Evaluating 802.11ac Features in Indoor WLAN: An Empirical Study of Performance and Fairness

Lito Kriara
Disney Research Zurich
Switzerland
lito@disneyresearch.com

Edgar Costa Molero
ETH Zurich
Switzerland
cedgar@student.ethz.ch

Thomas R. Gross
ETH Zurich
Switzerland
thomas.gross@inf.ethz.ch

ABSTRACT

We present a thorough and extensive experimental performance characterization of the achievable data throughput, jitter, and fairness of the IEEE 802.11ac standard for indoor Wireless Local Area Networks (WLANs) using real testbed deployments and statistical analysis. 802.11ac achieves higher throughput by incorporating wider channels, more spatial streams, and denser modulation compared to the 802.11n standard. Through diverse testbed experiments we use multiple linear regression to gain insights on the influence of individual 802.11ac features and of their combinations on network performance and fairness for various link and interference scenarios. We further show that 802.11ac WLANs with wider channels can be fairer compared to 802.11a/n in dense environments with high interference.

1. INTRODUCTION

Traffic demands and the need for more bandwidth in wireless communication continue to increase. Therefore, IEEE 802.11ac [1] was introduced to address this increased need to carry wireless traffic. This standard has the potential to deliver multi-gigabit per second throughput by incorporating wider channels, more spatial streams, and denser modulation than 802.11n [2].

In this paper we present and discuss the design and evaluation of two indoor WLAN testbeds as a mean to collect measurements and characterize IEEE 802.11ac and its new features in terms of usage pattern, traffic and impact on network performance and fairness.

We set up two testbeds in different environments (i.e., office and home scenarios) to validate that the trend of our results applies in more than one specific use case. Using these testbeds we report on an investigation of all available 802.11ac features in current commercial hardware (i.e., channel width, number of streams, length of guard interval, and Modulation and Coding Scheme index) on throughput and jitter performance. We do not evaluate the impact of frame aggregation in this study, because previous studies have validated the positive impact of frame aggregation on WLAN performance [3].. Hence, we quantify the impact of the various features on WLAN performance for characterizing and understanding their behaviour under different interference scenarios and for future reference in optimising link adaptation techniques. Finally, we create two new smaller testbeds of four greedy clients with and without hidden nodes to evaluate the performance of 802.11ac and varying channel widths in terms of fairness and compare it to earlier standards (i.e., 802.11a [4] and 802.11n).

Our study describes in detail the WLAN testbed design and evaluation, as well as how to perform a comprehensive measurement collection for further characterization and statistical analysis of the data. We evaluate across all available key features of 802.11ac using two real testbeds, for different link scenarios (no interference and real–world interference). We present a concrete characterization of the features' individual and combinational influence on throughput and jitter performance using statistical analysis. We show that more concrete insights are drawn when different visualization techniques are used combined to overcome the limitations of each single one. Finally, we conduct a fairness study on 802.11a/n/ac across different channel widths and 802.11ac and wider channels are the fairest for supporting low jitter demanding applications.

2. METHODOLOGY

2.1 Testbed Deployment

We deploy two indoor 802.11ac WLAN testbeds; one testbed is set up in an office (§3.1) and the second one in a home environment (§3.2). Each node (clients and access points (APs)) in our testbeds is a laptop run-

WiNTECH'16, October 03-07, 2016, New York City, NY, USA

DOI: http://dx.doi.org/10.1145/2980159.2980167

ning Ubuntu 14.04 with kernel 3.16, the open source ath10k [5] wireless driver, and is equipped with a 3 × 3 antenna and a 802.11ac Qualcomm Atheros QCA9880 chipset–based mini PCI express card [6]. The transmission power is fixed to the default (i.e., 30 dBm for channels over 149, and 17 dBm for channels 32-44). Throughout the paper we use a testbed instead of simulation approach to gain realistic insights for the performance of an everyday WLAN.

The metrics we use to evaluate the testbeds' performance and fairness are application layer throughput and jitter on the receiver side. We report results in no interference and real–world interference scenarios for varying channel conditions. For the no interference scenario, we use the 5 GHz band and channels 149–161 (i.e., 149 as the primary channel) where no activity was detected by our spectrum analyser. All of our no interference experiments are conducted during night hours to minimize human interference. In the case of real–world interference, we use channels 36–44 in the 5GHz band, where additional uncontrolled APs are present. Real–world interference is totally uncontrolled, but the results are not biased (as we report the averages of multiple runs). In the case of the office scenario and real–world interference, we measured – using a spectrum analyser – another nine APs overlapping and operating at the same channels. We conduct our real–world interference experiments during working hours to account for human interference.

We use different channels for the no and real–world interference because channels over 149 are not available for public commercial use, so there is no interfering traffic. It is only possible to conduct real–world interference experiments using channels 32-44. The network performance is not the same for the two different channel setups, but in this study we overcome this testbed design issue by only comparing the measurements within one channel setup. The relative impact of each feature in the specific interference case is accurately compared.

2.2 Measurement Collection

Each data point in the results reported throughout this paper is obtained from averaging across three to five runs of the same experiment. For the characterization study each experiment lasts sixty seconds (§3) and for the fairness analysis two minutes (§4). During the experimental 802.11n/ac feature impact characterization study we use static feature settings, disabling any rate control algorithm. Therefore, a shorter duration of experiments does not affect our study since self-calibration of the rate control algorithm is not needed. For the fairness tests, the same default rate control (i.e., minstrel_ht) is enabled for all features other than the channel width, which we vary for evaluating its impact on fairness. Therefore the length of the experiment is longer to allow for rate control calibration. We use the Iperf tool [7] for UDP traffic generation from the clients to the AP. We set the packet size to the maximum (i.e., 2304 bytes), which was experimentally found to achieve optimal throughput performance.

In this study, we consider all features currently available in commercial 802.11ac hardware, i.e., channel bonding (CB), spatial streams (SS), guard interval (GI), and modulation and coding scheme indexes (MCS). For channel bonding we explore the options of 20, 40, and 80 MHz channel widths, for spatial streams we vary from one to three streams, and for the MCS index there are ten options according to the 802.11ac standard. Finally, for the guard interval option there is the 800 ns long (LGI) or 400 ns short guard interval (SGI). There are in total 180 combinations for each link and interference type. Given that our office testbed (§3.1) has eight nodes and we consider two interference scenarios, our dataset increases to almost 3000 different configurations. Finally, we perform outlier detection before we process the collected data to validate that our insights are not affected by outliers in the measurements (e.g., in case the client disconnect while a specific experiment is running). If outliers are found, we discard them and repeat the experiment.

2.3 Statistical Analysis

To gain a deeper understanding of the results of the extensive measurement campaign, we use statistical analysis. Specifically, we apply multiple linear regression [8], which is a predominant empirical tool in epidemiology, economics, and other sciences, to generate compelling evidence of causal relationships between parameters and data of controlled experiments. We use it to model the relationship between two or more explanatory variables, and their influence on a response variable by fitting a linear equation to the observed data.

Every value of the independent variable vector x is associated with a value of the dependent variable y. Given a data set $\{y_i, x_{i1}, x_{i2}, \cdots, x_{ip}\}_{i=1}^{n}$ for n statistical samples, a linear regression model assumes that there is a linear relationship between the response variable y_i and the x_i explanatory variable vector. Adding the unobserved error ε_i, the model is formed as follows:

$$y_i = \beta_1 x_{i1} + \cdots + \beta_p x_{ip} + \varepsilon_i, i = 1, \cdots, n$$

where y_i is the dependent variable and x_{ip} are the independent variables. β is a p-dimensional vector containing the regression coefficients. Statistical estimation and inference in linear regression focuses on the β.

In our case, the explanatory variables are the 802.11ac features as described before (i.e., CB, SS, GI, and MCS), and the response variable is the respective metric (i.e., throughput and jitter). Hence, we can use the sign of the β_i to infer whether the impact of the specific feature is positive or negative on the metric performance. Using normalized coefficients across each link allows us to compare the relative impact of each feature in the specific link scenario. However, the drawback of regression analysis is that the β is the result of normalizing with

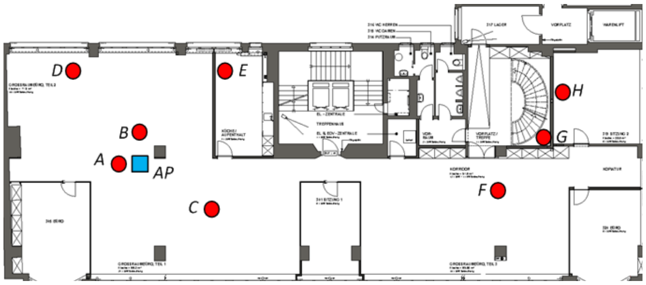

Figure 1: 802.11ac office testbed used for the throughput and latency measurements. The blue square indicates the access point and a red circle a client.

Link	RSSI	Line of Sight	Quality
A	-10 dBm	Yes	Strong
B	-14 dBm	Yes	Strong
C	-27 dBm	No	Strong
D	-40 dBm	No	Medium
E	-45 dBm	No	Medium
F	-57 dBm	No	Medium
G	-61 dBm	No	Weak
H	-75 dBm	No	Weak

Table 1: Average RSSI values for each link type in the office testbed, when using channel 149.

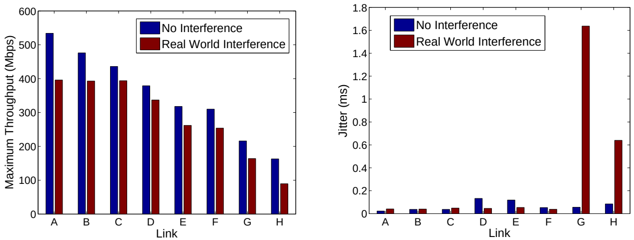

(a) Throughput (b) Jitter

Figure 2: Maximum throughput and corresponding jitter in case of no and real–world interference for all links in the office testbed.

the number of possible values this feature can have (e.g., ten for the MCS index and two for the GI). This means that the impact of MCS might seem lower than the GI, but the meaningful information we can still retrieve is the relative impact of the same feature across different link/interference scenarios. To overcome this issue we use more ways (e.g., tables) to analyze and visualize results throughout the paper to overcome the limitations of each one and gain more spherical insights.

To evaluate if multiple linear regression is applicable to our data, we examine the *p-value* of the F-test on the regression model. If the *p-value* is lower than the *alpha* threshold, then the model accurately describes the data. We set the *alpha* threshold to 0.05, which means that we want our model to describe 95% of our data. Running this test, we find that indeed the *p-value* is always lower than *alpha*, validating that multiple linear regression can properly fit our data, and therefore our multiple linear regression results can be trusted.

3. 802.11AC THROUGHPUT & JITTER PERFORMANCE CHARACTERIZATION

3.1 Office scenario

We deploy an indoor 802.11ac WLAN testbed in our offices, covering an area of $40 \times 15m^2$. The office testbed is depicted in Fig. 1 with the blue square indicating the AP and the red circles the clients . The average RSSI and other characteristics of each client evaluated in the office testbed are described in Table 1.

We follow the methodology described in §2 for the no and real–world interference scenarios. Fig. 2 depicts the maximum throughput (Fig. 2(a)) observed for each link in the office testbed (Fig. 1) across all possible feature combinations and interference scenarios (i.e., no and real–world interference). We also report the corresponding jitter observed for the maximum throughput yielding setting (Fig. 2(b)). Observing such results, we cannot gain much useful information on why the observed changes happen. Therefore, we apply multiple linear regression on our measurements. Results in Fig. 3 show the relative impact of single features and their combinations on throughput and jitter performance for both interference scenarios across different link qualities (i.e., A–H). We use the β computed by the multiple linear regression as described in §2.3 to generate the heatmaps of Fig. 3. Warmer colors (i.e., redder) mean that this feature combination has a higher positive impact on the performance; colder colors (i.e., bluer) indicate negative impact. The maximum (i.e., dark red) and minimum (i.e., dark blue) impact of a feature is ±0.5 as Fig. 3 shows. This means that for a given link type, the specific feature combination provides a gain or loss (up to 50%) in throughput/jitter performance. For example, red CB impact means that larger channel width increases throughput performance. Note that for the case of GI, redder GI indicates that enabling SGI increases throughput performance for the specific link and interference type combination. Green indicates that the specific feature combination does not have any impact (i.e., 0%) on the performance of the metric.

Fig. 3 shows that jitter performance is more susceptible to changes in parameter settings than throughput (darker and non-green colors noticed). Moreover, even though the guard interval is not an important factor for jitter performance, it is the feature with the strongest impact on throughput performance per feature setting. This is because GI has only two possible settings (i.e., SGI/LGI) compared to three, three and ten for CB, SS and MCS, respectively (as explained in §2.3). Tables 2

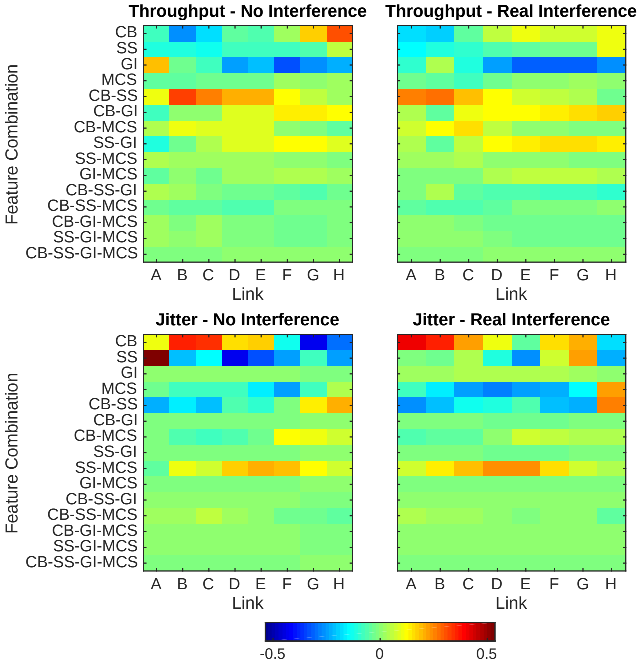

Figure 3: 802.11ac feature combination impact in case of no and real–world interference on throughput and jitter in an office scenario. The colormap shows the impact of the specific feature combination indicated by each color, normalized over each link scenario.

Link	CB	SS	GI	MCS
A	458 Mbps	295 Mbps	100 Mbps	479 Mbps
B	393 Mbps	432 Mbps	305 Mbps	475 Mbps
C	331 Mbps	436 Mbps	306 Mbps	402 Mbps
D	311 Mbps	379 Mbps	247 Mbps	361 Mbps
E	212 Mbps	310 Mbps	67 Mbps	306 Mbps
F	248 Mbps	216 Mbps	52 Mbps	318 Mbps
G	167 Mbps	211 Mbps	46 Mbps	215 Mbps
H	124 Mbps	162 Mbps	61 Mbps	162 Mbps

Table 2: Maximum throughput difference when one specific feature varies and all others remain constant in the case of no interference.

Link	CB	SS	GI	MCS
A	300 Mbps	304 Mbps	66 Mbps	395 Mbps
B	346 Mbps	388 Mbps	189 Mbps	392 Mbps
C	359 Mbps	394 Mbps	172 Mbps	393 Mbps
D	289 Mbps	317 Mbps	152 Mbps	289 Mbps
E	194 Mbps	262 Mbps	126 Mbps	261 Mbps
F	188 Mbps	254 Mbps	89 Mbps	249 Mbps
G	103 Mbps	114 Mbps	60 Mbps	103 Mbps
H	68 Mbps	89 Mbps	45 Mbps	88 Mbps

Table 3: Maximum throughput difference when one specific feature varies and all others remain constant in the case of real–world interference.

and 3 show the maximum gain in throughput of the optimal setting over the worst one when varying a single feature (all other settings remain constant). We see that indeed GI is not the feature with the maximum impact on throughput performance, but the feature with the maximum individual settings impact. The GI results indicate that enabling the SGI creates more losses as the link quality decreases, resulting in lower throughput. On the other hand, in an ideal link quality scenario with no interference, SGI increases throughput performance. GI is the time interval between symbols being transmitted and is used to ensure that distinct transmissions do not interfere with one another. Therefore, a shorter interval between the transmitted symbols may create interference and result in losses when the channel quality is not close to ideal.

Next, we see that wider channels have a high impact on jitter (Fig. 3). Better link qualities (i.e., A–E) have an increased jitter with higher channel width than the poorer links (i.e., F–H), especially in the real–world interference scenario. The reason is that wider channels increase the chances of better links dominating the medium constantly or opportunistically (Table 7). Therefore, the latency experiences high variation and consequently the jitter is increased.

As far as the SS is concerned, we see that on average more streams minimally deteriorate throughput performance because the total transmission power is divided between the multiple streams. This decreases the range and strength of the signal, and consequently incurs more losses. We also notice that more spatial streams reduce the jitter (i.e., bluer values on the SS row for jitter results). When trying to transmit data to the full capacity of the link, then more streams can faster and more efficiently transmit the data, using redundancy compared to one and space time block coding.

Moreover, we notice that the MCS index is not a dominant factor on the throughput performance, suggesting that the correct setting of the previously mentioned features is more crucial for the same MCS setting. However, this is due to the higher number of setting MCS can have (i.e., ten) compared to the other feature, as mentioned in §2.3. Tables 2 and 3 show that MCS index has a consistently high impact on throughput performance (spread across ten MCS index options), whereas the channel width with slightly lower impact but only three options (i.e., 20, 40, and 80 MHz) is shown to have more significant *impact per setting option* in Fig. 3. Still, higher MCS indexes reduce the jitter in both interference cases, because lower rates cause increased PER resulting in higher jitter too.

Tables 2 and 3 do not reflect the positive or negative impact of increasing the value of the specific feature. The tables show the maximum gain in throughput of the optimal setting over the worst one when varying a single feature (all other settings remain constant). We notice that the impact of interference is so high that

Link	RSSI	Line of Sight	Quality
I	-20 dBm	Yes	Strong
J	-41 dBm	No	Medium
K	-58 dBm	No	Medium
L	-70 dBm	No	Weak

Table 4: Average RSSI values for each link type in the home testbed.

even the maximum possible throughput gain reduces dramatically. Moreover, we see that channel width has a high impact in the no interference case, similar almost to the MCS index, whereas in the case of real–world interference, the number of spatial streams becomes the second most important feature following the MCS index. This result shows that in real–world interference (with co-channel and adjacent channel interferers), channel width is more heavily affected and provides less throughput improvement than in the absence of interference. The reason is the reduced transmission power for each of the 20 MHz channels when combined. On the other hand, the number of spatial streams maintains a similar or even higher throughput improvement with the optimal setting in the case of interference compared to the no interference scenario. Again, the reason lies in the redundancy added to multiple streams. Unlike channel bonding, spatial stream performance is also not directly affected by channel leakage and other interference parameters. Moreover, errors that may occur, because of high demand for the medium in a real–world interference scenario, can be handled by FEC and CCA mechanisms, which make the number of spatial streams more resilient to interference.

We also notice that interference (Tables 3) decreases the maximum impact of a feature on throughput performance by almost half for GI no matter the link quality compared to the no interference scenario (Tables 2). On the other hand, CB, SS and MCS experience a similar decrease in the maximum impact only for poor links (i.e., G–H). Interference has a lower impact (up to 25%) of CB, SS and MCS on good quality links (i.e., A–F).

Fig. 3 also shows that not only single parameter estimation is significant for achieving better performance but also combinations of multiple parameters. It is surprising that even though increasing *only* the channel width or *only* the spatial stream number has a negative impact on throughput, increasing the channel width *and* spatial stream number jointly has a highly positive influence on throughput performance and at the same time decreases jitter for both interference scenarios. This observation also supports the argument of jointly adapting all available features would result in higher throughput performance gain [9].

Zeng et al., mention that in a heterogeneous channel width environment, where legacy links operate at different channel widths and secondary channels, competition to access the medium becomes increasingly unfair and results in starvation of the larger channel width links [10]. The authors of that work focus only on legacy interference. However, our study shows that in scenarios with real–world interference (i.e., when both legacy and 802.11ac links are active), wider channels are not the feature that causes significant performance degradation for 802.11ac links under test.

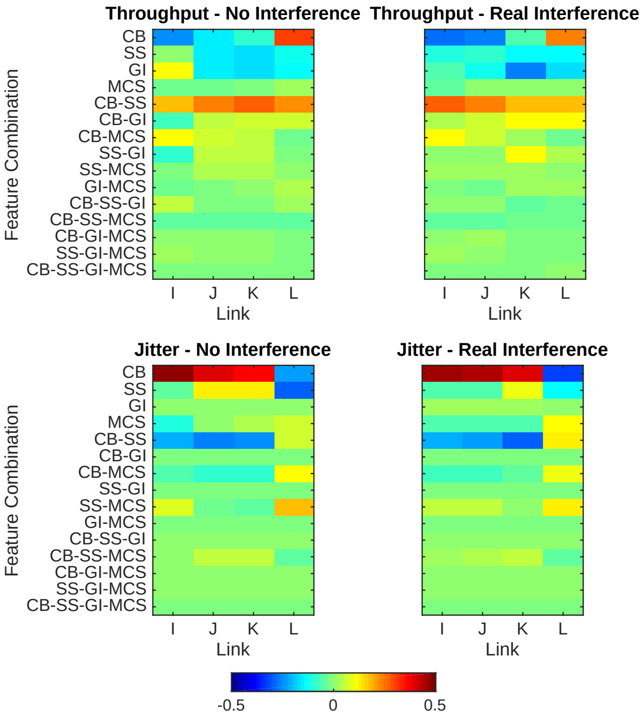

Figure 4: 802.11ac feature combination impact in case of no and real–world interference on throughput and jitter in a home scenario. The colormap shows the impact of the specific feature combination indicated by each color, normalized over each link scenario.

3.2 Home scenario

To validate that the results presented in §3.1 are consistent in more than one testbed, we repeat the same characterization with the same methodology in a home testbed covering an area of $18 \times 15m^2$ (mainly built of wood). The links evaluated are described in Table 4. Note that we examined multiple different areas of the home environment but there was not high variation in link qualities and therefore we only present four links.

We see that the trend of the office regression (Fig. 3) heatmaps is similar is followed also in the case of the home testbed (Fig. 4). However, the trend for each feature in the home testbed is much clearer and more monotonic and consistent across the different link qualities because of the lack of high human interference, as well as the material (wood) the home testbed is made of – compared to the office one (concrete, steel) – min-

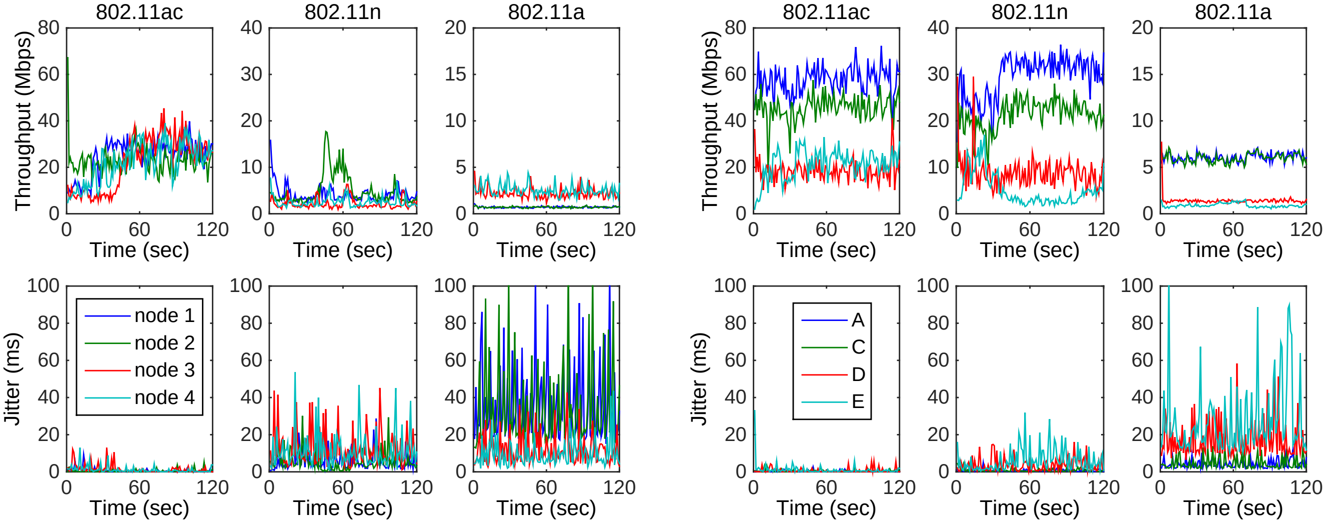

(a) No hidden nodes setup (b) Office Testbed setup

Figure 5: Bandwidth allocation and jitter performance, when four clients share the medium trying to transmit data to the AP. The 802.11ac/n/a standards are compared using the maximum possible channel width in each case (i.e., 80, 40 and 20 MHz, respectively).

imizing the impact of reflections and multipath. Moreover, we observe that even though the same trend in feature impact is followed by both no and real–world interference setups (Fig. 4), there is more diversity in the feature impact on throughput when the testbed is influenced by real–world interference compared to when there is no interference. In the case of real–world interference there is not only more diversity across the different link qualities, but also the impact is higher on throughput performance (i.e., darker colors) compared to the no interference case. So, not only the testbed space (home or office) is a crucial factor for designing a testbed, but also the testbed channel conditions.

4. FAIRNESS

For the fairness evaluation, we explore two different testbed setup scenarios using four nodes simultaneously transmitting UDP traffic to the AP for two minutes. In the first scenario, all competing nodes are within range of each other (no hidden terminals) with line of sight to the AP and an average RSSI of -21 dBm each. In the second scenario, we use four of the nodes in the office testbed shown in Fig. 1 (nodes A, C, D, and E) to span a diverse set of link qualities with hidden nodes. Fairness experimentation results in this section are trustworthy because we only adapt *one* feature setting at a time and all stations try to transmit simultaneously at their maximum possible rate like in a real world scenario (ignoring the overall network performance).

We compare throughput (bandwidth allocation) and jitter metrics per node per second of 802.11ac against 802.11n and 802.11a, to identify the fairest settings and the impact of the 802.11 standard evolution on fairness.

Fig. 5 presents the application layer throughput and jitter per second for the different setups and 802.11 standards. Note that the maximum possible channel width was used for each standard (i.e., 80, 40, and 20 MHz for 802.11ac, 802.11n, and 802.11a, respectively) along with enabled rate control for results shown in Fig 5. Even though the bandwidth allocation is not shared equally in most cases, jitter is minimized with 802.11ac, both for the no hidden nodes and the office testbed with diverse link qualities. The throughput fluctuation noticed between the four stations in the first few seconds (for both testbed setups) is due to the self-calibration of the rate control mechanism.

As expected, diverse link qualities with hidden nodes result in unbalanced bandwidth allocation (Fig. 5(b)), whereas the variance in achieved throughput is much smaller in the setup with no hidden nodes (Fig. 5(a)) because of the Clear Channel Assessment (CCA). CCA enables the stations to determine if the medium is occupied to prevent the station from attempting a transmission when they are within each other's reach, avoiding collision and interference errors.

We notice that even in an ideal scenario (Fig. 5(a)) the 802.11n and 802.11a standards result in higher variance compared to 802.11ac in both throughput and jitter performance between the links. We explain this result with the observation that a wider channel can support higher amounts of data transmitted simultaneously. Therefore, jitter decreases as we move from 802.11a (using 20 MHz) to the newer standards that support channels wider than 20 MHz. To validate the

Standard & Channel Width Combination	No Hidden Nodes Setup	Office Setup
802.11ac & 80 MHz	0.1160	0.1032
802.11ac & 40 MHz	0.5029	0.1994
802.11ac & 20 MHz	1.3244	0.0744
802.11n & 40 MHz	0.3545	0.6349
802.11n & 20 MHz	2.0044	0.1073
802.11a & 20 MHz	0.8940	0.1535

Table 5: Impact of the various standards and channel widths on the unfairness metric for throughput performance.

Standard & Channel Width Combination	No Hidden Nodes Setup	Office Setup
802.11ac & 80 MHz	0.1227	0.4696
802.11ac & 40 MHz	0.4234	0.6566
802.11ac & 20 MHz	0.2022	1.1487
802.11n & 40 MHz	0.3796	0.3479
802.11n & 20 MHz	0.1913	0.6655
802.11a & 20 MHz	0.1750	0.4435

Table 6: Impact of various standards and channel widths on the unfairness metric for jitter.

	A	**C**	**D**	**E**
80MHz	40.4%	31.9%	14.8%	12.9%
40MHz	42.3%	32%	19.9%	5.8%
20MHz	33.8%	29.7%	29.1%	7.4%

Table 7: Percentage of bandwidth allocation per link varying the channel width out to the total network throughput in the case of the office testbed setup and no hidden nodes setup.

hypothesis that wider channels improve fairness across the different standards independent of the node setup in the testbed, we perform the following experiment.

We use the same testbed setups (i.e., no hidden nodes and office testbed setup) with four nodes transmitting simultaneously towards the AP in channel 149 with no external interference. The settings we adapt are the channel width and the 802.11standard. We examine 802.11a/n/ac; each for all supported channel widths (i.e., 20, 40 and 80 MHz).

Existing fairness metrics, like Jain's fairness index [11] and the fairly shared spectrum efficiency (FSSE) metric [12], cannot reflect the proportional fairness for both bandwidth and jitter measurements, because they both focus on the bandwidth allocation fairness. Therefore, we define our own metric for computing the unfairness of the network to reflect proportional unfairness. This metric normalizes the standard deviation of the average throughput/jitter performance of each node in the network by dividing it by the average throughput/jitter across all nodes:

$$unfairness_{standard,cb} = \frac{std(Data_{standard,cb})}{mean(Data_{standard,cb})}$$

where *standard* and *cb* represent the 802.11 standard and channel width, respectively, and $Data_{standard,cb}$ is the 2–dimensional matrix of measurements for each experiment; the different columns are the various nodes, and the rows are the per second throughput/jitter measurements. The $mean(Data_{standard,cb})$ gives a vector of the mean throughput/jitter per node for the whole duration of the experiment. The lower the $unfairness$ metric, the fairer the system, since it reflects no high variation between the throughput/jitter performance of the various nodes (i.e., low standard deviation).

The impact of different standards and channel widths on fairness in terms of throughput and jitter performance is depicted in Tables 5 and 6, respectively. Lower unfairness values in the tables indicate better fairness performance of the 802.11 standard and channel width combination. Our hypothesis that wider channels result in a fairer system is validated, both in terms of throughput and jitter performance.

802.11a is fairer than 802.11n/ac in terms of throughput (Table 5) when using 20 MHz channels, as reported in [3]. Hence, we expected that 802.11ac would follow this trend and also perform less fair than 802.11n. However, we see that both 802.11n and 802.11ac become even fairer as the channel width increases independent of the setup of the nodes and the existence of hidden nodes. The same pattern is observed also in the case of jitter performance (Table 6). Only in the case of the office testbed setup we notice that diverse links can experience higher unfairness in throughput for wider channels (Table 5). This is expected since higher quality links may dominate the network (Table 7). We notice that the best two links (i.e., A-C) dominate the 63.5% and 72.3% of the total network bandwidth for channel width of 20 and 80 MHz, respectively, while at the same time links D and E occupy 36.5% and 27.7%.

Overall, we see that the combination resulting in the fairest performance when multiple nodes are sharing the medium is 802.11ac using 80 MHz channel width. This suggests that 802.11ac not only delivers high throughput, but is also fairer than the earlier standards due to channel bonding and its higher capacity.

5. RELATED WORK

Yu et al. combat the inter-cell interference of 802.11ac-based Multi-User MIMO (MU-MIMO) networks with multiple APs. This only allows a single access point to serve its clients at any given time, significantly limiting the network capacity [13]. The authors exploit the AP and clients' antennas for beamforming. On the same topic, Xiong et al., proposed MIDAS, a Multiple-Input Distributed Antenna System, to improve 802.11ac performance by challenging the access point antenna topology [14]. Both works were implemented and evaluated on a WARP platform instead of a real–world testbed with off-the-shelf hardware.

Zeng et al. conduct a thorough outdoor experimental characterization of 802.11ac on commodity hardware focusing on power consumption and channel width [10]. The authors also address the unfairness in the case of 802.11ac devices coexist *only* with legacy, and conclude that heterogeneous channel width environments (where legacy links operate at different channel widths and secondary channels) increase unfair competition to access the medium, which results to starvation of the larger channel width links. Similarly, Lee et al., focus on power efficiency and channel bonding, proposing an adaptation scheme that optimizes power consumption in 802.11ac by adapting the channel width [15].

Finally, Kriara et al. performed an indoor 802.11n WLAN performance characterization [3] and extended their work using insights from the characterization to perform efficient 802.11n link adaptation by minimizing the sampling overhead [9].

6. CONCLUSION

In this paper we discuss the design and evaluation of two indoor wireless testbeds used to perform a repeatable extensive experimental characterization of the jitter, throughput and fairness performance of 802.11ac WLANs under two different interference scenarios. We report our insights of how to overcome limitations of testbed deployments and visualization techniques. Unlike prior work, we include all commercially available 802.11ac features in our characterization study. We show how features have different impact on performance and that they are inter-related. Moreover, we report on 802.11ac performance and address the fairness issue in the context of 802.11ac in real–world interference conditions for the first time and show how wider channels improve network fairness.

Our characterization takes a thorough approach towards identifying the key features of 802.11ac that enable improving network performance under various channel conditions. We show that different features and feature combinations have a different impact on performance depending on the channel conditions. Moreover, we find that combinations of features can have the opposite impact on network performance when considered jointly, than when each feature is considered alone, adding value to holistic link adaptation. Finally, we see that a single visualization technique (e.g., regression) is not sufficient to provide deep and meaningful insights, but we gain a more spherical understanding of the measurements by combining different techniques.

In our fairness study, we compare 802.11ac to prior standards with varying channel widths when multiple clients are simultaneously active in the network. We explore the impact of hidden nodes in the testbed. However, no matter if there are hidden nodes in the network testbed or not, our results indicate that 802.11ac and wider channels improve both throughput and jitter performance. Therefore, 802.11ac and wider channels is the setting combination maximizing throughput and jitter fairness; efficient not only for video and other throughput-hungry applications, but also for applications demanding consistent packet delivery times.

7. REFERENCES

[1] IEEE 802.11ac. Enhancements for Very High Throughput for operation in bands below 6GHz. In *IEEE P802.11ac/D5.0*, 2013.

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

[3] L. Kriara et al. Characterization of 802.11n Wireless LAN Performance via Testbed Measurements and Statistical Analysis. In *IEEE SECON*, 2013.

[4] IEEE Std 802.11a. *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High-speed Physical Layer in the 5 GHz Band, Std., Sep.*, 1999.

[5] Qualcomm. Ath10k Linux driver.

[6] Compex. WLE900V5-18 NIC.

[7] Iperf version 2.0.0, http://iperf.sf.net.

[8] D. Freedman. *Statistical Models: Theory and Practice*. Cambridge University Press, 2005.

[9] L. Kriara and M. K. Marina. SampleLite: A Hybrid Approach to 802.11n Link Adaptation. *ACM SIGCOMM CCR*, April 2015.

[10] Y. Zeng et al. A First Look at 802.11ac in Action: Energy Efficiency and Interference Characterization. In *IEEE IFIP Networking Conference*, 2014.

[11] R. K. Jain et al. *A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer System.* DEC Technical Report 301, 1984.

[12] M. Eriksson. Dynamic Single Frequency Networks. *IEEE Journal on Selected Areas in Communications*, 2001.

[13] H. Yu et al. Combating Inter-cell Interference in 802.11ac-based Multi-User MIMO Networks. In *ACM MobiCom*, 2014.

[14] J. Xiong et al. MIDAS: Empowering 802.11ac Networks with Multiple-Input Distributed Antenna Systems. In *ACM CoNEXT*, 2014.

[15] O. Lee et al. WiZizz: Energy Efficient Bandwidth Management in IEEE 802.11ac Wireless Networks. In *IEEE SECON*, 2015.