Breaking Through the Full-Duplex Wi-Fi Capacity Gain

Saulo Queiroz*†
*Academic Dept. of Informatics
Federal University of Technology (UTFPR)
Ponta Grossa–PR, Brazil.
Email: sauloqueiroz@utfpr.edu.br

João Vilela
CISUC, Dept. of Informatics Engineering
University of Coimbra
Coimbra, Portugal
Email: jpvilela@dei.uc.pt

Roberto Hexsel
†Dept. of Informatics
Federal University of Paraná (UFPR)
Curitiba–PR, Brazil
Email: roberto@inf.ufpr.br

***Abstract*—In this work we identify a seminal design guideline that prevents current Full-Duplex (FD) MAC protocols to scale the FD capacity gain (i.e. 2× the half-duplex throughput) in single-cell Wi-Fi networks. Under such guideline (referred to as 1:1), a MAC protocol attempts to initiate up to two simultaneous transmissions in the FD bandwidth. Since in single-cell Wi-Fi networks MAC performance is bounded by the PHY layer capacity, this implies gains strictly less than 2× over half-duplex at the MAC layer. To face this limitation, we argue for the 1:*N* design guideline. Under 1:*N*, FD MAC protocols 'see' the FD bandwidth through *N*>1 orthogonal narrow-channel PHY layers. Based on theoretical results and software defined radio experiments, we show the 1:*N* design can leverage the Wi-Fi capacity gain more than 2× *at* and *below* the MAC layer. This translates the denser modulation scheme incurred by channel narrowing and the increase in the spatial reuse factor enabled by channel orthogonality. With these results, we believe our design guideline can inspire a new generation of Wi-Fi MAC protocols that fully embody and scale the FD capacity gain.**

***Keywords*-Full-Duplex Wireless, MAC Protocol Design, IEEE 802.11 WLANs, Performance Evaluation.**

I. Introduction

Recent works have demonstrated the feasibility of Self-Interference Cancellation (SIC) techniques, turning Full-Duplex (FD) radios into a reality e.g. [1]. Such radios are capable of receiving and transmitting simultaneously within the same frequency band, achieving a gain of 2× the half-duplex link capacity in theory (i.e. the FD gain). An important question raised by that achievement is whether it is possible to design a Medium Access Control (MAC) protocol that accomplishes the goal of *scaling* the FD gain in a single cell Wireless Local Area Network (WLAN). By surveying the MAC literature e.g. [2], [3], [4], one can find out it is hard to accomplish that scalability goal, since the contention overheads and the lack of spatial reuse can shrink the FD gain to 1.58× [5].

To tackle the limitation of current FD MAC protocols, we go a step further and identify a common design strategy we refer to as the 1:1 MAC design guideline. With the 1:1 design, an FD MAC protocol 'sees' the whole FD bandwidth through a *single* PHY layer. To maximize FD gains with such design, MAC protocols attempt to minimize the difference between the start time of two concurrent transmissions in the channel. This leads to gains bounded by the capacity of two nodes freely transmitting to each other in the channel. In fact, in a single-cell WLAN, the MAC throughput is bounded by the PHY layer capacity.Thus, doubling such capacity with FD radios may limit the maximum capacity gain achieved at the FD MAC layer to a value strictly less than 2× the half-duplex throughput. This suggests one needs to improve the capacity below the MAC layer more than 2× to give room for MAC protocols that actually approaches the FD gain.

In this paper we report novel results that break through the capacity gain leveraged by FD radios in single-cell WLANs. We accomplish this by arguing for an alternative FD MAC design guideline we refer to as 1:*N*. Under that, the MAC layer arranges the FD bandwidth into *N*>1 PHY layers. Each PHY layer is assigned to a portion of spectrum that is narrower than the available FD bandwidth and orthogonal to the other PHY spectrum portions. Similar design have been studied before from the perspective of MAC and/or radio architecture e.g. [6], [7], [8]. These works highlight the advantages of parallel narrow channels on a single radio but under the half-duplex constraint. To fully realize the FD gain over a wireless bandwidth allocated to concurrent narrow channels, one has to refer to the same kind of wide-band SIC design (e.g. [1]) assumed by current state-of-the-art 1:1 FD MAC proposals. We refer to such advance to report *unprecedented contributions towards the FD gain scalability in WLANs*.

Our first contribution is to show that, contrary to the popular assumptions and beliefs, it is possible to attain more-than-doubled capacity gains within an FD bandwidth i.e. *below* the MAC layer. Indeed, narrowing a channel relaxes receive sensitivity requirements enabling denser modulation schemes [9, Table 18–14]. Thus, spectrum usage improves. For instance, instead of occupying a 10 MHz channel with two (FD) transmissions, one can split it into two 5 MHz orthogonal FD channels and activate four concurrent transmissions. This can yield gains of ≈2.2× over a 10 MHz half-duplex link even considering guard-bands. We demonstrate this theoretically and through a proof-of-concept study with USRP platforms.

Our second contribution is to scale the novel FD gain *at* the MAC layer. We consider the ideal condition for an

This work was partially funded through project PTDC/EEI-TEL/3684/2014 (SWING2 - Securing Wireless Networks with Coding and Jamming), Project co-funded by COMPETE 2020, Portugal 2020 - Programa Operacional Competitividade e Internacionalização (POCI), European Union through Fundo Europeu de Desenvolvimento Regional (FEDER) and Fundação para a Ciência e Tecnologia (FCT).

1:1 FD Wi-Fi MAC protocol [10] and show its performance more than doubles under 1:*N*. This happens because channel orthogonality multiplies FD opportunities by increasing the spatial reuse factor. We believe these results instigate further research towards a solid FD IEEE 802.11 stack.

The remainder of this paper is organized as follows. In section II we present our system model and background. In sections III-A and IV we present the 1:*N* design guideline and its capacity model, respectively. In section V we present our results. In section VI we present conclusion and future work.

II. System Model and Background

We consider the design directives that a Wi-Fi compliant FD MAC protocol should follow to scale the FD gain. In this sense we focus on models to assess capacity upper-bounds *at* and *below* the MAC layer in a single-cell infrastructure IEEE 802.11 WLAN. For the MAC protocol study, the cell is composed of one Access Point (AP) and n STAtions (STA). STAs perform the standard CSMA/CA to initiate a transmission to the AP (uplink). The AP is assumed to always have a frame enqueued to its current transmitting STA. Then, the AP can establish an FD (down)link to the STA upon processing its incoming header. As we discuss in section IV, this corresponds to an ideal condition the capacity upper-bound of an FD Wi-Fi MAC protocol can be derived from.

For each MAC proposal we assume saturated traffic and ideal channel conditions [11]. These assumptions ensure we assess 'the most each MAC protocol can do' when provided with best conditions. Note, however, any MAC protocol under the design guideline we are about to present might actually perform better in noisy environments. This happens because the narrow Wi-Fi channels we rely on are less prone to noise, as we discuss in the section III-C. Also, we assume each compared MAC and PHY model suffers from the same level of negligible self-interference residue. Again, a successful (de)modulation process might be less demanding in terms of SIC requirements if performed over narrower channels instead of wide channels [1].

A. FD MAC WLAN Terminology

The ultimate goal of any FD MAC protocol is to take advantage of FD opportunities within a given wireless channel to maximize capacity. It means the protocol attempts to activate two overlapping transmissions to maximize channel utilization and so, throughput. In Wi-Fi compliant WLANs, the *Primary Transmitter* (PT) is the first node to start transmitting a data frame after winning a typical CSMA/CA contention round. The node PT transmits to is called *Primary Receiver* (PR). During the primary transmission, the FD MAC protocol may start a secondary transmission in the channel. In this case the sender and receiver are called *Secondary Transmitter* (ST) and *Secondary Receiver* (SR), respectively.

Basically, the FD opportunities can be classified into either *symmetric* or *asymmetric* dual-links [4]. In symmetric dual-links, PT and PR coincide with SR and ST, respectively (i.e. [PT=SR]⇄[PR=ST], where the direction of each arrow denotes the destination of a transmission). In asymmetric dual-links, there must be a *third* node involved in the secondary communication. Such node is either a SR or a ST. In the former case, the PR coincides with the ST i.e. PT→[PR=ST]→*SR*. Otherwise the PT coincides with SR, i.e. PR←[PT=SR]←*ST*. Note the two possible asymmetric dual-links are not different views of the same scenario since in one case an already *receiving* node starts transmitting while in the other an already *transmitting* node starts receiving.

B. Medium Access Control Challenges with Dual-links

The performance of an FD MAC protocol results from a balance between how effectively it exploits dual-links and the cost it takes towards that. Concerning asymmetric dual-links, the main challenge consists in assuring the secondary transmission does not collide with some possible ongoing primary transmission. Collisions may happen whenever the receiver node of a primary (secondary) transmission is within the interference range of a secondary (primary) transmission. In case of symmetric dual-links, the challenge consists in identifying a pair of nodes that have frames to each other. To maximize FD gains regardless of the type of dual-link, *any FD Wi-Fi MAC protocol attempts to minimize* $\Delta_t = t_{st} - t_{pt} \geq 0$. Particularly for our scenario, t_{pt} is the time at which a STA starts a primary transmission after winning a CSMA/CA contention round and t_{st} is the time at which the AP starts the corresponding secondary transmission.

III. The 1:*N* MAC Design Guideline

A. Novel Classification for FD MAC Protocols

In this work we identify a new category for the design of FD MAC protocols. With this novel category, MAC protocols are classified according to the way they exploit the available wireless FD bandwidth. In this sense, we identify a seminal trend we refer to as the 1:1 MAC design guideline [2], [3], [4]. Under the 1:1 guideline the MAC protocol 'sees' the FD bandwidth through a *single* PHY layer. Thus, the best-case of any 1:1 MAC protocol is bounded to the capacity of a dual-link before the contention overheads.

A reasonable way to overcome the performance limitation of 1:1 FD MAC protocols consists in, firstly, improving the capacity below the MAC layer. Toward that goal we advocate an alternative FD MAC design guideline we refer to as 1:*N*. Under this novel guideline, a MAC protocol sees the FD bandwidth through $N>1$ PHY layers. Each PHY layer is assigned to a sub-channel that is narrower than the whole available FD bandwidth and orthogonal to the sub-channel of the other PHY layers. Large N increases the number of concurrent transmissions but penalizes spectrum efficiency due to guard-bands. In this work we report proof-of-concept results from a case study for N=2 and leave other cases for future works.

B. Increased Spatial Reuse Factor

The 1:*N* design creates more FD opportunities than 1:1 by increasing spatial reuse factor, as shown in Fig. 1. In

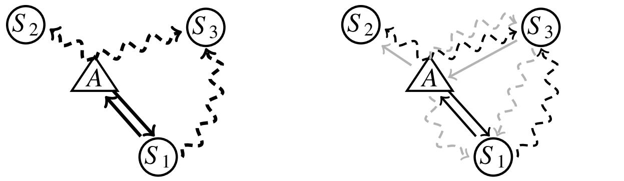

(a) 1:1 possible best-case scenario. (b) 1:N=2 possible best-case scenario.

Fig. 1: Best-case comparison: Under 1:N (b), the number of dual-links (couple of solid straight arrows) outperforms 1:1 (a) by a factor of N. Channel orthogonality (gray and black colors) overcomes interference (dashed waved arrows) to increase spatial reuse.

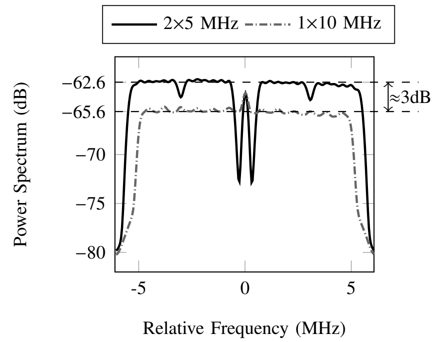

Fig. 2: Each concurrent 5 MHz-wide channel can outperform a single 10 MHz channel about 3 dB even under the same output power.

the 1:1 best-case scenario (Fig. 1a) a dual-link can increase throughput while avoiding that a hidden node (e.g. STA S_2) collides with an ongoing transmission (e.g. $S_1 \rightarrow A$). However, this sacrifices spatial reuse by interfering with all other STAs (dashed waved arrows) [5]. By arranging the FD bandwidth into N *orthogonal* narrower-channel PHY layers, the 1:N best-case scenario overlaps $N-1$ *additional* dual-links in the same space. This is illustrated on Fig. 1b for N=2, in which channel orthogonality (i.e. gray and black colors) *also* helps against collisions and enables one additional dual-link in the network.

C. Improved Signal to Noise Ratio

Prior works [12], [13] show experimentally that halving a single Wi-Fi channel increases the total energy in the bandwidth, yielding a Signal to Noise Ratio (SNR) gain of ≈3 dB. We enhance these tests to check whether the SNR statement holds when the total active bandwidth remains the same but the number (then the width) of channels changes. In each test, we set Wi-Fi signals to the same parameters. However, one scenario considers a 10 MHz-wide channel and the other considers two concurrent 5 MHz-wide channels. In Fig. 2 we plot the Power Spectrum (PS) of the strongest signals as reported by a couple of single-antenna Ettus USRP B210 platform. We estimate the PS samples and their average based on the Matlab's pwelch procedure. From the plots, one can see each narrow channel benefits from ≈3 dB gain over the wider channel. In fact, although both narrow channels occupy the same 10 MHz spectrum, they are employed independently. Thus, both the environmental and noise floors experienced within a channel does not account for the signal processing in the other.

D. Capacity Model Below the MAC Layer

The SNR improvements resulting from channel narrowing can translate into higher capacity for a Wi-Fi bandwidth. Consider an AWGN Wi-Fi channel measuring B (Hz) under a given SNR. According to the Hartley-Shannon theorem, the maximum information that can be modulated and carried over a half-duplex bandwidth B is C_{hd} Bits/Hz/s (Eq. 1). Assuming an FD radio and expressing the SNR in dB ($SNR_{dB} = 10\log(SNR)$), one derives Eq. 2 for the capacity limit C_{fd1} of FD MAC protocols under the 1:1 design.

With the 1:N guideline, the FD bandwidth B is equally divided among N narrow channels. Considering N=2, the 3 dB gain induced by channel narrowing, the guard-band g (Hz) and the FD capability assumed before, the total capacity $C_{fd:2}$ achieved within B is given by Eq. 3.

$$C_{hd} = B\log_2(1 + SNR) \quad (1)$$

$$C_{fd1} = 2B\log_2\left(1 + 10^{SNR_{dB}/10}\right) \quad (2)$$

$$C_{fd2} = 4\left(\frac{B-g}{2}\right)\log_2\left(1 + 10^{(SNR_{dB}+3)/10}\right) \quad (3)$$

IV. FD Wi-Fi MAC Protocol Capacity Upper-Bound

In this section we characterize the ideal condition to derive the capacity upper-bound of a Wi-Fi compliant FD MAC protocol. Then, we present a model to assess such capacity under both the 1:1 and 1:N MAC design guidelines.

A. Ideal Condition for Wi-Fi Compliant FD MAC protocols

To keep Wi-Fi compliance, a MAC protocol shall follow the CSMA/CA access method. In the context of FD radios, this means that *at least* the primary transmission initiates following a typical exponential backoff procedure. Since CSMA/CA is half-duplex by nature, some additional mechanism is required to admit a collision-free secondary transmission. The resulting time overhead to coordinate such a secondary transmission (i.e. Δ_t) is the key reason why MAC protocols' performance falls well below the FD gains [5]. Therefore, under an 'ideal FD condition', an Wi-Fi compliant MAC protocol maximizes the FD gain utilization by minimizing the time overhead Δ_t.

A naive way of characterizing the 'ideal FD condition' is assuming Δ_t=0 i.e. t_{st}=t_{pr}. This implies that the same backoff number is shared without overheads by a pair of arbitrary nodes at the beginning of each time slot. This is a too strong assumption for our scenario because conflicts with the random uniform behavior of the CSMA/CA backoff procedure. A reasonable alternative for this consists in assuming that *the PR always has a data frame enqueued to the PT*. In our scenario, this means that the minimum Δ_t corresponds to the time interval the AP needs to start the secondary transmission just after processing the incoming primary transmission's header H_1. The FD Wi-Fi MAC protocol presented in [10] enables the AP to do that in real-time. The authors explain that the WLAN throughput should increase by a factor of two – the expected FD gain over half-duplex – if, ideally, AP can always start a data frame transmission back to the current transmitting STA. Of course, $\Delta_t > 0$ might penalize the resulting FD Wi-Fi MAC throughput. This ideal condition is illustrated in Fig. 3.

In the Fig. 3, an arbitrary STA starts a primary transmission to the AP at the time instant t_0 upon winning a CSMA/CA contention (waiting time not illustrated). After receiving and processing H_1, the AP fetches a data frame and starts a secondary transmission to the corresponding STA at the time t_2. This defines the minimum Δ_t, which corresponds to $t_2 - t_0 > 0$ in the figure. Note, however, that FD becomes profitable only at t_3, the time at which useful data starts being transferred. To avoid collisions due to hidden terminals, both transmissions have to be finished simultaneously [4], then the maximum secondary payload L_2 (bytes) for the capacity upper-bound is resized accordingly. The other parameters on the Fig. 3 are helpful for the capacity model, as we explain next.

B. Capacity Limit Model

To compute the capacity limit of CSMA/CA under the ideal FD condition for each design guideline, we refer to the IEEE 802.11 capacity model proposed by Bianchi [11]. The model is twofold. Firstly it computes the probabilities τ and p that a CSMA/CA station transmits and collides, respectively. These probabilities are computed in the same way for our scenario, since the STAs contends for primary transmissions just as in half-duplex CSMA/CA. The second part of the model consists in a expression that computes the throughput S for IEEE 802.11 WLANs regardless of the channel access mode.

To assess S assuming an FD channel, we need firstly to characterize the possible events related to a primary transmission at the beginning of a time slot. In our case they correspond to same possible events of a CSMA/CA half-duplex channel, namely, 'success', 'collision' or 'absent' (empty slot). These events happen with probabilities P_s, P_c and P_i and take T_s, T_c and T_i absolute time units (e.g. μs), respectively. Of these, T_i is obtained straightforwardly from the standard waiting slot time [9]. Moreover, only the first event carries an expected amount of useful payload, that we denote as $E[L]$.

1) Probabilities of channel events: To compute P_s, P_c and P_i, recall that each one of all n STAs does transmit with probability τ and *does not* with probability $(1-\tau)$. Thus, channel is idle with probability $P_i=(1-\tau)^n$. A primary transmission succeeds if only a single STA transmits and the remaining $(n-1)$ STAs remain silent, what happens with probability $\tau(1-\tau)^{n-1}$. Since each of the n STAs has the same chance to succeed $P_s=n\tau(1-\tau)^{n-1}$. A collision happens if the channel is not idle and, at the same time, a primary transmission does not succeed i.e. $P_c=(1-P_i)(1-P_s/(1-P_i))$.

2) Duration and payload of a successful primary transmission: Let H_1 and L_1 be the PHY-MAC headers and payload sizes of a primary transmission, respectively. Similarly, H_2 and L_2 have equivalent meaning for a secondary transmission, as illustrated on Fig. 3. Also, let T_H and T_L be the time taken to transmit H_1 (or H_2) and L_1 under given control and data rates, respectively. Denoting as $SIFS$ plus T_{ACK} the total IEEE 802.11 standard time needed to acknowledge a data frame and

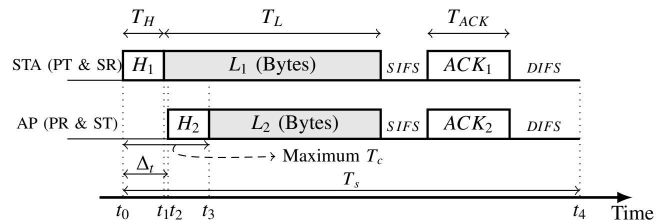

Fig. 3: Ideal FD condition for the performance of an Wi-Fi compliant FD MAC protocol. The AP (PR) always has a frame enqueued to the STA (PT). At time t_2 the AP (ST) starts sending a data frame to the STA (SR) upon receiving and processing the primary transmission header (during $[t_0,t_2]$).

δ as the physical propagation delay, the overall duration of a successful primary transmission is given by Eq. 4. Note that T_s also comprises $DIFS$ i.e. the minimum Wi-Fi standard time interval all STAs must wait before assuming channel is idle again and restarting the CSMA/CA count-down.

$$T_s = T_H + T_L + \delta + SIFS + T_{ACK} + \delta + DIFS \quad (4)$$

As one can also see on Fig. 3, under the ideal FD condition, a dual-link comprises two data frame transmissions. Therefore, the total expected payload carried within the channel event 'success' is defined as $E[L]=L_1+L_2$. Note that $L_2=L_1-f_L(H_2)$, where $f_L(H_2)$ is the amount of useful payload that the secondary transmission's data rate could send during the time interval comprising fetching and transmitting H_2 (i.e. $[t_1, t_3]$ on Fig. 3).

3) Duration of a collision: To detect a collision, the PT starts a timer just after pushing the last symbol header into the channel. If no signal is detected from the PR before the timer expires, the PT interrupts the transmission and assumes a collision. If the PT detects an incoming header, it finishes the reception to check whether the header comes from the ST. In case of a collision, the received header is unintelligible or is not the expected H_2. Only after this process, the PT interrupts the transmission. Therefore, in the worst-case, the duration of a collision is $2(T_H + \delta) + \Delta_t$. Of these values, Δ_t is exclusive of the FD technology. Particularly, in [10] authors claim a Δ_t=11μs to fetch an Wi-Fi-like frame and start an FD secondary transmission in real-time.

4) MAC guidelines saturation throughput: The FD CSMA/CA capacity formula S (Eq. 5), comes from the ratio between the payload and the time duration associated to each possible event in the channel.

$$S = \frac{P_s(L_1 + L_2)N}{P_s T_s + P_c T_c + P_i T_i} \quad (5)$$

Each value in the ratio are weighted by the corresponding channel event probability. This formula stands for both design guidelines. The difference is that N=1 for the 1:1 design. Hence, under the ideal FD condition, each CSMA/CA round triggers two transmissions across the whole channel. With the 1:N design, N>1 and each CSMA/CA round triggers 2×N narrow-channel transmissions under the same ideal condition. Also, all timing parameters rescale according to channel width just as the IEEE 802.11 standard mandates [9].

V. Results

In this section we report the performance gains both the 1:1 and 1:N MAC design guidelines present over the IEEE 802.11a half-duplex MAC protocol. Firstly, we consider the theoretical capacity each guideline delivers to the MAC layer. We also present a proof-of-concept study to validate the reported capacity gains. Then, to facilitate our performance scalability study at the MAC layer, we refer to the analytic model presented in Section IV-B. In all tests, both the PHY and MAC parameters are set according to the IEEE 802.11a standard [9]. Due to space constraints, we kindly refer the reader to our source codes [14] to reproduce all our results.

A. Novel Capacity Limit Below the MAC Layer

In Fig. 4a we plot the theoretical capacity upper-bound for the 1:1 and the 1:N=2 design guidelines against half-duplex across different SNRs (Eqs. 2 and 3, respectively). The total bandwidth is B=10 MHz so 1:N corresponds to two 5 MHz channels. Each 5 MHz channel is separated by a guard-band g=100 KHz, what can be achieved by actual filters e.g. [8]. We also plot the half-duplex capacity for comparison purposes (Eq. 1). Theoretically, the gain of any FD radio is bounded by 2× the half-duplex capacity. However, the SNR gains induced by channel narrowing breaks such expected gain even paying a 100 KHz guard-band overhead. We verified this result holds for a g up to ≈1 MHz.

To investigate whether the above-two FD gain preserves in practice, we propose a proof-of-concept study based on a pair of Ettus USRP B210 Software Defined Radio (SDR) platforms. Each radio is equipped with one antenna for transmission and one for reception. We compare a single 10 MHz Wi-Fi channel against two 5 MHz Wi-Fi channels. An *ideal* FD radio doubles capacity by entirely releasing the bandwidth for reception while transmitting. To mimic such behavior, we rely on an *out-of-band* FD test. Thus, in all FD scenarios, each radio has 10 MHz channel dedicated for reception and another 10 MHz for transmission. To mimic an ideal SIC circuit we separate the 10 MHz channels by 60 MHz. In the 1:1 design, a single PHY layer is attached to the whole available bandwidth. As a result, two 10 MHz transmission are supported, what should double the half-duplex capacity. With our 1:N case study, each 10 MHz channel is split into two orthogonal 5 MHz channels. Since the whole 10 MHz bandwidth is FD, four 5 MHz simultaneous transmissions are supported.

We set the highest modulation the IEEE 802.11 standard mandates under a Received Signal Strength Indication (RSSI) of −80 dBm i.e. QPSK 3/4 for 10 MHz and 16-QAM 1/2 for 5 MHz [9, Table 18–14]. This yields data rates of 9 and 6 Mbps, respectively. We produce Wi-Fi signals based on the gr-ieee80211 GNURadio module [15] and measured all bytes transferred through saturated links. Since SDR experiments are sensible to the CPU load and FD doubles such processing demands, we assess the half-duplex link from the best FD link. For each experiment we gather as many samples as needed to calculate mean throughput with a confidence of 95% and a relative error < 5%, following the statistical procedures of [16]. From the plots on Figs. 4b and 4a, one can see that the proof-of-concept experiments presents lower capacity in comparison to the theory. The reason is the latency introduced by the USB 3.0 connection between the USRP and the host PC. However, the 1:N design guideline outperforms the expected (1:1) FD gain (i.e. 2× the half-duplex capacity) in both cases. In the proof-of-concept experiments (4b), the (1:1) FD capacity doubles the half-duplex capacity as predicted in theory (Fig. 4a). However, the 1:N design guideline improves over half-duplex by about 2.2× and breaks through the expected (1:1) FD gain. It is worthy to remark that the capacity of actual *in-band* FD radios can be strictly less than 2× half-duplex's because of residual self-interference. *However, our findings suggest that the gains claimed by (1:1) FD radio proposals might be underestimated*. For instance, we believe that the best currently reported FD gain – 1.87× in an 80 MHz channel under P_R dBm of RSSI [1] – could be improved if performed over two 40 MHz FD channels set to efficient filters/guard-bands and to the densest Wi-Fi modulation scheme supported under P_R dBm.

B. Novel Capacity Limit at the MAC Layer

To check whether the FD gains can *scale* at the MAC layer, we report the saturation throughput for both the ideal 1:1 and 1:N FD Wi-Fi MAC protocols in comparison to the IEEE 802.11 half-duplex MAC protocol. All throughput results are computed in accordance to the analytic model of section IV-B. For half-duplex, we report results considering the basic access mode (DATA/ACK) and the Request-to-Send/Clear-to-Send (RTS/CTS) access mode. We do not consider the RTS/CTS handshake for FD MAC protocols since FD symmetric links naturally replace such mechanism [4]. Recall that the ideal 1:1 FD Wi-Fi MAC protocol corresponds to the best-case of the protocol proposed in [10], which theoretically doubles the half-duplex capacity by assuming that an AP can always start a data frame transmission back to the current transmitting STA. We assume an air propagation delay of δ=1μs, a bandwidth of B=20 MHz and N=2 (i.e. two 10 MHz channels for 1:N). All other timing parameters are set according to the IEEE 802.11a best-effort traffic class.

We verify that the FD MAC protocols outperform the half-duplex Wi-Fi across different data rates and frame payload sizes. Due to space constraints, on Fig. 4c we only report results for data rate of 48 Mbps in 20 MHz channels. This implies in at least 27 Mbps for 10 MHz channels [9]. Similarly, for these respective channel widths, we set control rates to 18 Mbps and 12 Mbps and MAC payload to 536 bytes. Larger payloads dramatically damages 2-way half-duplex performance upon collisions, specially as network grows (Fig. 4c). The 4-way handshake mitigates that by preceding data transmission with smaller RTS frames but the overall handshake slows all successful transmissions. In turn, with FD only a very small part of the primary transmission's payload is exposed to collision. This happens with no penalty to successful transmissions.

In addition to the FD advantages, the poor half-duplex

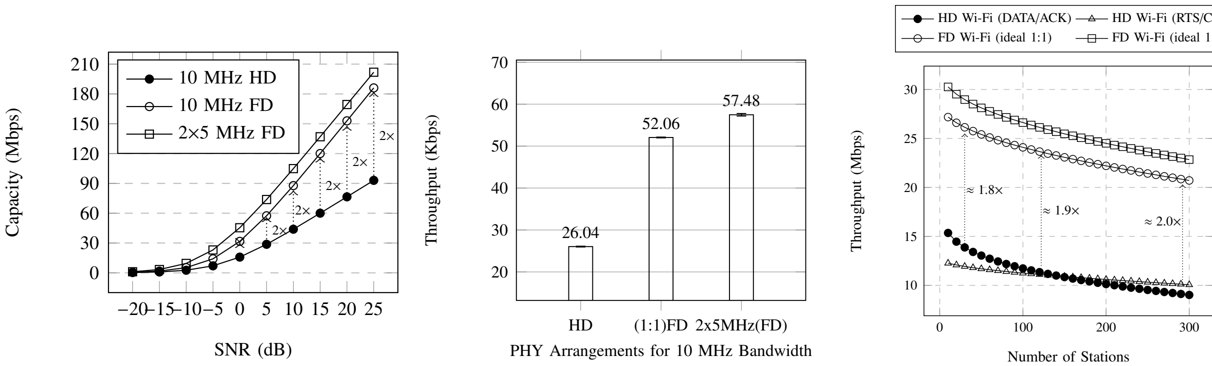

(a) Theoretical capacity below the MAC layer. (b) Proof-of-concept capacity experiments. (c) Throughput scalability at MAC layer.

Fig. 4: Full-Duplex (FD) vs. Half-Duplex (HD). Under 1:N, the capacity delivered to the MAC layer more than doubles in comparison to a HD channel (4b), as predicted in theory (4a). Similarly, the saturation throughput of an ideal 1:1 FD Wi-Fi MAC protocol [10] improves under the 1:N design guideline for an increasing number of nodes (4c).

performance over an increasing number of nodes causes the 1:1 FD CSMA/CA to be as higher as 2× the half-duplex performance (as of ≈290 nodes, Fig. 4c). However, as one can also see on Fig. 4c, such gain can be improved by conforming the 1:1 FD Wi-Fi to the 1:N design guideline. The channel orthogonality exploited by the 1:N design guideline enables higher spatial reuse, causing the number of concurrent transmissions to increase by a factor of N in the FD bandwidth. Also, the sum data rate of two narrow concurrent transmissions can be as high as the data rate of a single wide bandwidth transmission. Thus, under the same ideal conditions of the 1:1 FD Wi-Fi MAC protocol, the 1:N design guideline can scale a gain of 2× the half-duplex throughput. Although non-exhaustive, these results represent an unprecedented step towards the FD gain scalability in single-cell WLANs.

VI. Conclusion and Future Work

In this work we study the capacity limits of single-cell FD WLANs. We inquire what prevents current Wi-Fi compliant FD MAC protocols to fully profit from the theoretical double of throughput leveraged by FD radios. In addition to the overheads *at* the MAC layer, we realize this is also explained by the capacity bound imposed *below* the MAC layer. Thus, we propose a design categorization based on which MAC protocols are classified according to the way they 'see' the FD bandwidth. In this sense, we identify current FD Wi-Fi MAC protocols are classified into what we refer to as the 1:1 design guideline, meaning they 'see' the FD bandwidth through a *single* PHY layer and bound the best-case to a pair of transmissions in the FD channel. Instead, under the 1:N design guideline we advocate, MAC protocols 'see' the FD bandwidth through N>1 orthogonal narrow-channel PHY layers. Based on theoretical results and software defined radio experiments, we show it is possible to outperform the current assumed FD capacity gain at and below the MAC layer. In future works one can design novel mechanisms to exploit the increased spatial reuse factor of our proposal. We plan to study the impact of residual self-interference and guard-bands on throughput for large N. Also, we intend to study the 1:N design along with the MIMO technology.

References

[1] D. Bharadia, E. McMilin, and S. Katti, "Full duplex radios," in *Proc of the ACM SIGCOMM 2013*. ACM, 2013, pp. 375–386.

[2] S. Goyal, P. Liu, O. Gurbuz, E. Erkip, and S. Panwar, "A distributed MAC protocol for full duplex radio," in *Signals, Systems and Computers, 2013 Asilomar Conference on*, Nov 2013, pp. 788–792.

[3] M. Duarte, A. Sabharwal, V. Aggarwal, R. Jana, K. Ramakrishnan, C. Rice, and N. Shankaranarayanan, "Design and characterization of a full-duplex multi-antenna system for WiFi networks," *Vehicular Technology, IEEE Transactions on*, vol. PP, no. 99, 2013.

[4] N. Singh, D. Gunawardena, A. Proutiere, B. Radunovic, H. V. Balan, and P. Key, "Efficient and fair MAC for wireless networks with self-interference cancellation," in *Int Symp on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt'11)*, 2011, pp. 94–101.

[5] X. Xie and X. Zhang, "Does full-duplex double the capacity of wireless networks?" in *INFOCOM, 2014 Proceedings IEEE*, April 2014, pp. 253–261.

[6] S. Queiroz and R. Hexsel, "Translating full duplexity into capacity gains for the high priority traffic classes of IEEE 802.11," in *Applied Computing (SAC), 2015 ACM/SIGAPP 30th Symposium on*, 2015.

[7] J. Fang, K. Tan, Y. Zhang, S. Chen, L. Shi, J. Zhang, Y. Zhang, and Z. Tan, "Fine-grained channel access in wireless LAN," *IEEE/ACM Trans. Netw.*, vol. 21, no. 3, pp. 772–787, Jun. 2013.

[8] K. Krishna, B. Radunovic, V. Balan, M. Buettener, S. Yerramalli, V. Navda, and R. Ramjee, "WiFi-NC: WiFi over narrow channels," in *Proc 9th USENIX*, 2012, pp. 43–56.

[9] IEEE, "IEEE standard for information technology– part 11: Wireless LAN MAC and PHY specifications amendment 5: Enhancements for higher throughput," *IEEE Std 802.11-2012*, 2012.

[10] J. Mayank, J. I. Choi, T. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha, "Practical, real-time, full duplex wireless," in *Proc 17th Int Conf on Mobile Computing and Networking*. ACM, 2011, pp. 301–312.

[11] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 535–547, 2006.

[12] 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 Trans on Networking*, vol. PP, no. 99, 2012.

[13] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl, "A case for adapting channel width in wireless networks," *SIGCOMM Computer Communication Review*, vol. 38, no. 4, pp. 135–146, 2008.

[14] S. Queiroz, "Saulo queiroz's git repository," https://github.com/sauloqueiroz/thesis-codes.git, 2016, last accessed: Jun 27, 2016.

[15] B. Bloessl, M. Segata, C. Sommer, and F. Dressler, "An IEEE 802.11a/g/p OFDM receiver for GNU radio," in *Proceedings of the Second Workshop on Software Radio Implementation Forum*, ser. SRIF '13. ACM, 2013, pp. 9–16.

[16] G. C. Ewing, K. Pawlikowski, and D. McNickle, "Akaroa2: Exploiting network computing by distributing stochastic simulation," in *Int Society for Computer Simulation*, 1999, pp. 175–181.