Demystifying 802.11n Power Consumption

Daniel Halperin*, Ben Greenstein†, Anmol Sheth†, and David Wetherall*†

University of Washington and Intel Labs Seattle†*

Abstract

We report what we believe to be the first measurements of the power consumption of an 802.11n NIC across a broad set of operating states (channel width, transmit power, rates, antennas, MIMO streams, sleep, and active modes). We find the popular practice of racing to sleep (by sending data at the highest possible rate) to be a useful heuristic to save energy, but that it does not always hold. We contribute three other useful heuristics: wide channels are an energy-efficient way to increase rates; multiple RF chains are more energy-efficient only when the channel is good enough to support the highest MIMO rates; and single antenna operation is always most energy-efficient for short packets.

1 Introduction

802.11 continues to advance and plays a key role in enabling increasingly demanding networked applications like video conferencing, multiplayer 3-D games, and cloud-supported mobile augmented reality [11] on mobile computing devices. Even smartphones have already moved to using 802.11n, the latest version of the standard, which users generally prefer over wide area technologies (e.g., 3G, LTE, and WiMAX) because its capacity (up to 600 Mbps [2]) is orders of magnitude greater.

All of this communication costs energy, and thus it is important to use the network efficiently in order to maximize platform lifetime. This concern is particularly heightened for smartphones, where radio interfaces can account for up to 50% of the total power budget under typical use [9] and can quickly drain phones' limited batteries when transmitting at peak rates. The 802.11n radio we tested in our lab, for example, can draw 2.1 W in its high rate MIMO configurations, which would deplete a typical smartphone battery in under three hours and would emit nearly enough heat to burn a user's hand.[1] While the fraction of network power varies with the platform, it is likely to remain an important component as wireless capability increases and overall power consumption drops. The driving factors for energy efficiency are unlikely to change either, given concerns for sustainability, and that battery capacity grows at a rate of only 10% every two years [7].

However, there is precious little data to help designers understand how to use 802.11 cards in an energy efficient fashion in practice. Modern 802.11n cards now have many factors that affect active power consumption, including multiple rates, antennas and channel widths. Many of these factors are new with 802.11n and were not present in 802.11a/b/g NICs and thus are not covered in prior studies of these cards [4, 9]. Data sheets, standards [2], and analytical work [8] provide little insight, and there is no detailed reported power data for these cards and factors, to the best of our knowledge. One 802.11 channel width study [3] found that using large channels could worsen power consumption by 40%, but used a different technique than that of 802.11n.

Moreover, the basic strategies for minimizing energy expenditure are unclear. The latest power-saving proposals for 802.11 (e.g., APSD and PSP [2]) presume that power consumption is best reduced by a "race to sleep," i.e., transmitting at the highest bit rate possible, to allow more time to sleep. This is based on the presumption that the power consumption in active modes is orders of magnitude greater than in sleep modes because radio circuitry consumes substantially more power than is actually radiated. On the other hand, Shannon capacity implies that the radiated energy consumption per bit (which is the energy cost that cannot be reduced) *grows* with the bit rate [8]. This has led others to argue that a radio should transmit at the *slowest* rate possible to be most power efficient [6, 5].

In this paper, we report what we believe to be the first measurements of power consumption of an 802.11n NIC running in many different modes available in laptops today and presumably in smartphones in the near future. Our goal is to demystify the power consumption of 802.11 NICs by answering basic questions. For example: Does reducing transmit power yield significant savings? Is sleeping much more efficient than idling? Does transmitting cost more power than receiving? Are wider channels always less energy-efficient? Are faster rates always more energy-efficient? Is power roughly a multiple of the number of MIMO streams? The answers, as we will see, are no, yes, yes, no, no, and no.

In addition to yielding the answers to these basic questions, our results help us understand how to design energy-efficient wireless systems. We find that the "race to sleep" model of using the highest rate is roughly correct, but with some important caveats. These lead to simple recommendations that can improve energy-efficiency significantly. For example, receiving short packets with a single antenna (e.g., at 108 Mbps) can be more than twice as energy-efficient as receiving as fast as possible using all three antennas (e.g., 405 Mbps) to maximize sleep.

[1] A phone reaches this thermal limit when run at 2–3 W [12].

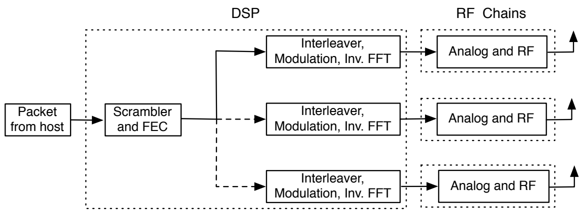
DSP
RF Chains
Packet from host
Scrambler and FEC
Interleaver, Modulation, Inv. FFT
Interleaver, Modulation, Inv. FFT
Interleaver, Modulation, Inv. FFT
Analog and RF
Analog and RF
Analog and RF

Figure 1: Transmit block diagram of a MIMO radio that supports up to 3 spatial streams. The DSP chip modulates packet data into multiple spatial streams. A dedicated RF chain per antenna includes the analog and RF components that produce the transmitted signals.

We describe the key 802.11 NIC operating modes we measure in §2. Our experimental setup is described in §3. Then we report basic power measurements in §4, followed in §5 by the implications of these measurements for minimizing the energy per bit that is transmitted or received. We conclude with next steps in §6.

2 802.11n Background

The IEEE 802.11n standard [2] offers wireless bitrates as high as 600 Mbps. 802.11n takes the spectrally efficient orthogonal frequency division multiplexing (OFDM) physical layer of 802.11a/g, and adds antennas to take advantage of multiple spatial paths in the radio frequency (RF) environment. The multi-antenna based "multiple input, multiple output" (MIMO)[2] techniques of 802.11n use multiple antennas to provide robustness and range to weak links through *spatial diversity* — sending *redundant* information along independent spatial paths. MIMO raises the maximum wireless bitrates through *spatial multiplexing*, techniques that send *independent* information along independent spatial paths.

The gains from spatial diversity and spatial multiplexing do not come for free. To use multiple antennas for transmission or reception, a device must have multiple instances of its RF analog processing hardware (called *RF chains*), one dedicated to each antenna. Most other components of wireless transmission and reception, such as RF oscillators and the digital signal processing (DSP) chip, need only be instantiated once and can be shared by all RF chains and antennas. Figure 1 shows a simplified block diagram of a MIMO transmitter with three transmit/receive antennas.

Multi-antenna systems have additional configuration options, which have implications to power consumption. If a receiver activates all its antennas and RF chains, it can receive weak signals more reliably and strong signals at much higher rates; however, the extra active RF chains consume more power. This suggests a need for understanding the energy costs of these techniques in practice.

Config.	Streams	Tx. Chains	Rx. Chains
SISO	1	1	1
SIMO	1	1	1, 2, or 3
MIMO2	2	2	2 or 3
MIMO3	3	3	3

Table 1: MIMO configurations that use different number of spatial streams and Tx/Rx RF chains.

In this work, we consider a commercial Intel 802.11n NIC with 3 transmit/receive antennas with a peak power draw of 2.1 W. Its configuration space includes the number of transmit spatial streams (up to 3) and RF chains used by the transmitter/receiver (up to 3). Additionally, 802.11n adds the option of using wide 40 MHz channels that offer twice the rate, and so the choice of 20 or 40 MHz channels is another component of the configuration. We also consider two legacy 802.11 knobs: link rate used across the spatial streams (6.5 to 65 Mbps), and transmit power (from -10 dBm to $+15$ dBm). Table 1 shows the relationship between the number of spatial streams and RF chains in use. We measure the NIC while it is transmitting, receiving, listening to an idle medium, and sleeping.

3 Power Measurement Setup

We perform power measurements for the Intel Wi-Fi Link 5300 a/b/g/n wireless network adapter. This is a commercial NIC available since mid-2008, and is representative of the state of the art. The 5300 supports 3x3 MIMO (up to three transmit and three receive streams).

We conduct measurements using two testbed nodes running the 2.6.33-rc7 Linux kernel and a modified version of Intel's open source `iwlagn` driver [1]. The nodes are situated close to each other with a very strong connection that supports all available 802.11n rates. The modified driver allows us to quickly reconfigure the transmitter and receiver configurations by dynamically disabling antennas and RF chains; changing channel widths between 20 and 40 MHz; and varying the number of spatial streams, the link rates, and the transmit power.

We measure power consumption of the NIC by instrumenting one of our two testbed nodes. The NIC we use is available in a mini-PCI Express form factor, which we connect to the desktop testbed node via an adapter that converts to a PCI Express interface. We place a current sense resistor (40 mΩ) on the 3.3 V power supply to the NIC. By measuring and recording[3] the voltage drop across the resistor we can calculate the current and thus the power consumed by the NIC. We send long (1500 byte) packets and average multiple samples during the packet payload to estimate the instantaneous power consumption of the NIC during transmission or reception; and similarly for sleep and idle receive modes.

[2] We use "single-input, single-output" (SISO) configurations with one antenna at each end, and SIMO/MISO as well.

[3] We use the National Instruments 6218 Data Acquisition Module, which is capable of sampling voltage at 250 kHz with 16-bit precision, and LabView software.

Device	Mode	Power Consumption
Intel 5300	Idle Rx	820 mW
Intel 5300	Sleep	100 mW
Intel 4965	Sleep	46 mW

Table 2: Comparison of Idle Rx and Sleep power states for two Intel 802.11n NICs.

Mode	Power Consumption
Idle Rx, 1 Ant	0.82 W
Idle Rx, 2 Ant	1.13 W
Idle Rx, 3 Ant	1.45 W
Rx SISO, 1 Ant	0.94 W
Rx MIMO2, 2 Ant	1.27 W
Rx MIMO3, 3 Ant	1.60 W
Tx SIMO	1.28 W
Tx MIMO2	1.99 W
Tx MIMO3	2.10 W

Table 3: Power consumption for various Rx and Tx modes. One number is shown per configuration because power use varies less than 10% for different rates and channel widths within one mode.

For the discussion in this paper, we use the notation in Table 1 to specify the number of transmit or receive antennas in use. When relevant for SIMO and MIMO2, we also specify the number of receive antennas used.

4 802.11n Power Consumption

In this section, we analyze the measurements described in §3 to answer basic questions about how much power is saved by sleeping, and how rate, channel width, spatial multiplexing and transmit power affect the power consumed in transmission and reception.

Low-Power Sleep. Table 2 presents a comparison of active and sleep power draw. The lowest active mode is Idle Rx, in which the receiver has a single antenna enabled and listens for the start of a transmission on an idle channel. In Sleep mode, the NIC disables most of its circuitry and puts the remainder in a low-power state, enabling a few interrupts for timers (e.g., to wake up for beacons) and for host communication. Idle Rx for a single antenna consumes about 820 mW, and sleeping while associated consumes only 100 mW.[4] *This suggests that indeed increasing sleep time (when workload permits) is an effective strategy to save power.*

Rate. Table 3 shows the power consumed when listening to an idle medium, reception using 1, 2, or 3 antennas, and transmission of single-, double-, and triple-stream packets. We present one number per configuration of transmit and receive antennas and spatial streams, as we observed that the power consumed when sending or receiving packets using different rates does not vary significantly. *These results suggest that simply increasing the link speed for the same antenna configuration does not turn on extra hardware.* The higher bits/sec DSP processing required to process the fastest rates does incur a small overhead while receiving. Receiving the fastest MIMO3 40 MHz rate (405 Mbps) uses 120 mW more than the slowest (40.5 Mbps).

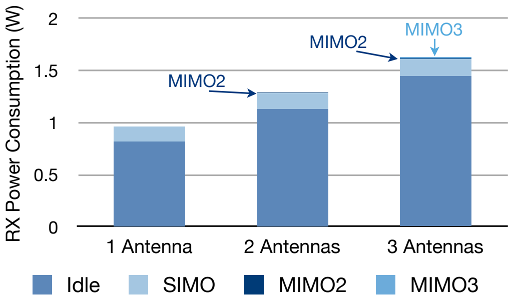

Figure 2: Incremental receive power consumption of single-stream, MIMO2, and MIMO3 processing.

Channel Width. We observed that using wider channels has a negligible impact on power consumption for the transmitter and receiver while more than doubling the data rate of the link. This presents the opposite conclusion from the study by Chandra et al. [3] that found wider channels to be significantly more costly. That work increased the NIC's clock frequency and OFDM parameters, but 802.11n simply adds subcarriers keeping clocks (already fast for processing high data rates) the same. *This implies that wide channels should be used where available, as they are always more energy-efficient.*

Multiple Spatial Streams. We now analyze the power consumption of transmitting multiple spatial streams from Table 3. Transmitting a single-stream signal consumes about 1.3 W, while a MIMO2 sender consumes 2.0 W and MIMO3 2.1 W. It is interesting to note that this constitutes a large jump from 1 to 2 streams and only a small gap between 2 and 3 streams. This suggests that a MIMO2 transmission engages nearly all of the hardware components used for MIMO3. In summary, for this NIC, *MIMO2 consumes approximately 53% more power than SISO when transmitting while MIMO3 consumes only 5% more than MIMO2.*

Multiple Receive Antennas. For each receive configuration in Table 3, we enable only as many antennas as spatial streams to minimize the number of active receive chains. We see that the use of multiple chains does indeed raise power consumption. Enabling each receive chain adds 330 mW to the power consumption. However, this increase of about 330 mW per chain represents only a factor of 1.3 (MIMO2) and 1.7 (MIMO3) rather than directly doubling or tripling costs.

Receiving Multiple Streams. Next, we analyze in Figure 2 the cost of receiving multiple spatial streams rather

[4] We expect the sleep power to further reduce as the Intel 4965, an older 802.11n NIC that supports only 2 spatial streams, consumes only 46 mW in Sleep mode. Likewise, sleep power would be additionally reduced in a phone by using a more power-efficient interface than PCIe.

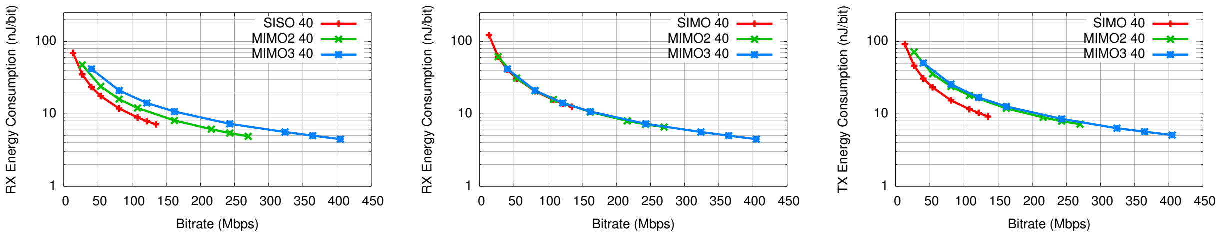

(a) Receiving with only necessary antennas enabled.

(b) Receiving with 3 antennas enabled.

(c) Transmitting.

Figure 3: Energy consumption per bit.

than simply having excess receive antennas enabled. To isolate the cost of spatial multiplexing as opposed to the cost of faster data rates, we use the 20 MHz 1x39 Mbps, 2x19.5 Mbps, and 3x13 Mbps configurations that each support an aggregate link rate of 39 Mbps. Receiving a single-stream packet adds 140–170 mW to the IDLE RX power. Surprisingly, we see that the incremental cost of MIMO2 and MIMO3 over SIMO is less than 10 mW, a negligible increase compared to the 330 mW cost of enabling the RF chains. *Thus, to save power in* RX *mode, it is much more important to minimize the number of antennas enabled then to reduce the number of spatial streams.*

Transmit Power. Finally, we conducted an experiment in which we reduced the transmit power from 32 mW to less than 1 mW. Reducing the radiated power by about 97%, the total consumption in TX mode dropped by about 140 mW total. This is significantly more than the reduction in transmit power, yet still only about 10% of the power consumed. This concurs with prior work on 802.11g [10], which suggests that *transmit power control offers little gain in practice since the savings represent a small fraction of total power.*

5 Racing to Sleep

The "race to sleep" heuristic makes the key assumption that the energy to transmit or receive packets decreases monotonically as data rate increases. To understand whether this assumption holds, we process the measurements in §4 to compute the per-bit energy cost as power consumption (W = J/s) divided by bitrate (Mbps), which results in an energy cost (J/bit) for each configuration and rate. In the following figures, we plot the energy cost against the data rate for RX mode using only as many antennas as spatial streams (Figure 3(a)); for RX mode using all antennas (Figure 3(b)); and for TX mode (Figure 3(c)). As the 20 MHz and 40 MHz lines for each configuration overlap in all but the lowest 20 MHz rate, we omit the 20 MHz modes from these graphs.

Fixed Antenna Configuration. In §4, we found that power consumption varies minimally within the same transmit or receive configuration. Therefore Figures 3(a), 3(b), and 3(c) show that a faster MCS using the same configuration always has a lower energy cost. These results support the existing 802.11 wisdom (for which, lacking MIMO, there is only a single configuration for a NIC) that it is always better to send faster. This even holds across all receive configurations when excess antennas are enabled (Figure 3(b)), because the incremental cost of MIMO processing is negligible.

Several Spatial Multiplexing Options. Conversely, the "race to sleep" strategy is *not* ideal when comparing across configurations that consume different amounts of power. Both Figure 3(a) and Figure 3(c) show that fast single-stream configurations are better than all but the fastest MIMO2 and MIMO3 rates, including some configurations that have a higher link data rate. And as these fastest MIMO2/3 rates are unlikely to work except on the best links, there will be scenarios when an efficient client will want to use fewer streams.

Packet Overhead. Thus far, we have considered the energy cost of transmission in terms of each bit of the payload. However, to translate this per bit cost into a cost per packet, we must account for packet overhead.

There are two important components of packet overhead. The first is padding. All 802.11a/g/n packets are sent as a series of OFDM symbols with data on each subcarrier; the unfilled last symbol is padded and devices must pay the energy cost of the extra bits. For example, a 10-byte ACK payload comprises 134 data bits after encoding and checksum, and this fits in a single SISO 39 Mbps symbol that can carry 156 data bits. There will be no energy savings using any faster rate, even those with a lower nominal energy cost. This single-symbol payload represents the shortest transmission possible, and using a faster rate only means more padded bits in the first and only symbol. Note that padding can be substantial: at the highest MIMO3 rate, a single symbol carries 1 620 bits.

The second source of packet overhead is the preamble. Packets sent using multiple spatial streams have longer preambles that enable the receiver to measure the channel between all pairs of transmit and receive antennas and decompose the streams. MIMO2 and MIMO3 preambles are 1 and 3 OFDM symbols longer than preambles for single-stream packets. At 405 Mbps, this effectively

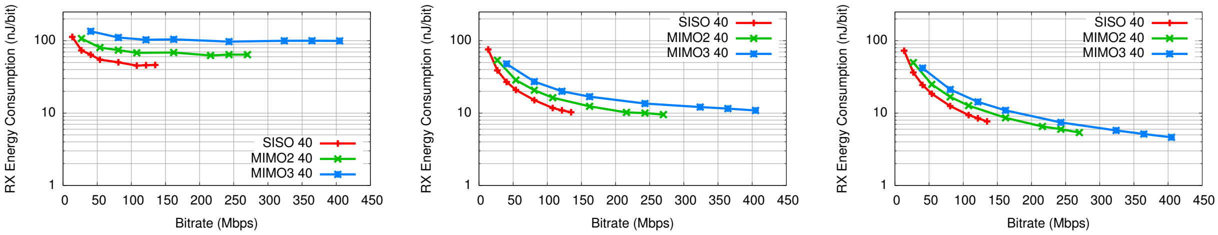

(a) Receiving a 100-byte packet.

(b) Receiving a 1500-byte packet.

(c) Receiving a 65 000-byte packet batch.

Figure 4: Amortized energy consumption.

lengthens the MIMO3 packet by 4 860 bits.

The relative costs of overhead vary with packet size, as the costs of longer transmissions are increasingly dominated by payload. We consider short packets (100 bytes), long packets (1500 bytes), and batches of packets (a single 65 000 byte packet using 802.11n aggregation). We compute the amortized energy consumption of the packet by computing the total energy to send or receive the packet, including overhead, and dividing this figure by the number of data bits. Figures 4(a), 4(b), and 4(c) show the energy cost of receiving these packets with excess antennas disabled. To save space, we omit the graphs for transmission as the results do not differ significantly.

Figure 4(a) shows clearly that for small payloads, the use of SISO mode is most efficient. Even moderate single-stream rates use less energy than all MIMO2 rates. The only situation in which it would be advantageous to use multiple antennas would be for a very weak SISO link operating at the worst rate. In this case the additional receive diversity from a second antenna might improve speed enough to reduce power use.

In, Figure 4(b) we see that when the payload length approaches 1500 bytes, the best single-stream, MIMO2, and MIMO3 rates all consume approximately the same amount of energy. MIMO2 is marginally better than the other configurations, yet again this holds only for the fastest rates. Even for large packets, most links will do best with single-stream transmissions.

Finally, Figure 4(c) shows that as workload increases to the point that large packet batches are being sent, the amortized energy use approaches the base power draw calculated in Figure 3(a). MIMO2 and MIMO3 rates prove more efficient than single-stream transmission. Even in these scenarios, however, the case for using multiple streams is limited to strong links.

In summary, *a good heuristic for optimizing energy consumption is to use the fastest single-stream rate possible, especially for shorter packets. Only for large packets and strong links is the use of multiple streams energy-efficient.*

6 Conclusion

We presented the first empirical characterization of the power consumption of an 802.11n NIC across a broad range of configurations. We have found that some configuration changes, e.g., doubling the bandwidth to double the bit rate, cost very little, while others, e.g., adding a transmit chain, cost a lot, and that some changes that seem to have everything to do with power, e.g., transmit power control, have very little effect on the power actually consumed. Most surprisingly, we found that SISO transmissions are often more power efficient than MIMO, particular when packets are short, despite transmitting at a lower bit rate and providing less time to sleep. It is our hope that the observations we presented will inform the design of future power efficient radios and protocols. In practice, optimal device settings will also depend on channel conditions and workload.

References

[1] Intel wireless WiFi Link drivers for Linux. http://intellinuxwireless.org/.

[2] IEEE Std. 802.11n-2009: Enhancements for Higher Throughput. http://www.ieee802.org/, 2009.

[3] R. Chandra et al. A case for adapting channel width in wireless networks. In *ACM SIGCOMM*, 2008.

[4] L. M. Feeney and M. Nilsson. Investigating the energy consumption of a wireless network interface in an ad hoc networking environment. In *IEEE INFOCOM*, 2001.

[5] C. Gunaratne et al. Reducing the energy consumption of Ethernet with adaptive link rate (ALR). *IEEE Trans. Comput.*, pages 448–461, 2008.

[6] R. R. Kompella and A. C. Snoeren. Practical lazy scheduling in sensor networks. In *ACM SenSys*, 2003.

[7] K. Lahiri et al. Battery-driven system design: A new frontier in low power design. In *ASP-DAC*, 2002.

[8] G. Miao, N. Himayat, Y. G. Li, and A. Swami. Cross-layer optimization for energy-efficient wireless communications: a survey. *Wireless Communications and Mobile Computing*, 9(4):529–542, 2009.

[9] T. Pering et al. CoolSpots: Reducing the power consumption of wireless mobile devices with multiple radio interfaces. In *ACM MobiSys*, 2006.

[10] F. Piazza et al. On the effects of transmit power control on the energy consumption of WiFi network cards. In *Springer QSHINE*, pages 463–475, 2009.

[11] M. Satyanarayanan, P. Bahl, R. Cáceres, and N. Davies. The case for VM-based cloudlets in mobile computing. *IEEE Pervasive Computing*, 2009.

[12] E. Walsh, R. Grimes, and P. Walsh. The performance of active cooling in a mobile phone. In *ITHERM*, 2008.