On link rate adaptation in 802.11n WLANs

Sriram Lakshmanan *, Shruti Sanadhya † and Raghupathy Sivakumar *
* School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA
† College of Computing, Georgia Institute of Technology, Atlanta, GA

Abstract—The IEEE 802.11n standard is gaining popularity to achieve high throughput in Wireless LANs. In this paper, we explore link adaptation in practical 802.11n systems using experiments with off-the-shelf hardware. Our experiments reveal several non-trivial insights. Specifically, (1) trivial extensions of algorithms developed for 802.11g provide minimal benefits in 802.11n systems; (2) in contrast to theoretical expectation, multiple antenna transmission does not always lead to higher throughput in practice; (3) both stream and antenna selection are essential to reap the full benefits of MIMO technologies. We use insights developed from experiments to develop a new metric for stream selection called the Median Multiplexing Factor (MMF). The proposed metric can be used to develop intelligent rate selection algorithms that can achieve high throughput with purely software changes.

I. Introduction

Link rate adaptation is a well studied problem in wifi networks using 802.11g/a technologies[1] [9], [11], [17], [6]. The adaptation algorithms react to channel conditions by adapting the operating point of the link with the goal of achieving the maximum possible rate for the conditions. The degree of freedom such algorithms exercise is the choice of the modulation and coding (MCS) scheme to be used for a given transmission. The highest order modulation that can satisfy a desired signal to interference and noise ratio (SINR) is chosen given the channel conditions. In the interest of not requiring any instrumentation of lower layers, the algorithms are also capable of estimating the channel conditions using coarse metrics such as packet delivery ratios. A trivial adaptation algorithm would involve using a lower order modulation when the packet delivery ratios are decreasing, and using a higher order modulation when the packet delivery ratio stays the same or increasing. The efficacy of the algorithms is evaluated based on how closely they approximate the achievable performance, and how fast they converge to this performance.

While 802.11g is by far the dominant technology in wifi deployments today, more recently the 802.11n standard has been ratified by the IEEE [2] for wifi networks. The 802.11n technology relies on multiple antenna elements at the transmitter and/or the receiver, employs algorithms to leverage the consequent spatial multiplexing and diversity benefits that such antenna arrays can provide, and in general offer the promise of significantly higher data rates. Several products that use the 802.11n standard are available in the market today. While link rate adaptation has to be performed in 802.11n networks, it is not clear that adaptation algorithms developed for 802.11g networks apply as-is for 802.11n links and networks.

At a high level, the number of degrees of freedom that can be exercised by a rate adaptation algorithm grows to *three* in 802.11n. Not only does the modulation and coding scheme continue to be a degree of freedom, but in addition the number of *spatial streams* to use and the specific *antenna elements* to use for those streams are also other degrees of freedom available. While it is true that adaptation algorithms developed for 802.11g do not consider the latter two degrees of freedom, it is reasonable to ask whether trivial extensions of those algorithms can be used in 802.11n networks. Interestingly, *commercially available 802.11n products do in fact use trivial extensions of rate adaptation algorithms developed originally for 802.11g*. Briefly, the trivial extensions include hard wiring the choice of antennas for a given number of streams, and eliminating several of the (# streams, modulation/coding) tuples to ensure that the rates of the different combinations retained are non-overlapping so that the linear adaptation of the 802.11g rate adaptation will work as-is.

However, in this work, we first investigate whether or not the performance of 802.11n links can be improved by using algorithms that truly exercise all three degrees of freedom. *We conduct the evaluation using an experimental test-bed that includes 802.11n products from four different vendors using Broadcom, Ralink, Atheros and Intel chipsets.* Based on the experimental results, we conclude that there is considerable room for improvement in performance by using adaptation algorithms better tailored to the properties of 802.11n links. Our experiments reveal several non-trivial insights on stream and antenna selection for practical 802.11n links. Interestingly, *we find that even when physical layer feedback is achievable, ideal adaptation is still non-trivial in 802.11n links*. We develop a novel metric for rate selection in multiple stream MIMO 802.11n links. This metric can be used to develop intelligent rate selection algorithms for 802.11n links with software-only changes to the WiFi devices.

The rest of this paper is organized as follows: in Section II we present some background material and motivation; in Section III we perform extensive experimental analysis of 802.11n links. In Section IV we present our metric for link adaptation. Finally, in Section V we discuss related work and conclude the paper in Section VI.

II. Background and Motivation

A. 802.11n Strategies and Rate adaptation

The IEEE 802.11n standard for High Throughput Wireless LANs incorporates several mechanisms to improve the throughput by the use of multiple antennas.

[1]In the rest of the paper we consider only 802.11g technology as the erstwhile standard, but our related arguments apply to 802.11a as well.

The key mechanism is *Spatial Multiplexing (SM):* where independent data streams are transmitted across multiple antennas at the transmitter (Tx) to the receiver on the same frequency. The receiver (Rx) uses the channel state information to separate out the combined data streams that arrive at its (multiple) antennas. When the channel is richly scattering and uncorrelated across antennas, the capacity of the link scales linearly with the minimum of the number of antennas at the Tx and Rx without increase in the spectrum requirement.

Link adaptation is the problem of adapting the link parameters that determine transmission rate to handle varying channel and network conditions [9], [11]. For a conventional 802.11 a/g link, a set of eight pairs of modulation and coding (MCS) are used to achieve a desired reliability of transmission (Bit Error Rate) across varying link SNRs. Consequently, they yield different physical layer rates that span from 6 Mbps to 54 Mbps. Link adaptation in such contexts, involves learning about channel conditions and using it to select the right MCS in run-time. The choice of the rate is determined by an estimate of the channel condition either using packet loss [11], [14], [3], delivery ratio [17], throughput [6] or Signal to Interference and Noise Ratios (SINR) [9], [12].

B. Motivating experimental study

The success of rate adaptation algorithms for legacy 802.11ag systems has prompted vendors to incorporate such approaches for 802.11n links as well. In fact the rate adaptation algorithms from two popular vendors Atheros [1] and [4] incorporate the collection of statistics by probing rates and using the success of MAC layer Acknowledgements for updating the rate in steps.

We conduct our experiments in an office environment using a Linksys WRT600N access point/router (based on a Broadcom chipset) and a client using an 802.11n miniPCI Desktop adapter from Sparklan (WMIR199N based on the Ralink RT2860 chipset). Our setup is shown in Fig. 1 a) [2]. The experiments are conducted in both the 5GHz band using a 20 MHz channel. All of the devices support transmission of up to two spatial streams. We use Iperf over UDP for the traffic. For the results, we determine the maximum rate obtained by increasing the Iperf rate till the throughput drops. Each run involves transmissions for thirty seconds and runs are repeated for increased confidence.
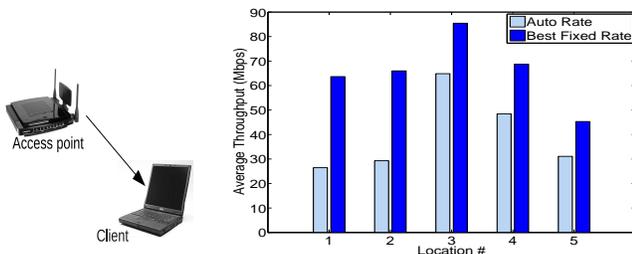
(a) Setup   (b) Throughput
Fig. 1. Scope for improvement over existing 802.11n

[2]The sub-optimality occurs even for nodes based on Atheros AR9160 and Intel4965agn chipsets in our testbed

We perform experiments for different client locations. For each location, we first use the Auto Rate control algorithm used by the Linksys router followed by experiments with fixed rates, by setting the MCS from 0 to 15. The resulting average throughput is presented in Fig. 1 b). The figures reveal that there is a significant gap (up to 2.7x) between the throughput with current rate control algorithms and the best throughput achievable, indicating *significant potential for improving the rate control algorithm.* Additionally, we observe that this gap exists even when using transmitters from other vendors, (i.e. up to 2.65x with the RT2860 Ralink chipset and up to 2.45x with the Intel 4965agn chipset and Atheros AR5416 chipset).

Thus link adaptation approaches that provide near-optimal performance for 802.11ag do not work well for 802.11n, even across different vendors.

Therefore, the key question we ask is: *What are the underlying factors which impact the rates of an 802.11n link and how can the link parameters be adapted in practice without requiring physical layer feedback?*

III. UNDERSTANDING MIMO 802.11N LINKS

We consider APs equipped with k_a antennas and r_a RF chains. Similarly the clients have k_c antennas and r_c RF chains. Typical values of these numbers vary across vendors with $2 \leq k_a \leq 16$ and $1 \leq r_a \leq 4$ and $2 \leq k_c \leq 3$, $2 \leq r_c \leq 3$ [5]. We perform experiments to analyze the performance of existing 802.11n devices across two dimensions, namely streams and antennas under non-interfered and interfered conditions.

A. Stream sustainability awareness

We use the same setup described in Fig. 1 a). In this experiment, we measure the impact of the number of streams chosen on the link throughput. We experiment with single stream rates MCS $0-7$ and two stream rates MCS $8-15$ and plot the maximum single stream rate and maximum two stream rate for each location in Fig. 2. In contrast to conventional wisdom that multiple streams leads to higher capacity than single stream, some locations yield a higher single stream rate than two stream rate!
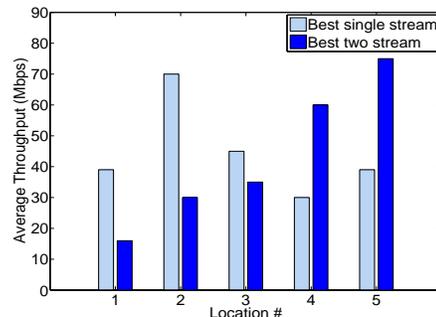
Fig. 2. Using maximum streams does not yield maximum benefits always

To understand this effect further, we study the gains for each modulation and coding pair in greater detail. We recall that MIMO enables increased throughput by spatial multiplexing of data streams [18]. Theoretically, the capacity increases linear to the minimum of the number of antennas at the transmitter and receiver. Hence, we are interested in determining whether
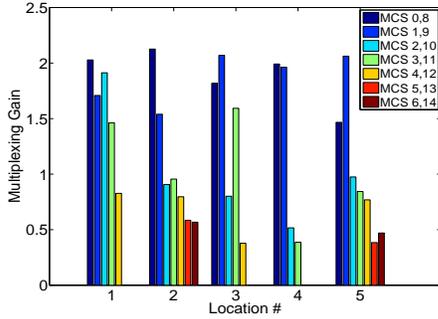
Fig. 3. Multiplexing gain varies across modulation and code for each location and is not always equal to number of streams

the spatial multiplexing gains are always achieved in practice and if so what the magnitude of improvements are. We define the spatial multiplexing gain (G) as

$$G(s, m, c) = \frac{Throughput(s, m, c)}{Throughput(1, m, c)} \quad (1)$$

where s is the number of streams, m is the modulation and c is the code rate.

We first observe the spatial multiplexing gain across the eight modulation and coding pairs present in MCS $0 - 15$ in 802.11n, for five locations. The results presented in Fig. 3 indicate that the multiplexing gains are not uniform across locations or across modulation and code sets, because of the large variation in channel properties.

We observe that for any AP-client location pair, lower modulations are more tolerant towards channel imperfections. i.e even if channel is not fully rich scattering, the multiplexing gain of two is observed for low modulations. However, high modulations do not provide the expected multiplexing gains because they are sensitive to channel imperfections such as lack of richness in scattering. Further, meeting the SNR requirements on both streams is challenging particularly when the channels are not full-rank. This has implications on the link adaptation possibilities in 802.11n links.

B. Interference effect and channel quality coupling

We study the impact of interference on 802.11n systems by generating interference using a linksys card whose carrier sense is disabled. The setup is illustrated in Fig. 4 a).

We study the average throughput across four MCS pairs that achieve the same PHY Rate. i.e. for example MCS 1 which uses QPSK and MCS 8 which uses BPSK both yield a 13 Mbps PHY Rate. The results of our first scenario are presented in Fig.4 b). As expected, the two stream MCS, which employs a lower modulation to achieve the same PHY Rate, performs much better than the single stream MCS. We repeat this experiment for another location and plot the resulting throughput in Fig. 4 c). However, in this case, the results contradict intuition. We observe that the single stream rate although operating at a higher modulation, also yields a higher throughput under interference. For instance, at point 1 in Fig. 4 c), MCS1 which uses QPSK yields higher throughput than MCS 8 which uses BPSK. We attribute this to the sensitivity of ill-conditioned MIMO channels to interference. i.e. If the

channels are not rich scattering and sufficiently de-correlated, interference affects two stream transmission significantly.

Our experiments reveal that, different from legacy 802.11 a/g systems, a degradation in channel quality and increase in interference power do not affect 802.11n links similarly. Conventional 802.11 ag rate adaptation calculates the Signal to Interference and Noise Ratio (SINR) or calculates packet losses irrespective of whether they stem from channel impairments or interference. While this works reasonably well for single stream links, the impact is very different in 802.11n links which involve both single and multiple streams. Specifically, the joint (MIMO) processing of the signals from multiple transmit and receive antennas allows a signal degradation to be overcome by decorrelating the streams at the receiver using the knowledge of the channel. However, the channel from the interferer to the receiver is unknown. Hence, instead of decorrelation across antennas, the joint processing amplifies the detrimental impact of interference. Consequently, 802.11n links are more susceptible to interference.

Implication: This result highlights why legacy rate adaptation protocols which rely on Packet Delivery Ratio of the individual rates or the Signal strength information cannot accurately determine the correct rate to use. Further, the interference must be estimated and differentiated from channel losses since the intended signal and interfered signal are subject to different effects after MIMO signal processing.

C. Maximizing multi-stream rates vs. maximizing SNR by antenna selection

The Linksys router used in the previous experiment has three fixed antennas. However, we are interested in studying the performance of more antennas than the number of streams. Using more antennas enables diverse channels to be achieved due to the multipath scattering nature of indoor propagation. Antenna selection has been discussed extensively in theory but its implications on link adaptation in 802.11n are not yet established. Hence, we perform experiments to study the potential of antenna selection for link throughout improvement. Since we do not have access to hardware that has multiple stream radios and multiple antennas, we create a setup with a desktop adapter using the Atheros AR9160 802.11n card in the AP mode using the open-source hostap driver. The desktop card has three antennas with 15cm long cables connecting them to the chipset. Hence, we create a regular square grid of 4*4 points each separated by 3cm which is half of the wavelength at 5GHz. We place the three antennas of the AP at different triplet of points in this grid and measure the average throughput for different antenna subsets. The results of our experiments for high SNR and moderate SNR conditions is presented in Fig. 5 a) for three strategies. The first is the fixed set of antennas at fixed corner points of the grid (denoted as Fixed Antenna in Fig. 5 a)). The second is the strategy that selects the antenna subset which yields the highest average SNR at the three antennas (denoted as Max. SNR in Fig. 5 a)). The third strategy involves a search across all possible triplets from 16 points and is denoted as Max.Rate in the figure. We also consider the three strategies of fixed antennas,

(a) Setup

(b) Moderate power

(c) High power

Fig. 4. Interference scenarios
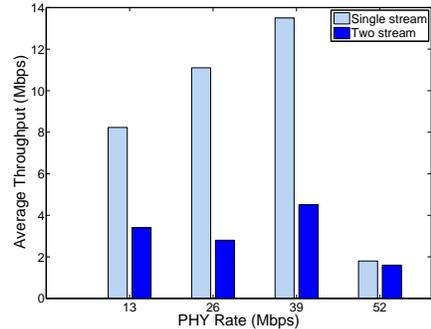
max. average SNR, highest rate subset under interference. The results plotted in Fig. 5 b.
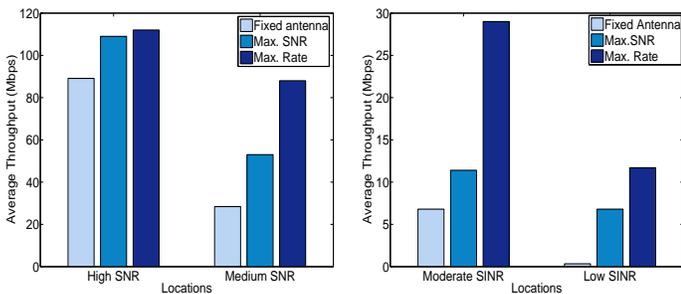
(a) Low to moderate SNR

(b) Low to moderate SIR

Fig. 5. Antenna subset impacts performance for low to moderate SINR

Our results reveal the following non-trivial insights: (1) Antenna subset affects performance and can improve throughput significantly under low to moderate SNR conditions. Our experiments with antenna choice in Fig. 5 reveal that antenna selection can improve throughput from 26 Mbps to 89 Mbps (improvement of 3.3x!). (2) For low and moderate interference antenna selection helps to offset the effects of interference. (3) Simple SNR based extension of antenna selection does not yield the best benefit and is very sub-optimal under interfered conditions.

IV. METRIC FOR FAST STREAM SELECTION

Simple Packet Delivery Ratio or SNR based metrics for each MCS do not capture the unique effects of 802.11n links such as spatial multiplexing gain of the given scenario, the sustainability and sensitivity of different modulations.

i. Metric: We recall from Fig. 3 that the multiplexing gains vary across locations and across modulation and code pair. Thus, whether a single stream or multiple stream yields better performance depends on the multiplexing gain of the middle region. Hence, we define a new metric to succintly capture whether a single stream or two stream is better. We call it the Median Multiplexing Factor (MMF), defined as the multiplexing gain factor of the median MCS for a given number of streams.

$$MMF(s) = \frac{\frac{Rate(MCS(3+(s-1)8))}{Rate(MCS3)} + \frac{Rate(MCS12(4+(s-1)8))}{Rate(MCS4)}}{2} \quad (2)$$

For the case of two streams, this can be simplified to the median of multiplexing gains of MCS 3 and 11, MCS 4 and 12.

$$MMF(2) = \frac{\frac{Rate(MCS11)}{Rate(MCS3)} + \frac{Rate(MCS12)}{Rate(MCS4)}}{2} \quad (3)$$

ii. Rationale: The above metric leverages the key observation that the richness of channel structure and the SNR can be reasonably accurately estimated by examining the performance of the median MCS in the set. Using just the performance of the lower MCS is insufficient as it does not reveal the highest rate that can be supported. Similarly, probing of highest rates may not always be sufficient (particularly under low-to-moderate SNR conditions). Hence, the median MCS probes (MCS: 3,4,11,12) for a given number of streams serve as better indicators than the extremes (MCS:1,7,8,15). However, the median MCS by themselves do not provide any information about the richness/structure of the current channel. At this stage we exploit the fact that richer channels result in higher delivery rates for multi-stream transmission and the number of parallel streams in a spatial multiplexed link can be uniquely given by the eigenmodes of the channel [18]. This in turn can be obtained by normalizing the multi-stream packet delivery rate to the single stream packet delivery rate to obtain the possibility of using higher streams. Specifically, when the rank of the channel between Tx and Rx is close to r, r data streams can be transmitted successfully using an MCS that depends on the minimum SNR of the streams. This is approximated by the ratio of the delivery rates as described in the equation.

iii. Experimental verification: We use results from experiments in our testbed to validate the use of MMF. Fig. 6 plots the MCS index of the highest rate for the five different locations, where MCS $0-7$ are single stream and $8-15$ are for two streams separated by the dotted line. We also plot the corresponding values of MMF for each of the locations in Fig. 7. We observe a close correlation between the highest rate MCS and the MMF. Further experiments indicate that MMF is a reasonably accurate indicator for most locations except at the edges of signal coverage, where the accuracy is reduced. Hence for a majority of locations in indoor scenarios, MMF picks the best rate with minimal probing overhead.

iv. Design Rule: (1) When $MMF > 1$, the best rate is a two stream rate, whereas when $MMF \leq 1$, the best rate is a single stream rate.
(2) When $MMF = 1$, the best MCS is closer to the mid
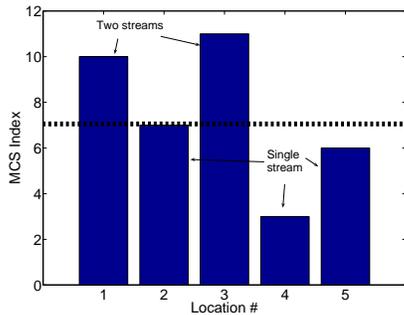
Fig. 6. Highest rate MCS for different locations
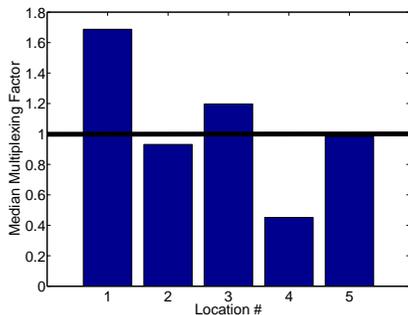
Fig. 7. Median Multiplexing Factor (MMF) for locations in Fig. 3 and 6

value of 7 (e.g. point 5 in Fig. 6 chooses MCS 6 and has a MMF value of 0.98, similarly point 2 chooses MCS 7 and has MMF of 0.93), the farther the MMF, the farther the best rate. (3) In high SNR scenarios, the MMF is closer to 2 and the highest two stream rates yield the best throughput.

Hence, the MMF can be used to intelligently index into the correct rate, without having to probe the rate of all MCS values. The MMF is motivated by the fact that multiplexing gains vary across modulations and codes used for a given channel condition. Further, the MMF metric extends easily to characterize the multiplexing gain with increasing number of streams. Our experiments reveal that the MMF can be used to yield the highest rate MCS under several channel and interference conditions.

V. RELATED WORK

Link adaptation has been studied extensively in the context of 802.11a/b/g networks. The sender picks the best transmission rate based on channel conditions, characterized through packet loss [11], [14], [3], delivery ratio [17], throughput [6] or Signal to Interference and Noise Ratio(SINR) [12], [9]. In the context of 802.11n link adaptation, [20] provides an experimental study highlighting the shortcomings of existing 802.11g rate adaptation algorithms when applied to 802.11n. Additionally, there are few works either requiring physical layer feedback [21], [16] or based on stream switching using extensive probing [15]. None of these works identify the stream selection dilemma and the interplay of antenna and streams for rate adaptation. Theoretical works have studied antenna selection using different channel models [22], [7], [10] and stream selection [8] to a smaller extent. Similarly, in [13] the authors compare 802.11n hardware platforms to study platform choices for testbeds, whereas in [19], the authors

focus on channel bonding and MAC diversity in 802.11n. Thus, in contrast to existing works, our work uses real-life experiments with 802.11n links to study link adaptation, without requiring physical layer feedback.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we study link adaptation in practical 802.11n systems using experiments. Our experiments reveal several non-trivial insights. Mainly (1) trivial extensions of algorithms developed for 802.11g provide minimal benefits in 802.11n systems. (2) Both stream and antenna selection are essential to reap the full benefits of MIMO technologies under interfered and non-interfered network conditions. We use insights developed from experiments to develop a new metric for fast and accurate stream selection called the Median Multiplexing Factor (MMF). As part of future work, we intend to develop a practical link adaptation algorithm using the above metric. We also intend to develop joint antenna and rate adaptation algorithms for 802.11n links.

REFERENCES

[1] Atheros Inc.: www.atheros.com.
[2] IEEE standard for Wireless Local Area Networks: 802.11n,http://www.ieee802.org/11.
[3] MadWifi Driver: http://www.madwifi.org.
[4] Ralink, http://www.ralinktech.com.
[5] Ruckus wireless inc.
[6] J. C. Bicket. Bit-rate selection in wireless networks. Master's thesis, M.I.T., Cambridge, MA, February 2005.
[7] Chen-Mou Cheng et al. Transmit antenna selection based on link-layer channel probing. In *Proc. of the IEEE WoWMoM*, 2007.
[8] M. F. Demirkol and M. A. Ingram. Stream control in networks with interfering mimo links. In *Proc. of the IEEE WCNC*, 2003.
[9] G. Holland, N. Vaidya, and P. Bahl. A rate-adaptive mac protocol for multi-hop wireless networks. In *Proc. of the ACM Mobicom*, 2001.
[10] Hyungsoo Kim et al. Efficient transmit antenna selection for correlated mimo channels. In *Proc. of the IEEE WCNC*, 2009.
[11] A. Kamerman and L. Monteban. Wavelan-ii: A high-performance wireless lan for the unlicensed band. *Bell Labs Technical Journal*, pages 118–133, 1997.
[12] T. S. Kim, H. Lim, and J. C. Hou. Improving spatial reuse through tuning transmit power, carrier sense threshold, and data rate in multihop wireless networks. In *In Proc of the ACM Mobicom*, 2006.
[13] Konstantinos Pelechrinis et al. Design and deployment considerations for high performance mimo testbeds. In *Proceedings of WICON*, 2008.
[14] M. Lacage, M. H. Manshaei, and T. Turletti. Ieee 802.11 rate adaptation: a practical approach. In *Proc. of the IEEE MSWiM*, 2004.
[15] I. Pefkianakis, Y. Hu, S. Wong, H. Yang, and S. Lu. Mimo rate adaptation in 802.11n wireless networks. In *Proc. of the ACM MobiCom (to appear)*, 2010.
[16] F. Peng, J. Zhang, and W. E. Ryan. Adaptive modulation and coding for ieee 802.11n. In *Proc. of the IEEE WCNC*, 2007.
[17] S. Wong et al. . Robust rate adaptation for 802.11 wireless networks. In *Proc. of the ACM Mobicom*, 2006.
[18] D. Tse and P. Vishwanath. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.
[19] Vivek Shrivastava et al. 802.11n under the microscope. In *Proceedings of the ACM SIGCOMM IMC*, 2008.
[20] Wonsoo Kim et al. An experimental evaluation of rate adaptation for multi-antenna systems. In *Proc. of the IEEE INFOCOM*, 2009.
[21] W. H. Xi, A. Munro, and M. Barton. Link adaptation algorithm for the ieee 802.11n mimo system. *Lecture Notes in Computer Science*, 4982:780–791, 2009.
[22] J. Yuan. Adaptive transmit antenna selection with pragmatic space-time trellis codes. *Proc. of the IEEE Transactions of Wireless Communications*, 5:1706–1715, 2006.