Recent Advances in Interference Management for Wireless Networks

Amr Nabil[†] Y. Thomas Hou[†*] Rongbo Zhu[†‡] Wenjing Lou[†] Scott F. Midkiff[†]

† Virginia Polytechnic Institute and State University, USA

‡ South-Central University for Nationalities, China

Abstract

Interference has been the central challenge for wireless networks. In wireless networking, the prevailing paradigm to handle interference is avoidance. Over time, many interference avoidance techniques have been proposed following this paradigm. Recently, research advances at the physical layer allow us to explore a new direction in interference management. The new direction is to allow interference to occur and exploit the desired information from interference, rather than avoiding interference completely. This new direction allows much higher utilization of radio channel and spectrum and opens the door for a whole new perspective on how interference should be managed in a wireless network. This article offers a timely overview of recent advances in this exciting area, with a focus on its application in wireless LAN (WLAN). We envision that the deployment of these new techniques will lead to dramatic change in wireless networking paradigm, with profound impact on the future research direction for the wireless networking community.

1 Introduction

Interference is a fundamental challenge in wireless networks. In this paper, we consider interference in its most general form: at a receiver, if its desired received signal is being overlapped by another signal, we say this desired signal is being interfered with and the interfering signal is considered as an interference. If the interfering signal comes from another node, then we call it *mutual interference*; if the interfering signal comes from the same node, then we call it *self interference*. At a receiver, the desired signal may be decoded successfully if and only if the interfering signal is below a certain threshold. Otherwise, the desired signal may not be decodable at the receiver and the corresponding transmission is considered unsuccessful.

In the wireless networking community, the classic (and still prevailing) paradigm to cope with interference is *avoidance*, which can be done in time, frequency, code, or space domains. Interference avoidance addresses potential

[*]For correspondence, please contact Prof. Tom Hou (thou@vt.edu).

interference by not allowing any other signal to overlap with the desired signal at a receiver in the underlying domain. Although interference avoidance is simple to implement, its throughput region is rather small.

Recent advances in interference management have broken away from this traditional avoidance approach. The new paradigm aims at *exploiting* and/or *canceling* interference rather than avoiding it. The basic idea is to allow multiple signals to be transmitted in the same channel at the same time. Then through special interference cancellation (IC) and/or decoding techniques, a receiver tries to decode its desired signal correctly. As expected, this new approach has the potential to improve wireless channel utilization and increase network and user throughput significantly.

The goal of this paper is to present an overview of the recent advances in this exciting area. We classify these new interference management techniques into two broad areas:

- **Solutions to Handle Self Interference:** Traditional wireless nodes can only operate in half-duplex mode. Under half duplex, a wireless node can either transmit or receive (in a channel) but not both at the same time. On the other hand, full duplex allows *simultaneous* transmission and reception of signals by the same transceiver in the same channel. The key to full duplex is to cut down self-interference power level below a certain threshold. Consequently, the receiver will be able to distinguish the desired received signal from self-interfering signal and hence, the desired received signal is successfully decoded. However, it is very difficult (once considered impossible) to cut down self-interference power to the noise level. Recently, a number of major advances has been made in this area. In Section 2, we review these advances.

- **Solutions to Handle Mutual Interference:** Mutual interference differs fundamentally from self interference in that a receiver does not have complete knowledge of the interfering signal as it is produced by a different node. In a WLAN, mutual interference can impede a receiver from being able to decode desired signals successfully. For over two decades, many designs have been proposed to overcome this problem. The crux in most of these solutions is to avoid mutual interference. Recently, new solutions have been proposed to exploit the interference rather than avoiding it. These techniques, once properly designed, may successfully extract and decode the desired signal at the receiver. In Section 3.1, we will review new techniques that only need a single antenna at a node.

  Another way to mitigate mutual interference is to exploit the capabilities of Multiple-Input Multiple-Output (MIMO) technology. Since its inception, MIMO has transformed wireless communications [1]. In particular, the so-called "zero-forcing" technique allows us to null potential interfering signal at an unintended receiver [2]. To conserve the precious resources of MIMO, the so-called interference alignment (IA) technique has been proposed and implemented. IA allows interfering signals to be overlapped in the same direction and thus be canceled efficiently using a fewer number of node resources . In Section 3.2, we will review advances in both IC and IA for WLAN with MIMO.

Due to space limitation, the coverage of this article is not meant to be exhaustive, but rather a sample of recent advances in this rapidly growing area. Other recent work may exist under each category (or beyond these categories) that are not covered in this article. Nevertheless, we hope this short article can bring to the attention of the general audience in the community on this new and exciting direction for future wireless networking.

The rest of the paper is organized as follows. In Section 2, we review recent advances in reducing self interference and achieving full duplex. In Section 3, we present new solutions for mutual interference, depending on whether or not the underlying techniques require MIMO capability.

2 New Solutions for Self Interference

Throughout the history of wireless communications, a transceiver has been limited by half duplex. Consequently, the mindset of most people has been framed to half duplex and it has been assumed by default that half duplex is an intrinsic property of a wireless transceiver. This mindset, however, is no longer valid, thanks to recent advances in full-duplex technologies. The main challenge in achieving full duplex is to eliminate the strong self-interference signal from the transmitter so that the (weak) desired signal from another node can be decoded successfully at the receiver. In a WLAN environment, a single-antenna node is transmitting a 20dBm signal while at the same time, an intended signal ($\sim$-65dBm) is being received at the receiver in the same channel. The background noise level is about -90dBm. The desired received signal can be decoded in the presence of background noise when there is no other stronger interfering signal. So, the full-duplex problem is how to cut the self-transmitted signal down to the background noise level (i.e., reduce the self-interference signal power by at least 110dB). If this is possible, then the receiver can consider the diminished self-interference signal as background noise and decode the desired received signal successfully.

Antenna Cancellation Technique: In [3], Choi *et al.* pointed out that just using RF and/or digital cancellation circuits cannot bring significant reduction in self-interference signal power. They proposed to employ an antenna cancellation circuit on the top of traditional RF and digital cancellation circuits to cancel self interference. This technique requires two transmit antennas and one receive antenna. One transmit antenna is placed at a distance d from the receive antenna, while the other transmit antenna is placed at a distance of d plus half wavelength of the transmitted signal on the opposite side from the receive antenna. Under this setting, the transmitted signal is sent over the two transmit antennas simultaneously so that at the receive antenna, they cancel each other due to the half wavelength distance. Due to this difference in propagation distance (half wavelength), power control is needed at one of the transmit antennas so that signal cancellation can be performed cleanly. Choi *et al.* showed that in conjunction with RF interference cancellation and digital cancellation, the proposed simple antenna cancellation scheme can cancel 60dB of self interference. Taking into account of propagation power loss between the two transmit antennas and the receive antenna (20-30dB), the total reduction in self-interference signal is around 80-90dB. This level of

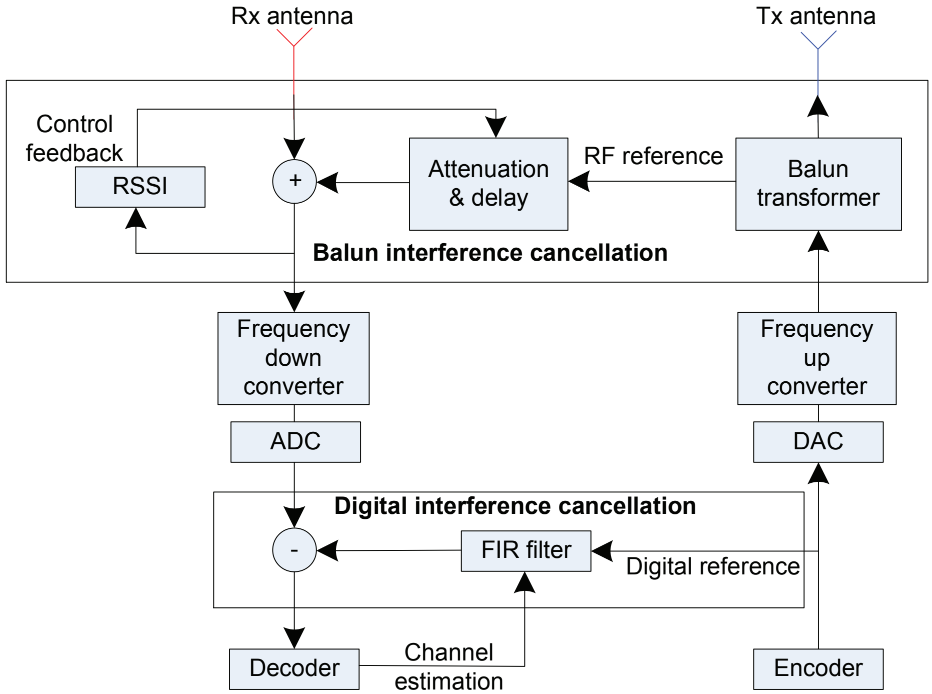

Figure 1: A schematic of the full duplex system proposed in [4].

self-interference cancellation is sufficient to enable full duplex for a low-power wireless technology such as 802.15.4 ZigBee, for which 90dB reduction is sufficient (as demonstrated in [3]). But it still falls short of the 110dB reduction that is required for WLAN (WiFi). Nevertheless, this work offers a glimpse that full duplex might be achievable and opens the door for a series of follow-up research in this area.

RF Cancellation with Balun Transformer: The antenna cancellation technique in [3] suffers from a serious limitation in that the operating wavelength of the transceiver cannot change once the location of the cancellation antenna (placed at an extra half wavelength distance from the receive antenna) is fixed. To remove this limitation, in [4], Jain *et al.* proposed a design without the cancellation antenna. Instead, there is only one transmit antenna, along with the receive antenna. The cancellation of the self-transmitted signal at the receive antenna is done by using an attenuated and delayed copy of the same signal at the transmit antenna. Such a cancellation circuitry with signal inversion is called balanced/unbalanced (Balun) transformer and is shown in Fig. 1.

Basically, a Balun transformer splits the transmitted signal into two parts: one part is sent to the transmit antenna while the other part is inverted in polarity and is used as an RF reference signal. At the receiver, the transmitted signal of the first part is then combined with an attenuated and delayed version of the RF reference signal, resulting in a cancellation of the self interference. The level of cancellation can be tuned by a control feedback circuit. In conjunction with the digital interference cancellation circuitry (in lower portion of Fig. 1), the Balun transformer (with active components) can achieve 85dB reduction in self-interference power. Similar to antenna cancellation technique in [3], the Balun transformer technique is still limited to a lower power wireless device (e.g., 802.15.4 ZigBee) but falls short of cancelling self interference for WiFi device (110dB).

A natural question to ask is how to compare the 2-antenna Balun full-duplex scheme to a 2×2 MIMO operation in half-duplex mode. In [4], the authors offered some comparison under different settings (power, SNR, and feedback). It is not entirely clear how meaningful such a comparison is (apple vs. orange) as full duplex is considered a fundamental transmission capability.

Full Duplex with Dual Transmit/Recieve Antennas with Extension to MIMO: Another approach to achieve full duplex is to employ dual transmit/receive antennas to cancel interference in the spatial domain. This idea was exploited in [5] by Aryafar *et al.*, as illustrated in Fig. 2(a). In this design, we employ dual (two) antennas for the Tx chain and another dual (two) antennas for the Rx chain at a node. With the use of an inverter on one branch of the transmit antenna, the outgoing signals from the two transmit antennas are of equal power but out of phase by π. Therefore, for any point on the cutting plane between the two transmit antennas, the two received powers cancel each other theoretically. To further improve cancellation performance, we can also employ two receive antennas on the cutting plane and is of equal distance to the center (see Fig. 2(a)). Again, an inverter is used at one of the two receive antennas so as to allow a second cancellation of any residual signals. It was shown in [5] that this dual transmit/receive antenna design allows to cancel up to 45dB of self-interference signal. In conjunction with the digital interference cancellation circuitry (similar to that in Fig. 1), it was shown that up to 75dB self-interference power can be canceled.

The extra price paid in this design is the use of dual transmit/receive antennas. But the advantage of this design is that it may be scalable (to achieve full duplex) for a MIMO node. Fig. 2(b) shows how to scale the design in Fig. 2(a) to realize a 2×2 MIMO. Note that the inner ring (with dual transmit/receive antennas) is a top view of Fig. 2(a). The second (outer) ring is another set of dual transmit/receive antennas. Combined together, we have an equivalent full-duplex 2×2 MIMO. In general, an $N \times N$ full-duplex MIMO can be built by having N rings of dual transmit/receive antennas.

Single-Antenna Full Duplex: All the full-duplex designs discussed so far were not able to achieve the ultimate goal of 110dB self-interference reduction that is needed to support WiFi devices. Further, all these designs require at least two antennas, which could be too demanding for small portable devices such as smartphones and tablets. In a very recent paper [6], Bharadia *et al.* reported that they had succeeded in building a single-antenna full-duplex system over an 80MHz bandwidth that meets the requirements of the 802.11 standard (i.e., reduction of self-interference power by 110dB). This is considered a major advance and is the first work that achieves full duplex for WiFi.

Figure 3 shows the design in [6], which consists of three parts: a circulator, an analog cancellation circuit, and a digital cancellation circuit. The circulator internally consists of three identical wave guides joined at the center to form a symmetrical Y-shape [7]. The junction of the three wave guides contains transversely magnetized symmetrical

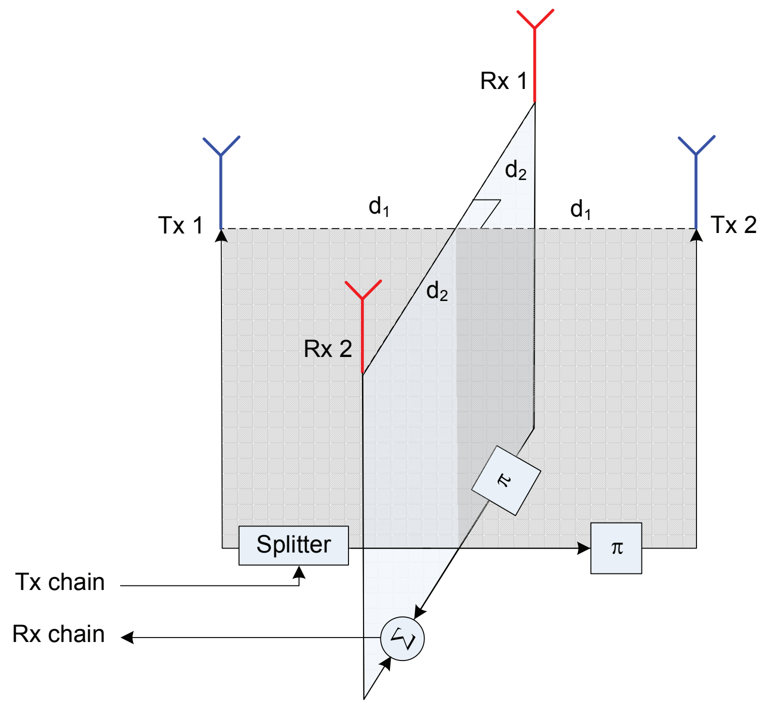

(a) 3D illustration for the two-level antenna cancellation.

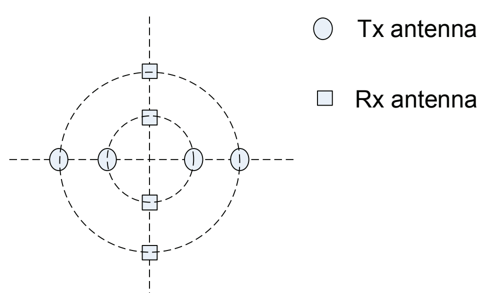

(b) Extension to 2×2 MIMO.

Figure 2: Full-duplex design with extension to MIMO as proposed in [5].

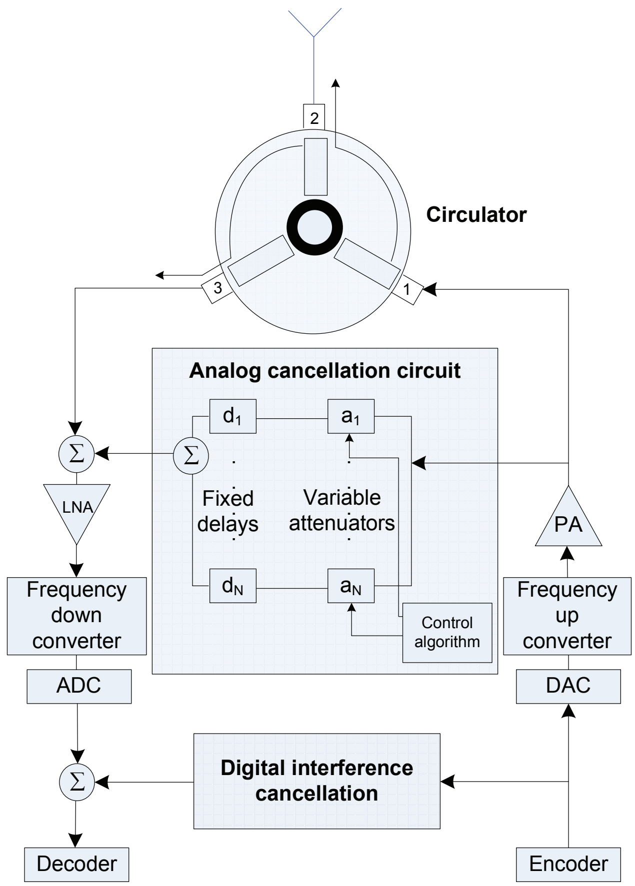

Figure 3: Full duplex radio block diagram as proposed in [6].

distribution of ferrite. The magnetized ferrite is an anisotropic medium. When a wave propagates through this special medium, the RF field distribution is asymmetrical. It was shown that the RF field can be displaced to one side so that most of the power is diverted into only one adjacent wave guide rather than both. This diversion results in passing the signal from one wave guide to the next without entering the third wave guide. As a result, any signal will be routed from one port of the circulator to the next while being isolated from entering the third port. In this design, the antenna is connected to the circulator via port 2. The circulator is connected to the TX and RX chains via ports 1 and 3 respectively. Signals are passed by the circulator from one port to the next, i.e. the transmitted signal from port 1 to port 2 and the received signal from port 2 to port 3. Consequently, the TX and RX signals, while being sent and received using the same antenna, are isolated inside the circulator. Nevertheless, the circulator cannot eliminate the leakage of the transmitted signal from port 1 to port 3 completely. It can cancel up to 15dB of self-interference signal power.

The second component in Fig. 3 is the analog cancellation circuit. It was designed to emulate and then cancel the self-interference signal, which is a delayed and attenuated version of the TX signal. The source of delay and attenuation comes from the circulator circuit. It is hard to determine this delay and Bharadia *et al.* proposed to define a range for this delay. To generate a self-interference-like signal, a copy of the TX signal before entering the circulator is taken as an input to the cancellation circuit. This copy is passed through parallel lines with different lengthes (to emulate different delays) and associated tunable attenuators. The parallel lines are separated equally, one half of them should have delays less than the lower bound of the self-interference delay range while the other half should have delays greater than the upper bound of the self-interference delay. Similar to the idea of sampling and interpolation in the digital circuit design, this setup is used in analog domain to cancel self interference. From Nyquist theorem, digital signals sampled at Nyquist rate contain all the signal information and can be used to recover the original signal. Here, the signals coming out of the parallel lines can be considered as those digital samples (from the interpolation algorithm's perspective). The best weights for the linear combination are determined and these weights are the values to which the tunable attenuators are set. As a result, this linear combination can construct a signal that is very close to the remaining self-interference signal that comes from port 3 of the circulator. This signal is then subtracted from the received signal to remove the bulk part of the self-interference signal. This operation can further cut around 60dB of self interference.

The last part is the digital cancellation circuit, which is similar to that in Fig. 1, is able to cancel 35dB of any remaining self-interference power. Putting three components together, the design in Fig 3 can achieve 110dB self-interference cancellation, which is what is needed for a practical full-duplex 802.11 transceiver. Using the proposed full-duplex system, Bharadia *et al.* were able to reach 1.87x of throughput compared to a half-duplex system.

Table 1 summarizes the self-interference cancellation techniques discussed in this section.

Reference	Main Technique	Strengths	Weaknesses	Performance (in dB)	WiFi Support
[3]	Employ two transmit antennas with proper distances to the receive antenna so that two self-interfering signals add up in a destructive manner	Simple	Limited to specific frequency (wavelength), require three antennas	80-90	No
[4]	Employ a Balun transformer inside the RF circuit to cancel self-interference through signal inversion	Independent of operating frequency	Require two antennas	85	No
[5]	Use phase shifting and dual transmit/recieve antennas to achieve two levels of self-interference cancellation	Independent of operating frequency, extensible to MIMO	Require four antennas	75	No
[6]	Employ a circulator to isolate Tx/Rx paths to/from a single antenna, followed by a special design of analog cancellation circuit	Use only one antenna, wide bandwidth	Complexity in analog cancellation circuit	110	Yes

Table 1: A summary of self-interference cancellation techniques.

3 Solutions to Mutual Interference

Mutual interference refers to the desired signal at a receiver being overlapped by a signal from *another* transmit node. A successful mutual-interference mitigation technique should be able to extract the desired signal from the composite received signal. Mutual-interference mitigation techniques can be categorized into two main categories: single-antenna solutions and multi-antenna solutions. We will examine single-antenna solutions in Section 3.1. Note that although we call them single-antenna solutions, it is by no means to infer that they can only operate on single antenna. In fact, each of these solutions can be extended to multi-antenna setting.[1] Section 3.2 examines multi-antenna solutions, which exploit the capabilities of MIMO to cancel mutual interference. We present two MIMO-based solutions to handle mutual interference. In our discussion, we focus on the so-called "zero-forcing" technique in MIMO and present its applications in nulling cross-technology interference for WLAN. We also present interference alignment (IA) and show how it can help conserve MIMO's resources for IC.

[1] The reason why we call them single-antenna solutions is to contrast to the multi-antenna solutions in Section 3.2, which rely on MIMO capabilities in their operation.

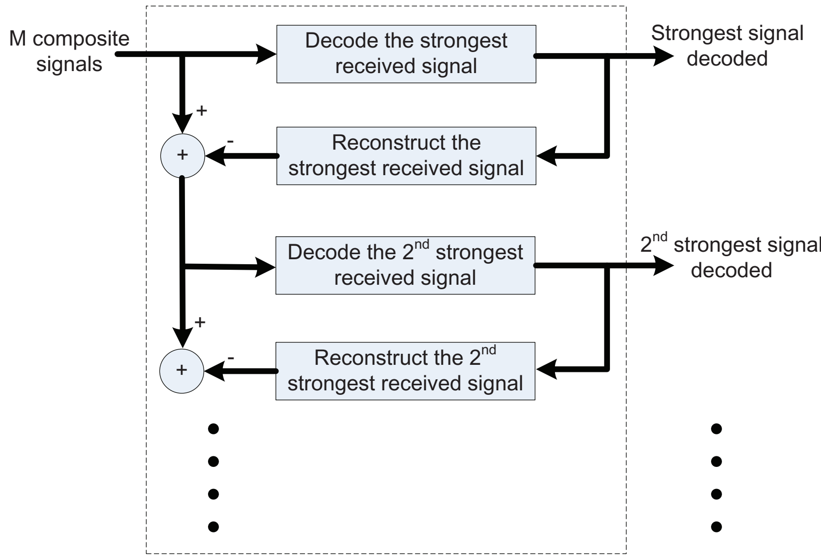

Figure 4: A schematic of SIC.

3.1 Single-Antenna Solutions

Over the years, many solutions have been proposed to address mutual interference. However, most of these solutions follow the classic interference avoidance paradigm. These solutions mainly stay at the MAC layer and focused on link layer protocol design to avoid interference. As discussed in Section 1, these interference-avoidance-based MAC solutions, although simple, do not exploit interference at the physical layer, and are known to lead to small throughput region. In this section, we review several new approaches that have been proposed to handle mutual interference by exploiting interference at the physical layer.

Successive Interference Cancellation (SIC): SIC is a well-known physical layer technique to exploit interference [8]. It allows a receiver to take multiple colliding packets from different transmitters and decode them iteratively. As shown in Fig. 4, for the composite signal, the receiver attempts to decode the strongest signal and considers all other signals as interference. If the strongest signal is successfully decoded (upon meeting certain decoding criterion), the receiver subtracts it from the original composite signal and then starts to decode the second strongest signal, and so forth. The process continues until all signals are successfully decoded, or the decoding criterion cannot be satisfied at certain stage.

Without loss of generality, denote the power levels of the M received signals at receive node j as $P_{ij}, 1 \leq i \leq M$, and are ordered in nondecreasing values, i.e., $P_{1j} \leq P_{2j} \leq ... \leq P_{Mj}$. Receive node j attempts to decode the M

signals successively according to the following algorithm:

$$\begin{aligned}
&\text{Step 1} && \frac{P_{Mj}}{\sum_{k=1}^{M-1} P_{kj} + \sigma^2} \geq \beta, \\
&\text{Step 2} && \frac{P_{M-1,j}}{\sum_{k=1}^{M-2} P_{kj} + \sigma^2} \geq \beta, \\
&\quad\vdots && \quad\vdots \\
&\text{Step } (M-i+1) && \frac{P_{ij}}{\sum_{k=1}^{i-1} P_{kj} + \sigma^2} \geq \beta, \quad (i = M, M-1, ..., 1).
\end{aligned} \tag{1}$$

The algorithm terminates when all M signals are successfully decoded or the SINR ratio is no longer greater than or equal to the decoding threshold β.

In [9], Halperin *et al.* proposed to employ SIC to address the hidden terminal problem in WLAN. They addressed two key practical issues, namely, detection of signal overlapping period and construction of received signal by taking into account of channel condition. Since it is unlikely that two packets will collide exactly from the first bit of their packets, the receiver can use the non-overlapping period (collision-free) in the first packet to determine the power level of this packet. To detect the collision (overlapping) period of the first and the second packets, the authors showed that the collision starting time can be found by a sharp increase in the amplitude of the received signal. Since the receiver knows the power levels of the composite signal (of two packets) and that of one of the packets (before collision), it can use SIC to decode the stronger signal first. Since the channel effect on the stronger signal must be taken into account before it is subtracted from the composite signal, the authors proposed to simply emulate the channel behavior by averaging over a few symbols of the received data (under the assumption of a constant channel). After reconstructing the received signal of the decoded packet by incorporating the channel condition, this signal is then subtracted from the composite received signal. This leaves behind only the signal for the other packet, which can now be decoded.

Analog Network Coding: In [10], Katti *et al.* employed analog network coding (ANC) to exploit collision and improve throughput. Their proposed idea can be illustrated in Fig. 5. In this figure, nodes A and C wish to send a packet to each other. If both of them send their packets at the same time, a collision will occur. Following interference avoidance paradigm, it will take 4 time slots to complete (schedule) packet transmission from A to C (via B) and C to A (via B), as shown in Fig. 5(b). Using ANC, both A and C transmit their packets in the same time slot and the two packets collide at B. Without decoding, B broadcasts the received composite analog signal in the next time slot. Upon receiving the composite signal at the end of the second time slot, node A subtracts its original transmitted signal from it and decode the remaining signal, which is the packet from C. Likewise, node C performs a similar operation and decodes the packet from node A. Based on the number of required time slots, ANC can double the throughput for the application scenario in Fig. 5(a).

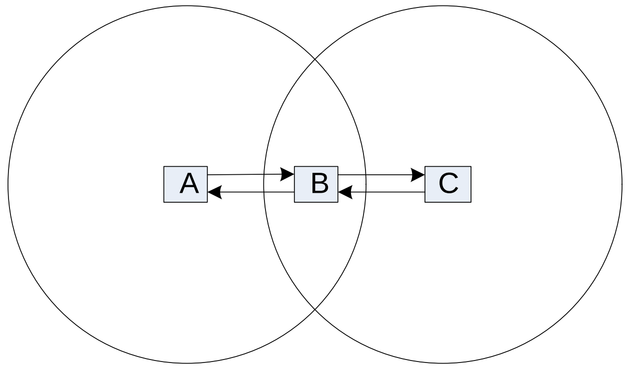

(a) A and C wish to send packets to each other via B.

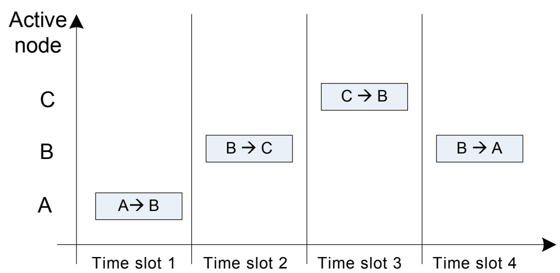

(b) Under interference avoidance scheduling, 4 time slots are needed to complete (a).

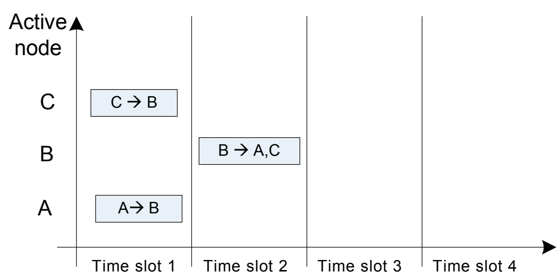

(c) Under ANC, only 2 time slots are needed to complete (a).

Figure 5: A schematic of analog network coding (ANC) as described in [10].

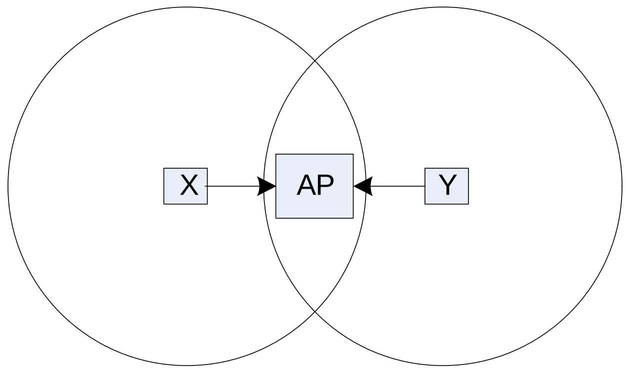

(a) The hidden terminal problem.

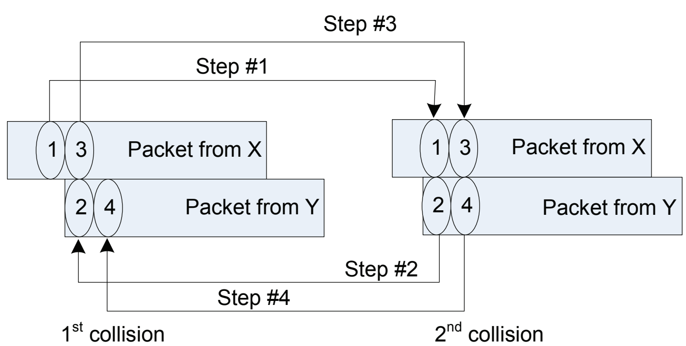

(b) ZigZag piece-meal decoding of two packets between two consecutive collisions.

Figure 6: An illustration of the ZigZag decoding algorithm as described in [11].

ZigZag: In [11], Gollakota and Katabi proposed a technique (codename "ZigZag") to exploit collision in the hidden terminal problem. Fig. 6 illustrates how this technique works. Suppose nodes X and Y are hidden from each other but can communicate with the Access Point (AP). Suppose nodes X and Y transmit their packets to the AP and result in a collision. Upon detecting the collision (either by NACK or timeout), nodes X and Y will retransmit their packets, which leads to a second collision. Under ZigZag, AP is able to determine the starting point (offset) when the two packets collide (overlap) in the two consecutive collisions. As shown in Fig. 6(b), the AP first decode the chunk of a packet that is free of collision in the first collision (i.e., chunk 1 of the packet in the first collision). Then it locates position of chunk 1 of the same packet in the second collision. Since chunk 1 from node X is already successfully decoded by the AP, the overlapping part with chunk 1 from node Y in the second collision (say chunk 2) can now be decoded by subtracting chunk 1 signal. Upon successfully decoding chunk 2 from node Y in the second collision, the AP can go back (ZigZag) to the first collision and decode chunk 3 following the same token and so forth. This back and forth decoding on chunks in the two packets continues until all chunks in the two packets are decoded successfully. It is not hard to see that by successfully decoding two packets in two collisions, the throughput is similar to transmitting two different packets through perfect scheduling.

3.2 Multiple-Antenna Solutions

MIMO Zero-Forcing and Cross-Technology Interference Cancellation: Due to the availability of multiple antennas, MIMO has intrinsic capability to perform IC. In particular, the so-called "zero-forcing" technique allows nulling undesired interfering signals at an unintended receiver [12]. Zero-forcing can be done either at a transmit node or a receive node. If done at the transmit node, it is called "zero-forcing precoding." Under zero-forcing precoding, assuming the channel state information (CSI) is known at the transmitter, the transmitter can pre-process the signals so that their interference is completely nulled at an unintended receive node. If done at the receive node, it is called "zero-forcing equalization." Under this technique, again assuming CSI is known at this receiver, the receiver can equalize the channel effect so that the interfering signals are completely nulled at the receive node. Regardless of whether zero-forcing is done at transmitter side or receiver side, there are some effective techniques to obtain CSI between transmit node and receive node as long as they are using the same transmission technology and protocol.

When transmission technology and protocol used by a transmitter is not recognizable at a receiver, IC of such a signal at the receiver becomes much more challenging, due to the lack of adequate means to obtain CSI between the transmitter and receiver. This is the so-called "cross-technology interference" in [13]. To cancel cross-technology interference, Gollakota *et al.* [13] proposed a Technology Independent Multi-Output (TIMO) design at a WLAN receiver. The basic idea of TIMO is to consider the received signal from a non-WLAN transmitter as if it was transmitted with the WLAN technology. Denote X as the WLAN signal and I as the non-WLAN (cross-technology) interferer signal. Denote h_{x1} and h_{x2} as the channel coefficients between the WLAN transmitter and the first and second receive antennas at the WLAN receive node, respectively; denote h_{i1} and h_{i2} as the channel coefficients between the non-WLAN interferer and the first and second receive antennas at the WLAN receive node, respectively. Denote Y_1 and Y_2 as the received signals at the first and second receive antennas at the WLAN receive node, respectively. Then we have:

$$\begin{aligned} Y_1 &= h_{x1}X + h_{i1}I \ , \\ Y_2 &= h_{x2}X + h_{i2}I \ . \end{aligned}$$

The desired WLAN signal X can be obtained by:

$$X = \frac{Y_1 - \beta Y_2}{h_{x1} - \beta h_{x2}} \ , \tag{2}$$

where $\beta = \frac{h_{i1}}{h_{i2}}$. The values of all parameters in (2) are known except β. Assuming that the desired WLAN signal X and the non-WLAN interfering signal I are independent, Gollakota *et al.* derived an expression to obtain β as a function of the power of WLAN signal X, h_{x1} and h_{x2}. They validated their TIMO idea for an 802.11n testbed with digital cordless phone, wireless monitor, microwave oven, and Bluetooth devices operating in the same 2.4GHz frequency band.

Interference Alignment and Cancellation: In the last few years, the information theory community has witnessed significant advances in IA. The basic idea of IA is to precode signals at different transmitters so that they are aligned in the same direction at a receiver. Once used correctly, IA can align interfering signals in the same direction without any overlapping with the desired signal, which is in an orthogonal directions and can be decoded cleanly at the receiver [14].

In [15], Gollakota *et al.* demonstrated a benefit of interference alignment and cancellation (IAC) technique in a WLAN environment. Their idea is illustrated in Fig. 7. Without IAC, the maximum number of packets that can be transmitted simultaneously over a wireless channel from the two clients to their respective APs is two. With IAC, Gollakota *et al.* showed that it is possible to transmit three packets simultaneously. Here, client 1 sends two packets (p_1, p_2) while client 2 sends a third packet (p_3) in the same channel. Each of the three packets is multiplied by a specific precoding vector (v_1, v_2, v_3, respectively) at its transmitter side. At AP 1, due to the precoding vectors, p_2 and p_3 are aligned in a direction perpendicular to the direction of p_1. As a result, p_1 can be successfully decoded. Then AP 1 sends the successfully decoded p_1 to AP 2 over a backend connection. Upon receiving p_1, AP 2 approximates the analog signal of p_1 going through the channel between AP 2 and client 1 (using the known channel coefficient H_{12}) and subtracts this approximated signal from the composite signal that it has received, which used to include $p_1 v_1 H_{12}$, $p_2 v_2 H_{12}$, and $p_3 v_3 H_{22}$. With only $p_2 v_2 H_{12}$ and $p_3 v_3 H_{22}$ left, AP 2 can now decode p_2 and p_3 using the basic capabilities of a 2×2 MIMO receiver.

Gollakota *et al.* showed analytically that IAC can be generalized for a greater number of antennas ($M > 2$). Using a MIMO system with M antennas per node, three or more APs, and at least 2 clients, IAC can achieve $2M$

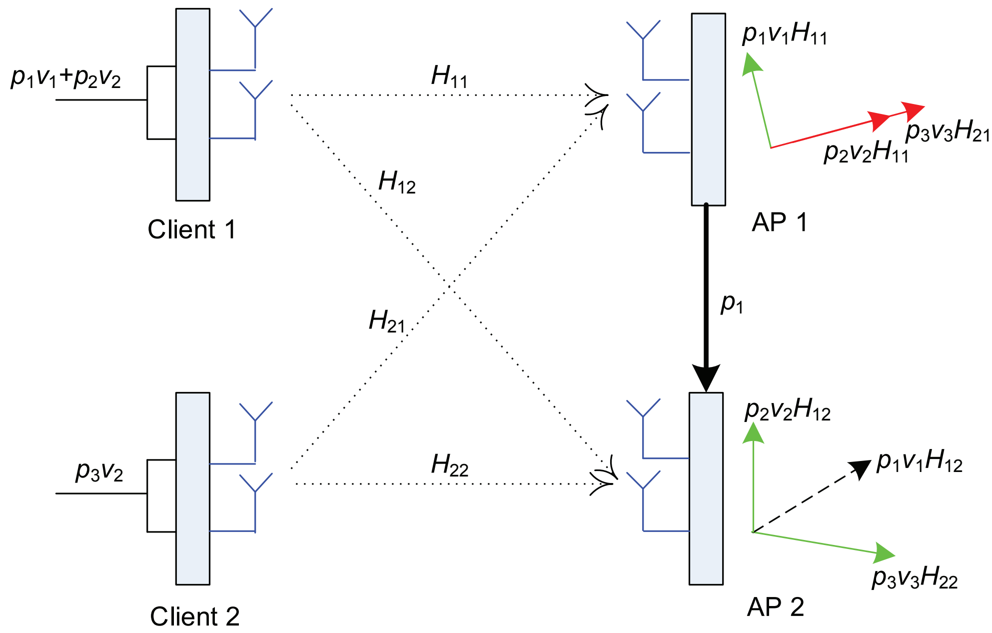

Figure 7: Interference alignment and cancellation (IAC) technique in [15].

packets on the uplink. For the downlink, $\max\{2M-2, \lfloor\frac{3}{2}M\rfloor\}$ can be achieved with M antennas per node and $(M-1)$ APs. This means that IAC almost doubles the throughput of MIMO systems when the number of antennas becomes large.

4 Conclusions

This paper offered a concise review of recent advances in interference management in wireless networks. The underlying theme of these advances is to pro-actively cancel and/or exploit interference rather than avoid it. Our review covered a number of new techniques in self-interference management and mutual-interference management. These techniques allowed to much higher utilization of radio channel and spectrum. They showed a whole new direction of how interference should be managed in a wireless network. We anticipate that advances in this area will continue in the future and these new enabling technologies will further push the performance envelop of wireless networks to a new frontier.

Acknowledgments

This research was supported in part by NSF under grants 1156318, 1343222, 1247830, 1064953, 1102013 and ONR under grant N000141310080. The work of R. Zhu was supported by the National Natural Science Foundation of China (No. 61272497, 60902053). Amr Nabil was supported in part by the VT-MENA program of Egypt.

References

[1] E. Biglieri, R. Calderbank, A. Constantinides, A. Goldsmith, A. Paulraj, and H. V. Poor, *MIMO Wireless Communications,* Cambridge University Press, 2007.

[2] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication,* Chapter 8, Cambridge University Press, 2005.

[3] J. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti, "Achieving Single Channel, Full Duplex Wireless Communication," in *Proc. ACM MobiCom,* Chicago, IL, Sep. 2010, pp. 1-12.

[4] M. Jain, J. Choi, T. Kim, D. Bharadia, K. Srinivasan, S. Seth, P. Levis, S. Katti, and P. Sinha, "Practical, Real-time, Full Duplex Wireless," in *Proc. ACM MobiCom,* Las Vegas, NV, Sep. 2011, pp. 301-312.

[5] E. Aryafar, M. A. Khojastepour, K. Sundaresan, S. Rangarajan, and M. Chiang, "MIDU: Enabling MIMO Full Duplex," in *Proc. ACM MobiCom,* Istanbul, Turkey, Aug. 2012, pp. 257-268.

[6] D. Bharadia, E. Mcmilin, and S. Katti, "Full Duplex Radios," in *Proc. ACM SIGCOMM,* Hong Kong, China, Aug. 2013, pp. 375-386.

[7] H. N. Chait and T. R. Curry, "Y-Circulator," *Journal of Applied Physics,* vol. 30, no. 4, April 1959, pp. 152-153.

[8] S. Verdu, *Multiuser Detection,* Chapter 7, Cambridge University Press, 1998.

[9] D. Halperin, T. Anderson, and D. Wetherall, "Taking the Sting out of Carrier Sense: Interference Cancellation for Wireless LANs," in *Proc. ACM MobiCom,* San Francisco, CA, Sep. 2008, pp. 339-350.

[10] S. Katti, S. Gollakota, and D. Katabi, "Embracing Wireless Interference: Analog Network Coding," in *Proc. ACM SIGCOMM,* Kyoto, Japan, Aug. 2007, pp. 397-408.

[11] S. Gollakota and D. Katabi, "ZigZag Decoding: Combating Hidden Terminals in Wireless Networks," in *Proc. ACM SIGCOMM,* Seattle, WA, Aug. 2008, pp. 159-170.

[12] Q. H. Spencer, A. L. Swindlehurst, and M. Haardt, "Zero-forcing Methods for Downlink Spatial Multiplexing in Multi-user MIMO Channels," *IEEE Trans. Signal Process.,* vol. 52, no. 2, Feb. 2004, pp. 461-471.

[13] S. Gollakota, F. Adib, D. Katabi, and S. Seshan, "Clearing the RF Smog: Making 802.11n Robust to Cross-Technology Interference," in *Proc. ACM SIGCOMM,* Toronto, Ontario, Canada, Aug. 2011, pp. 170-181.

[14] S. A. Jafar, *Interference Alignment: A New Look at Signal Dimensions in a Communication Network,* Now Publishers Inc, 2011.

[15] S. Gollakota, S. D. Perli, and D. Katabi, "Interference Alignment and Cancellation," in *Proc. ACM SIGCOMM,* Barcelona, Spain, Aug. 2009, pp. 159-170.

Biographies

AMR NABIL (anabil@vt.edu) received his B.S. degree in Electronics and Communications Engineering and M.S. degree in Engineering Physics, both from Cairo University, Giza, Egypt, in 2005 and 2009, respectively. He is currently working toward his Ph.D. degree in the Bradley Department of Electrical and Computer Engineering at Virginia Tech, Blacksburg, VA, USA. His research focuses on optimization and algorithm design for wireless networks.

Y. THOMAS HOU [F'14] (thou@vt.edu) is a Professor in the Bradley Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA, USA. His research focuses on developing innovative solutions to complex problems that arise in wireless networks. He was named an IEEE Fellow for contributions to modeling and optimization of wireless networks. He was on the editorial boards of a number of IEEE transactions. He is the Chair of IEEE INFOCOM Steering Committee.

RONGBO ZHU (rongbozhu@gmail.com) is currently an Associate Professor in College of Computer Science of South-Central University for Nationalities, China. He received his Ph.D. degree in Communication and Information Systems from Shanghai Jiao Tong University, China, in 2006. Dr. Zhu was a visiting scholar at Virginia Tech from 2011 to 2012. His research interests are performance optimization and protocol design of wireless networks.

WENJING LOU [S'01–M'03–SM'08] (wjlou@vt.edu) received her Ph.D. degree in Electrical and Computer Engineering from the University of Florida in 2003. She is a Professor in the Department of Computer Science at Virginia Tech. Prof. Lou's research interests are cyber security and wireless networks. She is on the editorial boards of a number of IEEE transactions. She is the Steering Committee Chair of IEEE Conference on Communications and Network Security (CNS).

SCOTT F. MIDKIFF [S'82–M'85–SM'92] (midkiff@vt.edu) received the Ph.D. degree in electrical engineering from Duke University, Durham, NC, USA, in 1985. He is a Professor and Vice President for Information Technology and Chief Information Officer with Virginia Tech, Blacksburg, VA, USA. His research interests include wireless and ad hoc networks, network services for pervasive computing, and cyber-physical systems.