

Wi-Fi Channel Bonding: an All-Channel System and Experimental Study from Urban Hotspots to a Sold-Out Stadium

Sergio Barrachina-Muñoz , Boris Bellalta, *Senior Member, IEEE* and Edward Knightly, *Fellow, IEEE*

Abstract—In this paper, we present WACA, the first system to simultaneously measure all 24 Wi-Fi channels that allow channel bonding at 5 GHz with microsecond scale granularity. With WACA, we perform a first-of-its-kind measurement study in areas including urban hotspots, residential neighborhoods, universities, and even a game in Futbol Club Barcelona’s Camp Nou, a sold-out stadium with 98,000 fans and 12,000 simultaneous Wi-Fi connections. We study channel bonding in this environment, and our experimental findings reveal the underpinning factors controlling throughput gain, including channel bonding policy and spectrum occupancy statistics. We then show the significance of the gathered dataset for finding insights, which would not be possible otherwise, given that simple channel occupancy models severely underestimate the available gains. Likewise, we characterize the risks of channel bonding due to other BSS’s, including their missed transmission opportunities and potential collisions due to imperfect sensing of bonded transmissions. We explore 802.11ax which imposes constraints on bonded channels to avoid fragmentation and defines different modes that can trade implementation complexity for throughput. Lastly, we show that the stadium, while seemingly too highly occupied for channel bonding gains, has transient gaps yielding impressive gains.

Index Terms—channel bonding, measurement campaign, spectrum occupancy, IEEE 802.11ax, Camp Nou stadium

I. INTRODUCTION

Channel bonding is a key mechanism for increasing Wi-Fi data rates as the maximum data rate increases in proportion to the total channel bandwidth. In Wi-Fi, while the *basic* channel width remains 20 MHz, the maximum *bonded* channel width has increased from 40 MHz in 802.11n to 160 MHz in 802.11ac/ax [1], and 320 MHz in 802.11be [2]. During this time, the standard has evolved to not only support wider bandwidths, but also to enable more sophisticated channel bonding policies: in 802.11n, only “static” channel bonding was allowed in which a fixed group of pre-configured channels must

always be bonded. Today, the standard enables a far richer set of capabilities including dynamic selection of channel width as well as bonding both contiguous and non-contiguous channels. In this paper, we provide the first experimental study of channel bonding policies and performance factors with a high-resolution all-channel measurement study.¹ In particular, we make the following contributions. First, we develop WACA as the first system to simultaneously measure all 24 Wi-Fi channels at 5 GHz with a 10 μ sec sampling rate. WACA employs multiple synchronized WARP software defined radios (SDRs) and has 24 antennas and 24 radio frequency (RF) chains to match the number of Wi-Fi channels. Our approach contrasts with prior work that uses one or several RF chains, thereby encountering *second* scale channel switching delays. Thus, no prior methods capture all channels at the channel-access μ sec timescale.

Second, using WACA, we conduct extensive measurement campaigns covering two continents, dense urban areas, and places of interest such as university campuses, apartment buildings, shopping malls, and hotels. We also perform measurements in one of the largest sports stadiums in the world, Futbol Club Barcelona’s Camp Nou. The shortest campaign took 20 minutes and the longest covered more than 1 week, and the total number of samples exceeds 10^{11} . In all cases, we record signal strength on all channels at 10 μ sec sample rate, from which we infer the epochs for which each channel is occupied [4]. Notice that the traces compose an unprecedented dataset to find insights otherwise not possible. We highlight Wi-Fi as the main research application of the dataset but studies on different technologies working in the 5 GHz band could also be conducted, e.g., cognitive radios or coexistence between Wi-Fi and LTE and/or Internet of Things (IoT) technologies, to mention a few examples.

Third, we evaluate the performance of a channel bonding basic service set (BSS) using the aforementioned high-resolution traces of channel activity. As the channel occupancies are highly dynamic, we group epochs according to their average utilization. While the stadium measurements have band occupancies as high as 99%, even dense urban scenarios with many competing BSS’s yielded maximum occupancies near 45%. Beginning with the non stadium scenarios, we first experimentally study the foundational elements controlling the gains of channel bonding. Our experiments assess multiple aspects: contiguous vs. non-contiguous channel bonding, pri-

Manuscript received MONTH DAY, YEAR; revised MONTH DAY, YEAR; accepted MONTH DAY, YEAR; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor XXX. Date of publication MONTH DAY, YEAR; date of current version MONTH DAY, YEAR. The work from S. Barrachina-Muñoz and B. Bellalta has been partially supported by Cisco, WINDMAL PGC2018-099959-B-I00 (MCIU/AEI/FEDER,UE), and SGR-2017-1188. The research of E. Knightly was supported by Cisco, Intel, and by NSF grants CNS-1955075 and CNS-1824529. (Corresponding author: Sergio Barrachina-Muñoz).

Sergio Barrachina-Muñoz is with the SMARTECH department, Centre Tecnològic Telecomunicacions Catalunya (CTTC), 08860 Barcelona, Spain (e-mail: sbarrachina@cttc.cat).

Boris Bellalta is with the Wireless Networking research group, Universitat Pompeu Fabra, 08018 Barcelona, Spain (e-mail: boris.bellalta@upf.edu).

Edward Knightly is with the Rice Networks Group, Rice University, TX 77005 USA (e-mail: knightly@rice.edu).

¹This paper is an extension of Wi-Fi All-Channel Analyzer [3].

mary channel selection, inter-channel correlation, models for channel occupancy, bandwidth deprivation, hidden nodes, and 802.11ax channelization constraints.

We find that the increased flexibility of non-contiguous channel bonding compared to contiguous can yield a throughput improvement of up to $5\times$. Unfortunately, the scenarios for these gains occur quite rarely yielding modest average gains of less than 10%. Moreover, selection of the primary channel is critical and improves average throughput by over 60%. We next study correlation among channel occupancies and show that it aids further channel bonding. Namely, when the primary channel becomes available, it increases the likelihood that other channels are available too. We also show that simple channel occupancy models severely underestimate the available gains for channel bonding. Thus, modeling occupancy behavior by introducing inter-channel correlation remains an important avenue for future work.

We next turn to potential detrimental effects involving other in-range BSS's. First, when one BSS channel bonds, others must defer: is channel bonding simply a zero-sum-game transferring bandwidth from one BSS to another? We fortunately find that this is not the case, as the experiments indicate that the channel occupancy dynamics enable the channel bonder to gain at least 45 times what the other BSS's lose. Likewise, hidden nodes could also reduce the performance of channel bonding, as failed transmissions can occur if neighboring BSS's do not sense and defer to a bonded transmission. Our analysis shows that such a scenario of hidden other-BSS traffic would be quite costly, and would often result in single channel transmission outperforming channel bonding, even under relatively low channel occupancy.

Next, we consider the channelization restrictions imposed by IEEE 802.11ax along with its three channel bonding modes: static channel bonding, dynamic channel bonding, and preamble puncturing. Unfortunately, we find that static channel bonding gives up the majority of the available throughput gains. We also find that 802.11ax's channelization restrictions reduce throughput by up to 19% for dynamic channel bonding. In contrast, the increased flexibility of preamble puncturing limits the reduction to less than 6%.

Finally, we deploy WACA in a sold-out sports event with over 98,000 attendees. Partnering with the stadium's network management team, we find that up to 12,000 Wi-Fi clients were simultaneously connected via thousands of access points (APs). Despite a peak occupancy over all channels of 99%, we observed significant transient occupancy lulls of less than 30% yielding gains of up to $5.3\times$. Qualitatively, we find that the lulls in network traffic are inversely related to the match's lulls in action, e.g., half time yields a traffic boom. Nonetheless, shorter time-scale lulls in occupancy, which benefit channel bonding, can be observed at all times.

II. RELATED WORK

To the best of our knowledge, there are no related works in the literature assessing channel bonding through traces simultaneously capturing the whole 5 GHz Wi-Fi band. However, the community has separately covered spectrum measurements and channel bonding through different approaches.

Prior work performed spectrum measurements for Wi-Fi traffic, e.g., [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15]. Example objectives include creating interference maps [15], assessing interference behavior [16], surveying Wi-Fi usage [10], quantifying spectrum occupancy in outdoor testbeds [5], [6], designing efficient scanning methods [7], [11], modeling spectrum use [8], opportunistic spectrum access [9], [14], dynamic channel selection [17], and assessing real-world network behavior by examining data from thousands APs [12]. Unfortunately, no such prior work provided simultaneous measurement across the entire 5 GHz band, which we require for our channel bonding study. While some papers do provide multi-channel measurements, e.g., [18], [14], [9], [5], they do so via sequential scanning, thus taking on the order of seconds to change from one channel to the next, orders of magnitude beyond the transmission time scale for channel bonding. Namely, WACA measures all channels simultaneously using SDRs having a sampling rate of $10\mu\text{sec}$. Moreover, prior work does not consider (for example) stadiums, the potential costs of other BSS's deferring or colliding, the impact of fine-grained channel correlation, nor a diverse set of policies.

Throughput gains of channel bonding have been demonstrated previously in testbeds. In particular, IEEE 802.11n static channel bonding has been shown to be affected not only by link signal quality, but also by the power and rates of neighboring links [19]. Likewise, intelligent channel bonding management was shown to benefit from identifying the signal strength of neighboring links and interference patterns [20]. High bandwidths were shown to be vulnerable to increased thermal noise or the power per Hertz reduction [21], [22]. Finally, an experimental study of IEEE 802.11ac channel bonding (both static and dynamic) showed that unplanned primary channel selection and bandwidth allocation may severely degrade the throughput of links operating at wider channels [23]. Nonetheless, existing experimental results have targeted only on one or few controlled links at most. In contrast, we develop WACA as a monitoring system in order to measure multiple BSS's in various operational settings, thus tackling how a channel bonding BSS interacts with surrounding BSS's in a broad set of channel bonding policies and scenarios.

Simulation studies and analytical models have also been employed to study channel bonding, e.g., early simulation studies demonstrated throughput gains of channel bonding compared to single-channel transmission [24], [25]. Likewise, analytical models have been proposed to study channel bonding, especially through Markov chains [26], [27], [28], [29], [30]. For instance, high-density deployments are evaluated in [27], showing exposure to unfairness in groups of overlapping BSS's. Analytical models for unsaturated traffic have been also proposed [29], [31]. As well, different channel bonding policies were introduced and modeled for spatially distributed scenarios [32], [31]. Lastly, primary channel selection was studied in [33], [34], [35], [36], traffic demands in [37], and QoS constraints in [38]. However, such work does not have an experimental validation as presented here. In this regard, WACA datasets can be used for validation of the aforementioned methods under real-world occupancies in future work.

III. WI-FI ALL-CHANNEL ANALYZER

A. Overview

Our objective is to simultaneously capture activity on all Wi-Fi channels, i.e., all 24 basic (non aggregated) 20 MHz channels in the 5 GHz band that permit channel bonding. In principle, this could be achieved with an off-the-shelf spectrum analyzer. However, most spectrum analyzers are not capable of dealing with the required bandwidth of this objective, i.e., they cannot simultaneously measure the entire Wi-Fi 5-GHz band: 645 MHz ranging from channel 36 to 161 (i.e., from 5170 to 5815 MHz). Moreover, wide-band spectrum analyzers that can cover this bandwidth lack resolution to analyze basic channels within the band. Likewise, one could envision a system comprising 24 off-the shelf Wi-Fi cards as sniffers, one per basic channel. Unfortunately, such a system would be quite unwieldy and would introduce a challenge of ensuring synchronicity among wireless cards: restricting the delay between channel measurements to the order of nano/microseconds is unfeasible due to the hardware interrupt latency and jitter from the different peripherals [39], [40].

Thus, we design WACA to simultaneously measure power (and I/Q signals if required) on all 5-GHz basic channels. Key benefits of WACA include the simplicity of experimental procedures (from deployment to post-processing), a dedicated RF chain per channel (covering the whole band and easing hardware failure detection), and the ease of configuration empowered by the WARPLab framework [41]. As for the latter, WACA is high-level programmed in MATLAB since it relies on the WARPLab Reference Design. Low-level code modifications (e.g., FPGA and embedded development) are not required to start using the platform.

B. Building Blocks

The key building blocks of WACA are (i) six WARP v3 programmable wireless SDRs [42], (ii) six FMC-RF-2X245 dual-radio FMC daughterboards, (iii) 24 5-GHz antennas (one per RF chain), and (iv) one Ethernet switch to enable communication between the WARPLab host (e.g., PC) and the WARP boards. The preminent building block is WARP, a scalable and extensible programmable wireless platform to prototype advanced wireless networks. The FMC-RF-2X245 module dual-radio FPGA Mezzanine Card (FMC) daughterboard extends the capability of WARP v3 boards from 2 to 4 RF chains, so each WARP is fed with 4 antennas. Therefore, by combining 2 stacks of 3 WARP boards each with their corresponding FMC-RF-2X245 daughterboards, we realize 24 RF chains (with one 5-GHz antenna each) enabling us to assign a single RF chain per basic channel. Let us emphasize that we rely on a 24 RF deployment in order to avoid the delays entailed by channel switching, which are to the order of 10 ms from one channel to the next. Finally, the Ethernet switch enables the communication from the WARP nodes to the WARPLab host. Fig. 1a shows the assignment of the RF chains to each basic channel allowed for bonding and Fig. 1b depicts the physical realization of WACA.

C. Measurement Methodology

An iterative procedure is followed for collecting power samples. Namely, in each iteration, WACA first simultaneously measures the power at each basic channel during T_m and then takes T_{proc} to process and forward the measurements to the WARPLab host. Both tasks are sequentially performed until the end of the measurement campaign. Table I shows the main WARPLab parameters used for measurements. Notice that all parameters are fixed except N_{it} , used for determining the duration of the measurement campaign.

Parameter	Value
Active RFs (channels 36 to 161), $\mathcal{R}$	$\{1, 2, \dots, 24\}$
Iteration's measurement duration, T_m	1 s
Iteration's processing duration, T_{proc}	~ 9 s
Duration of a complete iteration, T_{it}	~ 10 s
Original no. of power samples per channel per second, n_s^*	10^7
Downsampled no. of power samples per channel per second, n_s	10^5
Time between consecutive samples, T_s	10 μ s
Number of iterations, N_{it}	125-59500

TABLE I: WARPLab setup.


WARP boards install the MAX2829 transceiver, which has a fixed 10 Msps received signal strength sampling rate. Accordingly, since the measurement duration in an iteration is $T_m = 1$ second, the number of consecutive samples captured per basic channel per iteration is $n_s^* = 10^7$. Then, in each iteration, we store $|\mathcal{R}| \times n_s^* = 24 \times 10^7$ samples. Nonetheless, to decrease the amount of required memory, we downsample the gathered samples in each iteration by a $100\times$ factor, thus reducing the data size per iteration from around 60 MB to 600 KB.² Essentially, while the transceiver measures 1 power sample every 100 ns by default, we keep just 1 sample every $T_s = 10 \mu$ s. Notice that the resulting time scale is also suitable given Wi-Fi timings. Indeed, the short interframe space (SIFS) is the smallest interframe space and takes 16 μ s ($> T_s$).

As for the duration of processing and forwarding (period in which no data is collected), T_{proc} entails a significant yet unavoidable delay overhead with respect to the total duration of an iteration $T_{it} = T_m + T_{proc}$. Specifically, for the host PC used in all the scenarios in the dataset (Intel Core i5-4300U CPU @ 1.9 GHz x 4 and 7.7 GiB memory) and the parameters listed in Table I, $T_{proc} \approx 9$ s for $T_m = 1$ s. Once initiated, WACA operates by itself and no human intervention is required. Notice that in the experiments later presented, we assess periods (or epochs) composed only of consecutive samples from unique iterations. That is, we never mix samples from different iterations, thus overcoming the processing and forwarding gaps.


D. Validation

Before deploying WACA for the measurement campaigns, we perform an extensive set of in-lab controlled experiments

²The iteration size may vary because of the MAT files generation with MATLAB, which compresses numeric data by determining the minimum required amount of memory (e.g., 8 bit instead of double).



(a) IEEE 802.11ac/ax channelization at the 5-GHz and assignment per RF.



(b) Deployment schematic.

Fig. 1: The WACA all WiFi channels spectrum analyzer.




Fig. 2: Validation testbed at the main events room of Universitat Pompeu Fabra's Communication campus.

for validation, comprising over 6×10^8 power samples collected in WACA cross validated with controlled and known transmissions from commercial APs. We conduct the validation of all the boards by first measuring the power perceived against transmission bandwidth.³ Then, we explore the spectrum behavior when setting up different off-the-shelf channel bonding configurations. The measurements were gathered in an empty and large event room (about 300 m²). Figure 2 shows the complete deployment. In particular, we deploy WACA including the PC hosting the WARPLab framework (lower part of the figure), 3 PCs receiving *iperf* traffic (right part of the figure), and 3 PCs sending *iperf* traffic with the corresponding 3 APs enabling the *iperf* connections (left part of the figure). The AP models we use are Asus RT-AC87U (AS), TP-Link AC1750 (TP), and Linksys WRT 3200 ACM (LS).⁴ Remarkably, all of these APs are only capable of performing static channel bonding.

1) *RSSI vs. bandwidth*: We measure the RSSI in dBm⁵ at a fixed distance (3.6 m) in four basic channels at the same time for 20, 40, and 80 MHz transmissions. The general procedure followed in the experiment can be summarized in 4 steps: *i*) the AS AP is set up with bandwidth $b \in \{20, 40, 80\}$ MHz and primary channel $p = 124$ since negligible activity

³Further validations on power vs. central frequency (and RF) and power vs. distance can be found at [3].

⁴Notice that we use different AP models to build a heterogeneous deployment where BSSs' have different capabilities, as it is the most common case in real-world scenarios.

⁵We follow the MAX2829 transceiver data sheet to convert from 10-bit RSSI values (from 0 to 1023) to dBm values as done in [43].

was detected; *ii*) an *iperf* client is placed close to the AP and associates to it; *iii*) an *iperf* server is located at distance $d = 3.6$ m from the AP and also associates to it; finally, *iv*) the *iperf* communication is triggered and WACA captures the resulting RSSI at channels 116, 120, 124, and 128, thus covering 80 MHz. Notice that it is enough to use just one WARP board with the corresponding daughterboard to cover the 4 basic channels of interest. Theoretically, when doubling the bandwidth, the transmission power gets reduced by a half (3 dB). We corroborate this fact by looking at the similar reduction factor in the power detected per Hz in Fig. 3.

2) *Controlled testbed evaluation*: We now measure the activity of 3 overlapping BSSs', all implementing static channel bonding, under different traffic regimes. As shown in Fig. 2, we separate the APs from their corresponding *iperf* servers by $d = 4.8$ m. The WACA platform was placed equidistantly from the central AP and STA also at $d = 4.8$ m. We analyze two particular setups where none and all the BSSs' saturate. The load of every BSS is 20 and 150 Mbps in the first and second setup, respectively. Table II collects the setups' details. Essentially, in the unsaturated scenario, all the BSSs' share the same primary channel (48) but are set with different bonding capabilities. The second setup assigns different primary channels and bandwidth capabilities to each BSS. Note that we changed the primary channels from the previous experiments due to AP hardware restrictions. Nonetheless, we also confirmed that the external interference at the new band of interest was negligible.

WACA collected 1 second of consecutive RSSI samples during the execution of each setup to measure the spectrum occupancy. Namely, we just compare the RSSI against the CCA level and say that a given channel is occupied iff




Fig. 3: Transmission bandwidth effect on power.




Fig. 4: Spectrogram evolution of the controlled scenarios. Slots represent the 10-ms mean occupancy.

Regime	AP	36	40	44	48	Thr. [Mbps]
Unsaturated $\ell = 20$ Mbps	AS				p	20
	TP				p	20
	LS				p	20
Saturated $\ell = 150$ Mbps	AS		p			100
	TP	p				59
	LS				p	16

TABLE II: Setup of controlled experiments. Letter **p** indicates the primary channel whereas colors represent the allocated bandwidth to each AP.

$RSSI > CCA$. Note that, throughout this paper, we compare the RSSI directly with the CCA no matter the type of signal being received. This approach is much more restrictive than relying on carrier sense since CCA thresholds are much smaller than energy detection (-82 and -62 dBm, respectively). This lets us simplify the experiments and consider worst-case occupancy values.


Fig. 4 shows the occupancy evolution at each basic channel for the unsaturated and saturated setups. We confirm that the former setup does not reach full occupancy in any channel, whereas it is the contrary for the later. We also observe the effect of the adjacent channel interference in those channels that should present a priori the same occupancy pattern (e.g., channels 44 and 48 in the saturated setup) since APs statically transmit in 20, 40 and 80 MHz only. In the saturated scenario, the load ($\ell = 150$ Mbps) is high enough to make channels 36 and 40 occupied almost all the time (i.e., occupancy at the first 40-MHz band is always close to 1) while not being able to successfully deliver all the load as indicated by the throughput values in Table II. This has an important consequence for the LS AP. Essentially, since static channel bonding is applied, the LS AP does not benefit from having its primary channel (48) far from the rest of APs. Consequently, its throughput is drastically reduced given that AS and TP introduce inadmissible activity in the first 40 MHz band.

IV. MEASUREMENT CAMPAIGNS

Using WACA, we perform extensive measurement campaigns covering two continents, dense urban areas, and multiple hours of samples in places of interest such as university campuses, apartment buildings, shopping malls, hotels and the Futbol Club Barcelona (a.k.a Barça) stadium (Camp Nou), one of the largest sports stadiums in the world. The measurements were taken from February to August 2019 in Houston, TX,

Id	Location	Type	Duration
1_RVA	Rice Village Apart., HOU	Apartment	1 day
2_RNG	RNG lab at Rice, HOU	Campus	1 day
3_TFA	Technology for All, HOU	Com. center	1 day
4_FLO	Flo Paris, HOU	Cafe	1 hour
5_VIL	Rice Village, HOU	Shopping mall	20 min
6_FEL	La Sagrera, BCN	Apartment	1 week
7_WNO	WN group office, BCN	Campus	1 day
8_22@	22@ area, BCN	Office area	1 day
9_GAL	Hotel Gallery, BCN	Hotel	1 day
10_SAG	Sagrada Familia, BCN	Apartment	3 days
11_FCB	Camp Nou, BCN	Stadium	5 hours

TABLE III: Measurement locations.



(b) Runchart of the normalized occupation for apartments (upper) and universities (lower) highlighting day (yellow) and night (blue). The three curves in the 10_SAG location correspond to three different days of measurements.

Fig. 5: General spectrum occupancy trends.

USA, and Barcelona, Spain. The shortest experiment took 20 minutes and the longest covered more than 1 week. The list of locations is shown in Table III.

Here, we overview the complete dataset: there are 153,033 1-second iterations accounting for 42 hours, 30 minutes, and 3 seconds of actual measurements in the 24 20-MHz channels of the 5 GHz WiFi band. We first assess the entire activity record of each location at band $\mathcal{B}$, where $\mathcal{B}$ is a predefined set of contiguous channels. For instance, U-NII-1 band is defined as $\mathcal{B} = \{1, 2, 3, 4\}$, corresponding to basic channels 36, 40, 44, and 48 (see Figure 1a).⁶ Fig. 5a shows the normalized mean idleness of each band, i.e., the mean number of samples that were found idle in each channel of band $\mathcal{B}$. We observe that the spectrum is idle most of the time in all scenarios except the stadium, indicating that the 5 GHz band is still profoundly underutilized even in densely populated areas.

Fig. 5b shows the daily temporal evolution of 4 example locations (2 apartments in the upper subplots, and 2

⁶The rest of bands are sequentially composed of groups of four consecutive 20-MHz channels. So, the next band is U-NNI-2 with channels $\{5, 6, 7, 8\}$ (basic channels 52, 56, 60, and 64), and the last one is U-NNI-3 with channels $\{21, 22, 23, 24\}$ (basic channels 149, 153, 157, 161).

university campuses in the lower subplots). For the sake of representation, we plot the normalized occupancy of the whole band averaged in periods of 10 minutes. Concretely, we normalize with respect to the highest 10-minutes average occupancy encountered in each location. We clearly observe higher activity at working hours in the campus locations and a much less variable pattern in the apartment locations. In any case, from the low spectrum utilization observed in Fig. 5a, significant opportunities for channel bonding can be expected.

V. EVALUATION METHODOLOGY

A. Overview

Our objective is to study the performance that a (virtual) fully backlogged channel-bonding BSS, w , would obtain if it encountered the channels recorded in the measurement campaigns described above. Therefore, we overlook the potential sources of the spectrum activity and use the power traces as a dummy map of idle/busy samples. Accordingly, we disregard the characteristics of the possible neighboring nodes, including the number of nodes or any synchronization issues. The performance is a function of a number of factors such as which primary channel the bonding BSS selects, which channel bonding policy it employs, and the spectrum occupancy. Such factors are explored in subsequent sections. We assume a constant MCS for computing the throughput. Accordingly, the number of STAs in the BSS does not affect the results since w acts like a whole entity. Here, we describe the overall methodology for all experiments.

The datasets captured by WACA at scenario s are represented by a 2-dimensional matrix Y^s of size $(N_{it}^s \times n_s) \times |\mathcal{R}|$, where N_{it}^s is the number of iterations of scenario s , and any element $y_{t,c}^s$ represents the power value at temporal sample t in basic channel c . From Y^s , we generate a binary matrix X^s of same size through an occupancy indicator function, where any element $x_{t,c}^s$ represents whether channel c was occupied at temporal sample t (1) or not (0). Formally, $x_{t,c}^s = (y_{t,c}^s > CCA : 1, 0), \forall t, c$, where the clear channel assessment (CCA) is set to -83.5 dBm (or 150 10-bit RSSI units), corresponding to the common CCA threshold -82 dBm plus a safety margin of -1.5 dBm. While IEEE 802.11ac/11ax introduce different CCA levels for the primary and secondary channels, in this work, we consider a more restrictive approach by assuming the same threshold in order to fairly compare different channel bonding policies. The mean occupancy at band $\mathcal{B}$ in scenario s is simply defined as

$$\bar{o}_{\mathcal{B}}^s = \frac{\sum_t \sum_{c \in \mathcal{B}} x_{t,c}^s}{N_{it}^s n_s |\mathcal{B}|}. \quad (1)$$

To provide meaningful experiments, we separately consider two 160-MHz bands composed of 8 basic channels: the U-NII-1&2 and part of the U-NII-2c sub-bands, $\mathcal{B}_{1\&2} = \{1, 2, 3, \dots, 8\}$ and $\mathcal{B}_{2c} = \{9, 10, 11, \dots, 16\}$, respectively. These sub-bands cover from channel 36 to 64 and from channel 100 to 128, respectively (see Fig. 1a). Notice that these are the only sub-bands that allow to perform 160-MHz transmissions in the IEEE 802.11ac/ax channelization. Moreover, we focus on epochs (or periods) of duration $T_{per} = 100$ ms

(containing 10^4 temporal samples each) for which the mean occupancy at such sub-bands is at least 5%, i.e., $\bar{o}_{\mathcal{B}} \geq 0.05$, where $\mathcal{B} \in \{\mathcal{B}_{1\&2}, \mathcal{B}_{2c}\}$.

B. State Machine

We develop a discrete state machine that characterizes how the channel-bonding BSS responds to each power sample (or temporal sample) t according to the current state $S(t)$, and channel-bonding policy π , following the 802.11 standard. Given that the channel-bonding BSS is fully backlogged, the set of possible states is $\mathcal{S} = \{Busy, DIFS, BO, TX/RX\}$. State *Busy* indicates that the primary channel is busy, *DIFS* represents the period before initiating the backoff process, the backoff counter is decreased during *BO* state, and *TX/RX* represents the actual frame transmission-reception (including the control frames RTS, CTS, and ACK, the DATA frame, and the SIFS periods in between). We represent the channel-bonding BSS w as an AP and one or multiple clients that would perceive exactly the same spectrum activity as WACA captured in the measurement campaigns and must contend accordingly. To focus on channel bonding effects, we do not consider collisions within the channel bonding BSS, but only collisions that can occur due to other BSS's (see §VII-B).




Fig. 6: Example of the transitions between states.

The set of basic channels selected for transmitting a frame depends on the spectrum occupancy and on the selected channel-bonding policy π . Fig. 6 illustrates an example of the transitions between states. We use Wi-Fi parameters according to IEEE 802.11ax as shown in Table IV. However, empty slots have a non-standard duration $T_{slot} = T_s = 10 \mu s$ rather than $9 \mu s$ (802.11's default value) to align the duration of an idle backoff slot with the sample duration.⁷ Hence, whenever the channel-bonding BSS is in the backoff process at state *BO*, every idle sample at the primary channel p results in a backoff counter decrease of one empty slot. Then, the duration of the inter-frames (e.g., SIFS) and control packets (e.g., RTS) are rounded to multiples of T_s .

After running the state-machine through all the temporal samples in the epoch, we compute the throughput Γ as the number of bytes in the successfully transmitted data packets n_d divided by the duration of the epoch, i.e., $\Gamma = (n_d L_d) / T_{per}$.

C. Channel Bonding Policies and Response

A channel-bonding policy π selects the set of basic channels to aggregate at the end of the backoff provided that the

⁷WACA collects 1 RSSI sample per channel every $10 \mu s$ to reduce unmanageable memory sizes. We selected $10 \mu s$ rather than other values for the convenience of working with round numbers. While this is a limitation as for fine-grained estimations, the insights presented in this paper hold for other inter frame durations as shown in the supplementary material [4].

Param.	Description	Value
CCA	CCA threshold	-83.5 dBm
MCS	MCS index	9 (256-QAM 5/6)
b	basic channel bandwidth	20 MHz
L_d	Length of a data packet	12000 bit
$\mathcal{B}$	Allocated set of basic channels	$\mathcal{B} \subseteq \mathcal{R}$
p	Primary channel	$p \in \mathcal{B}$
N_a	Max. no. of agg. packets per frame	64
T_e	Duration of an empty slot	10 μ s
T_{SIFS}	SIFS duration	20 μ s
T_{DIFS}	DIFS duration	30 μ s
T_{PIFS}	PIFS duration	30 μ s
T_{RTS}	RTS duration	50 μ s
T_{CTS}	CTS duration	40 μ s
T_{BACK}	Block ACK duration	50 μ s
TXOP	Max. duration of a TXOP	5 ms
CW_{min}	Min. contention window	16
m	No. of backoff stages	5

TABLE IV: Trace-driven setup.

primary channel is available. Namely, *contiguous* channel bonding can select a set of channels both above and below the primary channel, provided they are consecutive. In contrast, *non-contiguous channel bonding* can combine all available channels at the time the primary channel becomes available. In §VIII, we consider additional channelization constraints imposed by 802.11ax.

How will other BSS's respond to the channel bonding BSS? For most experiments, we consider that they will defer their transmissions. We call it *hinder* interaction model. Namely, the channel bonding BSS needs the channels to be available only when its countdown timer expires. If the bonding BSS does transmit but the trace indicates that a channel would have been occupied at some point during the transmission, we consider that such other BSS's will sense the bonding BSS and defer. The exception is in §VII-B in which we consider that other BSS's are hidden, do not sense the channel bonder, and cause a collision. We call the latter *hidden* interaction model.⁸

VI. EVALUATION OF BONDING GAINS AND ORIGINS

Here, we experimentally study the foundational elements of channel bonding including contiguous (CO) vs. non-contiguous (NC) channel bonding, load, primary channel selection, and cross-channel correlation. Due to the disparate nature of the stadium scenario, we focus our study on the non-stadium measurement campaigns to provide insights on a more homogeneous dataset. We leave the stadium's study for a dedicated section (§IX).

A. Contiguous vs. Non-Contiguous Aggregation

For each transmission, non-contiguous channel bonding can utilize additional channels as compared to contiguous, by “skipping over” the busy channels to find the next unused one. Here, we explore the gains of this flexibility as well

⁸Assuming that surrounding nodes will not vary their traffic patterns as a consequence of our virtual BSS's activity is a clear limitation. However, we anticipate that the *hinder* and *hidden* interaction models provide factual insights on worst and best case scenarios.

as (in rare cases) the losses by comparing the throughput of contiguous and non-contiguous channel bonding in three load regimes: low ($\bar{o}_B \leq 0.1$), medium ($0.1 < \bar{o}_B \leq 0.2$) and high ($\bar{o}_B > 0.2$), respectively.




Fig. 7: Throughput ratio of contiguous vs. non-contiguous. The bar chart inset depicts the mean aggregated ratio for low (L), medium (M), and high (H) occupancy regimes.

Fig. 7 shows the the throughput ratio of contiguous to non-contiguous channel bonding $\Gamma_{CO}(p)/\Gamma_{NC}(p)$, where $\Gamma_{\pi}(p)$ is the throughput achieved by policy π when selecting primary p in a given period. We plot⁹ the ratio for all possible primaries in $\mathcal{B}_{1\&2}$ and $\mathcal{B}_{2c}$.

The data reveals two remarkable phenomenon. First, contiguous outperforms non-contiguous in 2.5% of the cases (albeit with a modest throughput difference of 1.9%). But since non-contiguous is more flexible, how can it ever do worse? The answer is that the two policies result in different instants for transmission attempts. The contiguous policy occasionally (and quite randomly since the traces are the same) ends up with more favorable attempt instants. Nonetheless, in most cases, non-contiguous obtains higher throughput. For example, in many periods, at least one 20-MHz channel is idle during the whole period, which will always yield a gain for non-contiguous, but only sometimes yields a gain for contiguous bonding. In some cases, the difference can be quite high (e.g., a ratio of approximately 0.2 observed in low load). The origins of such extreme cases are the selection of the primary channel which we explore next. Second, the bar chart inside Fig. 7 reveals that both contiguous and non-contiguous channel bonding perform quite close on average for all occupancy regimes (low, medium, and high), and especially for the latter, as high load results in far fewer bonding opportunities overall.

Finding: Non-contiguous almost always outperforms contiguous channel bonding and their throughput differences are occasionally over a factor of 5. Nonetheless, their average throughputs are quite similar, which may ultimately favor contiguous channel bonding, since it is simpler to implement.

B. Primary Channel Selection

Selection of the primary channel is critical for channel bonding as it is where the backoff procedure is run. Namely,

⁹We use the mean occupancy in the x-axis only for plotting purposes. Notice that trace-driven estimations capture the whole spectrum evolution, including mean and any other statistics on occupancy.




Fig. 8: Normalized throughput of the best primary channel (*best*) and mean throughput for every primary channel (*mean*).

the transmitter must wait for a transmission opportunity on the primary channel and only then can explore adding channels. For contiguous channel bonding, there is an additional “edge effect”. For example, having the first channel as primary only allows bonding of higher numbered channels, whereas having a middle channel allows both higher and lower, provided they are contiguous in both cases. To explore this issue, we define the best-throughput $\Gamma^*(\pi)$ of policy π as the throughput achieved when selecting the best primary channel $p^* \in \mathcal{B}$, i.e., the primary channel that maximizes throughput in each period when implementing π . So, apart from assessing the throughput when considering all possible primary channels, we also evaluate an upper bound on performance by focusing only on the best primary channel.

Figure 8 shows the throughput gain with respect to the throughput achieved by single channel transmission, i.e., $\hat{\Gamma}_\pi = \Gamma_\pi / \Gamma_{SC}$, where Γ_{SC} and Γ_π are the throughput achieved by single channel transmission and channel bonding policy π , respectively. Such normalized throughput is presented in two ways for each period: the *mean* value $\hat{\Gamma}$ in the legend refers to the mean normalized throughput achieved by selecting each of the possible 8 primary channels per sub-band, and the *best* value $\hat{\Gamma}^*$, to the normalized throughput achieved by the primary p^* providing the highest throughput.

First, observe that both contiguous and non-contiguous channel bonding outperform single channel by a factor of at least 3 to over 6. Second, observe the region where contiguous channel bonding has the lowest gains. Strikingly, while selecting the best primary channel boosts throughput, these lower-gain cases cannot be entirely eliminated with better primary channel selection. Thus, in these cases, the available spectrum does not have a consistent structure and only non-contiguous bonding can exploit the gaps (and there are indeed many gaps as the load is low here). Next, observe the high variance of the throughput due to the spectrum activity distribution. In particular, throughput differences for similar occupancy values can be up to $1.9\times$ and $1.7\times$ for contiguous and non-contiguous, respectively. Lastly, despite these first two findings, selecting the best primary channel indeed provides substantial gains and can raise throughput by up to 68% and 64% for contiguous and non-contiguous, respectively. While non-contiguous channel bonding might seem impervious to primary channel selection as it can bond any channels, recall




Fig. 9: Distribution and empirical CDF of the best primary channel correlation ξ .

that even non-contiguous requires an idle primary channel to begin transmission.

Finding: Selection of the best primary channel improves average throughput by over 60%. But for contiguous channel bonding, it cannot overcome high variance and lower gain scenarios due to the channel occupancy dynamics of the spectrum band.

C. Inter-Channel Occupancy Correlation

When a transmitter finds its primary channel available for transmission and attempts to bond channels, it aims to find secondary channels available. Thus, correlation among channel occupancies can be expected to help channel bonding performance. Here we study the spectral correlation of the best primary channel with other channels and compute the mean correlation coefficient as

$$\xi = \sum_{p \in \mathcal{B}, p \neq p^*} \rho(p^*, p) / (|\mathcal{B}| - 1), \text{ with} \quad (2)$$

$$\rho(p^*, p) = \frac{\mathbb{E}[(p^* - \mu_{p^*})(p - \mu_p)]}{\sigma_{p^*} \sigma_p}$$

where $\mathcal{B} \in \{\mathcal{B}_{1\&2}, \mathcal{B}_{2c}\}$, thus $|\mathcal{B}| = 8$, $\rho(p^*, p)$ is the Pearson correlation coefficient of the temporal occupancy of channels p^* and p , as a function of the means, μ_{p^*} and μ_p , and standard deviations, σ_{p^*} and σ_p , of channel p^* and p , respectively. Notice that this assessment of spectral correlation by simultaneously measuring all channels is achieved for the first time by WACA.

Fig. 9 shows both the distribution of ξ vs. mean occupancy $\bar{o}_B$ and the empirical cumulative distribution function (CDF) of ξ . We observe that ξ ranges approximately from -0.10 to 0.60, so we define three correlation levels for ease of analysis: low ($\xi \leq 0.1$), medium ($0.2 \leq \xi < 0.4$), and high ($\xi \geq 0.5$). While we find that most of the periods (62%) fall inside the low correlation range, 17% of them present medium or higher correlation, indicating a significant amount of correlated epochs. A key origin of correlation among channels in the data set may be that the APs are already using channel bonding. Indeed, even though the dataset does not provide header information, some subsets of traces have multiple contiguous channels having almost identical busy/idle evolution. This suggests that channel bonding was employed during some portions of the measurement campaign.

Figure 10 shows throughput given selection of the best primary channel, with the sub-figures highlighting intervals




Fig. 10: Evolution of the best-primary normalized throughput vs. correlation. We use transparencies at the points to highlight denser regions.

with low, medium, and high correlation. We observe that most of the measured channel occupancies show minimal inter-channel correlation. However, the figures for periods of medium and high correlation indicate that with higher correlation among channel occupancy, the lower performance periods are increasingly avoided, especially for contiguous channel bonding. Interestingly, this effect holds regardless of the mean occupancy $\bar{o}_B$.

Finding: There are significant periods (17%) presenting medium to high correlation among channel occupancies. Such periods improve the performance of channel bonding regardless of the load. This can provide a hint as to what ideally would happen if multiple BSS's employ channel bonding, since channel bonding itself creates channel correlation.

D. A Markov Model for Channel Occupancy?

In §VI-C, we found that even though channels could be considered to be uncorrelated on average since activity is expected to be normally from separate BSS's, traces exhibited a significant amount of epochs of inter-channel occupancy correlation. With this experiment, we aim at pointing out the importance of capturing inter-channel correlation by comparing the original traces against two simple models. In particular, we first consider a Markov model that characterizes each channel as an independent two state (occupied, not occupied) Markov chain with exponential holding times for each state. We compute the mean transition rates (and hence mean occupied and not-occupied time) for each channel from the traces.




Fig. 11: Original vs. synthetic traces.

We then compare performance between the parameter-matched Markov model and the actual traces. As a baseline, we also consider a *uniform i.i.d.* model in which each temporal sample is occupied or not according to an i.i.d. Bernoulli distribution, with each channel again having the mean band occupancy matched from data. Our objective is not to develop a statistical occupancy model, but to study the extent to which simple models can or cannot characterize the behavior we observe.

We study below downlink throughput for these two models as compared to the traces when applying contiguous channel bonding. Since the throughput is impacted by the primary channel selection, we evaluate the selection of the 16 possible primary channels: 8 in sub-band $B_{1\&2}$ and 8 in B_{2c} .

Figure 11a shows the raw throughput (considering every possible primary) of the original traces and the synthetic model-generated occupancy. While both models capture the general trend of throughput decaying with increasing occupancy, the difference between datasets is evident in the form of the scatter plots.

Fig. 11b and Fig. 11c depict the mean relative error (MRE) of the models' throughput with respect to the original traces grouping by load and correlation regime, defined by the thresholds presented in §VI-A and §VI-C, respectively. We observe that the i.i.d. model is completely misleading with a mean error around 50% for high loads and/or high correlation. Unfortunately, although the Markov model is able to capture the mean on and off times for each channel, we observe that the mean error is still significant (up to 16% and 37% on average for high load and high correlated periods, respectively). Worse, particular periods assessed through the Markov model lead to large errors up to 33%, 58%, and 62% for low, medium and high occupancy regimes, respectively.

While we expected the higher relative errors to arise for high load periods, it is surprisingly critical the impact of not capturing correlation on both models. In fact, in contrast to the correlation distribution of the original traces shown in Fig. 9,

both models present a distribution similar to a delta function centered at $\xi = 0$. Given that the original mean occupancy is kept for all the periods generated by the models, results suggest that missing correlation is the main cause of outraging the models' accuracy.

Finding: unfortunately, simple channel occupancy models severely underestimate the available gains for channel bonding. Thus, modeling occupancy behavior by introducing inter-channel correlation remains an important avenue for research.

VII. WHITHER OTHER BSS'S

Thus far, we have evaluated how much throughput is gained from channel bonding. However, when one BSS gains access to multiple channels, nearby BSS's operating on other channels will have to defer or could yield hidden terminal collisions. We consider both cases as follows.

A. How Much Are Others Hindered?

When other BSS's defer to a channel bonder, it will add to their channel access latency and could decrease their throughput. Here, we study how much all other BSS's are hindered due to channel bonding and introduce a bandwidth deprivation metric ω as the aggregate bandwidth that would have been active (i.e., was being used by neighbor BSS's) during the transmissions of the channel bonding BSS. This metric can serve as an upper bound to how much others are hindered by viewing that all of these attempts by other BSS's would be lost vs. deferred. More formally, let

$$\omega = \frac{(\sum_{f=1}^{n_f} \sum_{t \in \mathcal{T}_f} \sum_{c \in \mathcal{B}_{tx,f}} x_{t,c}) b T_s}{T_{\text{per}}}, \quad (3)$$

where n_f is the number of frames transmitted, $\mathcal{T}_f$ is the set of consecutive temporal samples used for transmitting frame f , $\mathcal{B}_{tx,f}$ is the set of channels used for transmitting frame f , T_s is the duration of a temporal sample (10 μ s), $b = 20$ MHz is the bandwidth of a basic channel, and $T_{\text{per}} = 100$ ms is the observation duration. We normalize by $T_{\text{per}} = 100$ ms rather than the transmission time of each policy to provide a fair comparison. That is, ω represents how much raw bandwidth is deprived to surrounding BSSs in absolute terms [MHz], without considering how many transmissions are performed by each policy.

Fig. 12 shows the distribution (boxplot) of the bandwidth deprivation in MHz and Mbps for every primary channel using the occupancy categories defined in §VI-A. In order to estimate the corresponding data rate, we assume that all occupied samples are comprised of 20-MHz data frames transmitted also at MCS 9. The figure illustrates the intrinsic consequences of bonding channels without considering how the others might use the spectrum during the transmission. Namely, a channel bonder may find that most channels are free at backoff termination, thus subsequently occupying most (or all) of them. However, it may occur that right after starting the transmission, external activity appears in some of the channels being used, thus hindering surrounding BSS's.

Second, we observe that the differences in bandwidth deprivation for non-contiguous vs. contiguous channel bonding




Fig. 12: Distribution of bandwidth and data rate lost for other BSS's for every primary channel.

are quite modest, even when non-contiguous channel bonding achieves higher throughput. The key reason is that most high occupancy periods concentrate activity in some 20 MHz channels, often leaving at least one 20 MHz channel idle during nearly the entire period. In contrast to contiguous, non-contiguous channel bonding can always realize small throughput gains by bonding those idle channels regardless of the primary channel allocation. Also, despite the low average impact, there are multiple outliers in which the bandwidth lost by other BSS's is significantly higher than average, thus manifesting the worst-case external effect.

Finally, we can ask if channel bonding is simply a zero-sum-game: do the gains of a bonding BSS simply subtract from other BSS's throughput? To address this question, we assess whether the average throughput gained by channel bonding exceeds the worst-case prevented transmissions by other BSS's, using single channel transmission as a baseline. Thus, we define the ratio

$$\kappa_\pi = \frac{\bar{\Gamma}_\pi - \bar{\Gamma}_{\text{SC}}}{\bar{\omega}_\pi - \bar{\omega}_{\text{SC}}}, \quad (4)$$

where $\bar{\Gamma}_\pi$ and $\bar{\Gamma}_{\text{SC}}$ are the mean throughput of policy π and SC, and $\bar{\omega}_\pi$ and $\bar{\omega}_{\text{SC}}$ are the corresponding mean throughput deprivation. A value $\kappa_\pi = 1$ would indicate that the trade-off between the bonding gain and external throughput deprivation is actually a zero-sum-game, whereas $\kappa_\pi > 1$ would indicate net gains. Results indicate that both contiguous and non-contiguous channel bonding yield substantial net gains over all occupancy regimes. Specifically, $\kappa_{\text{CO}} = 95.2, 68.9, 49.4$ for contiguous, and $\kappa_{\text{NC}} = 91.9, 65.7, 45.5$ for non-contiguous, under low, medium, and high occupancy regimes, respectively. In this regard, we observe a potential trend to a zero-sum game as mean occupancy increases to extreme values.

Finding: In the worst case, a channel bonding BSS can hinder neighbors by 100's of Mb/sec. However, the *average* hindrance is quite low and neighbors can defer instead of not transmitting to reduce the impact. Moreover, the data rate gained by bonding far exceeds the worst-case data rate lost by neighbors – enabling the channel bonder to gain at least 45 times what the other BSS's lose – since neighbors are not always backlogged during bonding epochs.

B. The Hidden Cost of Hidden Nodes

Here, we continue our study of potential detrimental effects of channel bonding by considering the case of hidden nodes.




Fig. 13: Distribution of the norm. best-primary-throughput.

We again impose a worst-case scenario on the measurements as follows: channel access occurs as previously with contention occurring on the primary channel and channel bonding adding channels according to availability and the policy, contiguous or non-contiguous. However, here we reconsider the measured activity that occurs during the bonded transmission. While in the previous subsection, we considered that the other BSS's would have deferred, here, we consider that other BSS's are hidden from the bonded transmission and still transmit and result in frame loss. Thus, this scenario considers that the RTS/CTS mechanism is not used or fails and disallows capture effects in which the higher signal-to-noise ratio receiver correctly decodes, e.g., [44]. Thus, in these two ways, this experiment provides an empirical upper bound as to the damage that hidden terminals could do to a channel bonded transmission. In particular, we assume that a frame f is lost if temporal samples are active in any of the channels used for transmitting f . Specifically, f is lost if the following expression holds,

$$\sum_{t \in \mathcal{T}_f} \left(\left(\sum_{c \in \mathcal{B}_{\text{tx},f}} x_{t,c} \right) \geq 1 : 1, 0 \right) \geq \alpha |\mathcal{T}_f|, \quad (5)$$

where $\mathcal{T}_f$ is the set of consecutive temporal samples required for transmitting frame f , and the term $\sum_{c \in \mathcal{B}_{\text{tx},f}} x_{t,c}$ denotes the number of occupied channels at time t . We apply a factor $\alpha = 0.01$ to exclude spurious noise as cause of loss. Thus, we avoid for instance losing an entire frame when just a single temporal sample was found active during the transmission.

Fig. 13a and Fig. 13b depict the throughput distribution under hinderling (without hidden terminals) and hidden-terminal scenarios, respectively. First, observe that throughput is significantly reduced in comparison to the case with no hidden terminals. For example, for contiguous channel bonding in high load, mean throughput gains over single-channel transmission are reduced from $5.2 \times$ to $2.6 \times$ (50%).

Second, there are significant periods where channel bonding results in extremely low performance, even for low occupancy values. In fact, there are scenarios in which single-channel access outperforms bonding (blue areas where $\hat{\Gamma}^* < 1$ in the figure). In particular, single channel outperforms contiguous channel bonding in 18%, 8% and 8% of the cases for low, medium and high occupancy, respectively. Even worse, single channel outperforms non-contiguous channel bonding in 38%, 14%, and 15% of cases respectively. This illustrates that external traffic patterns lead to accentuated chances of collisions.

Third, due to its more aggressive nature, non-contiguous channel bonding performs slightly worse than contiguous both in terms of average throughput and how often it is outperformed by single channel transmission. This, together with the fact that contiguous performs similarly on average to non-contiguous without hidden terminals, indicates that gains provided by non-contiguous can be relatively small considering its greater risk and more complex design.

Finally, also in contrast to Fig. 13a, Fig. 13b indicates that the decaying trend of mean normalized best-throughput with load when neighbors defer does not hold with hidden terminals. Indeed, we observe a slight increase of the normalized best-throughput as the occupancy increases as fewer bonding opportunities also diminish the risk of collisions.

Finding: While channel bonding outperforms single-channel transmission in most cases, it is vulnerable to hidden node interference, with non-contiguous channel bonding being impacted the most. To avoid the aforementioned worst-case throughput penalties, nodes can use RTS/CTS, implement capture (to better receive even while a hidden terminal transmission overlaps in time), or increase channel sensing sensitivity, i.e., reduce the threshold for sensing channel activity to reduce the frequency of hidden terminal transmissions.

VIII. IEEE 802.11AX CONSTRAINTS

Here, we consider additional constraints imposed by 802.11ax on channel bonding. Namely, 802.11ax defines three modes of channel bonding, each of which restrict channel bonding to a subset of channels [1]. For example, U-NNI-1 and U-NNI-2 restrict channel bonding groups to $\mathcal{C}_{\mathcal{B}_{1\&2}} = \{\{1, 2\}, \{3, 4\}, \{5, 6\}, \{7, 8\}, \{1, 2, 3, 4\}, \{5, 6, 7, 8\}, \{1, \dots, 8\}\}$ so that (for example) a 40 MHz bonded channel that combines channels 2 and 3 is not allowed. Notice that we do not use the IEEE 802.11 channel indices (e.g., 36, 40, 44,...) in the example above for ease of exposition. While reducing throughput when only a single BSS bonds, such restrictions can reduce fragmentation if multiple BSS's bond. Moreover, such restrictions simplify implementation. Thus, here we study the cost of 802.11ax's channelization restrictions when a single BSS bonds and compare to the non-restricted cases explored this far.

Static Channel Bonding (SCB): employs a *fixed* channel width and therefore must always use the same set of channels. It can be viewed as an “all-or-nothing” policy: for example, if static channel bonding uses 8 consecutive basic channels (160 MHz), it must wait until all 8 channels are free in order to transmit using the entire 160 MHz. Although simple to




Fig. 14: Best primary throughput ratio of the 11ax policies w.r.t. contiguous and non-contiguous channel bonding.

implement, it is vulnerable to excessive deferral while waiting for all channels to become available.

Dynamic Channel Bonding (DCB): picks the largest available set of contiguous basic channels found free subject to both the primary channel being included and the total set of channels being an allowed set under the restrictions such as above. For example, if channel 4 is selected as a primary channel, dynamic channel bonding allows bonded transmissions with channels $\{3, 4\}$, $\{1, 2, 3, 4\}$, or $\{1, \dots, 8\}$, as well as single channel transmission on channel 4. Other combinations such as $\{4, 5\}$ are not permitted.

Preamble Puncturing (PP): allows puncturing wide channels, i.e., aggregating non-contiguous basic channels. Even though the total channel is still subjected to the above channelization constraints, preamble puncturing is more flexible than dynamic channel bonding since it enables non-contiguous channel combinations, provided the primary channel is included [1], [45]. For example, a transmission using primary channel 3 could transmit on 60 MHz using channels $\{1, 3, 4\}$, which would be a punctured subset of $\{1, 2, 3, 4\}$. Nonetheless, not every combination of punctured channels is allowed in the standard. For instance, $\{1, 3\}$ is not permitted. Observe that both static and dynamic channel bonding are restricted special cases of contiguous channel bonding studied previously, whereas preamble puncturing is a restricted special case of non-contiguous channel bonding.

Here, we compare the throughput of both the (non-restricted) contiguous and non-contiguous policies with the three 802.11ax policies, all using selection of the best primary channel. Fig. 14 shows the resulting 802.11ax throughput normalized to the unrestricted case for the respective category of contiguous or non-contiguous.

First, observe that 802.11ax static channel bonding suffers severe throughput degradation compared to contiguous channel bonding. Namely, the desirable simplicity of static channel bonding comes with a high performance cost. For example, under heavy load, static channel bonding only obtains 39% of the throughput compared to contiguous channel bonding, indicating that too much time is spent deferring to realize significant gains from wide channels. Moreover, static channel bonding also encounters poorly performing outliers, well below the mean. For example, the 1%-percentile performance is less than 10% for medium and high loads, even less than single channel allocation.

Second, comparing 802.11ax dynamic channel bonding to contiguous channel bonding (middle bars), we observe that it

achieves 87% to 81% of the average throughput of contiguous channel bonding for low and high load respectively, with 1%-percentiles over 65% for any load regime.

Finally, comparing 802.11ax preamble puncturing to non-contiguous channel bonding (right set of bars), the former obtains 96% to 94% of the throughput of the latter in low and high load respectively. Further, throughput differences are kept smaller also in the most divergent scenarios, reaching 1%-percentiles over 77% for all the loads.

Finding: IEEE 802.11ax channelization restrictions have a modest throughput cost compared to non-restricted. For example, standard-compliant preamble puncturing obtains over 90% of the throughput of (non-restricted) non-contiguous channel bonding. Unfortunately, the simplest standard policy, static channel bonding, performs quite poorly.

IX. A SOLD OUT STADIUM

Finally, we assess if channel bonding can provide throughput gains in areas with persistently high load over all channels. To study such a scenario, we deployed WACA in the press box of the Camp Nou stadium during a football game with over 98,000 spectators present. Measurements were taken on August 4, 2019, from 17:24 to 22:30 accounting for a total duration of 5 hours and 6 minutes. On that date, the Joan Gamper trophy was held, which pitted the local club (Futbol Club Barcelona) against the visiting club (Arsenal Football Club). Fig. 15a shows a photograph of the setup.

We also obtained data from the stadium's network management team which indicated that up to 12,000 Wi-Fi clients were simultaneously connected. Downlink and uplink traffic (provided by Camp Nou's management) is depicted in Fig. 15b and the spectrogram in Fig. 15c. Each slot in the spectrogram represents the occupancy of each channel averaged in 1-second




Fig. 15: Wi-Fi activity at the Camp Nou stadium.




Fig. 16: Gains for the bonding BSS assuming others defer.

periods. We observe that most channels were highly occupied during the measurements. Moreover, we can observe the users' behavior induced nonstationarity of the traces as the match progressed. Namely, while activity is always high, there is a notable reduction during play (first and second half) compared to activity before, between halves, and right after the game time. We also observe that the majority of the channels are crowded most of the time with periods reaching mean band occupancy values $\bar{o}_B$ rising up to 99%. In fact, 22% of the periods have $\bar{o}_B > 0.8$. Nonetheless, we pose the question as to whether the remaining occupancy gaps can be exploited by channel bonding to boost aggregate throughput.

Fig. 16 shows normalized throughput under best primary channel selection when applying contiguous and non-contiguous channel bonding. Under loads exceeding 90%, both policies have throughput close to single channel transmission, indicating that there is simply minimal remaining margin for gain. Nonetheless, in transient epochs where load is low, both policies exploit throughput gains. For example, in the stadium during epochs of 20-30% load, the gain over single channel throughput is $3.69\times$ and $5.30\times$ for contiguous and non-contiguous respectively. In fact, we observe that non-contiguous significantly outperforms contiguous in a wide regime of “non-extreme” loads given the heterogeneous activity distribution over channels. For example, at 40-50% average occupancy, contiguous channel bonding has average throughput gain of $2.38\times$ whereas non-contiguous has $3.97\times$, a gain of 67%. Recall that for the non-stadium experiments in Figure 13a, the gains during epochs of 20-30% load (considered inside “high” load for the urban scenarios) were $5.19\times$ and $5.34\times$ for contiguous and non-contiguous bonding. Comparing to the stadium, non-contiguous channel bonding performs remarkably similarly (5.30 in stadium vs. 5.34 otherwise). In contrast, contiguous channel bonding performs worse in the stadium (3.69 in stadium vs. 5.19 otherwise) even under the same average load. The reason may be the minimal inter-channel correlation at the stadium. Indeed, the stadium management confirmed that no channel bonding was implemented in any of the stadium's APs.

Lastly, as shown in §VII-B, the story can turn quite negative if competing BSS's do not defer to the channel bonder, i.e., if other BSS's are hidden. In such a worst case, all bonded transmissions with subsequent activity are considered lost. Figure 17 shows the results of this scenario for the stadium traces. As indicated by the blue shadowed areas, most of the periods




Fig. 17: Hidden terminal scenario with the ratio of epochs that single channel outperforms bonding shown by the blue curve.

achieve higher throughput with single channel transmission rather than using channel bonding. Indeed, single channel transmission outperforms both channel bonding policies in at least 80% of the epochs over all occupancy regimes. Moreover, comparing similar loads to the other data sets (e.g., 20-30% occupancy), the stadium performs far worse with hidden terminals. For instance, comparing Figure 17 and Figure 13b, we find that for 20-30% average occupancy, even with worst-case hidden terminals, the average throughput gain of contiguous channel bonding was $2.66\times$, whereas for the stadium, the gain is reduced to $0.59\times$, i.e., throughput is lost. This indicates that the stadium traffic and channel occupancies yield greater hidden terminal risk than urban environments, even for the same traffic load. Thus, similarly to the hindering case, the results suggest that negligible correlation is detrimental in the presence of hidden terminals.

Finding: Even under the stadium's extremely high average load, short durations of lower load can be exploited to yield significant throughput gains. However, the risk is high as hidden terminals could drive the throughput to levels even worse than without any channel bonding.

X. CONCLUSION

In this paper, we introduce WACA, an all-channel Wi-Fi spectrum analyzer for simultaneous measurement of all 24 20 MHz channels that allow channel bonding at 5 GHz. We present extensive measurement campaigns covering two continents, diverse areas, and many hours of signal strength samples. We use the measurements to explore key aspects of channel bonding spanning from statistics of the channel occupancies to channel bonding standards and policies.


ACKNOWLEDGMENT

The authors appreciate the generosity of Futbol Club Barcelona for letting them take measurements in the Camp Nou stadium. We would also like to thank Chia-Yi Yeh, Albert Bel, Álvaro López, and Marc Carrascosa for their contribution in taking the measurements.

REFERENCES

- [1] “IEEE P802.11ax/D4.0 High Efficiency WLAN,” 2019.
- [2] “IEEE 802.11be Extremely High Throughput, 802.11-18/1231r6,” 2019.
- [3] S. Barrachina-Muñoz, B. Bellalta, and E. Knightly, “Wi-Fi All-Channel Analyzer,” in *14th International WiTECH'20.*, ACM, 2020.

- [4] S. Barrachina-Muñoz, B. Bellalta, and E. Knightly, "WACA: Wi-Fi All-Channel Analyzer." https://github.com/sergiobarra/WACA_WiFiAnalyzer, 2021.
- [5] M. Rademacher, K. Jonas, and M. Kretschmer, "Quantifying the spectrum occupancy in an outdoor 5 GHz Wi-Fi network with directional antennas," in *2018 IEEE WCNC*, pp. 1–6, IEEE, 2018.
- [6] T. Taher, R. Attard, A. Riaz, D. Roberson, J. Taylor, K. Zdunek, J. Hallio, R. Ekman, J. Paavola, J. Suutala, et al., "Global spectrum observatory network setup and initial findings," in *2014 9th International CROWNCOM*, pp. 79–88, IEEE, 2014.
- [7] S. Subramaniam, H. Reyes, and N. Kaabouch, "Spectrum occupancy measurement: An autocorrelation based scanning technique using USRP," in *2015 IEEE 16th Annual WAMICON*, pp. 1–5, IEEE, 2015.
- [8] M. Wellens, J. Riihijärvi, and P. MähöNen, "Empirical time and frequency domain models of spectrum use," *Physical Communication*, vol. 2, no. 1-2, pp. 10–32, 2009.
- [9] V. Kone, L. Yang, X. Yang, B. Y. Zhao, and H. Zheng, "The effectiveness of opportunistic spectrum access: A measurement study," *IEEE/ACM Transactions on Networking*, vol. 20, no. 6, pp. 2005–2016, 2012.
- [10] S. W. Day and A. Wagstaff, "Activity in the WiFi bands - An Objective Approach to Estimation," in *ARMMS Conference*, Mass, 2014.
- [11] F. Salahdine and H. El Ghazi, "A real time spectrum scanning technique based on compressive sensing for cognitive radio networks," in *2017 IEEE 8th Annual Ubiquitous Computing, Electronics and Mobile Communication Conference (UEMCON)*, pp. 506–511, IEEE, 2017.
- [12] S. Biswas, J. Bicket, E. Wong, R. Musaloiu-e, A. Bhartiya, and D. Aguayo, "Large-scale measurements of wireless network behavior," in *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, pp. 153–165, 2015.
- [13] M. López-Benítez and F. Casadevall, "Modeling and simulation of time-correlation properties of spectrum use in cognitive radio," in *2011 6th International ICST CROWNCOM*, pp. 326–330, 2011.
- [14] V. Kone, L. Yang, X. Yang, B. Y. Zhao, and H. Zheng, "On the feasibility of effective opportunistic spectrum access," in *Proceedings of the 10th ACM SIGCOMM*, pp. 151–164, 2010.
- [15] M. Höyhty, A. Mämmelä, M. Eskola, M. Matinmikko, J. Kalliovaara, J. Ojaniemi, J. Suutala, R. Ekman, R. Bacchus, and D. Roberson, "Spectrum occupancy measurements: A survey and use of interference maps," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 4, pp. 2386–2414, 2016.
- [16] S. A. Hanna and J. Sydor, "Distributed sensing of spectrum occupancy and interference in outdoor 2.4 GHz Wi-Fi networks," in *Global Communications Conference (GLOBECOM)*, pp. 1453–1459, IEEE, 2012.
- [17] S. A. Hanna, "A 3-state hypothesis test model for cognitive radio systems," in *2014 IEEE International Symposium on Dynamic Spectrum Access Networks (DYSpan)*, pp. 291–302, IEEE, 2014.
- [18] S. Yin, D. Chen, Q. Zhang, M. Liu, and S. Li, "Mining spectrum usage data: a large-scale spectrum measurement study," *IEEE Transactions on Mobile Computing*, vol. 11, no. 6, pp. 1033–1046, 2012.
- [19] L. Deek, E. Garcia-Villegas, E. Belding, S.-J. Lee, and K. Almeroth, "The impact of channel bonding on 802.11n network management," in *Proceedings of the Seventh Conference on emerging Networking Experiments and Technologies*, pp. 1–12, 2011.
- [20] L. Deek, E. Garcia-Villegas, E. Belding, S.-J. Lee, and K. Almeroth, "Intelligent channel bonding in 802.11n WLANs," *IEEE Transactions on Mobile Computing*, vol. 13, no. 6, pp. 1242–1255, 2014.
- [21] M. Y. Arslan, K. Pelechrinis, I. Broustis, S. V. Krishnamurthy, S. Addepalli, and K. Papagiannaki, "Auto-configuration of 802.11n WLANs," in *Proceedings of the 6th International Conference*, pp. 1–12, 2010.
- [22] M. Y. Arslan, K. Pelechrinis, I. Broustis, S. Singh, S. V. Krishnamurthy, S. Addepalli, and K. Papagiannaki, "ACORN: An auto-configuration framework for 802.11n WLANs," *IEEE/ACM Transactions on Networking (TON)*, vol. 21, no. 3, pp. 896–909, 2013.
- [23] Y. Zeng, P. H. Pathak, and P. Mohapatra, "A first look at 802.11ac in action: Energy efficiency and interference characterization," in *2014 IFIP Networking Conference*, pp. 1–9, IEEE, 2014.
- [24] M. Park, "IEEE 802.11ac: Dynamic bandwidth channel access," in *Communications (ICC)*, pp. 1–5, IEEE, 2011.
- [25] M. X. Gong, B. Hart, L. Xia, and R. Want, "Channel bonding and mac protection mechanisms for 802.11ac," in *2011 IEEE Global Telecommunications Conference-GLOBECOM 2011*, pp. 1–5, IEEE, 2011.
- [26] B. Bellalta, A. Faridi, J. Barcelo, A. Checco, and P. Chatzimisios, "Channel bonding in short-range WLANs," in *European Wireless 2014; 20th European Wireless Conference*, pp. 1–7, VDE, 2014.
- [27] B. Bellalta, A. Checco, A. Zocca, and J. Barcelo, "On the interactions between multiple overlapping w lans using channel bonding," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 2, pp. 796–812, 2015.
- [28] A. Faridi, B. Bellalta, and A. Checco, "Analysis of dynamic channel bonding in dense networks of w lans," *IEEE Transactions on Mobile Computing*, vol. 16, no. 8, pp. 2118–2131, 2016.
- [29] M.-S. Kim, T. Ropitault, S. Lee, and N. Golmie, "A throughput study for channel bonding in ieee 802.11ac networks," *IEEE Communications Letters*, vol. 21, no. 12, pp. 2682–2685, 2017.
- [30] M. Yazid and A. Ksentini, "Stochastic modeling of the static and dynamic multichannel access methods enabling 40/80/160 MHz channel bonding in the VHT WLANs," *IEEE Communications Letters*, vol. 23, no. 8, pp. 1437–1440, 2019.
- [31] S. Barrachina-Muñoz, F. Wilhelmi, and B. Bellalta, "To overlap or not to overlap: Enabling channel bonding in high-density WLANs," *Computer Networks*, vol. 152, pp. 40–53, 2019.
- [32] S. Barrachina-Muñoz, F. Wilhelmi, and B. Bellalta, "Dynamic channel bonding in spatially distributed high-density WLANs," *IEEE Transactions on Mobile Computing*, vol. 19, no. 4, pp. 821–835, 2020.
- [33] S. Khairy, M. Han, L. X. Cai, Y. Cheng, and Z. Han, "A renewal theory based analytical model for multi-channel random access in IEEE 802.11ac/ax," *IEEE Transactions on Mobile Computing*, 2018.
- [34] A. Zakrzeska and L. Ho, "Dynamic channel bandwidth use through efficient channel assignment in ieee 802.11ac networks," in *2019 IEEE 90th Vehicular Technology Conference*, pp. 1–6, IEEE, 2019.
- [35] S. Barrachina-Muñoz, F. Wilhelmi, and B. Bellalta, "Online primary channel selection for dynamic channel Bonding in high-density WLANs," *IEEE Wireless Communications Letters*, vol. 9, no. 2, 2020.
- [36] Z. Khan and J. J. Lehtomäki, "Interactive trial and error learning method for distributed channel bonding: model, prototype implementation, and evaluation," *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 2, pp. 206–223, 2019.
- [37] S.-s. Lee, T. Kim, S. Lee, K. Kim, Y. H. Kim, and N. Golmie, "Dynamic channel bonding algorithm for densely deployed 802.11ac networks," *IEEE Transactions on Communications*, vol. 67, no. 12, 2019.
- [38] M. Han, S. Khairy, L. X. Cai, Y. Cheng, and F. Hou, "Capacity analysis of opportunistic channel bonding over multi-channel WLANs under unsaturated traffic," *IEEE Transactions on Communications*, 2019.
- [39] G. Buttazzo, C. Di Franco, and M. Marinoni, "Design and analysis of target-sensitive real-time systems," *Software: Practice and Experience*, vol. 46, no. 9, pp. 1181–1200, 2016.
- [40] L. Farzinash and M. Kargahi, "A scheduling algorithm for execution-instant sensitive real-time systems," in *2009 15th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications*, pp. 511–518, IEEE, 2009.
- [41] N. Anand, E. Aryafar, and E. W. Knightly, "WARPlab: a flexible framework for rapid physical layer design," in *Proceedings of the 2010 ACM workshop on Wireless of the students, by the students, for the students*, pp. 53–56, ACM, 2010.
- [42] "Rice University Wireless Open Access Research Platform." <http://warpproject.org>, last accessed January 30, 2020.
- [43] S. Jang, K. G. Shin, and S. Bahk, "Post-CCA and reinforcement learning based bandwidth adaptation in 802.11ac networks," *IEEE Transactions on Mobile Computing*, vol. 17, no. 2, pp. 419–432, 2018.
- [44] J. Lee, W. Kim, S.-J. Lee, D. Jo, J. Ryu, T. Kwon, and Y. Choi, "An experimental study on the capture effect in 802.11a networks," in *ACM international WINTECH*, pp. 19–26, 2007.
- [45] Q. Qu, B. Li, M. Yang, Z. Yan, A. Yang, D.-J. Deng, and K.-C. Chen, "Survey and performance evaluation of the upcoming next generation WLAN standard - IEEE 802.11ax," *Mobile Networks and Applications*, vol. 24, no. 5, pp. 1461–1474, 2019.



Sergio Barrachina-Muñoz holds a Ph.D. in Information and Communication Technologies (2021) by Universitat Pompeu Fabra (UPF). His thesis research focused on developing autonomous learning techniques for improving next-generation Wi-Fi networks through efficient spectrum access. Sergio is currently working as a postdoctoral researcher in the SMARTECH department at Centre Tecnològic de Telecomunicacions de Catalunya (CTTC).



Boris Bellalta is an Associate Professor in the Department of Information and Communication Technologies (DTIC) at Universitat Pompeu Fabra (UPF). He obtained his Ph.D. in Information and Communication Technologies from UPF in 2007. His research interests are in the area of wireless networks, with emphasis on the design and performance evaluation of new architectures and protocols.



Edward W. Knightly (S'92–M'96–SM'04–F'09) received the B.S. degree from Auburn University and the M.S. and Ph.D. degrees from the University of California at Berkeley. He is currently the Sheafor-Lindsay Professor and the Chair of the Department of Electrical and Computer Engineering, Rice University. His research interests are in experimental wireless networking, urban-scale testbeds, 60-GHz, terahertz, and VLC bands, wireless security, and performance evaluation. He is a Sloan Fellow.