Examining Relationships Between 802.11n Physical Layer Transmission Feature Combinations

Ali Abedi
University of Waterloo
ali.abedi@uwaterloo.ca

Tim Brecht
University of Waterloo
brecht@cs.uwaterloo.ca

ABSTRACT

To increase throughput the 802.11n standard introduced several physical layer transmission features including a short guard interval, wider channels, and MIMO. Since obtaining peak throughput depends on choosing the combination of physical layer features (*configuration*) best suited for the channel conditions, the large number of configurations greatly complicates the decision. A deeper understanding of relationships between configurations under a variety of channel conditions should simplify the choices and improve the performance of algorithms selecting configurations. Examples of such algorithms include: rate and channel width adaptation, frame aggregation, and MIMO setting optimization.

We propose a methodology for assessing the possibility of accurate estimation of the frame error rate (FER) of one configuration from the FER of another. Using devices that support up to 3 spatial streams (96 configurations), we conduct experiments under a variety of channel conditions to quantify relationships between configurations. We find that interesting relationships exist between many different configurations. Our results show that in 6 of the 7 scenarios studied *at most five configurations* are required to accurately estimate the error rate *of all remaining 91 configurations* and in the other scenario *at most 15 configurations are required*. Although we show that these relationships may change over time, perhaps most surprising is that relationships have been found over periods of up to *one hour*. These findings suggest optimization algorithms should not need to measure the FER of many configurations, but instead can sample a small subset of configurations to accurately estimate the FER of other configurations. To demonstrate this possibility, we make simple modifications to the Minstrel HT rate adaptation algorithm to exploit relationships and observe improvements in throughput of up to 28%.

1. INTRODUCTION

Advancements in the 802.11 standard have made gigabit per second wireless communication possible by offering physical layer transmission features such as denser modulations, channel bonding, and MIMO, in addition to MAC layer frame aggregation. While these advancements help achieve higher data rates, efficient link adaptation in 802.11 networks is more challenging because of these options.

We use *rate configuration* (or simply *configuration*) to refer to a particular combination of physical layer transmission features such as the the modulation and coding scheme, channel width, short/long guard interval (SGI/LGI), and the number of spatial streams. The 802.11g standard offers only 8 configurations. That number has increased to as many as 128 in 802.11n networks (our work examines the 96 configurations available on devices with 3 spatial streams) and up to 640 in 802.11ac networks. Because the 802.11 standard does not specify how to choose physical layer transmission features, optimizing these choices is an active area of research. Examples of such research include: channel bonding [4], rate adaptation [9, 6], energy efficiency [17], QoS analysis [12], and STBC/SDM settings [9].

In rate adaptation studies, the combination of physical layer features used for transmission is chosen by sampling available configurations to determine their effective throughput. However, sampling can incur significant overhead [9] because probe packets are usually sent without frame aggregation. This conservative approach is used because probing often requires testing rates that may fail and the failure of a large number of frames that have been aggregated can negatively impact application performance.

In this paper, our hypothesis is that since several physical layer transmission feature combinations (rate configurations) share common features, (e.g., half use SGI and half use LGI), relationships may exist between the *average frame error rate (FER)* of different configurations. If it is possible to estimate the FER of one configuration from the measured FER of another configuration, algorithms that adapt configurations to changing channel conditions can be simpler and more effective. In this paper, we first develop a methodology for characterizing the relationship between the FER of different configurations. Then we conduct experiments in a variety of settings and report on the relationships we observe. The contributions of this paper are:

- We design a methodology for evaluating relationships among rate configurations that can be used in mobile environments with WiFi and non-WiFi interference.
- We characterize relationships under a variety of channel conditions and study changes in relationships over time. Interestingly, we find that large numbers of relationships exist, over surprisingly long periods of time,

MSWiM '16, November 13-17, 2016, Malta, Malta

DOI: http://dx.doi.org/10.1145/2988287.2989159

even in the presence of mobility and interference.

- Using our methodology, we find that LGI provides higher throughput than SGI in several scenarios. This is contrary to the notion that the LGI may not be required in indoor environments [5, 14, 15, 13].
- By using relationships between configurations, we demonstrate that it is feasible to improve throughput obtained using the Minstrel HT rate adaptation algorithm by up to 28%.

2. METHODOLOGY

Our relationship analysis methodology consists of the following phases: (1) collect data, (2) compute the FER for each rate configuration, and (3) compute relationships between the FER of different configurations. As demonstrated in Section 4, these steps can be used to characterize relationships between configurations.

2.1 Data Collection

To analyze the relationship between two rate configurations, the frame error rate (FER) of these rate configurations must to be measured under identical channel conditions. Hence, previous work [10] [9] has conducted experiments at night without any movement in the environment using the 5 GHz band, while also ensuring that the only interference is controlled interference (e.g., co-channel and adjacent channel interference). Additionally, theses studies use an unmodified rate adaptation algorithm. Therefore, only those rate configurations chosen by the RAA are examined and may not properly cover all configurations.

Previous work [2] has argued that repeating 802.11 experiments with identical channel conditions is difficult. More importantly, experiments from environments with only controlled interference and without mobility are unsuitable for understanding relationships between rate configurations in commonly encountered environments that include mobility and uncontrolled WiFi and non-WiFi interference.

In contrast to previous approaches, we design an experimental methodology for collecting FER information in any environment *and* that also properly covers all configurations. We have used a similar methodology previously [1] to collect representative traces for an 802.11 trace-based simulator. Our methodology does not require repeatability and can therefore be used in uncontrolled environments (including human movement and mobile devices operating in the 2.4 GHz band with WiFi and non-WiFi interference). With our technique, all configurations are sampled in a round-robin fashion. This process is continually repeated to collect information about changes in FERs over time.

Figure 1 shows a data collection example using a device with n rate configurations. Frames are sent with different configurations (denoted $R_1, R_2, ..., R_n$). The fate of each packet is denoted with 1 or 0, representing success or failure. Each sequence of n sampled configurations forms a round.

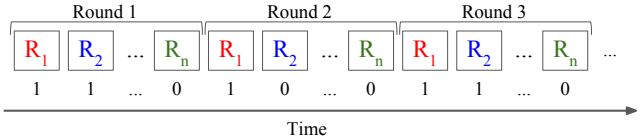

Figure 1: Round-robin data collection methodology

Since configurations in a round are subject to the same channel conditions (they are in the same channel coherence window), when interference does occur all configurations in a round experience the same conditions [2]. Changes, on average, impact each configuration equally. We implement round-robin sampling by modifying the device driver of the sending device used to collect data.

2.2 Frame Error Rate (FER) Computation

We now determine the number of packets required to compute the average FER. We use the following formula for calculating the minimum number of samples required to determine the population mean with a specified level of confidence, when the population standard deviation is known:

$$k > \left(\frac{z * \sigma}{MOE}\right)^2 \quad (1)$$

For a 95% confidence level, $z = 1.96$. The sample size k is maximized when the standard deviation (σ) is maximized. The value of σ is maximized (i.e., $\sigma = 0.5$) when half of the frames fail and the other half succeed. Using a 10% margin of error (MOE) and confidence level of 95% the minimum sample size required is 97 frames.

Since it takes approximately 43 ms to complete a round for all 96 rate configurations and we need a minimum of 97 observations (i.e., rounds), it takes about 4.2 seconds to collect enough samples to compute an average FER. If the channel access is delayed by WiFi and non-WiFi interference, it takes more than 43 ms to complete a round and more than 4.2 seconds to conduct enough observations to compute the FER. Therefore, we calculate the average FER using a 10 second window, this means the number of samples used in all experiments is 232.

2.3 Relationships and their Computation

Many methodologies could be used to assess the relationship between two rate configurations. We first define what we mean by *relationship* and then describe the methodology used in our study. Section 8 describes several methodologies that seem appropriate but are not suitable. Note that there exist several different connotations of the term relationship and what a relationship between rate configurations might mean. It is important to understand that for the purposes of this paper, we are concerned strictly with the relationships as denoted by the following definition.

2.3.1 Relationships

Relationship Definition:
We say that there exists a relationship between rate configurations R_1 and R_2 ($R_1 \mapsto R_2$) if the FER of rate configuration R_1 can estimate the FER of rate configuration R_2 with some expected degree of accuracy. In this case, we call R_1 the *estimator* and R_2 the *estimated* configuration. Note that relationships may or may not be reflexive.

To provide the intuition behind our methodology, we first present an example of a relationship analysis between two transmission rate configurations. Using two stationary devices and the 2.4 GHz band, we collect data for all 96 rate configurations and compute the average FER over 10 second windows using the techniques described in Sections 2.1 and 2.2. Figure 3 shows the FERs for two of the 96 configurations, namely 2S-I6-LG-20M=104 (configuration R_1) and 2S-I7-LG-20M=117 (configuration R_2). The transmission rate configuration notation is described in Figure 2.

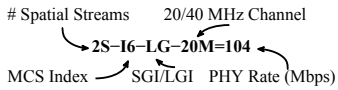

Figure 2: Transmission rate configuration notation

The changes in the FER over 30 minutes can be seen in Figure 3. One can see that the FERs of these two rate configurations seem to change together, suggesting the existence of a relationship between them.

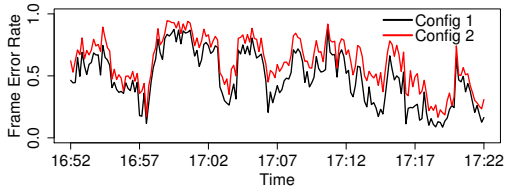

Figure 3: FER of two rate configurations changing over time

Another way to examine the relationship between the FERs of two rate configurations (irrespective of time) is to use a scatter plot with the FERs for one configuration along the x-axis and the FERs of the other configuration along the y-axis. We remove the time component because our goal is to determine relationships that persist over time *even in the presence of changing channel conditions*. If at a point in time, t, in Figure 3 the FERs of configurations R_1 and R_2 are e_1 and e_2, respectively, these two points are represented on the scatter plot (Figure 4) by one point with x and y values of e_1 and e_2, respectively.

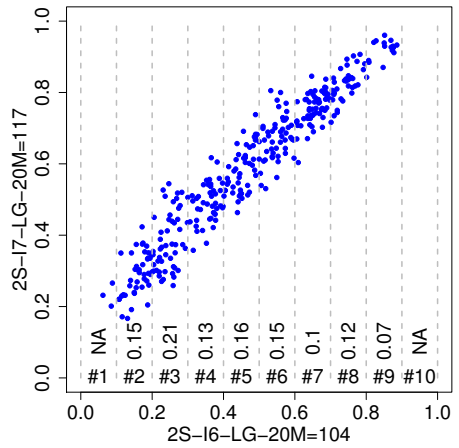

Figure 4: FERs, bins and estimation power for $R_1 \mapsto R_2$

Rate configuration R_1 can *accurately estimate* rate configuration R_2 if the FER values of R_1 are mapped to *a relatively small range* of FERs of R_2. To determine whether or not rate configuration R_1 is an estimator for rate configuration R_2, we divide the data in the scatter plot into bins, illustrated by dashed lines in Figure 4. For our analysis, we have chosen 10 bins, other values are possible and this is discussed in Section 8. Ideally, the dispersion of points (i.e., variation) in the vertical dimension is low in each bin, indicating that a reasonably accurate estimation of the FER of rate configuration R_2 from FER of rate configuration R_1 is possible. In this instance, an accurate estimate the FER of R_2 when the FER of R_1 is 0.85 is possible, because the dispersion of the points in bin #9 is relatively low. However, accurate estimation of the FER of R_2 when the FER of R_1 is 0.25, is not possible, because the corresponding FERs in bin #3 for R_2 have a fairly wide range (0.22 – 0.57).

2.3.2 Estimation Power

Ideally, we would like to summarize the strength of the relationship between two rate configurations with a single quantity. We start by quantifying the variation of data points within each bin. Statistical dispersion, which determines how "stretched" or "squeezed" the distribution of data points are, is one such suitable measure. There are several measures of statistical dispersion that could be used. Range and standard deviation are two common measures. However, these measures are highly sensitive to outliers, which may be common because of potentially high variation in frame error rates over time. On the other hand, other measures such as mean absolute and quartile deviation are more robust to a small number of outliers. We have considered several robust measures of dispersion including *median* absolute deviation (MAD), *mean* absolute deviation, and interquartile range (IQR). In this paper, we have chosen to use the interdecile range which is the difference between the first and the ninth deciles (i.e., the first 10% and last 90%). This measure provides the characteristics desired for this study such as excluding some but not too many outliers. We found that other measures can provide undesirable or misleading values (discussed in more detail in Section 8).

We now provide an example of how we apply the interdecile range to the FER data to obtain a single quantity that represents the strength of the relationship between two rate configurations. As depicted in Figure 4, in each bin the vertical dispersion of the values is calculated based on the interdecile range which is written above the bin number. The interdecile range values quantify the dispersion in each bin and provide a measure of the relationship between two rate configurations based the variation of FER in each bin. Note that bins with fewer than 5 values are ignored (labeled as NA, for not applicable), as they do not contain enough data points to provide a reliable measure of dispersion. We describe the importance of bins labeled NA and how we account for them in Section 2.3.3.

To examine relationships between 96 × 96=9,216 pairs of configurations, we define *estimation power* which aggregates dispersion values from all bins for each pair of configurations.

Estimation Power: EP(R_1,R_2)
The estimation power of a relationship between rate configurations R_1 and R_2 is a measure of the expected ability of the FER of R_1 to estimate the FER of R_2. It is calculated as the fraction of bins with an interdecile range below a specified threshold.

The total number of bins excludes those that do not have a sufficient number of data points for the interdecile range to be deemed reliable (we use 5). We use a threshold of 0.2 and discuss both choices in Section 8. In the example data in Figure 4, the estimation power of rate configuration R_1 to estimate rate configuration R_2, EP(R_1,R_2)= $\frac{7}{8}$.

2.3.3 Variability Indicator

The estimation power (EP) is valuable for quantifying the relationship between two rate configurations. However, we found it beneficial to be able to differentiate types of relationships based on why they exist. We define a new metric called the *Variability Indicator* (VI) that quantifies this variation. Note that the variability indicator is not used to quantify the relationship nor to indicate the lack of a relationship but to interpret and understand the EP.

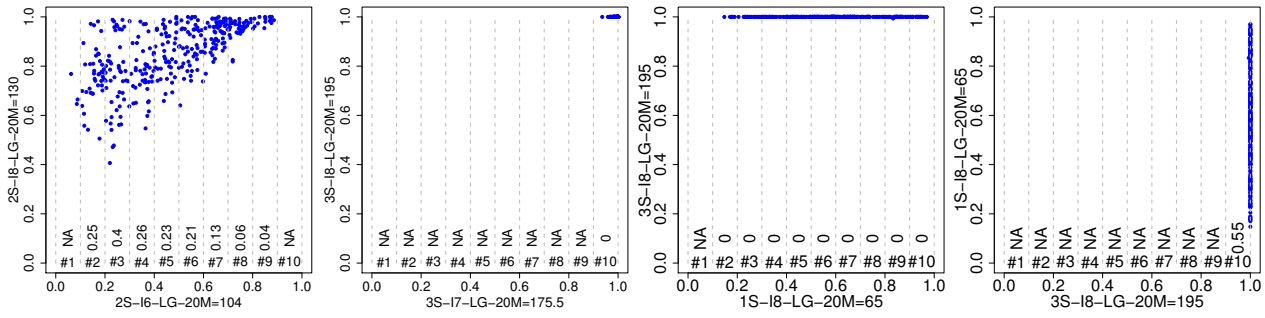

(a) EP(x, y)=$\frac{3}{8}$, VI(x)=0.55, VI(y)=0.29

(b) EP(x, y)=1, VI(x)=0.01, VI(y)=0.00

(c) EP(x, y)=1, VI(x)=0.55, VI(y)=0.00

(d) EP(x, y)=0, VI(x)=0.00, VI(y)=0.55

Figure 5: Estimation Power (EP) and Variability Indicators (VIs)

Variability Indicator (VI):
The variability indicator is a measure of the variability of the frame error rate of a rate configuration. The metric we use is the interdecile range of the FERs of a given configuration over the course of an experiment.

As discussed in Section 2.3.2, interdecile range is a suitable metric for quantifying the variation or dispersion of frame error rates. The variability indicator helps us to better understand the underlying reasons for the existence or nonexistence of a relationship between two configurations. It is important to note that *relationships can exist regardless of the value of the variability indicator*. In Section 4.1, we use this measure to help explain the estimation power.

2.3.4 Understanding Our Metrics

To better understand the estimation power metric and the need for the variability indicator metric, we review some scatter plots for other pairs of configurations from the same experiment as used in Figure 4. The caption for each subfigure shows the values for the estimation power and variability indicator metrics. The EP metric indicates the ability of the rate configuration on the x-axis to estimate the FER of the configuration on the y-axis. The VI metric is shown for the configurations on both the x and y axes.

Figure 5a shows that within several bins the vertical dispersion of FERs is relatively high. This results in the low estimation power metric ($\frac{3}{8}$) which means that the x-axis configuration (2S-I6-LG-20M=104) is not able to estimate the y-axis configuration (2S-I8-LG-20M=130). Figure 5b and 5c show an example of one rate configuration (3S-I8-LG-20M=195, on the y-axis) with a constant FER (i.e., VI(y) = 0.00). The constant FER makes it possible to accurately estimate this configuration from all other configurations regardless of their variation in FER. For example, the estimator configuration in Figure 5b has low variation (i.e., VI(x) = 0.01), while the estimator configuration in Figure 5c has relatively high variation (i.e., VI(x) = 0.55).

Figure 5d demonstrates why an estimator configuration with a constant FER (i.e., FER is always 1) cannot estimate a configuration with highly variable FER. In this case, highly dispersed data points are all gathered in one bin (i.e., #10) making an accurate estimation impossible, as indicated by the EP value of 0. To emphasize that EP is a directional metric, Figure 5d shows the same configurations as Figure 5c, except we have switched the estimator and estimated configurations. As discussed in Section 8, symmetric measures that are oblivious to the direction of the relationship (e.g., correlation coefficient and R^2 obtained from a statistical regression), are not suitable for this study.

3. EXPERIMENTAL ENVIRONMENT

We have created a small test bed for conducting experiments. This test bed is housed within lab and office space in a building on a university campus.

Our access point and stationary clients are desktop systems, each containing a TP-Link TL-WDN4800 dual-band wireless N PCI-E adapter. These cards use the Atheros AR9380 chipset and support up to three streams (i.e., a 3x3:3 MIMO configuration). This device uses the Ath9k (Atheros) device driver. For mobile experiments, we use a laptop configured to use a TP-Link TL-WDN4200 dual-band wireless N USB adapter. This adapter contains an Ralink RT3573 chipset and also supports a 3x3:3 configuration. This device uses the rt2800usb (Ralink) device driver.

To fully utilize the network infrastructure, we use iperf [7] to generate UDP traffic from the access point to a client at as high a packet rate as possible. We have modified the Ath9k device driver to implement a rate configuration selection algorithm that transmits using each configuration in a round-robin fashion as explained in Section 2.1. To record much of the information reported in this study, we use highly detailed information obtained directly from the Ath9k driver. Since we are interested in physical layer relationships, MAC layer frame aggregation is disabled to increase the efficiency of the data collection mechanism. In future work, we plan to investigate frame aggregation.

3.1 Different Scenarios Studied

We conduct experiments under a variety of channel conditions including stationary and mobile devices both with and without interference. In some experiments, we use the 5 GHz band to examine channels that are free of interference. In this case, we use a spectrum analyzer to ensure that there is no WiFi or non-WiFi interference. For other experiments, we use the 2.4 GHz band to ensure that the channel is exposed to different types of WiFi and non-WiFi interference. We intentionally selected channel 6 which overlaps with the channel used by the campus WiFi network to test our ability to study relationships in uncontrolled environments.

In mobile experiments, a laptop (i.e., receiver) is carried at walking speed for 15 minutes. The mobile experiments are conducted in two environments (referred to as *office* and *hallway*) which we designed to exercise a variety of channel

conditions. In the office environment, no line-of-sight exists between the AP and client for most of the experiment, as the signal is blocked by obstacles such as metal cabinets, cubicle partitions and walls. In these experiments, the distance between the AP and client ranges from 1 meter to about 20 meters. In the hallway experiments, a line-of-sight exists between the AP and client, and the distance between them changes from 1 meter to 60 meters.

In the stationary experiments, the AP and client are placed in different rooms in an office environment with no line of sight. All experiments are conducted during the day with people moving in and between offices (this can cause signal attenuation and influence multipath propagation).

To better understand the experimental scenarios, we present some statistical characteristics of the collected data. We classify each of the 96 rate configurations into three categories. The first two, FER < 0.1 and FER > 0.9, indicate that all frame error rate measurements for that configuration are bounded by these values. The variability indicator is low for these categories. The final category captures all configurations that do not fit into the first two categories. Table 1 shows, for each scenario, the number of configurations in each category. We observe that in the stationary 5 GHz experiment, which is our most stable environment, the FER of 71 configurations (out of 96) are either less than 0.1 (i.e., most frames are received successfully), or greater than 0.9 (i.e., most frames are lost). The same stationary experiment (with constant transmission power) using the 2.4 GHz band shows different behavior due to WiFi and non-WiFi interference; even the most robust modulation and coding schemes experience errors in presence of interference. Moreover, 11 configurations nearly almost always fail in this scenario.

	Office					Hallway	
	Stationary			Mobile		Mobile	
Scenario	1	2	3	4	5	6	7
Band (GHz)	2.4	5	5†	2.4	5	2.4	5
FER < 0.1	0	51	0	0	0	0	0
FER > 0.9	11	20	24	0	4	12	9
$0.1 \leq$ FER ≤ 0.9	85	25	72	96	92	84	87

Table 1: Characteristics of scenarios († = TX power cycling)

To introduce more variability in the FER for the 5 GHz stationary experiment and to increase the variability indicator in that scenario, we conduct another experiment where the transmission power is changed. Transmission starts at the default maximum setting of 30 dBm and is decreased by 1 dBm every 30 seconds until it reaches 0 dBm. It then increases transmission power by 1 dBm until it reaches 30 dBm and repeats in a round-robin fashion. The table shows that in this experiment the FER of the majority of rate configurations are now variable and even the most robust configurations experience some errors.

As mentioned previously, the mobile experiments were designed so that a variety of channel conditions are observed during the experiment. The data in Table 1 confirms that a majority of configurations experience a variable FER during the experiment. Note that in three of these four scenarios, the FER of some configurations are constantly above 0.9, even though the distance between the AP and the mobile client is about one meter during some points in the movement trajectory. A closer inspection of the raw data showed that these are the rate configurations that result in the highest physical data rates which never find the perfect channel conditions they need.[1]

4. CHARACTERIZATION RESULTS

In this section we utilize the proposed methodology to examine relationships among 802.11n rate configurations.

4.1 Examining Relationships

We first examine the *estimation power (EP)* and *variability indicator (VI)* of different configurations. Since 96 configurations are supported in our 802.11n cards, 9,216 (i.e., 96×96) relationships can be examined in each experiment.

Figure 6a illustrates the relationships between all 96 configurations for the 5 GHz office scenario with stationary devices (Scenario 2). The large square heat map (the EP heat map) shows the estimation power of a configuration on the x-axis for estimating the FER of the configuration on the y-axis. Each row and column in this heat map represents a rate configuration. Note that there are 96 configurations on each axis, but labels are removed as they are unnecessary for the high level view we start with. Later we present subsets of such heat maps in order to have a closer look at some particular relationships. The colors encode ranges for the estimation power (EP) as follows: high ($EP \geq 0.7$) in green, medium ($0.5 \leq EP < 0.7$) in yellow and low ($EP < 0.5$) in red. As depicted in Figure 6a the estimation power of all pairs of configurations are very high. To understand why, we study the variability of the FER of these rate configurations.

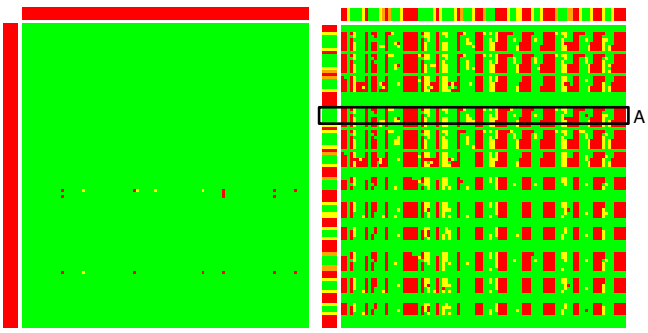

(a) Constant Power (b) TX Power Cycle

Figure 6: Office: stationary, 5 GHz, 96x96

The thin heat maps at the top of and to the left the EP heat map represent the variability indicator for the estimator and estimated configurations and will be referred to as VI heat maps. The two VI heat maps are always identical but they are shown on both axes for readability. For the VI heat maps the colors encode ranges for the variability indicator as follows: very high ($VI > 0.75$) in green, high ($0.5 < VI \leq 0.75$) in yellow, medium ($0.25 < VI \leq 0.50$) in orange, and low ($VI \leq 0.25$) in red.

Interestingly, in Figure 6a, the variability indicator is quite low for all rate configurations. This indicates that the FERs do not change significantly. In this scenario, the main reason for so many strong relationships is because the FER of all configurations are mainly constant and are, therefore, easy to estimate. Note that in some of these cases one configuration may consistently fail (e.g., 3S-I8-SG-40M=450)

[1]This may be because, at short distances, the multipath field is insufficiently rich for the three antennas in the small USB WiFi adapter.

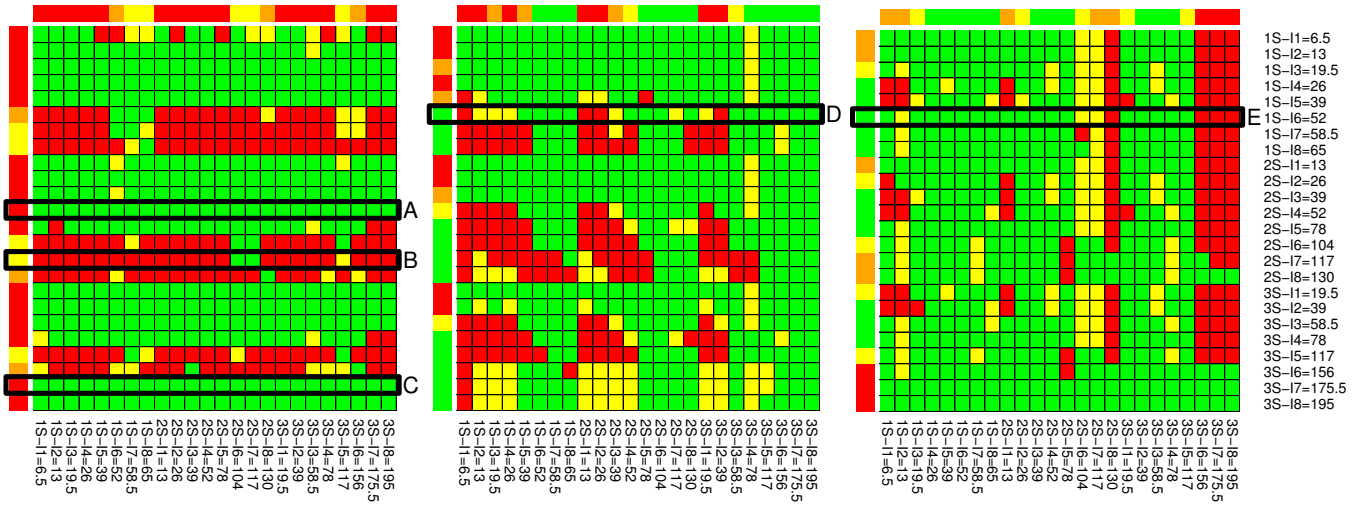

(a) Stationary, 2.4 GHz (b) Mobile, 2.4 GHz (c) Mobile, 5 GHz

Figure 7: Office, 24x24

while the other is consistently successful (e.g., 1S-I1-LG-20M=6.5). Nevertheless, this is a strong relationship.

In an attempt to see if any relationships exist when the FERs are variable, we program the AP to change transmission power as described in Section 3.1 (Scenario 3). The results are presented in Figure 6b. The differences between the VI heat maps in Figure 6a (which contain only red) and Figure 6b (which also contain some green, and some yellow) indicate that the changes in transmission power increase the variability of FERs. The EP heat map in Figure 6b shows that many rate configurations are still strongly related despite some FERs being highly variable. For example, although the few configurations in the rows outlined by the rectangle labelled "A" have highly variable FERs as depicted by green cells in the VI plot, we observe that there are many green cells in these rows in the EP heat map. Each green EP heat map cell indicates that the corresponding x-axis configuration can accurately estimate the y-axis configuration. As can be seen in the rectangle "A", many rate configurations can accurately estimate the y-axis configurations in those rows. Note that despite changing the transmission power, the variability of the FERs of some configurations is still low. By examining details of the collected FER data, we find that these configurations consistently fail in Scenario 2 which uses the maximum transmission power. As a result, lowering the transmission power (in Scenario 3) does not affect the FER of these configurations.

To examine some relationships in more detail, we now consider a subset of all rate configurations in the office scenarios. Figures 7a, 7b and 7c, show relationship results for the stationary (2.4 GHz) and mobile (2.4 and 5 GHz) scenarios, respectively (i.e., Scenarios 1, 4, and 5). In these scenarios, we examine only those configurations that use a long guard interval (LGI) and 20 MHz channels. This restricts the number of configurations to 24 (8 MCSes $\times$ 3 spatial streams) and the heat maps to 24×24 pairs of configurations.

In all graphs in Figure 7, we see patterns of colors suggesting the existence of relationships between estimation power and combinations of physical layer transmission features. The green cells (indicating a high estimation power) are not scattered randomly on the EP heat maps, but rather clustered in specific patterns based on the physical layer configurations indicated on the x and y axes. We now examine these plots in more detail to better understand these results.

In the stationary scenario, by comparing the relationships in the 2.4 GHz and 5 GHz bands (Scenario 1 in Figure 7a and Scenario 2 in Figure 6a), we observe more red cells in the 2.4 GHz scenario indicating a decrease in the number of related configurations. The major difference between these experiments is the lack of interference in the 5 GHz band. Figure 7a shows that some configurations, such as 2S-I4=52 (see rectangle "A"), can be estimated accurately by many configurations. The variation of the FER of this configuration is low (i.e., indicated by the red cell in the left VI heat map). Therefore, it is easy to estimate. On the other hand, configurations such as 2S-I7=117 (see rectangle "B") that have higher variation of FER (indicated by the yellow square in the left VI heat map), can be more difficult to estimate. In Figure 7a, the variability indicator is low for two groups of configurations: (a) configurations that consistently succeed (in this scenario) such as 2S-I4=52 (rectangle "A") and (b) configurations that consistently fail (in this scenario) such as 3S-I7=175.5 (rectangle "C"). Other configurations, which are not in either of these groups, experience variable FER such as 2S-I7-117 (rectangle "B").

Figures 7b and 7c show the relationships for a scenario where the client (i.e., receiver) device is moving at walking speed in an office environment using the 2.4 and 5 GHz bands (Scenarios 4 and 5). We consider these scenarios to study the effect of mobility on relationships. These are highly variable environments due to mobility. This can also be seen by the number of green cells in the VI heat maps when compared to their stationary counterparts. Figures 7b and 7c illustrate that strong relationships exist between many rate configurations. Interestingly, in both scenarios, for the estimated (y-axis) configurations with highly variable FERs (i.e., green or yellow cells in the left VI heat map), there are several (x-axis) configurations that can accurately estimate them. For example, configuration 1S-I6=52 (rectangles "D" and "E") has a highly variable FER in both scenarios as indicated by the corresponding green cells in the left VI heat maps. However, as depicted in Figures 7b, and 7c there are several green cells in the 1S-I6=52 rows ("D" and "E"), indicating the existence of several estimators for

		Office: Stationary						Office: Mobile				Hallway: Mobile			
		Scenario 1		Scenario 2		Scenario 3†		Scenario 4		Scenario 5		Scenario 6		Scenario 7	
		2.4 GHz		5 GHz		5 GHz		2.4 GHz		5 GHz		2.4 GHz		5 GHz	
		60 minutes		60 minutes		60 minutes		15 minutes		15 minutes		15 minutes		15 minutes	
EP Threshold	Stat	$\lvert*\mapsto R\rvert$	$\lvert R\mapsto*\rvert$	$\lvert*\mapsto R\rvert$	$\lvert R\mapsto*\rvert$	$\lvert*\mapsto R\rvert$	$\lvert R\mapsto*\rvert$	$\lvert*\mapsto R\rvert$	$\lvert R\mapsto*\rvert$	$\lvert*\mapsto R\rvert$	$\lvert R\mapsto*\rvert$	$\lvert*\mapsto R\rvert$	$\lvert R\mapsto*\rvert$	$\lvert*\mapsto R\rvert$	$\lvert R\mapsto*\rvert$
0.7	Min	1	32	88	91	18	33	6	4	13	19	11	19	4	17
	Max	95	66	95	95	95	95	61	77	95	59	95	80	95	92
	Min SC	5 – 11		1		1		3		2		3		2	
1.0	Min	1	19	88	91	1	33	6	3	13	19	10	19	3	17
	Max	95	41	95	95	95	63	60	68	95	59	95	68	95	55
	Min SC	5 – 15		1		5		4		2		3		3	

Table 2: Summary of relationships. † indicates TX power cycling

this configuration. The red cells in the 1S-I6=52 rows correspond mostly to the estimators with low variability FERs (i.e., indicated by red cells in the top VI heat map), as it is difficult for their relatively constant FER to estimate a variable FER. We have highlighted a few interesting scenarios and now present an overview of the relationships.

4.2 Overview of Relationships

To provide a high level view of the number and strength of different relationships, we summarize the results from all experiments in Table 2. One measure of interest presented in this table is the count of the number of other rate configurations that can be used to estimate a particular rate configuration R, denoted $|*\mapsto R|$. Another metric we present is the count of the number of other rate configurations that a particular rate configuration R can estimate, denoted $|R\mapsto*|$. To quantify if R_1 can estimate R_2, we use a threshold for the estimation power. In the heat maps shown in this paper we have used a threshold of 0.7 to indicate a very strong relationship. As a result, we also use this threshold when computing the data presented in Table 2. The intuition is that if the estimation power of R_1 when estimating R_2 is greater than or equal to 0.7, we presume that R_1 can estimate R_2. In addition, to study the impact of that threshold on our results we also include computations using the most conservative threshold possible (1.0). Recall that this means that the dispersion metric in *all bins* must be no greater than the dispersion threshold (i.e., 0.2).

Table 2 shows a variety of information for the seven scenarios examined. For each scenario, we present the Min, Max and Min SC values for the defined measures $|*\mapsto R|$ and $|R\mapsto*|$. The Min and Max values are of interest in understanding the number of relationships that exist between different configurations (Min SC will be described later). The greater the number of relationships, the more likely we are to be able to accurately estimate the FER of one configuration from the FER of one or more other configurations.

We start by focusing on the values obtained with a threshold of 0.7. In the worst case (i.e., the lowest value) across all scenarios, there is only one estimator ($|*\mapsto R|$) for a particular rate which occurs in Scenario 1 (see the row labelled "Min" and column $|*\mapsto R|$). When examining the number of configurations that can be estimated by a single rate ($|R\mapsto*|$), the minimum value is 4, which occurs in Scenario 4 (see the row labelled "Min" and column $|R\mapsto*|$).

We now consider the minimum number of rate configurations that are required to be able to estimate (i.e., cover) all other rate configurations. This problem is equivalent to the *minimum set cover* problem which is NP-hard [8]. Interestingly, the minimum set cover in six (i.e., Scenarios 2 – 7) of the seven scenarios was small enough (the largest of these was 3), that we were able to find them using a brute force approach that examines all sets up to size 4. These small values indicate that there are a number of strong relationships that could potentially be exploited. In Scenario 1, we have used a heuristic search to determine that the set cover size is between 5 and 11. The results for all scenarios are shown in the row labelled Min SC in Table 2.

We now focus on the values obtained in Table 2 using the most conservative threshold of 1.0. Comparing the values obtained using the two different thresholds shows that in most instances the values do not change significantly. Furthermore, the size of the minimum set cover is unchanged in three of the seven scenarios and increases only slightly in the other four. Note that the precise numbers in these tables depend on the various parameters used to determine whether relationships exist or not. The impact of these choices is discussed in Section 8.

It is interesting to note that there are large numbers of strong relationships between a variety of different configurations in the experiments conducted using the interference free, 5 GHz band. This can be seen in Figure 6 and from the data provided in Table 2. Perhaps more compelling, however, are the surprisingly large numbers of relationships in the 2.4 GHz and mobile experiments. We believe that these are intriguing results considering: (1) The complexity of the interactions of the large number of possible configurations of physical layer features. (2) The constantly changing channel conditions due to WiFi and non-WiFi interference in the uncontrolled 2.4 GHz experiments (Scenarios 1, 4 and 6). (3) The continual fluctuations in signal quality due to mobility (Scenarios 4 – 7). (4) The relatively long period of time over which these relationships hold (e.g., data for the stationary experiments covers 1 hour).

4.3 Changes in Relationships Over Time

To study if and how relationships might change over time, we start with a simple demonstration by dividing the one hour stationary 2.4 GHz experiment (i.e., Scenario 1) into four 15 minute experiments and perform the same statistical analysis over each time window. Table 3 presents the relationship data for each sub-window (EP threshold = 0.7). We observe that various measures of relationship change with time. For example, the size of the minimum set cover which was calculated to be between 5 and 11 over the full 1 hour experiments is 1, between 5 and 7, 5, and 1 for the four 15 minute sub-windows. These are significant reductions that demonstrate that a greater number of strong relationships may exist over shorter time intervals and that relationships between rate configurations vary over time. The time window over which relationships will be computed in practice will depend on the application in which they are being used.

However, because we have found evidence that relationships can be found over periods of 15 to 60 minutes in the scenarios we have examined, applications should not have to recompute relationships too frequently.

	0–15 Min		15–30 Min		30–45 Min		45–60 Min		Overall	
Stat	$\|*\mapsto R\|$	$\|R\mapsto *\|$	$\|*\mapsto R\|$	$\|R\mapsto *\|$	$\|*\mapsto R\|$	$\|R\mapsto *\|$	$\|*\mapsto R\|$	$\|R\mapsto *\|$	$\|*\mapsto R\|$	$\|R\mapsto *\|$
Min	5	40	1	43	3	19	10	40	1	32
Max	95	75	95	77	95	72	95	95	95	66
Min SC	3		5 – 7		5		1		5 – 11	

Table 3: Changes in relationships over time, scenario 1

5. STUDYING PHY LAYER FEATURES

In this section, we use our methodology to better understand the efficacy of short and long guard intervals and then analyze the relationships that exist because of this feature.

5.1 Efficacy of LGI and SGI

Some people have argued that, for indoor environments, the 800 *ns* guard interval (LGI) used in 802.11 protocols prior to 802.11n was more conservative than necessary [5, 14, 15, 13]. As a result, the shorter 400 *ns* guard interval (SGI) was introduced in the 802.11n standard. Because we are not aware of any empirical studies that examine if there is a need for the LGI in indoor 802.11n networks, we utilize our methodology to study this issue.

Suppose we examine the FER of all rate configurations that differ only in whether they use LGI or SGI (i.e., other features are the same). If their FERs are nearly identical, it indicates that shrinking the guard interval from 800 to 400 *ns* for these particular rate configurations does not have an adverse effect on the FER in the scenarios examined. Therefore, SGI should be used instead of LGI, since it provides higher throughput in situations where the FERs of LGI and SGI are the same.

The FERs of a pair of rate configurations, 1S-I2-LG-20M=19.5 and 1S-I2-SG-20M=21.7, are shown in Figure 8. The only difference between these configurations is the length of their guard intervals. The data was collected using the office, mobile scenario using both the 2.4 and 5 GHz bands (i.e., Scenarios 4 and 5). When using the 5 GHz band, the FER of the SGI configuration is generally higher than the LGI configuration. We also observe this behavior for many pairs of rate configurations that differ only in their use of SGI or LGI (not shown here) in the 5 GHz band office and hallway scenarios (i.e., Scenarios 5 and 7). Interestingly, when using the 2.4 GHz spectrum (i.e., Scenarios 4 and 6), the ratio of the FERs are close to equal.

We believe different results are seen in the 2.4 and 5 GHz bands because many building materials reflect 5 GHz signals orders of magnitude better than 2.4 GHz signals [16]. Thus, the delay spread can be longer for 5 GHz signals, increasing the FER of SGI rate configurations because reflected signals arrive after the SGI and interfere with next symbol.

We now study if the observed difference in FER is significant enough to result in the inferior performance of the SGI configuration. The short guard interval provides a throughput increase of at most 11% when compared with the throughput of the corresponding LGI configuration. If the FER of the SGI configuration is too high, the extra throughput achieved from the SGI can no longer compensate for the higher FER. The blue dashed line (labeled "Threshold") shows this threshold. Points above this line indicate

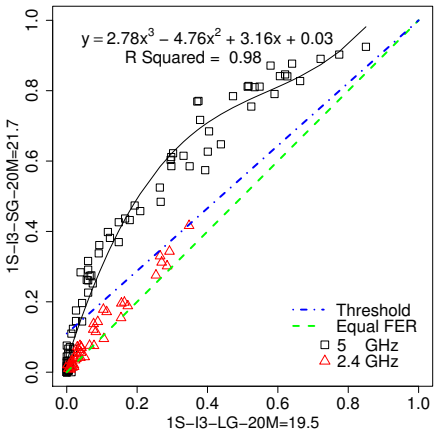

Figure 8: Office: mobile, 2.4 and 5 GHz

that the throughput of the LGI configuration is higher than the SGI configuration. As seen in Figure 8, in the 5 GHz band, many data points are above this line, indicating that the LGI configuration provides higher throughput than the SGI configuration. We observed similar results for other pairs of rate configurations and scenarios. Therefore, in the 5 GHz band, using the LGI could provide higher throughput for some configurations.

5.2 Relationship Between LGI and SGI

We now more closely examine the existence of relationships between LGI and SGI rate configurations. Figure 8 shows a non-linear relationship for the 5 GHz experiment. We found that a quadratic regression model (shown at the top of the figure) fits the data very well (i.e., $R^2 = 0.98$). To verify that LGI and SGI rate configurations are related for other combinations of physical layer transmission rate features and other scenarios, we compute the estimation power from LGI to SGI configurations and from SGI to LGI configurations for all other scenarios and rate configurations (i.e., 48 configurations). Figure 9 shows the estimation power for all configurations for Scenarios 1, 2, 4, and 5 (from top to bottom, respectively). The graphs plot the rate configuration on the x-axis and the estimation power value on the y-axis for both the LGI $\mapsto$ SGI and LGI $\mapsto$ SGI relationships. Note that we only label every 2nd rate configuration on the x-axis. The results show that the relationships are strong between all pairs of configurations that differ only by the guard interval length. Similar results were observed for the other scenarios. In the next section, we utilize the relationships between LGI and SGI configurations to demonstrate how a rate adaptation algorithm might utilize such relationships to optimize transmission rate feature configurations.

6. IMPACT ON APPLICATIONS

Many rate adaptation algorithms (RAAs), including the Minstrel HT algorithm implemented in the widely used Ath9K driver, rely on sampling to determine the best rate configuration. Unfortunately, the overhead of sampling is significant [9]. In order to reduce this overhead, the sampling frequency can be reduced by sampling a smaller subset of rate configurations and utilizing relationships between FERs to estimate the FERs of the remaining rates.

To demonstrate the impact of utilizing relationships in rate adaptation, we modify Minstrel HT to estimate the FER of LGI configurations from the SGI configurations (in-

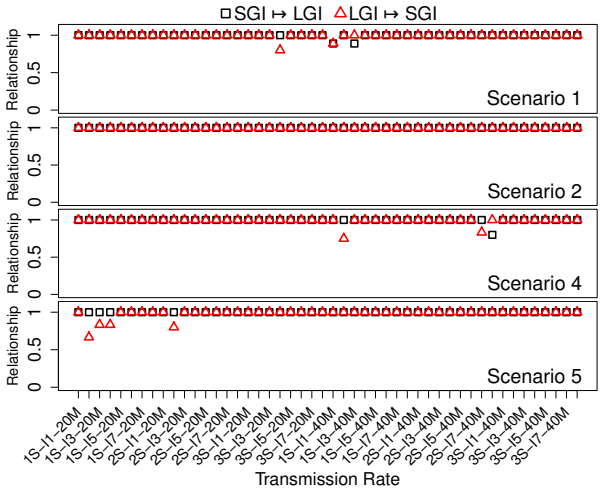

Figure 9: EP of SGI/LGI rate configurations

stead of sampling LGI configurations). In this way, we reduce the frequency at which different configurations are probed. We refer to this modified version as Minstrel HT Relationship. Our estimation function assumes that the FERs of two rate configurations are identical if the configurations only differ in the guard interval. Note that the RAA may still select LGI configurations, it just does not sample them. Although in Section 5.1 we show that in the 5 GHz bands the FER of the SGI configurations are sometimes higher, the performance results indicate that our estimation function is sufficiently accurate for this illustration. We run each experiment multiple times and present the average throughput with 95% confidence intervals in Figure 10. The two pairs of outer bars show that Minstrel HT Relationship increases throughput when compared with the vanilla version by 28% and 17%, in the stationary 2.4 GHz and mobile 5 GHz scenarios, respectively. The middle bars in each grouping show the throughput obtained using Minstrel HT when the probing frequency is reduced (Low Sample Rate) to that used by Minstrel HT Relationship. The gaps between the "Low Sample Rate" bars and the "Relationship" bars demonstrate that throughput is improved as a result of using relationships and does not come only from reducing the number of probed configurations.

These findings demonstrate the potential power of exploiting relationships among rate configurations in algorithms that optimize the selection of physical layer transmission features. Designing and comprehensively evaluating an RAA to fully utilize relationships is outside the scope of this paper and is left for future work.

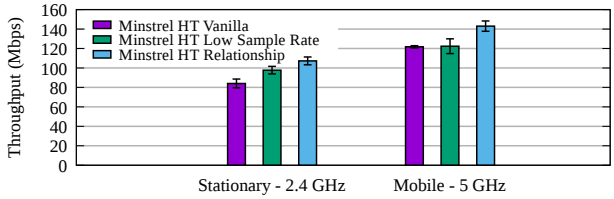

Figure 10: Impact of using relationships in RAAs

7. RELATED WORK

To better understand 802.11n networks, researchers have characterized these networks from different perspectives, such as the relationship between physical layer features and throughput [10], jitter and energy efficiency [17]. We now review some of these 802.11n characterization studies. Although some of these studies are not directly related to our work, they help to understand different viewpoints related to the characterization of 802.11n networks.

LaCurts et al. [11] use over 1,400 access points to empirically study the correlation between sender-side SNR and throughput. The results show that SNR is not sufficient to determine the best transmission rate for a network although with sufficient training, SNR can be a good predictor of the throughput for a specific link. Halperin et al. [6] have conducted a similar study and found that due to the frequency selective fading problem,[2] the SNR is an inaccurate metric of the quality of the communication channel as it does not represent the quality of the signal received by each OFDM sub-carrier. They propose and evaluate a packet delivery prediction model that uses the effective SNR which is derived from the Channel State Information (CSI). The shortcoming of this technique is that in practice many 802.11n chipsets do not report the channel state information [3, 9] and a mechanism is required to transfer CSI data (only available at the receiver) back to the sender.

To our knowledge, the closest work to ours is the study conducted by Kriara et al. [10]. They characterize the positive and negative impacts of each physical transmission feature on throughput, FER and jitter. The authors argue that to maximize throughput all transmission features should be optimized jointly. In other words, if the transmission features are optimized one by one, a sub-optimal throughput might be achieved. This finding emphasizes the need to limit the search space when selecting the best rate configuration, as a divide and conquer approach does not work.

To limit the search space when optimizing the rate configuration, Kriara et al. [9] utilize sender side RSSI information to reduce the number rate configurations that need to be probed. In some scenarios, they are able to improve rate adaptation by reducing the overhead of probing many configurations. The data collection methodology used in [9] requires repeatability across experiments. Therefore, all experiments conducted to measure the relationship between RSSI and transmission features are done in controlled environments with no mobility and no interference other than those introduced by the experimenters. Additionally, their methodology only obtains information about rates if and when they are selected by the RAA. We postulate that this is why their algorithm does not obtain benefits in some of the uncontrolled environments they have evaluated. On the other hand, our data collection methodology examines all rates equally and can be used in controlled and uncontrolled environments without any restrictions.

In contrast with all previous work, we characterize the relationship among rate configurations in order to find how accurately the FER of one rate configuration can estimate the FER of another configuration. We show how these relationships can be used to provide new insights into the behaviour of different physical layer features (e.g., SGI and LGI). Additionally, we demonstrate how relationships might be used to reduce the search space when choosing the best combination of physical layer transmission features.

[2]OFDM sub-carriers suffer from different and independent fading, rendering the average SNR over all sub-carriers a very coarse and inaccurate metric for the quality of the received signal.

8. DISCUSSION

Our results are limited to the scenarios and hardware used for these experiments. However, the number of strong relationships between many configurations across the scenarios examined suggests that interesting relationships are likely to exist under a variety of channel conditions.

Our methodology utilizes several parameters such as: the number of bins (10), the minimum number of data points required for the dispersion metric to be considered reliable (5), the threshold for the interdecile range for a bin to be considered acceptable for accurate estimation (0.2) and the value used to consider the estimation power metric to be good enough to consider the relationship as strong (0.7 for the heat maps). We chose these parameters based on visual inspection of significant amounts of data and by trying to find a set of parameters and a methodology that are fairly intuitive. We believe that the chosen parameters are fairly conservative. However, the best choice of these values will depend on the purpose of the particular characterization study or the application to which it is being applied. Topics for future research include: the choice of parameters, good estimator functions and studying their accuracy.

To choose the most suitable methodology for characterizing relationships we have considered many different methodologies including but not limited to correlation, conditional entropy, and parametric and non-parametric regression. As outlined in Section 2.3 our definition of relationship is, *by necessity* a directional measure. Unfortunately, many techniques such as correlation (e.g., the Pearson product-moment correlation coefficient), and R^2 of regression provide the same value for $R_1 \mapsto R_2$ and $R_2 \mapsto R_1$ relationships which is not true of the relationships we believe are interesting for this study. In addition, correlation only detects simple linear correlations; however, we found non-linear relationships between rate configurations. For these key reasons, correlation and R^2 were not used in this study. While conditional entropy does consider the direction of the relationship, it is not defined for some special but critical cases we observe. For example, it is not defined when there is no variability in the FER of the estimator rate configuration. Therefore, it cannot be used to quantify estimation power.

We chose to focus on the 802.11n standard because it is widely used, it supports both 2.4 and 5 GHz bands and because the open source Ath9k driver made it possible to implement our data collection methodology quickly and easily. Although we believe that many of our findings will apply in 802.11ac networks because they use share many physical layer features with 802.11n networks, studying 802.11ac networks is left for future work.

9. CONCLUSIONS

In this paper, we design a methodology for evaluating the relationships between the FER of different physical layer transmission feature combinations (rate configurations). We find that in all seven scenarios examined, a surprisingly small number of rate configurations can estimate the FER of all other configurations. Interestingly, although we demonstrate that relationships can change over time, relationships are observed in uncontrolled environments over periods of up to one hour. Finally, by utilizing a small fraction of relationships, we provide a simple illustration of how the throughput of the widely used Minstrel HT algorithm can be increased by up to 28% in the uncontrolled environments tested. This demonstrates that there are significant opportunities for utilizing relationships between rate configurations in designing algorithms that must chose the best combination of physical layer features from a large number of possibilities.

10. ACKNOWLEDGMENTS

Funding for this project was provided in part by a Natural Sciences and Engineering Research Council (NSERC) of Canada Discovery Grant and a Discovery Accelerator Supplement. The authors thank Andrew Heard for many fruitful discussions and help with some initial experiments during the early phases of this work as well as Martin Karsten, Kamal Rahimi Malekshan and the anonymous reviewers for their feedback on previous drafts of this paper.

11. REFERENCES

[1] Abedi, A., and Brecht, T. T-RATE: A framework for the trace-driven evaluation of 802.11 rate adaptation algorithms. In *MASCOTS* (2014).

[2] Abedi, A., Heard, A., and Brecht, T. Conducting repeatable experiments and fair comparisons using 802.11n MIMO networks. *SIGOPS Oper. Syst. Rev. 49*, 1 (2015).

[3] Biswas, S., Bicket, J., Wong, E., Musaloiu-E, R., Bhartia, A., and Aguayo, D. Large-scale measurements of wireless network behavior. In *SIGCOMM* (2015).

[4] Deek, L., Garcia-Villegas, E., Belding, E., Lee, S.-J., and Almeroth, K. The impact of channel bonding on 802.11n network management. In *CoNEXT* (2011).

[5] Gast, M. S. *802.11n: A Survival Guide*. O'Reilly, 2012.

[6] Halperin, D., Hu, W., Sheth, A., and Wetherall, D. Predictable 802.11 packet delivery from wireless channel measurements. In *SIGCOMM* (2010).

[7] IPerf. http://sourceforge.net/projects/iperf/.

[8] Karp, R. Reducibility among combinatorial problems. *Complexity of Computer Computations* (1972), 85–103.

[9] Kriara, L., and Marina, M. Samplelite: A hybrid approach to 802.11n link adaptation. *ACM SIGCOMM Computer Communication Review* (2015).

[10] Kriara, L., Marina, M., and Farshad, A. Characterization of 802.11n wireless LAN performance via testbed measurements and statistical analysis. In *SECON* (2013).

[11] LaCurts, K., and Balakrishnan, H. Measurement and analysis of real-world 802.11 mesh networks. In *IMC* (2010).

[12] Liu, R. P., Sutton, G., and Collings, I. A new queueing model for QoS analysis of IEEE 802.11 DCF with finite buffer and load. *Wireless Communications, IEEE Transactions on* (2010).

[13] Lorincz, J., and Begusic, D. Physical layer analysis of emerging IEEE 802.11n WLAN standard. In *8th International Conference Advanced Communication Technology* (2006).

[14] Prasad, R., Dixit, S., van Nee, R., and Ojanpera, T. *Globalization of mobile and wireless communications*. Springer Netherlands, 2011.

[15] Van Nee, R., Jones, V. K., Awater, G., Van Zelst, A., Gardner, J., and Steele, G. The 802.11n MIMO-OFDM standard for wireless LAN and beyond. *Wireless Personal Communications* (2006).

[16] Wilson, R. Propagation losses through common building materials, 2.4 GHz vs 5 GHz. http://www.am1.us/Protected_Papers/E10589_Propagation_Losses_2_and_5GHz.pdf. Accessed: 2016-05-12.

[17] Zeng, Y., Pathak, P., and Mohapatra, P. A first look at 802.11ac in action: Energy efficiency and interference characterization. In *IFIP* (2014).