

# A Campus-Wide Testbed over the TV White Spaces<sup>\*</sup>

Ranveer Chandra, Thomas Moscibroda, Paramvir Bahl, Rohan Murty<sup>a</sup>, George Nychis<sup>b</sup>, Xiaohui Wang<sup>b</sup>

Microsoft Research, Microsoft Corporation, Redmond, WA, USA

<sup>a</sup>Department of EECS, Harvard University, Cambridge, MA, USA

<sup>b</sup>Department of ECE, Carnegie Mellon University, Pittsburgh, PA, USA

*We have deployed a wireless network that operates in the white spaces of the TV band spectrum and covers most of Microsoft campus in Redmond, WA. Since the campus is large (approximately 1 mile x 1 mile), there are several shuttles that move employees from one building to another. We have used the white spaces network to enable a key productivity scenario on campus - Internet connectivity in campus shuttles. We have modified one such shuttle to operate over the TV band white spaces. A white space radio in this shuttle communicates with two base stations deployed on buildings on campus. Inside the shuttle, we bridge the white space connection to Wi-Fi, so that an employee's laptop that does not have an integrated white space radio can nevertheless connect to the Internet, using Wi-Fi within the shuttle, and white spaces between the shuttle and the base stations. To the best of our knowledge, this is the first network deployed over the TV white spaces.*

## I. Introduction

The FCC's white space ruling has triggered tremendous excitement around wireless networking over the TV white spaces. The ruling, which allows unlicensed devices to operate in the unoccupied TV bands, has opened up many new application scenarios ranging from rural broadband connectivity to faster, reliable connections within the home. Many of these scenarios depend on the two unique benefits of TV white spaces: significant amount of new unlicensed spectrum (up to 300 MHz depending on the region), and good propagation characteristics due to the lower frequencies. Several technology companies, including silicon vendors, hardware manufacturers and software companies have eyed these opportunities and started developing technologies to efficiently use this portion of the spectrum.

In this paper, we report on our effort to enable one such new, compelling scenario - providing network connectivity in all outdoor areas of a large university or industrial campus. Efficiently enabling this scenario with existing technology is not possible: Using cellular networks is expensive and offers low bandwidth. Furthermore, client devices are no longer directly connected to the corporate network, which introduces additional overhead, thus reducing performance. Prior attempts to solving the problem using unlicensed spectrum have shown that Wi-Fi has limited range and suffers from losses, especially in the

face of mobile clients. For example, it was reported in [6] that even deploying three to four Wi-Fi Access Points on the rooftop of every two-storeyed campus building led to several coverage holes in the network. In contrast, white spaces are a cheaper alternative for providing more bandwidth to students and employees in all parts of the campus and under the direct control of the university or employer's IT department.

However, building a campus wide white space network is non-trivial among other reasons due to the lack of previously deployed systems that can be followed as an examples. For instance, the range of a white space base station at the FCC-permissible parameters is unknown in practice. A previous simulation-based study expects coverage of 33 km at 4 W EIRP [1], which proves to be a drastic overestimate. Similarly, the impact of wireless microphones (mics) is not known. In the FCC's Second Order from September 2010, the FCC provides two mechanisms to protect mics [2]. First, two TV channels are exclusively reserved for mics. Secondly, and this is of great importance to parts of our work, microphone operators can nevertheless reserve other TV channels (besides the two mentioned above) for exclusive use. White space devices need to vacate a channel that is reserved by a mic within 400 meters of its location (if it using the geolocation database) or, if it detects using low-threshold sensing, whenever it senses the mic at -107 dBm.

We have deployed a white space network that covers most of Microsoft campus in Redmond. Since the campus is large (approximately 1 mile x 1 mile), there

<sup>\*</sup>More details can be found on the KNOWS project website: <http://research.microsoft.com/knows>

are several shuttles that move employees from one building to another. We have used the white spaces to enable a key productivity scenario on campus - Internet connectivity in the campus shuttles. Specifically, we modified one shuttle to operate over the white spaces. A radio in this white-space enabled shuttle communicates via empty TV channels with two base stations on buildings on campus. Inside the shuttle, we bridge the white space connection to Wi-Fi, to ensure that an employees' laptop can connect to the Internet even without being able to communicate over white spaces directly. The laptop uses Wi-Fi within the shuttle, and white spaces between the shuttle and the base stations. Through this trial deployment we show the following:

- White spaces enable long distance communication, much more than Wi-Fi. Only two base stations were sufficient to provide coverage to nearly all of campus.
- TVs can be protected solely using the geo-location service. We successfully demonstrated no interference to KOMO TV. This was shown to a delegation of senior executives from Fisher Communications, the parent company of KOMO TV.
- Wireless microphones can be protected solely using the geo-location service. Entries in the service are populated either manually or automatically through a device that we call MICProtector.
- Signals from indoor microphones attenuate fast, and reach up to -114 dBm within a few hundred meters.

## II. Deployment Setup

We received an experimental license from the FCC to operate on all the TV band white spaces within our campus. This license allows us to transmit at 4 W EIRP from fixed nodes, and 100 mW EIRP from mobile devices. We are allowed to transmit from mobile nodes in the UHF (512 to 698 MHz) and VHF (174 to 216 MHz) frequencies. We coordinated with the Society of Broadcast Engineers (SBE) for the subset of TV channels we could operate on.

We use custom-built hardware to operate in the white spaces. The hardware is built on a WiMax IEEE 802.16d chipset, and is capable of operating below 1 GHz. The specific frequencies it can operate on are tunable through a policy file. We set it to operate only on the upper VHF and UHF white space frequencies,




Figure 1: A snapshot of our campus, and the route taken by our shuttles for most experiments in the paper. The highlighted portion of campus is 1 mile long and 0.75 miles wide.




Figure 2: We attached a 25 to 1500 MHz VHF/UHF antenna to the shuttle hitch. This connects to the radio in the shuttle using an RG-8 cable.

i.e. 174 to 216 MHz and 470 to 698 MHz. This hardware is also capable of transmitting at up to 10 W although we only transmit up to 4 W EIRP as permitted by the experimental license.

We deployed 2 nodes for measurements in our campus, which spans 1 mile by 0.75 miles, as shown in Figure 1. We placed a 25 to 1500 MHz discone antenna with a 2 dBi gain on the rooftop of a 4-storeyed building (shown as a pushpin on the map). A 100 foot RG-8 cable was used to connect the antenna to our radio, which we placed in the server room on the second floor of the building. The RF cable added a 2 dB attenuation which was offset by the gain of the antenna.

We modified a campus shuttle to operate over white spaces (Figure 2). We connected the radio to the shuttle's battery, and tested our system with two different antennae - a VHF/UHF antenna to the hitch of the shuttle, as shown in Figure 3. The performance of the smaller antenna was good in the UHF spectrum, while the larger antenna performed well in both the UHF and VHF bands.



Figure 3: The PC and white space radio are placed next to the driver. The Windows 7 laptop records the measurements, and also geo-tags it with the GPS reading.

We have deployed a five node network using the radios from Shared Spectrum. Two radios are setup as base stations on rooftops of campus buildings. We setup a third client in an on-campus shuttle (Figure 2), and use the remaining two nodes as nomadic clients, which we move to different parts of our campus. We have also tested our deployment with radios from Adaptrum Inc, Lyrtech and the Microsoft KNOWS prototype [3, 8].

To have real users of our white space network, we use the PC in the shuttle to serve as a Wi-Fi Access Point. This also solves the legacy problem - riders on the shuttle can use their regular Wi-Fi enabled devices (e.g. laptops) to connect to the Internet (and Intranet) using their Wi-Fi connection to the Wi-Fi access point within the shuttle, which then connects to the rooftop access points via the long range white spaces. This enables us to see the benefits of white spaces by having users access the Internet for free over this spectrum even before commercial white space devices hit the market.

The base station in our network transmits at 4W EIRP, while the clients use transmit power control, and use a maximum of 100 mW EIRP. Using BPSK modulation at the Base Station and clients, and using the above settings, we were able to successfully ping the base station from all points in our route in Figure 1 over VHF. Over UHF, the coverage was more limited, and an approximate convex hull of our coverage will be shown in Figure 8.

**Network Stack Enhancements:** Unlike Wi-Fi and other wireless protocols that expose an Ethernet like interface to the upper layers, the FCC regulations require white space radios to operate differently. They need to determine spectrum availability by communicating to a geo-location service over IP (or sensing from the radio). We have modified the Windows 7 network stack to add this support. Our implementation

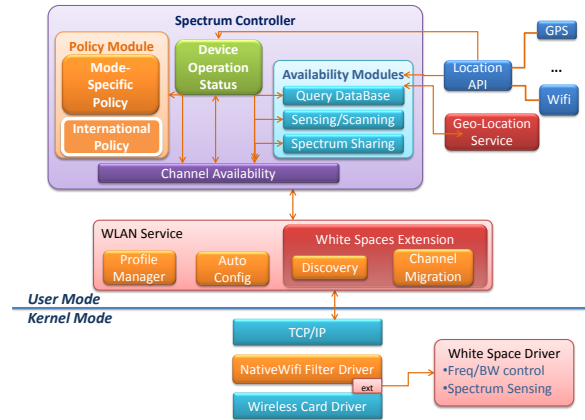


Figure 4: We have modified the Windows network stack on the client and Base Station to support white spaces. Most components reside in the user space, while driver modifications are needed to support different bandwidths, and switch frequencies in the presence of primary users.

(shown in Figure 4) has three main components. First, the Spectrum Controller is a user level service that determines the white spaces that can be used by the device. It currently does so by communicating with a geo-location service (for which it needs location), but the software is also provisioned to support spectrum sensing or using feedback from nearby nodes. In addition, the spectrum that can be used depends on the policy of the country or the region. This is taken care of by the policy module of the Spectrum Controller. The main API exposed by this component is the set of channels that can be used by the device. The second component is extensions to the existing wireless service in Windows 7 that manages the associations of the Wi-Fi card. We have enhanced the service to make seamless transitions to the white space network when the Wi-Fi network is unavailable. This component is also responsible for switching to a different part of the spectrum if the Spectrum Controller signals its current channel as unusable. Finally, the third component in the stack is a set of modifications to the wireless driver. We wanted to keep these changes to a minimum. Our current modifications include support for adapting channel widths and for buffering packets when switching across channels.

### III. Coverage over the TV White Spaces

For the purpose of measurements in this section, we used 518 MHz to test the white space coverage over UHF frequencies, and 177 MHz to test coverage over VHF frequencies. Both corresponding TV channels are available in our campus location.

Each radio is controlled using a PC, to which we



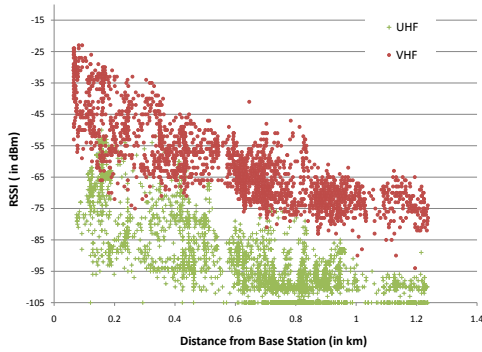


Figure 5: The raw received power at the shuttle at different distances from the base station. The noise floor of our receiving radio was -105 dBm when transmitting over 1.75 MHz of the spectrum. As we see, both the UHF and VHF spectrum have excellent propagation characteristics.

also attached a GPS unit (Figure 3). The GlobalSat’s BU-353 GPS is based on the SiRF Start III chipset and its location accuracy is within three meters 95th the time. We use the unit to record the GPS reading once every second. For our measurements, we follow the route shown in Figure 1 in our white-spaces enhanced shuttle. It usually takes 45 minutes to an hour for the shuttle to cover the above route, which implies more than 2500 data points (one reading per second) for every configuration.

Due to their lower frequencies, we expect the UHF and VHF spectrum bands to have better signal propagation than Wi-Fi’s ISM bands. For example, for the same transmission power, the Friis formula [5] predicts 4 times the range at 600 MHz over 2.4 GHz Wi-Fi. The practically important question is whether this theoretically predicted range does really translate into a corresponding increase in network coverage when used in a campus-like environment with obstacles, buildings, etc.

We present the variation of received signal strength with the distance from the base station in Figure 5. These measurements reveal interesting observations. First, as expected, the signal is weaker on average when increasing the distance. However, there is significant variation in the received power for the same distance, which we attribute to the number of obstructions in the path. For example, the shuttle receives a 10 dB stronger signal at the line-of-sight building on the other side of the freeway in Figure 3, than another building at a similar distance from the base station on the same side as the freeway. Second, VHF signals are much less affected by obstructions. They have a much smoother fall off than UHF signals. In fact, the propagation is much better than what is predicted by the Friis formula. According to the formula, signals

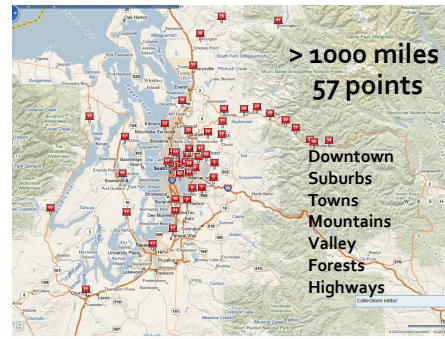


Figure 6: A 1500 mile path over which we measured UHF spectrum at 57 different locations. The push pins indicate points of measurement.

at 518 MHz should experience approximately 10 dB more attenuation than 177 MHz VHF. Although we observe similar attenuation at shorter distances, the attenuation is significantly more on increasing the distance.

Finally, we note that UHF frequencies can propagate more than 700 meters before hitting the noise floor of our system of -105 dBm. Therefore, the propagation characteristics of the white space frequencies seem promising in providing network coverage in large campuses.

#### IV. Avoiding Interference to TV Broadcasts

As mentioned, the shuttle has a GPS, using which the PC periodically communicates its location to the base station. The base station queries our geo-location service for the available channels, and configures the network to operate on a clean channel. The geo-location service hosted at <http://whitespaces.msresearch.us> reads the TV tower data from the FCC database and terrain data from NASA, and performs sophisticated propagation modeling to predict the TV channels available at a given location. For a detailed description of the design of the geo-location database, we refer the reader to our DySPAN 2011 paper by Murty et al. [7].

##### IV.A. Accuracy of Propagation Modeling

We measured the UHF spectrum across the state of Washington. Since spectrum availability varies by population density and terrain features, spectrum measurements were taken during the months of July-August 2009, across a driving path of 1500 miles, at a set of 57 diverse locations including large cities, downtowns, suburbs, between large buildings, mountain ranges, forests, valleys, at the edge of water bod-




Figure 7: Comparison with ground truth for different propagation models. We lose fewer white spaces by using terrain based Longley Rice propagation.

ies, and areas of different population densities.

Using a spectrum analyzer fitted with a UHF antenna, we measured the signal strength for all 30 UHF channels by restricting the channel on the spectrum analyzer at the center frequency of every TV channel and a bandwidth of 6 MHz. We then measured the channel power of the spectrum across this 6 MHz band. We mark a channel as being occupied if the corresponding channel power is less than or equal to -81 dBm. (a limitation imposed by our spectrum analyzer hardware when measuring channel power values), and available otherwise. Using the same threshold, we compare these findings with those predicted by different models and high resolution terrain data (measured at 100 m intervals across the planet’s surface). If a channel is occupied in the ground truth data but free in the models, we flag it as a false positive; if it is available in the ground truth data but not by the model, it is a false negative.

In our data set of 57 locations (shown in Figure 6) and 30 channels each, none of the models gave any false positives. In the free-space model, we used a very conservative path-loss exponent of 2. Using an exponent of 3, we measured several false-positives. So, they all met the safety requirement. To quantify efficiency for each location, we express the number of false negative channels as a fraction of the total white spaces available at that particular location. We use these measurements to compare the following four well-known propagation models: Free Space, Egli, and Longley-Rice (L-R) with terrain, and L-R (without terrain), the commonly used models for white space analysis. Of these, L-R with terrain is the most complex model since it takes into account climactic effects, soil conductivity, permittivity, the Earth’s curvature, and surface refractivity. It also uses terrain elevation data as input and as a result is computationally intensive.

The results, shown in Figure 7, represent a CDF of the fraction of channels lost because of false nega-

tives. The median loss rate for L-R with terrain is only around 8 available channels. To us this was a surprising result since it suggests that with careful modeling it is possible to 1) not lose too many white spaces and at the same time 2) have very few or no false positives. In contrast, all other models result in many white spaces being wasted.

Therefore, the loss of white spaces by not sensing is low when using L-R model with terrain data. However, one observation with important systems implications is that only complex and computationally-intensive models based on terrain data are able to achieve a satisfying accuracy; using simplistic propagation models results in a waste of white spaces.

**Demo to KOMO TV:** Executives from Fisher Communications, the parent company of KOMO TV, visited our campus in January 2010 to ensure that our transmissions did not interfere with their transmissions. During their visit, we set a TV receiving over the air KOMO TV and were operating on a nearby channel and no interference was seen by the visitors. They were very happy with the visit and were willing to collaborate with us on better use of their spectrum.

## V. Avoiding Interference to Wireless Microphones

Our campus has a large number of wireless microphones (mics). There are lectures in buildings, sometimes more than one, and our campus usually hosts several other events as well. We briefly present our solution to protect mics in this section. For a more detailed study, we refer the reader to our CoNext 2012 paper at Nychis et. al. [9].

### V.A. Occurrence of MICs

Wireless microphones operate in virtually all the available TV channels. We analyzed the mic usage in 4 different buildings in our campus – two department buildings and two conference centers. Our results showed that (i) every building had a few rooms setup with wireless mics, (ii) every such room was overprovisioned with the number of mics (since wireless mics run on battery), and (iii) every mic has been carefully set to operate on a different mic channel. Through conversations with operators we learnt that mic frequency allocation is performed by experts once every year.

The two department buildings we explored belonged to the research and marketing group respectively. The former building had 8 locations set up with wireless mics, and per-location there were 3, 4 or 8

mics, for a total of 32 mics. Each mic is tuned to operate in a different mic channel, and collectively, the 32 mics operate in 12 different TV channels. The marketing building is significantly bigger and has 22 locations set with wireless mics, for a total of 101 mics that are allocated over 12 TV channels (7 of which were common with the research building).

The two conference buildings had a higher density of mics. The big building had 80 mics for regular use, all on different mic channels, and spreading across 16 TV channels. The smaller building had 39 mics on different mic channels and across 10 TV channels.

We note that not all mics are in use all the time. Depending on the schedule, anywhere from one to all locations was active in every building. For every location, the number of mics was fixed, although the channels used by the mics varied depending on the particular handsets that were in use, for example, the ones with the most remaining battery lifetime.

## V.B. Protecting MICs

We use two mechanisms to support wireless mics:

**Manual updates:** We provided mic users with an API to add an entry for a mic as a primary user in our geo-location service. Authorized users could add an entry for the mic's frequency, transmit power and the location and duration of the event. Our system then treats the mic transmitter similar to a (very) low power TV tower for the specified duration. These entries can be expired after a time out period.

**MICProtector:** We provide an alternative to manually updating the primary user database. Our technique consists of a small device, which we call the MICProtector, which is plugged in close to the mic receiver. The device detects the presence of the mic and updates the back-end database using a web-API. It uses an alternative technology, such as 3G, Wi-Fi or Ethernet to connect to the geo-location service. When the Updater does not detect the mic for a predefined amount of time (in our system, 5 minutes), the entry is deleted from the database.

Note that this solution does not simply move the difficulty of low-threshold sensing from the white space client device to the MICProtector. Because the MICProtector is powered and close to the mic system, it does not require complex low threshold sensing. This reduces the number of false negatives/positives as well as hardware cost. Hence, the MICProtector approach is very effective, non-intrusive, and relatively inexpensive way to prevent interference with wireless microphones. It also does not require legacy wireless mic users to change their microphones or overhaul

their deployments. Instead, this approach requires a fixed cost - at most one MICProtector is needed for every wireless microphone user.

The concept of a separate device to protect the primary users is not new. The FCC distributed the converter box for backward compatibility after the DTV transition. Similarly, mic manufacturers such as Shure and Motorola are proposing the use of a separate device to protect mic s as well. One could argue that implementing sensing at the white space clients may be cheaper than handing out MICProtector to the mic users. However, we note there is far fewer wireless mics (on the order of a few hundred thousand in the US including illegal mic s) compared to tens of millions of potential white space clients if it were to become as successful as Wi-Fi. Furthermore, many mic s are co-located at a venue, such as an auditorium, in which case multiple MICProtector may not be required. For these reasons, the trade-off of reducing significant complexity from clients (no sensing required) at the cost of a low-complexity (no low-threshold sensing required) device for mic s seems preferable.

## V.C. Experience on Campus

Although we provided the MICProtector option to the mic users, we did not have to physically use it in any location, except during testing. Even so we did not receive any complaints from mic users about interference from our white space deployment. We believe this is mostly because most mic usage is indoors and there is heavy RF shielding material used in buildings, which we describe in more detail in the next section. Furthermore, mic receivers are also usually kept close to the mic transmitter, and this too reduces the likelihood of interference with the mic.

## VI. Region Blocked by MICs

It is commonly believed that a single wireless mic can potentially block off a complete TV-channel within a large area. For example, a simple free-space path-loss model would predict (even for a path-loss exponent as high as 4) that a mic transmitting at 14 dBm would block off the channel within an area of roughly 2.5 kilometers range round the mic, at the FCC-mandated sensing threshold of -114 dBm. And even a much more conservative propagation model, such as the Egli model, predicts the channel to be blocked within a range of 1.3 kilometers. Hence, theoretical models predict that even a single wireless mic located near the center of our campus would essentially block off an

entire TV channel, rendering it unusable throughout the campus.

If, indeed, a single wireless mic blocked off an entire TV channel on campus, this would cause a significant problem in system capacity, because in areas with many wireless mics (such as our campus), there would remain only very few available TV channels. Fortunately, this is not the case. The reason is that wireless mics are almost always used indoors, typically in conference rooms. Our measurements show that signal emissions from mics within conference rooms are significantly shielded from the outside, such that the remaining signal strength that leaks outside the buildings is low.

We studied these further using measurements on our campus. We placed a 4W transmitter at 518 MHz inside two buildings, and measured the received signal strength around the route highlighted in Figure 3. Figure 8 shows the region that is blocked off at -114 dBm by a 10 dBm mic in the topmost building. The shaded region is a manually estimated convex hull for all regions where the mic would be heard at more than -114 dBm. This is a small fraction compared to the entire coverage of the base station, i.e. locations from where the client could successfully ping the base station when operating under the FCC rules, i.e. client transmitting at 100 mW and the base station at 4W. We quantify this further in Figure 9. As we can see, a 10 dBm transmitter will mostly block off less than 300 meters around it.

We notice that the signal is sufficiently attenuated by the building to cause the channel to be blocked off for white space use only within the near vicinity of the building. The strong attenuation seen in the above measurement is not accidental, but is fundamental to many campuses and to urban areas in general. The reason is that modern buildings use so-called low emissive (Low-E) windows and doors to prevent heat from passing through windows in order to save energy. Windows achieve Low-E typically with a metallic coating which reflects the long-wave rays that bring heat and UV. In our context, the issue with Low-E coated windows is that they significantly add to the attenuation of the mic’s signals.

In summary, these results imply that while the specific attenuation faced by an indoor mic signal can vary greatly depending on the precise location and building structure, the attenuation caused by the building is typically high. Hence, we conclude that the amount of white spaces lost due to wireless mics even on a busy campus with numerous conference rooms is small. On the other hand, both FCC mandated mic

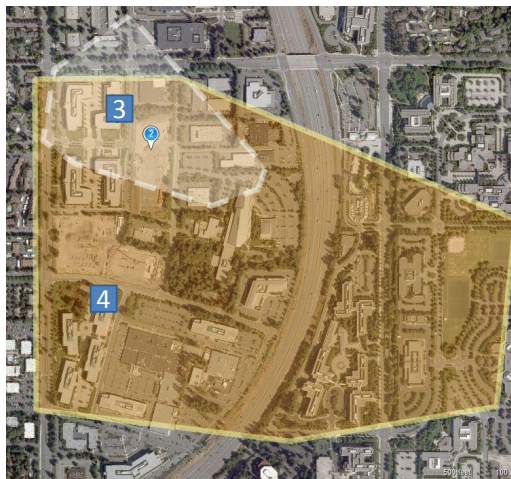


Figure 8: The small region blocked off by a wireless mic (Pushpin 3) when compared to the overall coverage of the base station (Pushpin 4).

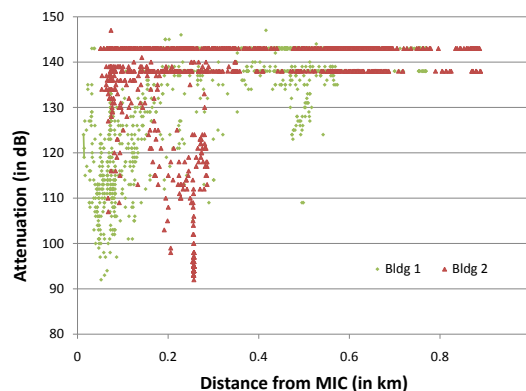


Figure 9: The attenuation of the signal when the transmitter is placed with the building. For a 10dBm transmitter, a 124 dB attenuation (to get to -114 dBm) nearly always occurs with 400 meters.

protection approaches will lead to significant amounts of white spaces being unnecessarily being lost, because the base station range is nevertheless higher than the mic protection range.

### VI.A. Local Spectrum Asymmetry (LSA)

This observation leads to interesting open research problems. Given that the base station range is large, but each mic blocks off a channel only in a small portion within the base station’s coverage range (see Figure 8) a kind of Local Spectrum Asymmetry (LSA) problem can arise: It may not be possible or efficient to serve all clients connected to the same base station on the same channel, because different clients may be close to mics that operate on different frequencies. For this reason, we conjecture that efficient solutions that can scale to a large number of clients distributed across the entire campus will ultimately require new kind of MAC-layer protocols in which one base sta-



tion can simultaneously service multiple clients on different channels. In [4], we analyze the impact of LSA more precisely and propose a preliminary solution to the problem in the form of a new channel-assignment protocol. But, additional research is required to fully understand and solve this problem.

## VII. Impact

Our deployment was operational on October 16th 2009 and to the best of our knowledge this was the first deployed white space network. Our demonstrations have shown that (i) it is possible to send high speed data transfers over a single TV channel, (ii) we do not interfere with TV broadcasts, and (iii) our system is agile to any changes in spectrum availability. To demonstrate (iii), we show that our system moves to another available TV channel when we introduce a wireless MIC near a white space device on the same channel.

This one of its kind network has gained attention worldwide and spectrum regulators from several countries have visited the Microsoft campus to see the network in operation. Some visitors have included the FCC Chairman as well as the FCC Chief Counsel, the TRAI Chairman from India, and officials from IDA (Singapore), SARFT (China), and ANATEL (Brazil).

## References

- [1] IEEE 802.22 Working Group on WRANs, <http://www.ieee802.org/22/>.
- [2] Office of Engineering and Technology Invites Proposals from Entities Seeking to be Designated TV Band Device Database Managers. November 2009.
- [3] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White Space Networking with Wi-Fi like Connectivity. In *SIGCOMM*, 2009.
- [4] R. Chandra, T. Moscibroda, and X. Wang. The Impact of Heterogeneous Spectrum Availability in a Campus White Space Network. *Manuscript*, 2010.
- [5] H.T.Friis. A Note on a Simple Transmission Formula. In *Proc IRE*, 1946.
- [6] R. Mahajan, J. Zahorjan, and B. Zill. Understanding WiFi-based Connectivity from Moving Vehicles. In *IMC*, 2007.

- [7] R. Murty, R. Chandra, T. Moscibroda, and P. Bahl. SenseLess: A Database-Driven White Spaces Network. In *IEEE DySpan*, 2011.
- [8] 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.
- [9] G. Nychis, R. Chandra, T. Moscibroda, I. Tashchev, and P. Steenkiste. Reclaiming the White Spaces: Spectrum Efficient Coexistence with Primary Users. In *ACM CoNEXT*, 2011.