[SRC] Poster: Regression-based Characterization of 802.11ac Indoor Performance

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

Lito Kriara
Disney Research Zurich
lito@disneyresearch.com

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

ABSTRACT

We investigate the applicability of IEEE 802.11ac for entertainment low-latency show control systems (to orchestrate free-riding vehicles in a theme-park environment) by experimentally characterizing the indoor throughput and jitter performance of 802.11ac using statistical analysis. 802.11ac is the successor of 802.11n that provides higher throughput by incorporating wider channels, more spatial streams and denser modulation. We show that multiple linear regression provides valuable insight in the influence of 802.11ac's independent features and their combinations on performance, for various links and interference scenarios. Finally, we show that 802.11ac could be used for not only delivering high throughput for multimedia streaming but also supports applications requiring minimal jitter variance in the setup investigated.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design–Wireless communication

General Terms

Experimentation, Measurements, Performance, WLAN

Keywords

802.11ac, Analysis, Characterization, MIMO

1. INTRODUCTION

A recent survey reported that *"by 2016 the amount of mobile devices' traffic, relying on Wireless Local Area Networks (WLANs) for Internet access, will exceed the traffic from wired devices"* [1]. IEEE 802.11ac [2] was introduced to address this increased need in wireless traffic. It has the potential to deliver multi-gigabit per second throughput, by incorporating wider channels, more spatial streams, and a denser modulation than 802.11n.

MobiCom'15, September 07-11, 2015, Paris, France
ACM 978-1-4503-3619-2/15/09.
http://dx.doi.org/10.1145/2789168.2795173.

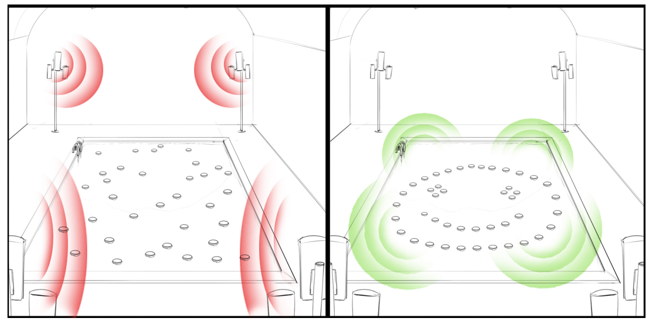

Figure 1: Concept art (© Disney): Future indoor 802.11ac–based show control system with free-riding vehicles.

We report on an investigation of the available 802.11ac features on throughput and jitter performance across a wide range of link qualities, via extensive real–world testbed experimentation. We do that so to determine if 802.11ac is a standard that could possibly support low latency show control systems in the future (Fig. 1). As a first step in this direction, we perform a thorough real–world evaluation and characterization study of the 802.11ac performance in an indoor environment using already available hardware in an 802.11ac office testbed.

Zeng et al. conduct an outdoor characterization of 802.11ac focusing on power consumption and the impact of channel width [3]. In this paper, we present a more thorough evaluation of the standard by considering all available features and their combinational influence on performance.

2. METHODOLOGY

2.1 Testbed setup

We deploy an indoor 802.11ac WLAN testbed, covering a $40 \times 15m^2$ area, as shown in Fig. 2; the blue square indicates the access point, and the red circles represent the clients of the WLAN. Each node is a laptop running Ubuntu with kernel 3.16, the open source ath10k wireless driver [4], and is equipped with a 3×3 Qualcomm Atheros QCA9880 chipset–based mini PCIe. The transmission power is always fixed to the default (i.e., 30 dBm for channels 149-161, and 17 dBm for channels 32-44).

The metrics we use to evaluate the testbed's performance are application level throughput and jitter on the receiver side, measured between a client and access point using the

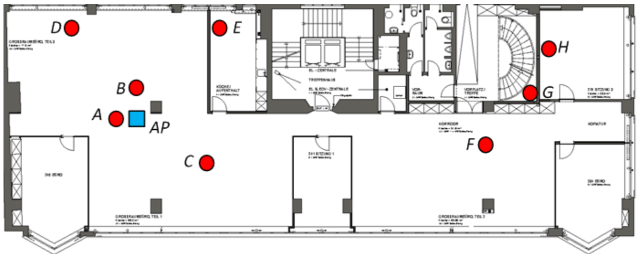

Figure 2: 802.11ac testbed for the throughput and jitter measurements. Blue square: access point, red circle: clients.

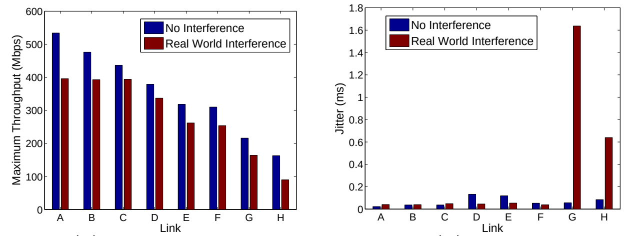

(a) Throughput (b) Jitter

Figure 3: Maximum throughput and corresponding jitter in case of no and real world interference for all links in the testbed.

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.

Iperf traffic generator tool [5]. Note that MAC layer packet loss related statistics were not available using the ath10k driver at the time the experiments were conducted. We demonstrate results in "no interference" and "real-world interference" scenarios, with varying channel conditions. For the "no interference" scenario we use during night hours the 5 GHz band and channel 149, where no activity was detected by a spectrum analyzer. For "real-world interference" we use channel 36 in the 5GHz band, where another 9 access points are operating during working hours to account for human interference, too. Table 1 shows the different link characteristics.

We consider all available 802.11ac features in our characterization study, namely channel bonding (CB), spatial streams (SS), guard interval (GI), and the modulation and coding schemes (MCS) index. For channel bonding, we explore the options of 20, 40 and 80 MHz channel width; for spatial streams we vary from one to maximum three streams available in our hardware, and for the MCS index, we explore the ten options as described in the 802.11ac standard. Finally, for the guard interval option there is the long (LGI) or short guard interval (SGI). To sum up, there are 180 combinations for each link/interference type. Given that our testbed has eight nodes, and we consider two interference scenarios, our possible setting number increases to almost 3000 different configurations. Therefore, we explore the use of statistical techniques on this dataset to gain meaningful insights.

2.2 Multiple Linear Regression

Multiple linear regression [6] is an empirical tool in epidemiology, economics and other sciences and is used to search for evidence of causal relationships between parameters and observational data of controlled experiments. We use it to model the relationship between two or more explanatory variables, as well as their influence on a response variable, by fitting a linear equation.

The explanatory variables are 802.11ac features (i.e., CB, SS, GI and MCS), and the response variable is our respective metric (i.e., throughput and jitter). Hence, we can use the coefficients of the linear equation to infer, for example, whether the impact of wider channels is positive or negative on the throughput performance by the sign of the CB coefficient. For normalized coefficients across different links, we can also confirm that a specific feature has a higher impact than another in certain scenarios.

To evaluate whether multiple linear regression is applicable to our data, we examine the significance (*p-value*) of the regression model's F-test. If the *p-value* is lower than an *alpha* threshold, then the model can accurately describe the data. We set the *alpha* threshold to 0.05 (i.e., the model should describe 95% of the data), and 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. EXPERIMENTAL EVALUATION

We follow the methodology described in the previous section for the "no interference" and "real world interference" scenarios. Fig. 3 depicts the maximum throughput (Fig. 3(a)) observed for each link in the testbed (Fig. 2), under both setups of interference, as well as the corresponding jitter observed (Fig. 3(b)). With this kind of visualization we cannot gain much insight on why does an observed change (e.g., between different interference scenarios) happen. Why does jitter increase rapidly for weak links? Are there setting combinations to minimize jitter, but maximize throughput performance?

To answer such questions, we use multiple linear regression on our measurements. Results in Fig. 4 show the impact of single features and their combinations on throughput and jitter performance, under "no" and "real world" interference. Warmer colors (i.e., redder) mean that this feature combination has a higher positive impact on the performance when its value is increased; colder colors indicate negative impact. E.g., red CB means that larger channel width increases throughput performance. On the other hand, colder colors (i.e., bluer) would indicate negative impact and a decrease of throughput for larger channel width. Green shows that the specific feature combination does not have a significant impact on the performance of the metric. Note that for the case of GI, a redder color GI indicates that enabling SGI increases throughput performance.

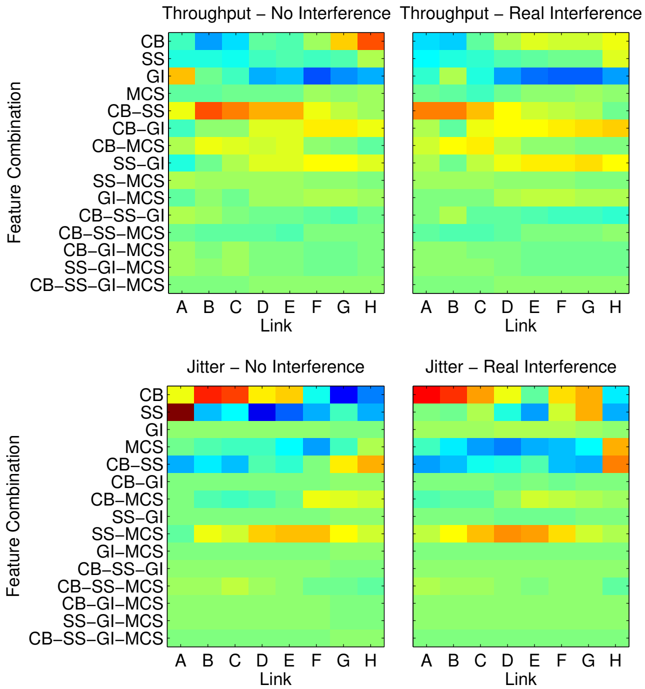

Figure 4: 802.11ac feature combination impact in case of "no" and "real world" interference on throughput and jitter.

Fig 4 shows that jitter performance is more susceptible to changes in parameter settings than throughput (darker 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. As expected, enabling the SGI creates even more losses as the link quality decreases, resulting in lower throughput, whereas in an ideal link quality scenario with no interference, GI proves to be increasing throughput performance.

Next, we notice that wider channels result in decreasing throughput performance, because the signal transmission power is divided across a larger number of subcarriers. This setup degrades the link quality of the specific scenario, when the remaining of the settings are set to high values, similar to the 802.11n case [7]. Wider channels have an even higher impact on jitter variance. Better link qualities (i.e., A–E) have an increased jitter with higher channel width, compared to the poorer links (i.e., F–H), because better links have fewer losses, so more delayed packets. We see this changing for the real world interference case, where the losses are so high that most packets do not arrive at all, so delay cannot be measured, explaining the high jitter of link G (Fig. 3).

As far as the SS is concerned, we see that on average more streams deteriorate throughput performance, because the transmission power is divided between the multiple streams, decreasing the range and strength of the signal, and consequently incurring losses. Surprisingly, we see that at the same time, more spatial streams reduce the jitter, because when trying to transmit data to the full capacity of the link, more streams can more efficiently transmit all the data.

Moreover, the MCS index does not look like a dominant factor on throughput performance, suggesting that the correct setting of the previously mentioned features is more crucial for the same MCS setting. However, we see that higher MCS indexes reduce the jitter, in both interference cases. Jitter increases with higher MCS indexes in the case of poor links, similarly to the channel bonding.

We also see that not only single parameter estimation is significant for achieving better performance but also combinations of multiple parameters. Fig. 4 shows the combinations that are more significant and their impact. Is it 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 impact on throughput performance and at the same time decreases jitter, for both interference scenarios. Similarly, we see that jointly increasing the number of SS and MCS index increases the jitter, even though this combination has no significant impact on throughput performance. This is due to the reduced transmission power, when more streams are used. Hence, the link quality is reduced even more, introducing higher jitter.

Finally, we notice that on average the two-way interactions (e.g., CB–SS, CB–GI) increase throughput performance when the settings are higher, whereas increasing the setting of single features has the opposite impact, surprisingly. This observation also supports the argument of jointly adapting all available features would result in higher throughput performance gain, as also shown in [8].

4. CONCLUSION

We performed an extensive experimental characterization of the throughput and jitter variance performance of 802.11ac WLANs. Unlike previous approaches we consider all 802.11ac features in this study. We show that multiple linear regression analysis can provide valuable insight in analyzing network performance measurements, and we validate that jointly adapting feature setting may increase performance. Finally, the results on jitter obtained in this testbed indicate that 802.11ac can be employed not only for delivering high throughput but may also be useful for applications that demand low jitter variance.

5. REFERENCES

[1] Cisco. *Visual Networking Index (VNI), The Zettabyte Era - Trends and Analysis*, 2013.

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

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

[4] Qualcomm. Ath10k Linux driver.

[5] M. Gates et al. *Iperf version 2.0.0*, 2004.

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

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

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