Concise Paper: Semi-Synchronous Channel Access for Full-Duplex Wireless Networks

Xiufeng Xie and Xinyu Zhang
University of Wisconsin-Madison
Email: {xiufeng,xyzhang}@ece.wisc.edu

***Abstract*—Full-duplex radios are often envisioned to double wireless link capacity. Substantial work has focused on redesigning the radio hardware to achieve this theoretical gain. From a network-protocol perspective, however, it remains an open problem how to exploit full-duplex radio, and how much gain it can achieve in practical multi-cell wireless LANs. In this paper, we propose FuMAC, a channel access protocol tailored for full-duplex radios to optimally exploit their unique capabilities. FuMAC addresses a unique tradeoff between PHY-layer full-duplex transmission and MAC-level spatial reuse, through a semi-synchronous channel access principle. Its design is enabled by a novel self-interference cancellation mechanism called Active Antenna Cancellation. We verify FuMAC using software-radio implementation combined with large scale simulation. The results demonstrate that conventional MAC protocols severely under-utilize full-duplex's potential. In contrast, FuMAC can achieve more-than-doubled throughput gain over half-duplex wireless LANs and significantly outperform alternative full-duplex MAC designs, while maintaining a much higher level of fairness.**

I. Introduction

Full-duplex radio implementation has driven substantial research efforts in recent years [1]–[3]. Theoretically, full-duplex technology can double the wireless link data rate. However, it remains an open problem how full-duplex PHY-layer can be integrated with a MAC protocol and how much performance gain it can achieve in a wireless network with multiple contending links.

The focus of this paper is to exploit the full-duplex potential through a MAC protocol that enhances 802.11 multi-cell wireless LANs. We identify a unique design tradeoff in full-duplex networks, *i.e.*, full-duplex offers extra transmission opportunities, but at the cost of worse spatial reuse: A full-duplex transmission pair adds one more transmitter to a half-duplex link, which expands the interference range and tends to suppress other transmission attempts around it.

We propose *FuMAC* to address this unique tradeoff. Prior work advocated a mix of half- and full-duplex transmissions in a full-duplex MAC [4]. However, through an analytical model, we find that the throughput-optimal strategy is to persistently, instead of opportunistically, leverage full-duplex. This requires a full-duplex pair to synchronize their transmissions whenever possible, yet they still need to respect the independent/asynchronous carrier sensing and channel access specified in CSMA. Thus, we call this strategy *semi-synchronous contention*. We show semi-synchronous contention not only optimizes full-duplex potential, but also circumvents the notorious exposed terminal problem. Note that the hidden terminal problem is known to be solved naturally by full-duplex [1].

The principle of semi-synchronous contention in FuMAC necessitates a revisit of the conventional MAC primitives in CSMA. In particular, we incorporate a transmitter-side collision detection, and a collision intensity sensing (CIS) mechanism, to realize efficient channel access and collision resolution. At the PHY layer, these mechanisms are facilitated by active antenna cancellation (AAC), a novel realization of full-duplex self-interference cancellation through 802.11ac MU-MIMO technology.

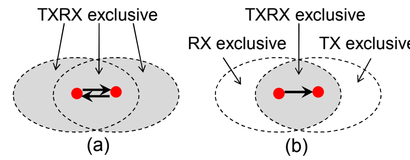

Fig. 1: Spatial reuse in (a) full-duplex and (b) half-duplex networks. Dotted lines denote interference range. RX (TX) exclusive region can be reused by other TX (RX). But TXRX exclusive region cannot be reused by others.

To sum up, FuMAC makes the following contributions: *(i)* A semi-synchronous contention mechanism that optimizes full-duplex potential; *(ii)* A MU-MIMO based self-interference cancellation mechanism to support the MAC primitives; *(iii)* A software-radio implementation, combined with large-scale simulation, to verify the feasibility of FuMAC, and its 2.85$\times$ throughput gain over 802.11 half-duplex MAC.

II. Unique Tradeoff: Full-duplex vs. Spatial Reuse

In this section, we investigate the unique tradeoff in full-duplex MAC design, *i.e.*, the extra transmission opportunities provided by full-duplex versus loss of spatial reuse. We focus on the *bi-directional transmission* paradigm, *i.e.*, concurrent transmission/reception between a pair of full-duplex nodes. A full-duplex node can also serve as a wormhole relay between two other nodes, but this approach is only applicable when those two nodes are hidden from each other, which is relatively uncommon in practical topologies [5].

A. Full-Duplex Loses Spatial Reuse

Since a full-duplex link expands the interference range compared with a half-duplex link, the network supports fewer concurrent links, thus diminishing full-duplex's benefits. We formalize this intuition by characterizing the spatial occupation of a full-duplex link by an exact *exclusive region* around the corresponding pair of transceivers[1]. See [7] for the proof.

Lemma 1. *Define the exclusive region of a full-duplex link as the union of two circular regions around the two transceivers, with radius* $R = \frac{(1+\Delta)r}{2}$*, where* r *and* $(1+\Delta)r$ *are the transmission and interference range, respectively. The necessary and sufficient condition for interference-free full-duplex transmission is that no exclusive regions overlap.*

[1]This is in stark contrast to half-duplex networks where only a lower-bound of the exclusive region can be obtained [6]

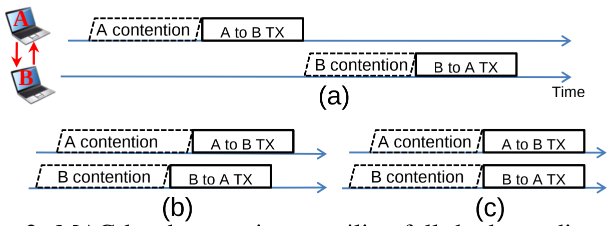

Fig. 2: MAC-level strategies to utilize full-duplex radios: (a) Half-duplex access; (b) Opportunistic full-duplex access; (c) Semi-synchronous access.

Lemma 1 allows us to characterize the exact spatial usage of each full-duplex link. In contrast, for half-duplex, the "exclusive region" of half-duplex networks does not preclude the existence of other concurrent transmitters/receivers [6], as shown in Fig. 1. In other words, *we can pack more half-duplex links than full-duplex ones, in the same network area*. Given the tradeoff, is it worthwhile to always utilize the extra transmission opportunities provided by full-duplex radios?

B. Semi-Synchronous Full-Duplex Transmission

We answer the above question through an asymptotic comparison between different MAC-layer strategies to exploit full-duplex. For a full-duplex link, a MAC protocol may schedule transmission from both sides so that their packets *(i)* never overlap: This essentially reduces the radios to half-duplex mode (Fig. 2(a)). *(ii)* opportunistically overlap: This strategy allows the two sides to run 802.11 CSMA independently and leverage full-duplex only if their transmissions happen to overlap (Fig. 2(b)). *(iii)* persistently overlap: A transmission starts only if both sides of the link gain channel access, *i.e.*, the transmissions are always synchronized at packet level (Fig. 2(c)). We denote the above three strategies as *half*, *asyn* and *syn*, respectively. By computing the mean area of the exclusive region defined in Lemma 1 for these strategies, we can analytically compare their asymptotic throughput in a CSMA multi-cell wireless LAN (see [7] for detailed proof).

Theorem 1 *Suppose all nodes have saturated traffic and denote R_s as the mean throughput achieved by strategy s, then we have the following asymptotic relation:*

$$R_{syn} \geq R_{asyn} \geq R_{half} \tag{1}$$

Our analysis concludes that *syn* strategy is optimal to exploit full-duplex transmission under the tradeoff with spatial reuse. However, this strategy requires a pair of full-duplex nodes to synchronize their channel access, *i.e.*, finishing random backoff at the same time, which violates the asynchronous nature of current distributed CSMA protocols. A main objective of this work is to introduce a novel MAC mechanism, *semi-synchronous contention* (Semi-SYN), to resolve this problem unique to full-duplex, which allows both sides to access the channel in a synchronized manner while still performing carrier sensing and channel contention independently.

III. FuMAC Design

In this section, we proceed to design FuMAC, a protocol that realizes the semi-synchronous contention principle for full-duplex multi-cell wireless LANs. FuMAC's MAC layer

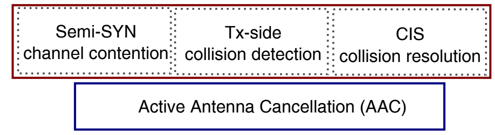

Fig. 3: Architecture and components of FuMAC.

Algorithm 1 Semi-synchronous channel access.

```
/* backoff based on CIS (Sec. III-C) */
If carrier sensing and backoff succeeds
  Start transmission
If an idle node A receives preamble from the other side
  B of its full-duplex link
   /* Semi-synchronous transmission */
   If node A senses an idle channel
     If node A has packets pending
       Stop backoff; Start transmission immediately
     Else node A emits busy-tone
   /* Silence as implicit collision notification*/
   Else node A remains silent
/* Transmitter-side collision detection */
If collision detected
   Collision resolution based on CIS (Sec. III-C)
```

is comprised of three core components as illustrated in Fig. 3: Semi-SYN channel contention (Sec. III-A), Tx-side collision detection (Sec. III-B), and collision-intensity sensing (CIS) based collision resolution (Sec. III-C). These MAC components build on top of a novel full-duplex PHY implementation called Active Antenna Cancellation (Sec. III-D), which runs over 802.11ac MU-MIMO hardware.

A. Semi-Synchronous Channel Contention

1) Basic contention mechanism: Algorithm 1 summarizes the semi-synchronous contention mechanism. FuMAC prioritizes synchronous full-duplex transmission opportunities while adhering to the CSMA primitives. Before transmission, each node independently performs CSMA, *i.e.*, running carrier sensing and backoff to gain channel access. When a node wins the channel contention, it initiates the transmission, while its intended receiver transmits a data packet back immediately as long as it senses no alien transmitters. Otherwise the receiver remains silent. Note that an AP can form a bi-directional transmission pair with different clients. Thus, each client needs to check the destination field of the AP's packet to decide whether it is the intended receiver and backward transmitter.

If the initial transmitter hears no response from its receiver, it infers that either a collision occurred or the receiver senses a busy channel occupied by another full-duplex transmission pair. In either case, the receiver cannot receive the packet properly. Therefore, the initial transmitter aborts transmission immediately to release the channel if it does not receive any backward transmission within a short time threshold.

In case the receiver can access the channel but has no data packet in queue for the initial transmitter, it will still emit a busy tone back. This prevents the initial transmitter from falsely aborting its transmission and wasting channel time. The busy tone itself does not waste any channel time or worsen spatial reuse, because it only suppresses alien interferers,

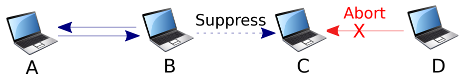

Fig. 4: Semi-SYN contention circumvents exposed terminals.

which should not be allowed to transmit anyway. As a side benefit, the busy-tone also eliminates hidden terminals [1].

With these primitives, FuMAC maintains the asynchronous CSMA mechanism, but can still enforce persistent bi-directional full-duplex transmission. More specifically, both sides of a full-duplex link perform carrier sensing and random access independently, but they are allowed to transmit simultaneously if either gains channel access. This ensures efficient channel access discussed in Theorem 1, but fairness remains an issue, which will be addressed separately in Sec. III-C.

After finishing transmission, a node A needs to await ACK from its peer B. When B's packet is longer than A's, node A has to wait until B finishes transmission, which wastes channel time. Therefore, both sides of the full-duplex link should align their packets to fully utilize the channel. Specifically, the node that initiates full-duplex bi-directional transmission embeds packet length information in its header, and the other node adjusts its packet size to match the length.

2) Circumventing exposed terminals: Consider the exposed terminal topology in Fig. 4. For a half-duplex network, spatial reuse allows up to two active links B→A and C→D. However, the exposed terminal problem in CSMA suppresses one transmission opportunity C→D that should have been enabled without interference. In contrast, under FuMAC's semi-synchronous contention, the lost transmission opportunity is compensated by the backward transmission A→B. When the full-duplex link A↔B is active, C can neither transmit or receive packets, *i.e.*, no exposed terminal exists. Meanwhile, any node (D) that attempts to contact C, who is suppressed by A and cannot transmit back, will abort transmission, thus avoiding collision.

B. Transmitter-Side Collision Detection

Prior work showed the feasibility of receiver-side collision detection [8]. Upon detecting a collision, the receiver can notify the transmitter to abort transmission. But transmitting and decoding the notification packet incurs nontrivial delay. In FuMAC, we advocate transmitter-side collision detection by upgrading the 802.11 packet preamble. Each transmitter is assigned a unique PN-sequence as a signature, which is embedded in the preamble. In case when transmitters fall in the same collision domain, if a full-duplex node overhears an alien signature while transmitting, then it knows there must be another concurrent transmitter nearby, *i.e.* collision occurs.

To address more general cases where multiple nodes collide or colliding nodes are hidden terminals, we make the following observation. When collision happens, intended receivers cannot decode transmitters' signatures, and thus remain silent. Following semi-synchronous contention, the colliding transmitters abort transmission as they hear nothing from their intended receivers. Therefore, collision can still be *implicitly* detected without any notification overhead.

To summarize, under transmitter side collision detection, a transmitter aborts transmission immediately if *(i)* it hears signature from an alien interfering transmitter *or (ii)* it does not hear a header of backward transmission or busy-tone from its full-duplex peer receiver.

C. Collision Resolution with Collision Intensity Sensing

Given the collision detection capability, an important question is: is it optimal to migrate CSMA/CD from Ethernet to full-duplex wireless LANs? The key mechanism of CSMA/CD is Binary Exponential Backoff (BEB), which effectively resolves collisions: When collision occurs, only colliding transmitters double their backoff window size, while others' remain intact. However, this may cause severe short-term unfairness [9], temporarily starving the collided nodes and worsen their packet delay. Moreover, a single collision cannot reflect the overall contention level in the network. When collision probability is low, BEB may overreact by forcing a doubling of contention window size, which unnecessarily degrades performance.

FuMAC overcomes BEB's limitations by using *collision intensity*, the number of collisions a node senses within unit time, as a metric for adapting transmission aggressiveness. Rather than passively react to collision, nodes actively determine the collision intensity that leads to optimal throughput, based on time-averaged collision statistics. Then, they can adapt transmission aggressiveness to approach the optimal collision intensity. Below we give the throughput-optimal collision intensity c^* for a full-duplex wireless LAN, which will be utilized in FuMAC's collision resolution.

Theorem 2 *In a full-duplex WLAN with persistent contention, the throughput-optimal collision intensity for every node is:*

$$c^* = \frac{e^{\tau} - 1 - \tau}{t_s + \tau T_t} \tag{2}$$

which is insensitive to the number of nodes N as long as N is not too small. Here τ solves the equation $1 - \tau = (1 - t_s/T_{col})e^{-\tau}$; t_s denotes the duration of a time slot; T_t denotes the packet duration and T_{col} denotes the duration of a collision. (Proof available in [7]).

1) Adapting transmission aggressiveness based on collision intensity sensing (CIS): Below we discuss how should a node evaluate its collision intensity and adjust its transmission aggressiveness accordingly. First, each node can evaluate its collision intensity by computing a moving average of the number of collision events within unit time. Then it compares current collision intensity $c(t)$ and the known optimal collision intensity c^*. If $c(t) > c^*$, the node decreases its access probability P_{cs} by dividing a factor of β ($\beta > 1$), otherwise it increases P_{cs} by multiplying the same factor. In this way, every node can approach the optimal collision intensity and thus optimal network throughput.

2) Leveraging queue-length evolution for fairness: Theorem 2 assumes neighboring nodes have similar collision intensity, and hence similar c^*. However, this may not always hold in a multi-cell wireless LAN. For example, in the well-known flow-in-the-middle (FIM) topology [9], a transmitter in between two others has fewer channel-access opportunities, and thus it should be more aggressive, *i.e.*, risk a larger collision intensity to ensure fairness.

Algorithm 2 Collision resolution based on collision intensity sensing (CIS)

1. Each node performs semi-synchronous channel contention before transmission (Sec. III-A).
2. On detecting a collision happened around or to it, each node updates its sensed collision intensity c in a certain sliding time window.
3. Adapt $P_{cs}(c)$:
   if $\bar{c} < c^*$
   Increase aggressiveness: $P_{cs} \leftarrow P_{cs}\beta, \beta > 1$
   else
   Decrease aggressiveness: $P_{cs} \leftarrow P_{cs}/\beta$
4. Adapt $P_{ql}(Q)$ based on current queue length $Q(t)$, following Eq. (4).
5. Update access probability P_a based on P_{cs} and P_{ql}, following Eq. (3).
6. Adjust backoff window (CW) size with CW = $\frac{2}{P_a} - 1$.

To remedy this issue, we enhance CIS with transmitter's queue length information. The basic idea is intuitive: the demand-supply differential of a link can be reflected by its MAC-layer packet queue length, thus we can use current queue length of a link to control its aggressiveness. Take the FIM topology as an example, transmitter in the middle should have more packets accumulated in its queues, which can be used as a hint to increase its transmission aggressiveness.

In the actual protocol operations, a node determines its overall access probability P_a through a multiplication of two factors: $P_{cs}(c)$ which follows the above CIS mechanism; and $P_{ql}(Q)$ which is determined by current queue length.

$$P_a(c, Q) = P_{cs}(c)P_{ql}(Q) \tag{3}$$

We use the following sigmoid function for $P_{ql}(Q)$, which reflects channel congestion status based on queue length [10].

$$P_{ql}(Q) = \exp(Q)/(\exp(Q) + C) \tag{4}$$

where C is a constant determining the sensitivity, its value is chosen around 500 to make the corresponding contention window size span across all options $(2^{k+1} - 1, k = 0, ..., 9)$ defined in 802.11. Note that in actual protocol implementation, the channel access probability needs to be translated into backoff window (CW) size as in [11]: CW = $2/P_a - 1$. Algorithm 2 summarizes the above collision resolution mechanism.

D. 802.11ac-Compatible Active Antenna Cancellation

FuMAC is facilitated by a simple and novel full-duplex PHY implementation called active antenna cancellation (AAC), built on top of 802.11ac MU-MIMO hardware. Unlike passive antenna cancellation [1], AAC does not rely on antenna placement. Moreover, it performs self-interference cancellation on a per-subcarrier basis, and thus works for wideband OFDM systems like WiFi. With enough number of transmit antennas, AAC alone can enable full-duplex transmission for 802.11ac devices without hardware modification, as shown in Sec. IV-A1 later. It can be further integrated with analog self-interference cancellation circuits [12] to achieve even higher performance.

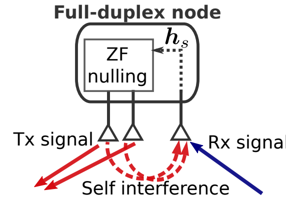

Fig. 5: Active antenna cancellation (AAC) over MU-MIMO.

1) AAC mechanism: MU-MIMO allows a multi-antenna transmitter to send multiple concurrent data streams to different users, which can be realized by Zero-Forcing Beamforming (ZFBF). ZFBF minimizes the mutual interfere between data streams by steering the null points of each data stream towards unintended users.

By nature, a full-duplex radio also requires its transmit antenna(s) to minimize self-interference to its receive antenna(s). This similarity inspires us to design AAC that enables a node to actively steer the null point of its transmitted signals over its receiving antenna using ZFBF, so as to minimize self-interference as shown in Fig. 5.

Assume a full-duplex node has one receive antenna and M transmit antennas. Let $\boldsymbol{h}_s$ denote the channel state vector from the AP's transmit antennas to its own receive antenna, and $\boldsymbol{h}_r$ that to its peer receiver. Let $\boldsymbol{H}_s = [\boldsymbol{h}_r, \boldsymbol{h}_s]^T$, we can use pseudo inverse to compute a *beamforming weight matrix* $\boldsymbol{W} = [\boldsymbol{w}_1, \boldsymbol{w}_2, \cdots, \boldsymbol{w}_M]$ that satisfies the ZFBF constraint:

$$\boldsymbol{W} = \boldsymbol{H}_s^{\dagger} = \boldsymbol{H}_s^{*}(\boldsymbol{H}_s\boldsymbol{H}_s^{*})^{-1} \tag{5}$$

In this way, this node can serve at most M spatial streams without interfering with each other, in which one special stream is steered to its own receive antenna but with all symbols equal 0, thus the self-interference is nulled.

More generally, the AAC framework can be extended to implement full-duplex MIMO. Suppose both sides of the full-duplex link has E receive antennas, then at least $2E$ transmit antennas are needed for each node to null self-interference to its own E Rx antennas and serve E concurrent data streams to its peer. We leave this to future work.

2) AAC implementation: We have implemented a prototype of AAC on top of a 802.11ac-compatible MU-MIMO OFDM system built on the WARP software radio platform, in which the transmitting module uses ZFBF precoding as in Eq. (5) to simultaneously perform self-interference nulling and data transmission. The precoded data streams are then modulated by OFDM and prepended with the 802.11ac packet preambles [13].

When the preamble transmitted by the node itself is detected, the channel state information $\boldsymbol{h}_s$ estimated based on LTF [13] is sent to Tx module of the same node for ZFBF, which requires no over-the-air feedback. Note that a receiving node cannot estimate $\boldsymbol{h}_s$ when it intends to transmit back because the channel estimation will be interfered by the ongoing transmission it is receiving. To circumvent this problem, an AAC node can perform self-interference nulling using the $\boldsymbol{h}_s$ estimated from its latest channel estimation preamble, which is sent the last time when this node won the channel contention. This is because the self-interference channel $\boldsymbol{h}_s$ between Tx and Rx antennas on the same node remains relatively stable owing to the very close distance between them.

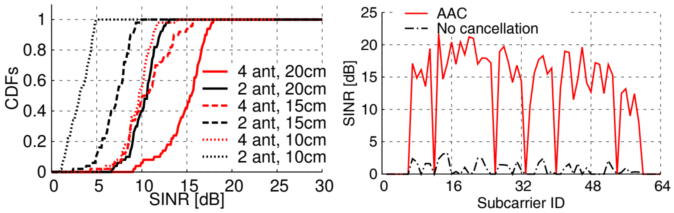

(a) Under different Tx antenna number & interfering distance.

(b) Receiving SINR on 64 OFDM subcarriers (20MHz channel).

Fig. 6: Performance of AAC implementation.

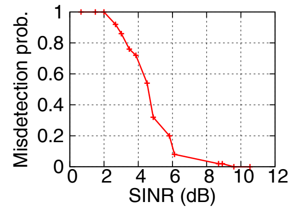

Fig. 7: Full-duplex collision detection with AAC.

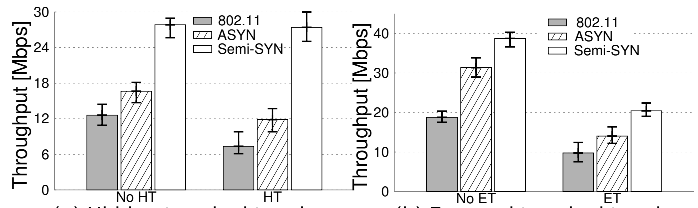

(a) Hidden terminal topology

(b) Exposed terminal topology

Fig. 8: Throughput comparison between different contention mechanisms, with and without hidden/exposed terminals.

IV. Performance Evaluation

In this section, we first evaluate the feasibility of FuMAC's AAC mechanism using a prototype implementation over the WARP software-radio platform. Since WARP's processing/interface latency prevents us from implementing the real-time MAC layer, FuMAC's MAC layer performance is evaluated on a custom-built C++ discrete-event simulator.

A. Verifying Active Antenna Cancellation (AAC)

1) Evaluating AAC full-duplex performance: To examine whether AAC can effectively suppress the self-interference without dedicated interference cancellation circuit, we trigger bi-directional transmissions between two full-duplex nodes, each of which is constructed by 2 clock-synchronized WARP boards and has up to 4 Tx antennas and 1 Rx antenna. The carrier frequency is 2.4GHz and bandwidth is 20MHz.

Fig. 6(a) plots the CDF of the full-duplex link SINR with varying Tx antenna number and self-interference distance D_s which is the distance between the Rx antenna and the closest Tx antenna. We observe that, under the same D_s, using 4 Tx antennas always achieves better performance than 2. Also, the SINR increases with larger D_s. When the D_s is only 10cm, 4 Tx antennas can ensure the SINR of most packets is greater than 6.5dB, the minimum for decoding 802.11 packets [14], while 2 Tx antennas requires at least 15cm distance to satisfy the same requirement. Existing passive antenna cancellation mechanism based on antenna placement can cancel around 10 dB of self-interference [1], but it does not apply in OFDM system. In contrast, AAC efficiently cancels self-interference and achieves around 15 dB of SINR on all OFDM subcarriers as shown in Fig. 6(b), where notches are pilot and DC subcarriers.

2) Transmitter-side collision detection: To evaluate the full-duplex transmitter-side collision detection proposed in Sec. III-B, we vary the SINR of full-duplex link by controlling the transmit power. Fig. 7 shows that the collision misdetection rate drops very close to 0 when overhearing SINR is about 6dB. For the implicit collision detection, receiver of the colliding node can always detect the collision if it cannot decode the signature of any transmitter. On the other hand, if one signature dominates the other, the receiver can decode the stronger one and perform backward transmission to it. Collision is resolved in this case even though the nodes does not explicitly detect it.

B. FuMAC Performance

We emulate FuMAC's MAC layer using a C++ discrete-event simulator, in which we configure the MAC/PHY parameters following typical values in 802.11: time slot is $9\mu s$ and packet size is 1.5KB. In simulating FuMAC's implicit collision detection, we assume a timeout period of 10 MAC slots.

1) Effectiveness of semi-synchronous contention: To verify the superiority of semi-synchronous contention, we compare it with the two alternatives discussed in Theorem 1: *half* strategy used in 802.11 which is naturally half-duplex, and *asyn* strategy. To isolate the impact of collision resolution mechanism, we assume BEB is used to resolve collision for all cases. We also assume both uplink and downlink traffic are saturated. Fig. 8(a) and Fig. 8(b) plot the network throughput of different strategies under benchmark topologies each containing two links, with and without hidden/exposed terminals.

Theoretically, the capacity gain of full-duplex over half-duplex is upperbounded by 2. However, Fig. 8(a) shows that semi-synchronous contention can have a throughput gain greater than 2, even without hidden terminals. This is attributed to the implicit collision detection mechanism in FuMAC, which effectively shortens the duration of each collision. In case when transmitters are hidden terminals, collision misleads half-duplex transmitters to become extremely conservative as they will have longer backoff duration, while FuMAC achieves almost 7 times of throughput gain. Similarly, in exposed terminals, 802.11 DCF underutilizes channel due to unnecessary waiting and backoff. Combining efficient channel contention with the bi-directional transmission, FuMAC achieves more than $3\times$ throughput gain in this case. In addition, FuMAC has much higher throughput than the *asyn* strategy in all cases, which is consistent with our analysis.

2) System level evaluation: We evaluate the system level performance of FuMAC in random networks of various scales, in comparison with legacy 802.11 and one existing full-duplex MAC protocol, FD-MAC in [4], the only existing work that shares similar idea to synchronize transmission from both sides of a full-duplex link. To simulate randomly deployed multi-cell WLANs, we generate Poisson distributed APs with a given number in a fixed area, and clients are uniformly distributed within the range of each AP.

Simulated throughput performance is shown in Fig. 9(a).

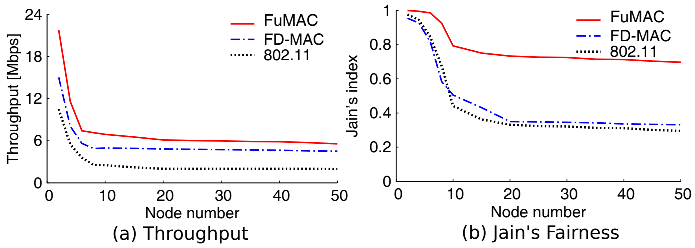

Fig. 9: Throughput and fairness in large multi-cell WLANs.

We observe that average throughput of all protocols decreases with increasing node density, but full-duplex protocols always have significant throughput gain over half-duplex protocol. *FuMAC achieves a throughput gain of around 2.85 over half-duplex scheme*. This is mainly attributed to a combination of semi-synchronous channel contention and CIS mechanisms that fully exploit full-duplex capabilities at the MAC level. FuMAC also outperforms FD-MAC, which wastes some full-duplex transmission opportunities due to its inefficient way to synchronize channel access — it uses half-duplex data packets to synchronize the backoff of both sides of the link, thus data transmission of both sides only partially overlap. Fig. 9(b) shows the Jain's fairness index of these protocols. Fairness of all protocols decreases with node density, due to increasing chance of asymmetrical contention within the topology. However, FuMAC maintains significantly higher fairness compared to other protocols. Summarizing, these results show that FuMAC can achieve more-than-doubled throughput gain over half-duplex networks, while maintaining a high level of fairness in realistic large-scale topologies.

V. Related Work

Recent research on full-duplex mainly targets the implementation of full-duplex transceivers. Choi *et al.* [1] realized single-channel full-duplex over 802.15.4 radios through passive antenna cancellation combined with analog self-interference cancellation, which inspired substantial work (*e.g.*, [3], [15]). Recently, Bharadia *et al.* [12] implemented the first full-duplex WiFi radio with a single antenna.

Little work has been devoted to the impact of full-duplex on higher layers. Singh *et al.* [16] proposed a MAC protocol for ad-hoc networks, based on the full-duplex wormhole relaying mechanism [1]. RCTC [5] realized a full-duplex mode selection scheme between bi-directional transmission mode and wormhole relaying mode. Their results show that the full-duplex transmission opportunities in wormhole relaying mode are few. Goyal *et al.* proposed a CSMA based full-duplex MAC protocol [17], which directly reuses the random back-off mechanism of 802.11, thus inherits all its disadvantages. FD-MAC [4] allows a pair of transceivers to cooperate with shared random backoff thus synchronize their full-duplex transmission. However, it encounters severe fairness problem in multi-cell topologies like FIM, which should be more common in full-duplex networks. Moreover, as shown in Sec. IV, FD-MAC cannot fully exploit full-duplex transmission opportunities due to its inefficient way to realize synchronized channel access. Janus [18] uses a centralized mechanism to schedule transmission within a single cell, but it fails in multi-cell networks. To our knowledge, FuMAC is the first *CSMA-compatible* MAC that utilizes full-duplex radio for transmitter-side collision detection and derives a unique semi-synchronous contention principle for multi-cell WLANs.

VI. Conclusion

In this paper, we provide theoretical and experimental evidence showing the importance of a MAC protocol specifically designed for multi-cell full-duplex WLANs. In particular, we investigate the optimal strategy to address a unique tradeoff between full-duplex and spatial reuse. Based on insights gained from the investigation, we propose FuMAC, a MAC protocol designed to fully exploit the unique features of full-duplex radio. We build FuMAC on top of a novel full-duplex PHY realization, Active Antenna Cancellation (AAC), which is compatible with 802.11ac MU-MIMO hardware. In our evaluation, AAC is verified through a prototype built on WARP software radio, and extensive simulation shows that FuMAC achieves more-than-doubled throughput gain over half-duplex protocols, while maintaining a higher level of fairness.

References

[1] J. I. Choi, K. Srinivasan, P. Levis, and S. Katti, "Achieving Single Channel, Full Duplex Wireless Communication," in *MobiCom*, 2010.

[2] M. Duarte, A. Sabharwal, V. Aggarwal, R. Jana, K. K. Ramakrishnan, C. W. Rice, and N. K. Shankaranarayanan, "Design and Characterization of a Full-duplex Multi-Antenna System for WiFi networks," *CoRR*, vol. abs/1210.1639, 2012.

[3] E. Aryafar, M. A. Khojastepour, K. Sundaresan, S. Rangarajan, and M. Chiang, "MIDU: Enabling MIMO Full Duplex," in *MobiCom*, 2012.

[4] A. Sahai, G. Patel, and A. Sabharwal, "Pushing the Limits of Full-duplex: Design and Real-time Implementation," *CoRR*, vol. abs/1107.0607, 2011.

[5] W. Zhou, K. Srinivasan, and P. Sinha, "RCTC: Rapid Concurrent Transmission Coordination in Full Duplex Wireless Networks," in *Proc. of IEEE ICNP*, 2013.

[6] P. Gupta and P. Kumar, "The Capacity of Wireless Networks," *IEEE Transactions on Information Theory*, vol. 46, no. 2, 2000.

[7] X. Xie and X. Zhang, "Semi-Synchronous Channel Access for Full-duplex Wireless Networks," http://goo.gl/FndUtT, 2014.

[8] S. Sen, R. R. Choudhury, and S. Nelakuditi, "CSMA/CN: Carrier Sense Multiple Access with Collision Notification," in *ACM MobiCom*, 2010.

[9] J. Liu, Y. Yi, A. Proutiere, M. Chiang, and H. V. Poor, "Towards Utility-Optimal Random Access Without Message Passing," *Wirel. Commun. Mob. Comput.*, vol. 10, no. 1, 2010.

[10] J. Lee, Y. Yi, S. Chong, B. Nardelli, E. W. Knightly, and M. Chiang, "Making 802.11 DCF Optimal: Design, Implementation, and Evaluation," *CoRR*, vol. abs/1207.3740, 2012.

[11] M. Heusse, F. Rousseau, R. Guillier, and A. Duda, "Idle Sense: An Optimal Access Method for High Throughput And Fairness in Rate Diverse Wireless LANs," in *ACM SIGCOMM*, 2005.

[12] D. Bharadia, E. McMilin, and S. Katti, "Full Duplex Radios," in *ACM SIGCOMM*, 2013.

[13] "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," *IEEE Std. 802.11ac Draft 3.0*, 2012.

[14] X. Xie, X. Zhang, and K. Sundaresan, "Adaptive Feedback Compression for MIMO Networks," in *Proc. of ACM MobiCom*, 2013.

[15] M. Jain, 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. of ACM MobiCom*, 2011.

[16] N. Singh, D. Gunawardena, A. Proutiere, B. Radunovic, H. Balan, and P. Key, "Efficient and Fair MAC for Wireless Networks with Self-Interference Cancellation," in *Proc. of WiOpt*, 2011.

[17] S. Goyal, P. Liu, O. Gurbuz, E. Erkip, and S. Panwar, "A Distributed MAC Protocol for Full Duplex Radio," in *IEEE Asilomar Conference on Signals, Systems and Computers*, 2013.

[18] J. Y. Kim, O. Mashayekhi, H. Qu, M. Kazandjieva, and P. Levis, "Janus: A Novel MAC Protocol for Full Duplex Radio," Stanford University, Tech. Rep., 2013.