Analyzing 802.11n Performance Gains

Konstantinos Pelechrinis*, Theodoros Salonidis†, Henrik Lundgren†, Nitin Vaidya‡
*Dept. of Computer Science and Engineering, UC Riverside
kpele@cs.ucr.edu
†Thomson, Paris, France
{theodoros.salonidis, henrik.lundgren}@thomson.net
‡Dept. of Electrical and Computer Engineering, UI Urbana-Champaign
nhv@uiuc.edu

ABSTRACT

We examine the *peak* performance of an 802.11n system with respect to the achievable throughput and we decompose the gains observed due to the various PHY/MAC layer features offered from 802.11n. Our results show that while the MIMO PHY and channel bonding offered from 802.11n increase the PHY data rate, packet aggregation and block acknowledgment are required to achieve throughput closer to the latter.

1. INTRODUCTION

802.11n draft is currently being finalized with the objective to create a high throughput 802.11-based standard. For this reason, 802.11n offers several enhancements over legacy WiFi systems. In particular, **(a)** it utilizes MIMO PHY layer, **(b)** offers possibility for using 40 MHz channel (channel bonding) and **(c)** allows for the use of packet aggregation and block acknowledgment (PABA). All these features aim at improving the higher layer system's performance (i.e., achievable throughput).

In particular, the MIMO PHY layer utilizes multiple antenna elements that can be combined to achieve either higher data rates (Spatial Division Multiplexing - SDM - mode) or higher range (Space Time Block Coding - STBC - mode). Additionaly, using, 40MHz channels, increases the maximum possible data rate as compared to the conventional 20 MHz bands of legacy 802.11 systems. Finally, PABA increases the number of packets transmitted with a single medium access, reducing at the same time the number of ACK packets required to acknowledge a specific amount of DATA packets.

In this study, we first seek to experimentally examine the maximum achievable performance of 802.11n. Previous efforts on experimental characterization of 802.11n have reported throughput values that are far from the maximum data rate viable (i.e. 300 Mbits/sec for a 2x3 802.11n system). Thus, we seek to examine in more detail the actual *peak* performance of 802.11n, for both SDM and STBC modes of operation. This will serve as a benchmark for further studies. To accomplish our goal, we use commodity cards that can operate at the 5 GHz band. This is important, since this specific band is currently not widely used for WLAN depoyments, giving us the capability to perform *controllable* experiments to a large extent. We proceed by isolating and decomposing the gains observed due to the different 802.11n factors. We would like to emphasize on the fact that, as our measurements show, even on a lossless isolated link, 802.11n protocol overheads are significant. Thus, the gains possible with the MIMO PHY layer technology cannot be completely realized in current systems.

Understanding the interplay between the features that constitute 802.11n's gains and the achievable performance is important towards the design of a dynamic system, which adaptively enables/disables the various protocol characteristics. In particular, our initial study shows that despite the fact that MIMO PHY and channel bonding increase the data rate and consequently the throughput, in order for the end user to enjoy a performance close to 200 Mbits/sec, the protocol overheads need to be reduced[1]. The latter is achieved by using PABA.

The rest of this extended abstract is organized as follows. In Section 2 we present related studies. In Section 3 we provide a description of our experimental setup. Section 4 presents and discusses our experimental results. Section 5 concludes and provides our future directions.

2. RELATED STUDIES

Recently, there are a few experimental studies that try to characterize the performance of 802.11n links, providing an initial characterization of the real world performance of 802.11n products. In particular, Shrivastava *et al* [1] were the first to experiment with 802.11n. They are mainly focused on multi-user scenarios and they show that the CSMA/CA policy employed from the draft, can significantly hurt the performance under the presence of multiple active links. Moreover, using a previously introduced interference model [2], they try to capture the interactions between channel bonding and interference. In addition, they also perform experiments on isolated links on the 2.4 GHz band, with various combinations of 802.11n features. In the same direction, the authors in [3] compare the empirical performance between an 802.11g system and an 802.11n system. Finally, in [4] guidelines for deploying a testbed able to realize the MIMO benefits are being provided by the authors. However, the (isolated link) results reported in these studies, are far below the maximum possible data rate of 802.11n. In contrast, with these works, we are mainly interested in achieving the maximum performance of 802.11n, and we consider both modes of MIMO operations, that is, SDM and STBC.

3. EXPERIMENTAL SETUP

[1]These overheads with 802.11n exist, due to the requirement for backward compatibility with 802.11 legacy devices.

Our experimental set up consists of 2 nodes. Each node is a mini-ITX form factor PC, with 1.7GHz Intel CPU, 512MB RAM and 80GB HDD, running Ubuntu 8.04.1 with kernel version 2.6.24-21. Each device is equipped with one Ralink *RT2880* card that supports MIMO-based (802.11n) communications and carries three 5-dBi omnidirectional antennas mounted on the RT2880 card. The Linux driver of the *RT2880* can support both STBC and SDM modes and can operate at the 5 GHz band.

In our experiments we disable the rate adaptation algorithm. We create a link with zero packet losses at the maximum supported data rate and we measure the achievable MAC layer throughput for various combinations of the 802.11n features. Given that there are no losses, this is also the maximum possible throughput under the specific configurations. The maximum data rate for the 4 combinations of channel width and MIMO modes can be found in Table 1 (the short Guard Interval option is used). We use the `iperf` traffic generator to backlog the link with 1470 bytes UDP packets.

	20MHz	**40MHz**
SDM	144	300
STBC	72	150

Table 1: Maximum data rates (Mbits/sec) for various combinations of channel width and MIMO modes.

Eliminating losses on the link was not a trivial task. In particular, the system appears to be very sensitive to a number of factors. The spacing l, between the antennas placed on a linear topology as well as the length of the link D, affect significantly the packet delivery ratio. Even small variations at the values of l and D can dramatically change the link quality. In general, the antenna spacing needs to be greater then half wavelength λ in order for the antenna signals to be uncorrelated. In addition, the wireless environment, and in particular the richness in scattering, is another important factor to consider. Finally, we initially experimented at the 2.4 GHz band. However, eliminating interference in these frequencies is hard, due to its wide usage for WLAN deployments.

For our lossless link we set $D = 1.3$ meters, while the antennas are arranged on a linear topology with a spacing of $l = 7.8$ cm between them. Note that since we operate on channel 60 (5.3 GHz), this antenna spacing is greater $\frac{\lambda}{2} = 2.8$ cm. Nodes are placed inside a meeting room, rich in scattering. Furthermore, we have verified that there is no other WLAN operating on the same and the neighboring frequencies, and thus, our experiments are interference free.

4. 802.11n BENCHMARK PERFORMANCE

In this section we present our experimental results towards the benchmark performance[2] of 802.11n.

Our experimental procedure starts by activating 802.11a and measuring the achievable throughput; following we add step by step additional features of 802.11n and observe the difference in the performance. A pictorial summary of our results is given in Figure 1.

[2]This is measured using a specific wireless card, but we believe that it is representative of the 802.11n performance.

802.11a performance: With 802.11a the maximum data rate supported is 54 Mbps. As we can observe from Figure 1, with 802.11a we are able to reach 37 Mbps on the lossless link created in our lab.

MIMO PHY: The first step towards examining the performance of 802.11n involves employing only the MIMO PHY layer offered from the new draft standard. Thus, we enable 802.11n on our wireless NICs, while keeping channel bonding and PABA disabled. Using the corresponding maximum data rates (shown at Table 1), we are able to achieve throughput performance of 61.9 Mbits/sec and 42.3 Mbits/sec with SDM and STBC mode respectively, as we can see at box (2) of Figure 1. Even though there is a clear improvement as compared with 802.11a (especially for the SDM case), the achieved throughput is far from the PHY rate. Examing the reasons behind this, we define the **Effective Throughput Ratio (ETR)** to be the ratio of the measured throughput T_m over the PHY data rate R. Therefore, we have:

$$ETR_{SISO} = 0.69 > ETR^{STBC}_{MIMO} = 0.58 > ETR^{SDM}_{MIMO} = 0.43 \quad (1)$$

Only 58% of the PHY layer data rate is obtained from the 802.11n when STBC is being employed. SDM achieves much higher throughput, however, it is just 43% of the maximum possible. The reason for this performance can be found in the 802.11n MAC layer overheads. In particular, for the case of lossless link, ETR can be expressed as:

$$ETR = \frac{T_m}{R} = \frac{T^{DATA}_{air}}{T^{DATA}_{air} + T_{overhead}} = \frac{\frac{L}{R}}{\frac{L}{R} + T_{overhead}} \quad (2)$$

where T^{DATA}_{air} is the air time of the DATA packet (MAC layer payload of size L) and $T_{overhead}$ captures the various 802.11n protocol overheads adopted from legacy 802.11 for backward compatibility (timings related with DIFS, SIFS, ACK, contention, etc.). We have experimentally validated that, $T_{overhead}$ is constant for the various modes of operation (SDM, STBC) and channel widths, and depends only on whether PABA is used or not. From Equation (2), we can see that ETR is a decreasing function with respect to the data rate R, for the same packet size L and $T_{overhead}$; thus, $T_{overhead}$ *introduced from 802.11n is more profound for higher data rates*, verifying the trend observed in Inequalities (1). In contrast with legacy 802.11 that supports much lower data rates, the air-time of the packet for the high rates supported from 802.11n is comparable to the protocol overhead; therefore, the latter becomes the dominant factor that limits performance at the 802.11n high data rates.

MIMO PHY + Channel Bonding: Channel bonding combines two or more adjacent channels in order to form a wider one, to support higher data rates. 802.11n offers the possibility of using channels of 40MHz width. Adding this feature to our previous configuration increases the maximum data rates supported by 108%. However, the throughput achieved is only increased by 27% for the case of SDM (79.1 Mbits/sec) and 46% for the case of STBC (61.9 Mbits/sec). ETR is further reduced to $ETR^{STBC}_{MIMO} = 0.41$ (30% reduction) and $ETR^{SDM}_{MIMO} = 0.26$ (40% reduction).

The reduction of ETR is much larger for the SDM case as

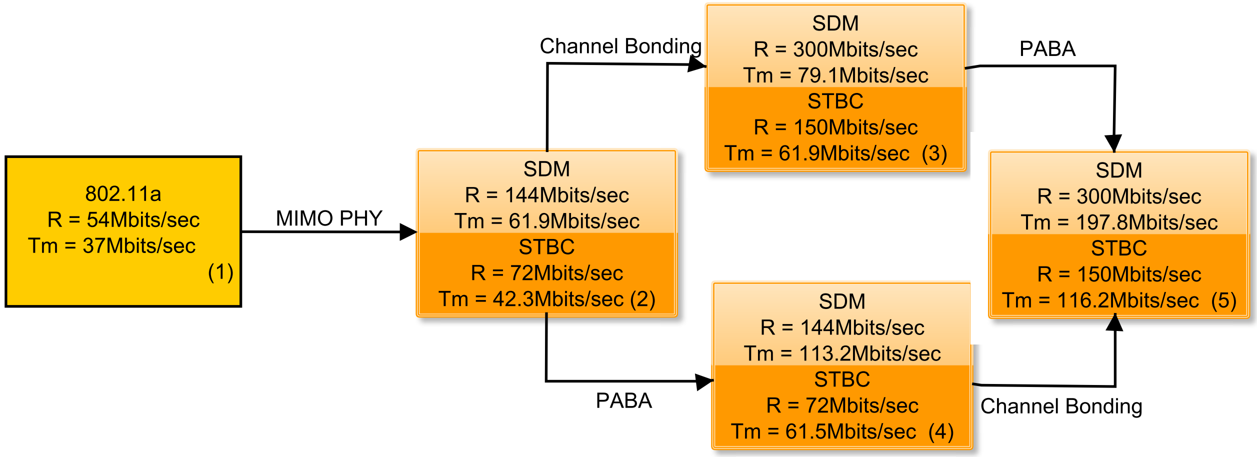

Figure 1: Break down of the 802.11n performance.

compared with the STBC. This should be expected based on the shape of the ETR function given from Equation 2. In particular, one can analytically verify that as we increase the data rate ($R^{40MHz}_{SDM} = 300Mbits/sec > R^{40MHz}_{STBC} = 150Mbits/sec$), we should expect a larger decrease in ETR.

With channel bonding we can improve the performance seen by the end user. However, we are still far from the PHY layer capabilities in terms of bandwidth; being more precise, we are even further - in relative terms - as compared to the case where 20 MHz channel is used. This is because employing channel bonding does not affect $T_{overhead}$ but only increases R. Thus, recalling again equation 2, the effect of $T_{overhead}$ is even more intense when channel bonding is being employed.

MIMO PHY + PABA: As seen above, the overhead introduced from 802.11n limits the peak throughput T_m. One way to boost the performance is by reducing $T_{overhead}$ for a fixed data rate. Then, ETR will increase and thus T_m will also increase.

PABA achieves this by amortizing the protocol overheads over a sequence of packets, without increasing the data rate at the PHY layer. Indeed, after enabling PABA on our cards, we measure a throughput of 61.5 Mbits/sec with STBC and 113.2 Mbits/sec with SDM (box (4) - Figure 1). The overhead is significantly reduced, since the above values correspond to $ETR^{STBC}_{MIMO} = 0.85$ and $ETR^{SDM}_{MIMO} = 0.78$.

MIMO PHY + Channel Bonding + PABA: In the previous scenario we observed that PABA reduces the 802.11n protocol overhead, and boost the achievable throughput. However, the limiting factor for the above configuration is the maximum possible data rate (just 72 Mbits/sec with STBC and 144 Mbits/sec with SDM). As a consequence, next we add channel bonding in the above configuration to increase the throughput even more.

The throughput observed is even higher and in particular, 197.8 Mbits/sec are obtained with SDM and 116.2 Mbits/sec with STBC mode. The ETR is reduced as compared with the case where only MIMO PHY and PABA is used - as one can also expect from Equation (2). However in absolute terms, we are able to obtain slightly more than 5 times the maximum throughput achieved with legacy 802.11a.

5. CONCLUSIONS AND FUTURE WORK

Our results show that the overhead introduced to 802.11n by the 802.11 MAC protocol family is significant even on an isolated lossless link. These are more pronounced in the higher transmission rates supported from the draft, as compared to legacy 802.11. MIMO PHY layer and channel bonding offered from 802.11n, can increase the data rate. However, in order to be able to achieve throughput *close* to the actual PHY layer capabilities (in the order of 200 Mbits/sec) we need to use enhancements at the MAC layer that reduce the protocol overhead (e.g. PABA).

In this paper, we mainly followed an experimental approach focusing on the peak performance of 802.11n. These results will serve as a benchmark for our future studies. In particular, we seek to examine in more detail the performance sensitivities mentioned in Section 3. Taking care of these sensitivities is important in many realistic scenarios. For example, when mobility is considered, parameters such as D, the exact environmental characteristics or the relative antenna positioning can change rapidly. Thus, we seek to examine in more detail the above effects and also extend our study on lossy links. Our final goal, is the design of a dynamic system able to adapt various 802.11n parameters in such a way that will utilize as much of MIMO PHY as possible.

Acknowledgments

We would like to thank Ralink Technologies for providing us the source code and support for the Linux driver for the RT2880 AP/client.

6. REFERENCES

[1] V. Shrivastava, S. Rayanch, J. Yoon, and S. Banerjee. 802.11n Under the Microscope. In *IMC*, October 2008.

[2] A. Mishra, V. Shrivastava, S. Banerjee, and W. Arbaugh. Partially overlapped channels not considered harmful. In *ACM SIGMETRICS*, 2006.

[3] V. Visoottiviseth, T. Piroonsith, and S. Siwamogsatham. An Empirical Study on Achievable Throughputs of IEEE 802.11n Devices. In *WinMee*, June 2009.

[4] K. Pelechrinis, I. Broustis, T. Salonidis, S.V. Krishnamurthy, and P. Mohapatra. Design and Deployment Considerations for High Performance MIMO Testbeds. In *WICON*, November 2008.