aggregation for 1024Byte frames transmitted at physical layer rate of 6Mbps, while we would allow transmission of up to 10 frames of the same size while transmitting at 54Mbps. Note that our algorithm also employs other checks to contain the aggregated frame sizes to the MTU which would not allow a 10MByte frame size.

Figure 5(b) shows the maximum permissible frame size for an aggregated frame for different WMM [3] traffic classes in 802.11e. As shown, transmit opportunities for the classes are set as: 2048μsecs for Best effort (BE), 1024μsecs for background (BK), 3008μsecs for video (VI), and 1504μsecs for voice (VO) traffic. Using calculations based on our algorithm we observe that at the highest level, frame size is limited by the MTU. Below the MTU, maximum frame size is limited by the estimated transmission time of the frame and corresponding *txoplimit* for the traffic class.

Finally, Figure 5(c) shows the number of frames that will be aggregated for different frame size and physical layer rate combinations. These results are based on a fixed *txoplimit* of 2048μsecs, and an MTU of 2272Bytes. We observe that in an extreme case our algorithm does not permit any frame aggregation, while in another it could allow for the aggregation of up to 70 backlogged frames. Practically, in such cases the aggregation will also be limited by the size of the transmit buffer.

D. Impact on Delay

The total delay (D) experienced by a frame is given as-

$$D = Q + T, \tag{5}$$

the sum of the queueing delay (Q) and transmission delay (T). Using Little's Law, the average amount of queuing delay experienced in a system can be attributed to the number of backlogged frames L in the system as: $Q = \frac{L}{\lambda}$. In this system, λ denotes the arrival rate of frames from the higher layer to the MAC. From the above equations, we deduce that the overall delay (D) is evaluated as-

$$D = T + \frac{L}{\lambda} \tag{6}$$

When studying the delay impact of aggregation, we consider the following conditions,

$$D = \begin{cases} T & \text{for } L = 0 \\ T + \frac{L_{aggr}}{\lambda} & \text{With Aggregation} \\ T + \frac{L_{no-aggr}}{\lambda} & \text{With No Aggregation} \end{cases}$$

Since our system does not aggregate or wait for aggregation in the absence of backlogged frames ($L = 0$), delays achieved below saturation are only limited by transmission delay and hence similar to a standard MAC mechanism. In the presence of backlogged frames in the system, the number of residual frames in the queue L, are directly proportional to the efficiency of the MAC mechanism in transmission. As observed from the theoretical results, and as will be shown through experimental evaluations, aggregation recommended by our algorithm always produces better throughput per link, thus resulting in a higher service rate (μ) of frames from the transmit queue. Considering poisson arrival times, and exponential service times, the average number of frames in the queue based on an M/M/1 model can be evaluated as: $L = \frac{\rho^2}{1-\rho}$, where $\rho = \frac{\lambda}{\mu}$. Thus a higher service rate (μ) due to aggregation will lead to a smaller ρ, and hence lesser frames being backlogged. This results in reduction of the average length of the backlogged queue (L) to L_{aggr} which is always less than that without aggregation $L_{no-aggr}$. Hence, for any given λ, the average long term delay in the system will be the least if we *always* aggregate *only* in the presence

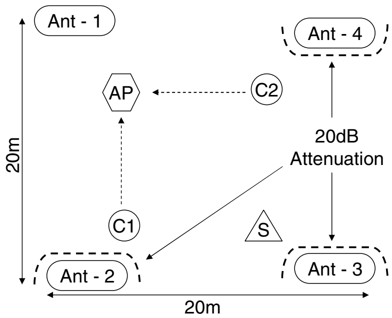

Fig. 6. Experiment setup on the ORBIT radio grid. Figure indicates scaled relative position of entities for measurements on the ORBIT grid. *C1* and *C2* represent the clients sending traffic to the AP. The four noise injection antennae (Ant - *) are located at the four corners of the grid with only Ant - 1 used to pump noise at the receiver running on the AP. *S* is a sniffer used for inter-frame delay measurements.

Parameter	*Value*
Experiment Duration	120 secs
Averaging Duration	Per sec
Operation Mode	802.11a Infrastructure
Frequency	5.18GHz
Chipset	Atheros AR5212
Driver	MadWiFi 0.9.4
Tx-Diversity/Virtual-CS	Disabled

Fig. 7. Common experimental parameters used with ORBIT nodes. Other parameters such as channel rate and packet sizes which may vary with experiments are mentioned explicitly.

of backlogged frames, thus eliminating any performance-delay tradeoff.

E. Baseline 2-Frame Setup

In the following section we will provide experimental results that would justify the use of our BQ algorithm. For the purpose of evaluation, we use a static 2-Frame aggregation already built in the MadWiFi driver. Baseline information on the working of this aggregation feature and its impact on scientific experiments have been explored in a previous study [6]. We leverage this two frame aggregation scheme since it helps us to demonstrate the benefits of the BQ algorithm even under extreme packet sizes and worst case operating conditions. Typically, benefits of our BQ aggregation will be higher with more number of aggregated frames. All aggregation decisions in the driver are controlled through private *ioctl* calls.

VI. Experimental Results

A. Testbed And Topology

We use the ORBIT testbed facility [8] which consists of 400 802.11 wireless nodes arranged in a $20m \times 20m$ grid. As shown in the Figure 6, four noise injection antennae are incorporated in the testbed that allow controlled injection of AWGN noise at desired power and frequency. Other experiment parameters are as shown in Figure 7.

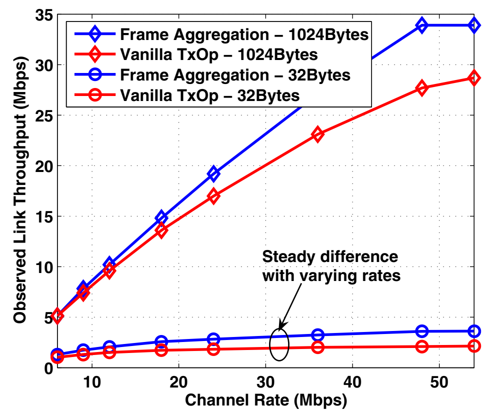

Fig. 8. Performance of aggregation against no - aggregation for extreme packet sizes with varying transmission rates. For small packet sizes, difference in performance due to aggregation is constant. However, gains with channel rate vary with large frame size.

B. Ideal Link - Baseline

To verify the model proposed in Section III, we perform baseline experiments with a high SNR channel between two nodes. In this case, the link is setup between C1 and AP as shown in Figure 6. Figure 8 plots the saturation throughput for a LOS link operating at different transmission rates ($6Mbps - 54Mbps$) and frame sizes. As determined by the theoretical model, aggregation produces a significant improvement with a high SNR channel. Even for the slowest of transmission rates and largest of frame sizes within MTU, where PHY/MAC overheads have the least impact, we see some performance benefits with aggregation. For a frame size of $1024Bytes$, we observe savings of almost 30% over a no-aggregation, no-txop throughput of $27Mbps$ which conforms to our modeling. Similarly, the benefits with vanilla *txop* are close to 10%, thereby conforming with our model. As observed in our modeling, we see better performance benefits at higher physical transmission rates and smaller frame sizes. On a related note, a baseline comparison of average delays not shown here had an improvement of 16% and 42% for packet sizes of 1024Bytes and 32Bytes respectively when operating under these conditions.

C. Noisy Environments

Analysis and theoretical performance estimation showed that at increased noise levels, aggregate number of MAC retries with no aggregation are more, which could lead to degraded performance. To validate our modeling, we setup an experiment which has a single sender *C1* sending UDP data frames to a receiver running at *AP* as shown in Figure 6. Noise is injected at the receiver through antenna *Ant-1* in a controlled manner using an $AWGN$ noise generator. Maximum aggregated packet size is chosen as recommended by the BQ algorithm at the MTU ($1024bytes \times 2$). UDP throughput is measured for a fixed transmission rate of 54Mbps.

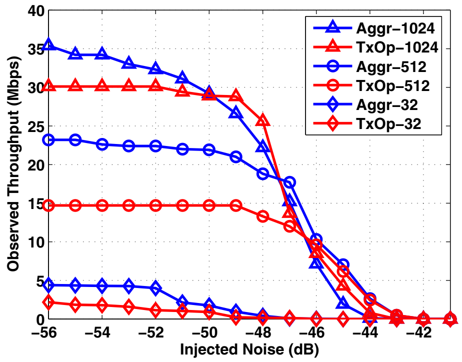

(a) Throughput Vs Noise at 54M

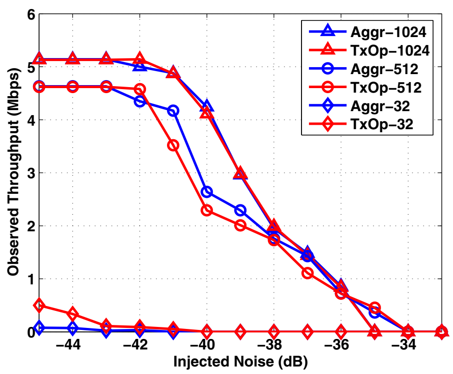

(b) Throughput Vs Noise at 6M

Fig. 9. Performance of links with varying levels of noise at the receiver in the presence and absence of frame aggregation. Results, show that performance with aggregation is always superior to that without aggregation.

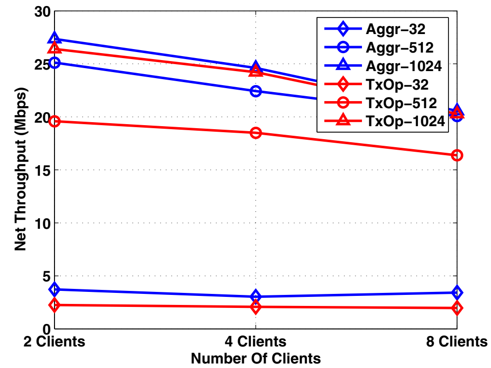

Fig. 10. Performance of links with varying levels of contention in the presence and absence of frame aggregation. Results, show that performance with aggregation is always superior to that without aggregation in the presence of high contention.

	Video		FTP
$Rate\|Param \rightarrow$	Mean$\Delta\mu s$	$Stddev(\Delta)\mu s$	$Time(s)$
$9M - Aggr$	653	1283	137.3
$9M - NoAggr$	664	1186	145.8
$18M - Aggr$	657.8	1210.9	67.8
$18M - NoAggr$	642	1190.2	73.6
$36M - Aggr$	600.4	1110.7	35.2
$36M - NoAggr$	600.55	1118.0	40.6
$54M - Aggr$	569.6	1036.4	26.7
$54M - NoAggr$	850.5	2234.3	31.7

Fig. 11. Performance of the joint video, FTP link performance test. Results show that as channel rate improves, FTP transfer time decreases. We also observe that in all cases aggregation provides superior performance for file transfer while not impacting the video streaming.

Results of this experiment are as shown in Figure 9(a). It can be observed that as the noise increases, throughput drops rapidly with injected noise. However, for all noise levels at which a link can be sustained, we see that the aggregate throughput seen with aggregation is consistently better (up to 56% better) than that without aggregation. We also observe that as expected, the performance benefits are higher with aggregation for smaller frame sizes.

Benefits of aggregation are the least at slowest of rates and large frame sizes. To determine a worst case performance with aggregation, we fix the transmission rate of the sender at 6Mbps which is the slowest rate for the $802.11a$ mode. Results are as shown in Figure 9(b). We observe that even at the slowest of rates and largest of frame sizes permitted by the BQ algorithm, aggregation always performs better on account of savings of small frame re-transmissions. These observations corroborate our analysis and validate our subsequent BQ algorithm that recommends aggregation up to the valid MTU or permitted *txop*.

D. High Contention

To determine performance in the presence of high contention, we measure observed throughput with multiple senders sending traffic to a single access point. UDP transmission rates at all senders are fixed at $54Mbps$. Results are plotted for $2, 4$, and 8 senders with frame sizes of $32, 512$ and 1024byte. As shown in Figure 10 we observe that across all frame sizes and all client combinations, the performance achieved with frame aggregation is significantly superior. A smaller number of frame transmissions with aggregation ensures that the time spent in channel contention is significantly reduced leading to better performance. As before, maximum savings of up to 25% are achieved for the smallest frame size.

E. Case Study: Combined Streaming - FTP

Real time applications such as audio and video are most affected by varying delays at receivers. Typically, real time audio applications have stricter delay constraints of less than 50msecs. Previous studies have shown that the measured variance in delay (jitter) can almost double with aggregation [6]. Though from an experimental measurement stand point, we would see a significant increase in the amount of jitter, most audio and video decoders are built with de-jitter buffers. These buffers help in alleviating variation in delays seen at receivers for *packetized* real time traffic. Using typical buffer sizes with popular video streaming, we will determine the impact of frame aggregation on delay variance from an application perspective.

The experiment emulates a usage scenario, where a client is downloading a 100MB data file using FTP, while also streaming a live video using the VLC [5] player. The FTP buffer size is maintained at the default 16KBytes. We measure the downlink video download time, and corresponding performance for the video streaming. Results are as shown in Figure 11. The streaming video performance is measured as Δ the difference in inter-frame times as seen at the video player. The results plot the mean and standard deviation of Δ in μseconds which are indicative of the video quality seen at the receiver. Figure 11 shows that even with a simultaneous FTP download, standard deviation of Δ with and without aggregation is around 1200μ secs which is significantly lower than typical buffer delay ($50 - 100$msecs even for real time traffic like video conferencing). It is also seen that the FTP performance improves significantly with aggregation in all cases. This improvement is particularly higher for higher rates, since aggregation has relatively low MAC/PHY overheads at these rates. Thus, we observe that using aggregation allows significant improvement with FTP file transfer while also maintaining video streaming performance thereby enhancing user experience.

F. Discussion

In this section we mention a few other factors that make our aggregation scheme important in current WLAN usage such as public internet hotspots.

(1)Internet Packet Sizes: Our BQ algorithm always advocates aggregation of small frame sizes in a WLAN setting. These results are particularly important since previous internet traffic studies [17] have shown that more than half of the packets are smaller than 200 bytes, which could lead to considerable control overhead and degraded efficiency over wireless edges. Another trend observed in a recent measurement study [19] showed that 40% of the packet sizes are around 40Bytes, further justifying the use of our scheme in such settings.

(2)Indirect Fairness Implication: Apart from the goal of ensuring inter-operability with other drivers, being 802.11e compliant also allows the access point to have a direct control on client fairness. Conforming to the *txoplimit* prevents any station from hogging the channel with large aggregated frames, thus allowing each of the clients to get a fair share of channel time. Such a channel time based fairness is typically important since it works irrespective of the transmission rate of the channel, and the packet sizes used by the sender. Also, since the WMM parameters are sent our by the access point (AP) as a part of beacons, depending on the load and typical delay requirements, the AP could periodically broadcast newer *txoplimit*. Such a mechanism also allows for APs to cater better to specific services, since the *txoplimit* can be explicitly controlled for different traffic classes.

VII. Conclusion and Future Directions

Instead of making the aggregation decision explicitly delay/noise adaptive, this paper proposes the BQ algorithm which makes it load adaptive while accounting for channel noise, and making the solution *802.11e* compliant. Such an approach yields good delay performance while also achieving superior performance in the presence of channel noise and contention. Our proposed algorithm also allows for inherent fairness control by limiting aggregation at individual stations, and works irrespective of other MAC mechanisms such as rate and power control. By leveraging the static 2-frame aggregation scheme already present in the MadWiFi driver, we show that aggregating as per the BQ algorithm, even if it is only limited to two frames always yields significantly superior performance over the plain use of *txop* in mid-sized WLANs. Future work involves MAC extensions and evaluations for multi-hop ad hoc networks.

Acknowledgements

Research supported in part by US NSF grant # *CNS-072505*

References

[1] Click Modular Router. *http://read.cs.ucla.edu/click/*.
[2] IEEE 802.11a standard. *http://standards.ieee.org/*.
[3] IEEE 802.11e standard. *http://standards.ieee.org/*.
[4] Multiband atheros driver for wifi. *http://madwifi.org/*.
[5] Video LAN Player. *http://www.videolan.org/vlc/*.
[6] G. Bhanage, R. Mahindra, I. Seskar, and D. Raychaudhuri. Implication of MAC frame aggregation on empirical wireless experimentation. In *IEEE Globecom 2009 Wireless Networking Symposium*, Honolulu, Hawaii, USA, Nov 2009.
[7] C. cheng Chen, H. Luo, E. Seo, N. Vaidya, and X. Wang. Rate-adaptive framing for interfered wireless networks. In *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, pages 1325–1333, May 2007.
[8] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh. Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols. In *IEEE WCNC* , March 2005.
[9] A. Iyer, G. Deshpande, E. Rozner, A. Bhartia, and L. Qiu. Fast Resilient Jumbo Frames in Wireless LANs. In *Globecom*, Dec 2009.
[10] J. Lorchat and T. Noel. Energy saving in IEEE 802.11 communications using frame aggregation. In *IEEE GlobeCom*, Dec 2003.
[11] K. Lee, S. Yun, I. Kang, and H. Kim. Hop-by-Hop Frame Aggregation for VoIP on Multi-Hop Wireless Networks. In *IEEE ICC*, May 2008.
[12] Y. Kim, S. Choi, K. Jang, and H. Hwang. Throughput enhancement of IEEE 802.11 WLAN via frame aggregation. In *Vehicular Technology Conference, 2004. VTC2004-Fall. 2004 IEEE 60th*, pages 3030–3034, September 2004.
[13] S. Kuppa and G. Dattatreya. Modeling and analysis of frame aggregation in unsaturated wlans with finite buffer stations. In *Communications, 2006. ICC '06. IEEE International Conference on*, volume 3, pages 967–972, June 2006.
[14] P. Lettieri and M. Srivastava. Adaptive frame length control for improving wireless link throughput, range, and energy efficiency. In *IEEE INFOCOM*, volume 2, pages 564–571, Mar-2 Apr 1998.
[15] Y. Lin and V. Wong. Wsn01-1: Frame aggregation and optimal frame size adaptation for ieee 802.11n wlans. In *IEEE GLOBECOM*, pages 1–6, 27 2006-Dec. 1 2006.
[16] T. Maufer. A field guide to wireless lans: for administrators and power users. *Prentice Hall Series in Networking*, 2003.
[17] S. McCreary and K. Claffy. Trends in wide area IP traffic patterns - A view from ames Internet exchange. *Proceedings of the 13th ITC Specialist Seminar on Internet Traffic Measurement and Modelling, Monterey, CA*, 2000.
[18] M. W. Ritter. The future of wlan. *in ACM Queue*, 1(3):18–27, 2003.
[19] R. Sinha, C. Papadopoulos, and J. Heidemann. Internet packet size distributions: Some observations. *ISI-TR-2007-643*, 2007.
[20] T. Li, Q. Ni, D. Malone, D. Leith, T. Turletti, and Y. Xiao. Aggregation with Fragment Retransmission in Very High-Speed WLANs. In *IEEE/ACM Transcations on Networking*, March 2006.
[21] S. Yun, H. Kim, H. Lee, and I. Kang. 100+ voip calls on 802.11b. *IEEE Journal on Selected Areas in Comm.*, 25(4):689–698, May 2007.