

A Dual Technology Femto Cell Architecture for Robust Communication using Whitespaces

Sayandeep Sen¹ Tan Zhang¹ Milind M. Buddhikot² Suman Banerjee¹ Dragan Samardzija² Susan Walker²

¹University of Wisconsin Madison, USA

{sdsen, tzhang, suman}@cs.wisc.edu

²Alcatel-Lucent Bell Labs, USA

{milind.buddhikot, dragan.samardzija, susan.walker}@alcatel-lucent.com

Abstract—We propose *WhiteCell* — a dual technology femto cell architecture for wide-area wireless users that addresses some of the critical challenges being faced by the cellular industry today in better meeting surging demands. *WhiteCell* extends the traditional notion of a femto cell in which small, low-power, short-range access points are placed in homes to improve coverage, connectivity and spatial re-use in the cellular operator’s own frequency and technology. In *WhiteCell*, each indoor access point is equipped with the ability to communicate over two separate blocks of spectrum — the cellular operator’s own spectrum using the same technology as before, and the recently released swath of TV whitespace spectrum that allows opportunistic use under specific FCC guidelines in the US. The two spectrum blocks and their corresponding technologies complement each other very well. While whitespace spectrum allows us to add significant capacity to the otherwise constrained cellular spectrum, the cellular frequencies allows the system to support some minimal expectation of performance guarantee that whitespace alone cannot, due to license exclusivity. While this extension is conceptually simple, it provides dramatic performance gains for both the cellular operator and the end users who are putting increasing demands on the limited cellular spectrum. This paper describes the overall *WhiteCell* architecture, a system implementation, and various challenges addressed in efficiently utilizing whitespace spectrum including a collaborative approach in spectrum sensing, as well as in efficiently transitioning traffic across this dual technology structure. In addition, the paper demonstrates the significant performance advantages of the architecture through detailed evaluation of our *WhiteCell* prototype.

I. INTRODUCTION

Spectrum is the scarce resource in wireless communication systems and is considered especially so in wide-area cellular data communication networks. In terms of numbers, it is anticipated that cellular data traffic volume will reach several exabytes per month by 2014 (1 exabyte = 1 million terabytes), roughly equaling the traffic volume in the entire global Internet back in 2006 [1, 12] putting even more stress on the limited available spectrum.

Other than improving the underlying technologies for wireless communication, there are two popular approaches to address this challenge: (i) enable greater spatial re-use through the deployment of many low-power and short range “mini base stations” called femto cells or small cells, and (ii) perform cellular traffic offloading to other unlicensed spectrum bands, e.g., using WiFi APs [20]. In this paper, we advocate an approach that combines both of the above to achieve significantly superior performance than either. More specifically, we

define an architecture called *WhiteCell* in which we propose the deployment of low-power and short-range mini base stations, each equipped with radio interfaces operating using two complementary technologies — the cellular operator’s own licensed spectrum technology, and technology based on the newly available TV whitespace spectrum block (Figure 1). The mobile clients that connect to these base stations use a “bonded” wireless link spanning both spectrum blocks and technologies, in order to be able to communicate with the *WhiteCell* base stations. The base stations maintain backhaul data connectivity to the Internet through common broadband services available in the indoor environment, e.g., DSL, cable modem, fiber, etc.

A. Role of dual technology design in *WhiteCell* and uniqueness

In typical multi-interface and multi-technology systems, the primary goal has been to utilize all such spectrum simultaneously to gain the advantage of the aggregate bandwidth. While this possibility exists in the *WhiteCell* architecture, this is not the primary reason for the multi-interface design. In general, we prefer to use *only* the whitespace interface for all communication as much as possible. However, as the FCC ruling in the US mandates, whitespace spectrum can only be used opportunistically. A whitespace transmitter is required to vacate its channel of operation as soon as the “primary” user of the channel (e.g., a TV broadcaster or a wireless microphone) re-appears. At that instant, the whitespace transmitter needs to find an alternative whitespace channel to operate on. As prior work (*WhiteFi* [6]) as well as our experiments have shown, finding an efficient alternative whitespace channel may either take order of tens of seconds (sometimes as much as 60 to 120 seconds in presence of high interference) or the transmitter may need to settle on a significantly sub-optimal alternative. Such delays or loss in performance is typically quite disruptive to end users.

In *WhiteCell*, we therefore use the existing licensed spectrum technology of the operator to serve as a *stop-gap communication link* between the base station and mobile clients, *until a new whitespace channel is identified* and made ready for continued communication. This technology, being licensed, can provide a minimum performance guarantee to mobile clients connected to the *WhiteCell* base station, while a better

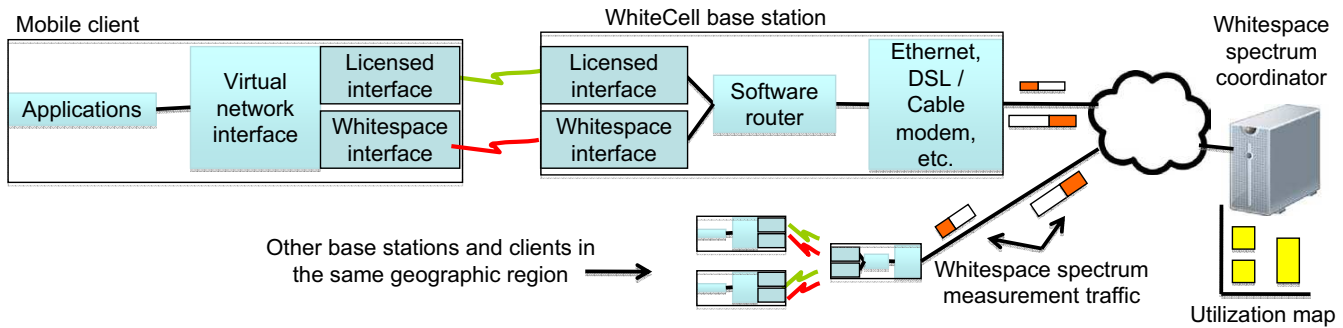


Fig. 1. WhiteCell architecture overview and components.

whitespace channel is identified for communication.

In addition to the above use, we also use the licensed interface in WhiteCell as a control channel between the mobile clients and the base station. Finally, given that two blocks of spectrum are available through this design, we also allow the option of using both the licensed and the whitespace interfaces simultaneously for striping multiple client flows across them, if the operator so chooses.

A number of recent efforts, such as WhiteFi [6], Deb. et. al [14], and Zheng et. al. [19, 25, 30], have designed general-purpose, whitespace only communication solutions addressing some channel access and contention resolution issues. In addition, the concept of traffic offloading using alternative technologies (most commonly using WiFi [20]) is also quite well known. Our design of WhiteCell is a stylized version of these above techniques, customized to meet the specific requirements of a high bandwidth femto cell architecture. For instance, our whitespace communication solution in WhiteCell includes specific channel sensing and spectrum coordination methods that are possible due to the likely presence of multiple WhiteCell base stations in a certain geographic area. Similarly, our traffic offloading is typically from whitespace to licensed spectrum and back. While traditional traffic offloading is client-initiated, in our case such offloading is base station initiated. Most importantly, our dependence of whitespace technology as the main communication channel and the use of licensed technology as a “temporary bridge over troubled whitespace waters,” makes its design relatively unique among recent related work.

While researchers are investigating the benefits of utilizing whitespaces in outdoor settings, in this work, we show that it can help improve the performance of the indoor femto cells. Specifically, we highlight the unique challenges in utilizing the TV whitespace in indoor femto cell design such as accurately identifying available spectrum after attenuation of the building and high temporal variation of spectrum availability due to wireless microphones.

B. Rationale compared to some alternative designs

In designing WhiteCell, we considered and decided against some other alternatives that are also reasonable design choices, some of which are in use today. We comment on them next.

- *Licensed-only femto cells*: The femto cell architecture popular today uses a single licensed radio interface to com-

municate between femto base stations and mobile client, to provide additional coverage into hard to reach locations and to provide improved throughput to hotspots. In this paper, we seek much greater improvements that are scalable with spectrum availability. Specifically, we chose a dual interface design as it not only allows us to leverage free whitespace spectrum (when available) but also provides an insurance of un-interrupted operation using the licensed spectrum bands in case of disruption or degradation in quality of the free whitespace spectrum.

- *Unlicensed-only femto cells*: It was possible to design a WhiteCell base station to only have a single whitespace interface. However, we decided to add a licensed interface to serve as the stop-gap and a back-up communication link with some performance guarantee during periods of disruption on an existing whitespace channel due to return of a primary, or any other such reason.

- *Why whitespace and not WiFi as the unlicensed technology*: In principle, our proposed architecture could be designed to use the existing and popular WiFi technology as the unlicensed technology, instead of our current choice of whitespace technology. In fact, it is possible to use both WiFi and whitespace technologies simultaneously as unlicensed communication options. In WhiteCell, we focused on whitespace technology as the choice for unlicensed communication for two reasons. First, it provides a new spectrum block hitherto unavailable for unlicensed and opportunistic use and is currently lightly loaded compared to the overgrazed WiFi spectrum. Second, integrating TV whitespace communication into this architecture requires us to solve new technical challenges of spectrum access and use that will be relevant even in a three technology solution (licensed, WiFi, and whitespace) of the future.

C. Our key contributions

In this paper, we make the following important contributions:

- **Optimizing whitespace communication for use in WhiteCell**: We design channel access and contention resolution mechanisms for whitespace communication customized to the unique design of the WhiteCell architecture. Our design addresses multiple usual challenges that occur in efficient design of the whitespace networking including: (i) the spatio-temporal variations in the

availability opportunities of whitespace spectrum, (ii) the flexibility and advantages of operating over whitespace channels with variable widths (ranging from 5 MHz to 20 MHz), (iii) the requirements of protecting primary users in the TV whitespace band. In addition, we leverage unique opportunities that are present in the WhiteCell architecture, which includes: (i) opportunities of collaborative spectrum sensing in a shared femto cell environment, and (ii) coordination in whitespace spectrum allocation among multiple WhiteCell base stations through the mediation of a natural point of centralization in this system.

- **Use of licensed spectrum to bridge transient problems in whitespace spectrum:** Whitespace spectrum can have transient problems for a secondary user, e.g., due to the return of a primary or due to activity of other secondary users. In such scenarios, WhiteCell uses the licensed interface in as a temporary communication link between the base station and clients, while the whitespace interface attempts to identify an alternative efficient channel in which to resume communication. Since connectivity to the client over the licensed band is more reliable, the whitespace interface has more time to identify the most efficient channel for subsequent communication. This approach also keeps the overall load on the licensed spectrum from WhiteCell base stations relatively low, thus reducing the interference for the operator's macro cell base stations.
- **Implementation and full evaluation of WhiteCell prototype:** We have completed a full WhiteCell prototype, including a whitespace communication radio system. While our whitespace radio system has some design similarities to prior approaches, e.g., WhiteFi [6], the exact hardware and the software techniques have significant differences. We also conducted detailed experiments with this platform to evaluate the benefits of this system. Our results indicate a minimal performance gain of 53% over other alternative designs with whitespace technologies.

Roadmap: In the following sections, we first briefly provide some background in the TV whitespaces and the relevant FCC ruling. We then present the overall WhiteCell architecture and the design of its constituents, including a centralized *Spectrum Co-ordinator*, dual-link *WhiteCell* and mobile clients. We also describe the use of a control channel, techniques for whitespace spectrum coordination, collaborative sensing, and flow scheduling over the multiple links. We then describe the system implementation of WhiteCell on our specific Software Defined Radio platform. In the last four sections, we compare WhiteCell with other related work, describe our future plans and finally, conclude the paper.

II. BACKGROUND ON TV AND OTHER WHITESPACES

The term spectrum whitespaces was first introduced in the FCC Spectrum Policy Task Force report [28] to describe spectrum that is allocated to a primary user but unused or under-utilized on a various space and time scales.

There are several spectrum bands with whitespaces, primary examples being Digital TV (DTV) band, public safety bands, radar, government and satellite communication bands. Currently, TV whitespaces are of intense interest due to both its excellent propagation characteristics as well as the sustained availability.

In the USA, TV whitespaces refers to the unused portions of TV broadcast spectrum – specifically lower VHF channels 2-6 (54-88 MHz), upper VHF channels 7-13 (174-216MHz) and UHF channel 14-51 (470-698MHz) with the exception of channel 37 reserved for radio astronomy. The FCC in U.S.A. published a ruling on November 4, 2008 [10], permitting the unlicensed devices to use the TV whitespaces on a do-no-harm basis i.e. any secondary use must not interfere with primary incumbents which include over-the-air TV broadcasts and wireless microphone transmissions. To this end the FCC mandated that the unlicensed devices should be able to detect the presence of TV stations with a received signal strength as low as -114 dBm and wireless microphones with a received signal strength as low as -126 dBm. Also, if a primary is detected, the unlicensed user is expected to vacate the band within 2 seconds.

In a subsequent ruling published on 23 September 2010 [11], the FCC relaxed these sensing requirements. Specifically, devices that have access to their geolocation and an Internet connection do not need to perform spectrum sensing and instead can obtain information about available channels by contacting a TV spectrum occupancy database. The ruling also reserved two TV channels –one below and one above channel 37 across entire U.S.A. for exclusive use by the wireless microphones, thus eliminating the need for sensing microphones.

However, we believe that spectrum sensing is still necessary for the following two reasons. First, it can increase the accuracy of determining spectrum occupancy at different regions and reduce the burden at the spectrum database. Moreover, the measurements collected through sensing can be sent back to the database as *feedback*, to continuously improve the accuracy of the database.

Second, the FCC envisions the spectrum database to be queried on a coarse timescale in the order of hours (*48hours* in [11]), to ensure the scalability of the database. However, free whitespace channels are expected to be shared by multiple secondary users, leading to short-term variability in quality. In this context, the presence of a local spectrum sensing capability can help the operator to pick the best TV channel for operation.

III. DESIGN OF WHITECELL

In this section, we present our proposed *WhiteCell* architecture that integrates spectrum whitespaces and cellular femto cells to improve cellular capacity and whitespace communication.

Our discussion is primarily in the context of DTV whitespace reuse. However, the technology challenges identified and our solutions apply to other whitespaces.

The WhiteCell architecture has three main logical components: (1) a *Spectrum Co-ordinator* that sits in the infrastructure and guides whitespace spectrum allocation to different WhiteCell base stations, (2) a set of *WhiteCell* femto cells or base stations and (3) client devices that are connected to various femto cells.

In this section, we discuss how this architecture addresses some of the main issues in solving the technical challenges. They include:

- *Collaborative spectrum sensing for primary detection:* We leverage multiple femto cells within a geographic area to cooperatively identifying the presence of primaries.
- *Coordinated whitespace channel selection across base stations:* In contrast to selecting operating channels for the base stations in a distributed manner [6], we decide to use a centralized coordinator to assign an appropriate channel for each femto cell. We present the gain of the centralized scheduling in the evaluation section.
- *Dealing with whitespace disruptions:* We use the licensed link as a backup to temporarily handle disruptions in any existing whitespace channel. This reduces the chance of connection disruption at the client side. Additionally, it provides the whitespace channel selection process more time, to better measure and identify an efficient channel.

At a high level the end-to-end operations in WhiteCells proceeds as follows: (1) The WhiteCell base station uses the licensed channel for all control functions. In particular, it uses this channel to periodically announce the current whitespace channel for data communications. (2) When a client arrives, it waits for beacons to discover and connect to the base station via the licensed channel. Once it authenticates itself, it switches between the licensed channel and the whitespace channel dictated by the base station for all data communications. (3) Both clients and base stations continuously monitor the whitespace channel for activities of primary users. If any primary activity is detected, all communication is seamlessly transported to the licensed channel. Any time a client is unable to communicate on the whitespace channel, it goes back to the primary channel to inform the base station of such failure, and the base station is responsible for instructing these clients about a new whitespace channel of operation. (4) The spectrum coordinator routinely tracks the whitespace channel in which each femto cell base station is operating. In this process, it also tracks the interference levels in these channels as reported by the clients and base stations. (5) Once disruption in whitespace operations are detected and clients are temporarily moved to the licensed channel, the base station initiates the process of identifying a better whitespace channel of operation. It does this in coordination with the spectrum coordinator. (6) Upon receiving the request of the new channel, the coordinator provides the base station with a set of good channel candidates in whitespaces. The base station then performs some measurements on these candidate channels to pick the best channel, and instructs all its clients to move their operations on this newly selected channel. Note

that the coordinator, though logically being viewed as a single entity, could be implemented on a cluster of computers and even on the cloud to prevent it from being overloaded.

1) *Collaborative spectrum sensing for primary detection:* Spectrum sensing is always a challenging task that requires high accuracy in our whitespace-based femto cell architecture. We use a hierarchical approach for coordinating spectrum sensing efforts.

At a local level, each client and base station continuously monitor its own whitespace channel to detect primary activities. These localized sensing operations utilize both simple energy detection as well as feature detection, e.g., through detecting the pilot tone from the spectrum of a DTV signal. We briefly comment on the specific primary detection algorithms used in Section IV.

At a global level, the coordinator periodically contacts a TV band database(TVDB) [18,22] for the information about primary users in different geographic region as per the new FCC ruling [11]. The coordinator further gathers measurements from each base station in the ground for validation and possibly provides the measurement feedback to the database to enhance its accuracy.

The density of *WhiteCell* base stations is likely to be high. Hence, collating information from multiple base stations, can help enhance primary detection accuracy. However, not all *WhiteCell* base stations can provide channel quality information relevant to a specific *WhiteCell* base station. To determine the set of *WhiteCell* base stations which can collaborate, the coordinator records the coordinates of each base station and also computes the degree of correlation in spectrum observations between different femto cells that are in close proximity of each other. Such a global measurement effort is known to dramatically improve sensing accuracy [15]. Thus, when a base station reports primary activity, the spectrum coordinator can pass this information along to other femto cells in the same vicinity, whose observations are highly correlated with the detecting femto cells.

These femto cells are asked to stop activity in their current whitespace channel and are required to find a new channel of operation. Additionally, when femto cells request new channels of operation, all such channels with detected primary activity are eliminated from the candidate list.

Additionally, the spectrum coordinator can report back the detection of the new primary user to the TVDB, helping it enhance the accuracy.

The frequency of collecting measurements can be customized based on the properties of primary users. For example, TV station broadcasts can for all practical purposes assumed to be always present, while wireless microphones are expected to be transient. Researchers have suggested using machine learning techniques learn the operational patterns of primary users to tailor the frequency of running detection algorithms [30]. In our implementation we assume that the frequency and periodicity of detecting primaries is an input to the system, guided by FCC's rules of agility with which spectrum needs to be vacated.

2) *Coordinated whitespace spectrum selection across base stations*: To facilitate the channel selection process, the coordinator collects whitespace channel measurements from each base station, which includes the estimated channel quality and its own transmit power. In addition, when a base station is idle, the base station collects measurements from other whitespace channels (not the one it is operating on) as well. The measurement process involves making two *WhiteCell* base stations transmit simultaneously on the same channel and then measuring the fractional increase in error rates due to the particular user. Using this pairwise error information, the coordinator incrementally builds up a conflict graph to determine the degree of interference from different *WhiteCell* base stations at a given *WhiteCell* base station.

As distant whitespace channels have different propagation characteristics, we extrapolate the estimate of interference by a factor of $20\log(f_1/f_2)$, where f_1 is the channel for which we have an estimate of the amount of error and f_2 is the channel for which we intend to calculate the degree of interference. This extrapolation is done based on a formula proposed in [14].

When an individual base station queries the coordinator for good whitespace channel options, the latter can use its conflict graph estimates to suggest multiple candidates, and their expected quality. The channels in this list are 6 MHz wide, which is the smallest unit of whitespace TV channel. The base station conducts its own measurements on each candidate channel and picks the best contiguous frequency band with different bandwidths. We use a simple noise floor based heuristic to determine the quality of a channel. We intend to implement more sophisticated heuristic like MCham [6] as part of our future work.

We will characterize the performance of the various aspects of spectrum sensing algorithm in our experiments section.

3) *Dealing with whitespace disruptions*: The clients and the *WhiteCell* keep monitoring the quality of the in-use whitespace channel. If the clients detect the quality to drop below a threshold, they can request the base station to change the operational channel. In response, the base station can (i) switch this particular client alone to the licensed link or (ii) switch all its clients to the licensed link as it attempts to find a better whitespace channel. If the base station is the one that detects the drop in channel quality then it simply switches all its clients to the licensed link while it finds a better whitespace channel.

The mechanism for channel switching involves the base station sending multiple messages over the licensed channel to inform the clients about switching to the licensed channel. The clients which correctly receive these messages follow suit. Any client that misses all these messages is disconnected and automatically reverts to the licensed link. Finally, when the base station finishes selecting a new whitespace channel of operation, it instructs all its clients to resume communications on this new whitespace channel.

The above approach ensures that the switching of channels does not lead to a disruption of an already ongoing communication. We describe our methodology for achieving seamless

connectivity in Section IV.

4) *Optional data striping across both links*: In addition to our proposed design in *WhiteCell*, some operators may also choose to use the available licensed link as a regular data communication channel, in tandem with the whitespace channel. Hence, we also implement a scheduling system by which flows may be mapped proportionally to the whitespace and the licensed links to reap the expected throughput gains. Since the gains in such a design are obvious (from a single femto cell perspective), we do not report further on this component of our system in this paper. Instead, we focus primarily on the use of the licensed link to enhance the robustness of the whitespace link and for all control communication purposes.

IV. IMPLEMENTATION OF WHITECELL

The design and organization of our prototype hardware for whitespace sensing and communication is mostly similar to that of KNOWS [23] used in WhiteFi. However, due to commercial unavailability of KNOWS platform, we had to build a new platform from scratch. For the sake of completion, we next describe the distinct features of our prototype whitespace radio. Our implementation of *WhiteCells* leverages a Software Defined Radio (SDR) platform called Wideband Digital Radio (WDR).¹ The WDR provides two main capabilities: (1) *Spectrum sensing* capability allows the radio to be tuned to any center frequency between 30 MHz to 7.5 GHz and a channel width of 5, 10, 15, 20 MHz to capture spectrum in terms of I/Q samples. (2) *Frequency translation* capability takes an input RF signal in ISM(UHF) band and converts it to the lower(higher) UHF(ISM) band.

The WDR uses an off-the-shelf dual-band 802.11 a/b/g transceiver - MAX2829 as the analog front-end to capture the spectrum. Such chipset currently supports configurable sampling rates of up to a maximum of 64 MSamples/sec and a resolution of 12 bit/sample at very low price. The captured spectrum samples need to be processed to perform functions such as spectrum sensing and (receive) baseband signal processing for data communications.

The WDR provides two modes of operation to process spectrum samples: (1) *As a Hardware Radio*: Using onboard Xilinx Virtex-5 FPGA for substantial I/Q processing and then selectively communicating information via IP packets over gigabit Ethernet, (2) *As a Software Defined Radio*: Extracting the I/Q sample stream in the form of packets and sending them over on-board gigabit Ethernet to a host that implements software signal processing such as SORA [29] or GNU Radio. A snapshot of our hardware prototype is shown in Figure 4.

To make a functional whitespace node capable of detecting primary users as well as communicating in whitespace channels, we implement an abstraction layer in the core FPGA engine which provides two logical paths into WDR shown in Figure 2. (1) A control path over which IP packets can be sent to control the radio. The key parameters of WDR,

¹The WDR platform is a precursor to the radio that will be deployed by the NSF GENI program .




Fig. 2. Logical organization of a whitespace transceiver.




Fig. 3. Logical organization of a WhiteCell

Channel width in MHz	Throughput achieved in Mbps
5	7.1
10	20.9
20	33

TABLE I

TABLE SHOWING ACHIEVABLE UDP THROUGHPUT WHEN OPERATING WITH A CENTER FREQUENCY OF 642MHz AND USING 54 MBPS FIXED RATE. THE PLOT SHOWS AVERAGE OF 20 RUNS.

such as the active antenna, TX/RX direction, center frequency, channel width, oscillator frequency for frequency translation to be controlled. (2) The data path over which actual I/Q samples or their processed version can be exchanged between the WDR and a host (e.g.: PC, baseband board, DSP board). The host computer controls the WDR using an application level device driver that uses the control path. The latency of changing WDR parameters using our application level driver is order of tens of microseconds, which is very small.

The WDRs in our possession currently support 2 antennas each in transmit and receive paths but only one of the 4 antennas can be active at any time. This allows half duplex communication and antenna diversity. Future versions of WDR will allow 2 antennas in each direction to be used concurrently allowing full duplex as well as MIMO communication.

We used the frequency translation functionality of WDR to implement WiFi in the DTV whitespaces by translating the ISM band (2.4 GHz) RF signals to the DTV bands. Since the available whitespace is in multiple of 6 MHz, we modified the WiFi device driver with support for 5, 10, 15 and 20 MHz communication [6,9,17]. We benchmarked the performance of our platform and found that the achieved data rates, shown in Table I are comparable to the data rates reported in literature [9].

We utilized the spectrum sensing capability of the WDR to detect the presence of the TV and the wireless microphone(Figure 2). For this we tuned the RF frontend with a spectrum width of 20MHz and generated 1024 FFT bins with a resolution of ~ 20 KHz/bin (20000/1024 KHz). These FFTs were packetized and forwarded to the host PC over an Ethernet link. The calculation of FFTs in the FPGA provided significant performance gains over KNOWS hardware where raw energy samples are sent over to the PC for calculating FFTs in software.

Spectrum sensor

The spectrum sensor consists of an application level driver running on the host PC which takes the FFT encapsulated in Ethernet packets as input. The WDR sends energy samples every 4 milliseconds. We averaged multiple energy snapshots to ensure a high degree of accuracy and applied a feature based identification of TVs and wireless microphones [16]. We show the waveform detected for the ATSC TV signals and the wireless microphone signals in Figure 5(b,c). We have experimented with this setup at multiple locations in two US states and have found that it is able to detect TV and microphone signals with signal strength as low as -110 dBm. Our TV detection algorithm searches for the presence of a pilot signal at a specific frequency (about 310 KHz from the lower edge of a frequency band).

In Figure 6, we plot the average accuracy of detecting the TV signal with our system as a function of the number of energy snapshots required. The actual energy of the TV signal was -110 dBm. As can be seen from the plot, collecting 50 spectrum snapshots would provide us a detection accuracy of 95%. Moreover, the duration of taking 50 spectrum snapshots is 0.2s ($50 * 4\text{milliseconds}$), which is trivial compared to the sensing period of 1s required by FCC.

For detecting wireless microphones, we search for a specific slope of energy floor over a band of 50 KHz. We achieve similar accuracy for detecting wireless microphones and omit the results for the sake of brevity. We are aware of the gap in the detection threshold between our algorithms and FCC's regulation(-114dBm for TV and -126dBm for MIC). This is due to the performance limitation of the analog front-end of the WDR. We are currently in the development of the second generation of the WDR which is expected to provide better sensing performance.

WhiteCell dual radio transceiver

In absence of a real cellular base station, we emulated the operation in the licensed channel using a WiFi radio. We will provide a detailed description of our emulation setup in Section IV. As described before, we use frequency translation functions in WDR to map the signal generated by a WiFi card operating on 2.4GHz band for operation in whitespace channel. To prevent interference with the above card, we operate the WiFi card acting as the licensed radio in 5GHz band. All the WiFi radios are configured to run in Pseudo-IBSS




Fig. 4. TV whitespace spectrum sensing logical architecture and actual snapshot.




Fig. 5. Sample snapshots of the spectrum captured.




Fig. 6. Accuracy in detecting TV signal as a function of the number of energy snapshots. Each datapoint is an average of 20 runs.

mode of MadWiFi drivers, wherein no management packets are sent. As 3G cellular interfaces generally operate on 5 MHz channels, we restrict the WiFi interface emulating the cellular interface to operate on a 5 MHz channel as well.

WhiteCell base station

Figure 3 illustrates logical organization of the *WhiteCell* base station. The base station application is implemented in a Linux based PC, where the two WiFi NICs for access and the Ethernet backhaul link to the internet are linked together as ports of a router. We manipulate different router settings such as determining outgoing ports for specific flows etc., using standard Linux networking commands.

When the *WhiteCell* intends to switch the service of a specific client from one interface to another, it first sends beacons to the client asking it shift its active interface, described next. It then modifies relevant routing table entry to change the designated port for the flows destined for that client, once it finds that the client traffic has arrived on the new port.

WhiteCell Client

Figure 3 illustrates the logical organization of a *WhiteCell* client. The *WhiteCell* client software is also implemented in a Linux based PC. To achieve switching of active interface can be done without breaking the end-to-end connection, we create a virtual interface using the Linux *bonding* driver. The bond driver exposes a virtual interface **bond0** to the client device and internally maps the traffic destined for the virtual interface to one of the two WiFi cards. A useful property of the bond driver is that it always uses a single IP and MAC regardless of which of the two underlying interfaces is actually used for communication. Finally, the actual interface on which the traffic should flow can be selected by writing to a `"/proc/` filesystem interface exposed by the bonding driver. The client side driver thus, waits for instruction from *WhiteCell* base station to switch between two interfaces. As the IP and MAC addresses are unchanged, the switch is transparent to the ongoing flows.

Switching the active interface: The switching of interfaces is initiated by the beacons sent by the *WhiteCell* base station. The *WhiteCell* base station sends multiple beacons in succession to ensure that all clients are able to switch. In our experiments we have found that around 4 beacons sent with an inter packet delay of 20 milliseconds are almost always enough

to ensure reception at clients. In case of non-reception of beacons, the clients either keep receiving service over licensed backup channel if they were in that band to begin with or once their non-communication timer expires and they revert back to the backup channel.

Channel quality assessment: We need to measure the channel quality to compare among a candidate set of channels to select a channel of operation. Authors in [6] have suggested the MCham metric, which calculates the aggregate bandwidth that a WhiteFi AP and its clients would receive if it were to operate in each of the potential whitespace channel to find the optimal channel of operation. Instead of implementing the MCham heuristic we use a simple noise floor based heuristic for the sake of simplicity. For channels with different width we scale the quality metric by the fraction of width. The exact heuristic is presented below. For two candidate channels C_1 and C_2 , with width W_1 and W_2 respectively, pick C_1 if $W_1 * \mu_1 \leq \alpha * W_2 \mu_2$, where μ_1, μ_2 are the average noise floor for C_1 and C_2 and α is an empirically selected error threshold. We benchmark the performance of our heuristic in Section V-B.

In-band channel monitoring is also necessary to trigger switch to a better channel when the quality of active channel deteriorates. For this, we modify the driver to collect using a system call (*ioctl*) low level statistics such as packet error rates. The *WhiteCell* base station and clients maintain a running average of these parameters to track the channel quality and request a channel switch when the indicators drop below a certain threshold.

V. EXPERIMENTAL RESULTS

We now present detailed experiments with our *WhiteCell* system to evaluate our design choices and understand the performance advantages of such design. We do this in two parts. In Section V-A we first present the overall performance advantages of *WhiteCell* operating with a spectrum coordinator, base stations, and clients, and describe the performance of generic TCP and UDP flows and also of VoIP clients running over this system. In Section V-B we present some more detailed analysis with our specific design choices.

Experimental setup: In our experiments we used multiple *WhiteCell* units (base stations and clients) each equipped with a whitespace radio (implemented through frequency translation of a Atheros-based WiFi system on the WDR platform

operating in the 470 to 698 MHz range) and a “licensed radio” emulated by using an 802.11a WiFi radio operating on channel centered at 5.18 GHz. Internally the whitespace radio’s WiFi chipset was tuned to the 2.4 GHz band to avoid any interference with the other WiFi radio in this setup. We validated this prior to the experiments. To provide a similar throughput achievable in the 3G cellular link, we configured the WiFi radio to operate in 5MHz bandwidth as mentioned before. We also restricted the bandwidth of all whitespace radios to 5MHz to ensure a fair comparison between our design and the alternative design with dual whitespace interfaces. We used the SampleRate rate adaptation algorithm [7] for the whitespace radio as implemented in the MadWiFi driver. In different experiments we had to detect regular TV channels operating over-the-air. We also used a wireless microphone [3] for some of our experiments as a primary user to be protected. Finally, we used other secondary whitespace transmitters in-the-air to emulate activity from various other secondary users in the vicinity.

A. System performance

We first present overall performance results of using WhiteCell when compared to some alternative designs.

1) *Advantages of licensed channel as backup*: In Figure 7 we compare WhiteCell with a distributed alternative involving only whitespace radios embedded into femto cells. More specifically, we consider each femto cell and client to be equipped with two whitespace radios and no licensed radios. We call this alternative, *wspace-wspace*. One whitespace radio is used for active communication, while the other whitespace radio is constantly scanning the spectrum for other channels. Whenever a disruption in the active whitespace channel occurs, communication is immediately switched to the second radio using the recently best channel observed by the second radio. Since the goal here is to resume communication as quickly as possible, the second radio is assumed to not be able to initiate a measurement process again once the disruption happens, and instead looks into its recent history of channels to pick the best one. Once it does so, the first radio in the detector (say, the base station) has to inform the corresponding radio in its peer (say, the client) to make this channel switch and indicate the new channel of operation through a beacon message. Occasionally such beacons might get lost due to high level primary activity that caused this disruption. In such a design, it is beneficial to trigger such channel switch operations rarely, i.e., only when performance in the current channel falls approaches our threshold and a switch is absolutely needed.

In the scenario presented, other independent whitespace radios were engaged in random communication across different parts of the whitespace spectrum over time using a single TV channel at any given time. We did this by using two such radios that picked and changed their channels of operations over time (roughly every 12 seconds), to emulate a number of such whitespace radio in the region.

The figure shows a timeline of TCP throughput achieved by WhiteCell (with a whitespace radio and a licensed radio)

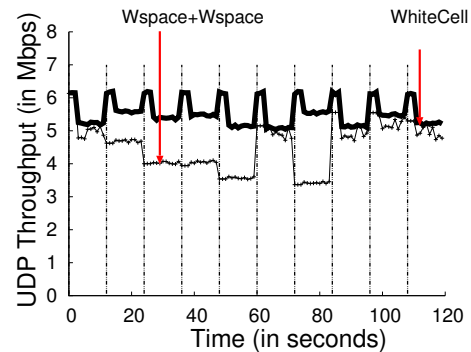


Fig. 7. Benefits of WhiteCell due to a steady licensed channel and due to careful channel selection through coordination, under interference from other secondary users and some primary activity. The figure shows variation in UDP throughput of an Iperf [5] session as a function of time.




Fig. 8. Benefits of licensed backup channel. The results are based on 20 runs

compared to *wspace-wspace*. In this experiment, a primary user (say, sporadically appearing microphones) appear in the whitespace channel of operation at regular intervals causing whitespace communication in that channel to fail. The performance of WhiteCell is superior to such an alternative due to multiple reasons. First, at the instant of a switch, WhiteCell can immediately transition its traffic on the licensed traffic providing no loss in performance to the client. Furthermore, the whitespace radio in WhiteCell can find a better whitespace channel because (i) it can coordinate its future channel of use through the coordinator and (ii) it has adequate time to measure the new channel alternatives before deciding on one of them. In contrast, the *wspace-wspace* system has to quickly pick an alternative channel and has to do so with somewhat stale information about channel characteristics. Measurements made about a certain channel in the long past is not always a good indicator of the future because other active secondary users (unlike primary users) may communicate only intermittently and in short bursts. This stale information leads to a poorer channel selection for its future communication.

Finally, the peaks in performance observed in Figure 7, corresponding to operation in licensed band is due to better quality of the 5 GHz channel used for our experiments compared to 2.4 GHz channel used for translation unto whitespaces. Figure 8 shows multiple runs of these experiments for both UDP and TCP flows comparing the performance of WhiteCell and *wspace-wspace*, again in presence of different interferers in the whitespace channels. The median performance gains due to the licensed backup and due to the ability to coordinate

spectrum selections is 4.0 and 2.1 Mbps for UDP and TCP respectively.

2) *Benefits of coordinated spectrum selection*: We next examine the benefits of coordinated spectrum selection through our coordinator and how this helps in improving the overall performance in WhiteCell in Figure 9. For this, we compare three different schemes to pick a channel once a disruption in the current channel occurs. They are: (i) *Random selection*: In this scheme, the base station picks a random set of whitespace channels, conducts measurements in each of them and selects the best channel among this random set; (ii) *History-based selection*: Here the base station picks a set of recently observed good channels (using say, a second whitespace radio) and then conducts new measurements to pick the best channel from among this set; and (iii) *Coordinated selection*: This is WhiteCell’s channel selection scheme in which the coordinator who specifies a set of candidate channels and the base station performs measurement to select the best channel among them. In each case we assume the set of candidate channels is the same to all of these selection processes, and that none of the candidate channels have any primary user active.

In this experiment, there are two other independent secondary whitespace radios transmitting on randomly picked channels changing channels roughly every 12 seconds (to emulate a number of other such radios communicating independently on different channels).

Figure 9 indicates the performance gains of using coordination between the base stations mediated through the coordinator. In these experiments, the number of candidate channels were varied from 1 to 5 for each scenario and the results plotted. Note that as the number of candidate channels increase, the time taken to complete measurements in each of them increase as well as so a large number of candidates is not practical. As expected, a history-based scheme outperforms the random scheme, while the coordinated scheme (used in WhiteCell) provides additional gains through channel selection.

B. Microbenchmarks

We present results to justify the various design choices made in WhiteCell through detailed performance micro-benchmarks.

1) *Sensing accuracy and collaboration*: We conducted experiments in two different states (in urban areas of a city in the east coast of the US and in a medium sized city in the Midwest of the US) to understand the accuracy of detecting primary users. For TV signals, we experimented with 30 TV channels, and obtained ground truth statistics from (i) a freely available TV channel occupancy database [4] and (ii) also by using a TV tuner card to detect all TV channels present. For microphones, we used our own, so ground truth was readily available.

Using our feature detection techniques and using only a single whitespace receiver at a single location, we were able to accurately identify the presence of a TV signal and a microphone signal more than 98.3% and 93.2% of the instances.

Sensing inaccuracies stem from unpredictable RF shielding in indoor settings, our sensor in some cases, is unable to detect signals from primary users, which would otherwise be easily detected at a nearby location. Hence, while the detection accuracy of the single sensor is good, we observed further performance improvements were possible through collaborative sensing. In our setting, due to the close proximity of at least one client to each base station, such collaboration is, indeed, natural.

Figures 10 and 11 show the improvement in detection accuracy for TV signals and microphones as we used collaboration between multiple whitespace nodes (e.g., a WhiteCell base station and its clients). The X-axis indicates the additional number of detectors used, and the Y-axis identifies the accuracy. In each of these experiments, we assumed that the detection is successful, if any one of the sensors returned the presence of the primary. Figure 10, shows that the overall accuracy of detecting a TV signal is around 88% when two sensors collaborate. The accuracy goes above 98% with five or more collaborating sensors. In Figure 11, we find that a single sensor can detect a microphone within 68% of times when operating within a radius of 250 meters. We also find that with two collaborating sensors (say, two nearby clients or base stations) the detection accuracy went up by 85%, and with three such additional sensors accuracy reached 98%.

The high degree of accuracy gained with multiple sensors justifies our decision of equipping whitespace nodes with collaborative sensing capability.

2) *Variation in channel correlation with distance*: In order for collaboration to be effective, we also studied how detection of signals correlated with distance between collaborating sensors. To do this, we fixed the location of one whitespace spectrum sensor and move another one to distances of upto 50 meters in 10 meter increments along 4 different directions, continuously collecting energy samples at both sensors. We present our observations in Figure 12. As can be seen from the plot, increasing distance leads to decrease in measurement correlation. We use a threshold of 0.5 for using collaboration between multiple sensors, which corresponds to a distance of 50 meters. In typical settings, we expect at least the communicating client(s) and the corresponding base station to be within this range, allowing them to effectively collaborate. If there are other base stations in vicinity, such base stations and their clients can also potentially collaborate. This observation allows us to also restrict the set of potential collaborators for the purpose of sensing to those within geographical vicinity, in turn reducing the workload of a coordinator.

3) *Efficacy of noise floor based metric*: We next describe the correlation between noise floor and channel performance. We setup a link between two nodes operating on a 5 MHz whitespace channel. We then incrementally increase the amount of interference by increasing the transmit power of an interferer node with carrier sense disabled in units of 5dBm. We run Iperf to measure the UDP throughput and the number of packets in error. We show the variation in UDP throughput as a function of increasing noise floor observed




Fig. 9. Timeline showing the improvement due to the coordination mechanism proposed in WhiteCell. The number of candidate channels in each case were varied from 1 to 5 before picking the best channel through measurements in each case.






Fig. 10. Improvement in primary TV signal detection accuracy as a function of additional number of collaborating sensors, for the scenarios where a single radio was unable to detect it at the given location. Results for a city in Midwest US (Loc A) and another in east coast US (Loc B) across a 2 mile radius in each case.

Fig. 11. Improvement in wireless microphone detection as a function of additional collaborative sensors, for scenarios where a single radio was unable to detect it at the given location. We experiment with upto 4 sensors spread across an area of 250 meters around the microphone.

Fig. 12. Variation in degree of correlation as a function of distance of separation and channel separation. The measurements were collected for a period of 10 hours, with inter-measurement gap of 1 msec.

Normalized Performance Difference	Number of Instances (out of 200)	Count of Erroneous Prediction	Cumulative Estimation Accuracy (in %)
< 0.2	40	3	92.5
< 0.4	78	1	96.6
< 0.6	42	0	97.5
< 0.8	36	0	97.9
≤ 1.0	4	0	98

TABLE II

TABLE SHOWING THE PERFORMANCE OF OUR NOISE-FLOOR BASED CHANNEL QUALITY ASSESSMENT HEURISTIC IN SELECTING CHANNELS. IN OUR SELECTION PROCESS, WE COMPARE TWO CHANNELS (C_1 AND C_2) AND ESTIMATE THE BETTER CHANNEL BASED ON OUR PROPOSED HEURISTIC. WE SHOW THE COUNT OF INSTANCES WITH NORMALIZED UDP THROUGHPUT BETWEEN THESE TWO CANDIDATE CHANNELS WHEN OUR ESTIMATES PICK THE CORRECT CHANNEL. WE USED $\alpha = 0.2$. FOR EXAMPLE, WHEN CHANNEL C_1 IS BETTER THAN C_2 BY 0.2, WE CORRECTLY PICK IT 92.5% OF THE TIME.

at the client (caused by increasing signal strength of the interferer) in Figure 13(a), and the increase in packet error rate in Figure 13(b). As can be seen from the plot, there is a distinct correlation between the observed noise floor and performance. This justifies our selection of noise floor based metric to determine the quality of different whitespace channels. The improvement in performance around -85 dBm is explained by




Fig. 13. Variation in achievable UDP throughput (in Mbps), error rate (in %) as a function of worsening noise floor of a 20 MHz TV band channel.

the fact that for the transmit power used by the transmitter the SNR at the receiver becomes above the sensitivity threshold.

To measure the accuracy of the heuristic proposed in our design section, we carry out the following experiment. We randomly pick a pair of center frequencies and channel width. We tune in the transmitter, receiver and interferer to both of the chosen frequencies and width one at a time. We carry out a UDP transfer between transmitter and receiver, in presence of interference. The transmission power of the interferer is set randomly to ensure a variety of experiment condition. We measure the UDP throughput over 20 seconds. We also monitor the average and variance of noise-floor during the

period of experiment. Using the noise-floor values we predict the better of the two channel, and match the prediction with the actual throughput. We present our observations in Table II. We find that out of 200 runs of this experiment the heuristic had an accuracy of 98% in predicting the better of the two channels. We also note that the heuristic made errors only when the quality of the two channels being compared were really close to each other. This is proved by the fact that the achievable throughput on both channels was within 10%-20% in all the error cases. The above observation leads us to the conclusion that, noise-floor based heuristic is a good predictor of channel measurement.

VI. RELATED WORK

In this section, we discuss related work in some detail.

A. Whitespace network architecture

The WhiteFi [6] provides WiFi like connectivity over whitespace channels. The designed system does not provide guarantees of disruption free communication and uses whitespace bands as the only band of communication. In contrast we advocate, simultaneous operation of two interfaces for capacity augmentation. Another point of difference of our system design is the use of coordinated spectrum allocation and collaborative sensing. The authors of in SenseLess [22] implement a TV whitespace database (TVDB) [11] that provides list of available whitespace channel at a given location. It primarily relies on propagation models that are inaccurate and do account for operational ground truth. In contrast, we integrate notion of sensing to bolster the spectrum occupancy maps.

Deb et al. [14] present a theoretical framework which accounts for scenarios where channel conditions vary slowly, only interference from radios in the network is modeled and all interference scenarios resolved using interference graph calculations. The authors rely on a simulation to analyze their approach. Also, like WhiteFi, they use one radio only for control and data channel operations.

In [14], authors follow a centralized approach of spectrum allocation, where the *Spectrum Co-ordinator* decides which spectrum block should be used for by each *WhiteCell*. In contrast, in our implementation, the *Spectrum Co-ordinator* only provides a set of "hints" letting the *WhiteCell* determine the optimal channel based on ground measurements.

B. Spectrum sensing and quality assessment

Given broad topic of spectrum sensing has been a subject intense theoretical and practical exploration in recent years, we refer reader to a good survey paper. [31].

In the case of DTV whitespace, sensing presence of DTV transmitters at signal thresholds such as -116 dBm and -126 dBm wireless microphone detection [10], is a very hard problem. Our current prototype implements a feature detection on DTV signal spectrograms – an approach similar to ones in [13, 23, 26] for primary detection. While more complex detectors

that exploit additional signal properties such as cyclostationarity [31] improve accuracy, the computational complexity and detection latency increases significantly. In our architecture, we can augment our baseline sensor by selectively requesting at different times a subset of (underloaded) *WhiteCell* base station to perform such expensive sensing to reduce the sensing burden.

The technique of collaborative sensing, which aggregates sensing measurements from multiple sensors to draw inference is very old and has applications in diverse areas of autonomous robotics, navigation, military and aerospace communication and mobile systems [27]. It has been shown that such collaboration when extended to spectrum sensing improves sensing accuracy and threshold and lowers false alarms [15]. Our unique architectural approach with a centralized *Spectrum Co-ordinator* exploits these gains.

The authors of [6] report a novel *Mcham* metric that estimate quality of whitespace channel by estimating number of *WhiteFi* secondary transmitters present. Unlike this metric, our simple and easy implement *noise floor* metric accounts for *all* (not just WhiteFi like) transmitters and energy sources. Also, correlating our metric at various nearby *WhiteCell* base stations provides a highly accurate estimate of channel quality.

Most of the prior works are either analytic in nature or conduct small scale experimentation. Our work represents first effort to address system level issues of intensive processing on the overall system performance.

C. Multi-interface networking

The concept of exploiting multiple links, called by various names such as multi-interface networking, link bonding, aggregation, striping etc. has found application in various areas.

It has been used for increasing link capacity in context of DSL modems [2]. In the context of wireline link bonding, problem is simple as links are independent (i.e no cross-channel interference) and are always available.

It has been used for contexts such as integration of cellular and WiFi networks [8] for maintaining seamless coverage and performing vertical hand-off for preserving end-user sessions.

In case of Mobile Routers, it has been used for providing backhaul to a group of mobile devices [24] and also, for increasing capacity of such backhaul [21]. In contrast, we use multi-interface capability to tide over disruption caused by loss of whitespace channel during spectrum sensing period. Though the implementation mechanisms we use may be used in other settings, our objectives are different.

VII. FUTURE WORK

Our future plans for the *WhiteCell* architecture are as follows:

- *Other spectrum band*: We plan to study exploitation of whitespace in other bands such public safety, radar, government and satellite communications. The primary characteristics are different and more challenging than

DTV band. We expect the power of our WhiteCell architecture and capability of our frequency agile hardware will be a great asset.

- *Large scale testbed:* We plan to create a large scale *WhiteCell* testbed in our institution to study complex system issues in the context of whitespace networking.

VIII. Conclusions

We believe that load on licensed spectrum continues to grow in a significant manner. Operators today are looking to reduce this load in many different ways. First, they have introduced the basic licensed femto cell architecture for improving spatial re-use in their licensed bands. Second, they are also exploring ways in which various unlicensed bands, such as WiFi can offload some of their client traffic. The availability of the new whitespace TV spectrum provides a new avenue for relief on this highly loaded spectrum. One routine concern among operators with unlicensed technologies is their inability to provide any performance guarantees, primarily because of the presence of various transmitters which fall outside the control of the operator. Our design in WhiteCell tries address this challenge. By using a licensed link as a backup, this system provides a safety net whereby poor performance in an unlicensed channel can be quickly handled without affecting a client's performance. The extra time gained while traffic is moved into a licensed band allows the communicating nodes to find another good channel for future communication. The coordination in the system also allows for improved performance across the entire whitespace band.

In designing WhiteCell, we did not try to optimize each and every feature of our system and we believe that many such optimizations can be part of future efforts. In particular, it might be possible that not one, but two different unlicensed technologies (WiFi and whitespace) be both used in such a system for reducing the load on licensed spectrum. However, the design and integration of whitespace spectrum into such a system did pose some interesting challenges and we made a first attempt at solving them.

Overall, we think that the biggest contribution of this work is to provide an acceptable trade off between the requirements and expectation of cellular operations of performance guarantees of licensed spectrum with the new wave of opportunities that are presented by the swath of unlicensed whitespace spectrum.

ACKNOWLEDGMENT

Milind Buddhikot's work was partially supported by his US NSF award 0831762. Sayandeep Sen, Tan Zhang and Suman Banerjee are supported in part by the following grants of the US NSF: CNS-1040648, CNS-0916955, CNS-0855201, CNS-0747177, CNS-1064944, and CNS-1059306.

REFERENCES

[1] Cisco visual networking index: Global mobile data traffic forecast update, 2009-2014. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html.
 [2] Ieee 802.3ad link aggregation group.

[3] Shure uhf-r wireless,model-ur24s. <http://www.shure.com/americas/products/wireless-systems/uhfr-systems/index.htm>.
 [4] Tv fool. <http://www.tvfool.com>.
 [5] NLANR/DAST : Iperf - the TCP/UDP bandwidth measurement tool. <http://dast.nlanr.net/Projects/Iperf/>, Accessed 2007.
 [6] Paramvir Bahl, Ranveer Chandra, Thomas Moscibroda, Rohan Murty, and Matt Welsh. White space networking with wi-fi like connectivity. In *SIGCOMM*, 2009.
 [7] John C. Bicket. Bit-rate selection in wireless networks. In *Master's thesis, Massachusetts Institute of Technology*, 2005.
 [8] M. Buddhikot, G. Chandranmenon, S. Han, Y.W. Lee, S. Miller, and L. Salgarelli. Integration of 802.11 and third-generation wireless data networks. In *INFOCOM 2003*, 2003.
 [9] Ranveer Chandra, Ratul Mahajan, Thomas Moscibroda, Ramya Raghavendra, and Paramvir Bahl. A case for adapting channel width in wireless networks. In *SIGCOMM*, 2008.
 [10] Federal Communications Commission. Fcc press release, fcc adopts rules for unlicensed use of television white spaces, November 2008.
 [11] Federal Communications Commission. Fcc press release, unlicensed operation in the tv broadcast bands, second memorandum opinion and order, September 2010.
 [12] Federal Communications Commission. National broadband plan, connecting america. <http://www.broadband.gov>, 2010.
 [13] C. Cordeiro, M. Ghosh, D. Cavalcanti, and K. Challapali. Spectrum sensing for dynamic spectrum access of tv bands. In *CrownCom*, 2007.
 [14] Supratim Deb, Vikram Srinivasan, and Ritesh Maheshwari. Dynamic spectrum access in dtv whitespaces: design rules, architecture and algorithms. In *MobiCom*, 2009.
 [15] A. Ghasemi and E.S. Sousa. Collaborative spectrum sensing for opportunistic access in fading environments. In *DySPAN*, 2005.
 [16] M. Ghosh, V. Gaddam, G. Turkenich, and K. Challapali. Spectrum sensing prototype for sensing atsc and wireless microphone signals. In *CrownCom*, 2008.
 [17] Ramakrishna Gummadi and Hari Balakrishnan. Wireless networks should spread spectrum based on demands. In *HotNets*, 2008.
 [18] Spectrum Bridge Inc. Enabling universal spectrum access. <http://www.spectrumbridge.com/Home.aspx>.
 [19] Vinod Kone, Ben Y. Zhao, Lei Yang, Haitao Zheng, and Xue Yang. On the feasibility of effective opportunistic spectrum access. In *IMC*, 2010.
 [20] K. Lee, I. Rhee, J. Lee, S. Chong, and Y. Yung. Mobile data offloading: How much can wifi deliver? In *ACM CoNEXT*, 2010.
 [21] Robert Melervy, Yan Sun, and Thomas F. La Porta. Multiple backhaul mobile access router striping. In *IEEE ICC 09*, 2009.
 [22] Rohan Murty, Ranveer Chandra, Thomas Moscibroda, and Paramvir (Victor) Bahl. Senseless: A database-driven white spaces network. *IEEE Transactions on Mobile Computing*, 2012.
 [23] S. Narlanka, R. Chandra, P. Bahl, and J.I. Ferrell. A hardware platform for utilizing tv bands with a wi-fi radio. In *LANMAN 2007*, 2007.
 [24] P. Rodriguez, R. Chakravorty, J. Chesterfield, I. Pratt, and S. Banerjee. Mar: a commuter router infrastructure for the mobile internet. In *ACM MobiSys '04*.
 [25] Ashwin Sampath, Lei Yang, Lili Cao, Haitao Zheng, and Ben Y. Zhao. High throughput spectrum-aware routing for cognitive radio based ad-hoc networks. In *CROWNCOM*, May 2008.
 [26] Stephen J. Shellhammer, Sai Shankar N, Rahul Tandra, and James Tomcik. Performance of power detector sensors of dtv signals in ieee 802.22 wrans. In *ACM TAPAS Workshop*. ACM, 2006.
 [27] Duncan Smith and Sameer Singh. Approaches to multisensor data fusion in target tracking: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 18:1696–1710, 2006.
 [28] Spectrum Policy Task Force. Spectrum Policy Task Force Report. *Federal Communications Commission ET Docket 02-135*, 2002.
 [29] Kun Tan, Jiansong Zhang, Ji Fang, He Liu, Yusheng Ye, Shen Wang, Yongguang Zhang, Haitao Wu, Wei Wang, and Geoffrey M. Voelker. Sora: high performance software radio using general purpose multi-core processors. In *NSDI*, 2009.
 [30] Lei Yang, Lili Cao, and Haitao Zheng. Proactive channel access in dynamic spectrum networks. In *CrownCom 2007*, 2007.
 [31] T. Yucek and H. Arslan. A survey of spectrum sensing algorithms for cognitive radio applications. *Communications Surveys Tutorials, IEEE*, 2009.