Usage of 802.11n in Practice: A Measurement Study

Naman Mishra, Computer Science and Engineering Department, IIT Bombay, namanmishra@cse.iitb.ac.in

Avinash Chaurasia, Computer Science and Engineering Department, IIT Bombay, avinashk@cse.iitb.ac.in

Arun Kallavi, Computer Science and Engineering Department, IIT Bombay, arunkallavi@cse.iitb.ac.in

Bhaskaran Raman, Computer Science and Engineering Department, IIT Bombay, br@cse.iitb.ac.in

Purushottam Kulkarni, Computer Science and Engineering Department, IIT Bombay, puru@cse.iitb.ac.in

***Abstract*—IEEE 802.11n offers several throughput enhancing features over its predecessor 802.11a/g. The two main features at the PHY layer are: MIMO and channel bonding, while the main throughput enhancing feature at the MAC/link layer is Frame Aggregation. While in theory, as well as in controlled experimental conditions, these features achieve throughput enhancements, the extent to which they are useful for in-the-wild deployments has not been studied thus far. This paper presents measurements from three sets of traces: one from a research conference, one from a busy airport, and the third from a dense classroom setting with extensive WiFi usage for classroom activities. Our findings are as follows: (a) the high data rates of 802.11n are not used substantially, although the presence of moderate rates is significant, (b) the use of the channel bonding feature is minimal in dense deployments, and (c) the percentage of bytes undergoing frame aggregation is considerable but the levels of aggregation are not very high. We also undertake controlled experiments which shed light on non-wireless, system bottlenecks. Specifically, we find that many clients are not equipped to handle high levels of frame aggregation. Worse, this interacts badly with the rate adaptation algorithm, significantly lowering the data rates being used. These issues need careful addressing before 802.11n features are used effectively.**

I. Introduction

The 802.11n standard [1] came out in 2009 with promise of significant throughput improvement: up to 600 Mbps raw data rate, up from a mere 54 Mbps in prior 802.11a/g systems. The main contributors to this raw improvement at the PHY layer are (a) the spatially multiplexed operation of MIMO, and (b) the channel bonding feature. The PHY also supports the short-guard-interval mode of operation, but this is a relatively minor feature in terms of throughput enhancement.

At the MAC/link layer, a main efficiency improving feature in 802.11n is that of frame aggregation. Specifically, the A-MPDU (Aggregated MAC Protocol Data Unit) mode of aggregation is widely supported in commercial WiFi hardware. In A-MPDU several link layer frames destined to the same receiver are aggregated at the PHY layer, with each frame being protected by an independent CRC. The A-MPDU frame aggregation feature is used with the corresponding Block-ACK feature where the receiver independently acknowledges each link layer frame sent in a prior aggregated A-MPDU frame. Such frame aggregation can significantly improve the MAC and link layer efficiency by cutting down on per-transmission header overheads and more importantly, channel contention overheads.

In theory, as well as in controlled experiments, it is easy to see the benefits of these enhanced features of 802.11n. But what is their usefulness in WiFi deployments in the wild? How often are they used and how effectively? These questions are the focus of this paper. The answers to these are important not only for network upgrade related questions (802.11a/g to 802.11n or the latest 802.11ac), but also more importantly for focusing efforts of future standards.

We start by analyzing data collected from a total of 27 hours of traces from three real deployment settings: a research conference, an airport at a busy time, and a dense classroom setting. To collect various statistics such as per-transmission aggregation, we have instrumented a commercial enterprise-grade 802.11n 2x2 Access Point. Our findings can be summarized as follows.

(1) Although the use of moderate data rates is substantial in all three cases, usage of high data rates of 802.11n (> 100 Mbps) is lower than expected. For instance, in the classroom, the AP-client distances are low enough to permit high data rates (e.g. 144 Mbps for 1x1 40MHz operation), as we verified through separate experiments. But we found a large fraction of 802.11n transmissions using data rates lower than this.

(2) Although a substantial percentage of transmitted bytes have undergone aggregation (68%, 29%, and 92% in the conference, airport and classroom setting respectively), the majority of the frames undergo no aggregation. So we lose out on the efficiency improvements due to frame aggregation most of the time.

(3) In the airport and conference settings, we see only small percentages of frames in channel bonded (40MHz) mode, but the classroom setting saw substantial usage of channel bonding. This reinforces the intuition that channel bonding is rare in dense, unplanned deployments.

(4) The estimated benefit of deploying 802.11n instead of 802.11g in terms of estimated additional transmission time with 802.11g is about 42% and 13% for conference and airport

978-1-4799-8439-8/15/$31.00 ©2015 IEEE

environments respectively. While this is noticeable, it is well below the order of magnitude improvement suggested by the raw PHY rate improvement.

(5) To explain the less-than-ideal usage of frame aggregation in our traces, we undertake several controlled experiments with a variety of clients having backlogged traffic. Surprisingly, even in these settings, we find less-than-ideal usage of frame aggregation. On probing this, we find that there are client-side bottlenecks in handling A-MPDU frames: MPDUs are lost between the client-side hardware and its driver! This interacts poorly with the rate adaptation algorithm as well, which treat these losses as channel losses. This leads to overall loss of achieved throughput, as well as less-than-ideal usage of frame aggregation in practice.

The introduction of features such as standards-mandated transmit beamforming, in 802.11ac, will likely improve the usage of higher data rates. However, other system bottlenecks such as client side issues need careful addressing, before the potential of 802.11n or the newer 802.11ac can be realized. Further, while there are research solutions available to address the poor interaction of rate adaptation and non-channel losses, these must be incorporated in future standards, for otherwise we stand to lose the potential of high speeds offered by these same standards, as observed in our measurements.

II. Related Work

Experimental studies in the recent past have focused on empirically characterizing 802.11n performance. The work in [2] concentrated on throughput measurements with various 802.11n configurations and usage scenarios: in isolation and in the presence of interference from another 802.11n link. The authors show degradation of throughput due to wider channel bandwidths. In [3], the authors studied 802.11n link quality at high data rates by taking Packet Delivery Ratio (PDR) as a metric. They also studied the effect of channel bonding on PDR and TCP goodput. Although both the studies offer insight into the performance of 802.11n, they have used controlled experiments, while our study looks at 802.11n deployments in the wild. While the prior studies answer what *could* be achieved using 802.11n, we answer what *is* indeed achieved in real deployments, using the 802.11n features.

Real-world trace-based measurements of prior versions of 802.11 have been done. In [4], the authors study 802.11b by analyzing wireless traces captured during SIGCOMM 2004. This is the closest to our work since it captures all activity "on the air" in the form of traces and analyzes them. However in contrast to this, our study focuses on 802.11n specific features.

Our controlled measurement on various 802.11n clients reports on the poor interaction between frame aggregation and rate adaptation: an aspect not reported in the above prior measurement studies.

III. 802.11n Features in the Wild

A. Trace Collection Setup

To study the performance gains due to 802.11n enhancements, in real-life deployments, we have collected wireless traces at various places. Three sets of traces have been captured — at a research conference (MobiHoc 2013), at Bangalore airport during a busy time, and in a lab-structured classroom of a departmental course. These locations represent dense and busy deployment settings with a large number of clients.

The conference and airport locations are likely to have a wide variety of applications having diverse usage of the network. In the lab-structured classroom, the clients were mostly using similar applications: recorded video material related to the lab activity, as well as general web access. All the three traces represent dense user scenarios; our study has not considered sparse environments. The first set of traces were collected over four days during the conference hours and spanned two locations at the same venue. The per-frame information was collected from the WiFi vendor's network management interface, which used a commercial enterprise-grade 802.11n 2x2 radio configured in promiscuous mode to capture all the traffic present "in the air". The sniffer radio was configured to switch between channel 1, 6 and 11 every five minutes, since we were not aware beforehand about the channel(s) the operational network would choose. PHY and MAC layer information was exported by the driver inside the sniffer, which was transferred on-the-fly through an HTTP connection to a laptop. The relevant PHY/MAC information exported by the driver include RSSI value, transmission rate, transmitter and receiver MAC address, frame length, type, subtype, and number of frames (MPDUs) in a transmission. With vendor support, specific sniffer instrumentation was carried out to get the frame aggregation information, as well as information about PHY layer channel occupancy (i.e. idle vs busy).

Trace	Frames (Million)	Airtime (Hours)	Access Points	Clients
Conference	8.5	23	25	182
Airport	0.75	0.8	49	137
Classroom	6.9	3	9	50

TABLE I: Summary of the collected traces

We note that the logs recorded only MAC and PHY layer information, with no IP or higher layer information available, and were part of the vendor AP's network troubleshooting functionality[1].

The second set of traces was captured at a busy airport. This activity was carried out for a short time duration (45 minutes) and captured much lower number of records as compared to the conference set, but being a setting that is very busy and having large number of users, it can provide useful information about 802.11n usage. The same methodology was used for collection as in the previous case.

The third set of traces was captured during a lab session of a departmental course. In this setup, one point of difference from the prior setups was that the same radio used for trace

[1]Further, we present only aggregate statistics, and do not analyse or report any individual client information.

collection also provided Internet access to the clients, doubling as an AP. We verified through extensive experimentation prior to trace collection that such live trace collection did not affect the performance of the AP (i.e. the AP was able to send/receive data at maximum throughput even while collecting traces).

Table I summarizes the traces. In counting the APs as well as clients, we have considered only those MAC addresses which send at least one data frame. An AP is a radio which also sends beacon frames. The number of radios in the channel suggest rather dense settings.

B. Channel Occupancy

From our instrumentation of the sniffer we get information about what fraction of the time the channel is detected busy (some energy seen on air) versus idle. This metric is computed every 2 seconds. Fig. 1 shows the CDF of the channel-busy percentage; the CDF is plotted considering all the 2-sec intervals for which trace was collected. We show this graph to get a sense of how busy the medium is. We see that in spite of the large number of radios, the channel is well below saturation most of the time; e.g. in the conference setting, the channel is less than 55% busy 90% of the time. We see that of the three traces, the airport setting shows the highest channel busy percentage.

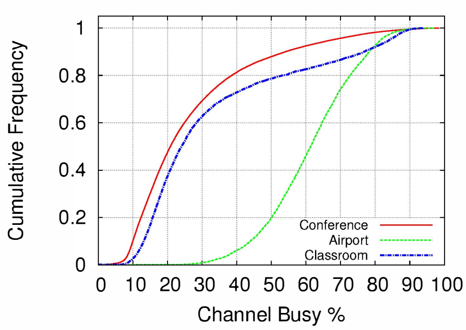

Fig. 1: CDF of channel-busy percentage

C. 802.11n vs Legacy 802.11 Traffic

The proportion of 802.11n and 802.11b/g traffic in terms of frames and bytes is shown in Figure 2. It helps in understanding the composition of the logs and the magnitude of 802.11n traffic we are dealing with. From Figure 2, it can be observed that although 802.11n frames are much less in number than legacy 802.11 frames, the number of 802.11n bytes is much larger as compared to 802.11b/g bytes except in the airport set of traces. This implies that the traces contain a significant amount of 802.11n traffic to analyze.

D. Usage of PHY Data Rates

Next we look into the PHY data rates used by the clients; especially we are interested in the 802.11n rates. We have plotted the distribution of rates across all 802.11n transmissions to find out which rates are the ones most frequently used. Figure 3 shows this distribution for all three sets. The rate distribution is very different for each of the three set of traces.

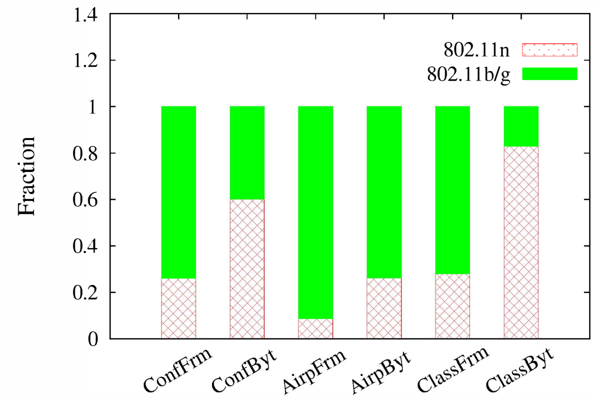

Fig. 2: Separation of 802.11n and 802.11b/g traffic in terms of frames and bytes seen

As can be seen from Figure 3a, moderate rates such as 58, 65, and 72 Mbps are dominant in the conference environment whereas relatively lower rates are predominant in the airport environment with a negligible fraction of rates $>$ 65 Mbps. The scenario is exactly opposite in the classroom environment where 150 Mbps alone accounts for 28% of the transmissions; in addition, proportion of rates $>$ 100 Mbps is significantly high. While Fig. 3 includes control and management frames in its statistics, we found that excluding these does not affect the overall conclusion as above.

We next look at the distribution of rates over transmitted bytes. From Fig. 3b, there isn't much difference in the pattern of the plots; the fraction of a few rates has gone up probably due to larger frames lengths being transmitted at those rates as a result of frame aggregation. Fig. 3c shows the distribution of air occupancy (only transmission) time across rates. The overall pattern remains the same here also except the marginal increase in share of low rates and reduction for high rates. Hence the same data rates dominating the frame and byte share are also dominating the airtime on the channel. Could the fact that we observed a high percentage of high data rates in the classroom traces be an artifact of the trace collection? After all, in the classroom, the traces were collected from the AP in service, while in the other cases, the traces were from a sniffer and not from the AP used for wireless access itself. Could it be that the sniffer missed the high data rate packets? We think this is unlikely. In the conference trace, the sniffer was in the same room as the AP, and was as far from the operational AP as the clients themselves. In the airport too, the sniffer was as far away from the operational AP as the clients, implying that we would have been in a position to capture high data rate transmissions (if any) from the clients to their APs. So its unlikely that trace collection was biased heavily toward reporting only low data rates.

We believe the following reasons contribute to the non-prevalence of high data rates in the conference and airport settings. First, these locations do not see much channel bond-

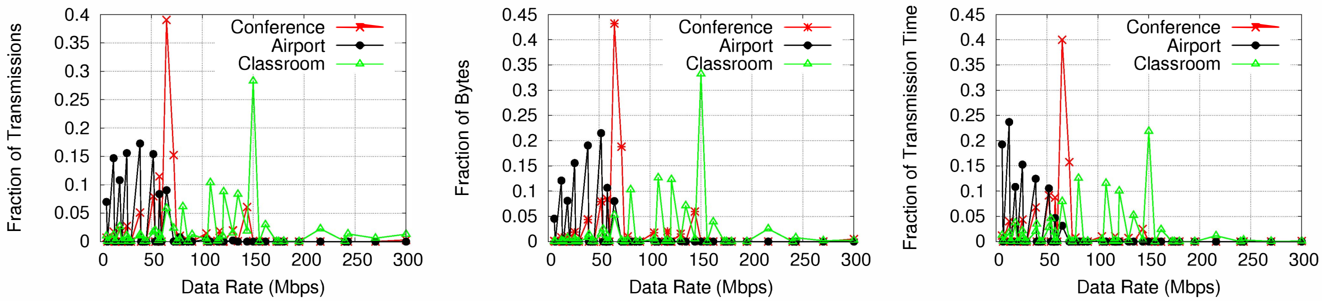

(a) Over No. of Transmissions

(b) Over Transmitted Bytes

(c) Over Transmission Time

Fig. 3: Rate distributions in all three settings

ing, as we will detail in the next section. Second, it is likely that in these settings most of the clients were smart-phones or tablets, most of which do not support beyond 1x1 MIMO configuration. It is worth noting in this context that recent research has indicated that 1x1 MIMO configuration is more energy efficient than using higher spatial multiplexing [5].

A third factor we surmise as a reason for relatively low data rates is the poor interaction of rate adaptation and frame aggregation related non-wireless losses. We elaborate on this using experiments in Sec. IV.

E. Channel Bonding Usage

We next find out the usage of channel bonding in 802.11n transmissions. Whether a transmission has been done using a 20 MHz or 40 MHz channel can be decided from the rate at which the frame was observed by the sniffer (AP)[2]. However due to a limitation in our sniffer AP instrumentation, we get only an integer floor estimate of the data rate of each transmission. This leads to ambiguity in 20MHz data rates versus 40MHz data rates in two cases: 13 Mbps and 216 Mbps. We have resolved this ambiguity by taking both these data rates as belonging to 40 MHz. Since the presence of these rates is not very high, either choice will not have a drastic effect on the results. The 802.11n standard specifies that an AP can operate in channel bonded mode only when there is no AP overlapping partially or fully with its secondary channel, or only partially overlapping with its primary channel. Given that the 2.4GHz band itself is only 80MHz wide, the likely cause of low usage of 40MHz channels in the conference and airport is the presence of other APs in overlapping 20MHz channels. In the classroom, although there are some other APs near the location of our deployment, they are weak in signal strength and only intermittently seen. Hence here the channel bonding usage seen is higher.

With increase in prevalence of 5GHz WiFi deployments, we may see better use of 40MHz channel bonding. The 802.11ac standard specifies channel bonding up to 160MHz, but the use of channels this wide may be difficult in dense settings, just as the use of 40MHz channel bonding is difficult in the current 2.4GHz band in dense settings.

[2]For the same MCS (Modulation and Coding Scheme), the sensitivity values for 20MHz and 40MHz operation are similar, which implies that in any setting, the sniffer is equally likely to capture either kind of transmission.

F. 802.11 Frame Aggregation Levels

We next investigate whether frame aggregation is being used significantly in practice and examine the typical aggregation levels. For this we plot the aggregation levels for 802.11n transmissions in each setting. Figure 4 shows the distribution of number of MAC frames (MPDUs) sent in a transmission opportunity for each of the three traces. We observe that majority of the 802.11n transmissions contain a single MAC frame (i.e. no aggregation) and only a small percentage of transmissions are undergoing any level of aggregation! Transmissions with no aggregation ranges from 64% in classroom traces to 91% in airport traces. This implies that usage of aggregation in terms of transmissions is not very high[3].

Since transmissions using aggregation are likely to be larger in size, we look at the distribution of aggregation levels with the corresponding fraction of 802.11n bytes transmitted. This too is shown in Figure 4. It can be seen that a larger fraction of bytes have been transmitted in aggregated frames. The figure shows that around 68% of bytes undergo some level of aggregation in the conference trace. This number is 29% for airport trace, and a much higher 92% for the classroom trace. These results indicate that a substantial amount of transmitted bytes undergo aggregation.

To further dig into the usage of frame aggregation, we look at the median level of aggregation in all three sets of traces. Considering all 802.11n frames, the median aggregation level is 1, i.e. no aggregation, in all three settings, as can be seen from Fig. 4. Therefore we look at median number of MAC frames per transmission, taking into account *only* those transmissions that undergo some level of aggregation. These values are as low as 3, 2, and 5 for the conference, airport, and classroom traces respectively. In comparison, as we observe in Sec. IV, for clients capable of sustaining high TCP throughput, we see aggregations levels as high as 15-30. So although considerable amount of bytes are undergoing aggregation, the median aggregation levels are not very high.

What does a specific aggregation level imply in terms of channel usage efficiency? Intuitively, the higher the aggregation level, the lower the effect of various fixed overheads. We can easily quantify this, by accounting for (a) the various

[3]The above result remained almost the same even when we considered only 802.11n DATA frames, excluding MGT and CTL frames.

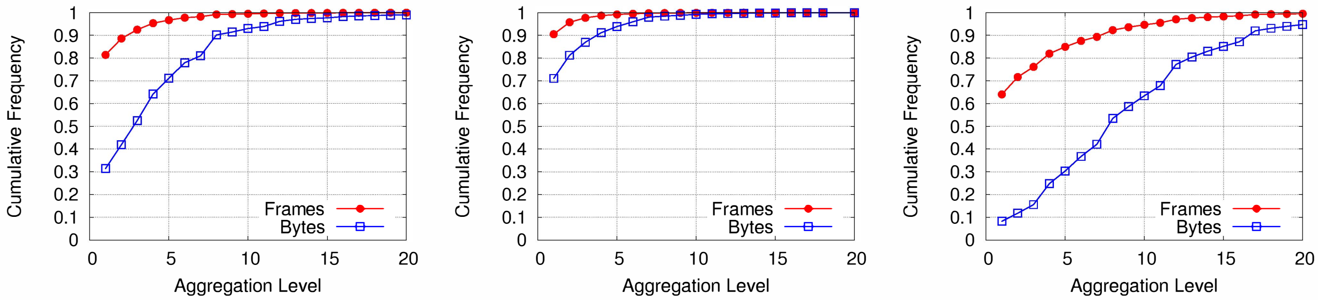

(a) Conference (b) Airport (c) Classroom

Fig. 4: Aggregation Levels over 802.11n transmissions

PHY and MAC overheads, (b) the DIFS, backoff, and SIFS overheads. The following equations gives the application level UDP and TCP throughputs after accounting for the above overheads; the throughput is a function of the UDP/TCP payload size P, the level of A-MPDU aggregation K, and the PHY data rate R.

$$T_{udp} = [101.5 + 36 + (70 + P) \times 8 \times K/R + 84]\mu s$$

$$Thrpt_{UDP} = [P \times 8 \times K/T_{udp}]Mbps \quad (1)$$

$$T_{tcpdata} = [101.5 + 36 + (82 + P) \times 8 \times K/R + 84]\mu s$$

$$T_{tcpack} = [101.5 + 36 + 82 \times 8 \times K/R + 84]\mu s$$

$$Thrpt_{TCP} = \frac{2 \times P \times K \times 8}{2 \times T_{payload} + T_{ack}} Mbps \quad (2)$$

For T_{udp}, $T_{tcpdata}$, and T_{tcpack}, the first term is DIFS plus average backoff, the second is the PHY header, the third is the MAC header and payload transmission time, and the fourth is the SIFS plus bulk ACK overhead. Fig. 5 plots the UDP throughput efficiency (i.e. Throughput/R) as a function of K for various representative values of R. Here we have taken the UDP payload as 1464 bytes (Ethernet MTU minus the link, IP and UDP headers).

Clearly, any level of aggregation is better than no aggregation, and higher levels of aggregation are desirable. However, as we can see from Fig. 5, for our measured level of aggregation, the efficiency is poor especially at the high data rates of 802.11n, in comparison with what is ideal. Hence even though most bytes are in aggregated frames in our measurements, there is significant scope for better channel efficiency through better use of the 802.11n aggregation feature.

Unfortunately, even though frame aggregation is an important part of channel efficiency, algorithms for frame aggregation have not received much attention; this is partly likely due to the fact that the frame aggregation algorithm is typically deep within the hardware.

We have performed controlled experiments, with the aggregation algorithm as a black-box, to understand how the aggregation level differs for various workloads. During a downlink TCP throughput test involving a single client, when it achieves close to the maximum possible throughput, we have observed frame aggregation levels of 15-30 (see Sec. IV). However, for single client HD video streaming (encoded at

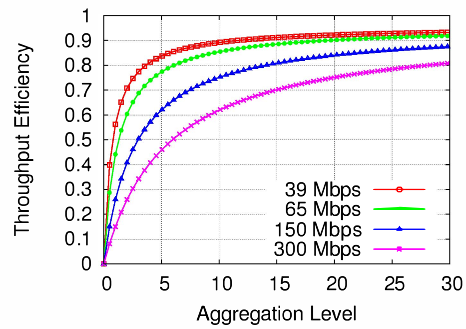

Fig. 5: Throughput efficiency vs aggregation levels

6.7 Mbps), the average aggregation level was just 2-3. As to how these values would change with increasing load (e.g. multiple client downloads) is an interesting open question.

G. Gains due to Frame Aggregation, 802.11n vs 802.11g

A comparison of the raw bit-rates of 802.11g versus 802.11n (54 Mbps vs 600 Mbps) suggests an order of magnitude improvement. To what extent to the various MAC overheads affect this improvement in practical 802.11n deployments? We can quantify this using our traces. We first quantify the gains due to frame aggregation by estimating the hypothetical transmission time of each workload had frame aggregation not been present. This calculation has been done by taking into account the additional header overheads as well as contention overheads (DIFS, backoff, SIFS) that would otherwise have to be incurred for each subframe in the aggregated frame. The contention overhead computation assumes a contention window of $CW_{min} = 16$, and a slot size of $9\mu s$. Table II shows the additional number of transmissions that would have to be done and the additional transmission time if there were no aggregation.

Extending the above hypothetical computation, we can estimate the additional transmission time that would be required for each of the workload if it had been transmitted using 802.11g. This can be an indicator of the benefit obtained, in terms of transmission time, by using 802.11n instead of 802.11g.

In the earlier calculation above, we had assumed that the disaggregated frames would have been transmitted at the same

Trace	Addnl. num. tx w/o frame-aggr. (%)	Addnl. tx time w/o frame-aggr (%)	Addnl. tx time with 802.11g (w/o frame-aggr or 802.11n data rates) (%)
Conference	53	32	42
Airport	19	10	13
Classroom	184	88	140

TABLE II: Estimated Additional Transmissions and Airtime in the Absence of Aggregation

Cases /Clients	Dell Inspiron-15 Laptop (AR9565 card)	Nexus phone (2.4Ghz)	Nexus Tablet	Samsung S3 (2.4Ghz)	macbook pro	Nexus Phone (5Ghz)	Samsung S3 (5Ghz)
400ns with aggr	34.3	45.6	30.2	19.0	79.0	36.4	28.6
400ns without aggr	20.1	23.1	19.1	18.8	28.1	23.6	21.4
802.11g	19.1	19.8	19.3	17.7		19.8	17.7

TABLE III: iperf averaged throughput in Mbps. **Expected throughput:** for a 1x1 client with aggregation is 56 Mbps, without aggregation is 23 Mbps; for a 2x2 client with aggregation is 97 Mbps, without aggregation is 28 Mbps. Expected throughput for 802.11g is 23 Mbps.

rate as the original (aggregated) frame. However, to estimate the transmission time with 802.11g, only 802.11g rates should be taken into consideration for transmission. One such way is to assign a 802.11g rate to a particular 802.11n rate according to their sensitivity values.

Here the methodology followed for determining a corresponding 802.11g rate is as follows: the mapping of 802.11n as well as 802.11g rates with their sensitivity values are taken from those of a commercial AP [6]. For every 802.11n transmission in the trace, the sensitivity value for its data rate is looked up from the mapping. This sensitivity value gives a lower bound on the RSSI at the receiver. The corresponding 802.11g rate is the one having highest sensitivity value smaller than the sensitivity value for 802.11n rate. This is the hypothetical 802.11g data rate that would have been possible at the same SNR. To take an example, suppose our trace has a frame at an 802.11n data rate of 78 Mbps (MCS=12). This rate has a sensitivity of -76dBm. Thus the receiver must have received it at least at -76dBm. If it had been an 802.11g receiver, it would have been able to achieve at least 48Mbps, which has a sensitivity of -78dBm (and not 54Mbps, which requires an RSSI of at least -75dBm).

Considering this method, we obtain the results for additional transmission time as given in Table II. Very high rates, e.g., $>$ 100 Mbps significantly contribute to time savings since transmission time at these rates will be much lower than the highest possible rate in 802.11g (54 Mbps). The other contributing factor is the high aggregation levels. From this we see that the classroom traces would have highest transmission time among the three if all 802.11n frames had been transmitted using 802.11g instead: the data rates used as well as the aggregation levels are the highest in this data set. On the other hand, the airport data set saw mostly low data rates and low levels of aggregation, making the transmission time for 802.11g close to that for 802.11n.

What do the above results imply? While in all three scenarios the benefits of 802.11n are noticeable, the gains due to 802.11n in terms of air-time saved, are well below that suggested by the order of magnitude improvement in raw bit rate. As indicated earlier, a main culprit in this is the low level of frame aggregation. Next we explore controlled experiments which measure the level of frame aggregation under backlogged conditions, for various types of clients.

IV. Frame Aggregation in Controlled Experiments

While a low level of aggregation in the conference and airport settings could be attributed to lack of adequate traffic, in the classroom, several students were downloading large-sized videos. So the low level of aggregation, well below that allowed by 802.11n, was somewhat surprising. To understand this better, we conducted several controlled experiments, using a 2x2 enterprise-grade AP, and different clients. With vendor support, we instrumented the AP as earlier to collect statistics of the A-MPDU aggregation in each transmission. We used the *iperf* tool to run TCP throughput tests in the downlink direction. Each throughput test was run for 1-minute. We fixed the AP to operate at 20MHz (i.e. disabled channel bonding) in these experiments.

Table III shows the measured download throughput of various clients, averaged over three readings. The table also mentions the expected throughput in the caption, computed using Eqn. 2, using $K = 10$ and $R = 72$ (for 1x1 clients) or $R = 144$ (for 2x2 clients). Note that we expect the highest possible data rate as the AP and client were within a foot of one another. The table also shows the measured throughput readings for a scenario where we turned off downlink A-MPDU aggregations (i.e. aggregation level is 1) using an AP interface for the same. It also shows the 802.11g throughput (the AP was turned into a 802.11g-only AP by configuration) with the same clients, for comparison.

We see from Table III that when frame aggregation is enabled, the measured throughput is significantly lower than the expected throughput for all clients except the MacBook Pro. However, with aggregation disabled, as well as for the 802.11g case, the measured throughput more closely matches the expected throughput.

To probe this further, we next look at the data rates used by the AP in the experiments. Fig. 6 shows the CDF of the data rates used for each downlink transmission, for three of the

Cases /Clients	Dell Laptop (AR9565 card)	Nexus phone (2.4Ghz)	Nexus Tablet	Samsung S3 (2.4Ghz)	macbook pro	Nexus Phone (5Ghz)	Samsung S3 (5Ghz)
400ns with aggr	4.9	4.4	8.7	17.9	-	13.4	18.2
400ns without aggr	1.4	1.7	8.3	13.1	-	7.7	22.0
Ping Experiment	0	0	2.1	2.3	-	-	-

TABLE IV: Retransmission percentage values

clients, for the cases with and without aggregation enabled. We see that when aggregation is enabled, there is significant usage of lower (than highest for that client) bit rates, especially for the laptop and Nexus tablet. We observed similar behaviour (lower data rates when frame aggregation enabled) for the other clients as well: these are not shown in the graph for clarity.

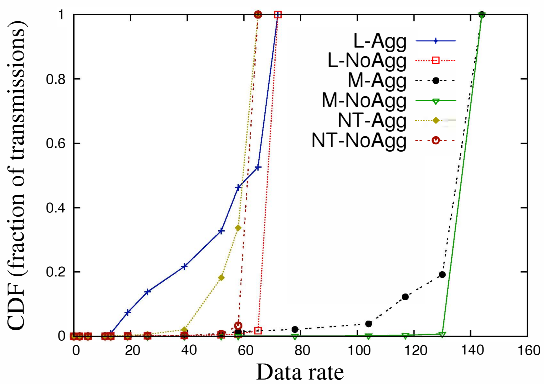

Fig. 6: CDF of data rates for three clients (L-laptop, M-macbook, NT-NexusTablet)

Why does frame aggregation result in usage of lower bit rates? To probe this, we used traces collected by a sniffer (placed next to the AP & client in the earlier experiments), to look for retransmitted MAC sequence numbers. Table IV gives these numbers for the same set of experiments as earlier. (Our sniffer was a 1x1 machine; so the numbers for the 2x2 MacBook Pro experiments are not available). We see that the retransmission percentages are rather high, especially when aggregation is enabled. In the same setting, the retransmission percentages are near zero for a "ping" experiment. Note that for ping packets, there is naturally no opportunity for frame aggregation.

From talking to vendors as well as looking at driver code, a possible explanation for the above observation becomes apparent. The A-MPDU de-aggregation is done on the receiver side in the WiFi hardware, which then passes up to the driver, individual MPDUs using interrupts. So as the level of aggregation increases, the rate of interrupts from the hardware increases too. It is likely that MPDUs are being lost between the client hardware and the driver. This is consistent with our other two observations: negligible retransmission rate in the case of ping, and lower data rates under higher levels of aggregation. The latter is due to the rate adaptation algorithm confusing losses between the hardware and the driver for channel losses.

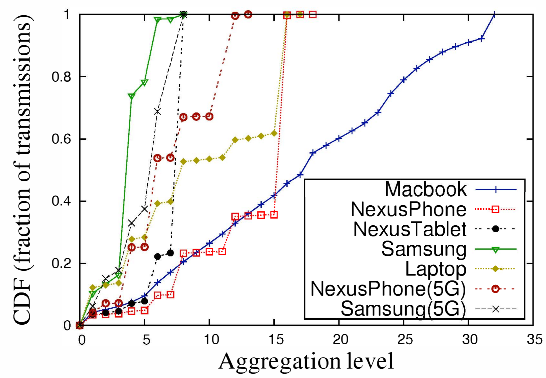

Fig. 7: CDF of frame aggr. for clients in 2.4Ghz band

Fig. 7 plots the CDF of the frame aggregation levels, for a subset of the clients. We see that the level of aggregation has good correlation with the throughput in Table III: some clients are able to sustain a higher level of aggregation and hence a higher download throughput, and others not.

We wish to note that we confirmed the above behaviour with measurements on two other vendors' APs as well, indicating that the problem is not due to the AP. We also observed similar behaviour when we turned off short guard intervals (again using an AP interface for the same). And as shown, we see similar behaviour in the (relatively clean) 5GHz band as well, indicating that the observations are not due to unforeseen interference in the 2.4GHz band.

In summary, we have found in our experiments that several latest clients, including smart-phones and tablets are not really geared toward best use of the important 802.11n feature of frame aggregation. In fact, there is poor interaction between client ability to process hardware interrupts, the rate adaptation algorithm, affecting the overall throughput attained. Unfortunately, any tweaks to the frame aggregation algorithm (e.g. based on feedback from client ability) appears difficult if not impossible, as this algorithm is embedded in the WiFi hardware, with little control at the driver. Future implementations must provide appropriate hooks onto this important feature, to enable exploration of appropriate frame aggregation techniques.

V. CONCLUSION

802.11n has introduced several throughput enhancing features compared to its predecessors 802.11a and 802.11g. The important features are: higher data rates supported by MIMO spatial multiplexing, channel bonding for 40MHz operation, and MAC/link layer frame aggregation. While on paper, or even in controlled experiments, the benefits of these features are clear, there has thus far not been a systematic study of "Are these features actually used in practice?" or "How effectively are the features used?". This paper presents measurement results from traces collected at three dense WiFi deployments: a conference, a busy airport, and a WiFi-enabled classroom. We find that except in the classroom, the high data rates of 802.11n as well as channel bonded mode find rare use. In all traces, the majority of the frames undergo no aggregation, and the median aggregation level is low (5 or lower) even among the frames which are aggregated. This translates to poor channel efficiency.

The implications of our results are manifold. The relatively rare use of the main throughput enhancing features of 802.11n points to other non-wireless bottlenecks. Apart from the obvious wired network bottleneck scenario, we find that often times the client itself can be a bottleneck and impediment to the use of high levels of frame aggregation. The client-side bottleneck introduces non-wireless losses, which interacts poorly with the rate adaptation algorithm, and exacerbates the situation. We do not view our measurement study as the final word; to the contrary, we believe it stresses the need for further careful and extensive real-world measurement studies of 802.11n (of which there has been a dearth) to ensure effective use of the various throughput enhancing features of current and future standards.

VI. Acknowledgements

This work was supported in part by AirTight Networks.

References

[1] "IEEE standard 802.11-2012." [Online]. Available: http://standards.ieee.org/findstds/standard/802.11-2012.html

[2] V. Shrivastava, S. Rayanchu, J. Yoonj, and S. Banerjee, "802.11n under the microscope," in *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, ser. IMC '08, New York, NY, USA, 2008, pp. 105–110. [Online]. Available: http://doi.acm.org/10.1145/1452520.1452533

[3] K. Pelechrinis, T. Salonidis, H. Lundgren, and N. Vaidya, "Experimental characterization of 802.11n link quality at high rates," in *Proceedings of the fifth ACM international workshop on Wireless network testbeds, experimental evaluation and characterization*, ser. WiNTECH '10, New York, NY, USA, 2010, pp. 39–46. [Online]. Available: http://doi.acm.org/10.1145/1860079.1860087

[4] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan, "Measurement-based characterization of 802.11 in a hotspot setting," in *Proceedings of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, ser. E-WIND '05, New York, NY, USA, 2005, pp. 5–10. [Online]. Available: http://doi.acm.org/10.1145/1080148.1080150

[5] M. O. Khan, V. Dave, Y.-C. Chen, O. Jensen, L. Qiu, A. Bhartia, and S. Rallapalli, "Model-driven energy-aware rate adaptation," in *Proceedings of the Fourteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, ser. MobiHoc '13. New York, NY, USA: ACM, 2013, pp. 217–226. [Online]. Available: http://doi.acm.org/10.1145/2491288.2491300

[6] "Cisco Aironet 1530i Access Point receive sensitivity values." [Online]. Available: http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-1530-series/data_sheet_c78-728356.pdf