MIPS: MIMO Power Save in 802.11n Wireless Networks

Ioannis Pefkianakis, Suk-Bok Lee, Songwu Lu
UCLA Computer Science
{pefkian, sblee, slu}@cs.ucla.edu

ABSTRACT

A fundamental limitation of applying MIMO in battery-constrained, 802.11n-enabled portable devices, is excessive power consumption of multiple active RF chains. We start our study in 802.11n MIMO power save (MIPS) by seeking to identify the factors that determine the most energy efficient chain setting. Surprisingly, our experimental results show that a) the fastest RF chain setting may not be the most energy efficient and b) the most power-hungry RF chain setting may not be the least energy efficient. The winner in fact derives from a unique interplay among MIMO gains, SNR, application data source rate, power consumption and can dynamically change in time. To this end we design PollChain, a transmitter-side MIPS, that opportunistically polls the best candidate chain settings to identify the most energy efficient one. PollChain applies novel adaptive probabilistic and frame aggregation bounding mechanisms to limit polling overhead. Our evaluation shows energy savings up to a factor of 3.8 in controlled scenarios over fixed chain settings. In field trials, PollChain yields savings from 9.5% to 26.2%.

Keywords

MIMO, Power Save, IEEE 802.11n

1. INTRODUCTION

The recently ratified 802.11n standard has opened the way for fully leveraging Multiple-Input Multiple-Output (MIMO) technologies in wireless LANs. With supported rates up to 600Mbps, it will be possible to effectively unwire enterprises and enable mobile devices to serve all the bandwidth intensive applications as video conferencing, multiplayer games, and content distribution. 802.11n has already started to appear in APs, laptops, WiFi-enabled portable devices and is expected to equip at least 87% of the 802.11n capable Smartphones by 2014 [1]. However, the high power consumed by the complex MIMO RF chain circuits (up to 3.7 Watt for MIMO over SISO in our platform) has prevented the first battery-powered 802.11n portable devices from implementing MIMO and has necessitated the development of new power supply standards [3].

To overcome this technological barrier it is imperative to apply MIMO power save (MIPS), which opportunistically turns on/off available RF chains[1] at runtime. Across this line, we start our study by asking a simple question. *What are the factors that determine the most energy efficient RF chain setting*? Our experimental study shows two interesting observations. First, *the fastest RF chain setting may not be the best*, in terms of energy efficiency. This contrasts with legacy 802.11a/b/g, where power consumption is best reduced by transmitting at the highest bit rate possible, to allow more time to sleep. Second, *the most power hungry RF chain setting may not be the worst*. The winner chain setting results from a very unique interplay among MIMO gains, SNR, application data source rate, power consumption and can dynamically change in fine time granularity. Specifically, MIMO wins only if its effective goodput gains which are upper-bounded by offered source rate, can compensate the MIMO circuit power consumption. Our measurements of energy consumption at the transmitter side, show that the selection of the best chain setting can yield energy savings up to 40% for SIMO over MIMO and up to 143% of MIMO over SIMO when the underlying rate adaptation is MiRA [6].

To this end we design PollChain MIPS, a transmitter-side 802.11n MIMO power save solution, which seeks to identify and set the most energy efficient transmission chain setting at runtime. PollChain MIPS uses timers and events to poll the best candidates of the available chain settings and finally selects the most energy efficient one. It balances between two conflicting objectives. From one hand, it seeks to be adaptive to MIMO channel and data source rate dynamics. On the other hand, it tries to limit polling overhead, which can significantly degrade system's performance as shown by our experiments. To achieve this, it applies a novel adaptive probabilistic polling scheme, whose goal is to limit transmissions to chains which consistently offer low performance, while remaining adaptive to source rate and

[1] In this paper, we use antennas and RF chains interchangeably, with a slight abuse of notation.

MIMO channel dynamics. It also seeks to further amortize polling overhead by bounding 802.11n frame aggregation, without compromising goodput performance. Along with PollChain, we design sequential-based and threshold-based MIPS, which they apply dynamic and fixed thresholds respectively, to switch to the most energy efficient chain setting.

We evaluate PollChain along with MIPS alternatives and fixed chain settings, under controlled static, mobile settings and field trials. We set up various data source rate scenarios and we are experimenting with both legacy 802.11a/b/g and MIMO 802.11n rate adaptation (RA) designs. PollChain is proven to be the most efficient, with energy savings up 15.3%, 26.2% over fixed chain settings and MIPS alternatives respectively, in field trials. In controlled settings, energy savings go up to a factor of 3.6, 3.8 over fixed chains and MIPS alternatives respectively, when they run over Atheros RA.

The contributions of this work are threefold. First, we experimentally study the average and runtime per-bit energy consumption of more than 25 wireless links in an indoors campus setting. Differently from earlier work [10, 17], we illustrate the tradeoff between MIMO goodput gains, which come at a cost of increased power consumption. Second, we design PollChain MIMO power save, which applies adaptive probabilistic polling to identify the most energy efficient chain setting. Along with PollChain, we also examine sequential-based and threshold-based MIPS alternative solutions. Finally, we implement MIPS along with various legacy 802.11a/b/g and MIMO 802.11n rate adaptation algorithms on an 802.11n standard-compliant platform and evaluate them under various controlled settings and realistic field trials.

The rest of the paper is organized as follows. Section 2 introduces the background on 802.11n power consumption, while Section 3 describes our experimental platform and setup. Section 4 analyzes the factors that contribute to energy consumption of different RF chain settings and Section 5 describes the design of MIPS. Section 6 presents our implementation and evaluation effort, while Section 7 compares our proposed transmitter-side MIPS, with the receiver-side spatial multiplexing power save proposed by 802.11n standard. Section 8 evaluates different directions to save energy in 802.11n systems. Finally Section 9 discusses the related work and Section 10 concludes the paper.

2. 802.11N POWER CONSUMPTION

The IEEE 802.11n standard adopts Multiple-Input Multiple-Output (MIMO) technology to support higher data rates under the same transmit power budget and bit-error-rate performance requirements as a Single-Input Single-Output (SISO) system. MIMO PHY uses multiple transmit (M_t) and receive (M_r) antennas to support two MIMO modes of operation. *Spatial Diversity* transmits a single data stream from each antenna, leveraging the independent fading over multiple antenna links to enhance signal diversity. *Spatial Multiplexing* (SM) transmits independent and separately encoded spatial streams from each of the multiple antennas, to boost performance. However, the performance gains of MIMO come at a cost of increased power consumption, since the circuit complexity of MIMO structures is much higher than that of SISO structures.

The total power consumption P_W along a signal path, includes the power consumption of all the power amplifiers P_{PA} and the power consumption of all other circuit blocks P_c [4] and is given by the equation:

$$P_W = (1 + \gamma) \cdot P_{out} + P_c \tag{1}$$

The first term represents the power consumption at power amplifiers P_{PA}, which depends on the transmit power P_{out}. Factor γ depends on the drain efficiency of the power amplifier. The circuit power consumption P_c is proportional to the number of transmit M_t and receive M_r RF chains and is given by the equation:

$$\begin{aligned} P_c \approx\ & M_t \cdot (P_{DAC} + P_{Mix} + P_{Filtx}) + 2 \cdot P_{Syn} + \\ & M_r \cdot (P_{LNA} + P_{Mix} + P_{IFA} + P_{Filrx} + P_{ADC}) \end{aligned} \tag{2}$$

where P_{DAC}, P_{Mix}, P_{LNA}, P_{IFA}, P_{Filtx}, P_{Filrx}, P_{ADC} and P_{Syn} are the power consumption values for the digital-to-analog converter, mixer, low noise amplifier, intermediate frequency amplifier, active filters at the transmitter and receiver side, analog-to-digital converter, frequency synthesizer respectively. As our study is mainly focused on transmitter-side power save, power consumption is estimated based on the first term of equation 2.

Current commodity 802.11n systems including our experimental platform, support up to three RF chains. In our study we consider SIMO 1x1, MIMO 2x2, MIMO 3x3 settings, which use one, two, three active transmitting RF chains respectively. Note that these settings refer to the number of active chains at the transmitter side and are independent of the active chains on the receiving device. One requirement though, is that the number of RF chains activated on the receiving device is equal or greater than the number of transmitting spatial streams.

3. EXPERIMENTAL SETTING

We next describe the experimental platform and setup for our MIMO energy consumption study.

Experimental Platform: We conduct all our experiments on a programmable 802.11n platform, which uses Atheros AR5416 2.4/5 GHz MAC/BB MIMO

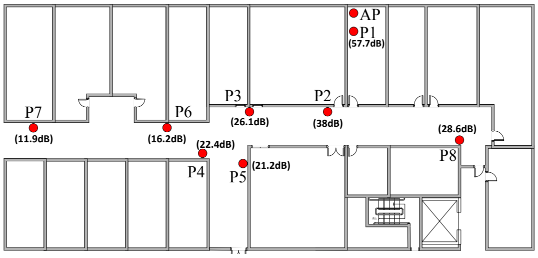

Figure 1: Experimental floorplan.

chipset and supports single stream (SS) and double stream (DS) MIMO modes. Our platform supports up to 130Mbps and 300Mbps data rates for 20MHz and 40MHz channel operations, respectively. Frame aggregation and BlockAck are also available. Frame aggregation amortizes protocol overhead over multiple frames. It packs several data frames, called MAC Protocol Data Units (MPDUs), into an aggregate frame (called A-MPDU). BlockAck acknowledges an A-MPDU frame. Finally our platform has three available antennas.

Our platform provides per-frame control functionalities, which facilitate our study in MIMO power save. Upon receiving a BlockAck, the MIPS module gets feedback including the total number of MPDUs in the transmitted A-MPDU (called as $nFrames$) and the number of MPDUs received with errors (called as $nBad$). If the entire A-MPDU is lost, the number of hardware retries (called as $retries$) is also available. We can then compute *Sub-Frame Error Rate* as $SFER = \frac{nFrames \times retries + nBad}{(retries+1) \times nFrames}$. SFER feedback is used to calculate the throughput performance of each chain setting. Per-antenna SNR information is available to the MIPS module as well. Finally our platform allows for switching transmission chain on a per-AMPDU granularity.

Experimental Setup: We conduct our experiments in a campus setting whose floorplan is presented in Figure 1. Spots P1 to P8 represent different locations where the clients are placed. In all the scenarios, we initiate traffic of 1.5KByte packets from our platform to 802.11n clients. Although transmitter's chain settings can vary during our experiments, on the receiving side the active RF chains are fixed to three. We run the same set of experiments with Linksys WPC600N 802.11a/b/g/n and Airport Extreme wireless adapters using Broadcom chipset. The results presented in the paper are from Linksys adapter.

To measure power consumption at the transmitter side, we connect in 5V DC power input of our platform, a commercial power meter. Its measurement accuracy is +/- 1.5%. Our power meter is also able to log power consumption information on per-second granularity, facilitating our study on runtime energy consumption.

4. MIMO 802.11N ENERGY CONSUMPTION UNDER THE MICROSCOPE

We start our study in MIMO 802.11n energy consumption by asking a simple question. What are the factors that determine the most energy efficient RF chain setting? Interestingly, our case study shows that a) *the fastest chain setting may not be the best*, b) *the most power hungry chain setting may not be the worst*. Taking one step further, our runtime analysis uncovers that the best and worst chain can dynamically change in fine time granularity. We elaborate on these findings in the following sections.

4.1 A Case Study

We first evaluate the performance of different chain settings at a low SNR location (P7 - 11.88dB). The results presented in Figure 2(a) reveal a tradeoff between per-bit energy consumption (E_B) and application data source rate. We observe that the faster in terms of effective goodput MIMO 2x2, 3x3 settings, consume up to 22.4% more energy than 1x1 at low (5Mbps) source rates. On the other hand, the more power hungry MIMO 2x2, 3x3 settings can yield up to 115.3% energy savings over 1x1 for medium (10Mbps) and high (15Mbps, 20Mbps) source rates. The observed behavior results from the interplay between power consumption (P_W) and achieved goodput (G_A), formulated as:

$$E_B = \frac{P_W}{G_A} \tag{3}$$

Differently from legacy 802.11a/b/g, the achieved goodput G_A is not only a function of PHY transmission rate, while in contrast to theory, power consumption P_W does not proportionally increase with the number of active chains.

Goodput Performance: *MIMO diversity gains* can lead to significant goodput gains of MIMO over SISO settings. For a fixed transmission rate R and transmit power P_{out}, increasing the number of active transmit chains M_t results in reduction of Sub-Frame Error Rate.[2] Lower SFER raises *effective goodput* G_E formulated as:

$$G_E(R) = \frac{DATA \cdot nFrames_R \cdot (1 - SFER_R)}{T_{overhead} + \frac{DATA \cdot nFrames_R}{R}} \tag{4}$$

where $DATA$ is the payload size of a MAC-layer frame, and $T_{overhead}$ is the various 802.11n protocol overhead (related to $DIFS$, $SIFS$, $BlockAck$, etc.). For our case study scenario, Figure 2(c) shows that MIMO 3x3 presents goodput gains up to 30.7% over MIMO 2x2 and up to 214.6% over SIMO 1x1 for high volume data

[2] Theoretical upper-bounds for the relation between BER and M_t are illustrated in [4].

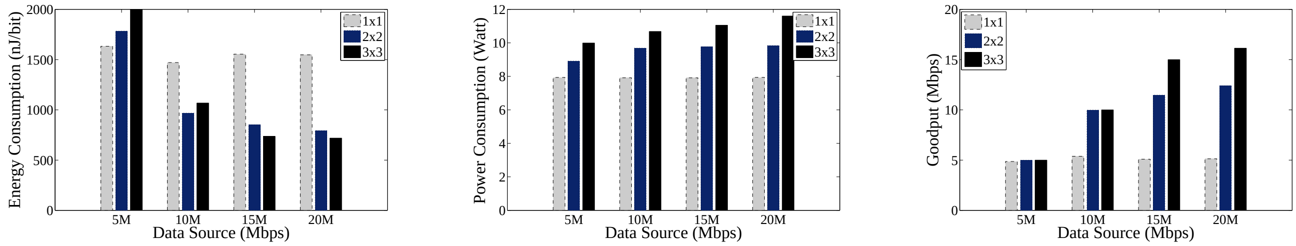

(a) Energy consumption (nJ/bit). (b) Power consumption (Watt). (c) Goodput performance (Mbps).

Figure 2: Energy consumption in low SNR region (location P7 - 11.88dB) for increasing data source rate.

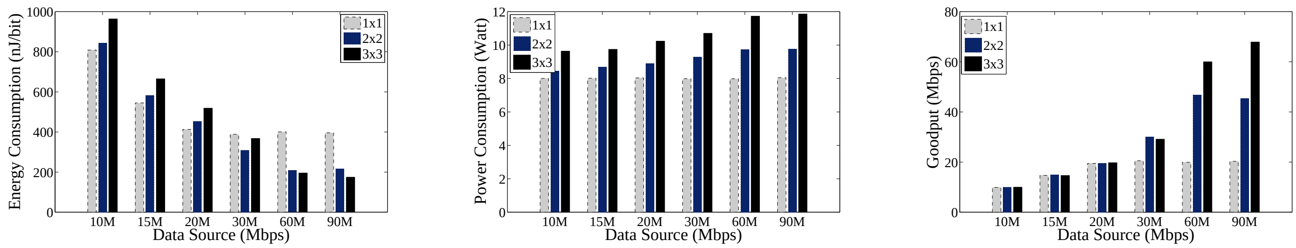

(a) Energy consumption (nJ/bit). (b) Power consumption (Watt). (c) Goodput performance (Mbps).

Figure 3: Energy consumption in medium SNR region (location P5 - 21.23dB) for increasing data source rate.

Rate (Mbps)	Goodput 1x1 (Mbps)	Goodput 2x2 (Mbps)	Goodput 3x3 (Mbps)
13.5SS	9.33	9.37	10.98
27SS	6.09	15.82	18.35
27DS	-	15.68	18.18
40.5SS	0	0.84	17.94

Table 1: Diversity gain in low SNR region (location P7 - 11.88dB).

source. Fixed rate experiments at location P7 presented in Table 1, verify MIMO diversity gains of 3x3, which can support transmission rates three times higher comparing to 1x1. Deactivating RF chains in this low SNR setting, results in an increase up to 115.3%, 15.4% of per-bit energy consumption for SIMO 1x1, MIMO 2x2 over MIMO 3x3 respectively. However, when application data source rate is small enough to be accommodated by less number of active chains, SIMO 1x1 or MIMO 2x2 are the winning settings with savings up to 22.4%. When the *achieved goodput* $G_A(R) = min\{G_E(R), SourceRate\}$ is upper-bounded by the offered source rate, power consumption is the factor to determine the most energy efficient chain setting.

Power Consumption: The power consumed by MIMO circuitry is linearly increased with the number of active transmit chains M_t based on equations 1, 2. However, our measurements at location P7 presented in Figure 2(b), reveal a gap between theory and practice. If we subtract 5.6 Watt which is approximately the power consumed by our platform in sleep mode, the power consumption for 1x1, 2x2, 3x3 is 2.34 Watt, 4.2 Watt, 6 Watt respectively for 20Mbps source rate, which is not proportional to M_t. The additional circuit power consumed by multiple active RF chains is still significant, so although high goodput MIMO settings have idle periods under low volume sources, they are still more power hungry. Figure 2(b) shows that 1x1 always consumes less power than MIMO 2x2 and 3x3 settings.

In summary, *the winning chain setting is always the one with the least number of active chains, which can accommodate the offered source rate*. In the following section, we characterize our findings under different SNR regions and we come up with thresholds that can determine the winning chain setting.

4.2 Understanding the Tradeoff between Power Consumption and Goodput

Our case study reveals an interesting interplay between MIMO gains, source rate and power consumption, which determines the most energy efficient chain setting. This tradeoff varies significantly in different SNR regions. Studying the per-bit energy consumption of a medium SNR location (P5 - 21.23dB) presented in Figure 3(a), we observe that SIMO 1x1 can be up to 25.6% more energy efficient than MIMO settings, for up to 20 Mbps data source rates. However, for higher source rates, the goodput gains up to 101.1% of MIMO

Chain Setting	Energy Consumption (nJ/bit)	Goodput (Mbps)	Power Consumption (Watt)
1x1	76.80	107.30	8.25
2x2	53.51	172.13	9.21
3x3	65.67	170.09	11.17

Table 2: Energy consumption in high SNR region (45.83dB) for high volume data source.

Chain Setting	SNR≤15 (dB)	15<SNR≤23 (dB)	23<SNR≤31 (dB)	SNR>31 (dB)
1x1	Src≤8M	Src≤25M	Src≤60M	Src≤115M
2x2	8M<Src≤15M	25M<Src≤50M	Src>60M	Src>115M
3x3	Src>15M	Src>50M	-	-

Table 3: Classification of winning chains based on SNR (dB) and data source rate (Mbps).

over SIMO (Figure 3(c)), compensate for the power consumed by the additional active chains (Figure 3(b)).

In high SNR regions, MIMO spatial multiplexing high transmission rates are proven to be the fundamental factor to determine the most energy efficient setting. For a high volume data source, higher rates lead to significant goodput gains and as a result to lower per-bit energy consumption as shown in Table 2. Note that 3x3 and 2x2 MIMO settings yield similar goodput performance as our platform supports up to double stream transmission rates. In case three streams and up to 450Mbps rates are available, 3x3 will be the most energy efficient chain setting for high data source rates.

By studying more than 25 wireless links in SNR range [10.18dB, 57.74dB], we classify the winner chain settings based on SNR and application data source rate in Table 3. In summary, as soon as SIMO can accommodate the offered data source rate, it is more energy efficient comparing to MIMO settings. Winner between MIMO 2x2 and 3x3 settings depends on the tradeoff of power consumption and MIMO gains. One should be cautious in applying the above thresholds. First, SNR-BER relations may vary with different propagation environments [5]. Moreover, SNR values reported on the wireless driver can differ significantly among vendors and also depend on the configured transmit power of the device. Finally, the thresholds of Table 3 will be different when more RF chains become available and additional spatial streams are supported (our platform supports up to double stream rates).

Our extensive experiments show that the energy savings upon switching to the appropriate chain setting can be significant. In Table 4 we present the increase of per-bit energy consumption of every chain setting over the others. We observe that the wrong selection of chain setting can result in up to 143% increase in per-bit energy consumption, over the total energy consumed by our platform.[3] In the scenarios where SIMO 1x1 is the

[3]The power consumption of our platform can be up to 12 Watt when all the available RF chains are active.

Chain Setting	1x1 over (%)	2x2 over (%)	3x3 over (%)
1x1	-	1% ∼ 15%	10% ∼ 40%
2x2	4% ∼ 95%	-	2% ∼ 33%
3x3	13% ∼ 143%	7% ∼ 38%	-

Table 4: Increase of per-bit energy consumption of each chain setting over the others.

most energy efficient chain setting, the maximum gains were observed at the maximum data source rate which can be accommodated by this setting. This is because higher source rates reduce the idle time of higher goodput MIMO chain settings and as a result increase their power consumption. In the scenario where MIMO chain settings win, the maximum gains were observed in low SNR regions where diversity gain is crucial.

4.3 Energy Consumption at Runtime

Our study so far has been focused on average performance of different chain settings under various SNR and data source rate scenarios. For a fixed location and data source rate, there is one chain setting which gives the best average energy consumption depending on its achieved goodput and power consumption performance. However, our runtime analysis reveals that, *the most energy efficient chain setting can change in fine time granularity*.

We study the energy consumption in time for all the SNR regions and source rate settings presented in Table 3. The main challenge of these measurements is to synchronize the power meter, which logs the power consumption information in per-second granularity, with our 802.11n platform which logs the goodput performance. To achieve that, we first use scripts to concurrently start/stop logging performance data and we also keep timestamps to verify fine scale synchronization. As it is not possible to conduct experiments for the different chains in parallel, we run them back-to-back in very close time proximity.

A 60 second trace of chains' runtime per-bit energy consumption at location P7 (case study setting) for a high source rate, is presented in Figure 4(a). Although 3x3 is the average winner with energy savings 12.1%, over MIMO 2x2 and 23.8% over SIMO 1x1 as summarized in Table 5, the winning chain can change over time. Specifically, 3x3 wins in 61.7% of the total snapshots, while 2x2 and 1x1 win in 33.3% and 5% of the snapshots respectively. This behavior can be mainly attributed to goodput dynamics presented in Figure 4(c), rather than to power consumption, which is almost constant with standard deviation no greater than 0.07 Watt (Figure 4(b)). For our 60 second trace, there are still 11.7% cases where 2x2 yields higher goodput than MIMO 3x3. We argue that the observed goodput dynamics are mainly attributed to channel and not to rate adaptation behavior, as the results presented

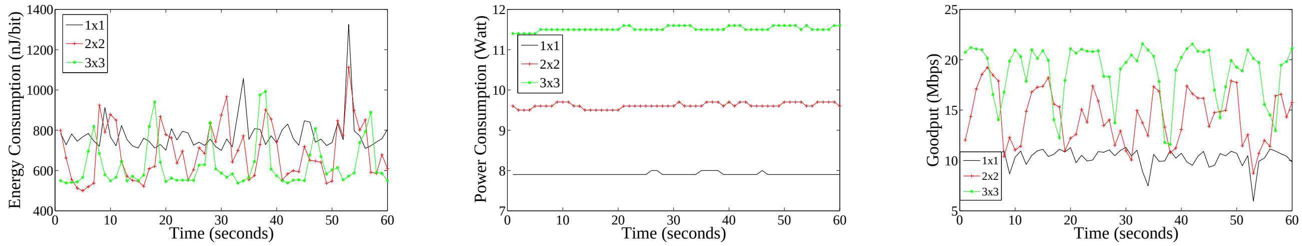

(a) Energy Consumption (nJ/bit). (b) Power Consumption (Watt). (c) Goodput Performance (Mbps).

Figure 4: Energy consumption in time for a 60 second trace in low SNR region (location P7).

Antenna Setting	Energy Consumption (nJ/bit)	Goodput (Mbps)	Power Consumption (Watt)
1x1	777.02	10.28	7.91
2x2	693.16	14.37	9.61
3x3	627.84	18.85	11.52

Table 5: Average energy consumption for the 60 second trace (location P7).

are per-second averages, while rate adaptation dynamics happen at time scale of a few milliseconds.

The key insight learned from the runtime analysis of per-bit energy consumption, is that MIMO power save has to periodically poll different chain settings to discover the most efficient one.

5. DESIGN MIMO POWER SAVE

MIMO power save (MIPS) seeks to identify and set the most energy efficient transmission chain setting at runtime. From one hand, MIPS needs to remain adaptive to dynamic MIMO channel and variable application data source rate. Adaptability though, may result in a large number of frame transmissions at high energy consumption chain settings. In Section 5.1 we first explore a polling-based MIPS scheme. Based on a novel adaptive probabilistic polling mechanism, polling-based design evaluates only the best candidate transmission chain settings. To further amortize polling overhead, it applies a new mechanism to bound 802.11n frame aggregation. In Section 5.2 we study MIPS alternatives, which apply fixed and dynamic thresholds to switch to the most energy efficient chain setting.

5.1 PollChain MIPS

PollChain MIMO power save is based on an *adaptive probabilistic polling* scheme to identify the most energy efficient transmission chain setting at runtime. It keeps track of current chain's and channel's performance, by maintaining moving average throughput, power consumption and SNR statistics. When polling is triggered, PollChain returns a sequence of candidate chains to be polled. The candidate chain settings are evaluated based on their achieved throughput and power consumption performance and finally the most energy efficient transmission chain is selected. The pseudocode of PollChain MIPS is presented in Procedure 1.

PollChain needs to address the following issues. 1) When polling is triggered? 2) What chain setting to poll and in what order? 3) What MIMO-mode, rate to poll? 4) How long polling will last? 5) How polling outcome will be evaluated? We elaborate on the first two issues in Section 5.1.1 and we discuss the last three questions in Section 5.1.2.

5.1.1 *Adaptive Probabilistic Polling*

PollChain MIPS tries to balance between two conflicting objectives. From one hand, it seeks to be adaptive to MIMO channel (Section 4.3) and data source rate dynamics and on the other hand it tries to limit the two-level polling overhead. The first level of polling, involves MIPS switching to a different chain to evaluate its performance. Secondly, upon switching to a new chain setting, rate adaptation algorithm needs to identify the best transmission rate usually using probing mechanisms [6–9]. To ensure adaptability, PollChain does not only use timers but also events to trigger polling. To exclude consistently low performance chain settings from polling, it applies an adaptive probabilistic polling mechanism.

Polling Triggers: PollChain MIPS triggers polling and subsequent chain evaluation using both time- and event-driven mechanisms. Time-driven polling is necessary to update stale energy performance state. To avoid oscillating between different chains, timer should be greater than the polling interval which we discuss in Section 5.1.2 and is set to one second in our prototype. Event-driven polling seeks to capture rapid channel dynamics and is triggered based on both SNR and SFER feedback from the last transmitted A-MPDU. To track actual SNR changes and to subdue random SNR fluctuations, PollChain maintains the weighted average SNR and its standard deviation:

$$\begin{aligned}\overline{SNR}(t) &= (1-\alpha)\cdot\overline{SNR}(t-1)+\alpha\cdot SNR(t)\\ \sigma_{snr}(t) &= (1-\alpha)\cdot\sigma_{snr}(t-1)+\alpha\cdot|SNR(t)-\overline{SNR}(t)|\end{aligned}$$

where $\alpha = \frac{1}{4}$ in our prototype. A sudden increase

Procedure 1 PollChain MIPS: Input (BlockAck, Power), Output (chain)

```
update_stats(BlockAck, SNR, Power, chain)

if !isPoll then
    if poll_event(SNR, BlockAck) then
        (isPoll, chain, chainSeq) ← bruteforce_polling(chain);
    else if poll_timer_expired(timer) then
        (isPoll, chain, chainSeq) ← probabilistic_polling(chain);
    end if

else if isPoll && chain_ewnd_expired() then
    (best_chain, p[chain]) ← poll_outcome(chain, best_chain);
    (chain, chainSeq) ← next_polling_chain(chainSeq);

    if is_chainSeq_empty(chainSeq) then
        chain = best_chain;
        isPoll=false;
        poll_timer_reset(timer);
    end if
end if

return chain
```

in SNR ($SNR(t) - \overline{SNR}(t) \geq 2 \cdot \sigma_{snr}(t)$) accompanied by zero SFER or a fast decrease in SNR ($\overline{SNR}(t) - SNR(t) \geq 2 \cdot \sigma_{snr}(t)$) accompanied by excessive SFER (90% in our prototype), implies channel improvement/degradation respectively and triggers polling. Upon an event, all the available chains will be polled starting from the highest MIMO setting available (3x3 for our platform). By using both loss and SNR statistics, PollChain minimizes premature polling attributed to rate adaptation dynamics. Rate adaptation dynamics, involve rapid change in loss (but not in perceived SNR) induced by false probe transmissions to low goodput rates. In an example scenario illustrated in Table 6, RRAA [9] and SampleRate [8] transmit 28% and 6.5% of the total frames respectively, to very lossy rates.

Probabilistic Polling: PollChain assigns a polling probability $p[i]$ to every available chain setting i. When timer expires, it will poll i with a probability $p[i]$, starting from the settings with the greatest number of available chains (3x3 for our system). PollChain finally adapts $p[i]$ based on the polling outcome. If polling fails at chain i by yielding higher per-bit energy consumption than the current lowest one, the polling probability is decreased by a factor k_{decr} as $p[i] = p[i]/k_{decr}$. Successful polling will increase probability as $p[i] = p[i] \cdot k_{incr}$. Applying MIMD in adapting polling probability can lead to faster convergence to the most energy efficient chain setting under intense MIMO channel and source rate dynamics. The next issue is how to set k_{decr}, k_{incr} factors. k_{decr} reflects the penalty of polling a high energy consumption chain and is set as $k_{decr} = \frac{E_B(i')}{E_B(i)}$ where i is the current most energy efficient chain and i' is the chain where polling failed. To ensure that probabilistic polling remains adaptive even

	SampleRate		RRAA	
Rates (Mbps)	Rate Distr. (%)	SFER (%)	Rate Distr. (%)	SFER (%)
54SS	10	0.07	0	-
54DS	1	0.31	3	0.61
81SS	34.5	4.92	55	2.34
81DS	48	6.84	14	13.52
108SS	5	90.76	28	97.17
108DS	1.5	51.06	0	-
Goodput (Mbps)	61.75		51.72	
SFER (%)	10.11		30.61	

Table 6: Probing loss scenario for RRAA and SampleRate at location P5.

after many consecutive failures, we bound the probability as $p[i] = max\{p[i], P_{min}\}$ where P_{min} is set to 0.1 in our implementation. Finally k_{incr} is set to two, increasing exponentially the probability upon successful polling.

Prioritized Polling: PollChain polls a chain i only if for an expected power consumption value, its highest/loss-free effective throughput (equation 4) can result in a lower per-bit energy consumption than the current best one. So by starting polling from the highest throughput chain settings (3x3 in our platform) as described above, it can avoid polling low-throughput, hence energy inefficient chain settings.

Adaptive probabilistic, prioritized polling is not the only mechanism that PollChain uses to amortize polling overhead as we describe in the following section.

5.1.2 *Cost-Effective Polling*

MIMO power save does not modify existing rate adaptation, but seeks to mitigate its overhead. When MIPS polls a chain setting, rate adaptation re-initiates rate selection to find the new best-goodput rate. The identification of the new best rate, can involve significant number of transmissions at lossy rates as we show in our case study of Table 6. To mitigate rate adaptation's probing overhead, MIPS applies a novel frame aggregation bounding scheme. To further reduce polling overhead, it also sets the polling interval defined as the time that polling lasts, at the minimum required to identify chain's performance. We next elaborate on these issues.

Bound Frame Aggregation: RA's probing frames should be small, to mitigate loss in case of premature probing at low goodut rates. However, protocol overhead is significant for small A-MPDUs. MIPS frame aggregation bounding scheme seeks to balance between these two conflicting objectives. It is based on the observation that there is a point where further increase in A-MPDU size does not lead to significant goodput gains. In Figures 6, 7 we plot the loss-free effective gooput given by equation 4 as a function of aggregation level $nFrames_R$ of our platform's supported SS, DS rates respectively, for 1.5KByte MPDUs.[4] We call the point

[4]As maximum A-MPDU size is 64KBytes, maximum

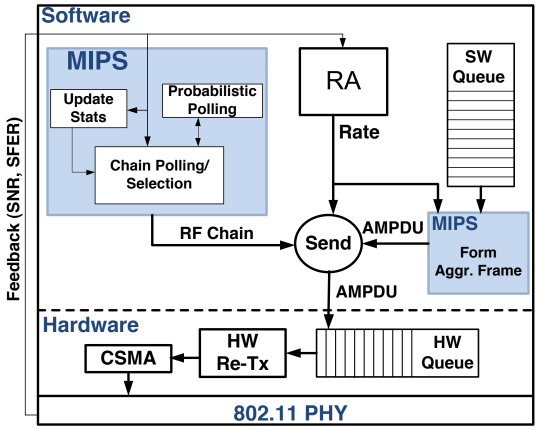

Figure 5: MIPS architecture.

where further increase in $nFrames_R$ results in goodput gain smaller than a given threshold $Thrsh_B$ as *Aggregation Saturation Point.* For a given rate R we bound aggregation level $nFrames_R$ by setting $Thrsh_B = 3\%$.

Polling Interval: Polling should be long enough for rate adaptation to identify and evaluate the best-goodput rate, while it should be short enough to limit overhead from polling an energy inefficient chain setting. We set the polling interval T_P as:

$$T_P = T_{RA} + T_E \tag{5}$$

where T_{RA} is the time rate adaptation algorithm needs to identify the best rate and T_E represents the time required to evaluate the performance of the best rate.

T_{RA} depends on rate adaptation design. RRAA evaluates every rate option for approximately 15ms. So in the worst case scenario under a stable wireless channel, T_{RA} must be set to $255ms$ given that all the available rate options of our platform are 17 for 40MHz channel bandwidth. However, as usually the best-goodput rates among different chains settings can differ up to three rate options as discussed in Section 4.1, we can set $T_{RA} = 45ms$ for RRAA. MiRA uses consecutive one A-MPDU transmissions to identify the best rate. In a loss-free channel, the transmission time of one A-MPDU depends on its size and transmission rate. For the maximum A-MPDU size bounded by our proposed frame aggregation mechanism, A-MPDU transmission time is approximately 4.3ms on average for all the rates. So for MiRA T_{RA} is set to 12.9ms.

MIPS maintains chain's throughput and power consumption performance as $\overline{Thr_A} = \frac{3}{4} \cdot \overline{Thr_A} + \frac{1}{4} \cdot Thr_A$, $\overline{P}_W = \frac{3}{4} \cdot \overline{P}_W + \frac{1}{4} \cdot P_W$ where Thr_A and P_W are the current throughput and power consumption respectively, updated every 20ms in our implementation. PollChain's prototype uses 6 samples to update power consumption

$nFrames_R$ can be no greater than 42 MPDUs.

Procedure 2 Sequential-based MIPS: Input (BlockAck, Power), Output (chain)

update_stats(BlockAck, Power, chain)

if timer_expired() **then**
  if $\overline{srcRate} > \frac{P_W^{chain+}}{P_W^{chain}} \times \overline{Thr}_A^{chain}$ **then**
    (chain) ← next_higher_chain(chain);
  else if $\overline{Thr}_A^{chain} \geq \beta \times \overline{srcRate}$ **then**
    (chain) ← next_lower_chain(chain);
  end if

  reset_timer();
end if

return chain

and throughput moving averages, so T_E is set to $120ms$. We further assess the impact of polling interval in Section 6.2.2.

Chain Evaluation Metric: Finally, PollChain estimates the per-bit energy consumption of a chain setting based on the equation 3 as $E_B = \frac{\overline{P_W}}{\overline{Thr_A}}$. The chain setting with the minimum E_B is selected for transmission.

5.1.3 Putting Everything Together

Figure 5 presents the overall architecture of MIPS. First, per-AMPDU SNR and SFER feedback updates the throughput and SNR statistics of MIPS, while it is used for event-triggered polling as well. The chain polling/selection module will either trigger polling or it will select the most energy efficient chain after polling ends. The polling outcome also updates the polling probability. In addition to MIPS, RA module uses SNR and SFER feedback to select the best-goodput rate. Finally, MIPS form aggregation module builds an A-MPDU from the available frames in software queue. As frame aggregation bound is different for each transmission rate as discussed in Section 5.1.2, RA feeds form aggregation module with the selected rate. Finally, the formed A-MPDU jointly with the selected rate and chain are pushed to hardware queue for transmission.

5.2 Alternative MIPS Designs

In addition to PollChain MIPS, we design and evaluate two alternative design approaches. In *Sequential-based* MIPS, current chain setting may switch to the next higher/lower transmission chain based on dynamic thresholds. *Threshold-based* MIPS utilizes the fixed thresholds presented in Table 3 to determine the best transmission chain setting. Both MIPS alternatives run on top frame aggregation bounding scheme. We next describe our proposals in detail.

Sequential-based MIPS: Sequential-based approach applies dynamic thresholds such that the new selected chain setting can minimize energy consumption. It moves to the next higher (in terms of available chains)

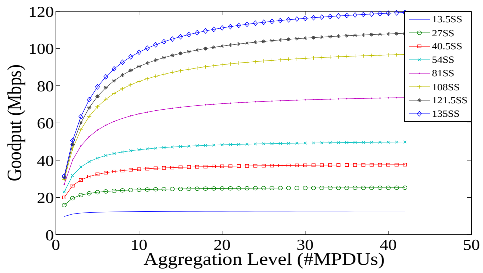

Figure 6: Goodput evolution as a function of aggregation level: SS mode.

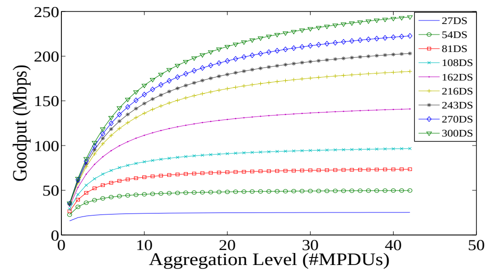

Figure 7: Goodput evolution as a function of aggregation level: DS mode.

setting $i+$, only if its achieved throughput $\overline{Thr}_A^{i+}$ is significantly greater than current chain's throughput, so as to outweigh the additional power consumption. The moving upward condition can be expressed as:

$$\overline{Thr}_A^{i+} > \frac{P_W^{i+}}{P_W^{i}} \times \overline{Thr}_A^{i} \quad (6)$$

However, the throughput performance of chain $i+$ is not known until we poll this chain. To overcome this limitation, sequential-based MIPS estimates the offered application data source rate as $\overline{srcRate} = \frac{3}{4} \cdot \overline{srcRate} + \frac{1}{4} \cdot srcRate$, where $srcRate$ is the current source rate, updated every 20ms in our implementation. The fact that $\overline{srcRate}$ is significantly higher than the current achieved throughput, implies that more chains should be activated to accommodate the high volume data source and it is a call for moving upward to the next higher chain.

Sequential-based MIPS moves to the next lower chain setting, if the achieved throughput of the current chain i can accommodate the offered source rate:

$$\overline{Thr}_A^{i} \geq \beta \times \overline{srcRate} \quad (7)$$

where $\beta = 0.95$ in our prototype. The intuition behind this design choice, is that if the current chain is sufficient for the existing source rate, then the more power efficient next lower chain may be able to accommodate the offered source rate as well.

The moving upward/downwrard conditions are evaluated periodically, every one second in our prototype. The pseudocode of sequential-based MIPS is presented in Procedure 2. The main differences between sequential-based and PollChain MIPS are twofold. Differently from applying polling to evaluate candidate chain settings, sequential-based MIPS decides the next transmission chain based solely on the performance of the current one. Sequential-based approach moves one chain option at a time, while PollChain seeks to jump directly to the most energy efficient chain.

Threshold-based MIPS: Threshold-based approach divides the space in four SNR regions and decides the most energy efficient setting based on the offered source rate, according to Table 3. Source rate $\overline{srcRate}$ is estimated similar to sequential-based approach. The main limitation of threshold-based algorithm is that Table 3 needs to be revised for different environments and hardware.

6. IMPLEMENTATION AND EVALUATION

In this section, we describe MIPS implementation on a programmable AP platform and evaluate its performance using both controlled experiments and field trials.

6.1 Implementation

We have implemented our proposed MIMO power save designs, along with both MIMO 802.11n (MiRA [6], Atheros MIMO RA [7]) and legacy 802.11a/b/g (SampleRate [8], RRAA [9]) rate adaptation algorithms, in the firmware of a programmable 802.11n platform. The first challenge we face, is measuring power consumption P_W in the wireless driver. Although transmit power P_{out} is known, it is not possible to measure power consumed by the various circuit blocks P_c. The measured power consumption in our platform lies in the region [7.8Watt, 8.3Watt] for SIMO 1x1, [8.3Watt, 9.9Watt] for MIMO 2x2 and [9.4Watt, 12.0Watt] for MIMO 3x3, with the highest power values to occur for the highest source rates. We import tables with power consumption, source rate mappings for all the available chain settings. Source rate $\overline{srcRate}$ is estimated as discussed in Section 5.2.

The second challenge is data source rate estimation in the driver. Source rate volume directly affects frames' availability in the software queue and as a result platform's transmit rate. So, we estimate source rate as the transmit rate calculated upon the formation of the A-MPDU and before it moves to hardware queue. Significantly higher transmit rate than achieved throughput, implies that the current chain setting cannot support the offered application source rate and as a result higher chain settings should be evaluated.

Finally, as SNR feedback from the receiver to the

Setting	Thresh. based	Seq. based	1x1	2x2	3X3
Uniform	3.4∼27.7%	0.3∼31.5%	1∼11.5%	3.6∼11.7%	8.3∼25.1%
Bursty	0.9∼244.1%	2.9∼256.6%	2.1∼279.6%	1.7∼25.1%	5.2∼24.8%
Mobility	2.7∼39.4%	7.7∼8.2%	6.1∼9.6%	up to 5.6%	up to 10.1%
F. Trial	0.1%	15.3%	16.5%	9.5%	26.2%

Table 7: Power savings of PollChain over the other algorithms.

transmitter is not currently implemented in the commodity 802.11n devices, we measure the average SNR of all the active chains from the received BlockAck frames. As PollChain MIPS uses SNR as a coarse-grained indicator to trigger polling and not to select the actual transmission chain, the absence of receiver's SNR feedback does not significantly affect its performance.

6.2 Evaluation

In this section we compare the proposed MIPS with fixed chain settings, in controlled static, mobile scenarios and field trials. In controlled settings, we configure our platform at 5GHz band on channel 36 which was interference-free during our experiments. Channel is set at 40MHz. We run controlled experiments at midnight when the building is empty and we repeat them under various traffic scenarios. Uncontrolled field trials represent realistic scenarios where a lot of devices interact in various ways. We first use MiRA [6], which has been shown to outperform both the state of the art legacy 802.11a/b/g and MIMO 802.11n RA solutions, as the underlying RA algorithm. Then we study the interplay between rate adaptation and MIMO power save, by switching to RRAA [9], Atheros MIMO RA [7], SampleRate [8].

Our results show savings of PollChain in more than 90% of the tested scenarios over 1x1, 3x3, sequential-based MIPS and in more than 80%, 85% of the cases over 2x2, threshold-based MIPS respectively. The maximum savings were observed over SIMO 1x1, which consumes up to 3.8 times more energy when it runs over Atheros RA. Savings observed over MIMO 2x2, 3x3 can go up to 25.1%, 26.2% respectively. These gains are still close to maximum 38%, 40% presented in Table 4, despite the fact that in our tested traffic scenarios, source rate is not fixed but changes periodically to favor every available chain setting. PollChain yields 15.3% savings over sequential-based, while it performs similar to threshold-based MIPS in field trials, which is also proven to be a promising direction for MIMO power save. The energy savings of PollChain over the other solutions are summarized in Table 7. These savings are achieved over the total power consumption of our platform, which ranges from 5.6 Watt (sleep mode) to 12 Watt. We are expecting to see significantly higher energy savings in an 802.11n system which supports more spatial streams (up to four allowed by the standard) and RF chains.

6.2.1 MIPS Performance in Various Settings

We first evaluate each algorithm in controlled settings and field trials using MiRA for rate adaptation. We both evaluate linearly increasing and bursty UDP data source rates.

Static Clients for Linear Source Rates: In static settings, we initiate traffic from our 802.11n platform to one 802.11n client placed in various locations which cover all the different SNR regions presented in Table 3. We first set up a traffic scenario where source rate linearly increases, until it reaches the maximum that can be accommodated by the highest chain setting. For example at location P7, source rate starts at 5Mbps and increases 5Mbps every five seconds until it reaches 20Mbps. As every experiment lasts 2 minutes, source wraps around six times.

The per-bit energy consumption for the different algorithms is presented in Figure 8(a). PollChain MIPS is the average winner and can give savings up to 25.1% over fixed rate settings and up to 31.5% over the other MIPS designs. The savings are mainly attributed to high goodput performance of PollChain MIPS as shown in Figure 8(c), achieved at a relatively low power consumption as presented in Figure 8(b). Specifically, PollChain yields similar or higher up to 48% goodput in all the tested locations comparing to sequential and threshold-based MIPS. It also gives goodput gains up to 11.2% over both MIMO 2x2 and 3x3 in four (P1, P3, P6, P7) out of the six tested locations. Goodput gains can be attributed to PollChain opportunistically switching to the highest goodput chain setting, exploiting the transient gains which happen in fine time granularity as discussed in Section 4.3.

Only at location P5, a fixed chain setting (MIMO 2x2) gives 12.9% savings over PollChain MIPS. This is a result of polling overhead. PollChain MIPS transmits 10%, 14% of the total frames at the lower goodput 1x1 and 3x3 settings respectively. This results in a 9.6% goodput decrease over MIMO 2x2, which is proven to be the best goodput chain. PollChain also consumes 0.18 Watt more power than 2x2 because it transmits 14% of the frames at 3x3. However it still gives savings over the other algorithms which can go up to 31.5%.

The alternative sequential and threshold-based MIPS perform better on average than fixed chain settings, but they are not as efficient as PollChain. Their main limitations are twofold. First, the estimation of data source rate $\overline{srcRate}$ as the A-MPDU transmission rate described in Section 6.1, is proportional to the achieved throughput. As a result, sequential-based MIPS may prematurely trigger moving downward based on equation 7. For example at location P5 (21.2dB), it transmits 95% of the total frames at 1x1, despite the fact that

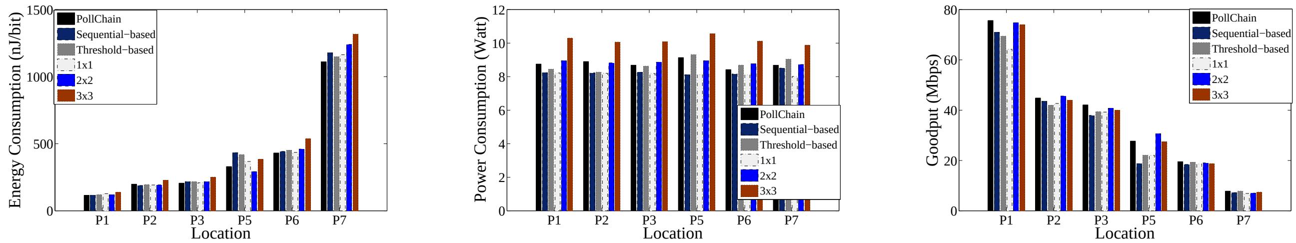

(a) Energy consumption (nJ/bit). (b) Power consumption (Watt). (c) Goodput performance (Mbps).

Figure 8: Per-bit energy consumption in various locations for linear data source.

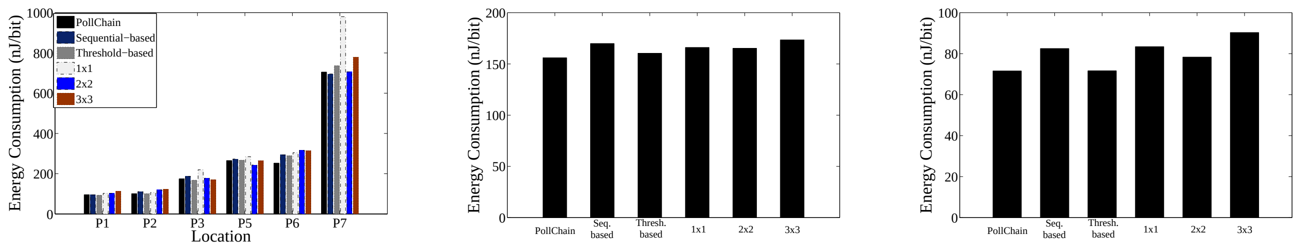

(a) Per-bit energy consumption in various locations for bursty data source. (b) Per-bit energy consumption in a mobility scenario. (c) Per-bit energy consumption in a field trial scenario.

Figure 9: Per-bit energy consumption at various static, mobile and field trial settings.

all chain settings can be the winners for different source rates according to Table 3. Inaccuracies of the source rate estimation have higher impact on threshold-based MIPS under different RA designs as we discuss in Section 6.2.2. An additional limitation for threshold-based MIPS is the use of fixed thresholds, which can only indicate the chain with the best average and not runtime performance. At location P5, threshold-based approach transmits 18.6%, 62.3% at SIMO 1x1 and MIMO 3x3 respectively, which yield lower goodput than 2x2 because of MIMO channel dynamics.

Static Clients for Bursty Source Rates: In our bursty data source scenario, the rate volume switches from very low, to medium, to very high periodically. PollChain MIPS is also proven to be the most efficient chain setting with savings up to 39% over fixed chain settings and up to 16.1% over MIPS alternatives. The highest energy savings on average are observed at the low SNR location P6. The savings can be attributed to goodput gains of PollChain MIPS over all the other settings, which can go up to 36.6% of fixed chain settings and up to 28.5% over MIPS alternatives. These goodput gains come at a relative low power consumption. PollChains consumes up to 1.07 Watt more power only over the low goodput 1x1, and sequential-based MIPS.

Mobile Clients: In our mobile setting, we walk a client from location P1 to P7 and back at approximately constant speed of 1m/s. Our experiments with both linear and bursty data sources reveal savings up to 10.1% over fixed chain settings and up to 39.4% over MIPS alternatives. The per-bit energy consumption for a linear source rate is presented in Figure 9(b). The savings for this scenario are up to 10.1% over fixed chain settings and up to 8.2% over MIPS alternatives. PollChain's event-driven polling described in Section 5.1.1, is triggered two times on average in our low pedestrian speed scenario. As a result savings are limited. We are expecting our mechanism to be more efficient under fast mobility cases, as vehicular networking scenarios.

Field Trials: We conduct uncontrolled field trials under realistic scenarios, where various sources of dynamics coexist in a complex manner. In our field trial, we use three static 802.11n clients at locations P1, P3, and P8. We also set up the bursty source rate scenario described above, from our 802.11n platform to the static clients. We finally configure our platform at 5GHz and set channel bandwidth to 40MHz. During our experiments, the physical environment was highly dynamic as people walk back and forth. Figure 9(c) shows that PollChain MIPS yields 16.5% savings over SIMO 1x1, 9.5% over MIMO 2x2, 26.2% over MIMO 3x3, and 15.3% over sequential-based MIPS. Threshold-based performs similar to PollChain MIPS and seems to have a lot of potentials to save energy when it operates over a robust RA algorithm. However, the fixed SNR-source rate thresholds of Table 3 can change under different environments and hardware. So threshold-based MIPS should be enhanced with learning mechanisms,

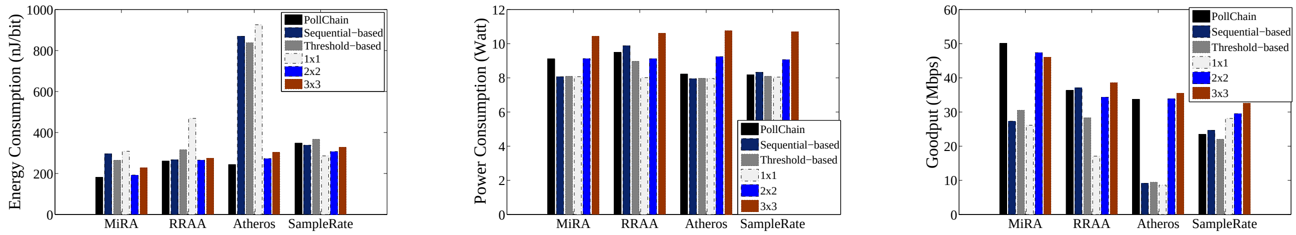

(a) Energy consumption (nJ/bit). (b) Power consumption (Watt). (c) Goodput performance (Mbps).

Figure 10: Per-bit energy consumption at location P4 for various rate adaptation algorithms.

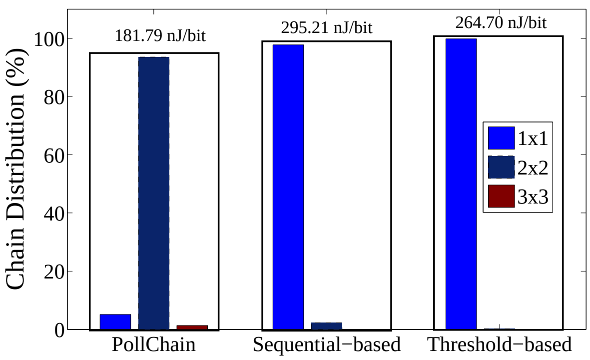

Figure 11: Chain distribution (%) for MiRA at location P8.

which will update them dynamically. This process may require polling mechanisms as well.

6.2.2 *The Interplay between MIPS and RA*

The performance of the underlying rate adaptation algorithm has a significant impact on the effectiveness of MIMO power save. In this section we study MIPS over both popular legacy 802.11a/b/g (SampleRate, RRAA) and MIMO 802.11n (MiRA, Atheros MIMO RA) RA solutions. Starting from legacy RAs, RRAA [9] estimates rate's loss ratio for a short-term time window and opportunistically moves by one rate option at a time, based on fixed thresholds. On the other hand, SampleRate [8] tries to directly switch to the rate that currently gives the lowest average frame transmission time. MIMO Rate Adaptation (MiRA) [6] which has been used so far in our experiments, opportunistically zigzags between MIMO modes, trying to discover the best-goodput rate. Finally Atheros MIMO RA [7] selects the best-goodput rate based on SFER statistics, among a set of candidate rates upper-bounded by a $maxRate$.

To illustrate our findings, we use the medium SNR location P4 and a bursty source rate, for our case study. The per-bit energy consumption presented in Figure 10(a), shows that PollChain MIPS is significantly more energy efficient than fixed settings and MIPS alternatives, for all RAs except SampleRate. The savings can go up to 279.6% over fixed chain settings, up to 256.6% over MIPS alternatives and are observed for Atheros MIMO RA. By studying the power consumption in Figure 10(b) and goodput performance in Figure 10(c), we observe that these high energy savings are attributed to goodput gains of PollChain over SIMO 1x1 and MIPS alternatives. Specifically, PollChain yields goodput gains up to 292.2% over 1x1, 255.5% over threshold-based MIPS and 268.9% over sequential-based MIPS for the Atheros MIMO RA scenario. To explain the observed performance, we analyze the loss statistics and rate distributions of the different chain settings. We find that frame losses make Atheros MIMO RA to transmit at very low rates in SIMO 1x1 setting. In fact, Atheros' rate upper-bound $maxRate$ is set to the lowest transmission rate 13.5Mbps, for the majority of the frame transmissions. As a result, 1x1 transmits 58% of its frame at 13.5Mbps, which leads to very low goodput performance. Threshold-based and sequential-based MIPS suffer from exactly the same problem. When they switch their transmission chain to 1x1, their throughput performance drops significantly to less than 10Mbps. As source rate $\overline{srcRate}$ is proportional to throughput performance, MIPS alternatives will be trapped to low goodput SIMO 1x1. In fact, they transmit more than 99.9% of their frames at 1x1. PollChain MIPS on the other hand, is able to overcome this limitation by polling and identifying the most energy efficient chain setting. Note that PollChain gives savings up to 24.5% over MIMO 2x2, 3x3 as well, by performing similar in terms of goodput but with much less, up to 2.5Watt power consumption.

PollChain MIPS gives similar or better goodput with relatively low power consumption for both MiRA and RRAA. For example when MIPS runs over MiRA, PollChain transmits 93.5% of its frame at MIMO 2x2 which is on average the most energy efficient chain setting, while the MIPS alternatives give sub-optimal chain distributions as shown in Figure 11. Its key advantage is the ability to periodically identify the most energy efficient chain setting and to overcome limitations of the underlying RAs.

PollChain is less energy efficient than fixed chains and sequential-based MIPS, when it runs over SampleRate. This is attributed to SampleRate's slow convergence to the best-goodput rate, which results from three design decisions.[5] a) It uses exponentially weighted moving average for updating rates' performance. b) It allows for rate change only every 2 seconds or upon four consecutive losses. c) It does not sample for 10 seconds rates which have faced four successive failures. In PollChain implementation over SampleRate we set the polling interval T_P to 150ms to mitigate overhead. This time window is proven very small for SampleRate to identify the best-goodput rate and leads to goodput degradation of PollChain over the other algorithms as shown in Figure 10(c). This finally results in higher energy consumption of PollChain over fixed chains and sequential-based MIPS. We can overcome this limitation by tuning the polling interval appropriately.

7. TRANSMITTER VS. RECEIVER-SIDE MIPS

The IEEE 802.11n standard [2] proposes Spatial Multiplexing (SM) Power Save, which differently from MIPS is a receiver-side solution. In the *static SM power save* mode, after the client notifies the AP that is now operating in SISO mode, it turns off all but a single chain, becoming essentially equivalent to an 802.11a/g client. In the *dynamic SM power save* mode, the client turns off all but one of its chains as well, but it can rapidly enable its additional chains when it receives a frame that is addressed to it. In this mode of operation, the AP typically sends a request-to-send (RTS) frame to the client, to activate its chains, prior to sending it a data frame. After transmission is completed the client can switch back to SISO mode. A comparison between MIPS with SM Power Save raises two critical questions. Transmitter or receiver-side power save has more potentials for significant energy savings? How these different approaches can work in concert to maximize system's energy savings?

MIPS vs. SM Power Save: Our experimental results show that transmitter-side MIPS has more potentials for high energy savings than receiver-side SM power save. First, transmitting frames from many active RF chains drains more power than receiving the same volume of traffic from the same number of active RF chains. The power consumption measurements for downlink and uplink traffic of a low volume source rate (25Mbps) presented in Table 8, reveal an increase up to 18.6% in power consumption[6] on the transmitting over receiving side. Second, the additional protocol overhead required from SM power save, reduces the achieved goodput and as a result increases system's per-bit energy consumption based on equation 3. Enabling RTS/CTS required by dynamic SM power save, at location P1 under high volume UDP source, for a client which supports up to 270Mbps rates, results in 151.2Mbps goodput. By turning off RTS/CTS we observe 19% goodput gain. These gains will be higher when higher rates become available. Finally, MIPS needs to be implemented only in one communicating side (transmitter) and does not make any assumptions about receiver-side settings. SM power save on the other hand, needs to communicate to the transmitter its current active RF chains.

	Power Cons. (1x1) (Watt)	Power Cons. (2x2) (Watt)	Power Cons. (3x3) (Watt)
Transmitter	8.13	8.65	10.00
Receiver	8.15	8.19	9.31

Table 8: Transmitter vs. Receiver Power Consumption.

Note that transmitter-side MIPS, is not only important for battery-constrained 802.11n portable devices, but also for 802.11n Access Points. With dual radios and 3x3 MIMO per radio, 802.11n APs consume up to 18 Watt, comparing to legacy 802.11a/b/g APs which have a maximum power draw of less than 13 Watt. The high power budget requires new technologies as IEEE 802.3at [3] to accommodate 802.11n power needs.

MIPS + SM Power Save: MIPS design principles can also be used for receiver-side power save. 802.11n describes only the SM power save mechanism, but the events which will trigger switching from SISO to MIMO and vice versa are still unspecified by the standard. Moreover, SM power save differently from MIPS switches only between two states, MIMO and SISO. However, as both throughput and SNR feedback is available at the receiver, SM power save can apply MIPS mechanisms and switch to the chain setting which yields the lowest per-bit energy consumption. Changes on the available active chains at the receiver side limit the number of chain settings that the transmitter can choose from, to send frames (e.g. a SISO receiver cannot accommodate a MIMO transmitter). So MIMO power save can be extended as an optimization problem, which seeks to identify the most energy efficient pair of chain settings. Addressing this problem is part of our future work.

8. EVALUATING ALTERNATIVES FOR 802.11N POWER SAVE

Our study so far, has been focused on energy consumption of different RF chain configurations. The available active chains can be used for spatial diversity

[5] Based on SampleRate's MADWiFi implementation.

[6] To calculate power increase 18.6%, we first normalize power measurements by 5.6 Watt which is our platform's power consumption in sleep mode.

MIMO Mode	Power Cons. (Watt) 108Mbps	Power Cons. (Watt) 81Mbps	Power Cons. (Watt) 54Mbps	Power Cons. (Watt) 27Mbps
SS	11.77	11.71	11.45	11.54
DS	11.80	11.60	11.51	11.60

Table 9: Power consumption for SS and DS rates at location P1, for a 3x3/5GHz/40MHz setting.

	20MHz			40MHz		
Source (Mbps)	Energy Cons. (nJ/bit)	Goodput (Mbps)	Power Cons. (Watt)	Energy Cons. (nJ/bit)	Goodput (Mbps)	Power Cons. (Watt)
30M	339. 02	28.70	9.73	335.59	29.47	9.89
130M	114.96	98.81	11.36	83.53	130.00	10.86
180M	115.69	98.36	11.38	65.67	170.09	11.17

Table 10: 20MHz vs. 40MHz at location P1, for a 3x3/5GHz setting.

or spatial multiplexing transmissions. Moreover, they can be configured to either 20MHz or 40MHz channels. In this section we examine the impact of MIMO mode and channel bandwidth in 802.11n power consumption.

MIMO Mode: Spatial diversity or single-stream MIMO mode, transmits the same data stream from all the available chains to enhance signal diversity, while spatial multiplexing transmits multiple, independent data streams to boost transmission rate. The number of data streams that can be used for transmission, is upper-bounded by the number of active RF chains. To identify the impact of MIMO mode in 802.11n power consumption, we measure the power drained by the same PHY transmission rates, implemented either as single stream (SS) or double stream (DS) rates (e.g. 54SS, 54DS). Table 9 shows a representative scenario at location P1, where our platform is configured at MIMO 3x3, 5GHz band while channel bandwidth is set to 40MHz. Application data source is set equal to the effective goodput of each transmission rate at location P1. We observe that the pairs of SS and DS rates yield similar power consumption performance, given that standard deviation varied from 0.06 to 0.2 Watt during our experiments. Even across an individual MIMO mode, power consumption is similar across different transmission rates. The small variations, can be attributed to different application data source rates used for the different transmission rate settings.

Channel Bandwidth: The wider 40MHz channels supported by 802.11n standard, can provide significantly higher rates (up to 300Mbps for our platform), comparing to 20MHz channels (up to 130Mbps for our platform). Higher transmission rates may yield higher achieved goodput G_A, and as a result significant energy savings as shown by our experiments. In Table 10 we present the energy consumption performance of 20MHz and 40MHz channels at a MIMO 3x3 setting, for a low, medium and high application data source rate. Rate adaptation is set to MiRA. At the maximum data source rate (180Mbps), 20MHz configuration consumes 76.2% more energy comparing to 40MHz. These energy savings result from 72.9% goodput gains of 40MHz over 20MHz channel bandwidth. Interestingly, our measurements also show that wider 40MHz channels do not come at a cost of increased power consumption. For all the application data source rate scenarios presented in Table 10, 40MHz consumes similar or less power than 20MHz.

In summary, power consumption is independent of the selected MIMO mode and 40MHz channels consistently give energy savings over narrower 20MHz configurations in our tested scenarios. As a result, MIMO power save, which selects the most energy efficient RF chain setting, seems to be the most promising direction to save energy in MIMO 802.11n wireless systems.

9. RELATED WORK

Energy consumption has been widely studied in legacy 802.11 wireless interfaces [12–17], but it still remains unexplored for the MIMO 802.11n systems. An initial effort in identifying factors as channel bandwidth, transmit power, transmission rates, antennas, MIMO streams, that contribute to 802.11n power consumption using commodity hardware, is presented in [10]. Similar to our study, authors measure power consumption for SISO and MIMO chain settings. Differently from our study, they calculate per-bit energy consumption as a function of PHY transmission rate and not as a function of achieved goodput performance. However, our experiments show that supported transmission rates may determine the winning chain setting only in high SNR regions. It is the interplay between SNR, MIMO gains and offered source rate which determines the achieved goodput and as a result the most energy efficient chain setting.

To our knowledge, there are only few design proposals for 802.11n MIMO power save. IEEE 802.11n standard [2] allows for Spatial Multiplexing (SM) power save, which has the client to interchange between MIMO and SISO modes. The events that will trigger mode switching are unspecified by the standard. Differently from MIPS, SM power save is a receiver-side MIMO power save solution. In [11] MIMO Power save is formulated as an optimization problem, which seeks to minimize the per-bit energy consumption for a given effective data rate requirement. However, this approach requires feedback from the receiver (PHY layer channel matrix, SNR or noise level), which is not available in commodity 802.11n systems.

Theoretical results illustrate the tradeoff between MIMO gains and power consumption as well. In [4], authors show that high energy saving is possible in MIMO, only for transmission distances larger than a given threshold. In a real system the tradeoff can be significantly different. First, application data source rate plays a key role in utilizing the MIMO gains. Second, power consumption is not proportional to the number of active transmit chains.

10. CONCLUSION

In this paper, we use an 802.11n standard-compliant programmable platform, to study energy consumption in 802.11n-enabled devices. The key insight gained, is that only the number of RF chains needed to roughly accommodate the offered data source rate should remain active. However, as the MIMO channel and the offered source rate can change in fine time granularity, MIMO power save should opportunistically poll the available chain settings to re-evaluate their performance. To this end we design PollChain MIPS, which uses an adaptive probabilistic polling scheme to identify the most energy efficient transmission chain setting at runtime. By measuring and not predicting the performance of different chain settings, PollChain is able to overcome any limitations of the underlying rate adaptation algorithms. We expect that MIPS takes one step further in incorporating MIMO in power-constrained mobile devices.

11. REFERENCES

[1] 802.11n to Feature in 87% of Wi-Fi-Enabled Smartphones in 2014. *http://www.abiresearch.com, January 2010.*

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

[3] IEEE Std. 802.3at-2009: DTE Power Enhancements Task Force *http://www.ieee802.org/3/at/, 2009.*

[4] S. Cui, A. J. Goldsmith, A. Bahai. Energy-efficiency of MIMO and Cooperative MIMO in Sensor Networks. *IEEE Journal on Selected Areas in Communications, vol. 22, no. 6, pp. 1089-1098, 2004.*

[5] A. Goldsmith, S. A. Jafar, N. Jindal, S. Vishwanath. Capacity Limits of MIMO Channels. *IEEE Journal on Selected Areas in Communications, vol. 21, no. 5, pp. 684 - 702, 2003.*

[6] I. Pefkianakis, Y. Hu, S. H.Y. Wong, H. Yang, S. Lu. MIMO Rate Adaptation in 802.11n Wireless Networks. *ACM MOBICOM'10.*

[7] M. Wong, J. M. Gilbert, C. H. Barratt. Wireless LAN using RSSI and BER Parameters for Transmission Rate Adaptation. *US patent* 7,369,510, 2008.

[8] J. Bicket. Bit-rate Selection in Wireless Networks. *MIT Master's Thesis*, 2005.

[9] S. H. Wong, H. Yang, S. Lu, V. Bharghavan. Robust Rate Adaptation for 802.11 Wireless Networks. *ACM MOBICOM'06.*

[10] D. Halperin, B. Greenstein, A. Sheth, D. Wetherall. Demystifying 802.11n Power Consumption. *USENIX HotPower'10.*

[11] H. Yu, L. Zhong, A. Sabharwal. Adaptive RF Chain Management for Energy-Efficient Spatial-Multiplexing MIMO Transmission. *ACM ISLPED'09*

[12] N. Balasubramanian, A. Balasubramanian, A. Venkataramani. Energy Consumption in Mobile Phones: A Measurement Study and Implications for Network Applications. *ACM IMC'09.*

[13] J. Liu and L. Zhong. Micro Power Management of Active 802.11 Interfaces. *ACM MobiSys'08.*

[14] A. Rahmati and L. Zhong. Context-for-Wireless: Context-Sensitive Energy-Efficient Wireless Data Transfer. *ACM MobiSys'07.*

[15] Y. Agarwal, R. Chandra, A. Wolman, P. Bahl, K. Chin, R. Gupta. Wireless Wakeups Revisited: Energy Management for VoIP over Wi-Fi Smartphones. *ACM MobiSys'07.*

[16] D. Qiao, S. Choi, A. Jain, K. G. Shin. MiSer: An Optimal Low-Energy Transmission Strategy for IEEE 802.11a/h. *ACM MOBICOM'03.*

[17] A.W. Papers. Power Consumption and Energy Efficiency Comparisons of WLAN Products. *Technical Report, Atheros Inc., 2003.*