Measurement Study of IEEE 802.11ac Wi-Fi Performance in High Density Indoor Deployments: Are Wider Channels Always Better?

Ljiljana Simić, Janne Riihijärvi and Petri Mähönen
Institute for Networked Systems, RWTH Aachen University
Kackertstrasse 9, D-52072 Aachen, Germany
E-mail: lsi@inets.rwth-aachen.de

***Abstract*—Wi-Fi is the dominant wireless indoor broadband solution and thus key for meeting the exponential traffic growth. The recent IEEE 802.11ac amendment enables PHY data rates exceeding 1 Gbps. However, it is not clear how this increased *per-link* performance, achieved especially via wider channels, translates to *network-level* performance. The latter is crucial for understanding the true potential of emerging Wi-Fi, as massive densification of network infrastructure is needed for keeping up with capacity demands. In this paper we present results from an extensive measurement study of the performance of IEEE 802.11ac Wi-Fi in a large 24-node indoor testbed. We investigate in detail the impact of channel width, network deployment density, and type/volume of traffic on the achieved network performance in dense indoor deployments. Our results show that using wide 80 MHz channels is beneficial only in dense networks with extremely high traffic loads, owing to strong adjacent channel interference (ACI) effects with narrower channels. We show that ACI not only reduces aggregate network throughput but causes severe unfairness among nodes, where up to half the nodes may experience starvation. Starvation occurs due to frequency flow-in-the-middle effects and the heterogeneous interference coupling among pairs of nodes typical of indoor deployments. However, our results also demonstrate that for bursty TCP traffic with loads typical of modern residential Wi-Fi deployments, there is no harm from using narrower 20 or 40 MHz channels. Finally, we study the scaling of IEEE 802.11ac deployments with increasing traffic load, which is highly relevant for the dimensioning of future networks. We show that ACI effects only become evident at very high loads which are beyond the per-node traffic demand expected in the foreseeable future. Therefore, *in practice* there is no network-level benefit from employing the wider channels enabled by IEEE 802.11ac, even in emerging highly dense indoor deployments.**

I. Introduction

Wireless technologies are evolving more rapidly than ever in order to keep up with the exponentially increasing requirements for capacity. Wi-Fi hotspots play a key role in this development, carrying ever increasing amounts of Internet traffic in especially residential and commercial indoor deployments. In an effort to improve the performance of Wi-Fi networks, standard amendments such as IEEE 802.11n and the recent IEEE 802.11ac are introducing both physical and MAC layer enhancements (e.g. support for modulation schemes up to 256-QAM, advanced frame aggregation features, and multiuser MIMO beamforming), and also enabling use of wider channel bandwidths [1], [2]. For an individual Wi-Fi link these enhancements promise throughput increases currently in excess of 1 Gbps in ideal conditions. However, especially the wide channel bandwidths also imply increased interference, thereby potentially reducing the *network-level* capacity. As massive densification of the network infrastructure is needed to keep up with increasing traffic demands, understanding in detail how the network deployment, configuration, and usage influence the overall network capacity is clearly crucial for evaluating the true potential of IEEE 802.11ac.

Wi-Fi technologies have been the subject of a large number of research studies in the literature, especially in relation to channel selection and the role of interference in limiting network capacity (see e.g. [3]–[8]). However, the vast majority of this work has focused on earlier Wi-Fi generations (especially IEEE 802.11a/b/g), which in particular could only operate using fixed channel bandwidths. While studies on both IEEE 802.11n [9], [10] as well as IEEE 802.11ac [11], [12] have started to appear in the literature, most of these have focused on characterising performance of the respective Wi-Fi links either in isolation, or in the presence of a small number of interferers. Thus their usefulness in understanding the network-level behaviour of the new Wi-Fi standard amendments, especially in dense deployments, is rather limited. Two initial studies towards denser scenarios have recently appeared [13], [14], but they also yield only limited insight on how the complex interplay of network density, traffic load, and channel configuration influences network performance. In particular, [13] studies the *uplink* performance and fairness with only a single access point (AP) active at a time, whereas the authors in [14] focus mostly on the design of mechanisms to improve fairness, whereas their limited baseline throughput and fairness measurements are only conducted with saturated TCP traffic in one fixed 6-AP deployment only. In [12] interference between IEEE 802.11ac traffic flows is also studied, but only with two AP/client pairs and only with saturated UDP traffic.

In this paper we report results from extensive measurements conducted in a large-scale IEEE 802.11ac indoor testbed enabling interference studies involving up to 24 nodes using diverse channel configurations and traffic models. Our focus

978-1-5386-2723-5/17/$31.00

TABLE I
IEEE 802.11AC TESTBED EQUIPMENT SPECIFICATIONS

	PCE-AC68	RT-AC87U
Peak data rate	1.3 Gbps	1.734 Gbps
Spatial streams	3 × 3 MIMO	4 × 4 MIMO
Channel width	20/40/80 MHz	20/40/80 MHz
DFS support	no	yes
Transmit power	18–22 dBm	19 dBm
Chipset	BCM4360	QT3840BC+QT2518B

is especially on understanding the *combined effects* of the chosen channel bandwidth, spatial density of the network deployment, and traffic load on both the network capacity and *fairness* in dense indoor deployments. In particular, we study network-level performance in both saturated and non-saturated traffic conditions. The latter is particularly important, since the capacity of the IEEE 802.11ac link in general far exceeds the capacity of the associated fixed line Internet connection, making it difficult to effectively saturate the wireless link with normal use. Our results show that using 80 MHz channels is beneficial under extremely dense deployments and traffic loads, with narrower channels resulting in lower network throughput as well as potentially extreme unfairness between clients due to adjacent channel interference (ACI) effects. This is in contrast to industry standard planning guidelines [15], [16] that explicitly recommend using narrow channels in dense deployments for maximum capacity. However, our results also show that for bursty TCP traffic loads taking into account backhaul capacity limitations typical of present-day and near future residential and commercial deployments, there is no harm (nor benefit) in using narrower channels.

The rest of this paper is organized as follows. Section II describes our experiment methodology, traffic modelling, and scenarios. In Section III we present and analyze our measurement results. Section IV concludes the paper.

II. EXPERIMENT METHODOLOGY & SCENARIOS

A. Testbed Architecture

Our testbed comprises twelve IEEE 802.11ac AP and client pairs. A client is an Asus PCE-AC68 wireless adapter connected to a PC, using Linux driver version 6.30.223.248. The APs are Asus RT-AC87U routers connected to a blade server. Table I gives the key specifications of our IEEE 802.11ac equipment. The overall architecture of our testbed consists of four blade servers (supporting three APs each) that are connected to a main workstation control PC via a high-speed Ethernet switch, to enable automated traffic generation and experiment control. Individual experiment execution is automated via bash shell scripts which enable the control PC to remotely open clients, transmit traffic from APs, and read the throughput results from the clients.

B. Traffic Modelling and Generation

To study the effects of traffic type on IEEE 802.11ac performance in dense networks, our measurements used a number of distinct traffic generation profiles, differing in terms of: (i) the temporal patterns of the traffic generated, (ii) data rates, and (iii) the transport protocol employed.

We generated two distinct classes of traffic pattern, hereafter called *constant bit rate* (CBR) and *bursty*. CBR traffic was generated for each measurement repetition for a time interval of 10 s constantly at a given data rate, enabling accurate estimation of the average throughput for any given AP/client pair. The objective of generating *bursty* traffic was to model as closely as possible typical Internet usage patterns in modern residential Wi-Fi deployments. For this, we used the *behavioural web traffic model* of Choi and Lamb as the basis [17], incorporating the more recent additions of object level modelling of HTTP traffic into the overall model as described in [18]. Finally, we updated the parameterization of the model in order to correspond with the increased per-user network traffic based on recent measurement studies [19]. The final traffic model is equivalent to a Semi-Markov ON/OFF model where the duration of the ON periods follows a Gamma distribution with the probability density function (PDF)

$$f(x \mid a, b) = \frac{1}{b^a \Gamma(a)} x^{a-1} \exp\left(\frac{-x}{b}\right), \tag{1}$$

and the OFF period duration is Weibull distributed with PDF

$$f(x \mid a, b) = \frac{b}{a}\left(\frac{x}{a}\right)^{b-1} e^{-(x/a)^b}. \tag{2}$$

The maximum likelihood estimates for the Gamma distribution of ON period durations (measured in seconds) were found to be $a = 0.99$ and $b = 3.43$, and the corresponding estimates for the Weibull distributed OFF period lengths are $a = 2.23$ and $b = 0.37$. In the ON state CBR traffic is sent at a given data rate, whereas in the OFF state no traffic is generated. Since the traffic load under bursty traffic is highly variable, much longer measurement durations are needed to reliably estimate the average achieved throughput. Therefore, in our experiments with bursty traffic each repetition lasted for 180 s.

Several different data rates were used to study the impact of traffic load on network performance. First, we define the *saturated* data rate as the offered load sufficient to saturate a single 80 MHz link using CBR UDP traffic. This was empirically found to be 800 Mbps and 850 Mbps for the single-room and four-room experiment scenarios in Section II-C, respectively (in our calibration measurements, the saturation

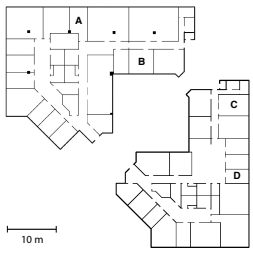

Fig. 1. Floor plan of our two-wing indoor office measurement environment: IEEE 802.11ac AP/client pairs were variously placed in rooms A, B, C, D for different experiment scenarios.

(a) 1 × 80 MHz channel: (36–48)

(b) 2 × 40 MHz channels: black=(36–40), blue=(44–48)

(c) 4 × 20 MHz channels: black=(36), blue=(40), green=(44), red=(48)

Fig. 2. Placement of IEEE 802.11ac AP/client pairs for the single-room experiment scenario, for different channel configurations (number of AP/client pair indicates the order in which the corresponding AP was turned on in experiments).

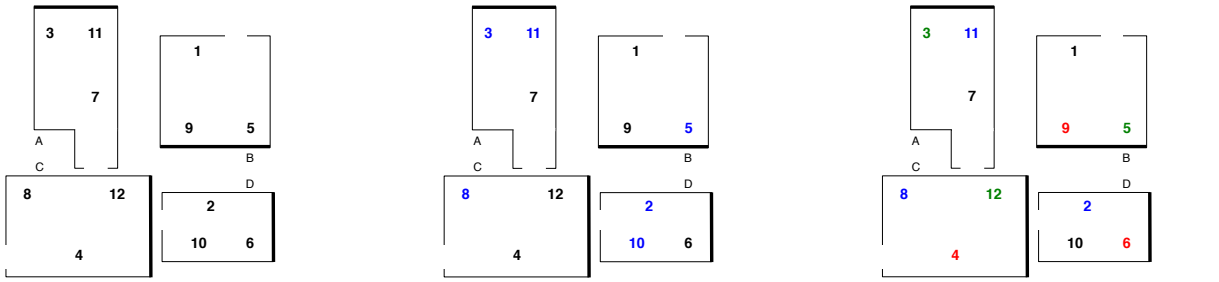

(a) 1 × 80 MHz channel: (36–48)

(b) 2 × 40 MHz channels: black=(36–40), blue=(44–48)

(c) 4 × 20 MHz channels: black=(36), blue=(40), green=(44), red=(48)

Fig. 3. Placement of IEEE 802.11ac AP/client pairs for the 4-room experiment scenario, for different channel configurations (number of AP/client pair indicates the order in which the corresponding AP was turned on in experiments).

throughput varied from 600 to 850 Mbps over the 24 different AP placements, due to MIMO and multipath effects). Second, we consider 35 Mbps as representative of modern residential broadband connection speeds [20], which for most applications will act as the bottleneck for the Wi-Fi traffic. In order to study the impact of traffic patterns on performance while keeping the overall data rate fixed, we also considered intermediate "matched" data rates between CBR and bursty traffic in our experiments, e.g. 15.5 Mbps CBR traffic matching the average offered load of 35 Mbps bursty traffic. Lastly, to consider the effects of future growth, we considered progressively increased offered load from 50 Mbps up to the single-link saturation throughput (*cf.* Section III-B).

The two transport protocols used were UDP and TCP, which carry the overwhelming majority of current Internet traffic. Our traffic models were implemented using *iperf* for the CBR UDP traffic and the *D-ITG* traffic generator [21] for all other traffic profiles, due to its flexible support for ON/OFF models.

C. Experiment Design and Scenarios

We performed all measurements in the typical single-floor indoor office environment shown in Fig. 1, whereby the twelve AP/client pairs were variously placed in four rooms for our different experiment scenarios. Specifically, we considered two different network layouts as shown in Figs. 2 and 3. The single-room scenario in Fig. 2 represents an extremely high network density, where we expect to observe the highest inter-AP interference. By contrast, distributing the AP/client pairs over four rooms in the scenario in Fig. 3 represents a lower (more realistic for the foreseeable future) interference case, by virtue of increased isolation between the Wi-Fi cells due to greater AP-AP separation and wall shielding. We note that we always fixed the AP-client separation to 1 m (and used only one client per AP), since our focus is on studying how inter-AP interference affects the achievable performance *network-wide*, rather than the performance within a single Wi-Fi cell.

For each indoor layout scenario, we performed experiments with different channel configurations. The total channel bandwidth used over the whole network was maintained at 80 MHz in each experiment, with the APs variously allocated to operate on: (i) 1 × 80 MHz channel, or (ii) 2 × 40 MHz channels, or (iii) 4 × 20 MHz channels. All experiments were performed over the four non-DFS channels 36, 40, 44, and 48 in the 5 GHz band. We note that our strategy in allocating APs to channels was always to maximize separation of co-channel APs; the specific AP-to-channel mapping in each scenario is indicated via colour in Figs. 2 and 3.

Finally, for a given indoor layout and channel configuration scenario, the experiment workflow for measuring the Wi-Fi network performance versus the number of operating APs, $N = \{1, \ldots, 12\}$, was as follows. Initially, one AP/client-pair is activated, the AP sends traffic to the client, as per one of the profiles specified in Section II-B, and the client (application-layer[1]) throughput is measured and recorded using *tshark*. This experiment is repeated for six measurement runs. Then the first and the second AP/client pairs are activated simultaneously (with small random delays to ensure random turn-on order to avoid capture effects), both APs send traffic to their associated clients, and the experiment is again repeated six times. The

[1]Direct processing of the capture files collected by *tshark* yields data rates that include overhead from protocol headers. To convert these into the reported application layer throughput, we multiplied them with the fractions of application layer payload size of the captured link layer frame length, i.e. $1472/1514$ and $1460/1514$ for UDP and TCP, respectively.

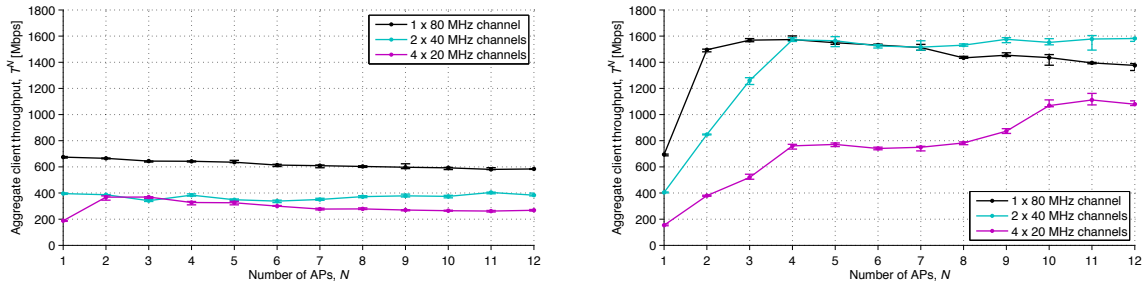

(a) single-room scenario of Fig. 2 (b) 4-room scenario of Fig. 3

Fig. 4. Aggregate client throughput vs. number of active AP/client pairs, for different channel configurations in two different indoor placement scenarios, for saturated CBR UDP traffic.

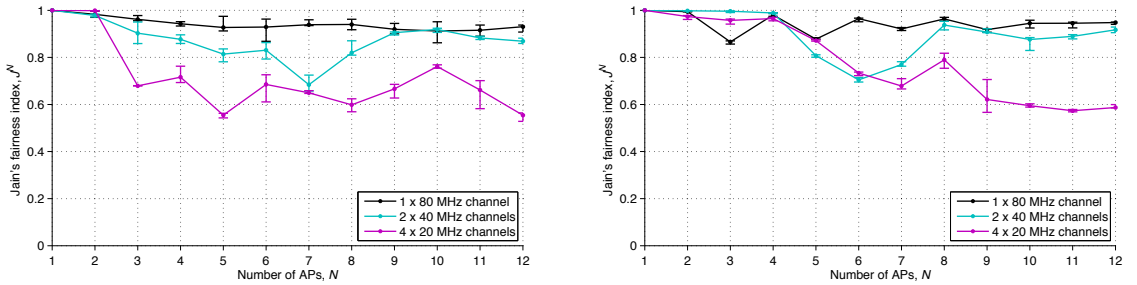

(a) single-room scenario of Fig. 2 (b) 4-room scenario of Fig. 3

Fig. 5. Fairness (Jain's index for achieved client throughput over network) vs. number of APs, for different channel configurations in two different indoor placement scenarios, for saturated CBR UDP traffic.

same procedure is followed, by increasing the number of AP/client pairs turned on at a time, until all twelve AP/client pairs are active. The exact order in which the AP/client pairs are turned on in each experiment scenario is indicated numerically in Figs. 2 and 3.

D. Evaluation Metrics

For each experiment our raw data consists of the measured throughputs $\{R^N_{i,j}\}^{j\in\{1,\ldots,6\}}_{i\in\{1,\ldots,N\}}$ between the i^{th} AP/client-pair, measured over six repetitions j with AP/client-pairs $1,\ldots,N$ being active simultaneously. In Section III we report results based on the throughput distributions of individual routers (i.e. results for $R^N_{i,j}$ for a fixed i and N), as well as results that characterize how the network as a whole is performing. For the latter purpose we study the behaviour of two primary metrics, the *aggregate throughput*

$$T^N_j \equiv \sum_{i=1}^{N} R^N_{i,j}, \tag{3}$$

and *Jain's fairness index* [22]

$$J^N_j \equiv \frac{\left(\sum_{i=1}^{N} R^N_{i,j}\right)^2}{N \times \sum_{i=1}^{N} (R^N_{i,j})^2}, \tag{4}$$

as the number N of active AP/client-pairs is varied. The values of Jain's fairness index range from 1 indicating perfect fairness (each AP/client-pair achieving the same throughput) to $1/N$ (where a single AP/client-pair is communicating successfully and all other pairs obtain zero throughput). For both T^N_j and J^N_j we report for each N the median over the repeated experimental runs j and the 25$^{\text{th}}$ and the 75$^{\text{th}}$ percentiles to show the variability over the experiments.

III. Results & Analysis

A. Saturated Traffic

Figs. 4 and 5 present our measurement results for saturated CBR UDP traffic, in terms of the aggregate client throughput T^N_j and Jain's fairness index J^N_j, respectively, versus the number of operating AP/client pairs N. We emphasize that this traffic profile represents the upper bound on per-AP offered traffic load and correspondingly, worst-case interference and contention in the network.

Fig. 4 shows that for the offered saturated traffic, the client throughput for the single AP case is around 650 Mbps, 400 Mbps, and 200 Mbps in case of 80 MHz, 40 MHz, and 20 MHz channels, respectively[2]. In the single-room scenario of Fig. 2, all co-channel APs are within carrier-sense range of each other. It follows that in Fig. 4a, the aggregate throughput for the 80 MHz channel case remains at the single-link saturation throughput in the case of one or two APs, and then decreases slightly for a greater number of contending APs, due to increased CSMA/CA MAC overhead of greater backoff intervals resulting from increased contention for the channel. By contrast, in the four room scenario of Fig. 3,

[2]We note that the slight variation in T^1 between Figs. 4a and 4b is due to the different placement of AP 1 within Room B for the four-room vs. single-room scenarios, and the resulting change in MIMO and multipath effects.

co-channel APs in different rooms are no longer necessarily within carrier-sense range of each other. Consequently, in Fig. 4b the aggregate client throughput for the 80 MHz channel case increases up until there are two active APs but levels off once the third AP is active. This is because AP 1 in Room B and AP 3 in Room A are still within carrier-sense range and thus contend for channel access time, whereas AP 2 accesses the channel independently as it is isolated in Room D of the other wing (*cf.* Figs. 1 and 3). Correspondingly, we always observe a somewhat lower fairness index J^N for odd N in Fig. 5b. Nonetheless, Fig. 5 shows that for the 80 MHz case the active APs generally get an equal share of channel access time, as reflected by a fairness index of around 0.90–0.95.

In general, as the number of active APs is increased, the aggregate client throughput increases linearly as long as there are *uncoupled independent* channels available, and then reaches a constant saturation throughput value for a greater number of APs. We might then expect that, in the case of APs operating on two 40 MHz channels or four 20 MHz channels, the aggregate throughput in the single-room scenario would increase linearly until the network density reached two and four APs, respectively. However, the curve for 2 $\times$ 40 MHz in Fig. 4a exhibits the same trend as the 80 MHz curve. This shows that adjacent IEEE 802.11ac channels are in fact not independent in the case of very high network density and traffic load, but rather that they are coupled by means of strong adjacent channel interference (ACI). In the case of four 20 MHz channels, the first and fourth channels are far apart enough in frequency to remain largely uncoupled due to high enough transmit spectrum mask attenuation. This is consistent with the aggregate throughput for $N = 2$ APs operating over 20 MHz channels in Fig. 4a being roughly double the single-link throughput, *cf.* the almost constant aggregate throughput for the 40 MHz or 80 MHz channel cases. Moreover, the different frequency separations among pairs of active APs for a given N, and the resulting heterogeneous environment of ACI and co-channel contention experienced by APs across the network, account for the significant reduction in fairness observed in Fig. 5 for narrower channels. Importantly, these dominant ACI effects explain why Fig. 4 and 5 show that, for saturated offered traffic, operation on a single 80 MHz channel is overall superior to splitting the APs across multiple narrower channels.

To study in more detail the effects of ACI on the achieved throughput and fairness, we consider the 4 x 20 MHz channel configuration and present in Fig. 6 plots of the per-client throughput R_i^N versus the number of active AP/client pairs N, for all $i \in \{1, \ldots, N\}$. Let us first consider in the single-room scenario the APs which experience the strongest ACI by operating on the two middle 20 MHz channels (i.e. $i = \{3, 7, 11\}$ and $i = \{4, 8, 12\}$ on channels 40 and 44, respectively). Fig. 6a shows that all but one ($i = 4$) of these APs achieve a near-zero throughput. Namely, in this extremely dense single-room scenario, up to *half* the APs experience *starvation*, resulting in a fairness index as low as 0.55 in Fig. 5b. To investigate the underlying mechanism causing

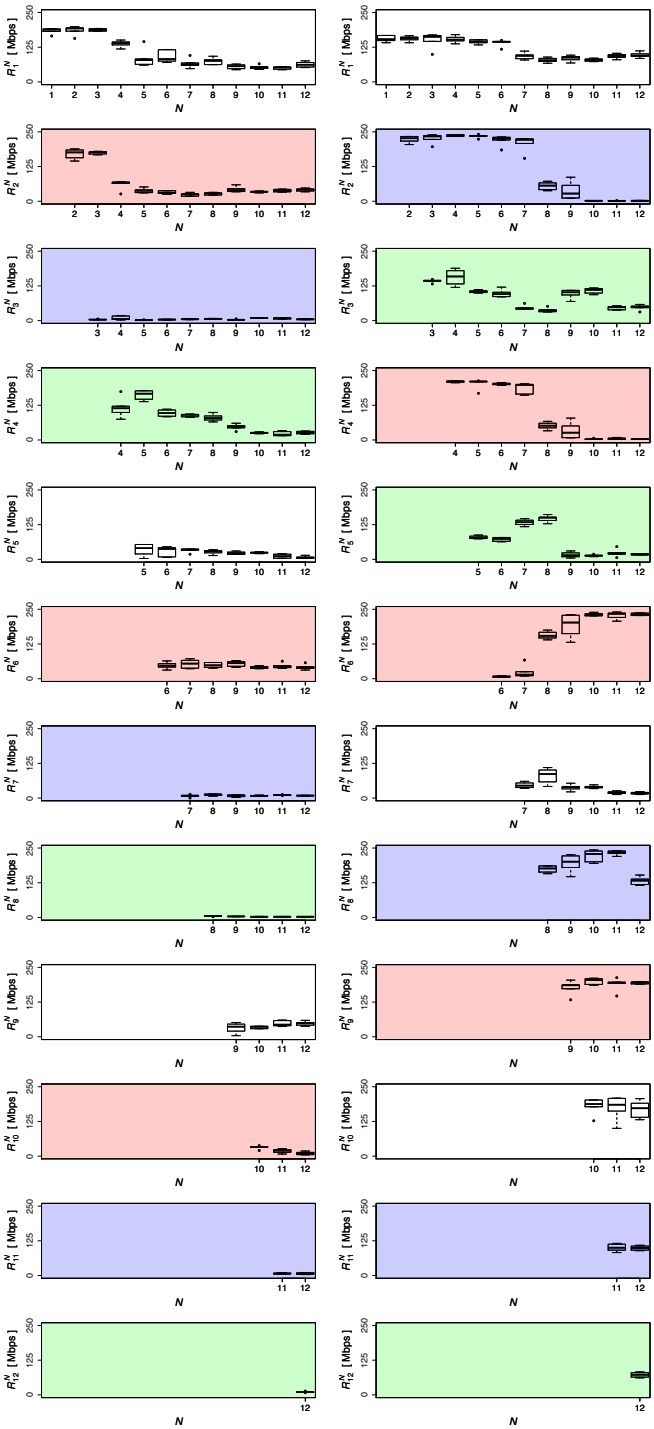

(a) single-room scenario of Fig. 2 (b) 4-room scenario of Fig. 3

Fig. 6. Individual client throughputs vs. number of active client/AP pairs, for saturated CBR UDP traffic in two different indoor placement scenarios, for the 4 x 20 MHz channel configuration (background colour of each plot denotes the channel assigned to that individual client/AP pair, as per Figs. 2 and 3).

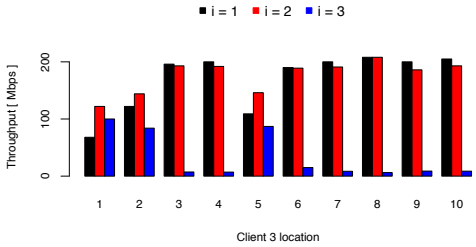

(a) Client throughput R_i^3

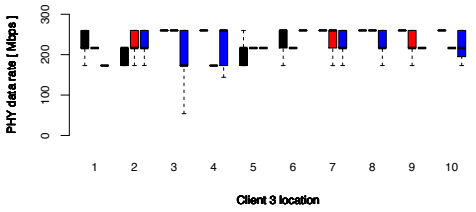

(b) Distribution of PHY rates for DATA frames

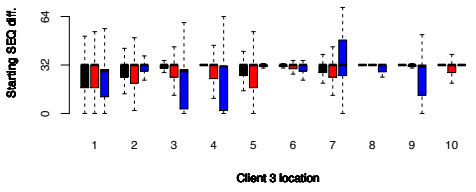

(c) Differences in starting SEQ numbers between successively monitored block ACKs

Fig. 7. Results of experiment studying starvation conditions for "frequency flow-in-the-middle" client 3 vs. its position on a 2 m line (saturated CBR UDP traffic in the single-room scenario, 4 x 20 MHz channel configuration).

this starvation phenomenon, we performed an additional experiment as follows. We conducted saturated offered traffic measurements with AP/client pairs $i = \{1, 2, 3\}$ being active on channels 36, 48, and 40, respectively, with three additional clients deployed in monitoring mode to capture data and control frames[3]. We performed 10 separate experiments, whereby we varied the location of client 3 in 20 cm linear increments over a total distance of 2 m, while keeping the locations of all other clients and all APs fixed. Importantly, we note that IEEE 802.11ac MIMO beamforming is by default enabled in our testbed, and is employed for downlink data transmissions from the APs to their clients. Fig. 7 presents the results, in terms of the mean application layer client throughput R_i^3 (Fig. 7a), the distribution of physical layer (PHY) data rates (Fig. 7b) for the frames transmitted from the APs to their clients, and the differences in the starting sequence numbers between successively monitored block ACKs (Fig. 7c), for each client/AP pair and location of client 3. Regarding Fig. 7c, we note that the block ACK size is constant at 32 frames.

[3]For IEEE 802.11ac this sniffing process is known to be only partially reliable, due to beamforming and chipset limitations, so that a large proportion of data frames is missed by the monitoring clients; most control (e.g. ACK) frames are captured since these are not transmitted in beamforming mode.

Thus in ideal conditions, the differences in starting sequence numbers are also 32. Higher numbers mean ACKs were missed by the monitoring clients. If an ACK is lost (i.e., not received by the AP), the AP will resend the entire window of frames, resulting in a difference in starting sequence numbers of zero. Differences between zero and 32 indicate lost data frames.

Fig. 7a shows that client 3 experiences starvation, as observed for this link originally in Fig. 6a, for all locations except 1, 2, and 5. Whenever client 3 starves, clients 1 and 2 each achieve close to their single-link saturation throughput; otherwise, we observe a diminished throughput for clients 1 and 2, due to ACI from link 3. Since the locations of all other nodes were fixed in our experiment, the starvation of client 3 must be caused by mechanisms which are affected by the changes in multipath fast fading conditions between client 3 and the other nodes. This corresponds to three possible causes: (i) changes in the downlink SINR (signal to interference and noise ratio) conditions; or (ii) changes in the SINR on the uplink (i.e. client to AP transmission of control frames, e.g. ACKs); or (iii) changes in the downlink CSMA/CA contention relationships among APs, due to AP 3 adjusting its beamforming antenna pattern towards client 3 in a way that makes it a "directional deaf node" with respect to the transmissions of other APs.

The first mechanism of changes in downlink SINR cannot be the major cause for client 3 starvation, since e.g. at locations 8 and 10, Fig. 7c shows that data frame errors are very rare and Fig. 7b shows that AP 3 maintains the highest PHY rate, but client 3 nonetheless starves. The second mechanism of changes in uplink SINR cannot be the cause of starvation, since Fig. 7c shows that lost ACKs are very rare (the ACK loss rate never exceeds 2% and is typically much smaller). Therefore, our results suggest that client 3 starvation is caused by the remaining feasible mechanism of changes in downlink contention conditions. Namely, whenever client 3 starves, the carrier-sense threshold of AP 3 is triggered by both AP 1 *and* AP 2 transmissions, amounting to link 3 being the starving "flow-in-the-middle" in the frequency domain. (AP 1 and AP 2 are three channels apart and thus never cause enough mutual ACI to contend for channel access.) It is then sufficient for AP 3 to become "deaf" to *either* AP 1 or AP 2 (due to its changed beamforming pattern towards client 3), to escape the starvation condition; AP 3 then mutually contends with the one AP it does hear, while the other adjacent channel AP unilaterally defers to AP 3's transmissions, resulting in the channel sharing behaviour evident in Fig. 7a at client 3 locations 1, 2, and 5. As further support for this explanation of the starvation behaviour, we additionally repeated this experiment with beamforming turned *off* for all APs; in this case (results not shown for the sake of brevity), client 3 *consistently* experiences starvation at all 10 locations. Similarly, we confirmed that qualitatively similar throughput behaviour as in the original experiment in Fig. 7a is also observed by changing the location of AP 3 instead of the location of client 3.

Downlink starvation can also be observed in Fig. 6b for the four-room scenario. However, the situation here is more com-

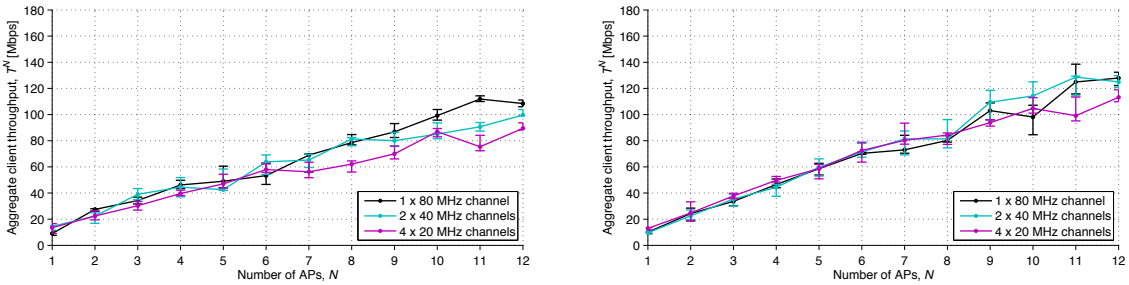

(a) single-room scenario of Fig. 2 (b) 4-room scenario of Fig. 3

Fig. 8. Aggregate client throughput vs. number of active AP/client pairs, for different channel configurations in two different indoor placement scenarios, for 35 Mbps bursty TCP traffic.

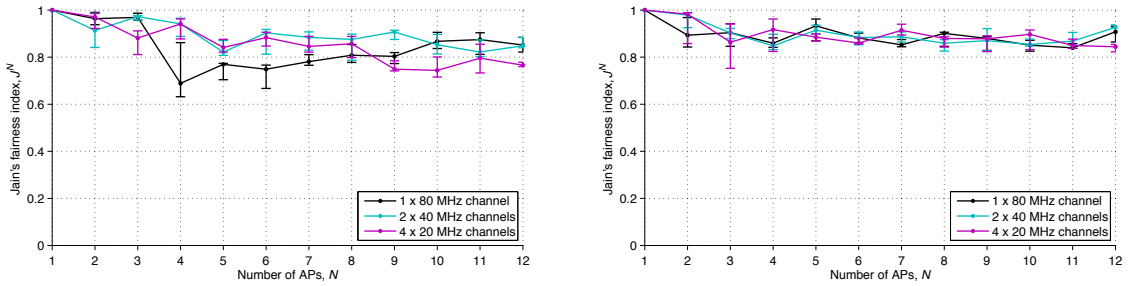

(a) single-room scenario of Fig. 2 (b) 4-room scenario of Fig. 3

Fig. 9. Fairness (Jain's index for achieved client throughput over network) vs. number of active AP/client pairs, for different channel configurations in two different indoor placement scenarios, for 35 Mbps bursty TCP traffic.

plex, due to the heterogeneous interference coupling among pairs of nodes in this multi-room indoor deployment. For example, let us consider APs $i = \{2, 4, 6, 8, 10\}$ which are deployed in Rooms C and D within the same wing. Activating AP 10 on channel 36 causes starvation of client 2 which is operating on channel 40 in the same room, i.e. $R_2^{N \geq 10} \approx 0$ in Fig. 6b; here client 2 is the frequency flow-in-the-middle with respect to the other two clients 10 and 6 in Room D. The starvation of client 2 has cascading effects, increasing the throughput of the co-channel client 8 in Room C, which in turn starves the adjacent channel client 4 in Room C, and finally increases the throughput for its co-channel client 6 in Room D. Another interesting example of the complex ACI interactions in the multi-room deployment is the change in throughput for clients $i = \{2, 4, 6\}$ observed in Fig. 6b once AP 8 is activated. Prior to this, client 6 achieves a very low throughput, presumably due to combined co- and adjacent channel interference on its uplink from clients 2 and 4. Once AP 8 is activated, client 2 and 4 throughputs are reduced, which in turn greatly increases the throughput of client 6. It follows that we observe relatively higher fairness in Fig. 5b for the 20 and 40 MHz channel cases for $N = \{4, 8, 12\}$, when the distribution of ACI is most even across the network owing to an even load of APs per channel. Finally, this variation in throughput fairness, and the underlying starvation of many nodes due to heterogeneous ACI effects, explains why the aggregate throughput in Fig. 4b for the 20 and 40 MHz channel cases does not steadily decrease for $N \geq 4$ as it does in the 80 MHz case, but instead increases.

B. Modern Residential Internet Traffic and Future Growth

Figs. 8 and 9 present our measurement results for 35 Mbps bursty TCP traffic, in terms of the aggregate client throughput T_j^N and Jain's fairness index J_j^N, respectively, versus the number of operating AP/client pairs N. We note that this traffic profile models typical modern residential Wi-Fi deployments, as discussed in Section II-B. Figs. 8 and 9 show that there is no significant difference in the network-wide performance achieved using the different channel width configurations, regardless of the network density. Importantly, this indicates that the strong ACI effects observed for the saturated traffic case in Section III-A are not present at these much lower per-link residential Internet traffic loads, even in the extreme case of 12 links operating in a single room.

For all channel configurations and for both indoor placement scenarios, the following trends are evident in Figs. 8 and 9. First, for both metrics we observe a much higher variability across the six experiment repetitions compared to the saturated CBR traffic case (i.e. higher error bars). Second, we observe a rather slight reduction in fairness in Figs. 9 as N is increased, down to around 0.8 and 0.9 for the single and four room cases, respectively. Third, we observe a slight flattening off of aggregate throughput in Figs. 8 as N is increased. Nonetheless, even for the extremely high network density of 12 links operating in a single room, these results demonstrate that IEEE 802.11ac scales rather well both in terms of aggregate throughput and fairness.

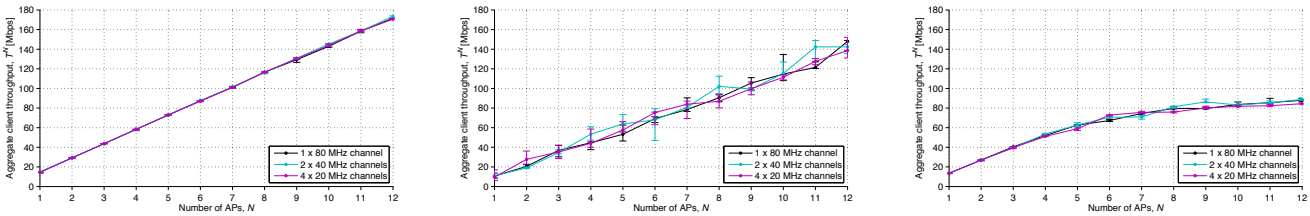

(a) UDP 15.5 Mbps CBR traffic (b) UDP 35 Mbps bursty traffic (c) TCP 15.5 Mbps CBR traffic

Fig. 10. Aggregate client throughput vs. number of active AP/client pairs, for different channel configurations in the single-room scenario of Fig. 2, for different traffic types with matched average offered traffic load.

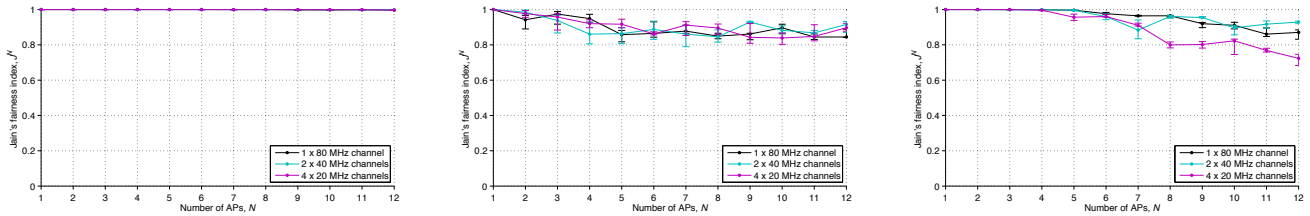

(a) UDP 15.5 Mbps CBR traffic (b) UDP 35 Mbps bursty traffic (c) TCP 15.5 Mbps CBR traffic

Fig. 11. Fairness (Jain's index for achieved client throughput over network) vs. number of active AP/client pairs, for different channel configurations in the single-room scenario of Fig. 2, for different traffic types with matched average offered traffic load.

To show the underlying causes of these behaviours, in Figs. 10 and 11 we consider individually the effects of TCP (Figs. 10c and 11c) and burstiness (Figs. 10b and 11b) against a baseline of UDP CBR traffic with a matched average load of 15.5 Mbps (Figs. 10a and 11a). Figs. 10a and 11a show that with UDP CBR traffic, the aggregate throughput scales linearly with increasing N, and that perfect per-client fairness is always achieved, i.e. $J^{N\in\{1,\ldots,12\}} = 1$. This confirms that the volume of offered traffic is not responsible for any of the three trends in Figs. 8 and 9 discussed above. Figs. 10b and 11b show that the variability in Figs. 8 and 9 is caused largely by the burstiness of the offered traffic. Figs. 10c and 11c indicate that the degradation in throughput and fairness for large N in Figs. 8 and 9 is caused by TCP, as TCP congestion control slows down its sending rate in response to increased MAC contention (i.e. an instance of the well-known problem of TCP over wireless [23]).

Finally, we consider whether network densification and traffic growth are expected to cause the ACI effects observed for the saturated traffic case in Section III-A to become dominant in *emerging* residential Wi-Fi deployments in the near future. To this end, in Figs. 12 and 13 we present our measurement results for CBR UDP traffic for increasing offered traffic volumes (up to the single-link saturation value for each channel width), in terms of the aggregate client throughput T_j^N and Jain's fairness index J_j^N, respectively, versus the number of operating AP/client pairs N. For this dimensioning study of real indoor deployments, we consider the four-room scenario (*cf.* the extreme interference single-room case). These results show that IEEE 802.11ac remains perfectly scalable and fair up to $N = 12$ for traffic loads of up to 100 and 50 Mbps for 40/80 MHz and 20 MHz channel deployments, respectively. We note that 50 Mbps CBR traffic is the equivalent average traffic volume as bursty traffic with the ON rate of around 113 Mbps, which exceeds the EU broadband policy target of 100% of households having 100 Mbps residential broadband access by 2025 [24]. Importantly, our results thus suggest that our conclusion, that there is neither network-wide harm nor benefit from using narrower channels, in even very dense IEEE 802.11ac residential deployments, is expected to hold well into the next decade. Namely, choice of channel width is in practice immaterial in IEEE 802.11ac deployments, even with significant densification of the wireless access network and taking traffic growth into account. Of course, the advantage of 80 MHz deployments remains that no channel allocation network planning is required and that such deployments are even more future-proof.

IV. Conclusions

In this paper we reported measurement results from extensive experiments in a large-scale IEEE 802.11ac indoor testbed, studying the combined impact of channel bandwidth, traffic profile, and AP density and placement on the overall network-level throughput and fairness. Our results show that wide 80 MHz channels are only beneficial in very dense deployments with extreme offered traffic volumes, due to significant ACI which couples narrower channels. However, for 35 Mbps bursty TCP traffic representative of modern residential Wi-Fi deployments, our results show that there is no harm in using lower bandwidth (20 MHz or 40 MHz) channels. We also showed that if the average offered traffic load per AP were increased to around 100 Mbps, IEEE 802.11ac with 20 MHz channels would still achieve scalable network performance with up to 12 active APs. From this we conclude that ACI effects are not expected to become dominant in the near future, even when taking into account

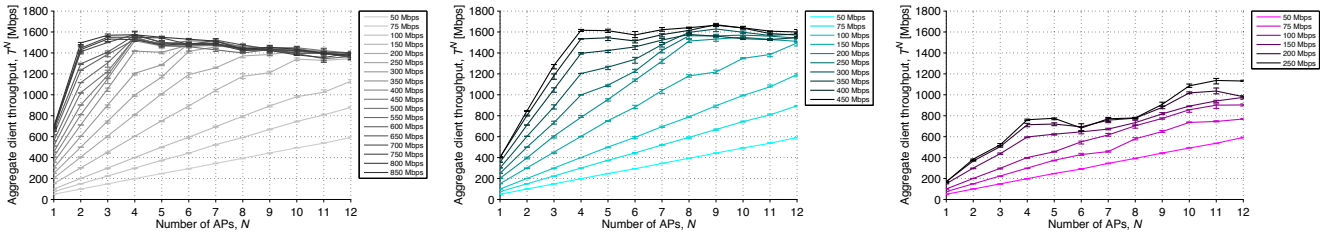

(a) 1×80 MHz channel (b) 2×40 MHz channels (c) 4×20 MHz channels

Fig. 12. Aggregate client throughput vs. number of active AP/client pairs, for different channel configurations in the four-room scenario of Fig. 3, for UDP CBR traffic with increasing offered load.

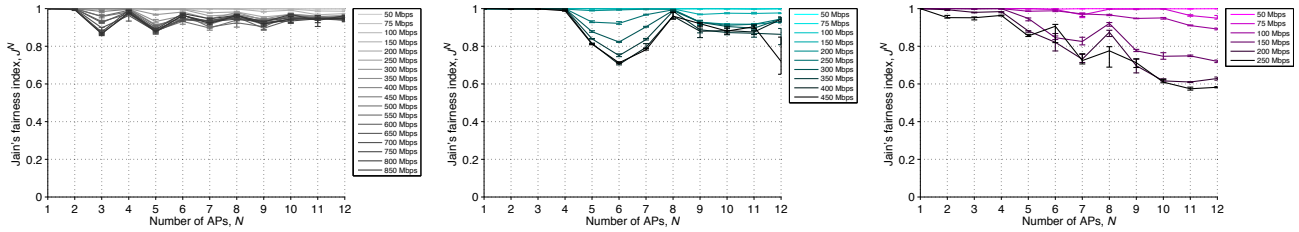

(a) 1×80 MHz channel (b) 2×40 MHz channels (c) 4×20 MHz channels

Fig. 13. Fairness (Jain's index for achieved client throughput over network) vs. number of active AP/client pairs, for different channel configurations in the four-room scenario of Fig. 3, for UDP CBR traffic with increasing offered load.

network densification and traffic growth in emerging Wi-Fi deployments. Our ongoing work is further studying the sensitivity of dense IEEE 802.11ac network performance to *heterogeneous* deployments with respect to channel width and traffic type and volume per AP.

ACKNOWLEDGEMENT

We thank Christos Tsirakis and Markus Stroot for help with setting up the testbed and Andra Voicu for useful discussions.

REFERENCES

[1] O. Bejarano, E. W. Knightly, and M. Park, "IEEE 802.11 ac: from channelization to multi-user MIMO," *IEEE Commun. Mag.*, vol. 51, no. 10, pp. 84–90, 2013.

[2] E. Perahia and R. Stacey, *Next Generation Wireless LAN: 802.11n and 802.11ac*. Cambridge University Press, 2013.

[3] M. A. Ergin, K. Ramachandran, and M. Gruteser, "Understanding the effect of access point density on wireless LAN performance," in *Proc. ACM MOBICOM*, Montreal, 2007.

[4] ——, "An experimental study of inter-cell interference effects on system performance in unplanned wireless LAN deployments," *Computer Networks*, vol. 52, no. 14, pp. 2728–2744, 2008.

[5] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot, "Measurement-based self organization of interfering 802.11 wireless access networks," in *Proc. IEEE INFOCOM*, Anchorage, 2007.

[6] A. Kashyap, S. Ganguly, and S. R. Das, "A measurement-based approach to modeling link capacity in 802.11-based wireless networks," in *Proc. ACM MOBICOM*, Montreal, 2007.

[7] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-based models of delivery and interference in static wireless networks," *ACM SIGCOMM CCR*, vol. 36, no. 4, pp. 51–62, 2006.

[8] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer, "Understanding congestion in IEEE 802.11b wireless networks," in *Proc. ACM IMC*, Berkeley, 2005.

[9] V. Shrivastava, S. Rayanchu, J. Yoonj, and S. Banerjee, "802.11n under the microscope," in *Proc. ACM IMC*, Vouliagmeni, 2008.

[10] S. Fiehe, J. Riihijärvi, and P. Mähönen, "Experimental study on performance of IEEE 802.11n and impact of interferers on the 2.4 GHz ISM band," in *Proc. ACM IWCMC*, Caen, 2010.

[11] M.-D. Dianu, J. Riihijärvi, and M. Petrova, "Measurement-based study of the performance of IEEE 802.11 ac in an indoor environment," in *Proc. IEEE ICC*, Sydney, 2014.

[12] Y. Zeng, P. H. Pathak, and P. Mohapatra, "A first look at 802.11 ac in action: Energy efficiency and interference characterization," in *Proc. IFIP Networking*, Trondheim, 2014.

[13] L. Kriara, E. C. Molero, and T. R. Gross, "Evaluating 802.11 ac features in indoor WLAN: An empirical study of performance and fairness," in *Proc. WiNTECH Workshop in ACM MobiCom*, New York, 2016.

[14] R. Karmakar, S. Chattopadhyay, and S. Chakraborty, "Channel access fairness in IEEE 802.11ac: A retrospective analysis and protocol enhancement," in *Proc. ACM MobiWac*, Malta, 2016.

[15] Cisco Meraki AP Documentation, "High Density Wi-Fi Deployment Guide," retrieved 19.12.2016. [Online]. Available: https://documentation.meraki.com/MR/Deployment_Guides/High_Density_Wi-Fi_Deployment_Guide_(CVD)

[16] Aruba Networks, "Very High-Density 802.11ac Networks Planning Guide," retrieved 19.12.2016. [Online]. Available: http://www.arubanetworks.com/assets/vrd/Aruba_VHD_VRD_Planning_Guide.pdf

[17] H.-K. Choi and J. O. Limb, "A behavioral model of web traffic," in *Proc. IEEE ICNP*, Toronto, 1999.

[18] J. J. Lee and M. Gupta, "A new traffic model for current user web browsing behavior," *Tech. rep. Intel corporation*, 2007.

[19] R. Pries, Z. Magyari, and P. Tran-Gia, "An HTTP web traffic model based on the top one million visited web pages," in *Proc. EURO-NGI*, Karlskrona, 2012.

[20] European Commission, "Quality of Broadband Services in the EU," 2015. [Online]. Available: https://ec.europa.eu/digital-single-market/en/news/quality-broadband-services-eu

[21] A. Botta, A. Dainotti, and A. Pescapè, "Multi-protocol and multi-platform traffic generation and measurement," in *Proc. IEEE INFOCOM Demos*, Anchorage, 2007.

[22] R. Jain, D.-M. Chiu, and W. R. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer system," Digital Equipment Corporation, Tech. Rep., 1984.

[23] G. Xylomenos, G. C. Polyzos, P. Mahonen, and M. Saaranen, "TCP performance issues over wireless links," *IEEE Commun. Mag.*, vol. 39, no. 4, pp. 52–58, 2001.

[24] European Commission, "Broadband Strategy & Policy," 2016. [Online]. Available: https://ec.europa.eu/digital-single-market/en/broadband-strategy-policy