

UC Santa Cruz

UC Santa Cruz Previously Published Works

Title

A Practical Approach to Rate Adaptation for Multi-Antenna Systems

Permalink

<https://escholarship.org/uc/item/7zp1z8j6>

Authors

Garcia-Luna-Aceves, J.J.

Nguyen, D.

Publication Date

2011-10-17

Peer reviewed

A Practical Approach to Rate Adaptation for Multi-Antenna Systems

Duy Nguyen and J.J. Garcia-Luna-Aceves
Department of Computer Engineering
University of California, Santa Cruz
Santa Cruz, CA 95064, USA
{duy,jj} @ soe.ucsc.edu

Abstract—Multi-antenna systems can provide greater throughput and range coverage than traditional single antenna systems. A key aspect of exploiting this new physical layer (PHY) is rate adaptation, which consists of finding the best rate for sending data packets. Unlike rate adaptation in single antenna systems, nodes have many choices apart from adapting different modulation types, and these choices include using spatial multiplexing or transmit diversity, types of guard intervals, and channel width. We present an evaluation and implementation of a new rate adaptation scheme for multi-antenna systems applicable to off-the-shelf wireless cards. Our rate adaptation scheme, rate adaptation for multi-antenna systems (RAMAS), is simple and practical, and eliminates the complexity of the rate adaptation approaches proposed for IEEE 802.11n in the recent past. Extensive experimental evaluation is used to show that RAMAS performs consistently better than many current IEEE 802.11n rate adaptation schemes with much less complexity, and that RAMAS is especially efficient in multi-user and interference-laden environments.

I. INTRODUCTION

The certification of the IEEE 802.11n standard [1] has paved the way for widespread adoption of the multi-antenna systems. Multiple input multiple output (MIMO) technology in IEEE 802.11n and IEEE 802.16 is likely to play an important role in future broadband wireless networks, given the performance advantages it can offer. For example, IEEE 802.11n MIMO systems can have more than 10 times throughput improvement over traditional IEEE 802.11 a/g systems by using four antennas and exploiting multi-path fading to achieve a significant throughput gain. One of the most important components of the IEEE 802.11n physical layer (PHY) is the ability to support rate adaptation, which consists of attempting to find the best rates for sending data packets in the noisy and interference-prone neighborhoods of the intended receivers.

Different bit rates in IEEE 802.11 are attained by means of different modulations, which are obtained by varying the amplitude, phase, and frequency of the waveform. Higher bit rates allow more data to be transmitted on high quality links but suffer from low throughput on lossy links; lower bit rates allow lower packet loss even in the presence of lossy links. With IEEE 802.11n, the availability of spatial multiplexing, transmit diversity, different types of guard intervals, and channel widths complicates the task of rate adaptation considerably. Due to this complexity, the IEEE 802.11n standard [1] defines a

modulation coding scheme (MCS) index with values that range from 1 to 76 to simplify rate selections for up to a maximum of four antennas. However, even with this simplification, selecting among the 76 MCS index values available in IEEE 802.11n is much more complex than using a dozen sampling rates as it is done in SISO 802.11 wireless systems.

Section II reviews some background and related work, with our focus limited to the most current and representative results published in the open literature. As our brief overview of prior work reveals, rate adaptation schemes can be classified based on whether explicit or implicit feedback to the transmitters is used. Explicit feedback requires the receiver to explicitly communicate the channel condition on the receiver's side back to the sender. Implicit feedback looks at acknowledgment (ACK) packets or other channel information (i.e., received signal strength indicator (RSSI)) to infer the channel conditions on the receiver's side.

Section III describes a new approach to rate adaptation in the context of IEEE 802.11n, which we call RAMAS (rate adaptation for multi-antenna systems). To simplify the job of adapting rates for IEEE 802.11n, RAMAS categorizes different types of modulation into a *modulation group*, and then categorizes spatial multiplexing, transmit diversity, types of guard interval, and channel width into an *enhancement group*. RAMAS then adapts these two groups concurrently. The combination of the modulation and the enhancement group is mapped back to the MCS for adapting rates in IEEE 802.11n. RAMAS is fully compatible with current and existing WiFi networks, as we have implemented it using open-source software and off-the-shelf wireless cards.

Section IV presents the results of extensive empirical tests comparing RAMAS with ARF [2], ATH9K [3], MISTREL [4], and MiRA [5], which are the leading approaches for rate adaptation reported to date for multi-antenna systems. The scenarios used in our study include indoor experiments, outdoor forest testing, multi-user environments, and limited mobility. In all cases, RAMAS outperforms the other schemes, even though its design is far simpler.

II. BACKGROUND

A. Explicit feedback approaches

Explicit feedback approaches can be viewed as receiver-driven rate adaptation, because the receiver dictates the rate

that should be used. The receiver obtains its current channel condition and relays this information back to the sender.

Receiver Based Auto-Rate (RBAR) [6] selects the bit rate based on signal-to-noise ratio (S/N) measurements. Upon processing a request to send (RTS) packet, the receiver calculates the highest bitrate and piggybacks this selected bit rate on the clear to send (CTS) packet. The limitation is that RBAR needs an accurate mapping between S/N value rates for different hardware.

Cross-layer wireless bit rate adaptation [7] uses confidence information from the physical layer to estimate the bit error rate (BER). The receiver sends a BER estimate to the sender in a link-layer feedback frame. The sender then uses this per-frame BER feedback to select the best transmit rate. The limitation with this approach is that it requires sending link-layer feedback frames at the lowest bit rate and in a reserved time slot.

Frequency-Aware Rate Adaptation (FARA) [8] computes the optimal choice of bit rate on each sub-band and sends it back to the sender in ACK packets. A sequence number is also added to acknowledgment (ACK) and data packets to prevent sender and receiver from going out of synchronization. However, the need to modify ACK packets and to synchronize makes this approach incompatible with the IEEE 802.11 standard.

B. Implicit feedback approaches

Implicit feedback approaches can be viewed as sender-driven rate adaptation, given that the sender adapts its rate by inferring the channel conditions on the receiver side.

The Automatic Rate Fallback (ARF) scheme is one of the earliest rate control algorithms designed for WaveLAN-II [2]. ARF operates at a default rate of 2 Mbps, and when it encounters a second missed acknowledgement (ACK) of data packets, the rate falls back to 1 Mbps. A counter is used to track the number of good and bad acknowledgement packets. When either the time expires or the number of successive good acknowledgements reaches 10, ARF upgrades the rate. The limitation of ARF is that it was designed for just a few data rates and does not work well with current IEEE 802.11 implementation.

The Sample rate adaptation (SRA) algorithm [9] begins by sending data at the highest bit rate. Upon encountering four successive failures, SRA decreases the bitrate until it finds a usable bitrate. At every tenth data packet, SRA picks a random bitrate that may do better than the current one. One drawback of this scheme is that it takes longer for it to return to an optimal rate after sudden changes in the environment. Minstrel [10], a widely deployed and popular Linux rate control, is an improved version of SRA, and takes into account the exponential weighted moving average statistics for sorting throughput rates. However, Minstrel still spends 10 percent of transmitted frames trying random rates when the current rate is working perfectly.

Robust Rate Adaptation Algorithm (RRAA) [11] uses a short-term loss ratios to opportunistically adapt the rates. Like

CARA [12], it employs an RTS filter to prevent collision losses from rate decrease. Enabling an RTS filter upon encountering failed transmissions might not work as well as simply transmitting the data at lower rates. Besides, this adds additional control overhead. Due to the nature of air interface, it is complex and difficult to predict the cause of packet collisions.

C. Explicit vs. Implicit approaches

In addition to incurring additional overhead by requiring the receiver to relay channel-state information back to the sender, the drawback with the explicit approach to rate adaptation is the possibility of using stale feedback when channel conditions change during data transmissions. If the channel coherence time is very short, the receiver may not be able to relay accurate channel-state information to the sender. In the worst-case scenario, the receiver ends up sending feedback information to the sender continuously, which occupies the channel with feedback packets and prevents the sender from transmitting data. However, explicit feedback works well if the channel conditions do not change frequently.

Indeed, each explicit and implicit approach has its own advantages and disadvantages. In RAMAS, we show that the implicit approach can be very effective and that ACK packets are all that is needed for adapting rates robustly.

D. What is New in MIMO IEEE 802.11n and Current Approaches?

The difference between IEEE 802.11n and its predecessor is the new enhancements for higher throughput such as multi-antenna transmissions (MIMO), shorter guard intervals, wider channel width, frame aggregation, and block ACKs.

The use of MIMO systems provides great performance improvements due to array gain, diversity gain and spatial multiplexing gain. Array gain refers to an increase in average receive SNR due to a coherent combining effect. Apart from depending on the number of transmit and receive antennas, channel knowledge at the transmitter and receiver is required in order to yield this array gain [13]. Transmit diversity is a technique used to overcome fading in wireless links and increase robustness; it allows transmitting the signal over multiple fading paths for redundancy (e.g., [14], [15]).

Spatial multiplexing allows nodes to send multiple streams of data independently and concurrently; hence, it can attain a significant increase in throughput. This is due to MIMO channels achieving a linear increase in capacity [16]–[18].

TABLE I: Modulation Group

m_index	Modulation	Coding Rate
0	BPSK	1/2
1	QPSK	1/2
2	QPSK	3/4
3	16-QAM	1/2
4	16-QAM	3/4
5	64-QAM	2/3
6	64-QAM	3/4
7	64-QAM	5/6

TABLE II: Enhancement Group

e_index	No. of Spatial Streams	Guard Interval	Channel Width
0	1	800 ns	20 MHz
1	1	400 ns	20 MHz
2	1	800 ns	40 MHz
3	1	400 ns	40 MHz
4	2	800 ns	20 MHz
5	2	400 ns	20 MHz
6	2	800 ns	40 MHz
7	2	400 ns	40 MHz
8	3	800 ns	20 MHz
9	3	400 ns	20 MHz
10	3	800 ns	40 MHz
11	3	400 ns	40 MHz
12	4	800 ns	20 MHz
13	4	400 ns	20 MHz
14	4	800 ns	40 MHz
15	4	400 ns	40 MHz

TABLE III: MCS Index Mapping

MCS Index	Modulation (M) and Enhancement (E)	Guard Interval (ns)	Channel Width (MHz)		
0	M0 :: E0-E3	800	400	20	40
1	M1 :: E0-E3	800	400	20	40
2	M2 :: E0-E3	800	400	20	40
3	M3 :: E0-E3	800	400	20	40
4	M4 :: E0-E3	800	400	20	40
5	M5 :: E0-E3	800	400	20	40
6	M6 :: E0-E3	800	400	20	40
7	M7 :: E0-E3	800	400	20	40
8	M0 :: E4-E7	800	400	20	40
9	M1 :: E4-E7	800	400	20	40
10	M2 :: E4-E7	800	400	20	40
11	M3 :: E4-E7	800	400	20	40
12	M4 :: E4-E7	800	400	20	40
13	M5 :: E4-E7	800	400	20	40
14	M6 :: E4-E7	800	400	20	40
15	M7 :: E4-E7	800	400	20	40
16	M0 :: E8-E11	800	400	20	40
17	M1 :: E8-E11	800	400	20	40
18	M2 :: E8-E11	800	400	20	40
19	M3 :: E8-E11	800	400	20	40
20	M4 :: E8-E11	800	400	20	40
21	M5 :: E8-E11	800	400	20	40
22	M6 :: E8-E11	800	400	20	40
23	M7 :: E8-E11	800	400	20	40
24	M0 :: E12-E15	800	400	20	40
25	M1 :: E12-E15	800	400	20	40
26	M2 :: E12-E15	800	400	20	40
27	M3 :: E12-E15	800	400	20	40
28	M4 :: E12-E15	800	400	20	40
29	M5 :: E12-E15	800	400	20	40
30	M6 :: E12-E15	800	400	20	40
31	M7 :: E12-E15	800	400	20	40

Additional performance gains are made possible in IEEE 802.11n through the use of reduced guard intervals, a wider channel width, frame aggregation, and block acknowledgments (ACK). Many of the rate adaptation schemes designed for IEEE 802.11 are not applicable to the enhancements available in IEEE 802.11n. As a result, they need to be redesigned to take full advantage of these enhancements. Recently, there have been a limited number of efforts trying to address these problems and challenges for rate adaptation in MIMO systems.

In [19], the authors address the complexity of deciding between spatial multiplexing and diversity by proposing the Demmel condition, which is based on instantaneous channel state and the computed minimum Euclidean distance at the receiver, as an indicator of spatial structure of MIMO channels [20]. The proposed rate adaptation is based on a two-dimensional look-up table of the average SNR per subcarrier

and the average Demmel condition number per subcarrier, which are exchanged through RTS and CTS packets [19]. Evaluation of rate adaptation based on the Demmel condition is limited to a small subset of features of IEEE 802.11n.

Ath9k [3] is an emerging open-source driver for IEEE 802.11n. It contains its own rate adaptation based on probing, and sorts the rates that provide the best throughputs, which takes into account of the packet error rates (PER). The PER for each rate is defined as the ratio of the number of bad frames over the number of frames transmitted. The PER for each rate is maintained using an exponential weighted moving average (EWMA). The main disadvantage of Ath9k is the random probing used and searching for the optimal rate. In some cases, rates are stuck in sub-optimal rates.

Minstrel High Throughput (HT) [4] extends Minstrel [10] for IEEE 802.11n rate adaptation by constructing two dimensional array tables with sample columns and MCS group rates. Because of the large sample space, Minstrel High Throughput divides the MCS group rate into smaller 8 sub-groups with varying number of streams and channel width. It then populates the samples table with rates randomly chosen from these groups and uses them for sampling rates. With the division of search space, Minstrel HT still suffers the same drawback as Ath9k due to its randomness in search.

MiRA [5] is based on zigzagging and sampling between intra- and inter-mode rate options or the mode selection between single spatial stream and double spatial streams. It begins by sampling rates between the intra-mode, until it cannot achieve any more throughput gains. Then, it switches to the inter-mode for further sampling. After this sampling period, it sorts through the sampled rates to find the best rate. It constantly monitors its current throughput for any sudden changes and upgrades to the next higher or lower intra-mode rate accordingly. Adaptive probing is also employed to lower the frequency of sampling bad rates in favor of good rates. The drawback of MiRA is that it expends resources exploring other rates when its first sample rate may be the best rate. MiRA, in its current state, is limited to single and dual spatial streams and does not address multiple spatial streams. For example, with three or more streams, we assume that it will now need to zigzag between every one of them. As a result, it may take longer to converge to the best rate.

Effective SNR [21] presents a delivery model by taking RF channel state as input and predicts packet delivery for the links based on the configuration of the Network Interface Controller (NIC). It takes advantage of the channel state information (CSI) either from feedback or estimated from the reverse path and computes its effective SNR by averaging the subcarrier BERs in order to find the corresponding SNR, where BER is a function of the symbol SNR and OFDM modulations. For MIMO, it computes effective BER averaged across both subcarriers and streams.

The drawback of using CSI is that SNR needs to be measured instantaneously, and feedback delay may not allow mode adaptation on an instantaneous basis [22]. CSI itself is an approximation of the wireless channel and has many param-

eters. To improve its accuracy, CSI may need to incorporate other information, such as higher order statistics of SNR and Packet/Bit Error Rate or both for robustness [22]. However, CSI may work well if the channel condition does not change instantaneously.

Given that simplicity in using implicit feedback is a key goal in our approach to rate adaptation, we address it by separating IEEE 802.11n rate adaptation into two groups: a modulation group and an enhancement group. We adapt these two groups concurrently. We adapt between the multiplexing and transmit diversity by relying on a simple observation that bit symbols are spread equally among antenna in spatial multiplexing. If the number of un-acknowledged packets is greater than the number of ACKs in block ACK packets, we consider switching it to diversity. When this occurs, it means that one or more of the multiple antennas is not doing very well. We will discuss this approach in the section III-B1.

III. MULTI-RATE ADAPTATION FOR MULTI-ANTENNA SYSTEM (RAMAS)

The amendment enhancements for higher throughputs to the IEEE 802.11 standard brings much promise as well as many challenges in adapting rates. For example, how do we switch between spatial multiplexing and transmit diversity? When should we fall back to using a smaller channel width (20 Mhz vs 40 Mhz)? Which guard intervals should we employ? For this reason, IEEE 802.11n divides and groups all different MIMO configurations into 76 MCS indices to help facilitate rate adaptation. However, MCS does not provide a monotonic relation between loss and index rates. This is why many rate adaptation schemes, including Ath9k [3], Minstrel HT [4], and MiRA [5], resort to variations of random sampling. Our goal is to show that one can take advantage of the monotonic relation between loss and modulation types and build a mapping with the new enhancement features of IEEE 802.11n. This mapping allows us to adapt rate in an orderly fashion instead of random sampling. With this motivation, we take a drastic approach by separating rate adaptation into two groups: (a) the modulation group, which consists of different types of modulation; and (b) the enhancement group, which consists of the new enhancement features for IEEE 802.11n such as spatial multiplexing, transmit diversity, guard interval, and channel width. Each group has its own rules for upgrading and downgrading indices, but they are adapted concurrently. These indices are then mapped back to MCS for adapting rates in IEEE 802.11n. Because modulation group consists of modulation in varied degrees of redundant bits, it is natural that we choose it greedily for the highest throughput. As for the enhancement group, we adapt them based on stream error detection and delivery ratio. We describe the modulation group first and follow by the enhancement group.

A. Adapting Modulation-Type Group

Table I lists a table of modulation schemes with its corresponding coding rate. For example, index 7 has a modulation type of 64-QAM with a coding rate of 5/6 (one redundant bit

Algorithm 1 Adapting Modulation Group

```

w = sampling of packets window
τε = success packets times acceptable error rates
τγ = credit threshold for promoting to the next rate
mindex = modulation group index as seen in Table I
credit γ = 0; retransmitPackets ρ = 0
successPackets σ = 0; errorPackets ε = 0
//comment: in addition to w, time window ω is required
while (σ + ε < w) do
  if (packetissuccess) then
    σ ++
  else if (packetiserror) then
    ε ++
  end if
  if (packetisretried) then
    ρ ++
  end if
end while
//comment: zero success packets or with many retries
if (σ == 0 || σ < ρ) then
  mindex = mindex - 1
  γ = 0
end if
//comment: downgrade modulation group index
if (ε > τε) then
  mindex = mindex - 1
  γ = 0
end if
//comment: within acceptable error threshold
if (ε ≤ τε) then
  γ ++
end if
//comment: ensure stability before upgrading
if (γ ≥ τγ) then
  mindex = mindex + 1
  γ = 0; ρ = 0
  σ = 0; ε = 0
end if

```

is inserted every six bits). It is natural that we want to increase to the highest index or the highest modulation type for the best throughput.

1) *Modulation-Group Adaptation Rules:* Adapting the modulation group is similar to rate adaptation in SISO systems. The separation into modulation group allows us to concentrate only on varying modulation types, which further reduces complexity and adaptation rules.

Modulation-group adaptation makes its decision by keeping track of the number of successes over a *rate adaptation window* w corresponding to the number of packets transmitted. First, it makes sure that packet transmission of w distinct packets (our minimum recommended value is 30), excluding retransmission, have been reached. Then, we enforce time window ω and make a decision when the time window




Fig. 1: Credit Threshold

concludes. We use $\omega = 100$ millisecond as an implementation guideline. The rationale for using a time window is to ensure proper reactivity in case the sender does not transmit w packets within ω seconds. This time window also allows us more flexibility in choosing w parameters.

During the transmission window, RAMAS keeps track of three counters, one each for three transmission cases: σ for packet success; ϵ for packet error; and ρ for packet retransmission.

- If the number of packet errors $\epsilon \leq \tau_\epsilon$, a credit is added to the credit counter γ . Once γ reaches the threshold τ_γ , then modulation group index is increased and γ is reset to 0. The credit counter allows us to increase more progressively to avoid erratic variations.
- If the number of packet errors $\epsilon > \tau_\epsilon$, index is decreased; the credit counter γ is reset to 0. Note that we decrease right away rather than subtracting a credit.
- Finally, if the number of success $\sigma < \rho$ (the number of re-transmissions), then the channel retransmits too much and the index is decreased.

At the end of the window, σ, ϵ and ρ are re-set to 0.

The exact method to increase and decrease index for modulation group is described in Algorithm 1.

Unlike many credit-based systems, which assign credits based on consecutive successes, our credit-based system is based on obtaining k successes out of N trials, where each trial is affected by the common “air interface” and contention errors.

Because stability plays in important roles in adapting rates, this Binomial-like property loosens the restrictions of consecutive successes and helps to facilitate the flow of credits and, at the same time, prevent rates from being “stuck” as experienced in ARF [2].

2) *Guideline for Setting Parameters:* Note that there are a few other parameters in Algorithm 1. First, what is the acceptable error threshold? Its purpose is to loosen requirement of consecutive successes and to facilitate the flow of credits. In our experiment, we find that requiring consecutive packet success can cause rates being “stuck” and throughput starvation. Between 15% to 25% acceptable error threshold, there is almost no impact on throughput (Our recommended value is 20%).

As for the credit threshold, its main objective is to ensure

Algorithm 2 Adapting Enhancement Group

```

 $e\_index$  = enhancement group index as as seen Table II
 $\alpha$  = Aggregate MPDU length
 $\beta$  = Aggregate MPDU ACKs length
 $\Phi$  = stream failure threshold
 $\Psi$  = packets delivery threshold for upgrading
 $\eta$  = number of streams;
 $\epsilon$  = stream errors transmission
//comment: for multiple streams
if ( $\eta > 1$ ) then
    //comment: check and register stream errors
    if ( $\alpha - \beta + \rho/2 > \beta$ ) then
         $\epsilon++$ 
    else
         $\epsilon = 0$ 
    end if
    //comment: making sure stream errors still persist
    if ( $\epsilon > \Phi$ ) then
         $e\_index = e\_index - 1$ 
         $\epsilon = 0$ 
    end if
    //comment: for upgrading index
    if ( $\frac{success}{attempts} > \Psi$ ) then
         $e\_index = e\_index + 1$ 
    end if
end if
//comment: for single stream
if ( $\eta == 1$ ) then
    if ( $\frac{success}{attempts} > \Psi$ ) then
         $e\_index = e\_index + 1$ 
    end if
end if

```

stability before upgrading the next modulation index. Figure 1 plots the impact of different credit threshold on throughput for various scenarios. The credit threshold is nothing more than an additional mechanism to prevent false positives. Values between 10 and 20 have almost no effect on throughput (Our recommended value is 10).

B. Adapting Enhancement Group

Table II shows the ranking of indices for the enhancement group. These are the new features and enhancements in IEEE 802.11n that allows node to achieve a much higher throughput. Number of spatial streams denotes the number of independent streams in transmissions. Guard interval denotes the space between each symbol and is enforced to reduce inter-symbol interference (ISI). Guard Interval can be switched between 400ns or 800ns depending on hardware support. The same applies to channel width with 20 Mhz and 40 Mhz channel. This enhancement group table covers up to a maximum of four antennas; however, our experiment uses off-the-shelf wireless cards that only supports 3 spatial streams and MCS indices up to 23. It is natural that we adapt the enhancement group index




Fig. 2: Stream Error Failure Threshold




Fig. 3: RAMAS Performance Decomposition.

to the highest index, in this case, 4 spatial streams, 400ns guard interval, and 40Mhz channel width.

1) *Spatial Multiplexing vs Transmit Diversity*: The selection between spatial multiplexing or transmit diversity can have an adverse impact on performance. Spatial multiplexing allows nodes to send multiple streams of data independently. Transmit diversity allows nodes to replicate data on different fading paths to increase robustness.

In [19], the authors point out that spatial multiplexing works best when there exists a multiplexing structure; otherwise, it may reduce the reliability and worsen the throughput. But how can we identify this multiplexing structure? In [20], the authors propose the Demmel condition as an indicator for figuring out this multiplexing structure. Based on instantaneous channel state and the computed minimum Euclidean distance at the receiver, the Demmel condition of the matrix channel can be obtained and provides sufficient condition for selecting between multiplexing and diversity. This Demmel condition needs to be relayed back to the sender by control packets.

Because RAMAS is centered on simplicity and implicit feedback, we strive to provide a simpler approach. The intuition revolves around spatial multiplexing and how symbols are spread out among each antennas. Therefore, RAMAS keeps track of the number of packets transmitted and acknowledged through spatial multiplexing. If the number of loss packets is greater than the number of ACKs, RAMAS deduces that one or more of the multiple antennas is having difficulty transmitting packets. If this event persists, RAMAS lowers the number of spatial streams until there is only one stream. Our approach allows nodes to gradually lower the number of spatial streams and switch to transmit diversity completely when the situation warrants it. The goal is to use spatial multiplexing

whenever possible because that is where most of the additional throughput capacity can be harvested.

2) *Enhancement-Group Adaptation Rules*: Adaptation for the enhancement group is centered on stream error detection. Our goal is to determine when one of the antennas is not doing well in terms of transmitting the packets. This is the case where the number of loss packets is greater than ACKs in a Block ACK bitmap. We rely on aggregate MAC protocol data unit (AMPDU) length and block ACK length for stream error detection. MPDU aggregation comes in various sizes, and more aggregation leads to higher throughput.

Note that the Enhancement Group adaptation has no packet window w or time window ω , instead, it is checked with the reception of block ACK packets. We divide it into two cases: multiple streams and single stream. Obviously, for single stream or complete transmit diversity, there is no need for stream error detection or further downgrading enhancement index since it is already at the lowest.

a) For multiple streams, RAMAS checks the following:


- If there are more errors than ACK packets during AMPDU transmission, one or more of multiple streams must be having problems transmitting packets. RAMAS relies on the AMPDU length α and block ACK length β for checking this stream error by comparing unACKed packets $\alpha - \beta + \rho/2$ and ACKed packets β . If this is the case, RAMAS increments the stream error counter ε by 1. Otherwise, it resets it to 0. Note that RAMAS considers two retransmitted packets ρ as one error because too many retransmissions can have adverse impact on performance.
- Given that stream errors can come and go, RAMAS should not act prematurely based on one single instance of a stream error. For this reason, a stream failure threshold Φ is implemented to ensure its persistence before downgrading the enhancement index and reducing the number of streams.
- Upgrading the enhancement index is based on the packet delivery ratio. Given that stream errors create more error packets than good packets, RAMAS relies on a high delivery ratio of Ψ to ensure stability for upgrading enhancement index.

b) For a single stream, or complete transmit diversity, RAMAS applies the same delivery ratio method as mentioned above for upgrading index.


The exact method to increase and decrease index for enhancement group is described in Algorithm 2

3) *Guideline for Setting Parameters*: Stream Error Failure Threshold is used to ensure that the stream error persists before we downgrade the number of streams. Figure 2 shows the stream error threshold and its corresponding throughput. Setting the stream error threshold to 1 suffers a great throughput degradation compare to other choices. The same applies to other extreme end of setting it to 15 since it is waiting too long to downgrade index.


In Algorithm 2, notice that we decrement the enhancement




(a) Map of Engineering 2 Building




(b) 2x2 MIMO High Throughput Locations



(c) 2x2 MIMO Low Throughput Locations



(d) 3x3 MIMO High Throughput Locations



(e) 3x3 MIMO Low Throughput Locations

Fig. 4: Experiments with Different Locations Around Our Engineering Building II

index instead of lowering it to the corresponding stream. Knowing that the stream is corrupted, the objective is to see if we can try anything else before resorting to complete transmit diversity, because transmit diversity yields the lowest throughput. This is why our algorithm begins by relaxing the guard interval, and channel width as well as lowering the number of streams before falling back to transmit diversity.

Stability plays an important role to adapting enhancement group. The objective is to ensure that there are no more stream errors before upgrading it to higher number of streams. In our experiments, we find that a high delivery ratio provides good indicator for upgrading enhancement index. If we set the acceptable delivery ratio too low, the protocol tend to encounter stream error again much sooner. Our recommended value for acceptable delivery ratio is 90% over the same recommended w sample window.

C. Mapping Modulation Group Index and Enhancement Group Index to MCS Index

We adapt modulation group and enhancement group concurrently and simultaneously. Each comes with its own rules and methods. These indices together are mapped back to the MCS index for adapting rates in IEEE 802.11n. Table III shows mapping to MCS index for up to four antennas even though our current software and hardware only support up to 3 antennas or MCS indices up to 23. For example, MCS Index 23 $M7 :: E8 - E11$ refers to a modulation index of 7 (or 64 - QAM) and enhancement group index of 8 to 11 in which these features' availability is depending on hardware support).

Figure 3 shows the performance decomposition of RAMAS. First, we modified RAMAS with no Binomial-like credit

system by setting parameter τ_γ to 1. With no credit system, RAMAS performs erratically. Second, we remove stream error detection and use only delivery ratio for upgrading and downgrading enhancement index. We observe that without stream error detection, the throughput with RAMAS drops sharply. Finally, with both credit system and stream error detection, RAMAS performs robustly.

IV. PERFORMANCE EVALUATION

Aside from comparing our protocol against two widely used and popular open-source rate adaptation ATH9K [3] and MINSTREL High Throughput [4], we implemented a modified ARF [2] and MiRA [5] for more rigorous baseline comparisons. We modified MiRA algorithm to have it zigzag through an extra triple streams for the 3x3 MIMO experiment.

Our experiments include wireless cards and router with Atheros wireless chipsets AR9380, AR5416+AR2133, AR5418+AR5133. We implemented RAMAS using open-source software from the Linux Kernel Wireless Stack and using the Wireless-Testing Git Tree [23]. Our implementation of RAMAS in the Linux Kernel Wireless Stack allows it to run on many different chipsets and independent of any specific hardware vendors.

All of the experiments are conducted in the 2.4 Ghz frequency due to its long range and being more suitable for many of our experiments, which range from forest settings to car mobility. All of our experiments are set to Greenfield mode so that only devices with IEEE 11n capability can participate. Because of the nature of real-world experiments, we run at least 7 repetitions for each of our data point and provide error bars. Our experiment scenarios include indoor fading, outdoor fading, limited mobility, and interference collisions.






Fig. 7: Multi-user Experiments



(a) GoogleMaps (b) Redwood Forest




(a) GoogleMaps




(b) Throughput

Fig. 6: Mobility 20mph (32km/h) from point A to B



(c) High Throughput Locations



(d) Low Throughput Locations

Fig. 5: Outdoor Redwood Forest Experiments

A. Indoor Fading Scenario

1) Setup: We place an AP router at a far end corner of the building shown as a star symbol in Fig. 4(a); then, we measure throughput at different locations. It is atypical to place the AP router near one end of the building, but our goal is to see the maximum coverage of 802.11n router. We vary AP router with 2 and 3 antennas as well as clients with 2 and 3 antennas.

2) Indoor Fading Results: Fig. 4 reports the result for various locations around our building. On one side of the building from location A to E, AP router does not have any problems extending coverage as opposed to location F to I where there are multiple thick walls.

The result for 2x2 MIMO scenario in which both client and AP equipped with two antennas is shown in Fig 4(b) and Fig. 4(c). We separate the results into high- and low-throughput locations for readability. ARF generally suffers from throughput degradation due to its conservative nature and the requirement of consecutive successes. We find that its rates tend to get stuck at lower throughputs.

ATH9K's performance is unpredictable and its throughput tends to depend on certain locations. For example, at location G, ATH9K performs better than MINSTREL and MiRA, and it performs exceptional well at location I (the farthest distance from the AP). At location D and E, its throughput fluctuates greatly. At a few occasions, we find that ATH9K becomes confused and ends up transmitting at Kbps. This unpredictable performance can be attributed to the way ATH9K samples

We provide the best or optimal MIMO throughput for indoor fading locations where we cycle through different MIMO configurations. We will discuss each scenario and its result in detail next.

random rates. It tends to keep using the rates that offer reasonable throughput, and it often fails to explore optimal rates.

MINSTREL and MiRA have similar performance at locations near AP but MiRA performs better than MINSTREL at other locations. Both MINSTREL and MiRA sample many rates and sort through these sampled rates to find the best throughput. MiRA is able to perform better at long distance fading because it is able to switch to other mode appropriately. However, MiRA's performance still trails that of RAMAS. We find that MiRA resorts to sampling every time its throughput slightly fluctuates. At times, it tends to sample other rates when its current rate is working perfectly fine since it needs to make sure that the current rate provides better throughput.

Fig. 4(d) and Fig. 4(e) repeats the same experiment around the building but for 3x3 MIMO where both clients and AP are equipped with three antennas, for a maximum of three spatial streams. For this scenario, we had to modify MiRA to make it support triple streams by zigzagging through another stream space. Because of the extra space that MiRA has to sample, in addition of frequent sampling due to small changes in throughput, it still does not perform as well as RAMAS.

RAMAS succeeds because it adapts rates in an orderly fashion instead of using random sampling. In the presence of minor errors or fluctuations, RAMAS lowers the number of spatial streams and modulation types gradually and appropriately to harvest the best throughput. Taking into account of the margin of errors, RAMAS's performance is close to the best throughput when we cycle through different MIMO configurations.

B. Outdoor Fading Scenario

1) *Setup*: Apart from indoor scenario, we carried our experiments out into the redwood forest due to its rich dynamic of multiple path wireless link in an area of 330x330 feet square (100x100 meter square) with varying degree redwood tree density, slopes, and steeps. We place our AP router on the edge of a parking garage structure overlooking the forest (marked as a star symbol in Fig. 5(a)).

2) *Outdoor Fading Results*: Fig. 5 reports the results for outdoor experiments performed at various location around the 330x330 feet square area (100x100 meter square) of the redwood forest. At location A and G near AP, all protocols have similar performance. At location B to F where there is a varying degree of fading distance, density of the redwood forest, and elevations, we find varied performance among protocols.

ATH9K performs well in location C and F but does not fare well at other locations. MINSTREL and MiRA have similar performance except at location F where MiRA is able to perform much better. RAMAS performs much better than the other protocols at location E and F where there are steep elevations, dense redwood trees, and no visible AP in sight. At these locations, we observe that MINSTREL and MiRA resorts to sampling very often due to the dynamic changing of multi-path links. RAMAS, on the other hand, maintains transmit

diversity and only varies modulation types. This component of RAMAS gives it an edge over all other protocols.

C. Mobility Scenario

1) *Setup*: Mobility is an important requirement in wireless networks. Therefore, we set up a limited mobility scenario on our campus. We test mobility with speed limit of 20 mph (32 km/h), driving from point A to point B as shown in Fig. 6(a). We could not drive too slow or too fast due to safety on public streets.

2) *Mobility Results*: Fig. 6(b) shows the result for a limited mobility scenario with mobile speed of 20 mph (32 km/h). Overall, all protocols have similar performance. We cannot conclude much except to defer it for further research in the future.

D. Interference Collisions

1) *Setup*: Interference collision is very common in wireless environment. Because of the scarce resources, many nodes have to compete for access to the medium or share the medium with other nodes. We have multiple nodes sending traffic to the target nodes on AP router. First, we set up a transmission between two nodes. Then we vary the number of interferers from 1 to 5. These interferers, in turn, vary their sending rates using 10, 50, and 100 Mbps. Finally, we create a scenario where all five interferers assume a different sending traffic rate from 10, 30, 50, 70, and 100 Mbps. These nodes serve both as interferers and users competing for resources. The objective is to see how protocols perform in an interference collision and multi-user environment where other nodes are competing for access to AP router.

2) *Interference Collisions Results*: Fig. 7 reports the result for the interference collision and multi-user scenario where multiple nodes are sending traffic to AP. In Fig. 7(a), we gradually add the number of interferers where each sending a fixed traffic of 10 Mbps to AP. When there are no interferers, ATH9K, MINSTREL, MiRA have similar performance. As we increase the number of interferers competing for resources at AP, the performance of ARF, ATH9K, MINSTREL and MiRA degrade drastically. We observe that these protocols responded prematurely to small traffic disturbance. RAMAS has no problem recognizing this small traffic disturbance and is able to perform significantly better than all other protocols.

Fig. 7(b) and Fig. 7(c) repeat the same experiments but have each interferers sending a higher fixed traffic of 50 Mbps and 100 Mbps respectively. With the fourth and fifth added interferers, RAMAS's performance gains become less significant. This is because the interference collision is so great that none of the rate adaptations can work normally.

Fig. 8 present a scenario where we have five interferers sending five different traffic rates to AP with 10, 30, 50, 70, and 100 Mbps. Again, we observe that all protocols tend to suffer with this interference collision scenario where there are many users competing for resources with varying degree of traffic among the interferers. RAMAS performs significantly better in this scenario. We observe that most protocols fail this




Fig. 8: Multi-user Experiments: Mixed Users with 10, 30, 50, 70, and 100 Mbps




Fig. 9: Jain's Fairness Index

scenario because they do not consider the impact of multiple users on rate adaptations.

Fig. 9 reports Jain's fairness index [24] for multi-user scenario. In our initial test, we find that nodes with more processing power tend to grab the biggest share of the bandwidth, i.e. laptop vs desktop. As a result, we provide the same machine architecture for all nodes. For two and three competing users, all protocols have similar fairness index. As we increase to four competing users, MINSTREL's fairness becomes unpredictable. It is not surprising that ARF performs as well as RAMAS in terms of fairness because, like RAMAS, ARF increases its rate index incrementally and progressively.

V. CONCLUSION

We presented RAMAS, a novel and simple multi-rate adaptation approach for multi-antenna systems. RAMAS simplifies the complexity of rate adaptation in IEEE 802.11n by separating it into two groups, namely a modulation group and an enhancement group. These groups are adapted concurrently and their indices are mapped back to an MCS index. This mapping allows us to adapt rates in an orderly fashion instead of chaotic random sampling. Simplicity is attained by using only implicit feedback. We implemented our scheme using open-source software and carried out an extensive performance evaluation including indoor building experiments and outdoor forest experiments, and accounting for driving speed mobility to heavy interference collision with many users competing for resources. The results show that RAMAS consistently

performs better than all the well-known prior schemes for rate adaptation in multi-antenna systems. Furthermore, RAMAS is simple and practical.

VI. ACKNOWLEDGMENTS

This work was supported in part by the Baskin Chair of Computer Engineering, and by the U.S. Army Research Office under grant W911NF-05-1-0246.

REFERENCES

- [1] IEEE 802.11n Working Group, *Enhancements for Higher Throughput*, 2009.
- [2] A. Karmar and L. Monteban, *WaveLAN-II: A High-performance wireless LAN for the unlicensed band*. Bell Lab Technical Journal, 1990.
- [3] ATHEROS Communications.
- [4] F. Fietkau and D. Smithies, "Linux Wireless Minstrel High Throughput."
- [5] I. Pefkianakis, Y. Hu, S. Wong, H. Yang, and S. Lu, "MIMO Rate Adaptation in 802.11n Wireless Networks," in *In Proc. ACM MOBICOM*, Chicago, USA, 2010.
- [6] G. Holland, N. Vaidya, and P. Bahl, "A Rate-Adaptive MAC Protocol for Multi-Hop Wireless Networks," in *In Proc. ACM MOBICOM*, Rome, Italy, 2001.
- [7] M. Vutukuru, H. Balakrishnan, and K. Jamieson, "Cross-Layer Wireless Bit Rate Adaptation," in *ACM SIGCOMM*, Barcelona, Spain, August 2009.
- [8] H. Rahul, F. Edalat, D. Katabi, and C. Sodini, "Frequency-Aware Rate Adaptation and MAC Protocols," in *In Proc. ACM MOBICOM*, 2009.
- [9] J.C Bicket, *M.S Thesis: Bit-rate Selection in Wireless Networks*. MIT Press, 2005.
- [10] Minstrel Linux Wireless, "linuxwireless.org/en/developers."
- [11] S. Wong, H. Yang, S. Lu, and V. Bharghavan, "Robust Rate Adaptation for 802.11 Wireless Networks," in *In Proc. MOBICOM*, 2006.
- [12] J. Kim, S. Kim, S. Choi, and D. Qiao, "CARA: Collision-aware Rate Adaptation for IEEE 802.11 WLANs," in *In Proc. INFOCOM*, 2006.
- [13] A. J. Paulraj, D. Gore, R. Nabar, and H. Bolcskel, *An Overview of MIMO Communications-A Key to Gigabit Wireless*. Proceedings IEEE, 2004.
- [14] S. M. Alamouti, *A simple transmit diversity technique for wireless communications*. IEEE J. Select. Areas Commun, vol. 16, 1998.
- [15] V. Tarokh, N. Seshadri, and A. Calderbank, *Space-time Codes for High Data Rate Wireless Communication*. IEEE Trans. Inform. Theory, vol 44, 1998.
- [16] A. Paulraj and T. Kailath, *Increasing Capacity in Wireless Broad-cast Systems Using Distributed Transmission/Directional Reception*. U.S. Patent 5 345 599, 1994.
- [17] E. Telatar, *Capacity of Multi-Antenna Gaussian Channels*. European Transactions on Telecommunications, 1999.
- [18] G. J. Foschini and M. J. Gans, *On Limits of Wireless Communications in a Fading Environment When Using Multiple Antennas*. Wireless Personal Communications, 1998.
- [19] W. Kim, O. Khan, K. T. Truong, S. Choi, R. Grant, H. K. Wright, K. Mandke, R. C. Daniels, R. W. Heath, and S. M. Nettles, "An Experimental Evaluation of Rate Adaptation for Multi-Antenna Systems," in *IEEE INFOCOM*, 2009.
- [20] R. W. Heath and A. J. Paulraj, *Switching Between Diversity and Multiplexing in MIMO Systems*. IEEE Trans. on Communications, 2005.
- [21] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Predictable 802.11 Packet Delivery from Wireless Channel Measurements," in *In Proc. of SIGCOMM*, New Delhi, India, 2010.
- [22] S. Gatreux, V. Erceg, D. Gesbert, and R. W. Heath, *Adaptive Modulation and MIMO Coding for Broadband Wireless Data Networks*. IEEE Communications, 2002.
- [23] Linux Wireless Stack, "wireless.kernel.org."
- [24] R. Jain, D.-M. Chiu, and W. Hawe, *A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer Systems*. DEC Research Report TR-301, 1984.