Characterization of 802.11n Wireless LAN Performance via Testbed Measurements and Statistical Analysis

Lito Kriara, Mahesh K. Marina and Arsham Farshad
The University of Edinburgh

***Abstract*—The 802.11n standard introduces a number of new MAC and PHY features to achieve high throughput and reliability. We conduct a comprehensive characterization of 802.11n performance with respect to its constituent features across a wide variety of scenarios with the aid of 802.11n wireless LAN testbed based measurements and statistical techniques including regression analysis. Our results show that different 802.11n features are interdependent when optimizing performance metrics such as throughput; the nature of interdependence as well as their relative impact are scenario dependent. We show the feasibility of online sender-side interference type detection, a key part of identifying the operational scenario for comprehensive 802.11n link adaptation, via a supervised machine learning based classifier. Finally, we highlight the unfairness problem of 802.11n networks that is linked to the frame aggregation feature.**

I. Introduction

Wireless LANs (WLANs) around the world are currently in the process of transitioning to use equipment based on the IEEE 802.11n standard [1], a high performance successor to the older 802.11a/b/g standards. The primary goal that led to the development of 802.11n was to achieve 100Mbps+ throughput above the MAC layer, nearly a 3-fold increase from that achievable through earlier 802.11a/g standards. A key physical (PHY) layer enhancement in 802.11n to realize this goal is the so-called MIMO, or the use of multiple antennas. Majority of 802.11n hardware available currently supports two antennas and two MIMO features, namely spatial division multiplexing (SDM) and space-time block coding (STBC) — SDM is aimed at improving throughput through the use of multiple concurrent data streams through different antennas, whereas STBC is for enhancing reliability by transmitting a single data stream across multiple antennas with redundancy. Besides MIMO, 802.11n also incorporates a channel bonding feature to permit doubling the channel width to 40MHz from the 20MHz width common in 802.11a/b/g systems. Other PHY enhancements include short guard intervals and replacement of a 802.11a/g bit-rate with a higher modulation and coding rate. Most of the above features along with the resulting bit-rates (in Mbps) are shown in Table I. LGI in the table refers to the long guard interval[1]. Note that SDM or STBC are not explicitly shown in the table — SDM option is implicitly chosen when more than one spatial streams are used, whereas STBC feature can be explicitly enabled when a single stream is used. Although PHY enhancements seem to be the key sources of improved performance with 802.11n, the actual throughput seen above the MAC layer is limited by the protocol overhead, more so than with 802.11a/b/g. So 802.11n includes two key MAC layer features called frame aggregation and block acknowledgements to improve the MAC efficiency by enabling multiple back-to-back frame transmissions upon each successful channel access.

MCS Index	Spatial Streams	Modulation Scheme	Coding Rate	20MHz w/ LGI	40MHz w/ LGI
0	1	BPSK	1/2	6.50	13.50
1	1	QPSK	1/2	13.00	27.00
2	1	QPSK	3/4	19.50	40.50
3	1	16-QAM	1/2	26.00	54.00
4	1	16-QAM	3/4	39.00	81.00
5	1	64-QAM	2/3	52.00	108.00
6	1	64-QAM	3/4	58.50	121.50
7	1	64-QAM	5/6	65.00	135.00
8	2	BPSK	1/2	13.00	27.00
9	2	QPSK	1/2	26.00	54.00
10	2	QPSK	3/4	39.00	81.00
11	2	16-QAM	1/2	52.00	108.00
12	2	16-QAM	3/4	78.00	162.00
13	2	64-QAM	2/3	104.00	216.00
14	2	64-QAM	3/4	117.00	243.00
15	2	64-QAM	5/6	130.00	270.00

TABLE I: (A subset of) 802.11n physical layer features and their corresponding bit-rates.

From Table I, we see that there are 32 different parameter settings for a link from a PHY perspective (16 MCS indices and 2 channel widths). Combined with the simplest MAC layer setting of whether or not to enable frame aggregation, we have 64 possible configurations to choose from to optimize the performance of a link. The best choice for a link, as with PHY rate adaptation in legacy 802.11a/b/g networks, maybe dependent on the scenario witnessed by the link at a given point in time in terms of channel and interference conditions. How to best select the settings for the various available features for a given 802.11n scenario to optimize link performance is the *802.11n link adaptation problem.*

The key to designing a comprehensive solution for the 802.11n link adaptation problem is an understanding of the impact of different 802.11n features on performance in different link scenarios as well as interdependencies among those features. *Gaining that understanding is the aim of this paper.* Briefly, our methodology to address this goal is as follows. Using an indoor 802.11n wireless LAN testbed, we experi-

[1]The short guard interval alternative is more efficient and is needed for 300Mbps bit-rates with 2 antennas but is supported in practice only for 40MHz channel width. So we limit our attention only to LGI.

mentally measure link performance with respect to different metrics (including throughput, packet loss and fairness) when using different settings for 802.11n features and under a wide range of link scenarios, including those that model adjacent channel interference. To gain insight from the large number of measurements so collected and to understand the relative impact of different 802.11n features on WLAN performance, we use regression analysis. Specifically, we use categorical regression [2] since 802.11n features are better viewed as categorical (nominal) variables. For understanding the interdependencies among various 802.11n features in different scenarios, we use the response surface methodology (RSM) [3].

Key findings of our study are:

- Regression based analysis is valuable in easing the characterization of the impact of different features on performance.
- The relative impact of different 802.11n features on performance (throughput, packet loss and fairness) is scenario dependent. For example, SDM is beneficial in terms of throughput only for high quality links and even that reduces in presence of adjacent channel interference (ACI), whereas channel bonding has a greater impact in scenarios with ACI.
- We find that different features are interdependent with respect to throughput and the nature of interdependence varies between scenarios. For instance, there is lesser degree of interdependence with poor quality links and in presence of interference because fewer set of features have majority of the impact on throughput.
- As a step towards practical and comprehensive link adaptation mechanism design, we show the feasibility of inferring interference type at a node online using throughput measurements and a supervised machine learning based classifier.
- We highlight the unfairness problem of 802.11n networks that is linked to the frame aggregation feature.

Our work improves upon earlier experimental studies of 802.11n networks [4], [5], [6], [7], [8], [9], [10] in two respects: (1) It is comprehensive in the set of features, metrics and range of scenarios considered. Table VI captures the focus of previous work. (2) In terms of the underlying goal — to capture important 802.11n features for performance optimization in different link scenarios and their mutual interaction. We also highlight the fairness issue with 802.11n that has not received much attention in the literature so far.

The rest of this paper is structured as follows. In the next section, we describe the various elements of our methodology to perform the characterization study as stated above. Section III presents our results as listed under the aforementioned key findings and discusses them. Section IV discusses related work and conclusions are provided in Section V.

II. Methodology

Our overall goal in this paper is to characterize the interaction between 802.11n features (frame aggregation, SDM vs. STBC, channel bonding, etc.) and their relative impact on link/WLAN performance across a wide range of scenarios, differing in channel and interference conditions. In this section, we describe the various elements of our methodology.

A. Indoor 802.11n Wireless LAN Testbed

In order to do the aforementioned characterization experimentally, we have deployed a 802.11n wireless LAN testbed in the Informatics Forum building at the University of Edinburgh. The testbed consists of 8 nodes in total of which 6 form an infrastructure 802.11n WLAN with one access point (AP) and 5 stations. Fig. 1 shows the locations of these 6 nodes on the building floor plan. The placement of these nodes was done to realize diverse set of link qualities as reported in the next section. As described in the next section, the other 2 testbed nodes are setup to be another co-located 802.11n WLAN to realize different interference conditions.

Each node in our testbed is actually a combination of a laptop and an embedded router board. Laptop is equipped with Centrino Duo 1.66GHz processor, 1GB RAM and Gigabit Ethernet interface and is setup to run Ubuntu 10.04 OS with Linux kernel version 2.6.32. The router board is a Ubiquiti RouterStation Pro[2] with 680MHz CPU, 128MB memory and 4 Gigabit Ethernet interfaces. The board hosts an 802.11n wireless interface card, specifically the MikroTik R52Hn 2x2 MIMO miniPCI card with an Atheros AR9220 chipset[3]. The miniPCI card on the board is connected to two dual band omnidirectional antennas. The laptop and router board of each node are connected through their Gigabit Ethernet interfaces and bridged via the Wireless Distribution System (WDS). We use this particular setup with the laptop acting as the traffic source/sink because we found the link throughput to be limited by CPU on the board when it is used as a traffic source/sink and at the same time operates at 802.11n top speeds. Note that this platform bottleneck issue has been reported previously in the literature [11]. Fig. 2 shows a picture of our testbed node. On the board, we use the open-source ath9k driver[4].

B. 802.11n Settings

We consider almost all 802.11n features — frame aggregation (FA), MIMO (SDM/STBC), channel bonding (ChB) and all available modulation and coding rates. The only exception is the short guard interval (SGI), which is only supported for 40MHz channels in the AR9220 chipsets. We disabled SGI in our experiments for consistency. Note that we consider the effect of enabling STBC for MCS indices 0-7 shown in Table I while MCS indices 8-15 in Table I always refer to the use of SDM. Similar approach was also taken in [8]. In order to separately see the impact of using SDM/STBC from the use of different modulation and coding rates, we use the term MCS in majority of the paper to only refer to the use of the eight modulation and coding rates shown in Table I and number them 0 – 7. The use of SDM or STBC features is explicitly shown separately.

To determine the channels of operation for our experiments, we surveyed the testbed area using the WiSpy spectrum

[2] http://www.ubnt.com/rspro

[3] http://www.mikrotik-store.eu/media/files_public/ejuissqr/R52Hn-Brochure.pdf

[4] http://linuxwireless.org/en/users/Drivers/ath9k

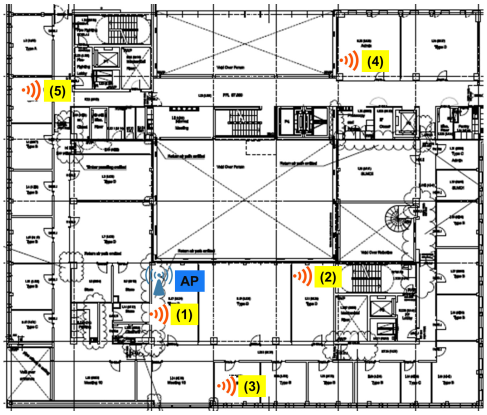

Fig. 1: Physical layout of nodes in the 802.11n WLAN testbed — the blue coloured node is the AP and red ones numbered (1)-(5) are stations.

Fig. 2: Picture of a node in our 802.11n WLAN testbed.

analyzer[5] to look for unused channels in both 2.4GHz and 5GHz bands since 802.11n can use either. We found that only channels 149 – 161 in the 5GHz band were free of any activity at all times, so decided to use only those channels for our experiments. Also, unless specified otherwise, the transmit power for 802.11n cards is at the default setting (25dBm).

For traffic generation, we use iperf[6] UDP traffic sessions between AP and one or more client stations. Packet size is fixed at iperf default value, which is 1500 bytes. Every experiment reported in this paper is repeated multiple times and the average value across those multiple experiment runs is taken as the measurement result.

C. Performance Metrics

We consider three metrics to quantify 802.11n link/WLAN performance: throughput, packet loss and fairness. Throughput of a link a running iperf UDP session is measured at the server (receiver) side. Aggregate throughput is used as the measure when multiple links in the WLAN are concurrently active. Packet (frame) loss is computed using MAC layer statistics at the sender side. Specifically, packet loss is measured as the difference between frames sent and successfully transmitted frames as a percentage of the frames sent. For quantifying fairness, we use the well-known Jain's index:

$$f(x_1, x_2, \ldots, x_n) = \frac{(\sum_{i=1}^{n} x_i)^2}{n \sum_{i=1}^{n} x_i^2}$$

where x_i is the throughput of the i^{th} link and n the total number of concurrently communicating links.

D. Statistical Techniques

1) Categorical Regression: We use categorical regression in sections III.B and III.E to study the relative impact of 802.11n features on performance. Standard multiple regression works best with continuous predictor variables. Categorical regression was developed as a method for linear regression for categorical predictor variables [2]. It relies on a method called optimal scaling that finds optimal numerical values to categorical values and in the process transforming categorical data into numerical data. The transformations of categorical variables are estimated simultaneously with the estimation of the regression coefficients using an alternate least squares procedure that maximizes the squared multiple regression coefficient, R^2, on the transformed variables. As a result, categorical regression results in transformed categorical variables that have values with numerical properties and are optimal for describing the relation between the response variable and predictors. Goodness of a categorical regression can be assessed with respect to a desired level of significance (typically, 0.05). If the p-value of the ANOVA is lower than the desired level then we consider the regression result to be statistically significant.

Pratt's measure [12] is a way to quantify the importance of each predictor and is seen as much more useful metric than the standardized regression coefficient. Pratt's measure for each predictor variable is computed by taking the product of its regression coefficient and its zero-order correlation (i.e., the correlation between the transformed predictor and the transformed response in the categorical regression). These products add to R^2, so importance values are usually divided by R^2 so that they add up to 1.

2) Response Surface Methodology: We use response surface methodology (RSM) [3] in section III.D to examine the interdependence among various 802.11n features with respect to a performance metric of interest like throughput. In RSM, the response variable y (e.g., throughput) is modeled as a function of the predictor variables $x_i, 1 \leq i \leq k$ as shown in the following equation:

$$y = f(x_i), i = 1, \ldots, k$$

Broadly speaking, RSM involves two phases. In the first phase, the function f is approximated as a quadratic function of the form:

$$y = \beta_0 + \sum_{i=1}^{k} \beta_i x_i + \sum_{i=1}^{k} \beta_{ii} x_i^2 + \sum_{i<j=2}^{k} \sum \beta_{ij} x_i x_j + \epsilon$$

In the second phase, optimization is performed on the approximated function to determine values for the predictor

[5]http://www.metageek.net/products/wi-spy/

[6]http://iperf.sourceforge.net

variables that optimize the response. RSM has been previously used in the wireless networking context. For example, Vadde et al. [13] have used RSM to optimize the interaction between routing and MAC layers in mobile ad hoc networks.

For our purpose of understanding mutual interaction among 802.11n features, we limit our attention only to the first phase of RSM. Roughly speaking, our focus is on the statistically significant $x_i x_j$ terms in the above quadratic functional form. Specifically, we examine the ANOVA table resulting from the first phase of RSM and look at each pairwise combination of predictors to see if their p-value is lower than a desired level of significance (0.05) and if so, we conclude that the interaction between the predictors in that pair to be statistically significant. We use the SYSTAT tool[7] for our RSM based study. Note that we do not use RSM to study relative impact of different features because it does not have a suitable measure of importance like Pratt's measure with categorical regression.

3) Machine Learning Classifiers: For our study on interference type classification in section III.C, we consider four commonly used yet very different supervised machine learning classifiers: naive Bayes, multinomial logistic regression, k-nearest neighbors and C4.5 decision tree. Here we briefly describe each of them. The Naive Bayes classifier is based on the Bayes rule of conditional probability. It makes use of all the features contained in the data, and analyses them individually as though they are equally important and independent of each other. Multinomial logistic regression (also known as maximum entropy classifier) is commonly used as an alternative to Naive Bayes classifier as it does not assume statistical independence of features. C4.5 decision tree belongs to a family of decision tree algorithms that decide on the response (class in our case) for a new sample based on the values of features in the available data. The k-nearest neighbors algorithm classifies based on closest training examples in the feature space. We use the Weka data mining tool[8] for Naive Bayes, k-nearest neighbors (IBk in Weka) and C4.5 (J48 in Weka). For logistic regression, we use IBM SPSS tool [9].

III. Results

A. Baseline Results

We begin our characterization study by verifying that links in our testbed have diverse link qualities, thus allowing us to experiment over a whole spectrum of channel conditions. For this we pick a particular configuration of values of 802.11n features for every link: frame aggregation and channel bonding are enabled, STBC is disabled and the default rate adaptation algorithm with ath9k driver (minstrel_ht) is used. Note that the minstrel algorithm chooses between all 16 MCS values shown in Table I and also adapts the aggregated frame length size via a hardcoded table depending on the chosen MCS value.

We measure the RSSI and throughput of different links, one at a time and in the absence of any interference. Results are shown in Fig. 3, which confirm that our node placement results in sufficiently different link qualities and throughputs across links. Our long-term RSSI measurements for these links

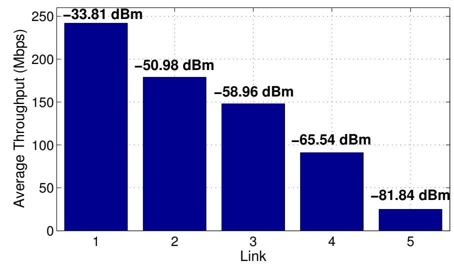

Fig. 3: RSSI and throughput variation across links in the testbed for a specific setting of values for 802.11n features.

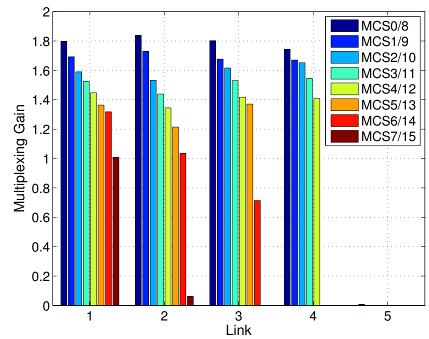

Fig. 4: Multiplexing Gain across different links and MCSs.

over several days (not reported here due to space limitations) additionally show that RSSI variation for each of these links remains within a few dB of the values shown in Fig. 3 even during day times when there is human mobility in the environment.

We now characterize the multipath environment in the testbed area and the opportunity available for spatial multiplexing by measuring the multiplexing gain [9] for different MCSs and links. For this experiment, we disable frame aggregation, channel bonding, STBC and the automatic rate adaptation algorithm. Results shown in Fig. 4 are along expected lines — multiplexing gain drops with worsening link quality and increasing modulation and coding rates.

Having done the confirmatory experiments focusing on link qualities and the environment, we now begin to look at the impact of other settings for 802.11n features while still using the minstrel rate adaptation algorithm and considering one link at a time with no interference. Since 3 features (frame aggregation, channel bonding and STBC) are considered each having two possible values (on/off), we have 8 possible configurations per link and 40 different configurations across all 5 links in the testbed. Results shown in Fig. 5 let us make certain observations such as frame aggregation is beneficial always regardless of link quality and channel bonding is helpful except when link quality is poor. We cannot assess the effect of choosing STBC versus SDM, however. This is due to the use of the automatic rate adaptation algorithm. But disabling the rate adaptation algorithm would mean introducing a new variable MCS (0-7) as discussed in the last section, which in turn has the effect of increasing the overall number of configurations

[7] http://www.systat.com/
[8] http://www.cs.waikato.ac.nz/ml/weka/index.html
[9] http://www.ibm.com/software/analytics/spss/

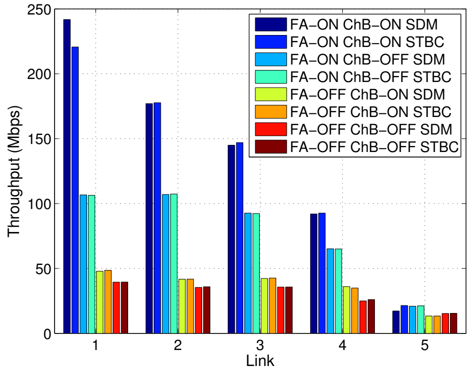

Fig. 5: Impact of frame aggregation, channel bonding and STBC/SDM on throughput with the default rate adaptation algorithm in use.

Scenario	Link 1	Link 2	Link 3	Link 4	Link 5
NI	0	0	0	0	0
CCI	0	$5 * 10^{-15}$	0	0	0
ACI	0	$2 * 10^{-10}$	0	$5 * 10^{-9}$	0

TABLE II: Significance of throughput categorical regression model for each of the scenarios.

Scenario	Link 1	Link 2	Link 3	Link 4	Link 5
NI	0	0	0	0	0
CCI	0	$8 * 10^{-11}$	0	0	0
ACI	0	0	0	0	$2 * 10^{-9}$

TABLE III: Significance of packet loss categorical regression model for each of the scenarios.

to analyze by 8-fold to 320 across all links. Note that this increase in possibilities is without having any interference in the scenarios; adding interference effects would further increase the possibilities by several fold, which motivates the need for an analysis approach that aids in easily understanding the impact of different features without the tedium of manually going through all possibilities.

B. *Throughput and Packet Loss Performance*

Led by the discussion at the end of the previous sub-section, we consider regression analysis as an effective approach to ease the characterization of the relative impact of different 802.11n features on performance in different scenarios. Given that the features under consideration are all categorical — nominal to be specific (e.g., frame aggregation ON or OFF) — categorical regression [2] is the most appropriate statistical analysis method for the problem at hand. See section II.D for a brief overview of categorical regression. For the regression based analysis in the rest of this sub-section and in section III.E, we use the widely used statistical analysis tool SPSS which implements categorical regression in the function named CATREG.

Now with the aid of categorical regression, it is less of a concern to expand the scenarios we consider to include interference effects. Specifically, we consider co-channel interference (CCI) and adjacent channel interference (ACI) conditions besides the no interference case that was the sole focus so far. This will effectively increase the number of scenarios being considered to 15 (3 types of interference x 5 different link qualities). To realize CCI and ACI effects we make use of 2 additional nodes mentioned in section II.A to create an interfering link belonging to a co-located 802.11n WLAN with a single station. To capture the worst case CCI and ACI effects, we place the interfering link in close proximity ($<$ 3 meters) to the AP shown in Fig. 1, effectively making every link under test in the testbed to experience strong CCI/ACI interference. To create CCI[10], we use same channel (149) for both the link under test (which can be one of the links between AP and stations (1)-(5) in Fig. 1) and the interfering link. To generate ACI, we assign adjacent channels to the link under test and the interfering link (channels 149 and 153, respectively when channel bonding is disabled and channels 149 and 157 otherwise). Similar measurement setup to create CCI/ACI effects was used in [8]. Note that we conduct the experiments which consider interference effects also during late nights to avoid human mobility related experimental noise. Also note that for CCI and ACI cases, we only focus on throughput and packet loss results for the link under test.

Results from applying categorical regression on throughput with respect to various 802.11n features (frame aggregation, etc.) for each of the 15 scenarios independently is shown in Fig. 6. Note that we show the results in terms of Pratt's importance measure (see section II.D) for natural interpretation of relative impact of various features. To associate confidence in the different regression models and verify their validity, we need to examine their significance levels, which need to be < 0.05 to be valid (as discussed earlier in section II.D). Significance levels for all throughput regression models are shown in Table II. Regression results and significance levels corresponding to the packet loss metric are shown in Fig. 7 and Table III, respectively. Note that significance levels for all models across both metrics satisfy the validity criterion. We have also manually verified this via detailed inspection of raw measurement results.

Note that in Figs. 6 and 7, "+" ("-") sign indicates positive (negative) impact of a variable on the metric in question. For example, "+" sign with frame aggregation for link 1 in Fig. 6 (a) means that enabling frame aggregation benefits throughput. Interpretation of these signs for MIMO and MCS are somewhat different. For MIMO, "-" sign suggests the use of STBC and SDM otherwise. Note that MCS in these figures represents only modulation and coding rates (8 different possibilities as shown in Table I) as already mentioned in section II.B and "+" sign for MCS suggests the use of higher modulation and coding rate and lower otherwise.

1) MCS Impact: From a throughput perspective, we observe that surprisingly MCS is the most important factor in the no interference case for all link qualities but its impact becomes negative with worsening link quality suggesting a

[10]We defer the study of another type of CCI where multiple links within the same WLAN are concurrently active until section III.E where we consider fairness.

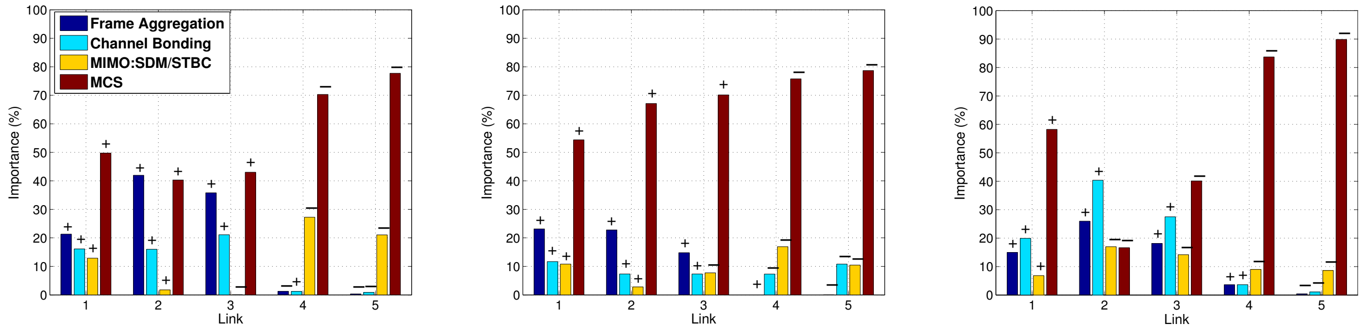

(a) No Interference (b) Co–Channel Interference (c) Adjacent Channel Interference

Fig. 6: Relative impact of 802.11n features on throughput performance in terms of Pratt's importance measure (section II.D) in different scenarios (channel quality and interference conditions).

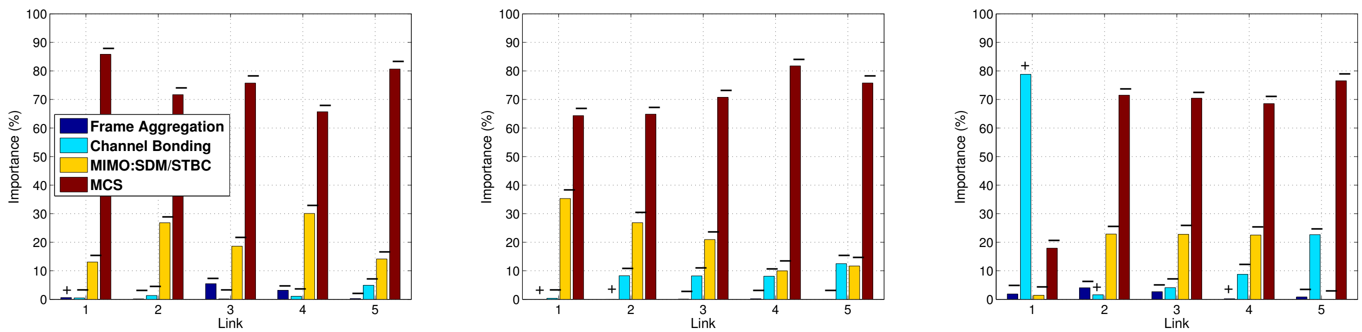

(a) No Interference (b) Co–Channel Interference (c) Adjacent Channel Interference

Fig. 7: Relative impact of 802.11n features on packet loss performance in terms of Pratt's importance measure (section II.D) in different scenarios (channel quality and interference conditions).

lowering of modulation and coding rate (Fig. 6 (a)). Latter is expected as higher modulation/coding rates require higher SNR which is not true for poor quality links. The importance of MCS in the CCI case relatively increases as a consequence of other factors such as frame aggregation becoming less important. In the presence of ACI, the effect of MCS becomes negative quickly with worsening link qualities because transmission activity from the interfering link in the adjacent channel increases the noise floor and reduces the SNR, making the lowered modulation and coding rate to be more effective. Results in Fig. 7 show that MCS is the dominant factor in reducing packet loss in almost all scenarios (combinations of interference types and link qualities) since lower modulation and coding rate increases link robustness.

2) MIMO (SDM/STBC) Impact: Benefit of SDM (which in our case is equivalent to using double streams) for throughput is limited to only very good quality links. Even that is less true in presence of ACI. STBC is more beneficial in such cases. This can be explained by the fact that transmission power is equally divided between the different antennas. With 2 antennas, this means power is effectively halved (reduction by 3dBm). Effective reduction of transmission power for each stream makes SDM more vulnerable when link quality is marginal. STBC is relatively more robust for marginal to poor quality links due to the redundancy it injects into a single stream. This also explains the beneficial effect of STBC on packet loss in most cases (Fig. 7).

3) Channel Bonding Impact: In terms of throughput, results show that use of channel bonding is less effective for poor quality links. This can be explained by a similar reasoning to that discussed for SDM vs. STBC above. Specifically, transmitting over a wider channel reduces per-subcarrier transmit power. This effect of channel bonding has also been observed in previous work [7], [8]. In presence of ACI, we observe that the importance of channel bonding relatively grows, especially for better quality links. This could be a result of the interfering link itself getting negatively affected because of the ACI from the link under test. From a packet loss perspective, channel bonding is generally harmful except for best quality link and ACI scenario when it has a very positive effect. We believe similar reason as the one just mentioned applies here as well.

4) Frame Aggregation Impact: Impact of frame aggregation on throughput is largely along expected lines — longer frames are not effective with worsening link quality. Results show that this fact holds even in presence of interference.

Packet loss results are relatively more interesting. They show that the impact of frame aggregation is almost negligble compared to that of MCS and STBC. This can be explained as follows. Because of the block ACK mechanism and selective sub-frame retransmission, frame aggregation can recover from losses efficiently. But on the other hand, using a high MCS or SDM could make losses go out of control of the ARQ mechanism with or without frame aggregation.

C. Differentiating Interference Types

Results in the previous sub-section suggest that relative impact of different features changes depending on the interference type a link experiences. For example, MCS becomes more important in presence of CCI whereas channel bonding influence grows in the ACI case. To exploit these observations in the context of a link adaptation mechanism, it is important for a sending node of a link to be able to detect the type of interference it is experiencing at a given point in time (this includes no interference case as well). Motivated by this, we assess the effectiveness of differentiating between interference types using a supervised machine learning classifier model. Classes for this model are different interference types (no interference, CCI, ACI) and feature vector could be $< FA, ChB, MIMO, MCS, Throughput >$. Using measurements like those we obtained for this characterization study, one could train a supervised machine learning classifier model providing it with feature vector and known class information. During the operational phase, a link could measure its current throughput and combine it with current settings for various 802.11n features and query that interference classifier model, which would output the most likely interference type, statistically speaking.

We implemented the above idea considering several different classifiers described in Section II.D. Results are shown in Table IV. We observe that logistic classifier provides with the best result with an average accuracy of 98% across all link qualities, suggesting that it is indeed possible for a sender node to indirectly infer the type of interference experienced by its links. Together with link quality measurements (which are relatively easier to obtain), the scenario of operation can be inferred and best settings for that scenario can be applied. Investigation of a 802.11n link adaptation algorithm based on this idea is a key aspect for future work. Also note that it is harder to infer the interference type for higher quality links, possibly because there are more combination of features that influence the performance differently when the link quality is better, making the feature space larger and classification harder.

Another interesting observation is that Naive Bayes classifier performs the worst among all the classifier algorithms considered. Since the assumption of independence between features in the feature vector is a unique aspect of Naive Bayes, its poor accuracy could be attributed to this assumption not holding true, which suggests interdependence among features. We explore this issue further in the next subsection.

D. Interdependence among 802.11n Features

To verify the existence of interaction among different features for optimizing throughput or packet loss metrics, we revisit the categorical regression results from section III.B and consider interpreting the importance values of each feature independently to choose an appropriate setting for that feature. For example, for ACI – Link 2 scenario, looking at Fig. 6 (c), we could take the importance values to imply choosing the

Classifier	Link 1	Link 2	Link 3	Link 4	Link 5
Logistic Regression	97.7%	97.7%	98.2%	99.1%	100%
k Nearest Neighbours	82.8%	87.1%	87.7%	88.8%	84.4%
C4.5 Decision Tree	63.8%	73%	80.3%	91.1%	82.2%
Naive Bayes	53.7%	53.9%	67%	86.76%	65.1%

TABLE IV: Interference type classifier accuracy for different link qualities and classifier algorithms.

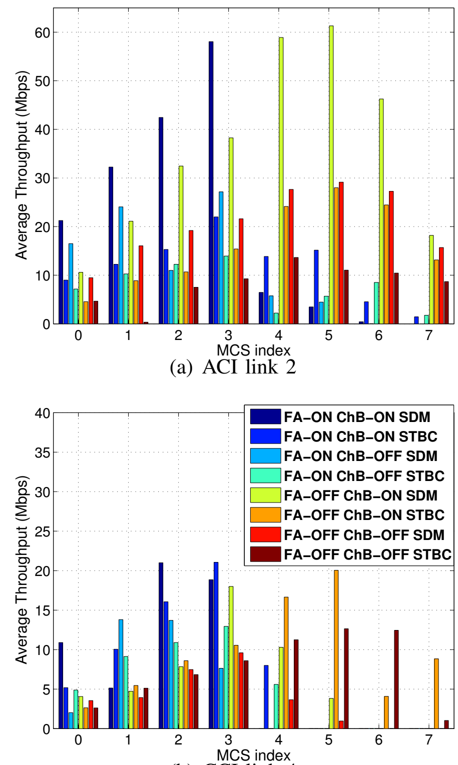

(a) ACI link 2

(b) CCI link 4

Fig. 8: Example scenarios for throughput suggesting interdependence among features.

features as follows: frame aggregation, channel bonding and STBC enabled, and a low to moderate modulation and coding rate. However, raw measurement results for this scenario shown in Fig. 8(a) suggest a different setting of features for obtaining optimal throughput: frame aggregation OFF, channel bonding ON, SDM and reasonably high modulation and coding rate. As another example, consider CCI - Link 4 scenario from Fig. 6 (b) and the corresponding raw measurement results showing the optimal configuration (Fig. 8(b)). Both these examples reinforce the fact that there exists potential interdependence among various features that prevents them to be treated in isolation when we aim to optimize performance.

In order to quantify whether interdependence among features exists we use the first phase of the response surface methodology (RSM) as described in section II.D to identify significant pairwise interactions among 802.11n features given a link scenario. Table V summarizes pairwise interactions found to be statistically significant after examination of the ANOVA table resulting from applying RSM. This table con-

	NI	CCI	ACI
Link 1	FA-ChB ChB-MIMO ChB-MCS FA-MIMO MIMO-MCS FA-MCS	FA-ChB ChB-MIMO ChB-MCS FA-MIMO MIMO-MCS FA-MCS	FA-ChB ChB-MIMO ChB-MCS FA-MIMO MIMO-MCS FA-MCS
Link 2	FA-ChB MIMO-MCS FA-MCS	FA-ChB ChB-MIMO MIMO-MCS FA-MCS	FA-ChB ChB-MIMO ChB-MCS MIMO-MCS
Link 3	FA-ChB MIMO-MCS	MIMO-MCS	ChB-MIMO MIMO-MCS
Link 4	ChB-MCS MIMO-MCS FA-MCS	MIMO-MCS	MIMO-MCS
Link 5	FA-ChB FA-MIMO MIMO-MCS	MIMO-MCS	MIMO-MCS

TABLE V: Pairwise interdependence among 802.11n features in different scenarios.

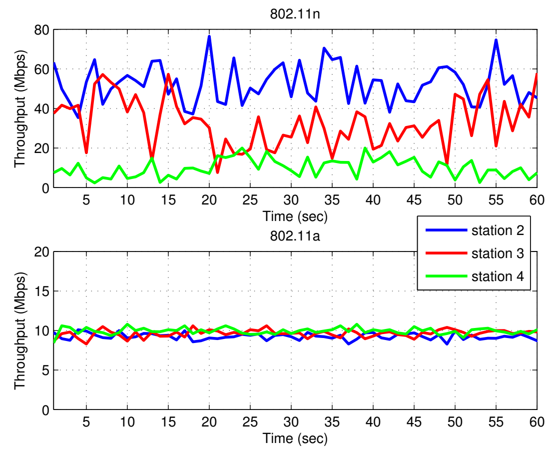

Fig. 9: Illustration of unfairness with 802.11n compared to 802.11a.

	Features					Interference	
	Number of Streams	STBC	ChB	MCS	FA	CCI	ACI
[4]	–	–	✔	✔	✔	✔	–
[5]	✔	–	✔	✔	✔	✔	–
[6]	–	–	✔	✔	✔	–	–
[7]	–	✔	✔	✔	–	✔	–
[8]	✔	✔	✔	✔	–	✔	✔
[9]	✔	–	–	✔	–	✔	–
[10]	✔	✔	✔	✔	–	✔	–

TABLE VI: 802.11n features varied and interference cases considered in previous work.

firms that there is interdependence among features in every scenario, indicating features that must be jointly selected. From a practical viewpoint, results in Table V suggest that for good quality links all features need to be selected together, whereas for marginal to poor quality links and in the presence of interference it is sufficient to consider interaction between only a subset of the features. This observation is consistent with regression results in Fig. 6, which show that few features have majority of the impact in poor quality links.

E. Fairness

Fairness in the context of 802.11n has not received much research attention. As an example to illustrate the fairness issue, Fig. 9 shows throughput share over time between links (2)-(4) in our testbed when they are simultaneously active in the 802.11n mode and 802.11a mode, respectively. Settings for this experiment are similar to the baseline results presented in Fig. 3. Unfairness with 802.11n is quite apparent, with higher quality links taking a greater share of the throughput at the expense of poor quality links. This is because higher quality links can use higher MCS values which in turn causes selection of larger aggregated frame sizes, thus making higher quality links occupy the channel for long periods.

Like with throughput and packet loss results (in Figs. 6 and 7), we have carried out categorical regression based analysis of the impact of different 802.11n features on fairness. For this we consider two scenarios, one where all links have similar and good link quality (see Fig. 10 a), and another corresponding to the node placement shown in Fig. 1. Results for fairness and aggregate throughput are shown in Figs. 10 (b) and (c). We see that relative impact of 802.11n features on fairness varies depending on the network scenario and is different from that on aggregate throughput. Due to space limitations, further discussion of these results is omitted.

IV. Related Work

The work of Shrivastava et al. [4] is the earliest attempt to analyze some of the key characteristics of 802.11n. Their study highlights the negative impact of legacy devices on 802.11n performance and also that channel bonding creates interference due to channel leakage. More recently, Deek et al. [8] have come to the same conclusion on the side effect of channel bonding. Both [8] and [7] focused on the impact of channel bonding. Arslan et al. [7] observed that channel bonding may be harmful even in the absence of interference for poor quality links; our results re-confirm this observation. Deek et al. [8] performed a thorough investigation of the impact of channel bonding under different types of interference and suggest the use of 20MHz channel separation in case of simultaneous transmissions between high quality links using 40MHz channels to counter the channel leakage issue mentioned above. Note that in [8] and [7], frame aggregation is disabled, which is a key shortcoming given that frame aggregation is a crucial feature affecting throughput in 802.11n networks [14].

Pelechrinis et al. [6] investigated the impact of packet size, channel width and transmission rate on 802.11n link performance. They observed that performance of SDM is highly sensitive with higher modulation and coding rates even in absence of interference, and that high rates are very susceptible to external interference and/or noise. They also suggest joint adaptation of MAC layer parameters as a way to avoid the throughput reduction with smaller packet sizes.

Pefkianakis et al. [5] discovered the non-monotonicity in the increase of throughput with increase in MCS index (Table

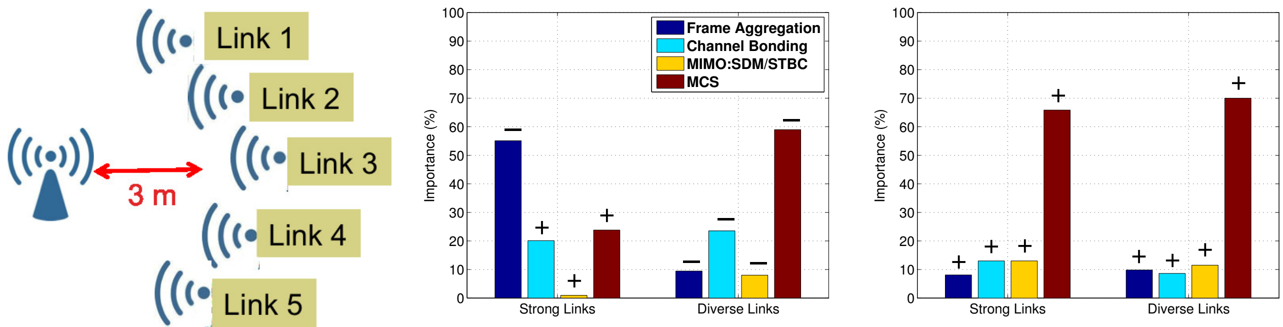

(a) All strong links (b) Impact on Fairness (c) Impact on Aggregate Throughput

Fig. 10: Relative impact of 802.11n features on fairness and aggregate throughput performance in different network scenarios.

I). They then proposed a new rate adaptation algorithm called MIMO Rate Adaptation (MiRA) algorithm that takes the non-monotonicity observation into account. Specifically, MiRA zigzags between single and double stream modes in the process of adapting the bit-rate. MiRA also use frame aggregation and block ACKs to detect and differentiate between collisions and channel errors. It does not, however, adapt the frame aggregation feature for the sake of throughput optimization.

More recently, Lakshmanan et al. [9] proposed a metric based on the multiplexing gain in order to adapt the bit-rate. Additionally, they observed that higher throughput can be achieved using a single stream for some link qualities and that 802.11n links are more susceptible to interference.

Nguyen et al. [10] proposed a 802.11n bit-rate adaptation algorithm termed RAMAS. It takes advantage of the monotonic relation between loss and modulation types to avoid random rate sampling and adapts the modulation scheme, number of streams, channel width and guard interval. They also consider fairness but their focus is limited to comparing RAMAS with other rate adaptation algorithms fairness. Table VI summarizes the 802.11n features studied or adapted and interference cases considered in previous work.

V. Conclusion

In this paper, we have experimentally studied how 802.11n features affect performance (throughput, packet loss and fairness) and interact with each other across a wide range of scenarios differing in channel and interference conditions. We employed categorical regression based analysis for easing characterization of relative impact of different features. We believe that this type of analysis should prove valuable even for other 802.11 standards in the making (e.g., 802.11ac). We have also assessed the potential pairwise interdependence among different 802.11n features in various link scenarios via response surface methodology. Our analysis showed that different features impact performance differently depending on the network scenario determined by channel and interference conditions; same is true about their mutual interaction. As a step towards practical and comprehensive 802.11n link adaptation, we showed the feasibility of identifying interference type online at sender side using throughput measurements and a supervised machine learning based classifier. We have also highlighted the unfairness problems peculiar to 802.11n. Expanding the scope of our characterization study to make it more comprehensive with additional scenarios and applications is an aspect for future work. Besides, in the future we intend to work on a detailed specification of holistic 802.11n link adaptation mechanism that incorporates interference differentiation component presented in this paper, leverages insights from our analysis and with wider applicability (e.g., multi-AP 802.11n network scenarios). We also plan to investigate the constrained optimization of throughput while ensuring fairness in the context of 802.11n link adaptation.

References

[1] IEEE 802.11n. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Enhancements for Higher Throughput, 2009.

[2] A. J. van der Kooij. *Prediction Accuracy and Stability of Regression with Optimal Scaling Transformations*. PhD thesis, Leiden University, 2007.

[3] R.H. Myers and D.C. Montgomery. *Response Surface Methodology*. John Wiley & Sons, Inc., 2002.

[4] V. Shrivastava et al. 802.11n Under the Microscope. In *Proc. ACM IMC*, 2008.

[5] I. Pefkianakis et al. MIMO Rate Adaptation in 802.11n Wireless Networks. In *Proc. ACM MobiCom*, 2010.

[6] K. Pelechrinis et al. Experimental Characterization of 802.11n Link Quality at High Rates. In *Proc. ACM WiNTECH*, 2010.

[7] M. Y. Arslan et al. Auto-Configuration of 802.11n WLANs. In *Proc. ACM Co-NEXT*, 2010.

[8] L. Deek et al. The Impact of Channel Bonding on 802.11n Network Management. In *Proc. ACM Co-NEXT*, 2011.

[9] S. Lakshmanan et al. On Link Rate Adaptation in 802.11n WLANs. In *Proc. IEEE INFOCOM*, 2011.

[10] D. Nguyen and J. J. Garcia-Luna-Aceves. A Practical Approach to Rate Adaptation for Multi-Antenna Systems. In *Proc. IEEE ICNP*, 2011.

[11] K. Pelechrinis et al. Design and Deployment Considerations for High Performance MIMO Testbeds. In *Proc. ACM WICON*, 2008.

[12] J. W. Pratt. Dividing the Indivisible: Using Simple Symmetry to Partition Variance Explained. In *Proc. 2-nd International Conference in Statistics*, page 245260, 1987.

[13] K.K. et al. Vadde. Optimizing protocol interaction using response surface methodology. *IEEE Transactions on Mobile Computing*, 5(6):627 – 639, June 2006.

[14] E. Perahia et al. *Next Generation Wireless LANs: Throughput, Robustness, and Reliability in 802.11n*. Cambridge University Press, 2008.