Experimental Evaluation and Comparative Study on Energy Efficiency of the Evolving IEEE 802.11 Standards

Stratos Keranidis[⊕], Giannis Kazdaridis[⊕], Nikos Makris[⊕],
Thanasis Korakis[⊕], Iordanis Koutsopoulos[⊎] and Leandros Tassiulas[⊕]
[○]Department of Electrical and Computer Engineering, University of Thessaly, Greece
[+]Centre for Research and Technology Hellas, CERTH, Greece
[∪]Department of Computer Science, Athens University of Economics and Business, Greece
{efkerani, iokazdarid, nimakris, korakis, leandros}@uth.gr, jordan@aueb.gr

ABSTRACT

Over the last decade, the IEEE 802.11 has emerged as the most popular protocol in the wireless domain. Since the release of the first standard version, several amendments have been introduced in an effort to improve its throughput performance, with the most recent one being the IEEE 802.11n extension. In this paper, we present experimentally obtained results that evaluate the energy efficiency of the base standard in comparison with the latest 802.11n version, under a wide range of settings. To the best of our knowledge, our work is the first to provide such a detailed comparative analysis on the performance of both standards. The followed power measurement methodology is based on custom-built hardware that enables online energy consumption evaluation at both the wireless transceiver and the total node levels. Based on in-depth interpretation of the collected results, we remark that the latest standard enables significant improvement of energy efficiency, when combined with standard compliant frame aggregation mechanisms. Our detailed findings can act as guidelines for researchers working on the design of energy efficient wireless protocols.

1. INTRODUCTION

IEEE 802.11 is currently considered as the default solution for implementing wireless local area network communications. The wide adoption of this standard by vendors of wireless devices offers high interoperability, which in combination with the provided ease of use and low deployment cost have resulted in its unprecedented market and everyday life penetration. While the base version of the standard was released in 1997, subsequent amendments have been proposed throughout the years, such as the widely adopted 802.11b, 802.11a and 802.11g. In 2007, the current standard IEEE 802.11-2007 [1] was released and merged several amendments with the base version.

In an effort to improve throughput performance of the base standard, the IEEE 802.11 standard working group started in 2003 to develop the IEEE 802.11n high-throughput (HT) extension of the base standard that was finally published in 2009. The most important improvement of the 802.11n on the Physical layer (PHY) is the ability to combine multiple antenna elements to achieve higher PHY bit rates and increased link reliability. In order to increase medium utilisation and exploit from the increased PHY bit rates, efficient Medium Access Control layer (MAC) frame aggregation mechanisms [2] are also supported. Several recent studies [3, 4] in the field of wireless networking have experimentally verified the improved channel efficiency and throughput performance that the 802.11n is able to deliver in comparison with 802.11a/g systems.

On the other hand, the recent penetration of 802.11n compatible chipsets in "smart" mobile devices has raised concerns regarding the energy efficient operation of the HT protocol. As the dramatically increased PHY bit rates, require the activation of multiple RF chains and complex baseband processing as well, 802.11n compatible chipsets induce significantly higher power consumption that grows with the number of active RF chains. Our experimental results that match other relevant studies [5, 6] as well, show that a modern 802.11n 3x3 MIMO compatible chipset is able to draw up to 2.45 W, in the case that all the various chipset components are constantly activated and the highest PHY bit rate performance is achieved. Especially in the case of smartphone platforms, the energy greedy profile of the supported state-of-the-art wireless technologies may induce up to 50% of the total platform power consumption [7], under typical use case scenarios. As the focus of researchers is usually on network performance, only a few works have presented detailed experimental results that characterise the energy consumption of 802.11n chipsets [5, 6, 8]. The work in [5] is restricted in characterising the power consumption profile of commercial 802.11n chipsets, while the rest two studies [6, 8] experimentally investigate the impact of standard compliant power saving mechanisms on the operation of 802.11n protocol.

In this work, we take a step further than relevant studies and characterise the energy efficiency of 802.11 compliant protocols, in comparison with the achievable network performance they are able to offer. Our study is not restricted in evaluating the performance of 802.11n chipsets, but presents detailed experimentally obtained measurements that compare the performance of the legacy 802.11a/g standard with the latest 802.11n version. The impact of various MAC layer enhancements, both vendor specific and standard compliant

e-Energy'14, June 11–13, 2014, Cambridge, UK.

http://dx.doi.org/10.1145/2602044.2602069.

ones, is also considered in the performance evaluation of both protocols. The obtained results are collected in realistic scenarios and under a wide range of settings, considering varying application-layer traffic loads, frame payload lengths and different network topologies that offer varying channel conditions as well. Finally the impact of the default 802.11 Power Saving Mechanism (PSM) on the performance of both protocols during periods of activity has also been investigated.

Accurate energy efficiency evaluation under real world scale and settings, is a rather complex task that requires the application of detailed evaluation methodologies, in combination with advanced power monitoring platforms. The followed power measurement methodology is based on custom-built hardware that was developed in our previous work [9] and enables online energy consumption evaluation at both the Network Interface Card (NIC) and the total wireless node levels. In-depth analysis of the extensive list of obtained measurements aided in identifying factors that affect energy consumption on commodity 802.11 hardware. Our detailed findings that can act as guidelines towards designing energy efficient wireless protocols, are summarised as follows:

- Activation of additional RF chains that enable MIMO communications results in remarkably increased power consumption (up to 2.5x) at the NIC level. However, our experiments have shown that the resulting increased PHY bit rates of 802.11n are able to increase energy efficiency at the NIC level by 33% during transmission (63% in reception), in comparison with the rates of 802.11a/g. Moreover, we observed that proper activation of the required number of RF-chains for each specific rate configuration can aid towards saving energy.
- Application of MAC-layer aggregation mechanisms is able to deliver substantially increased throughput, while also resulting in considerable energy savings. Power consumption experiments that consider consumption at the total node level have shown that the aggregation assisted 802.11n can improve energy efficiency by more than 80%, in comparison with the performance of the 802.11a/g standard.
- While transmitting MAC frames of low payload length and under high-SNR conditions, 802.11n increased energy efficiency at the node level by 90%, in comparison with 802.11a/g. This observation is related to the fact that the supported by 802.11n aggregation mechanisms enable delivery of high throughput performance (>100 Mbps), even when transmitting frames of 300 bytes payload.
- Considering low-SNR conditions, the Spatial Diversity mode of 802.11n offers increased MAC-layer Frame Delivery Rate (FDR) and throughput improvement by a factor of 4.6x, as observed in our experiments. In addition, we remark that the monitored throughput improvement did not induce significant energy costs, as the energy efficiency at the NIC level also increased by 58%.
- Experimentation with the default 802.11 PSM mechanism has shown that 802.11n is able to provide significant energy savings (> 75%), across varying traffic loads and without sacrificing application-layer throughput or jitter performance.

This paper is organized as follows. In Section 2 we present the evolution of the base 802.11 standard over the last years. Section 3 details the experimental setup and the followed power measurement methodology that is used in our experimental evaluation. In Section 4, we characterise the power consumption profiles of the two wireless chipsets that are used in our experiments, while in Section 5 we present extensive experiments that compare the performance of the 802.11a/g and 802.11n protocols in terms of network performance and energy consumption. Finally, in Section 6 we point out the conclusions reached through this work.

2. EVOLUTION OF IEEE 802.11

The aforementioned versions of IEEE 802.11 use different PHY layer specifications, but all follow the MAC architecture of the base protocol. The mandatory access scheme that has been specified by the legacy IEEE 802.11 standard is implemented through the distributed coordination function (DCF), which is based on the carrier sense multiple access with collision avoidance (CSMA/CA) mechanism. The large PHY and MAC layer overheads that are associated with the DCF process, result in a reduction of more than 50% of the nominal link capacity, which effect is more pronounced for higher PHY bit rates, as shown in [10]. The work in [11] has analyzed the throughput and delay limits of the IEEE 802.11 standard and has shown that for infinitely high PHY bit rate and a frame payload size of 1024 bytes, the maximum achievable throughput is upper bounded to 50.2 Mbps. Such observations highlighted that MAC layer enhancements need to be applied, in order to reduce the impact of the PHY and MAC layer overheads of the base standard.

In an effort to improve throughput performance, vendors of wireless products started integrating innovative techniques into their products, as early as 2003. Such techniques include the "Atheros Fast Frames" (*FF*) [12], which improves 802.11a/b/g performance, by combining two MAC frames into the payload of a single aggregated frame. However, application of vendor-specific techniques has been reported to result in hardware incompatibilities, or at least degraded performance for standard compliant devices, as presented in [13].

Along the same direction, the IEEE 802.11 standard working group introduced the 802.11n extension that offers both PHY and MAC layer enhancements over legacy 802.11 systems. Through the combination of multiple antenna elements and complex MIMO processing, 802.11n is able to achieve higher PHY bit rates (in Spatial Multiplexing mode) and increased link reliability through the exploitation of multipath transmissions and antenna diversity (in Spatial Diversity mode) [14]. Another significant feature is the application of channel bonding, which increases the channel bandwidth from 20 MHz to 40 MHz and thus doubles the theoretical capacity limits. Moreover, the available Modulation and Coding Schemes (MCS) were extended through the introduction of the new coding rate of 5/6, as well as through the decrease of the OFDM guard interval from 0.8 μs to 0.4 μs. Finally, the number of OFDM data subcarriers was increased from 48 to 52, towards improving spectral efficiency. Application of the aforementioned enhancements is

able to deliver the remarkably increased PHY bit rate of 600 Mbps (when using 4 antennas), resulting in performance improvement of more than 10x compared to legacy 802.11a/g systems.

In order to increase medium utilisation and exploit from the increased PHY bit rates, two different types of frame aggregation are provided, namely A-MSDU and A-MPDU aggregation. The former combines multiple higher layer packets into a single MAC layer frame with maximum size of 7935 bytes, while the latter combines multiple MAC layer frames to form an aggregated frame that cannot exceed the 65.536 bytes. In general, A-MPDU aggregation outperforms A-MSDU, which technique results in considerably degraded performance under low quality channel conditions and high PHY bit rates, as it was shown in [15]. Both frame aggregation mechanisms are enhanced by a block acknowledgment mechanism, which further reduces protocol overhead.

3. MEASUREMENT SETUP

In this section, we present the measurement setup that is used in the detailed experimental evaluation that aims at deriving a comparative performance analysis between the 802.11a/g and 802.11n standards. We start by describing the exact experimental setup, which is based on commercial wireless NICs that are representative of the state-of-the-art of each standard. In addition, we detail the followed power measurement procedure and the underlying hardware that are used to characterise the energy consumption performance of the considered protocols. The experimental setup that is used as the basis of our evaluation, consists of a single communicating pair of nodes that both feature the specifications listed in Table 1.

Component	Type
Motherboard	Commell LE-575X
CPU	Intel Atom D525 (1.8 GHz)
RAM	Kingston HYPERX DDR3 - 4GBs
Hard Drive	Samsung SSD - 64 GBs
Power Supply	60W - 12V
OS	Ubuntu 13.04
Wireless cards	Atheros 9380 / 5424
Wireless Drivers	madwifi-0.9.4 / backportsv3.12.1

Table 1: Node Specifications

Wireless communication is enabled through the Atheros AR5424 and AR9380 chipsets that implement the 802.11a/g and 802.11n protocols and are configured through the use of the *Mad-WiFi* [16] and *ath9k* [17] open source drivers accordingly. The wireless nodes are closely located in an indoor office environment at the University of Thessaly premises and are configured to operate in infrastructure mode, on the RF-isolated channel 36 of the 5 GHz band, in order to constantly guarantee un-interfered communication. We setup two different topologies, by keeping the AP node at the same physical location, while we move the STA between the two locations that are depicted in Fig. 1.

We use the Iperf [18] tool to generate traffic and collect network performance statistics. A typical experimental setup for experiments considering downlink transmissions, would be to run an Iperf client at the AP node, having also an Iperf server residing at the STA, receiving the traffic and collecting statistics. Moreover, we also exploit from the statistics that the applied Rate Adaptation algorithm of each wireless driver is able to export, in order to measure link reliability in terms of MAC-layer FDR, as calculated per each configured PHY-layer rate.

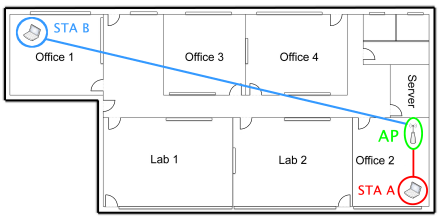

Figure 1: Experimental Topology

3.1 Experimental Setup

We first place the STA at location A to establish a line-of-sight high-SNR link (Signal to Noise Ratio (SNR) $\sim$ 35 dB), by configuring the transmission power of both nodes at the maximum level of 20 dBm. Towards executing experiments under low-SNR (SNR $\sim$ 15 dB) channel conditions, we move the STA to location B and reduce the transmission power of both nodes to the minimum available level of 0 dBm. The aforementioned SNR values correspond to the 802.11a/g link, while the 3x3 MIMO configuration of the 802.11n link provides approximately 5dB gain in each setup, by exploiting spatial diversity at the receiver through the Maximal-ratio Combining (MRC) technique [19].

Towards providing for a proper comparison setup between the two standards, we configure both transceivers to use the same channel bandwidth of 20 MHz and OFDM guard interval of 0.8 μs. Under this setup, we execute each discrete experiment in two phases, where in each phase either the 802.11a/g or the 802.11n protocols are configured through the use of the corresponding transceivers. As the AR9380 chipset is also able to operate in the 802.11a/g mode, we measure its performance under this configuration, across the various considered cases as well. Under these settings, the 802.11n compatible chipset supports the maximum PHY bit rate values of 65 Mbps, 130 Mbps and 195 Mbps for single, double and triple spatial stream configurations accordingly, while the 802.11a/g compatible chipsets support PHY bit rate values between 6 Mbps and 54 Mbps.

3.2 Power Measurement Methodology

In order to accurately measure the instantaneous power consumption, we follow a widely adopted power measurement procedure, which requires the placement of a high-precision, low impedance current-shunt resistor (R) of a known resistance value, in series with the power source and the power supply pin of the device to be measured. The exact measurement setup described above is presented in Fig. 2(a). By consistently measuring the voltage ($V_R(t)$) across the current-shunt resistor through proper voltage metering equipment, we are able to extract the instantaneous current draw of the device, based on Ohm's law. The instantaneous power consumption can be calculated as the product of the input voltage V_{IN} and the measured current draw:

$$P(t) = V_{IN} \frac{V_R(t)}{R} \quad (1)$$

Estimation of the total energy consumption during a specific experiment, necessitates the accurate sampling of the instantaneous power consumption during the total experiment duration. Total energy consumption can be calculated

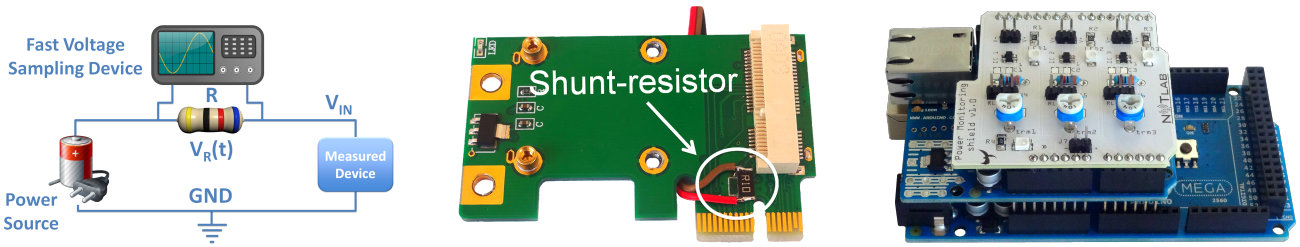

(a) Power Measurement setup (b) Modified mini PCI-e adapter (c) NITOS ACM card

Figure 2: NITOS ACM card and the accompanying hardware and software components

as the integral of the power consumption over the specified duration ($Dt = t_1 - t_0$), as follows:

$$E(Dt) = \frac{V_{IN}}{R} \int_{t_0}^{t_1} V_R(t)dt \quad (2)$$

However it should be made clear that through the voltage sampling equipment, only a finite number of samples of $V_R(\cdot)$ are acquired over $[t_0, t_1]$ at discrete time instances.

In our study, we consider power consumption at both the total node level, as well as at the level of the wireless NIC. As a result, we decided to equip both nodes with appropriate current shunt resistors that have been placed in series with the power supply of the NIC and the Atom-based node accordingly. In order to ease the interception of the NIC power supply pins and refrain from modifying each different type of card, we decided to insert the current-shunt resistor on a PCI-e to mini PCI-e adapter card that is compatible with both wireless cards. Fig. 2(b) illustrates the modified adapter card that is attached with a high-precision current-shunt resistor of 0.1 Ω. We also modified the nodes' power supplies by inserting a current-shunt resistor of 0.01Ω in series between the power supply and the node's motherboard. In order to accurately measure the voltage drop across the resistors, we used the NITOS ACM card that was introduced in our previous work [9]. The developed card, which is presented in Fig. 2(c), supports the high sampling rate of 63 KHz and features up to three input channels, thus providing for online power consumption monitoring at both the NIC and the total node level, in a joint way.

4. POWER CONSUMPTION PROFILING

This initial set of experiments has been designed to clearly describe the power consumption profile of each chipset and set the basis for the realistic performance evaluation that follows in Section 5. Based upon the high-SNR setup, we characterize the instantaneous power consumption of the two NICs across various operational modes and present the obtained results in Table 2.

We clearly observe that the later manufactured AR9380 chipset is highly optimised for energy efficiency, as it consumes far less than 50% power than the AR5424 chipset, under all operational modes in the single antenna configuration. While considering MIMO operation, we notice that the activation of additional RF chains remarkably increases power consumption, as several hardware components need to be activated, in order to provide for the complex baseband processing that MIMO communications require. However, it is interesting to note that only in the case that 3x3 MIMO transmissions are executed by the AR9380, its power consumption increases above the consumption levels

Chipset	AR5424	AR9380		
Antennas	1x1	1x1	2x2	3x3
Mode	Power Consumption (Watts)			
Sleep	-	0.12		
Idle	1.47	0.49	0.56	0.69
Receive	1.52	0.62	0.74	0.85
Transmit	1.97	0.98	1.75	2.45

Table 2: Power consumption of AR5424 and AR9380 NICs across different operational modes

of the AR5424 chipset. Considering the power consumption in the sleep mode of operation, we remark that the *MadWiFi* driver does not support the activation of the *Power Saving Mode (PSM)* for the 802.11a/g compatible chipset. On the other hand, the PSM mode can be activated for the AR9380 NIC through the *ath9k* driver and set the card in a low-power state (82% less than in idle mode), by disabling most of the NIC's circuitry.

In order to assess the impact of varying PHY bit rates on NIC energy efficiency, we next proceed by characterising the energy consumption per transmitted bit of information (E_B), under the various PHY bit rate configurations that each protocol supports. We calculate E_B, expressed in Joules/bit, as the division of the resulting power consumption (Joules/sec) under each operational mode, by the specified PHY bit rate value expressed in bits/sec. In Figures 3(a) and 3(b), we plot the obtained E_B across the available IEEE 802.11a/g compatible PHY bit rate configurations, for the AR5424 and the AR9380 chipsets accordingly. In the case that the AR9380 chipset is configured to operate in the IEEE 802.11a/g mode, we manually disable the excess RF chains and use the single mode antenna of operation, thus resulting in the significantly reduced E_B values. Based on the obtained results, we remark that instantaneous power consumption of both chipsets does not significantly vary between different PHY bit rate settings. On the other hand, as plotted in Figures 3(a) and 3(b) , higher PHY bit rate settings always result in lower E_B, mainly due to the decreased duration of the transmission or reception operations.

We next proceed to the characterization of the power consumption profile of the 802.11n compatible AR9380 chipset, which offers a wider range of available MCS configurations, as it features three RF-chains and supports up to 3x3 MIMO mode of operation. In Fig. 4(a), we plot the obtained E_B across the 23 available MCS configurations, in the case that all RF-chains are constantly enabled. We clearly notice that MCS configurations significantly impact power consumption, as imposed by the calculated E_B, which ranges from 47.1 nJ/bit (MCS0) to 1.53 nJ/bit (MCS23). This finding indicates the huge potential of energy expenditure minimi-

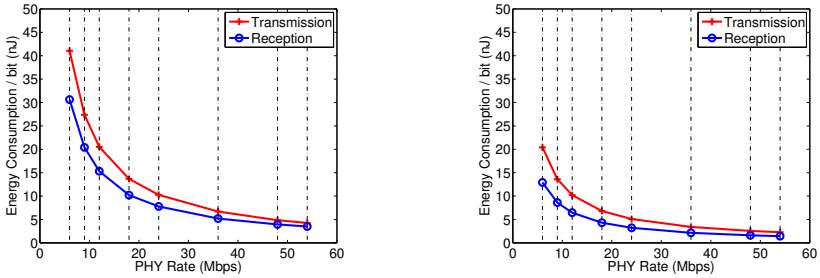

(a) AR5424 (b) AR9380

Figure 3: Energy consumption/bit of 802.11a/g NICs across available PHY bit rate configurations

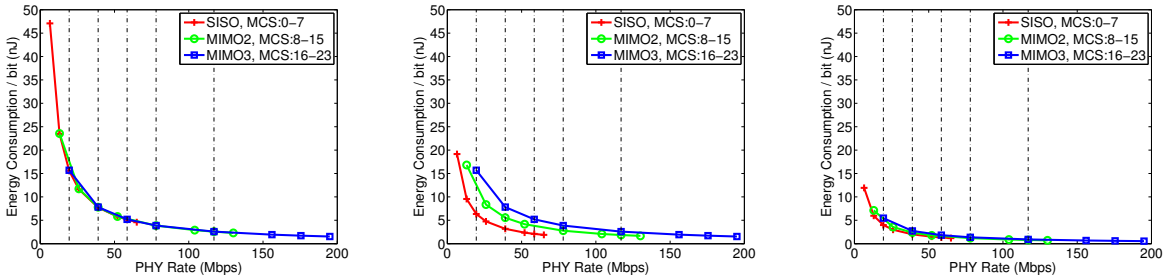

(a) Transmission - All RF chains constantly enabled (b) Transmission - Only required number of RF chains enabled (c) Reception - Only required number of RF chains enabled

Figure 4: Energy consumption/bit of 802.11n NIC across available PHY bit rate configurations

sation (up to 97%), through proper adaptation of MCS configurations, as employed in similar studies [20].

In the case that all RF-chains are continuously enabled and while operating in the SISO and MIMO2 modes, the excess antennas do not contribute to the PHY bit rate increase, but are only used to provide increased link reliability. However, under ideal channel conditions, the excess antennas are no longer required to improve link reliability and thus can be deactivated, towards reducing energy expenditure. Figures 4(b) and 4(c) present E_B measurements for transmission and reception accordingly, in the case that only the required number of RF-chains are enabled for each configured SS setting. We notice that proper activation of the required number of RF-chains (SISO, MIMO2) can significantly increase energy savings up to 60% for transmission (27% for reception), as for the MCS0 case, where E_B reduces to 19.15 nJ/bit (11.92 nJ/bit). Considering the 802.11n configuration of the AR9380 chipset, we observe that its instantaneous power consumption does not significantly vary between different MCS indexes within the same SS configuration. Based on this fact, we infer that the high diversity of E_B values that is plotted in Figures 4(b) and 4(c), is mainly due to the increased power consumption that activation of additional RF-chains results in.

Based on direct comparison of the E_B values plotted in Figures 3(b), 4(b) and 4(c), we aim at quantifying the energy savings that 802.11n can offer in contrast to the earlier 802.11a/g protocol. Considering the power consumption performance of the AR9380 chipset at the highest configurable PHY rates of each protocol (54 Mbps for 802.11a/g and 195 Mbps for 802.11n), we observe that energy savings of up to 33% can be attained during transmission (63% during reception), as E_B reduces from 2.27 nJ/bit to 1.53 nJ/bit (1.44 nJ/bit to 0.54 nJ/bit in reception). Concluding, we note that the remarkably higher rates of 802.11n protocol can offer significant energy savings when combined with proper adaptation of antenna modes, in comparison with the energy consumption performance of the earlier 802.11a/g protocol. However, the power consumption profile characterisation is only based on static PHY bit rate configurations and does not consider protocol overheads, MAC-layer parameters, such as aggregation, or link performance and their impact on energy efficiency. In the following section, we proceed by comparatively evaluating the performance of the two standards under realistic throughput experiments, considering various protocol parameters as well.

5. REALISTIC EXPERIMENTATION

In this section, we take a step further from characterising energy consumption under fixed modes of operation and conduct extensive realistic experiments to compare the performance of the two standards, by jointly considering application-layer performance and energy efficiency as well. We start by configuring the high-SNR setup to investigate the impact of varying application-layer traffic rate and frame payload length on the performance of each protocol. The obtained results are analyzed in Sections 5.1 and 5.2, for each scenario accordingly. In both scenarios, we manually configure the maximum available PHY bit rates of each standard and guarantee that these rates can be supported by the prevailing channel conditions, by constantly monitoring the achievable FDR and assuring that it never drops below 95% in all the conducted experiments.

Under these conditions, we start by measuring the throughput performance of each protocol without enabling any form of aggregation, while we next repeat the same experiments by explicitly enabling the *FF* and A-MPDU aggregation mechanisms. In each experiment, we also monitor the power consumption at both the NIC and the total node level, in order to assess the impact of the various configured settings on energy consumption. Next, in Section 5.2.2, we configure the low-SNR experimental setup to conduct series of

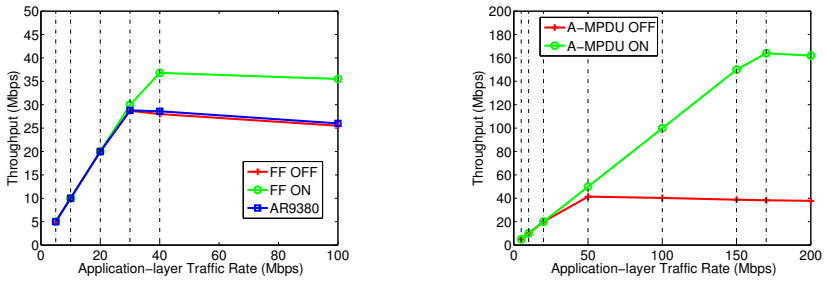

(a) AR5424 (b) AR9380

Figure 5: Throughput performance per NIC across varying Application-Layer Traffic load

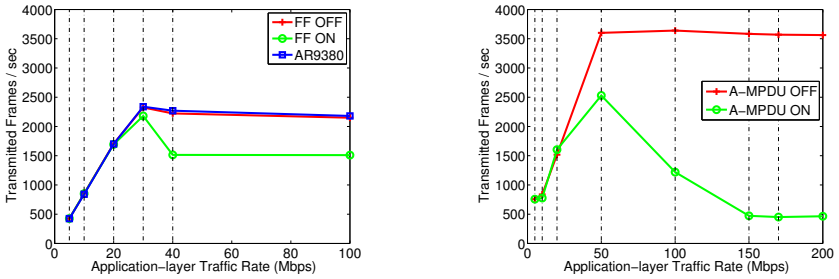

(a) AR5424 (b) AR9380

Figure 6: MAC-layer Frame Transmission rate per NIC across varying Application-Layer Traffic load

experiments under varying frame payload size and investigate the impact on performance of both protocols, under low link quality conditions. Finally, in section 5.3, we assess the potential energy savings that the application of the PSM mechanism is able offer, under varying application-layer traffic rates.

5.1 Varying Application-Layer Traffic load

We start by measuring the throughput performance under perfect channel conditions and across varying application layer traffic loads. The obtained results for the 802.11a/g and 802.11n protocols are illustrated in Figures 5(a) and 5(b) accordingly. In Fig. 5(a), we observe that below channel saturation, throughput performance is similar for both chipsets, while we also notice that the *FF* mechanism does not induce any impact. On the other hand and as soon as the load approaches 40 Mbps, we observe that the application of the *FF* mechanism offers approximately 31.4% increase in the maximum achievable throughput from 28 Mbps (*FF* disabled) to 36.8 Mbps (*FF* enabled). Based on detailed study of the *Mad-WiFi* driver code, we concluded that *FF* is only activated when the driver detects that the channel is approaching saturation through inspection of the transmission queue levels. We also verified our findings by monitoring the number of MAC-layer frames that are being transmitted in each time instant. In Fig. 6(a), we plot the collected results and observe that in the 40 Mbps load case, the frame transmission rate decreases from 2150 to 1510 frames/sec (30% decrease).

Based on the 802.11n compatible setup, we repeat identical experiments and plot the collected results in Figures 5(b) and 6(b) accordingly. In the case that A-MPDU aggregation is disabled, channel reaches the saturation point as soon as traffic load equals 50 Mbps, while in the A-MPDU enabled case, saturation is only reached at the traffic load of 170 Mbps. Similar observations were made regarding the activation of A-MPDU aggregation, which is only activated when the channel approaches saturation ($\tilde{5}$0 Mbps). Considering the 170 Mbps load case, we observe that A-MPDU aggregation increases throughput from 38.3 Mbps to 164 Mbps (4.3x increase) and decreases the MAC frame transmission rate from 3562 to 465 frames/sec. Our findings clearly verify that MAC layer improvements need to be applied, in order to exploit from the increased PHY bit rates that 802.11n offers.

Having extensively evaluated the throughput performance improvement that the 802.11n protocol can offer across the various considered traffic loads, we next investigate how the monitored improvement is related with the resulting energy consumption. Figures 7(a) and 7(b) illustrate the average power consumption of both the 802.11a/g compatible chipsets and the total Atom node across the various configured traffic load values. As expected, average power consumption at both the NIC and total node level increases at higher traffic loads, due to the increased frequency of frame transmissions at the NIC level and the increased rate of frames that are being processed at the node level. In Fig. 7(a), we observe that the NIC consumes between 1.55 W and 1.73 W, in both the *FF* enabled and disabled case, as the *FF* mechanism is not yet activated. As soon as the traffic load increases above 30 Mbps, *FF* is activated and average power consumption for the *FF* enabled case increases above the average monitored consumption for the *FF* disabled case, till it reaches the maximum value of 1.79 W. This observation comes due to the fact that in the *FF* enabled case, the NIC consumes more power on average as it operates in transmit mode for longer duration. Considering the power consumption of the total Atom node, we observe that the two different 802.11a/g based setups consume different amounts of power on average, due to the use different wireless chipsets and drivers. However, both setups witness an increase of approximately 0.5 W, as the traffic load increases from 5 Mbps to 30 Mbps. In the case of the *FF* enabled 802.11a/g setup, we observe that average power consumption at the total node level is decreased between the 30 Mbps (22.56 W) and the 40 Mbps (22.53 W), in spite of the through-

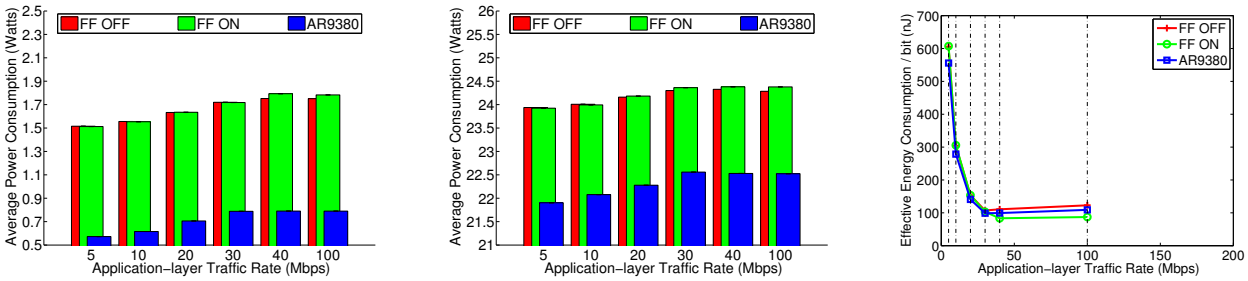

(a) NIC (b) Total Node (c) Total Node Effective Eb

Figure 7: Energy efficiency characterisation of 802.11a/g setup across varying Application-Layer Traffic load

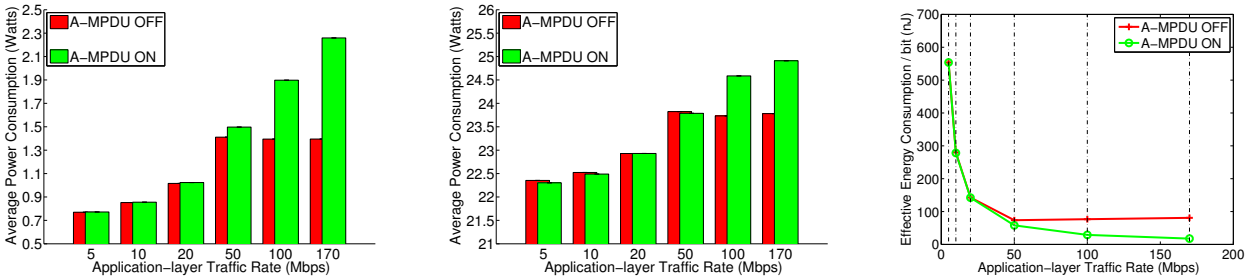

(a) NIC (b) Total Node (c) Total Node Effective Eb

Figure 8: Energy efficiency characterisation of 802.11n setup across varying Application-Layer Traffic load

put performance increase. This observation is related with the activation of the *FF* aggregation mechanism, which efficiently reduces the rate of MAC-layer frames that are being processed by the driver. Our findings are summarised in the Effective E_B representation in Fig. 7(c), which characterises the total node power consumption as a function of the resulting throughput and not as a function of the configured PHY bit rate. The obtained results clearly show that *FF* is able to reduce energy expenditure at the Atom-based node level, even up to 28% in the 40 Mbps case (119.5 nJ/bit FF OFF - 85.51 nJ/bit FF ON). Regarding the performance of the AR9380 equipped wireless node, we remark that in spite of its low power consumption profile, the resulting Effective E_B values are higher than the levels achieved by the application of the *FF* mechanism.

Similar results are obtained while evaluating the impact of A-MPDU aggregation on the power consumption of the AR9380 NIC and the Atom node, which are plotted in Fig. 8(a) and Fig. 8(b) accordingly. We clearly observe that A-MPDU aggregation results in significantly higher average power consumption, for traffic loads above 100 Mbps, both at the NIC as well as at the total node level, as a result of the achievable throughput gains. Comparing the consumption of the Atom node, as plotted in Figures 7(b) and 8(b), we observe that under low traffic loads ($<$ 50 Mbps) both protocols result is similar power consumption behaviour. This comes due to the fact the A-MPDU mechanism is not yet activated and as the high PHY rates of 802.11n are only able to reduce the average power consumption at the NIC level, the consumption of the total node is only minimally impacted. However, while considering traffic loads above 100 Mbps, we notice that the remarkably increased throughput performance that A-MPDU aggregation results in does not come at much higher energy costs. The Effective E_B representation in Fig. 8(c), summarises the above results and shows that A-MPDU aggregation can increase energy efficiency up to 78%. Finally, direct comparison of E_B values at the corresponding saturation points of each standard shows that 802.11n offers more that 80% reduction of E_B compared with the AR9380 based 802.11a/g standard.

5.2 Varying Frame Payload Length

Extensive throughput experiments were also conducted under varying frame payload lengths, in order to investigate how varying Payload size affects throughput performance and energy expenditure. In order to enable delivery of frames longer than 1500 bytes to the MAC layer, we configured the wireless NIC's Maximum Transmission Unit (MTU) size to the maximum supported value of 2304 bytes. We have to mention that in cases where the *FF* mechanism is enabled, aggregation of frames longer than 1700 bytes could not be handled by the driver, as the transmission duration exceeded the threshold of 4 ms that the 802.11 standard specifies as the maximum acceptable frame transmission duration. Having investigated the impact of varying frame payload length on performance under high SNR conditions, we next proceed by conducting identical experiments in the low-SNR experimental setup. The full list of obtained results are detailed in the corresponding sections that follow.

5.2.1 High SNR conditions

The throughput performance of the 802.11a/g and 802.11n protocols are illustrated in Figures 9(a) and 10(a) accordingly. We observe that under high-SNR conditions, increasing frame length values consistently result in improved throughput performance for both protocols. The throughput improvement between the lowest (300) and highest (2200) considered payload lengths varies between the factors of 3x and 5x for all the considered cases, except for the A-MPDU assisted 802.11n scenario, where the improvement is restricted

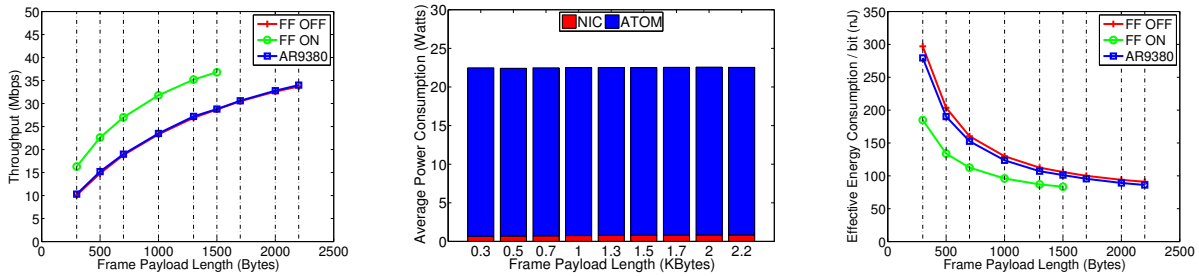

(a) Throughput (b) Average Power consumption (c) Total Node Effective Eb

Figure 9: Performance of 802.11a/g across varying Frame Payload Length values under high-SNR conditions

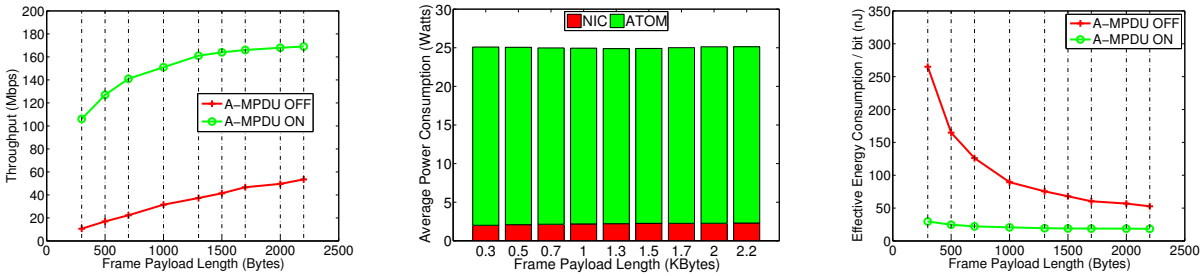

(a) Throughput (b) Average Power consumption (c) Total Node Effective Eb

Figure 10: Performance of 802.11n across varying Frame Payload Length values under high-SNR conditions

in the order of 1.6x. This observation is related to the fact that A-MPDU aggregation enables 802.11n to deliver high throughput performance (>100 Mbps) even at the lowest payload length, as A-MPDU frame size is only limited by the maximum number of subframes (64) and maximum A-MPDU length (65.535 bytes).

Figures 9(b) and 10(b) plot the power consumption of the AR9380 based setup, as it resulted during the operation of the 802.11a/g and A-MPDU-assisted 802.11n protocols accordingly. The full list of presented measurements have been collected using only the AR9380 based setup, in order to provide for direct comparison between the two protocols. Considering the minimum and maximum payload lengths for both protocols, we observe an increase of approximately 0.18 W in the consumption of the 802.11a/g configured NIC (0.3W 802.11n). The observed increase is directly related with the longer duration that the NIC remains in transmission mode, while achieving higher throughput. Regarding the consumption at the node level, in general we did not observe any significant consumption variation across varying payload lengths, which fact comes in contrast with the results obtained in the previous section, where even minor throughput improvement resulted in consumption increase at the total node level. In Fig. 10(b), we even observe that the average power consumption decreases when the frame payload size increases from 300 bytes (25.1W) to 1300 bytes (24.88 W), Considering also that the NIC's average power consumption is also increased between these two cases, we remark that the consumption increase at the total node level approximates 0.5 W. As a result, we reach the conclusion that considerable amounts of energy are consumed while each frame crosses the protocol stack and verify the findings of the work in [21]. This observation in combination with the high throughput gains that payload increase results in, indicate that the use of longer frames is preferable in both terms of network performance and energy efficiency.

Our findings are summarised in the Effective E_B representation in Figures 9(c) and 10(c). In comparison with the AR9380 based 802.11a/g setup, we remark that the A-MPDU assisted 802.11n is able to reduce the Effective E_B at the node level, from 279 nJ/bit to 29 nJ/bit (-90%) and also from 86 nJ/bit to 18 nJ/bit (-80%), when transmitting frames of 300 and 2200 bytes accordingly. Recent studies [22] of the packet size distributions in Internet traffic have shown that the most common packet lengths are of 576 bytes size, which fact highlights even more the energy savings that can be attained through the application of the 802.11n protocol. In addition, we remark that low frame lengths are usually preferable in the wireless domain, as they are able to provide increased FDR, especially when using complex modulation schemes that are susceptible to low-SNR conditions. This observation yields interesting insights and motivates further investigation regarding the performance of 802.11n across varying frame payload lengths and low-SNR conditions.

5.2.2 *Low SNR conditions*

Towards executing identical experiments under low link quality conditions, we establish the low-SNR experimental setup. In the following experiments, we measure the performance of 802.11a/g considering only the AR9380 setup and compare it against the A-MPDU assisted 802.11n case. Under this setup, the 802.11a/g configuration is able to sustain the PHY bit rate of 18 Mbps, while the 802.11n setup is able to use up to the MCS6 configuration in the Spatial Diversity mode. In addition, we measure the performance of less complex modulation schemes and more specifically the rates of 9 Mbps and 12 Mbps for the 802.11a/g configuration, while MCS4 and MCS5 are also configured for the 802.11n protocol. Characteristics of the various configured modulation schemes are listed in Table 3.

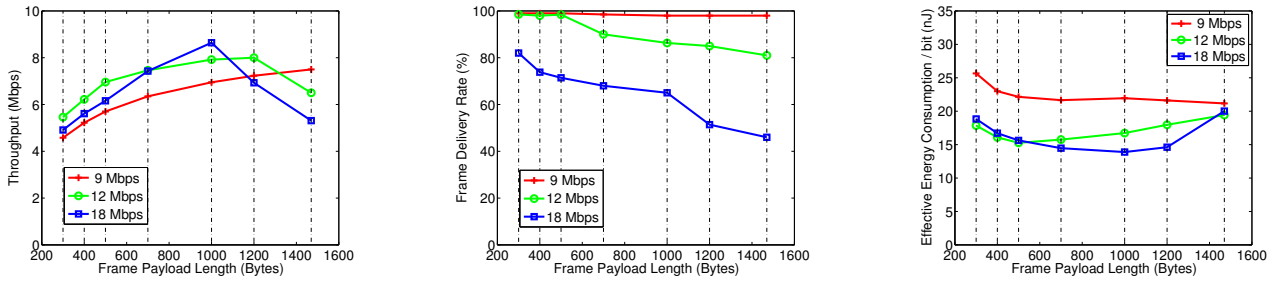

(a) Throughput (b) FDR (c) NIC Effective Eb

Figure 11: Performance of 802.11a/g across varying Frame Payload Length values under low-SNR conditions

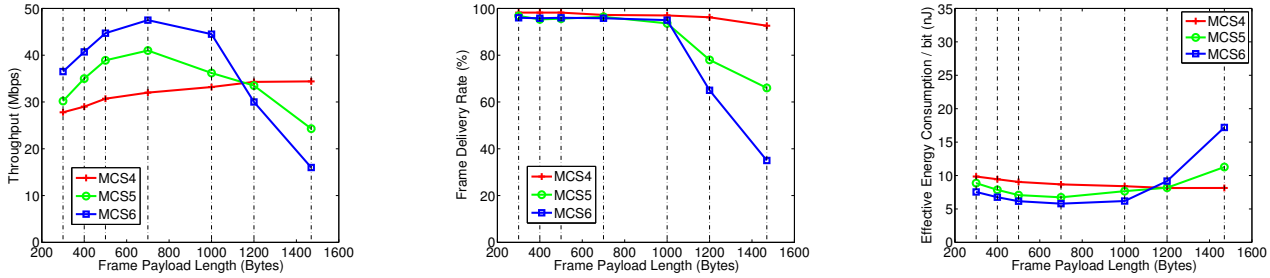

(a) Throughput (b) FDR (c) NIC Effective Eb

Figure 12: Performance of 802.11n across varying Frame Payload Length values under low-SNR conditions

MCS	PHY Rate (Mbps)	Modulation	FEC
9 Mbps	9	BPSK	3/4
12 Mbps	12	QPSK	1/2
18 Mbps	18	QPSK	3/4
MCS4	39	16-QAM	3/4
MCS5	52	64-QAM	2/3
MCS6	58.5	64-QAM	3/4

Table 3: Characteristics per configured MCS

We start by measuring the throughput performance across varying frame payload lengths, between 300 and 1500 bytes, which is the default MTU size for the wireless chipsets under consideration. In Figures 11(a) and 12(a), we illustrate the throughput performance that is achieved by each protocol. We clearly observe that in the default MTU case, the 802.11n protocol is able to deliver significantly higher throughput of 34.4 Mbps than the 7.5 Mbps of 802.11a/g (4.6x increase), by enabling the use of more complex and efficient modulation schemes. Moreover, we notice that only the lowest rate configurations of each protocol consistently provide higher throughput performance for increasing payload length, while in the rest configurations the maximum throughput is achieved under lower frame lengths. Detailed study of the throughput plots shows that proper payload length adaptation is able to provide up to 15% (8.64 Mbps - 1000 bytes payload) increase in the throughput performance of the 802.11a/g protocol and 38% (47.5 Mbps - 700 bytes payload) increase in the performance of 802.11n. The improved throughput performance is related to the increased FDR that lower payload length configurations are able to result in. In Figures 11(b) and 12(b), we depict the FDR performance of each protocol and highlight its relation with the complexity of each modulation type.

Energy consumption measurements were also conducted, in order to evaluate the energy efficiency of each protocol across varying payload lengths. Considering the fact that the AR9380 card is characterised by a totally different power consumption profile in the two setups, along with the highly varying FDR and throughput performance, we conclude that deriving the most energy efficient payload size per case is a rather complex task. Towards deriving concrete conclusions, we plot the Effective E_B representation at the NIC level in Figures 11(c) and 12(c) for each protocol. Regarding the Effective E_B at the NIC level as obtained between the two setups, we notice that the 802.11n setup is able to reduce energy consumption down to 5.78 nJ/bit (MCS6 - 700 bytes) and offer reduction of 58% in comparison with the 13.89 nJ/bit (18 Mbps - 1000 bytes) that the 802.11a/g can offer at best. The obtained results show that payload lengths between 500 and 1200 bytes are preferable in terms of energy efficiency for the operation of the 802.11a/g protocol, while in the case of 802.11n even lower frame sizes between 300 and 1000 bytes can further reduce energy expenditure. Concluding, we remark that it is important to design automated algorithms that jointly adapt the MAC frame payload length and the PHY bit rate, towards achieving higher throughput and lower energy consumption.

5.3 Experimentation with 802.11 PSM

Through this experiment, we aim at quantifying the potential energy savings of the 802.11 PSM during periods of network activity, by experimenting in network setups that are based on both protocols under consideration. The 802.11 PSM mechanism is designed to set the wireless NICs of stations (STAs) in a low-power state during periods of inactivity and periodically activate them to fetch cached data from the access point (AP). Considering an active network, the useful period during which the STA's NIC can remain deactivated, is directly determined by the inter-packet arrival time of traffic flows that are destined to the STA. While the STA's NIC is in sleep mode, all cached frames at the AP are being delayed till the next *Beacon Interval*.

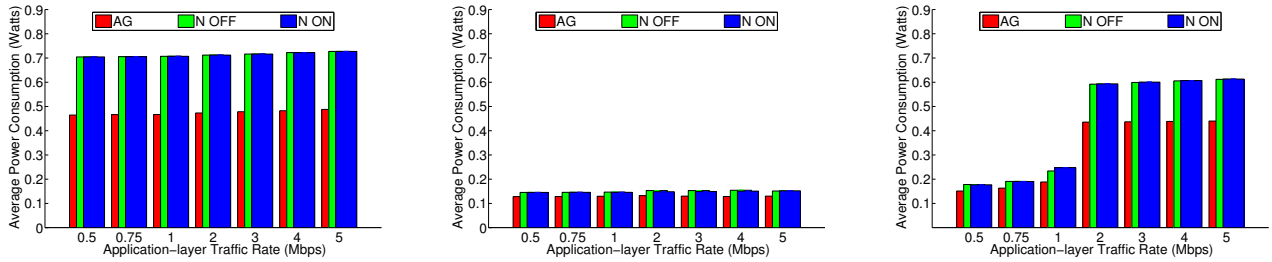

(a) PSM OFF (b) PSM ON - Timeout 1 ms (c) PSM ON - Timeout 15 ms

Figure 13: Power consumption of AR9380 NIC across varying PSM configurations

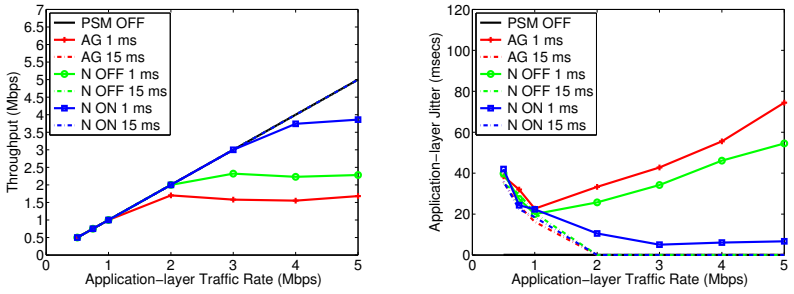

(a) Throughput (b) Jitter

Figure 14: Network performance of AR9380 NIC across varying PSM configurations

As the PSM mechanism is only able to affect the consumption of STAs, in this experiment we consider downlink transmissions and measure the impact of PSM on the energy efficiency of STAs (receivers), while also evaluating network performance metrics. We configure the high-SNR experimental setup that was described in Section 3 and equip both nodes with the AR9380 chipset that supports the 802.11 PSM mechanism. We also assign the default *Beacon Interval* of 100 ms at the AP node. Towards stressing the operation of the 802.11 PSM mechanism, we exploit the ability of the *ath9k* driver to tune the Timeout Period (*TP*) parameter, which configures the interval before the NIC goes back to sleep mode, in order to control the tradeoff between the induced delay and energy savings. In this experiment we vary the *TP* between the minimum value of 1 ms and 15 ms and investigate the impact of PSM on application-layer performance and NIC energy efficiency.

We start by measuring performance in the case that the PSM is deactivated and proceed with the next two phases, where the 1 ms and 15 ms TP intervals are configured. In each phase, we vary the application-layer traffic rate at the AP side, by configuring values between 0.5 Mbps and 5 Mbps and measure the network performance in terms of throughput and jitter, while also monitoring the energy consumption of the STA's NIC. Figures 13(a), 13(b) and 13(c) present the average power consumption of the STA's NIC in each phase, considering the application of each protocol accordingly. In the case that the 802.11n protocol is applied, we also distinguish between the A-MPDU assisted (N ON) and the non-assisted case (N OFF). We observe that the 1 ms *TP* configuration provides significant energy savings, in comparison with the PSM disabled case, that approximate at maximum 74% for the 802.11a/g setup (AG) and 79% for the 802.11n setup (N ON) . Considering the 15 ms *TP* configuration, a minimal reduction of energy savings is observed, across traffic load values below 2 Mbps. In the case that the traffic load exceeds the 2 Mbps value, the NIC rarely falls in sleep mode across all the considered cases, thus resulting in lower energy savings (not exceeding 15%).

Next, we characterise the impact of PSM on application layer performance, considering throughput and jitter, as plotted in Figures 14(a) and 14(b) accordingly. In general, we observe that the 15 ms *TP* configuration poses no impact on throughput and only minimal impact on jitter performance, thus not sacrificing network performance for saving energy. On the other hand, the 1 ms *TP* configuration that stresses the operation of the 802.11 PSM mechanism, provides more interesting results that clearly highlight the impact of the 802.11n's high PHY bit rates and A-MPDU aggregation on network performance. As depicted in Fig. 14(a), the increased PHY bit rates of the 802.11n improve throughput in comparison with the 802.11a/g setup, when considering traffic loads above 2 Mbps. Moreover, we observe that the application of A-MPDU aggregation further increases throughput performance, by combining several cached at the AP frames into a single A-MPDU frame, thus efficiently reducing frame losses due to buffer overflows at the transmitter. The impact of the 802.11n's increased PHY bit rates and A-MPDU aggregation in jitter performance are highlighted in Fig. 14(b), where we observe that the A-MPDU assisted protocol constantly enables on time frame delivery and results in remarkably reduced jitter.

Concluding, we remark that this last experiment has clearly demonstrated that the 802.11n protocol is able to provide both increased network performance and significant energy savings through the application of the PSM mechanism, during periods of network activity. Moreover, our results have shown that scheduling of sleep intervals in an adaptive to the prevailing traffic conditions and protocol parameters way, is able to bridge the gap between high network latency and low energy savings, as shown in [23].

6. CONCLUSIONS

In this work, we presented detailed experimentally obtained results that evaluate the energy efficiency of the base 802.11 standard in comparison with the latest 802.11n version, under a wide range of settings. In-depth analysis of the collected results has shown that the advanced features of the latest standard enable significant reduction of energy expenditure, across all the various considered scenarios. We envision that our findings will provide valuable insights to researchers working on the design of energy efficient wireless protocols.

7. ACKNOWLEDGEMENTS

The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 258301 (CREW IP project). I. Koutsopoulos, N. Makris and G. Kazdaridis also acknowledge the support of ERC08-RECITAL project, co-financed by Greece and the European Union (European Social Fund) through the Operational Program Education and Lifelong Learning - NSRF 2007-2013. S. Keranidis specifically acknowledges the support of AGILENET Marie Curie IRSES project.

8. REFERENCES

[1] IEEE 802.11-2007 Wireless LAN Medium Access Control and Physical Layers Specifications.

[2] D. Skordoulis, Qiang Ni, Hsiao-Hwa Chen, A.P. Stephens, Changwen Liu, and A. Jamalipour. IEEE 802.11n MAC frame aggregation mechanisms for next-generation high-throughput WLANs. *Wireless Communications, IEEE*, 15(1):40–47, 2008.

[3] K. Pelechrinis, T. Salonidis, H. Lundgren, and N. Vaidya. Experimental characterization of 802.11n link quality at high rates. In *Proceedings of ACM WiNTECH*, 2010.

[4] L. Kriara, M.K. Marina, and A. Farshad. Characterization of 802.11n wireless LAN performance via testbed measurements and statistical analysis. In *Proceedings of SECON*, 2013.

[5] D. Halperin, B. Greenstein, A. Sheth, and D. Wetherall. Demystifying 802.11n power consumption. In *Proceedings of SIGOPS HotPower*, 2010.

[6] I. Pefkianakis, Chi-Yu L., and Songwu L. What is wrong/right with IEEE 802.11n Spatial Multiplexing Power Save feature? In *Proceedings of ICNP*, 2011.

[7] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani. "energy consumption in mobile phones: A measurement study and implications for network applications". In *Proceedings of IMC*, 2009.

[8] M. Tauber and S.N. Bhatti. The Effect of the 802.11 Power Save Mechanism (PSM) on Energy Efficiency and Performance during System Activity. In *Proceedings of IEEE GreenCom*, 2012.

[9] S. Keranidis, G. Kazdaridis, V. Passas, T. Korakis, I. Koutsopoulos, and L. Tassiulas. Online Energy Consumption Monitoring of Wireless Testbed Infrastructure Through the NITOS EMF Framework. In *Proceedings of ACM WiNTECH*, 2013.

[10] G. Bhanage, R. Mahindra, I. Seskar, and D. Raychaudhuri. Implication of MAC frame aggregation on empirical wireless experimentation. In *Proceedings of GLOBECOM*, 2009.

[11] Y. Xiao and J. Rosdahl. *Throughput and delay limits of IEEE 802.11, IEEE Communications Letters, 2002.*

[12] *"Vendor Specific Improvements"*, http://goo.gl/ti8cxT.

[13] *"Wireless Incompatibilities"*, http://goo.gl/5c1cHq.

[14] IEEE 802.11n-2009, Amendment 5: Enhancements for Higher Throughput., 2009.

[15] B. Ginzburg and Kesselman A. Performance analysis of A-MSDU andA-MPDU aggregation in IEEE 802.11n. In *Proceedings of IEEE SARNOFF*, 2007.

[16] *"Mad-WiFi Wireless driver"*, http://madwifi-project.org/.

[17] *"Ath9k Wireless driver"*, http://goo.gl/VrHtj.

[18] *"Iperf"*, http://dast.nlanr.net/Projects/Iperf/.

[19] *"Maximal-ratio combining"*, http://goo.gl/ABj8PY.

[20] C. Li, C. Peng, S. Lu, and X. Wang. "Energy-based rate adaptation for 802.11n". In *Proceedings of Mobicom*, 2012.

[21] A. Garcia-Saavedra, P. Serrano, A. Banchs, and G. Bianchi. "Energy consumption anatomy of 802.11 devices and its implication on modeling and design". In *Proceedings of CoNEXT*, 2012.

[22] *"Packet size Distributions in Internet traffic"*, http://goo.gl/cBLxq6.

[23] K. Jang, S. Hao, A. Sheth, and R. Govindan. "Snooze: energy management in 802.11n WLANs". In *Proceedings of CoNEXT*, 2011.