Towards MIMO-Aware 802.11n Rate Adaptation

(Ioannis Pefkianakis, Suk-Bok Lee, Songwu Lu)

Advanced Computer Networks
CS577-2014 Fall

Presented by Tianyang Wang
Oct.28th, 2014

Introduction

1. Study MIMO 802.11n rate adaptation(RA) on a programmable access point(AP) platform.
2. The existing RA solutions are MIMO-mode oblivious
3. Design MiRA, a novel MIMO RA scheme that zigzags between intra- and inter-MIMO modes.
4. The experiments show that MIMO-mode aware designs outperform MIMO-mode oblivious Ras in various settings, with goodput gains up to 73.5% in field trials.
5. A window-based RA solution is also examined.

Background

A. IEEE 802.11n Standard

MIMO: PHY uses multiple transmit and receive antennas to support two MIMO modes of operation.

1. Spatial diversity: Transmit a single data stream from each transmit antenna.
2. Spatial multiplexing: Transmit independent and separately encoded spatial streams from each of the multiple transmit antennas.

Background

B. Experimental Platform and Setting

1. A programmable AP platform, which uses Atheros AR5416 2.4/5 GHz MAC/BB MIMO chipset.
2. Support SS and DS modes.
3. Available rates can go up to 130 and 300 Mb/s for 20- and 40-MHz channel operations.
4. Support frame aggregation with BlockAck

Background

AP
P1
P2
P6
P19
P4
P9
P3
P5
P18
P7
P17
P15
P8
P14
P13
P10
P16
P11
P12

Preliminary Case Study

TABLE I

RATE DISTRIBUTION, GOODPUT, AND SFER OF EXISTING RA ALGORITHMS AT P4

Rates (Mbps)	Atheros RA	RRAA	SampleRate	Fixed Rate Goodput (Mbps)	Fixed Rate SFER
MCS2 (40.5SS)				36.23	0.12%
MCS3 (54SS)	49%			49.08	0.20%
MCS9 (54DS)				48.87	0.12%
MCS4 (81SS)				72.94	0.07%
MCS10 (81DS)				72.64	0.06%
MCS5 (108SS)	51%			96.46	0.15%
MCS11 (108DS)		47%	89%	98.31	0.16%
MCS6 (121.5SS)		53%	4%	74.01	17.92%
MCS7 (135SS)			7%	36.56	54.61%
MCS12 (162DS)				128.46	4.31%
MCS12 (216DS)				5.71	96.73%
Goodput (Mbps)	71.40	85.36	91.95		
SFER	0.59%	13.24%	7.25%		

These result clearly indicate that the existing RA algorithm cannot be effectively applied in 802.11n networks.

Preliminary Case Study

Two factors play a critical role: nonnegligible, nonmonotonic relation between SFER and rate, and frame aggregation

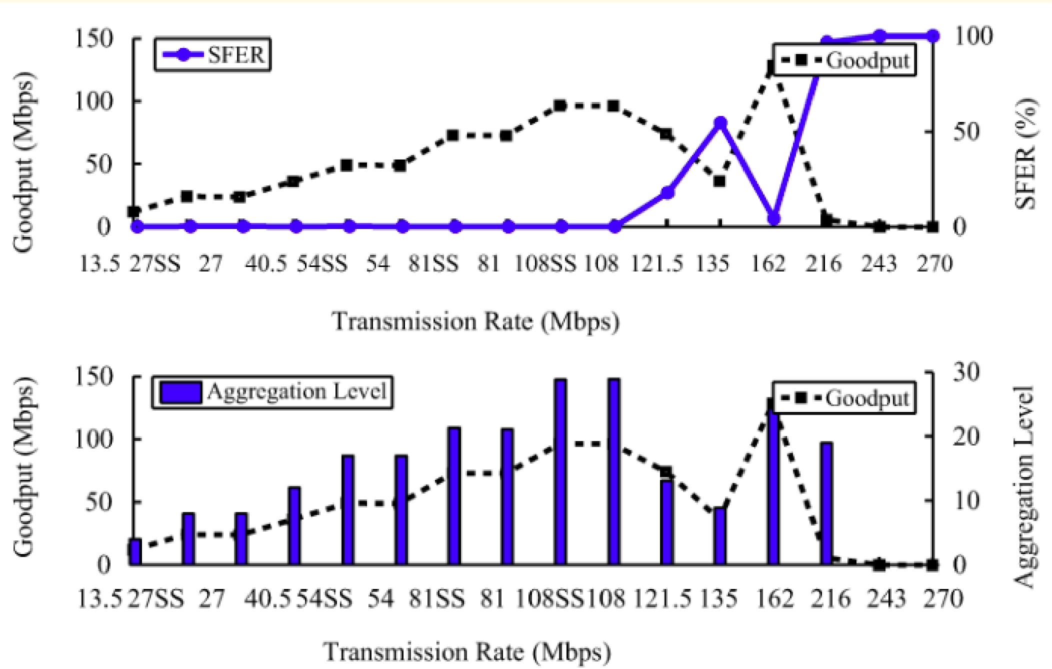

RRAA: Assume that SFER monotonically increase with rate.

Atheros MIMO RA: Assume monotonicity in that all rate above the current rate R have no smaller SFER.

SampleRate: Randomly samples diverse rates via probing, but suffers form stale statistics on the goodput and SEFR at a rate as it updates statistics only by probing these rate.

Preliminary Case Study

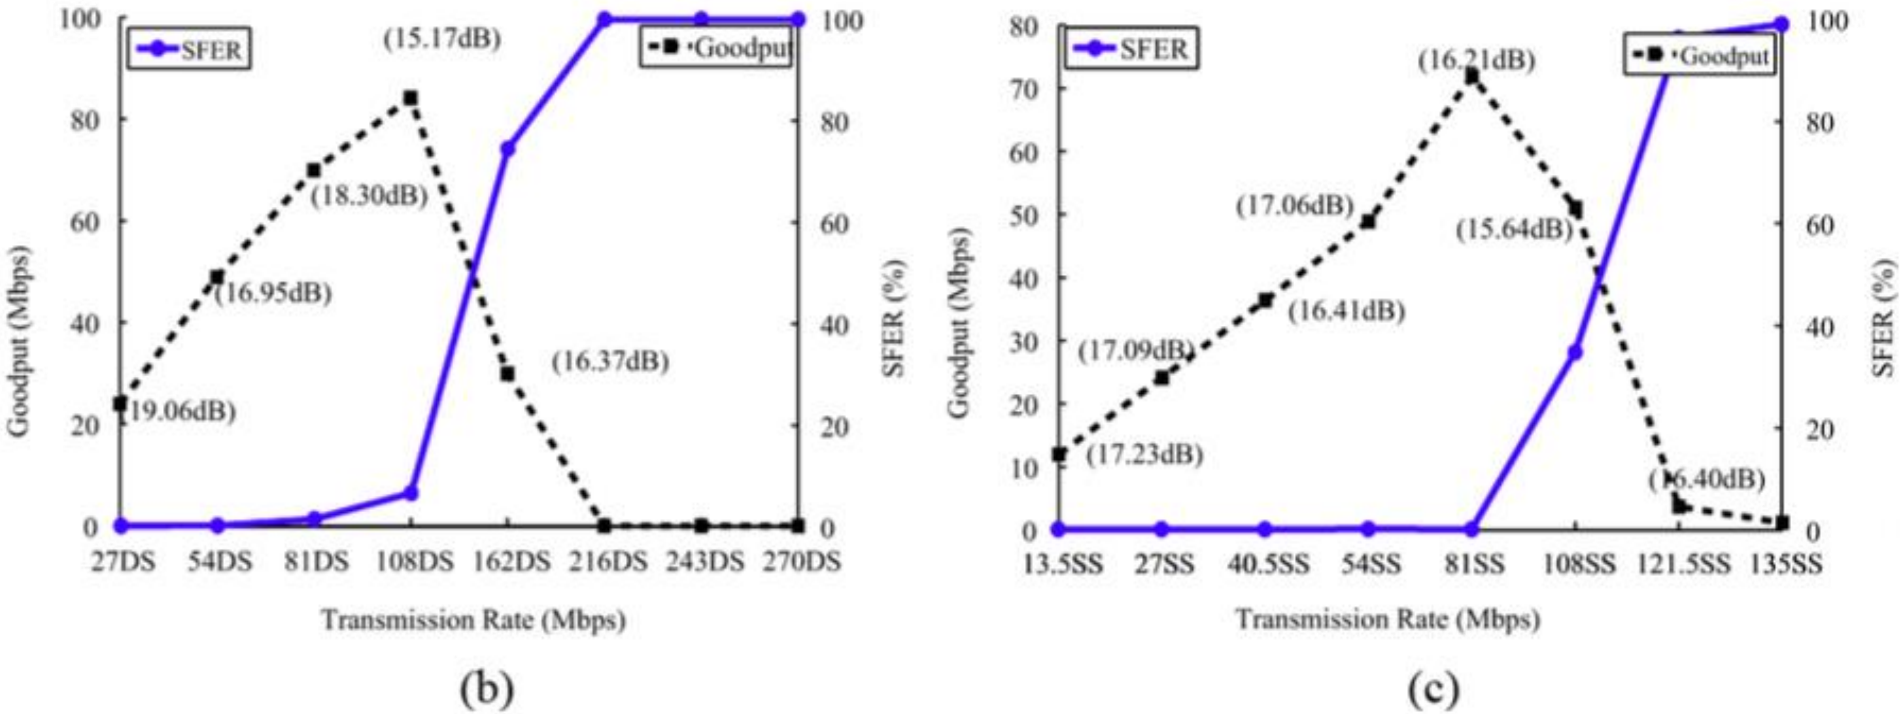

In contrast, their experiments reveal that the monotonicity between SFER and rate still largely holds in individual SS and DS modes.

An efficient rate adaptation design should be able to handle this nonmonotonic SFER behavior.

Preliminary Case Study

TABLE II
SFER NONMONOTONICITY W.R.T. RATE IN CROSS MODES

Location	$SFER_{121.5SS}$ (%) SNR (dB)	$SFER_{135SS}$ (%) SNR (dB)	$SFER_{162DS}$ (%) SNR (dB)
P3	0.39% 42.97 (dB)	7.99% 40.64 (dB)	0.33% 41.53 (dB)
P8	0.27% 29.69 (dB)	11.90% 30.80 (dB)	0.39% 31.22 (dB)
P4	17.92% 21.67 (dB)	54.61% 22.41 (dB)	4.31% 22.15 (dB)
P10	96.29% 17.38 (dB)	98.99% 16.75 (dB)	74.50% 17.79 (dB)

In low-SNR regions(<13dB), SS is more likely to outperform DS.
In high-SNR regions(>16dB), DS is more likely to outperform SS.

TABLE III
SFER W.R.T. DIFFERENT CROSS-MODE RATE PAIRS

Location	P10 SFER(%) SNR(dB)	P13 SFER(%) SNR(dB)	P14 SFER(%) SNR(dB)	P11 SFER(%) SNR(dB)	P7 SFER(%) SNR(dB)
MCS1 (27SS)	0.19% $17.10(dB)$	0.30% $14.93(dB)$	0.61% $12.96(dB)$	4.95% $12.34(dB)$	10.95% $7.03(dB)$
MCS8 (27DS)	0.23% $13.40(dB)$	0.31% $14.09(dB)$	0.52% $12.51(dB)$	17.79% $14.09(dB)$	25.143% $7.10(dB)$
MCS3 (54SS)	0.25% $16.1(dB)$	1.41% $12.34(dB)$	1.19% $12.87(dB)$	7.44% $10.60(dB)$	100% -
MCS9 (54DS)	0.25% $14.82(dB)$	0.72% $12.16(dB)$	9.23% $12.19(dB)$	16.73% $12.16(dB)$	100% -
MCS4 (81SS)	0.19% $17.05(dB)$	10.14% $11.95(dB)$	25.60% $11.58(dB)$	27.88% $11.95(dB)$	100% -
MCS10 (81DS)	1.54% $16.59(dB)$	10.03% $12.17(dB)$	37.04% $13.29(dB)$	37.15% $11.79(dB)$	100% -
MCS5 (108SS)	34.83% $16.13(dB)$	99.09% $11.64\ (dB)$	97.69% $13.15(dB)$	97.85% $11.64(dB)$	100% -
MCS11 (108DS)	6.68% $15.02\ (dB)$	82.88% $11.71(dB)$	93.60% $13.47(dB)$	98.24% $11.71(dB)$	100% -

However, one should be cautious in applying the above findings because they simply show the general trend rather than claim which specific mode is the winner in all cases.

WPI

Preliminary Case Study

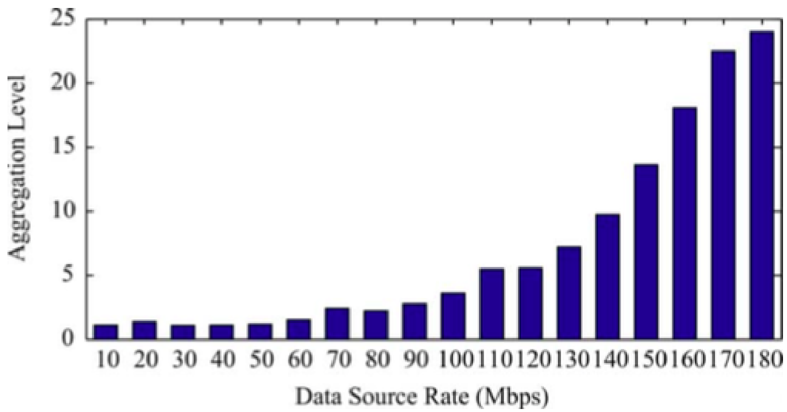

Fig. 4. Traffic source versus aggregation level.

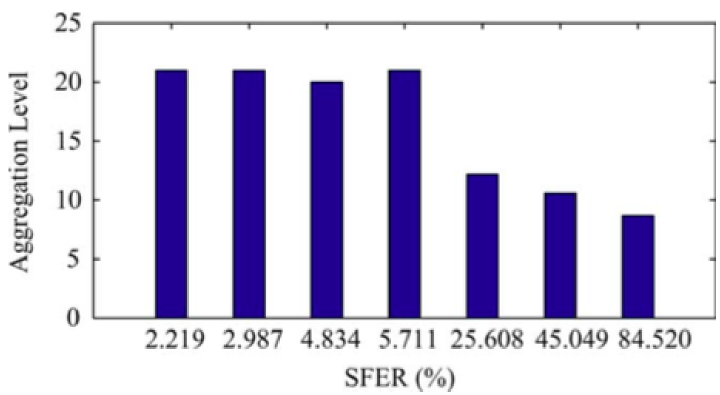

Fig. 5. SFER versus aggregation level.

Fig.4 presents aggregation level evolution with traffic source in a scenario where rate is fixed to 243 Mb/s and loss is smaller than 2%. However, higher SFER can have both positive and negative impact on frame aggregation.

Fig.5 plots the evolution of aggregation level with SFER in a setting where rate was fixed to 81SS and the data source was aggressive enough to ensure full software queue. And we see the high SFER dropped the average aggregation

Design

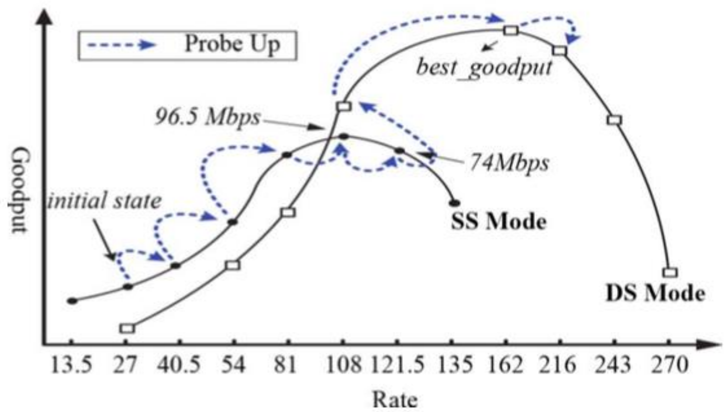

Fig. 6. Example for Zigzag RA: rate upward trajectory upon better channel.

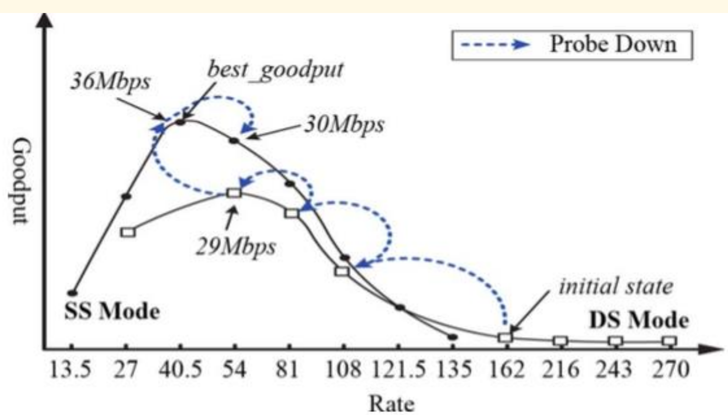

Fig. 7. Example for Zigzag RA: rate downward trajectory upon worse channel.

How It Works

MiRA zigzags between SS and DS modes.
It probes upward/downward within the current mode until it sees no further chance for goodput improvement. After intramode operations are completed, it then performs intermode RA by probing and changing rate to the other mode.
It uses probing-based estimation to identify the best goodput and adjust the current rate accordingly.

Design

TWO ISSUES, THE ZIGZAG RA SCHEME IN MIRA

1) How to decide which rates, in the same mode or across the mode, to probe

2) How to estimate the goodput based on the probing results while taking into account the effect of aggregation.

Design

The first issue:

1. ***Prioritized Probing***: Different from existing RA solutions, MiRA devises a novel, prioritized probing scheme to address MIMO-related cross-mode characteristics. It also applies adaptive probing to dynamically adjust the probing history in order to reduce excessive probing to bad rates.
2. ***Probing Triggers***: MiRA triggers probing and subsequent goodput estimation using both event-driven and time-driven mechanisms. Specially, it probes downward when $G_r(t) \leq \overline{G_r}(t) - 2 \cdot \sigma_r(t)$
3. ***Candidate Rates for Probing***: MiRA opportunistically selects the candidate set of rates to probe at a given time.
4. ***Two-Level Probing Priority***: MiRA ranks the sequence of rates to be probed within each mode and across modes using a two-level priority scheme. The first-level priority addresses intramode and intermode probing. The second-level priority manages probing order among candidate rates in the same mode.
5. ***Adaptive Probing Interval***: MiRA uses two mechanisms of loss-proporitional and binary exponential growth to adaptively set the probing intervals for three eligible rates: the two adjacent intramode rates and on intermode rate.

Design

The second issue:

1. Goodput Estimation: The moving average and deviations of the goodput at probe rate r are computed as follow

$$\overline{G_r}(t) = (1-\alpha)\cdot\overline{G_r}(t-1) + \alpha\cdot G_r(t)$$
$$\sigma_r(t) = (1-\beta)\cdot\sigma_r(t-1) + \beta\cdot|G_r(t) - \overline{G_r}(t)|$$

$$\overline{A_r}(t) = (1-\alpha)\cdot\overline{A_r}(t-1) + \alpha\cdot A_r(t)$$

$$G_r(t) = \frac{DATA\cdot\overline{A_r}(t)\cdot(1-SFER)}{T_{\text{overhead}} + \frac{DATA\cdot\overline{A_r}(t)}{r}}$$

However, they don't mention how they get the parameter α and β.

Design

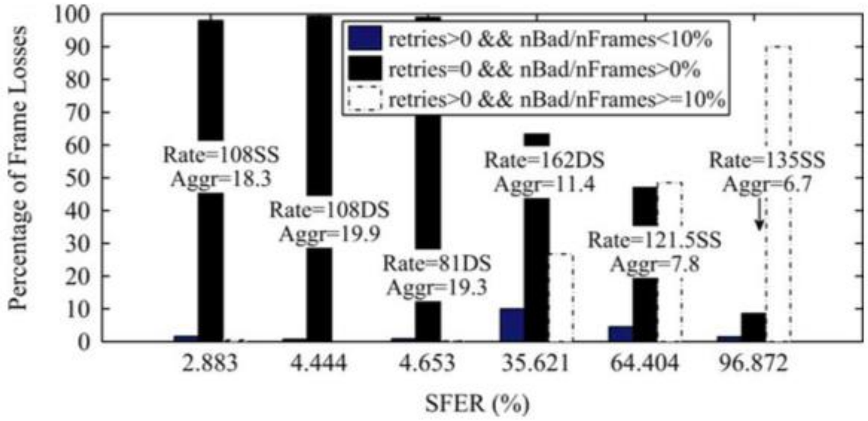

Fig. 8. Loss patterns without interference.

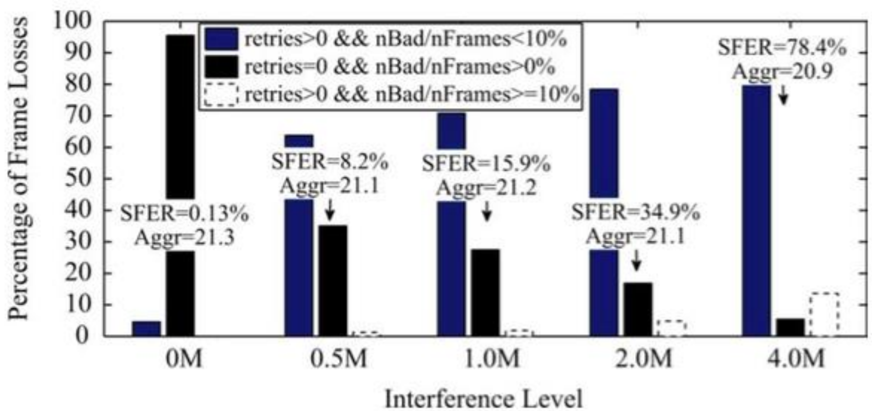

Fig. 9. Loss patterns with interference.

Handling Hidden Terminals:
A good RA design should differentiate between channel fading and collision losses

MiRA relies on repeated collision indications during a short timespan, rather than a single instance.

Window-Based 802.11n RA

Procedure 2: SingleModeRA: Input (BlockAck, MimoMode), Output (r)

```
getRSWndBounds(MimoMode, &maxRate, &minRate);

//Fast Reaction Mechanism
if txFailed(MimoMode, BlockAck) then
    move_RSWnd_down(maxRate, minRate);
else if down_timer_fired() &&
maxThr(maxRate, minRate) < lossfreeThr(minRate-1)
then
    //Timer Expired and Channel is Bad
    move_RSWnd_down (maxRate, minRate);
else if up_timer_fired() && $\overline{SFER}_{maxRate}$ < 15 then
    //Timer Expired and Channel is Good
    move_RSWnd_up (maxRate, minRate);
end if

update_probe_timer(BlockAck, r);

return best_RSWnd_rate(maxRate, minRate);
```

To overcome loss nonmonotonicity observed in cross-mode rates, WRA maintains and adjusts different RSWnds for both SS and DS modes.

Implementation

Implementation:

First, its probing mechanism requires frame transmission and rate control, which are two separate modules in the driver, to be synchronized on a per-AMPDU basis.

They maintain an additional binary state for each client, which is set upon collision losses and checked for each AMPDU transmission.

Second, the challenge is that the NAV for RTS cannot be directly set by the transmission module of the driver.

To reserve the wireless channel, the use Atheros' Virtual more Fragment interface, which consists of a virtual more-fragment and a burst-duration parameter.

Evaluation

TABLE V
GOODPUT GAINS OF MiRA OVER EXISTING RAs

	Atheros RA	RRAA	SampleRate
Static UDP	(3.4-82.3)%	(2.9-71)%	(1.1-104.5)%
Static TCP	(9.1-107.9)%	(5.9-37.5)%	(14.7-124.8)%
Mobility UDP	116.1%	30.2%	182.2%
Mobility TCP	72.5%	4.9%	94%
Hidden Terminal	(79.4-1094)%	up to 6.5%	(33.8-983)%
Field Trial	(46.35-67.4)%	(16-28.9)%	(19.4-73.5)%

The authors compare MiRA with RRAA, SampleRate, and Atheros MIMO RA.

Evaluation

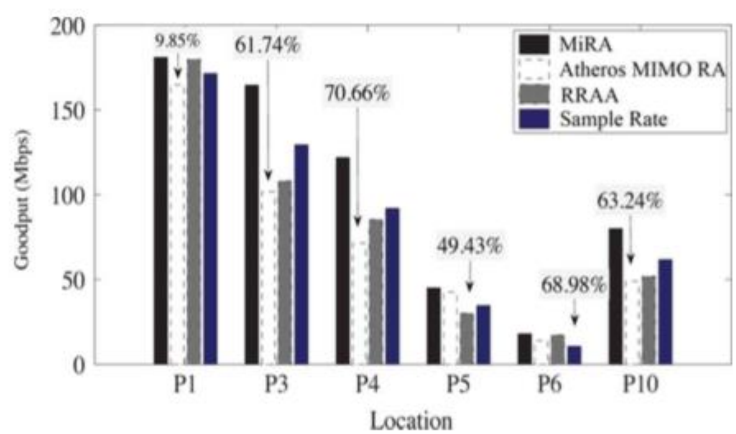

Fig. 10. 3 × 3/5-GHz/UDP static setting.

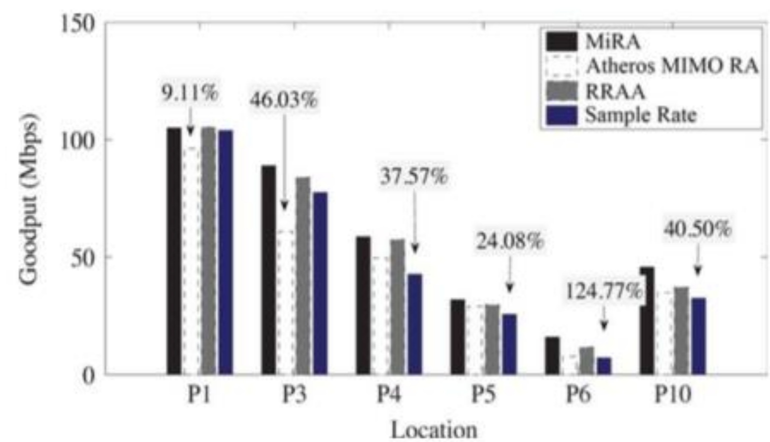

Fig. 11. 3 × 3/5-GHz/TCP static setting.

UDP goodput measured at six different locations with 3*3 antennas at 5-GHz band and the maximum MiRA goodput gains over the other design.

MiRA gives significant TCP goodput gain over others, up to 107.9% over Atheros MIMO RA, 37.5% over RRAA, and 124.8% over SampleRate.

Evaluation

Why MiRA has better goodput gains?

1. Effective Probing: Most existing Ras do not have any efficient mechanism to learn from short-term past channel's performance, which can lead to significant amount of transmissions in low goodput rates.
2. Handling SEFR NonMonotonicity: By zigzagging between MIMO modes, MiRA avoids to get trapped at lower rates in loss nonmonotonicity scenarios.

Evaluation

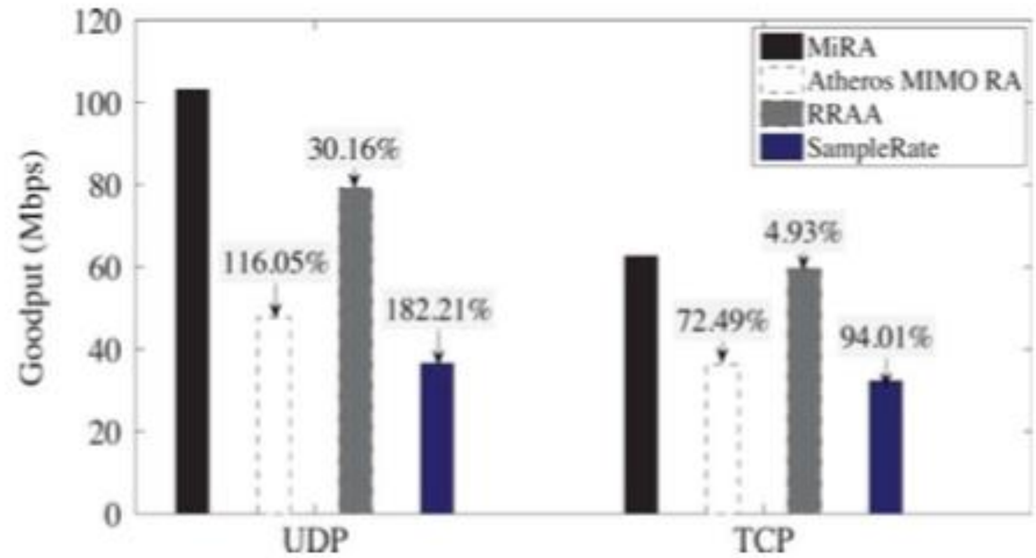

Fig. 13. Goodput in mobility setting.

Carry a client and walk from P1 to P6 and then come back at approximately constant speed of 1 m/s.

MiRA uses: 1) moving average to detect significant channel changes; 2) only on AMPDU to probe, which is transmitted in a relatively short period and typically contains enough samples; and 3) resetting statistical history upon rate changes.

Consequently, MiRA quickly adapts to channel dynamics due to mobility

Evaluation

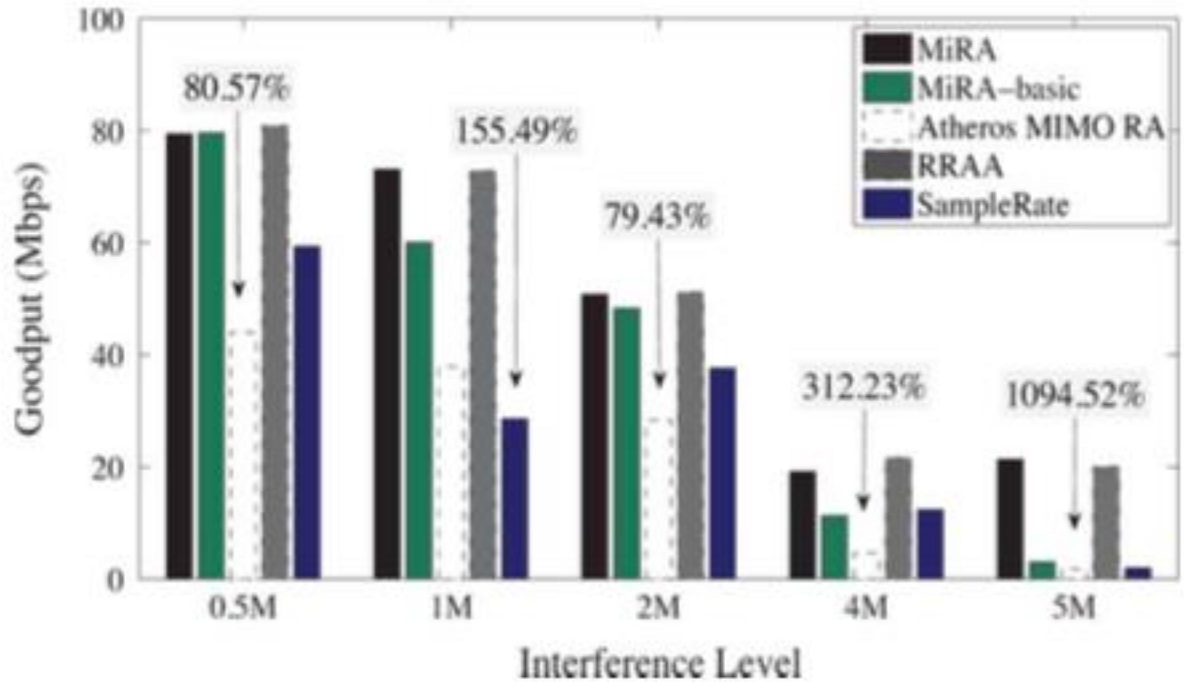

Fig. 14. Goodput in hidden terminal.

Evaluation

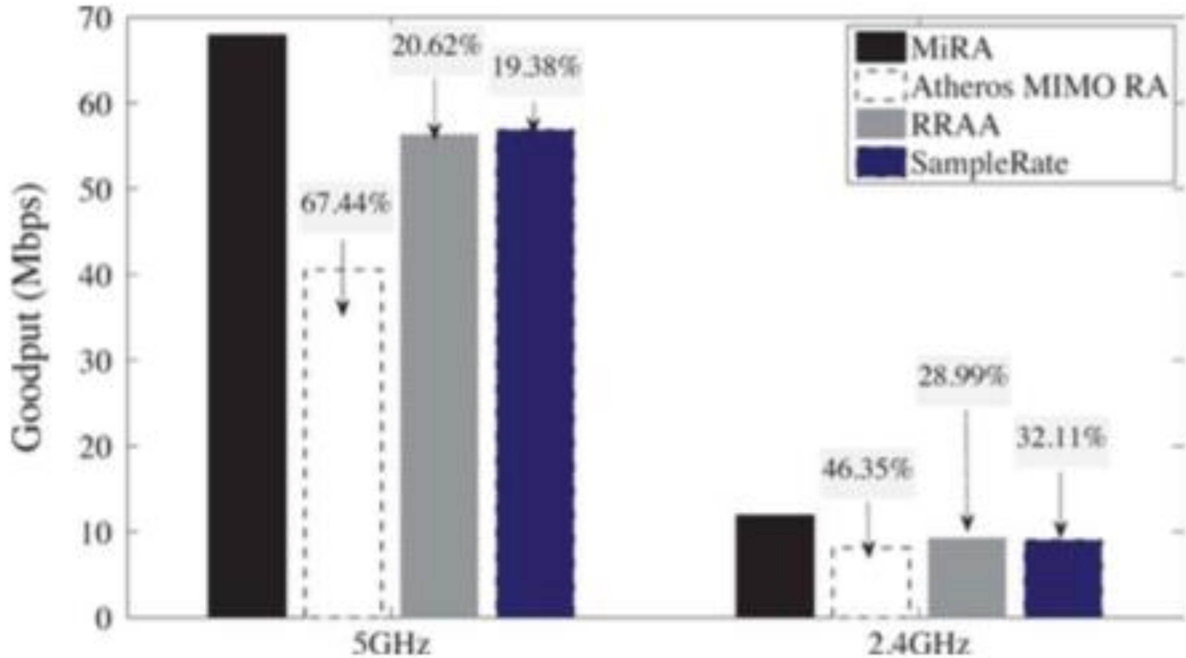

Fig. 15. MiRA performance in field trials.

Conduct uncontrolled field trials under realistic scenarios, where various sources of dynamics coexist in a complex manner. Use 3 static clients.
MiRA gives goodput gains up to 67.4% over Atheros MIMO RA, 28.9% over RRAA, and 32.1% over SampleRate.

Evaluation

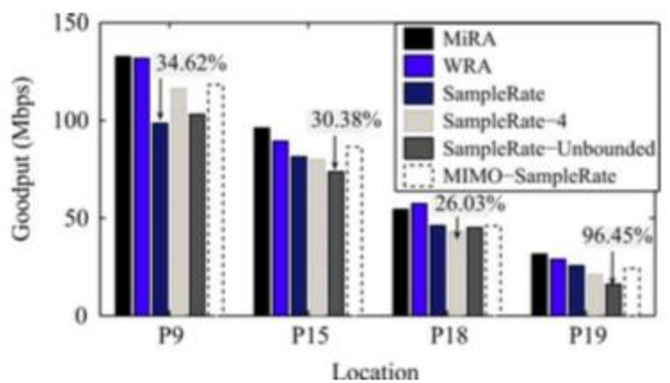

Fig. 16. MIMO RA alternatives in static settings.

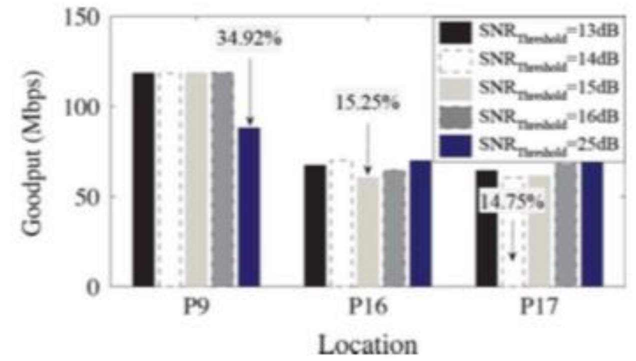

Fig. 17. MIMO-SampleRate versus different SNR thresholds.

Evaluation

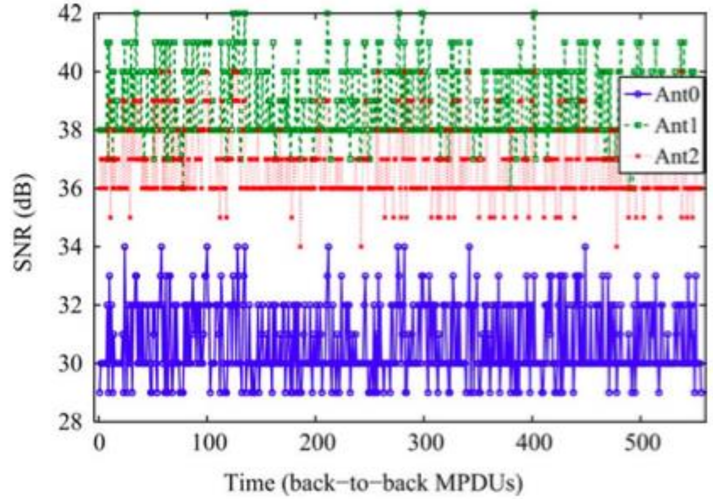

Fig. 19. Per-antenna SNR of control channel in a static RF-chamber setting.

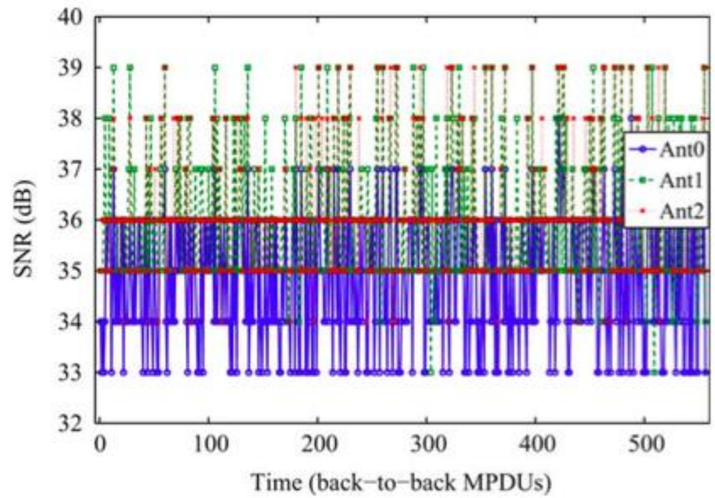

Fig. 20. Per-antenna SNR of extension channel in a static RF-chamber setting.

WPI

Related Work

There have been several rate adaptions, however, many of them target the legacy 802.11a/b/g networks or take a cross-layer approach, by using PHY-layer feedback to select the best-goodput rate. They are not designed for MIMO systems.

Though some schemes are designed for MIMO systems, they have not been widely adopted by commodity 802.11n systems due to their practical limitations.

Conclusions

The authors empirically study MIMO rate adaptation, using an IEEE 802.11n compliant, programmable AP platform.

Existing RA solutions do not properly consider characteristics of SS and DS, thus suffering from severe performance degradation.

Propose MiRA, a new zigzag RA algorithm that explicitly adapts to the SS and DS modes in 802.11n MIMO systems.

Also design and evaluate window-based and SNR-based MIMO RA solutions.

Experiments show clear gains of MIMO-mode-aware RAs.

Questions?