Beyond Full Duplex Wireless

Kannan Srinivasan[†], Steven Hong, Mayank Jain, Jung Il Choi, Jeff Mehlman, Sachin Katti, and Philip Levis
kannan@cse-ohio-state.edu, {hsiying, mayjain, jjungilchoi, mehlman, skatti}@stanford.edu, pal@cs.stanford.edu
Ohio State University, Stanford University

[†]Due to an editing mistake, Kannan Srinivasan was incorrectly removed from the version of this paper published in the official Asilomar 2012 proceedings. He is the correct first author of this paper. -Philip Levis

A. *Abstract*

Recent work has shown the possibility of implementing full-duplex wireless radios using commodity hardware. We discuss the possibility of extending full-duplex designs to support multiple input, multiple output (MIMO) systems. We explore how such a design could lead to a rethinking of wireless networks. We discuss various applications of full-duplex radios and the gains possible with those applications. We also discuss some of the challenges present in getting such radios and their applications to be a part of production networks.

I. INTRODUCTION

Recent results have shown that one can build practical radios which transmit and receive simultaneously on the same or nearby bands, using physical-layer techniques and RF designs, such as antenna cancellation [2], balun cancellation [13], paired RF transmit chains [4], [11], and circulators [10]. These full-duplex designs are efficient: some come within 8% of the performance of an ideal full duplex system [2].

The goal of this paper is to take a step back from the technical details of radio design and ask more general questions: what are the implications of self-interference cancellation and what are the benefits to wireless network design? Clearly, there are obvious immediate benefits that full duplex provides, namely higher throughput and fairness. But we believe, and argue in this paper, that self-interference cancellation will have a much greater impact on wireless network design beyond full duplex. Specifically, we show via design sketches that self-interference cancellation can both greatly simplify the design of the wireless control plane and allow devices to use spectrum much more flexibly.

Current antenna cancellation is not general enough to enable many of the above scenarios. It supports only a single data stream, making it incompatible MIMO radios, which are a critical part of the recent 802.11n and future standards. Further, the existing design cannot support asymmetric throughput on forward and reverse streams. This limitation means that the control channel uses resources equivalent to the data channel itself, which is highly wasteful in scenarios where the control channel needs very low throughput. We show the original design can be generalized to remove these limitations. Our design extends antenna cancellation to support multiple antennas and asymmetric forward and reverse channels. With it, we can make any n-antenna MIMO system full duplex, with configurable number of simultaneous transmit and receive antennas.

This paper is a first step towards exploring the potential of full-duplex wireless radios, and we hope it opens up a new research avenue for the community to tackle some long standing problems in wireless networks.

II. WIRELESS FULL DUPLEXING

This section describes the design of the current full-duplex system as described in the paper by Choi et al [3]. Their system uses a combination of three techniques, namely Antenna Cancellation, RF Cancellation and Digital Cancellation to achieve full-duplex operation.

Antenna Cancellation involves using two transmit and one receive antennas. For a wavelength λ, the two transmit antennas are placed at distances d and $d + \frac{\lambda}{2}$ from the receive antenna. Offsetting the two transmitters by half a wavelength causes their signals to add destructively and cancel one another. This creates a null position where the receive antenna hears a much weaker signal.

RF Interference Cancellation implements self-interference cancellation in the analog domain using a noise cancellation circuit. The transmit signal is fed to the circuit as a noise reference, which subtracts it from the received signal, after adjusting for phase and amplitude.

Digital Cancellation uses received digital samples after the analog-to-digital conversion in the receive path. The transmitted samples are stored in a local storage. The received samples are correlated with transmitted samples to determine the beginning of the transmitted packet and its phase in the received samples. The transmitted samples are rotated by the appropriate phase and subtracted from the received samples to almost completely remove the transmitted signal from the received signal. Existing work like Zig-Zag and SIC [7], [8] uses baseband interference cancellation techniques .

Figure 1 shows the block diagram of a full-duplex system implemented using the three self-cancellation techniques. This full-duplex system assumes reception and transmission each use a single radio. Using this design, Choi et al. have argued that the symmetric bi-directional data link could be used to solve some problems with wireless such as hidden terminals, congestion, fairness in WLANs and excessive latency in multihop networks.

III. GENERALIZED DESIGN

This section describes how to generalize antenna cancellation for non-symmetric operation; it extends the full-duplex

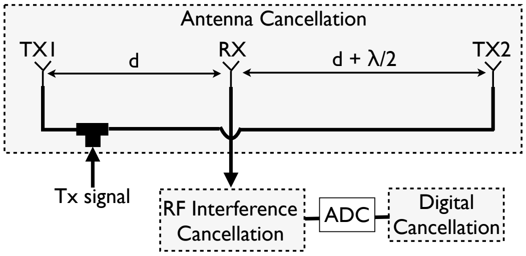

Fig. 1. Block diagram of existing full-duplex design with three cancellation techniques.

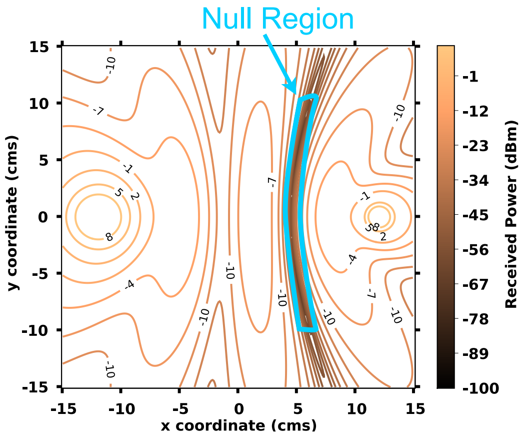

Fig. 2. Contour map showing freespace signal strength profiles for different transmit powers on two transmit antennas. Although the perfect null is present at a single point, there is a region ∼20cm in length which provides sufficient antenna cancellation for full-duplex operation.

design to allow for a high speed forward channel with a lower speed reverse channel. To this end, it extends the full duplex design to support multiple output multiple input (MIMO) operation in one direction, while supporting one or more streams in the reverse direction.

A. MIMO Overview and Challenges

To achieve their potential capacity gains, MIMO radios need to measure channel behavior at all antennas. A simple MIMO system uses only receiver side processing of data from multiple antennas. More complex systems communicate channel information back to the transmitter, enabling the transmitter to pre-process data to maximize improvements.

Typically, a receiver must wait until after a transmission, to give historic information. Although long-term channel characteristics can improve performance [5], the optimal feedback is instantaneous, something that has remained an open challenge. With the full-duplex capabilty, however, such real-time feedback is possible.

The challenge in designing a full duplex MIMO system comes from there being multiple antennas. A full duplex system needs null positions where *all* of the transmit antennas cancel. This cancellation is made even more challenging as the transmit antennas may be sending independent streams.

B. General Design

Figure 2 shows the contour map of received power with two transmit antennas using antenna cancellation [3]. There is a perfect null for the center frequency at a single point on the line between the two reverse antennas, but there is also a region of very strong destructive interference spanning approximately 20 cm from this point. Antennas placed anywhere in this region observe the 30-35 dB reduction in self-interference required for full duplex operation. MIMO antennas typically need to be spaced $\lambda/2$ apart for independent receptions. The 2.4GHz band, for example, allows up to 4 MIMO data antennas to be placed in the 20cm null region.

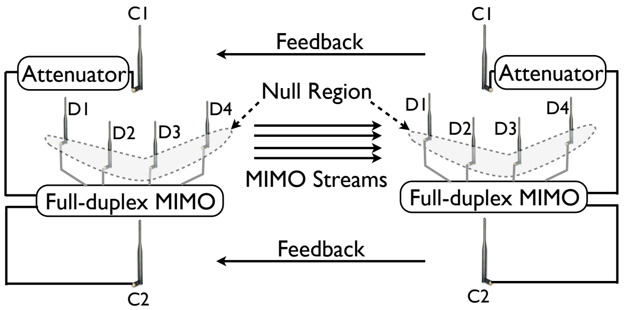

Fig. 3. Block diagram of a wireless full-duplex node with MIMO capabilities. Multiple data antennas (D1-D4) are placed in the null region of the two control antennas (C1, C2). Data antennas are used for regular data transmissions, and the control antennas can alternatively be used for real-time feedback or as a part of the data antenna array.

Figure 3 shows the design of a full duplex MIMO system with multiple data antennas and 2 control antennas. The data antennas are in the null region. At a receiver, the data antennas receive data and the control antennas transmit real-time feedback. Similarly, at a transmitter, the data antennas send data and the control antennas receive feedback. The symmetry of the antenna layout means that, just as the combined transmit signal of reverse antennas cancel at a data antenna, combining the receive signal of the control antennas cancels the signals of all of the data antennas.

Of course, the control antennas can also be used as part of the MIMO array. In that case, the data antennas and the control antennas are all in transmit or receive mode at the same time and the MIMO processing engine uses the control antennas as additional MIMO channels. This represents the trade-off between choosing the control antennas for feedback versus using them as additional MIMO channels.

It is also possible to extend this design to more duplex channels. Using a 3 dimensional arrangement of antennas can allow designing more general MIMO systems with multiple streams in both directions. One such arrangement would involve having one set of data antennas placed on a circle in a plane and another set of antennas placed on a line perpendicular to that plane and passing through the center of the circle.

The extension of the full-duplex design to MIMO systems shows that we can get a real-time in-band feedback or control channel in the reverse direction with a high speed forward direction data channel. This model relaxes the symmetry assumption presented in the existing full-duplex design, where the capacity of the forward and reverse directions are the same. Symmetrical duplexing works well for applications that use data duplexing, such as video conferencing. Asymmetric

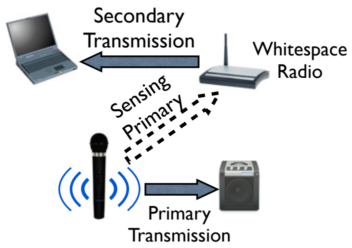

Fig. 4. Whitespace radios need to co-exist with incumbent primary transmitters. The whitespace radio senses a wireless channel before using it to avoid interfering with primary transmissions.

duplexing provides a low speed reverse channel, which may help realize the utopian view of theoretical wireless research of getting a perfect, zero latency feedback channel.

IV. Applications

The back-channel provided by full-duplex can be used for sending data or control traffic. Depending on the usage there are several applications possible with full-duplex systems. If the back-channel is used for sending data back, then such a system can be used to achieve wireless cut-through routing, mitigate hidden terminals, provide fairness in wireless local area networks (WLANs) and do real-time partial packet recovery [3]. On the other hand, if the back-channel is used for sending control traffic then such a system can be used in whitespaces, for immediate collision notification and for sending in-band channel status. The following subsections discuss these applications in detail.

A. *Opportunistic Spectrum Use (White Spaces)*

Much of the licensed spectrum is under-utilized: only 5.2% occupancy between 30MHz to 3GHz [1], [15]. For this reason, the FCC has passed a ruling in 2008 to allow for unlicensed (secondary) users to use licensed frequency bands as long as the licensed (primary) users have no perceivable interference [15]. The FCC requires that a secondary user be able to detect a primary signal that is as low as -114dBm. This requirement implies that current sophisticated solutions for a secondary user system cannot detect a primary user's presence while it's using a spectrum [1].

Figure 4 shows a setup with a secondary whitespace radio co-existing with a primary wireless device, such as a wireless mic. Without a full duplex antenna, secondary transmitters need to be very conservative in when they choose to send [12]. It is not necessarily safe for them to transmit even when the channel is sensed as vacant because they must account for the possibility that the primary might begin transmitting in the middle of their transmissions. This limits the utility extractable from vacant spectrum. By inferring the statistical properties of primary occupancy, smarter secondary strategies can be deviced, but the basic problem remains [12].

A full-duplex system can fundamentally alter this balance because the secondary transmitters can sense primary activity even while they are transmitting and quickly vacate the spectrum. This ability will allow for significantly more efficient use of the vacant spectrum.

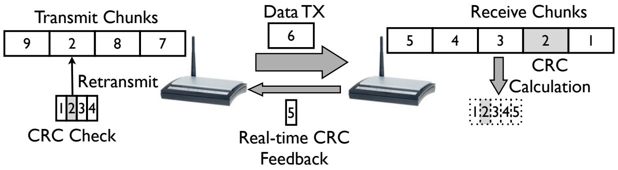

Fig. 5. Real-time error notification using CRC feedback over small blocks of data. The transmitter checks the CRC feedback for each block and retransmits blocks that have the wrong CRC. Erroneous blocks are marked grey.

Research on opportunistic spectrum use has also showed the effectiveness of cooperation among secondary nodes for more accurate sensing of primary activity [14]. This ability too is more easily engineered using a full-duplex system. A secondary receiver can use the full-duplex back channel to report periodically the channel state it observes on all the channels including the one that is currently used. This in-band shared information can be used by the secondaries to select channels with very low probability of being used by the primaries.

B. *Packet Error Notification*

The full-duplex back-channel can be used for notifying the transmitter about packet errors. This notification can be either explicit or implicit. For explicit notification, the receiver can send an abort packet back to the transmitter as soon as it encounters erroneous bits. This scheme works for both errors due to signal variation or due to collision from another node. An existing notification technique in the literature can reliably send this notfication back to the transmitter only if the notification is 18dB closer to the transmit power level [16]. With a full-duplex system, this notification can be sent even when it's 60dB lower than the transmit power level.

An implicit way to inform the transmitter of packet errors is for the receiver to simply transmit whatever it's receiving, back to the transmitter. This "mirroring" allows the transmitter to identify, if any, portions of the packet likely in error. This knowledge may be used by the transmitter to retransmit only the portions that it deems to be in error; a real-time partial packet recovery scheme.

An existing partial packet recovery scheme splits a packet into blocks and does a CRC on every block before sending the packet [9]. The receiver, after receiving the entire packet, sends the CRCs for all the blocks. This allows the transmitter to determine which blocks are in error and then send only the erroneous blocks. Figure 5 shows a similar technique implemented with full-duplex where the receiver sends the block CRCs as it's receiving data blocks. The transmitter receives these CRCs and can interleave retransmits in the middle of other data blocks. This saves the time equivalent to one packet transmission from the receiver and reduces the latency for getting retransmitted blocks.

C. *In-Band Channel Status*

In current wireless systems, a transmitter uses feedback from receivers for past transmissions to form a best guess of what the current wireless channel state may be. As wireless channels tend to be highly variable in nature, either systems

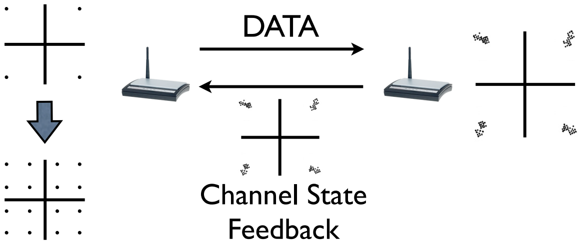

Fig. 6. Real-time feedback for rate adaptation. Receiver sends perceived constellation. Transmitter uses this feedback to adapt constellation real-time.

use conservative guesses to ensure a high packet success rate, or use higher layer mechanisms, such as retries. Essentially, not having real-time channel state information at the transmitter leads to sub-optimal wireless channel use.

The full-duplex back-channel may be used for sending real-time channel state as observed by the receiver. This real-time knowledge of the receiver's channel state is known in information theory as Channel Side Information at the Transmitter (CSIT) [6]. CSIT has been assumed to be unrealistic and used in many theoretic algorithms to show achievability of channel capacity.

With a full-duplex system, CSIT is now practical. Therefore, many capacity-achieving theoretic schemes such as waterfilling are now practical as well [6]. Waterfilling schemes provide a framework for a transmitter to change its transmit power, modulation and datarate, according to the channel state, to maximize link throughput. As an example, Figure 6 shows how a feedback from a receiver allows its transmitter to adapt modulation real-time. The receiver simply can send the received constellation periodically, while still receiving packets from the transmitter. This knowledge is useful for the transmitter to decide whether to use denser (sparser) constellation when the channel is good (bad). Current techniques allow a transmitter to change this modulation for every packet [17]. With the real-time feedback, a full-duplex transmitter can do this adaptation during a packet transmission. Specifically, this real-time adaptation can be used by wireless video streaming gadgets that, when ON, continuously send video streams to their receivers such as a TV set.

The channel state feedback being available at the transmitter is even more beneficial for MIMO systems. MIMO systems use channel state information to pick an optimal operating point between using multiple antennas for sending multiple streams, or for sending fewer streams more robustly, or with a higher data rate. A real-time feedback mechanism can thus increase the gains achieved with MIMO systems.

D. Duplexers and Spectrum Use

Many wireless systems today, such most mobile telephony protocols, are frequency division duplexing (FDD). FDD allows full duplex by dedicating pairs of channels, one for the uplink and one for the downlink. Because the two channels channels are close, radios require duplexing filters to prevent the transmitted self interference signal from saturating and desensitizing the receiver. These filters, however, are statically configured to operate over specific frequency bands and duplexing offsets. As a result, a new set of duplexing filters is required for each frequency band and each duplexing offset the radio must handle, limiting the generality and flexibility of the radio front end.

In this section, we posit that self-interference cancellation can be utilized to eliminate the need for multiple duplexers. Configuring the radio to operate on different channels of varying bandwidths and duplexing frequencies would then become a software exercise - greatly simplifying control and co-ordination in the process.

Recent work on Picasso [10] has demonstrated the feasibility of such an approach in the 2.4GHz band. Picasso combines self interference cancellation with a circulator to provide a single antenna system that can receive on one frequency band while simultaneously transmitting on an adjacent band without a duplexing filter. This has tremendous implications for spectrum allocation and spectrum planning. For instance, cellular spectrum fragmentation is likely to remain an issue globally because of short-sighted regulatory planning. The problem is compounded by the fact that even the same service providers own different fragments of spectrum in different regions, forcing mobile chipsets to accommodate a wide frequency range of operation in order to support roaming. Not only would a system like Picasso enable handset manufacturers to save costs by replacing the disparate chipsets with a single integrated solution, it would also facilitate global roaming and liberate consumers to more easily switch network operators, potentially driving improved quality of service due to increased competition between service providers.

Picasso is a general architectural solution that is not just limited to operation in the licensed bands. For example, Picasso shows it is possible to build multi-channel WiFi access points that functions with a single antenna and RF front end. The ability to simultaneously transmit and receive on different independent bands allows a multi-channel AP to leverage multiple separated spectrum fragments for different individual clients. With this ability, the AP does not have to worry about synchronizing transmissions and receptions across all the clients and can serve each one of them indepen- dently, greatly simplifying spectrum allocation and control.

Lastly, if we can build a mechanism that allows simultaneous transission and reception on arbitrary spectrum fragments with a single antenna, we can do much more than just exploit fragmented spectrum. Such a capability could also be used for radio sharing and coexistence. Portable devices today such as our smartphones must accommodate a growing list of separate ISM band protocols such as WiFi, WiFi-Direct, Zigbee, NFC, and Bluetooth. Current practice is to use a separate radio and antenna for each protocol but it is becoming increasingly difficult to find enough space to separately place all the antennas these radios would need. As the number of protocols each device must support continue to grow, it will become impractical to build separate radios and antennas to support each protocol.

Instead, a radio with a single antenna that allows simultaneous transmission and reception on arbitrary spectrum fragments could be shared among all the above protocols. WiFi would use one fragment, Zigbee would use another, and so on, saving valuable real estate on space-constrained devices. Each of the protocols could operate their own independent PHY/MAC protocols on the shared radio without interfering with each other, provided enough self interference cancellation could be achieved to ensure RF isolation amongst the protocols.

V. CHALLENGES

Previous sections argued that full-duplex has a potential that can change the fundamental assumptions of the current wireless networking paradigm. This section discusses how far the current full-duplex system is from this picture. It discusses the implication of non-ideal full-duplex system and the challenges in building a practical full-duplex radio.

A. Reliable Backchannel

The control applications of full-duplex use the reverse traffic capabilities of the system for sending real time control feedback. These applications assume a perfect duplex system where the feedback does not change the system behavior itself. However, in reality, it is unlikely to achieve such an ideal system due to the limits of engineering precision.

In the absence of perfect self-cancellation, the transmission of reverse traffic raises the noise floor at the receiver. The feedback system needs to incorporate this effect for applications like real-time channel state feedback. Even when the control channel is used purely for sensing, like for white spaces, the interference from the forward channel can lead to false positives. Also, the feedback channel is not lossless. For applications like packet error notification, losing feedback packets may lead to unexpected system behavior.

The interference pattern caused from self-interference may be quite different from conventional interference. For example, intermittent peaks can interfere with signal reception when the beginning or end of the self-interference is not correctly detected. Incorporating better interleaving in the modulating scheme can make the system more robust against interference spikes. In general, designing channel coding schemes to cope with self-interference, such that the data and feedback channel can be made more independent, is a potential challenge.

B. Real-Time Digital Cancellation

The current full-duplex design uses antenna, RF interference, and digital interference cancellation. This design is by no means optimal. For example, improved RF interference cancellation could achieve enough cancellation without using antenna placement, or a completely different cancellation technique could be used.

However, perfect cancellation cannot be achieved using only RF techniques, due to hardware imperfections such as antenna reflections and frequency selectivity as well as multipath effects. Rather, the goal of RF techniques is to decay the self-interference into the dynamic range of ADC, because the digital cancellation can remove the residual self-interference regardless of the multipath effects.

Therefore, digital cancellation is essential for achieving a full-duplex system. However, the existing full-duplex prototype uses software-defined radios to log samples and post-process them to subtract self-interference from the received samples [3]. Practical systems would require designing hardware with enough memory and processing power to implement real-time digital interference cancellation. The growing number of applications for flexible signal processing on raw digital samples, such as full-duplex, ZigZag [7] and SIC [8] should expedite the realization of such hardware.

VI. CONCLUSIONS

This paper has argued that the existing full-duplex radio can be generalized to provide real-time feedback channel to wireless networks including MIMO. It has exemplified some use cases where the feedback channel can enable co-existance within white spaces, efficient detection and recovery of packet losses, and capacity-achieving modulation using real-time CSIT. The real-time feedback channel enables rethinking how wireless coordination and control can be changed to a way which was generally assumed to be impractical. Although each of these examples can have a large impact, we believe these examples are merely the first steps towards understanding the true implications of full-duplex radio.

REFERENCES

[1] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with wi-fi like connectivity. *SIGCOMM Comput. Commun. Rev.*, 39(4):27–38, 2009.

[2] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti. Achieving single channel, full duplex wireless. In *MOBICOM*, 2010.

[3] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti. Achieving single channel, full duplex wireless communication. Technical Report SING-10-00, 2010.

[4] M. Duarte, C. Dick, and A. Sabharwal. Experiment-driven characterization of full-duplex wireless systems. In *IEEE Transactions on Wireless Communication*, 2012.

[5] D. Gesbert, M. Shafi, D. shan Shiu, P. J. Smith, A. Naguib, and S. Member. From theory to practice: an overview of mimo space-time coded wireless systems. *IEEE Journal on Selected Areas in Communications*, 21:281–302, 2003.

[6] A. Goldsmith. *Wireless Communications*. Cambridge Press, 2005.

[7] S. Gollakota and D. Katabi. ZigZag decoding: combating hidden terminals in wireless networks. In *SIGCOMM*. ACM, 2008.

[8] D. Halperin, T. Anderson, and D. Wetherall. Taking the sting out of carrier sense: interference cancellation for wireless lans. In *MobiCom*. ACM, 2008.

[9] B. Han, A. Schulman, F. Gringoli, N. Spring, B. Bhattacharjee, L. Nava, L. Ji, S. Lee, and R. Miller. Maranello: Practical partial packet recovery for 802.11. In *NSDI*, 2010.

[10] S. Hong, J. Mehlman, and S. Katti. Picasso: Flexible rf and spectrum slicing. In *SIGCOMM*, 2012.

[11] Y. Hua, P. Liang, Y. Ma, A. C. Cirik, and Q. Gao. A method for broadband full-duplex mimo radio. In *IEEE Signal Processing Letters*, 2012.

[12] S. Huang, X. Liu, and Z. Ding. Optimal transmission strategies for dynamic spectrum access in cognitive radio networks. *IEEE Transactions on Mobile Computing*, 8, 2009.

[13] M. Jain, J. I. Choi, T. Kim, D. Bharadia, K. Srinivasan, S. Seth, P. Levis, S. Katti, and P. Sinha. Practical, real time, full duplex wireless. In *MOBICOM*, 2011.

[14] S. M. Mishra, A. Sahai, and R. W. Brodersen. Cooperative sensing among cognitive radios. In *In Proc. of the IEEE International Conference on Communications (ICC*, 2006.

[15] A. Sahai, S. M. Mishra, and R. T. Spectrum sensing: Fundamental limits. *draft chapter for a Springer Book: Cognitive Radios: System Design Perspective*, June 2009.

[16] S. Sen, R. R. Choudhury, N. Santhapuri, and S. Nelakuditi. Moving away from collision avoidance: Towards collision detection in wireless networks. In *Hotnets*, 2009.

[17] S. Sen, N. Santhapuri, R. R. Choudhury, and S. Nelakuditi. Accurate: Constellation based rate estimation in wireless networks. In *NSDI*, 2010.