

White Space Networking with Wi-Fi like Connectivity

Paramvir Bahl[†], Ranveer Chandra[†], Thomas Moscibroda[†], Rohan Murty[‡], Matt Welsh[‡]

[†]Microsoft Research, Redmond, WA
{bahl, ranveer, moscitho}@microsoft.com

[‡]Harvard University, Cambridge, MA
{rohan, mdw}@eecs.harvard.edu

ABSTRACT

Networking over UHF white spaces is fundamentally different from conventional Wi-Fi along three axes: spatial variation, temporal variation, and fragmentation of the UHF spectrum. Each of these differences gives rise to new challenges for implementing a wireless network in this band. We present the design and implementation of WhiteFi, the first Wi-Fi like system constructed on top of UHF white spaces. WhiteFi incorporates a new adaptive spectrum assignment algorithm to handle spectrum variation and fragmentation, and proposes a low overhead protocol to handle temporal variation. WhiteFi builds on a simple technique, called SIFT, that reduces the time to detect transmissions in variable channel width systems by analyzing raw signals in the time domain. We provide an extensive system evaluation in terms of a prototype implementation and detailed experimental and simulation results.

Categories and Subject Descriptors:

C.2.1 [Computer-Communication Network]: Wireless communication

General Terms: Algorithms, Design, Experimentation

Keywords: white spaces, channel width, Wi-Fi, dynamic spectrum access, cognitive radios

1. INTRODUCTION

The unused portions of the UHF spectrum, popularly referred to as “white spaces”, represent a new frontier for wireless networks, offering the potential for substantial bandwidth and long transmission ranges. These white spaces include, but are not limited to, 180 MHz of available bandwidth from channel 21 (512 MHz) to 51 (698 MHz), with the exception of channel 37. On November 4, 2008, the FCC issued a historic ruling permitting the use of unlicensed devices in these white spaces [10]. In its ruling the FCC imposed an important requirement that white space wireless devices must not interfere with incumbents, including TV broadcasts and wireless microphone transmissions. This landmark ruling was a result of extensive tests performed by the FCC on white space hardware prototypes that were submitted by Adaptrum, Microsoft, Phillips and Motorola. These prototypes demonstrated feasible solutions for an accurate and agile sensing of incumbent signals [9].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'09, August 17–21, 2009, Barcelona, Spain.

Copyright 2009 ACM 978-1-60558-594-9/09/08 ...\$10.00.

Most of the prior research in UHF white spaces has focused on accurately detecting the presence of incumbent RF signals [14, 17, 18]. Recently, researchers have mentioned that they are beginning to look at the problem of establishing a wireless link between white space devices [8, 12]. Our research pushes the state-of-art to the next level by going beyond a single link. We identify the challenges of forming a UHF white space network and show how to overcome them by presenting techniques, algorithms, and protocols backed up by extensive evaluation over a prototype network as well as in simulations. We focus primarily on the problem of setting up a Wi-Fi like network consisting of an Access Point (AP) with multiple associated clients. We leave the case of evaluating multiple APs with multiple clients as follow-on work. Our solutions are complementary to the ongoing work in the IEEE 802.22 Working Group [1], as we discuss in Section 7.

To appreciate the networking problem, it is important to understand the differences between white spaces and the popular ISM bands where Wi-Fi devices operate. First, in both bands there is spatial variation in spectrum availability, but the impact of this variation is higher in white spaces than in ISM bands. This is because the FCC ruling requires non-interference with wireless transmissions of primary users (incumbents) (Section 2.1). Second, since the incumbents can operate in any portion of the white spaces, the network must be designed to handle spectrum fragmentation, with the possibility of each fragment being of different width. A UHF channel is narrow (6 MHz wide in the US), and prior research has shown that aggregating contiguous channels improves throughput [15, 21]. Consequently, the network must support variable width channels (Section 2.2). Third, RF transmissions in white spaces are subject to temporal variations because wireless microphones can become active at any time without warning. Our experiments show that even a single packet transmission causes audible interference during wireless microphone transmissions. Consequently, both the AP and its clients must disconnect and then rapidly reconnect using a different available channel (Section 2.3).

We have built WhiteFi, a UHF white space wireless network that adaptively configures itself to operate in the most efficient part of the available white spaces. In the following sections we describe three major innovations that allowed us to overcome the challenges in networking white space devices. Briefly, our contributions are:

- A novel spectrum assignment algorithm for managing variable bandwidth communications. Our algorithm is unique in the way it addresses the dual challenges of spatial variation of available spectrum and spectrum fragmentation. We introduce a new metric that leverages the available airtime measurements from each available UHF channel to predict the available airtime when using multiple channels.




Figure 9: Time to discover one AP at various locations.

quency and width to discover the APs. In this section, we consider two scenarios. First, we show the benefit of our algorithms as a function of contiguous width. Then, we evaluate the benefits in realistic settings, i.e., in metropolitan, suburban and rural settings.

Methodology: We set up two KNOWS devices as before, and configured one as an AP and the other as a client. In the beginning of the experiment, the AP started to beacon on a randomly chosen UHF channel and channel width. We then measured the time for the client to discover the AP using L-SIFT, J-SIFT and the non-SIFT baseline. Depending on the scenario, we artificially specified the spectrum at the AP and the client. The AP did not beacon on any of the occupied channels, and the client did not scan these channels for an AP.

Contiguous Channels: In this experiment, we set the spectrum map to have only one available fragment. We varied the number of UHF channels in the fragment from 1 to 30, since 30 is the total number UHF channels that are available to portable devices. In Figure 8, we plot the total time taken by L-SIFT and J-SIFT to discover the AP as a fraction of the total time taken by the non-SIFT baseline. When there is only one available UHF channel, the time taken by all the algorithms is the same. However, when we increase the width of the available fragment of spectrum, L-SIFT and J-SIFT perform much better than the baseline. As expected, L-SIFT outperforms J-SIFT initially (for narrow white-spaces) since it does not require the “endgame” of trying to find the proper placing of the AP channel. On the other hand, as exactly predicted by our analysis in Section 4.2, J-SIFT becomes more efficient for white spaces spanning more than 10 UHF channels (60 MHz).

Realistic Settings: We also measured the time to discover an AP in metropolitan, suburban and rural areas in the US. We used the methodology described for Figure 2 to obtain the spectrum maps post-DTV transition. We randomly placed the AP on an available channel and width and repeated the experiment 10 times for every locale. As shown in Figure 9, in metro areas, where there are fewer contiguous channels, J-SIFT is 34% faster than the baseline. In rural areas (more contiguous channels), we see that J-SIFT can discover APs in less than one-third the time taken by the baseline algorithm.

5.3 Handling Disconnections

We now quantify the time taken by WhiteFi to reconnect disconnected clients. We setup a client and an AP and started a data transfer between them. Then we switched on a wireless microphone near the client. This causes the client to disconnect, and it starts chirping on the backup channel. In our experimental setup, the AP switched to the backup channel once every 3 seconds, and picks up the chirp in at most 3 seconds. Immediately, the AP uses the spectrum assignment algorithm to determine the best available channel to operate on, and the system is operational again after a lag of at most 4 seconds.

5.4 Spectrum Assignment

We now evaluate WhiteFi’s spectrum assignment algorithm. For a detailed understanding of our algorithm, and to evaluate it under varied settings, we decided to use the QualNet simulator [3]. The need to use the simulator arose for two reasons. First, we were constrained by having a limited number of prototype devices, and second, we did not have an FCC license to transmit packets in the TV bands. Therefore, we evaluated our system (spectrum assignment, discovery and disconnection protocols) in a limited setting – on a testbed spanning one floor in our building, and a maximum transmit power of 1 mW.

Modifications to QualNet: We modified QualNet to support variable channel widths by appropriately scaling the OFDM symbol period, and various MAC layer parameters that were described in [15]. We also adjusted the channel noise levels based on the channel width. Furthermore, at every node, we explicitly drop packets that were sent at a different channel width. To ensure that a node appropriately contends with packets that are sent on overlapping channels of different widths, we modified the carrier sensing mechanism in QualNet such that a node spanning multiple UHF channels will transmit a packet only if no carrier is sensed on any of those channels. We also modified QualNet to support fragmented spectrum. Every node reads its initial spectrum map from a configuration file.

5.4.1 Simulation Results

We study the performance of WhiteFi’s spectrum assignment algorithm under various settings. First, we microbenchmark the *MCham* metric, and show that it is a good estimate of the expected throughput on a channel. Then, using large scale experiments, we show that WhiteFi performs reasonably well under: (i) varying amounts of background traffic on the channels, (ii) large amounts of spatial variation in spectrum availability, and (iii) when there is a lot of churn in background traffic. In all these experiments, WhiteFi performs nearly as well as an optimal algorithm. In the process, we also show the need for WhiteFi to adapt both the center frequency and the channel width.

Microbenchmark Setup: To verify that *MCham* correctly predicts the channel that will lead to the best throughput, we simulate a spectrum fragment of 5 adjacent UHF channels (26-30), each having one background client/AP-pair. There is one AP with one associated client, transmitting a link-saturating UDP flow. We vary the traffic intensity of the background nodes (from 0 to 50 ms inter-packet delay) and measure the effect on the *MCham* metric and client throughput when transmitting on the 5, 10, and 20 MHz channels centered at channel 28.

Accuracy of the MCham Metric: The results in Figure 10 show that the *MCham* metric accurately predicts which channel achieves the highest throughput for any given background intensity. For example, selecting a 20 MHz channel achieves best throughput until a background traffic intensity of roughly 18 ms inter-packet delay. Similarly, the *MCham* metric predicts that roughly at this level of background traffic, 10 MHz and 20 MHz become equally good, and the narrower 10 MHz channel surpasses the wider channel thereafter. Similarly, at about 24 ms inter-packet delay, 5 MHz starts achieving the highest throughput, which is accurately predicted by the *MCham* metric. We can conclude that the *MCham* metric yields a reasonably accurate prediction of which channel width will result in the highest throughput given a certain level of background traffic.

Setup of large-scale simulations: To better understand the behavior of WhiteFi in large-scale settings, the next three simulations consider the following basic setup. We place one AP in the middle of an area, and randomly distribute clients as well as background




Figure 10: $MCham$ value and resulting throughput of a 5, 10, and 20 MHz channel as a function of background traffic intensity. The $MCham$ metric accurately predicts which channel achieves highest throughput.

AP/client-pairs within transmission range of this AP (background clients are always deployed within transmission range of their respective background AP). The AP and clients are backlogged and transmit UDP flows (up- and downstream). Background nodes transmit constant-bit-rate (CBR) traffic at a pre-specified intensity. All experiments are repeated 5 times with different random placements of nodes, and results are averaged.

An underlying spectrum map is shared across all clients (except in the experiment in which we focus on the impact of spatial variation). Specifically, the spectrum map is taken from our real measurements in Section 2. There are 17 free UHF channels, and the widest contiguous white space is 36 MHz, i.e., there are multiple possibilities of selecting even 20 MHz wide channels for the AP.

In all experiments, we measure the per-client throughput of clients/APs. We consider the following baseline algorithms for comparison with WhiteFi. OPT 5 MHz denotes the throughput achieved when statically picking the best (across all non-incumbent) UHF channels. Similarly, OPT 10 MHz and OPT 20 MHz are the algorithms that statically pick the best possible 10 and 20 MHz channel, respectively. Finally, OPT is an ideal, omniscient algorithm that for every experiment run picks the channel with maximum throughput. The goal of the WhiteFi spectrum assignment algorithm is to approach OPT as closely as possible.

Impact of Background Traffic: Figure 11 shows how WhiteFi reacts to varying degrees of background traffic. Specifically, there are X background AP/client-pairs in the system, each being randomly assigned to one of the free UHF channels, and each sending at a packet interval delay of 30 ms.

The figure shows that WhiteFi achieves close to optimal performance for varying degree of background traffic. With little or no background traffic, WhiteFi performs as well as picking the widest available channel (OPT 20 MHz), which is optimal. As the traffic increases, the throughput achieved by OPT 20 MHz drops, and OPT 10 MHz becomes better (at about 10 background AP/client-pairs). Even at this point WhiteFi performs near-optimally, which shows that WhiteFi adaptively switches to narrower channels as needed. In fact, our evaluation shows that WhiteFi is always within 14% of the optimal value throughput OPT.




Figure 11: Impact of background traffic on throughput.




Figure 12: Impact of spatial variation on throughput.

An important observation is that due to fragmentation and background traffic, there is no single best center frequency and channel width that should be used in UHF white spaces. WhiteFi is capable of adjusting to the appropriate width and selects a near-optimal channel.

Impact of Spatial Variation: Figure 12 shows the impact of spatial variation on per-client throughput. In this experiment, there are 10 clients connected the AP, and one background client/AP-pair per UHF channel, transmitting at CBR with 30 ms inter-packet delay. Spatial variation is modeled as follows. Each client and the AP start with a common spectrum map. Then, for each client (and AP) and for each UHF channel i , we randomly flip the entry u_i with probability P . In the experiment, we vary P from 0 (no spatial variation) to 0.14 (large spatial variation).

It can be seen in the figure, spatial variation reduces achievable aggregate throughput. Because the AP needs to select a channel that is free at *all* clients, no contiguous free spectrum parts remain available for $P > 0.1$, and hence, the aggregate throughput reduces to the throughput of a single UHF channel (5 MHz). For low spatial variation, the throughput is much higher when selecting a 20 MHz wide (e.g. at $P = 0.01$) or a 10 MHz channel (e.g. at $P = 0.05$). Generally, the figure highlights the need for *adaptive* channel width in UHF white spaces: no single channel width (OPT 20 MHz, OPT 10 MHz, OPT 5 MHz) achieves close-to-optimal throughput in all cases. On the other hand, WhiteFi is near-optimal in all cases.

Impact of Churn: Finally, we want to understand the impact of churn (in terms of background traffic) on the throughput achieved by WhiteFi and the various baseline algorithms. There are a total of 34 background AP/client-pairs, two per free UHF channel. In order to model churn, we model background nodes using a simple discrete Markov chain with two states (A=active, P=passive). A background node in the active state transmits CBR traffic with 60 ms inter-packet delay. A node in the passive state does not transmit. We simulate this setting for various state transition probabilities, selecting them to cover the entire range of (1) likelihood of being in either state and (2) average state duration (see x-axis in Figure 13). The extreme cases are (i) all nodes are always in state P, (ii) nodes are in each state with equal likelihood and they remain in their current state for an average of 30 seconds, and (iii) all nodes are always in state A.




Figure 13: Impact of churn on throughput.




Figure 14: Experimental validation of WhiteFi’s spectrum assignment algorithm on a testbed with variable background traffic. Top figure shows the $MCham$ metric for each of the three channel widths. Bottom figure shows the throughput (averaged over 5 sec windows) for WhiteFi and OPT.

Figure 13 shows that WhiteFi performs near-optimally for varying degree of churn. For low churn and little background traffic, WhiteFi selects the widest channel. For high churn (e.g., state duration 45 seconds and passive probability $1/3$), always picking the widest channel (OPT 20 MHz) becomes the worst performing algorithm. Instead, WhiteFi is better than any static channel width choice. In fact, WhiteFi even outperforms OPT. In this experiment, this is possible because OPT is the optimal *static* channel selection throughout the entire execution of the simulation. Instead, WhiteFi is adaptive and can adjust to the current values of background traffic, changing its channel accordingly.

5.4.2 Results from our Prototype

To demonstrate the adaptability of WhiteFi’s spectrum assignment algorithm, we set up an experiment with an AP and a client in our building, which is Building 5 in Figure 1. The spectrum map of our building has the following free UHF channels: 26 to 30, 33 to 35, 39 and 48. Therefore, we have fragments of size 20 MHz, 10 MHz and two channels of 5 MHz to form a network.

Every client and AP using WhiteFi spends 1 second on every UHF channel to determine the airtime utilization using SIFT, as described in Section 5.1. All nodes feed their airtime to the AP, which computes the $MCham$ metric and decides on the channel to use for the network. We present the throughput of our system with time, and the corresponding $MCham$ value on the different spectrum chunks in Figure 14.

Initially, when there is no background traffic, the AP and client operate on the 20 MHz spectrum chunk between channels 26 and 30. Then at time 50 seconds, we introduce background traffic on channels 26 through 29. Correspondingly, the value of the $MCham$ metric for the 20 MHz fragment drops sharply, and the AP and its clients move to the 10 MHz spectrum fragment. As shown in the figure, this is also the fragment that has the best throughput. Then at time 100 seconds, we introduce background traffic on channels 33 and 34, and as before the value of the 10 MHz channel’s $MCham$ metric drops, and the system switches to channel 39 (any 5 MHz chunk could have been chosen). Then at times 150 and 200 seconds, we remove the background interference from channels 33 and 34, and from channels 26 through 29, respectively. Correspondingly, WhiteFi switches to the fragment with the best $MCham$ value, i.e. to the 10 MHz fragment at 150 seconds, and to the 20 MHz fragment at 200 seconds. We conclude from the above experiments that WhiteFi adaptively operates on the best part of the spectrum.

6. DISCUSSION AND FUTURE WORK

White space networking provides a unique opportunity for clean-slate network design, owing to the lack of existing standards. Our decision to build the WhiteFi prototype with a Wi-Fi card was motivated by several factors. Wi-Fi is a mature, well-understood technology that is inexpensive and easily available. Several wireless card vendors we have spoken with are considering pushing some version of Wi-Fi to the IEEE standards body for white space networking. Additionally, Wi-Fi enabled us to build a prototype quickly and focus on some of the higher layer issues that are somewhat agnostic to the existing physical and MAC protocols. However, we do realize that alternative designs are possible and might be used in future networks. We discuss a few of these below.

WhiteFi leverages the technique described in [15], which requires the AP and its clients to operate over the same contiguous chunk of spectrum. An alternative technique might use a PHY layer that operates over non-contiguous spectrum chunks. The AP can then operate over the entire bandwidth, decoding signals from the different clients who may be using different OFDM subcarriers. For AP-to-client communications, the PHY layer could either suppress or send a null signal on the subcarrier that the primary user is using [21]. In theory this is a reasonable idea but it poses two practical problems. First, leakage from adjacent subcarriers causes interference to the primary user. To avoid this interference, we would require a highly accurate bandpass filter of appropriate bandwidth but to the best of our knowledge researchers are still working on developing such sharp bandpass filters. Second, and more importantly, sending data over different subcarriers to an AP is difficult to implement for uplink traffic. We are not aware of any system that can decode packets sent simultaneously from multiple clients over non-overlapping subcarriers. This is an active research area and we are investigating the practicality of such a system.

Another issue is our choice of CSMA/CA, the medium access control (MAC) protocol for Wi-Fi, in WhiteFi. The research literature has several interesting proposals for MAC protocols, which can be broadly categorized under Listen Before Transmit (LBT) and Time Division Multiple Access (TDMA). Observing what is happening in the ISM bands we made the decision that WhiteFi must be able to co-exist with other unlicensed devices. The success of LBT protocols (e.g., Wi-Fi) in the ISM bands made it a natural choice for white space networking. We also believe that an alternative TDMA like MAC (e.g., Bluetooth) will not perform well in white spaces without significant modifications. Local interference from wireless microphones around the client or the AP would im-

packet scheduling and lead to poor performance. Furthermore, in UHF white spaces the clients and AP may be over a mile away, further aggravating the scheduling problem. Additional research is needed to understand these issues and is out of scope for this paper. Our initial results show that CSMA/CA is a reasonable choice for white space networking.

Prior work has proposed the use of control channels to reserve bandwidth and spectrum [12, 24]. While there are advantages to a control channel design, we believe that control channels can be compromised, thus bringing down the network. Also, control channel based solutions are prone to the range-mismatch problem [24]. We overcome these problems by not using a dedicated control channel. WhiteFi uses a backup channel in the white spaces (instead of 900 MHz spectrum as proposed by CMAC [24]) thereby avoiding the range mismatch problem. Also, WhiteFi does not use a static control channel. It dynamically adapts the backup channel to operate on spectrum that is not occupied by a primary user.

7. RELATED WORK

Prior work has mostly focussed on the problem of opportunistically forming a single link over UHF white spaces [8, 12]. This involves accurate sensing of the spectrum [14, 17, 18], reliable identification of incumbents, and radio agility on detecting a primary. However, to the best of our knowledge, no prior work has studied the problems of forming a Wi-Fi like network over white spaces.

WhiteFi builds upon our prior work on KNOWS [24], which uses a similar hardware platform and proposes a control channel based MAC protocol for ad hoc networks over white spaces. WhiteFi looks at the problem of forming an AP based network while reusing the Wi-Fi MAC and without using a control channel.

A complementary effort to WhiteFi is the IEEE 802.22 [1] working group's proposal for WRANs (Wireless Regional Area Networks) over UHF white spaces. It is intended to provide wireless broadband access to rural areas and neighborhoods. In contrast, WhiteFi considers a usage model similar to Wi-Fi, with one AP providing coverage to several possibly mobile users. Despite the difference in the scenarios, the techniques developed by WhiteFi, for disconnection, discovery and spectrum assignment, are also applicable in WRANs. For example, the 802.22 draft includes support for variable widths, although it does not specify how to use it.

A recent technology that enables unlicensed devices to co-exist with licensed users is SWIFT [21]. SWIFT pokes the primary user to learn about its presence. Unfortunately, this technology cannot be used over white spaces because the FCC does not allow "testing" the presence of an incumbent by "poking" at it with a transmission. Also, the incumbents of UHF white spaces do not back off.

8. CONCLUSIONS

In this paper, we have presented the design and implementation of WhiteFi, the first white space Wi-Fi like wireless network. We moved beyond the current state-of-art that considers a single link to building a real network with multiple links. In building WhiteFi we identified and described several unique challenges in operating a white space network and showed with extensive experiments how white space networks differ from ISM band Wi-Fi networks. WhiteFi contributes a new spectrum assignment algorithm that solves the dual challenges of spatial variation of available spectrum and spectrum fragmentation. We further described a new mechanism that quickly discovers APs operating anywhere in the 180 MHz white space, using any arbitrary channel width. We also described a new technique for handling disconnections where clients signal to the AP without interfering with ongoing wireless microphone transmissions. Underlying our solutions is

a new application of a signal recognition technique called SIFT, which quickly analyzes packets in the time domain, allowing fast AP discovery and managing disconnections due to temporal variations. We demonstrated WhiteFi in the context of our custom built prototype UHF hardware and QualNet simulations. As part of ongoing work, we are deploying WhiteFi over a campus wide white space network.

Acknowledgements

We would like to thank our shepherd, Brad Karp, for his help with the final version of the paper. His detailed comments were immensely helpful in improving the presentation of the paper. We are also grateful to the anonymous SIGCOMM reviewers for their insightful comments.

9. REFERENCES

- [1] IEEE 802.22 Working Group on WRANs, <http://www.ieee802.org/22/>.
- [2] In private conversations with Adaptrum.
- [3] Qualnet 4.5 simulator, <http://www.scalable-networks.com>.
- [4] TV Fool, <http://www.tvfool.com/>.
- [5] Universal Software Radio Peripheral, <http://www.ettus.com/>.
- [6] USRP FAQ, <http://www.gnuradio.org/trac/wiki/UsrpFAQ/Gen>.
- [7] FCC press release, FCC Acts to Expedite DTV Transition and Clarify DTV Build-Out Rules. November 2001.
- [8] Demonstration of a Prototype Dynamic Spectrum Access System, Philips Research. In *DySPAN demo session*, 2008.
- [9] FCC press release, Evaluation of the Performance of Prototype TV-Band White Space Devices. November 2008.
- [10] FCC press release, FCC Adopts Rules for Unlicensed Use of Television White Spaces. November 2008.
- [11] N. Ahmed, S. Banerjee, S. Keshav, A. Mishra, K. Papagiannaki, and V. Shrivastava. Interference Mitigation in Wireless LANs using Speculative Scheduling. In *MobiCom*, 2007.
- [12] R. E. D. Borth and B. Oberlie. Considerations for Successful Cognitive Radio Systems in US TV White Space. In *DySpan*, 2008.
- [13] G. Buchwald, S. Kuffner, M. Brown, and E. C. L. Ecklund. The Design and Operation of the IEEE 802.22.1 Disabling Beacon for the Protection of TV Whitespace Incumbents. In *DySpan*, 2008.
- [14] D. Cabric, A. Tkachenko, and R. W. Borden. Experimental Study of Spectrum Sensing based on Energy Detection and Network Cooperation. In *ACM 1st International Workshop on Technology and Policy for Accessing Spectrum (TAPAS)*, 2006.
- [15] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl. A Case for Adapting Channel Width in Wireless Networks. In *SIGCOMM*, 2008.
- [16] S. Inc. SLX series product manual.
- [17] H. Kim and K. G. Shin. Fast Discovery of Spectrum Opportunities in Cognitive Radio Networks. In *DySPAN*, 2008.
- [18] H. Kim and K. G. Shin. In-band Spectrum Sensing in Cognitive Radio Networks: Energy Detection or Feature Detection? In *MOBICOM*, 2008.
- [19] R. Murty, J. Padhye, R. Chandra, A. Wolman, and B. Zill. Designing High-Performance Enterprise Wireless Networks. In *NSDI*, San Francisco, CA, April 2008.
- [20] S. Narlanka, R. Chandra, P. Bahl, and I. Ferrell. A Hardware Platform for Utilizing the TV Bands with a Wi-Fi Radio. In *IEEE LANMAN*, June 2007.
- [21] H. Rahul, N. Kushman, D. Katabi, C. Sodini, and F. Edalat. Learning to Share: Narrowband-Friendly Wideband Wireless Networks. In *SIGCOMM*, 2008.
- [22] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra. Perceptual Evaluation of Speech Quality (PESQ)-A New Method for Speech Quality Assessment of Telephone Networks and Codecs. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2001.
- [23] K. Tan, J. Zhang, J. Fang, H. Liu, Y. Ye, S. Wang, Y. Zhang, H. Wu, W. Wang, and G. M. Volker. SORA: High Performance Software Radio Using General Purpose Multi-core Processors. In *NSDI*, 2009.
- [24] Y. Yuan, P. Bahl, R. Chandra, P. A. Chou, I. Ferrell, T. Moscibroda, S. Narlanka, and Y. Wu. KNOWS: Kognitiv Networking Over White Spaces. In *DySPAN*, 2007.