

Reclaiming the White Spaces: Spectrum Efficient Coexistence with Primary Users

George Nychis[†], Ranveer Chandra[§], Thomas Moscibroda^{*}, Ivan Tashev[§], Peter Steenkiste[†]

[†]Carnegie Mellon University, [§]Microsoft Research, ^{*}Microsoft Research Asia

ABSTRACT

TV white spaces offer an exciting opportunity for increasing spectrum availability, but white space devices (WSDs) cannot interfere with primary users, including TV channels and wireless microphones (mics). Mics are particularly challenging because their use is dynamic and it is hard to avoid interference since mic receivers are receive-only devices. For this reason the FCC and other regulatory agencies have made very conservative rules that require WSDs to vacate any TV channel that is used by a mic. However, our measurements show that mics typically require only 5% of a channel, wasting as much as 95% of the spectrum.

We present SEISMIC, a systems that enables WSDs and mics to operate on the same TV channel with zero audible mic interference. SEISMIC implements a MicProtector to measure the interference at the mic receiver and a signaling protocol to notify the WSD of impending interference. This allows the WSD to optimize its transmission (e.g. through subcarrier suppression) without impacting mics. We motivate and describe SEISMIC and present a detailed performance analysis that shows that SEISMIC can regain up to 95% of the spectrum in single mic scenarios, and up to 85% in many (10+) mic environments.

1. INTRODUCTION

The proliferation of wireless devices has led to an impending spectrum crisis [2]. To provide more spectrum, the FCC and spectrum regulators worldwide are exploring techniques to reuse unoccupied TV channels (white spaces) for data communication [8]. The FCC finalized its rules on September 23, 2010 [10], while the UK, Brazil, Finland, Singapore and other countries are working on white space rules as well. While such efforts have been significant, the rules in place to protect primary users threaten the very goal the white spaces are trying to achieve: additional spectrum availability.

This is particularly true regarding the ruling's handling of wireless microphones (mics), which along with TV broadcasts are primary users of this spectrum. In the First Order from December 2008, the FCC ruled conservatively towards mic protection by enforcing a white space device to vacate an entire TV channel (6MHz) in the presence of even a single mic (at most 500KHz). In the Second Order, the FCC made the ruling even more conservative by stating that two channels will now be exclusively reserved for mics. This is *in addition* to allowing licensed mics to operate in *any* channel under database (or sensing) and channel vacation protection.

Particularly in populated areas, these rules significantly reduce spectrum availability for white space devices. For example, in 12 of the largest 30 US cities, there are only 2 or fewer unoccupied TV channels, and in 21 of these cities, there are no more than 5. Dedicating two TV channels for

wireless microphones effectively eliminates white spaces in a large fraction of US cities. Moreover, as mics will still be able to operate as primary users in any of the other TV channels, the amount of white spaces in some locations of the remaining cities will effectively be reduced to zero.

It is clear that such conservative measures run counter to the goal of white space networking, and are likely to be a major impediment to its widespread adoption. Moreover, we show in this paper that such conservatism is unnecessary: full microphone protection can be achieved while still enabling white space devices (WSDs) to reclaim large fractions of the spectrum. It is well-known that in different parts of the spectrum, OFDM devices can use subcarrier suppression [22, 23, 25] to eliminate interference with narrow-band devices. In the white spaces, such an approach would potentially allow a WSD to coexist with a mic in the same TV channel, and use the remaining 95% of its spectrum [22].

Unfortunately, such solutions cannot easily be deployed in the white spaces for several reasons. First, existing techniques that adaptively determine the degree of required suppression require interfering with the narrowband devices (i.e., SWIFT [23]). This is unacceptable as it would cause harmful audible interference with the mic. Second, the amount of suppression required to protect the mic's transmission depends on the mic's received signal power, as well as the white space device's interference, at the mic receiver. Both values are unknown at the WSD, change over time, and cannot be estimated using channel reciprocity techniques [16] given mic systems are one-way (receiver never transmits).

To overcome the challenges, we present SEISMIC (Spectrum Efficient Interference-free System for MICs), a system that allows white space devices to coexist with mics and “recover” a close to optimal fraction of the spectrum in the TV channel, while fully protecting the mic in all circumstances. Our system design is based on an in-depth characterization of the impact of white space transmissions and RF interference on mic audio recordings. To allow cooperation between the mic system and WSDs, SEISMIC uses a simple device called a MicProtector to measure the interference at the mic receiver and a low-complexity signaling protocol to notify white space devices of *impending disruption*. Using this explicit signaling feedback, secondary users can suppress the

proper frequency to avoid disrupting the mic’s audio quality.

The SEISMIC approach combines several attractive properties. First and foremost, it is safe since the explicit feedback avoids harmful interference at the mic receiver, independent of the placement of mics and WSDs. Secondly, is purely reactive, restricting secondary white space communication only when needed and to the degree necessary. That is, SEISMIC optimizes spectrum usage by minimizing the number of subcarriers that are suppressed based on actual measured WSD signal levels. Finally, in spite of the addition of the new MicProtector device, SEISMIC is a very practical solution: It does not require replacing legacy mic equipment or advanced registration of mics before events, and cost is likely to be small (no low-threshold sensing!).

We find that SEISMIC allows white space devices to converge to within 25KHz of optimal suppression 72% of the time and within 75KHz 93% of the time, with zero-interference to microphones. We show that this allows a WSD to get up to 95% of the bandwidth of a channel that is completely lost with today’s restrictive FCC regulations, and even up to 85% of the channel in many (10+) mic scenarios. While these results are specific to mics in TV white spaces, we note that the SEISMIC design is more general and can be easily adapted for efficient coexistence with other primary users. For example, SEISMIC can enable WSDs to transmit at greater than 40 mW on channels adjacent to those occupied by TVs.

2. BACKGROUND & RELATED WORK

A wireless mic system consists of a mic transmitter and a mic receiver. The mic transmitter converts audio into RF using frequency modulation (FM). The mic receiver decodes the FM signal to retrieve the transmitted audio signal. Figure 1 shows the RF spectrum of the signal when six mics are idle. The signal consists of a center signal that carries data and two side tones, called *scquelch tones*. The mic receiver decodes the mic signal only when the scquelch tones are successfully received. This helps protect against garbled sound when there is interference, and prevents risk of audio amplifiers and speakers when the mic signal is low [1]. Table 2 summarizes the properties of the 6 mics used this paper.

Licensed mics operate under the FCC Part 74 rules in the US, and similar rules worldwide. This rule restricts the operation of wireless mics to at most a 200 KHz wide signal, and a max transmit power of 250mW, although most mics use a max of 10mW. Unlicensed mics are allowed to operate under the FCC Part 15 rules at a maximum transmit power of 50 mW. Worldwide, mics are allowed to operate on any unoccupied TV channel. In the US, the FCC recently modified the rules and reserved two TV channels exclusively for wireless mics [10]. These two channels vary by region, and cannot be used by WSDs. When two channels are not enough, organizers can reserve additional channels 30 days in advance.

Previously proposed solutions to protect mics:

1. Sense for wireless mics: This was amongst the first solutions proposed to avoid interference to mics [6, 11, 12, 20].

The FCC’s initial ruling set this threshold to -114 dBm over 200KHz, and the Second Order reduced it to -107 dBm. OFCOM is considering a sensing threshold of -126 dBm. Such a sensing-based approach has several drawbacks. First, spectrum sensing is an expensive operation in terms of cost an energy [10], and was the primary reason for removing this requirement in the Second Order. Second, the sensing threshold is extremely conservative, both because the signal propagation environment is unknown and because it is meant to protect the mic receiver, whose location is unknown by sensing the mic. Third, sensing for mics at below the noise level (as mandated by existing regulations) is prone to false positives [7]. Therefore, WSDs might end up vacating an entire TV channel even when there is no mic present in the vicinity. Finally, and most importantly, this approach is inefficient, since WSDs need to vacate an entire TV channel.

2. Mic Beaconer: To address the second and third concern above, mic companies, such as Motorola and Shure, have proposed the use of a separate beaconer device to reduce the sensing threshold [9, 13, 26]. The beaconer uses the first 500 KHz of every TV channel to signal the presence of a mic using a 250 mW signal. White space devices vacate every TV channel that has this beacon. This approach still suffers from the other two drawbacks from above – sensing needs extra hardware and entire channel needs to be vacated. Another shortcoming of this approach is that WSDs may vacate the channel even when their transmission does not interfere with the mic. We elaborate on this in Section 4.

3. Two-Reserved & On-Demand Reservation: This is the approach taken by the FCC in the Second Order. Reserving two TV channels for wireless mics (even when there are no mics in the vicinity), as specified in the FCC’s Second Order [10], significantly reduces the amount of white spaces in urban areas. In the top 30 urban areas, our analysis showed that 12 (40%) had only two unoccupied TV channels, 60% had three or less, and 70% had 5 or less¹, so two channels represents a significant fraction of the white space spectrum. This is a serious limitation since success in urban areas, where more WSD users are expected, is seen as a likely driver for white space device use in rural areas [10]. Pace availability is further reduced since event organizers can also reserve any TV channel for mic operation.

Despite this conservative approach that seemingly favors mic users, the audio community is concerned about the ruling as well [3]. The ruling leads to increased cost since most mic users will have to replace their existing equipment. This is because the 2 reserved channels will be geo-dependent as the first 2 free channels in the 180MHz of spectrum, however mic systems usually have a limited 40MHz front-end which is likely to not be in the range of the reserved channels. Furthermore, users who pay a large sum of money for their mic placement [11] will have to redo the mic placement. Mic operators, including those who handle big events such as the Super Bowl, are unhappy about having to reserve TV chan-

¹Geo-location database: <http://whitespaces.msresearch.us/>




Figure 1: Spectrum signature of 6 idle mics.

Microphone	Idle Width	Squelch Separation
A.Tech. ATW-T210	65KHz	30 dB
E.Voice BPU-2	62KHz	25 dB
Senn. E935	69KHz	30 dB
Senn. EW100	66KHz	30 dB
Senn. SK2000XP	69KHz	30 dB
Shure UR-2	65KHz	30 dB

Figure 2: Mics used in our study.




Figure 3: Suppression needed by WSD.

nels 30 days in advance [3]. The RF environment changes frequently, and they adjust frequencies until the last minute.

3. DESIGNING SEISMIC

The goal of SEISMIC is to maximize the amount of spectrum available for white space communication, while ensuring no interference to the primary users. The WSD can suppress a portion of frequency around a narrow-band primary user’s transmission to avoid interference with the primary user [4, 22, 23, 25]. This would suggest that if the secondary can learn about the exact transmission frequency and bandwidth of the primary user, the primary user could vacate a sufficiently large “guard-band,” and use the remainder of the channel without interfering. Here, a database or some beaconing-device could inform of such information.

Unfortunately, such non-adaptive, open loop solutions are very inefficient. The amount of spectrum to suppress depends on the SINR at the mic receiver, i.e. the ratio between the received signal strength from the primary transmitter and the collective interference power generated by the WSDs. Since, the mic receiver’s SINR is not known at the WSD, it needs to make the most conservative, worst-case assumptions. This results in suppressing too much spectrum. To show the inefficiency of such a static, open-loop solution we measure the amount of suppression needed with various SINR values when using a high-quality WSD prototype from Adaptrum and real mic system. Since the WSD does neither know the interference power it creates on the mic nor the received signal strength at the mic receiver, it has no idea where in Figure 3 it operates. Thus, it has to make worst-case assumptions to avoid interference: it must suppress the entire channel—which is exactly what the FCC requires.

3.1 Towards an Adaptive Solution

Several adaptive solutions have been proposed to determine the proper amount of suppression [14, 18, 19, 23]. They can be broken into two groups: (i) those that adapt based on how the the primary user reacts to interference [14, 23], and (ii) those that use channel reciprocity to estimate the interference the secondary user generates, while assuming the worst case for the (unknown) signal component (mic or TV).

Unfortunately, none of these solutions are suitable for co-existence in the white spaces. First, we cannot allow any disruption of primary users (both mics and TV broadcasts). Second, even if we were allowed to interfere temporarily, mics and TV broadcasts are passive transmissions; they do not back off. Third, channel reciprocity cannot be done with

primaries in the white spaces since all white space primaries are one-way systems in which the receiver never transmits. Hence, the WSD cannot estimate the interference component of SINR at the primary receiver. Even if you could solve that problem,² the signal component of the SINR is still unknown. The WSD would have to be conservative.

Thus, the challenge is how to devise a system in which the WSD can adapt their behavior based on the SINR-values (RF interference and signal strength) at the primary receiver; and to convey this information to the WSD in a *disruption-free* manner from the passive device. This should be done without adding significant complexity to the WSD or primary.

3.2 SEISMIC Design

In the design of SEISMIC, we try to approximate an ideal solution in which WSDs have explicit feedback about the SINR-values (RF interference and signal strength) at the primary receiver, allowing them to suppress the minimal number of subcarriers while avoiding any disruption of the mic system. Adaptive systems based on closed loop feedback must include three logical components: measurement, analysis, and adaptation. Since our system is distributed, we also need a signaling component to exchange information.

SEISMIC implements these components as follows:

1. **Measurement:** We introduce a low-complexity enhancement to the mic receiver called a *MicProtector*, that monitors RF interference and mic signal strength. Being co-located with the primary receiver, it accounts for all relevant factors when monitoring its SINR. (Section 5.2)
2. **Analysis:** With SINR, SEISMIC accordingly determines how the system can be optimized, i.e., how to minimize the number of suppressed subcarriers while avoiding audible interference. (Sections 4 and 5.2)
3. **Adaptation:** The WSD follows a protocol in which it adjusts both the number of suppressed subcarriers and the transmit power of its transmission. (Section 5.3)
4. **Signaling:** When needed, the *MicProtector* sends feedback to the WSD using a novel *signaling* mechanism (*strobing*) to warn of impending disruption. (Section 5.4)

While we depict the *MicProtector* as a standalone device, future mic systems can build such functionality in to the receiver. The option of using a standalone device is attractive since it allows deployment without replacing all mic systems. This is similar to the use of converter boxes to cope

²For example by deploying some “beaconer” device co-located with the mic receiver [9, 13, 26].

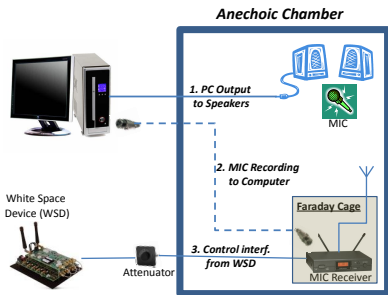


Figure 10: Anechoic chamber setup for audible interference tests.

with the DTV transition. Mic manufacturers have been willing to adopt an additional device to signal the presence of the mic [9, 13]. In private communication, mic operators at large events have been more willing to add this device than replace their existing mic systems because of the high cost of replacing equipment and of replanning frequencies.

In the next section we present a detailed measurement study on the the impact of RF interference on audio quality; this leads to the *analysis* component. The other components and their integration, are discussed in Section 5.

4. RF-INTERFERENCE ON MIC AUDIO

In this section, we provide the first in-depth analysis of how wireless data transmissions impact wireless mics. Such an understanding is critical towards *analysis* and proper *adaptation*. Using a controlled environment (§4.1), we introduce variable RF interference by independently controlling the *power*, *duration*, and *frequency*. We measure the amount of audible interference using the *Perceptual Evaluation of Speech Quality* (PESQ) metric [24] (§4.1).

4.1 Experimental Setup

Our experimental setup is illustrated in Figure 10. We use a PC to play sound samples. We place a wireless mic close to the PC speakers, mimicking a person speaking into the mic. To ensure that any audio disruption is caused only by WSD interference and not from other sources, we placed the PC speakers and the mic in an anechoic chamber. The mic receiver is connected to the PC using an XLR cable, where we save and process the resulting mic recording.

To study the impact of WSD interference on mic recordings, we used WSDs to transmit in a conducted setup to the mic receiver. The WSD was wired directly to the mic receiver to introduce the interference in a controlled manner. We also connected a spectrum analyzer to measure the RF spectrum and channel power. Since mic receivers have exposed antenna elements, we isolated them by placing the receiver in a Faraday cage. To control the power of the transmitted and received signals we placed two RF attenuators: between the WSD and its antenna and between the mic receiver and its antenna. We ran tests with two WSDs: Adaptor, and a USRP2 with a TV TX/RX (WBX) front-end.

We used PESQ to quantify the impact of the interference introduced by the WSD on the mic recording. PESQ is a

signal processing algorithm that provides an estimate for the Mean Opinion Score (MOS), a widely used measure of subjective sound quality, i.e. how humans perceive the quality of sound. PESQ outputs a number from one to five to mimic the MOS results. In our work we used the wideband version of PESQ algorithm called WB-PESQ, standardized in 2005 by ITU-T Recommendation P.862.2. For ease of explanation, we present *normalized* PESQ values. A score of 1 is perfect quality, and score of 0 represents heavy disruption.

4.2 Interference in Power

In the first set of experiments, we have the WSD generate interference continuously (worst case in *time*) and adjust both the power of the mic signal (P_m) and power of the white space interference (P_n) at the mic receiver. For reasons that will become apparent in the results, we focus on the amplitude of the white space interference in relation to two separate power components of the mic signal: mic peak power (P_m) and the power of the squelch tones (P_s) (§2). Therefore, the amplitude of the white space interference in relation to either component, *measured at the mic receiver*, is computed as: $WSN_m = P_m - P_n$ and $WSN_s = P_s - P_n$, where the WSN is in dB. So, if P_m is -30dBm, P_s is -60dBm and P_n is -40dBm, then WSN_m is 10dB and WSN_s is -20dB.

Noise Generation and Measurement: We use the variable attenuator to control the RF interference generated by the WSD. The interference level is set to 100mW (20dBm), and for each test we attenuate it in 5dB steps until the normalized PESQ value begins to approach 1. We then decrease the step size to 1dB for accuracy. We measure the power values (P_n, P_m, P_s) at the mic receiver using the spectrum analyzer. Since it is attached to the RF input ports, we can measure the power as close to the RF chain as possible, accounting for factors such as attenuation in the cables.

Power Results: Figure 4 shows the normalized PESQ score as a function of the white space interference level on the Sennheiser EW100 mic, marking the points at which the PESQ value becomes perfect with vertical grey lines. Our results show that if the WSN amplitude is greater than the mic signal peak (i.e., $P_n > P_m$), the interference is severe enough to cause the mic receiver to stop transferring audio (due to the squelch tones) as is seen from a normalized PESQ score of 0. However, once the peak is approximately 10dB above the white space noise (i.e., $WSN_m \geq 10dB$), the noise becomes less severe and the voice in the audio track becomes noticeable. Most surprisingly, *once the mic squelch tone power was 1dB above the white space noise (i.e., $WSN_s \geq 1dB$), the normalized PESQ score achieved a perfect value of 1. This is despite the fact that 19dB of RF interference still present in the operating band of the mic.*

We repeated the same experiment for the other mics. Figure 5 shows the same result holds: as soon as the white space interference level is a few dB below the squelch tones, we get a perfect PESQ score. The result even holds for the BPU-2 mic, which has squelch tones that are separated by 25dB from the mic signal peak (30dB for the other mics).




Figure 4: Audible interference, varying WSN. Areas where score reaches 1 are highlighted.




Figure 5: Relative amplitude of the WSN for each mic when a perfect score of 1 is reached.




Figure 6: Audible interference threshold (AIT) remains stable when varying mic signal atten.




Figure 7: Senn. EW100 requires 200KHz of non-interfering noise for no audible interference.




Figure 8: The frequency tolerance thresholds for all of the mics for no audible interference.

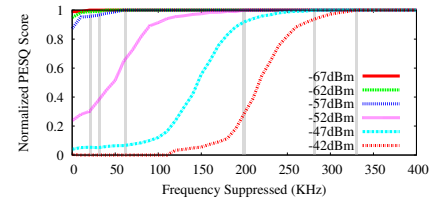


Figure 9: Varying WSD interf. power at mic receiver, suppression needed varies greatly.




Figure 11: Resulting interference with variable spacing.

To verify that this result is independent of the power of the mic, we repeated the experiments with the attenuator between the mic receiver and its antenna set to three different levels: 0dB, 20dB, and 40dB. The results, shown in Fig. 6, confirm this independence. For all mics and all three mic signal levels, the PESQ is perfect as long as the WSD interference is 1-2dB below the squelch tone signal levels.

Observing Capture: This result should not be a surprise given the well studied phenomenon of *FM capture*, which allows for zero reduction in audio quality, despite the possible presence of significant noise. For FM demodulation, frequency shift is measured by tracking the strongest frequency component in a limited band. As long as the main carrier power exceeds the noise by an amount which allows for clean tracking of the frequency shifts, then the FM receiver will “capture” the signal with zero noise [17]. Such behavior was acknowledged by the FCC in the First Order [5] (Paragraph 38): “FM receivers exhibit a ‘capture effect’ in which they respond to only the strongest signal received on a frequency and reject any weaker interfering signals.”

The capture effect is independent of the squelch tones, but the squelch tones are convenient in determining the allowable level of interference. While 1dB of separation may seem small, the actual power difference is significant (allowing capture) due to the decibel being a logarithmic unit.

4.3 Interference in Time

To evaluate the impact of the packet durations, we configure the USRP2 to mimic 802.11-like interference with

respect to symbol timings and OFDM subcarriers. With the USRP2’s master clock of 100MHz decimated by 8 and an FFT of size 64, we achieve an OFDM symbol time of: $(8/100MHz * 64) = 5.12\mu s$. To ensure the USRP2 can ramp up its transmitter, we use 3 successful symbols in length as our minimum to achieve a minimum interference duration of $15.36\mu s$, comparable to four 802.11a/g/n symbols which are $16\mu s$ in length, i.e. a very minimal “frame.” We ran experiments with all six mics and changed the timing of the interference by controlling the inter-frame gap using a sub-microsecond scheduler [21]. In Figure 11, the small inter-frame spacings (10s of μs , similar to 802.11 IFS and back-off) cause PESQ scores near zero, while even for spacings as high as 500ms, the normalized PESQ only reaches 0.7.

4.4 Interference in the Frequency Domain

We now evaluate how much spacing in the frequency domain is needed for a mic system to have zero audible interference. Interference is constant in time and the interference power is set to 10dB above the squelch tones. Initially, the WSD interferes across the entire TV channel (i.e., 6MHz). Then, we incrementally suppress frequency at the center of the mic’s band outwards in 5KHz steps. To get accurate measurements, we ensure that the power falloff is steep. We attenuate the mic signal and interference from the WSD so the power reaches the noise floor within 2KHz (Figure 12).

Frequency Domain Results: Figure 7 shows the impact of the amount of suppressed spectrum on the normalized PESQ score of the Sennheiser EW100. We see that the EW100’s audio quality is severely affected (PESQ=0) when there is less than approximately 110KHz of interference-free spectrum at the center of the mic signal. Beyond this point, audio quality begins improving and once there is 200KHz of free spectrum, there is zero audible interference on the mic. We perform this same experiment for all mics. The results, Figure 8, show that the minimal amount of interference-free spectrum ranges from 150KHz (ATW-T210) to 325KHz (BPU). Moreover, we note that models from the same man-

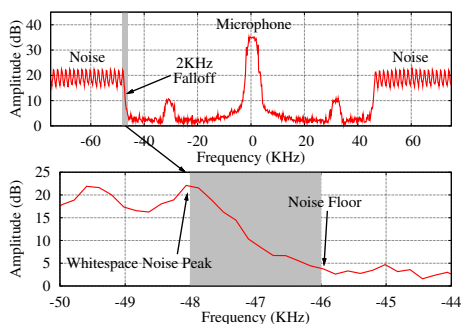


Figure 12: Suppressing 95KHz around a mic, illustrating the sharp falloff of the white space noise to reach the noise floor within 2KHz.




Figure 13: Overview of SEISMIC and MicProtector.

manufacturer can require different amounts of interference-free spectrum (e.g., Sennheiser’s E935, EW100, and SK2000XP).

Varying WSD’s power and power leakage: The proximity of the WSD affects its power, and thus power leakage past the suppressed subcarriers. To evaluate this, we repeated the frequency suppression test on the Sennheiser EW100, but vary the noise power at the receiver from -42dBm to -67dBm in 5dB steps. At each step, we sweep the amount of frequency suppressed from 0KHz to 400KHz in 5KHz steps and compute the normalized PESQ score.

We present the results in Figure 9, highlighting the point at which we achieve zero audible interference for the various noise powers. The results show that, as expected, the amount of frequency suppressed at the transmitter needed to achieve zero audible interference will be different depending on the noise power at the receiver. At the strongest power level, -42dBm, we need to suppress a little over 330KHz at the transmitter. At the lowest power, -67dBm, we only need to suppress 20KHz. Measured, although not shown, at -77dBm (in two more power steps), 0KHz needs to be suppressed for zero audible interference because the interference power is already more than 1 dB below the squelch tones (WSN_s).

5. SEISMIC: TOWARDS IDEAL COEXISTENCE WITH MICS

We explained in Section 3 how spectrum-efficient coexistence between mics and WSDs must be a feedback-driven, closed loop design that allows the WSD device to adapt based on the SINR properties at the mic receiver. In this section, we first revisit the SEISMIC design and then elaborate on the three key SEISMIC components in Sections 5.2–5.4.

5.1 System Overview

As discussed in Section 3, any spectrum efficient solution to the microphone coexistence problem in white spaces requires either additional hardware or changes to legacy systems. In our case, we use a simple device called a *MicProtector* which resides near the mic receiver (e.g., on top of the receiver in Figure 13), near an array of mic receivers common in productions (§5.6), or built in to future mic receivers. The MicProtector is responsible for both the *measurement* and *analysis* components in the closed loop control, in addition to providing feedback on the analysis to the WSD. To do so, the MicProtector monitors the interference power and mic signal (i.e., SINR), and employs a *Protection Threshold* to notify a WSD of *impending* disruption to the mic’s audio. Based on the study presented in the previous section, the protection threshold is set below the mic’s squelch tones.

To notify of impending interference, a low complexity pulse-based signaling mechanism (§5.4) is used to communicate with the WSD. We call this *strobing*, and it requires only carrier sense-like functionality. The strobes are transmitted in *control bands* surrounding the mic (see Figure 13), which we also use for measuring SINR. Since the strobes are raised in both control bands, the WSD can determine the mic’s operational band (i.e., frequency and bandwidth).

To ensure that a WSD never exceeds the *Protection Threshold* and causes an audio disruption, the WSD and MicProtector engage in a protocol. Whenever a WSD starts transmitting on a new frequency band, it does so at minimum power, and then increases this power gradually. As the WSD ramps up its power, *if* the WSD is in disruption-range of a mic, the interference level at the mic receiver will slowly approach the Protection Threshold at which point the WSD will be notified of *impending disruption*. With each impending disruption notification, the WSD suppresses additional frequency. In doing so the system approaches the “ideal state” of suppressing a minimal number of subcarriers.

5.2 Detecting Impending Interference

The MicProtector must accurately and quickly measure SINR to notify the WSD of impending disruption before it occurs. The technical challenge of doing so is that the mic signal is constant and the FM nature shifts power in the band, making interference estimation directly in the band difficult. Given our goal of enabling coexistence between wideband WSDs and mics, interference will be wideband. Coexistence between narrowband WSD and narrowband primaries (mics) is a completely separate challenge, which our work is not looking to address. We are assuming wideband WSD that have OFDM capabilities to perform subcarrier suppression. Narrowband devices are unlikely to use OFDM, and are not WSDs that could follow our protocol. We would suggest that such narrowband follow the channel vacation rule.

Under this assumption, we can accurately detect the level of interference independent of the mic’s signal by measur-




Figure 14: Overview of SEISMIC adaptation protocol.

ing the power directly outside of the band. Since the operational band is small ($\sim 200\text{KHz}$), estimation directly outside the band in our system is expected to be accurate: prior work [15] has shown frequency selective fading can be severe (30dB) across 20MHz frequency ranges but remains modest ($< 1\text{dB}$) for the smaller 200KHz frequency range.

To perform this measurement, we introduce *control bands* at the MicProtector which are 25KHz bands on both sides of the mic’s operational band (see Fig. 13). Using these bands, the MicProtector can accurately measure the interference power generated from WSDs in range. Given that noise is additive, measuring the interference power of multiple WSDs is handled through the measurement in the control bands. Noise will be cumulative in the SINR measurement.

The MicProtector must monitor the squelch tone power, as shown in §4.2, audible disruption is caused when the interference level reaches the squelch tones. To do so, it measures the power in the frequency area of the squelch tones, which are approximately at a $\pm 32\text{KHz}$ offset from the center of the mic’s band and subtracts the interference power.

Finally, the MicProtector must be able to warn a WSD of *impending interference*, i.e., there is a *Protection Threshold* below the squelch tones upon which the MicProtector starts signaling to the WSD. Ideally, if the mic signal were stable and there was no delay in WSD adaptation, this threshold could be placed exactly 1dB below the squelch tones. However, this is not the case. In the time it takes a WSD to adapt, the mic signal could drop due to changes in the environment or mobility; or the WSD’s signal may increase. Therefore, the protection threshold needs to be more conservative to protect against fluctuation. In Section 6, we show that using a conservative threshold of 10dB below the squelch tones achieves all these goals. However, we also show in our evaluation (§7.1) that even if we wanted to select an even more conservative threshold (e.g., 20dB below the squelch tones), the loss of white space reuse would not be huge, and significant spectrum gains can still be achieved.

5.3 Adaptation Protocol

The goal and challenge of the adaptation protocol is to reuse the surrounding frequency around a mic’s transmission without ever creating an audible disruption. Such a task is non-trivial. When first entering a channel, if a WSD were to transmit at full power without knowing mic placement or what SINR values it could create, it could easily exceed a mic’s protection threshold and create an audible disruption.

Algorithm 1 Adaptation Algorithm at WSD:

S : Spectrum used by WSD, initially the entire desired spectrum.
 P : Transmit power used by WSD, initially at minimum level.
 ΔT : Ramp up time interval.
 ΔS : Amount of additional spectrum suppressed in each iteration.
 ΔP : Power increment in each iteration.

Ramp-up:

- 1: while P below desired power level and $S \neq \{\}$ do
- 2: wait for time ΔT
- 3: transmit underlay signal on spectrum S using power P .
- 4: if strobe $M(F_{Mic})$ received then
- 5: Suppress an additional ΔS of spectrum around F_{Mic} .
- 6: else
- 7: Increase P by ΔP ;
- 8: end if
- 9: end while

To overcome this, SEISMIC exploits the FM capture effect in mic systems where RF interference below the squelch tones is disruption-free. From this, we design *underlay probe packets* to the mic system, which reside under the mic signal. Such packets implicitly ask the mic system: “is this frequency usage at this power level acceptable?”

To converge without causing a disruption when first entering a channel, the WSD begins at minimal power (P) and transmits a probe packet.³ After a probe transmission, the WSD waits ΔT for an impending interference notification. Without notification, the WSD increases its transmission power by ΔP and transmits another probe packet. The ΔT time between each step is dependent on the time it takes to reliably detect impending interference notifications. In our SDR-based implementation (§6), we require ΔT to be $320\mu s$. However, this time could be significantly reduced in a hardware implementation (10s of μs). For ΔP , we find 2dB to be a reasonable increment, ensuring interference is increased slowly without significantly increasing convergence time. Through evaluation, $\Delta P=2\text{dB}$ achieves 16ms average convergence time.

If a notification of impending interference is received (i.e., interference power reached protection threshold), the WSD *must* suppress ΔS frequency, or back down its power. Ultimately, ΔS will be dependent on the parameters of the WSD. Using subcarrier suppression for a discontinuous waveform, ΔS can be no smaller than the width of a subcarrier (i.e., suppressing in smaller steps is not possible). We use a ΔS of 25KHz in our USRP2 WSD implementation (§6), which also matches our Adaptrum industry WSD subcarrier size. Note that the larger ΔS is, the more likely the WSD will suppress un-needed frequency. The smaller ΔS is, the WSD will achieve a closer-to-optimal amount of suppression. This process continues until convergence, illustrated in Figure 14.

Several comments are in order:

- By design, if the initial minimal power level does not cause disruption at the mic, the protocol is guaranteed to ensure no mic disruptions. We discuss in Section 5.5

³If the signal strength at the mic receiver is very weak, the initial lowest power level could create audible disruption at the mic receiver. We address this scenario in Section 5.5.

how we can guarantee disruption-freedom in *all* cases.

- The protocol converges to an optimal state, or a close approximation, i.e., full power with minimal suppression.
- The protocol works even in the presence of multiple mics and multiple MicProtectors. Whenever a strobe signal $M(F_{Mic})$ is received, the WSD blocks off additional spectrum around the mic centered at $M(F_{Mic})$. I.e., there can be multiple “holes” in the spectrum used by the WSD.
- As we show in Section 6, the ramp up time interval can be implemented to be short; so convergence is fast.

There are two more details to the protocol. First, when a new mic enters the channel, it may be within disruption range of a WSD. To prevent disruption, when the MicProtector is initialized with the mic system, it sends out a special strobe pattern (§5.4) which acts as a reset. Detecting the reset forces *all* WSDs in to the probing and ramp-up phase.

The final question is when a WSD can reclaim suppressed spectrum as mics leave the channel. Given that the frequency with which mics enter and leave a channel is typically prolonged (e.g., a concert or a lecture), the process does not need to happen often or quickly. To reclaim spectrum, the WSD can simply re-initialize its transmission power, un-suppress frequency, and then restart the adaptation protocol using more spectrum. Notice that with this method of reclaiming spectrum, SEISMIC is inherently robust and conservative: WSDs react to impending interference in the most conservative way (by suppressing more spectrum), but can reclaim spectrum only by resetting their power level to the minimal level and restarting the entire protocol anew.

5.4 Strobing: Notifying Impending Disruption

The previous sections have shown that SEISMIC relies on a signaling technique from the MicProtector to notify of impending disruptions. The signal must be simple, robust, and spectrum efficient; yet able to convey the necessary information (mic’s operational band and center frequency). Specifically, requiring the support of a complex protocol (e.g., 802.11) limits WSD and mic system design. The signaling should also happen in-band to remain efficient and avoid the WSD needing to tune to another frequency.

To meet these goals, we introduce a technique we refer to as *strobing*. It adds minimal complexity on both sides. It only requires basic power generation at the MicProtector, and simple carrier sense-like power detection at WSD. Furthermore, with thoughtful placement of the strobe signals in the control bands, the strobes can convey the necessary information for the WSD to adapt in a spectrum efficient manner.

Stobes resemble On/Off-Keying (OOK) and Morse codes, in which the power of a tone is quickly raised and lowered (i.e., a strobe light) in a pre-determined pattern to convey a signal. Patterns are generated using alternations of on- and off-symbols, where an on-symbol is the presence of a tone and the off-symbol is the absence of the tone. On- and off-symbol lengths are fixed in time, and unique patterns are generated by alternating the power (or presence) of the tone,

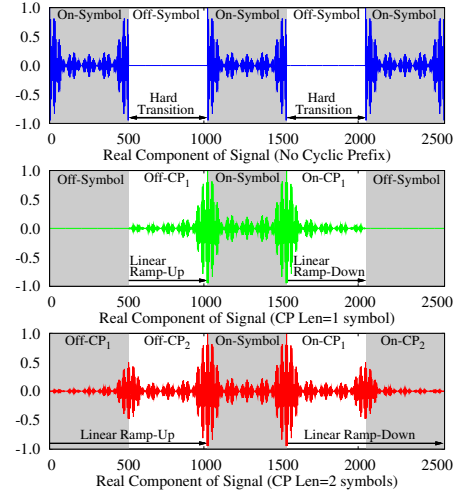


Figure 15: Strobes with varying cyclic prefixes in the time-domain.

for example: [1,0,1,0,1,0,...] or [1,1,0,0,1,1,0,0,...]. This is effectively changing the rate of the strobing, as a factor of the fixed symbol length. We provide a simple time-domain example at the top of Figure 15 with hard symbol transitions.

Simply, the presence of a strobe can signal of impending disruption from the MicProtector to a WSD. By strobing and simply changing the strobe rate, we create unique signals which do not mimic WSD behavior. Notice that by generating a strobe in the two control bands, both the center frequency and bandwidth of the mic are conveyed. The middle point between the strobes is the center frequency, and the distance between them is the bandwidth (see Figure 13).

Strobe Patterns: With a single MicProtector in a channel, generating two identical strobes in the control bands can properly convey location and width of the mic’s band. However, with more than one MicProtector strobing, where one band starts and another ends is difficult to determine (e.g., which strobe starts or ends a band?). In a planned environment where mics are spaced more than 500KHz apart, bands may be more easily distinguished. However, in unplanned environments where mics may be placed closer in frequency, their bands will be indistinguishable. To eliminate this problem, we use two different strobe patterns (i.e., rates) in the control bands: a start- and end-of-band pattern.

Safely Generating Strobes: The immediate concern of strobing is to ensure that the strobes sent by the MicProtector never interfere with the mic signal. We must ensure that the tones are generated in a way that no power is leaked in to the mic’s band. We find that a cyclic prefix is *critical* to ensuring this. From extensive evaluation using a nanosecond level sample capture, we find that hard symbol transitions create leakage in to the mic band and surrounding spectrum. Hard transitions are shown in the time domain at the top of Figure 15, and the resulting interference in the frequency domain is shown with the corresponding line in Figure 16. Clearly, the result shows interference in the mic band.

We find that from generating what we refer to as a linear power ramping cyclic prefix (LPR CP), we can eliminate this




Figure 16: Strokes with varying CP in the frequency-domain.

leakage in to the mic band. An LPR Off-CP is used to gradually scale the power up from an off-symbol to an on-symbol, and an LPR On-CP is used to gradually scale the power back down from an on-symbol to an off-symbol. The linear power scaling is done by scaling the complex samples in software (i.e., on the DSP or in the software of a software-defined radio) to avoid complications of fine-grained power control in hardware. The LPR CP is illustrated in the time domain in Figure 15, and the resulting frequency usage in Figure 16. As shown, using an LPR CP of size 2 removes this critical leakage, protecting from interference even at high TX power.

Strobe Detection: Strobe detection is similar to carrier sense-like functionality and pattern detection. After a probe, the WSD monitors the state of the channel broken down into 25KHz bins. The matching is done in an absolute manner, marking each bin as a 1 or 0 and over time matching a strobe pattern. This is very parallelizable in hardware.

5.5 Low-Power Mic Signals

If the mic signal is low and the squelch tones are barely above the noise floor, even a single (or multiple) new WSD’s transmitting probe packets at minimal power could create disruption at the mic. To avoid this, we introduce a final set of unique strobes which are proactively generated when the mic signal is low. The signal when the protection threshold is below the noise floor at the MicProtector, since this threshold represents the point at which interference above it threatens disruption. When a WSD detects a low-power strobe signal it vacates the channel. The low-power signal ends when the threshold goes above the noise floor.

5.6 Multiple WSDs & Mics

Multiple WSDs: So far, we have described the operation of SEISMIC in a scenario with a single interfering white space device. Clearly, in order for SEISMIC to be practical, it must also be robust in the presence of multiple WSDs. One may wonder whether the system is still sufficiently protective if many WSDs simultaneously ramp up their power.

Fortunately, *SEISMIC can handle any number of WSDs* and still guarantees full protection to the mic. Consider the following inductive argument. At some time T , there are n WSDs transmitting in proximity of the mic. Let the cumulative interference level I_T created by all these WSDs at the mic receiver be below the protection threshold. In the worst-case, all n WSDs simultaneously ramp up their transmission power once before the MicProtector is able to send a strobe signal. In this case, it is guaranteed that the new cumulative

interference level I_{T+1} is still below the mic’s squelch tones.

This is true because the adaptation in SEISMIC uses *multiplicative increases* in transmission power levels. By dB’s relative definition, any additive increase in dB corresponds to a multiplicative increase of power in mW. For example, an additive increase by 3dB corresponds to (roughly) 2x power in mW, and additive increase by 2dB is approximately 1.6x. Consequently, the interference power (in mW) of each of the n WSDs is increased at most by a multiplicative factor of ΔP ($\sim 1.6x$), and hence, the cumulative interference of all WSDs is $I_{T+1} \leq \Delta P \cdot I_T$. Thus, assuming the adaptation protocol is correct for a single WSD, it is also correct for n WSDs. In fact, observe that the more simultaneously transmitting WSDs, the smaller the variance and hence the more robust the protocol. The above computation only takes into account the effect of ramping up the transmission powers, but the exact same reduction from the n -WSD case to the 1-WSD case can also be made for the effects of mobility and/or fading.

Multiple Mics: SEISMIC also works in the presence of multiple mics (and thus multiple MicProtectors). In this case, the potential danger is that the strobing signals of one MicProtector could interfere with another mic in close proximity (thus causing disruption) because MicProtectors themselves do not actually follow the SEISMIC adaptation protocol before transmitting their strobing signals. Fortunately, it turns out that in practice, this is not an issue. Due to intermodulation interference⁴, proper frequency coordination in a location with multiple mics is essential, and coordination software for wireless mics ensure third-order and fifth-order harmonics are eliminated. A consequence of this coordination is that nearby mics will always be placed such that their frequencies are at least 500KHz apart, which leaves more than sufficient space for SEISMIC’s control bands.

5.7 Partial Deployment

An immediate concern of SEISMIC is deployment: the protocol relies on feedback from the MicProtector to suppress frequency. If a mic receiver does not have a MicProtector, a WSD will continue ramping up without suppression since it does not receive a notification of impending interference. An unlikely non-partial deployment solution to this problem would be to require *all* licensed mic systems to include the MicProtector, or risk WSD interference.

Instead, it is possible to partially deploy SEISMIC to enable protection over all mics (with or without the MicProtector), while also allowing coexistence with mic systems that have a MicProtector. Ultimately, the more mic systems that adopt the MicProtector over time, the more efficient the white space spectrum will become. To do so, licensed mics register in the database as being SEISMIC enabled or not. When entering a channel, the WSD consults the database to learn of *possible* mics within range. If all mic receivers are

⁴Intermodulation is a type of interference in which a receiver picks up two dissimilar frequencies that interact within the receiver’s electronics to produce sum and difference frequencies, including harmonics of these frequencies, which results in a whistling noise.

SEISMIC enabled, the WSD participates in the SEISMIC protocol. Otherwise, it must vacate the channel. A single non-SEISMIC mic system in the presence of many SEISMIC enabled systems reduces efficiency, however the average number of mics in range of a WSD is likely to be low.

6. IMPLEMENTATION

We implement a full prototype of the MicProtector and SEISMIC on the USRP2. We build a custom software stack for the key components: (i) measurement & analysis, (ii) strobe generation and detection, and (iii) the client with power ramping and suppression. We use the WBX daughterboard which is a full transceiver in the frequency range of 50MHz to 2.2GHz. We conduct all experiments over the air on TV channel 21, using an approved experimental license. Our SEISMIC parameters: $\Delta T=320\mu s$, $\Delta S=25\text{KHz}$, and $\Delta P=2\text{dB}$.

7. EVALUATION

We evaluate SEISMIC in several dimensions in this section. We begin with an evaluation of the protection threshold’s robustness. Then, live over-the-air experiments with a white space device running the SEISMIC protocol, and a mic system equipped with our MicProtector prototype. From this, we show the system’s efficiency and robustness to avoid disruptions in challenging scenarios. We conclude with a simulation using real mic placement data from 3 major events.

7.1 Impact of the Protection Threshold

The *protection threshold* at the MicProtector is set to allow a WSD to ramp up to proper suppression and operate without ever exceeding the power of the mic’s squelch tones. Without a buffer between this threshold and the squelch tones, variations in the mic’s signal power or the WSD’s interference level due to mobility, fading, etc., could cause the interference level to go above the squelch tones. There is an inherent trade-off when choosing this protection threshold. The lower the threshold, the more conservative the protection. The higher the protection threshold, the better the spectrum efficiency. To illustrate this, we perform a simple over-the-air experiment in which we vary the protection threshold and the interference from the WSD. As we see in Figure 17, for most cases the WSD can reuse most of the spectrum. When the WSD interference is high (e.g., -40 and -60 dBm) WSDs get reasonable spectrum only when the separation between the protection threshold and squelch tones is low.

Evaluating an Appropriate Protection Threshold: In determining the appropriate threshold, one has to consider the signal variation over the max transmission time of a WSD. The WSD cannot adapt during this time to ensure the interference does not exceed the squelch tones. To provide some insight, we performed a live mic experiment. Over a 60 second period, we walked to and from the mic receiver and swung the mic in fast movements to trigger quick signal variations. We calculated the maximum variation over 2ms and 10ms periods (i.e., max WSD TX times), and present

a CDF in Figure 18. This shows that over both periods the maximum variation we find is 4dB. The WSD could also be ramping up its power and probing, at a 2dB step. Despite the probes being much shorter in time ($\sim 10\text{s}$ of μs), we still account for this and now consider the minimum $4\text{dB}+2\text{dB}=6\text{dB}$. We add 4dB to account for other variations in the WSD power, and find 10dB to be sufficient.

7.2 Spectrum Efficiency Scenarios

Given a 10 dB protection threshold, we evaluate the spectrum efficiency achieved by SEISMIC under different scenarios. We cover the range of scenarios by varying the two components of SINR at the mic receiver: the mic’s signal at the receiver, and the WSD’s interference at the mic receiver. For the former, we vary the received power of the mic, and the latter is the power before frequency suppression. The resulting spectrum that can be used by the WSD is shown in Figure 19. When the mic’s squelch tone power is high, no suppression is needed and the entire 6 MHz of spectrum can be used. In most cases, only 250 KHz of spectrum needs to be suppressed. On increasing the WSD IBS power, it begins to overpower the mic signal, and therefore lesser spectrum is available for the WSD in order to protect the mic. Finally, the sharp cliff at the left occurs because SEISMIC is protecting the mic in low power (i.e., when protection threshold drops below the noise floor, causing WSDs to vacate - §5.3).

7.3 Live Experimentation with SEISMIC

To evaluate SEISMIC in a live setup, we use the MicProtector prototype paired with a Sennheiser mic system and our WSD running the SEISMIC protocol. We evaluate under two experimental setups: (1) moderate WSD interference (-70dBm) and mobile mic operation between distances of 10-30 feet, and (2) under a more challenging scenario with high WSD interference (-50dBm) and mobile mic operation between distances of 50-70 feet. Under mobility, we walk with the mic, lower the mic to hip level, raise it and speak in to it, turn our bodies, etc. A protection threshold of 10dB under the squelch tones is used, which we motivate in §7.1. Experiments are conducted for 5 minute time periods.

Results: The moderate WSD interference scenario where the mic is operated within 10-30 feet of the receiver is common in concerts (audio equipment is on/behind stage), and in lecture halls (mic receiver is near podium). Figure 20 illustrates the amount of spectrum that can be used by a WSD over a 30 second period. As shown, the available spectrum is both high and stable, despite the mobility of the mic. This is because the WSD suppression is adequate and the mic signal is strong. We also plot a CDF of the available spectrum in Figure 21 (Moderate). The gain is significant. The WSD used $>5.5\text{MHz}$ of spectrum for nearly 93% of its airtime.

The second scenario is more challenging. As we see in the time series of Figure 23, the mic’s squelch tones can at times be very low, e.g. at 12 seconds. Based on our protection threshold of 10 dB, and the noise floor of our USRP based MicProtector at -98 dBm, our system notifies of a low-power




Figure 17: Impact of protection threshold.




Figure 18: Variation in signal with mobility.




Figure 19: SEISMIC's spectrum gain.




Figure 20: Sample SEISMIC spectrum under moderate interference and mobile mic (10-30 ft).




Figure 21: Fraction of airtime with WSD having a given amount of spectrum with SEISMIC.




Figure 22: Spectrum inefficiency compared to ideal suppression as a fraction of airtime.

mic at -88 dBm (shown as a dotted line in the Figure).⁵ In these situations, such as at 12 and 24 seconds, the WSD vacates the entire TV channel. At other times, the WSD ramps up the power and uses the available spectrum. As we see in Figure 21, the WSD is able to use 5MHz of spectrum 75% of the time, and 4 MHz of spectrum 90% of the time. Even in this challenging scenario with heavy fluctuations of mic signal power, the protection threshold and SEISMIC protocol ensured zero audible microphone interference.

To highlight the efficiency of determining the proper number of subcarriers to suppress, we used information at the MicProtector to compute the optimal amount of frequency suppression and compared it to the spectrum the WSD actually used. Figure 22 shows that in the moderate scenario, the WSD is able to converge to within 25KHz of optimal suppression 72% of the time and within 100KHz 97% of the time. In the more challenging scenario (Figure 22), the WSD is able to converge to within 400KHz 89% of the time.

Finally, we evaluated the time it takes the WSD to reclaim spectrum after a mic leaves the channel. To do this, we allow the WSD to converge with the Sennheiser mic in the channel using SEISMIC, and then turn the microphone off. We repeated this experiment 25 times and found the average time to reclaim the spectrum 272ms.

7.4 SEISMIC's Efficiency with Many Mics

We now study the benefit of SEISMIC in heavy mic usage scenarios, which may typically be found in cities or on campuses. To evaluate these benefits, we obtained mic registration data for 3 major events: the 2008 NBA All Star Game (191 mics), the 2010 BCS Championship Bowl (108 mics), and the 2010 Worldwide Partner Conference (77 mics). The channel placement of these mics is shown in Figure 27.

Setup: To quantify SEISMIC's spectrum efficiency, we develop a simulation environment from the event data in which a MicProtector exists for every mic. Now we ask, if X WSDs are in this environment, each with various noise

⁵We note that in production systems, the noise floor over 400 KHz will be much lower, and can operate at lower squelch tones.




Figure 23: In the more challenging scenario, the WSD is able to avoid mic disruption, even at low mic powers, showing robustness.

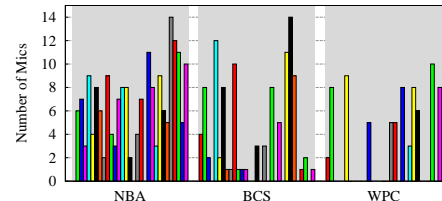


Figure 27: The number of mics per TV channel at each event, with the start and end of the white space channels highlighted.

powers at the mic receivers, what is the average amount of usable white space spectrum Y ? Two important characteristics are computed: (i) received squelch tone power at a mic's receiver, and (ii) WSD interference at each mic receiver. We weight mic signal strengths towards better-to-average (yet still have low signal mics), and generate the WSD powers uniform randomly between -110dBm and -20dBm. We account for WSD suppression to reduce interference on one mic, can contribute to reduced interference on another mic.

Results: Given that we know mic placement but not the active TV broadcasts, we evaluate SEISMIC by varying the number of channels available in each event. These results are presented in Figure 24, such that if only 10 channels were available in the area, SEISMIC achieves 44MHz compared to a max of 60MHz. The results show a significant gain in spectrum and promise for ensuring white space in highly dense areas where there are only 2 or 3 channels available.






Figure 24: Average SEISMIC client spectrum, Figure 25: Usable spectrum of the FCC ruling, Figure 26: CDF of the usable spectrum across given mic placement and channels available. perfect suppression, and avg. SEISMIC client. 100 SEISMIC enabled WSDs at each event.

Further exploring the possibilities of SEISMIC, we assume no TV broadcasts are active and present the resulting spectrum gains based on the mic placement data in comparison with vacating channels and “perfect” vacation of a mic’s operational band. As shown, SEISMIC can provide up to 21x the amount of spectrum over channel vacation (e.g., NBA event) and come close to “perfect” vacation. This is due to some WSDs not needing to suppress any subcarriers since their interference remains underneath the MicProtector’s protection threshold. To provide further insight in to this, in Figure 26, we show the available spectrum across the 1000 SEISMIC clients we simulate in each event. At the most dense event (NBA game), 50% of all clients have at least 130KHz. Only 5% of clients have less than 50MHz of spectrum. This highlights SEISMIC can significantly increase spectrum availability for white space networking.

8. CONCLUSION

SEISMIC can enable a significantly more spectrum-efficient use of the available white space spectrum in the TV bands, while coexisting in a disruption-free manner with mics. In particular, no channels would need to be reserved for mics. For this reason, we believe that the FCC should amend its ruling to allow WSDs to operate on the same TV channel as long as its power is below the squelch tones of the mic at the mic receiver; and that the white space protocols (e.g. IEEE 802.11af and IEEE 802.22), should be modified to ensure the power limits. Such changes are not unattainable. The FCC has shown its willingness to make changes to the ruling through its removal of the sensing requirement in the Second Order. To accomplish this, we have demonstrated SEISMIC to the FCC, including Chairman Genachowski, various mic operators who plan events such as the Super Bowl and mic manufacturers. In this context, it is encouraging to note that mic manufacturers such as Shure show great interest in a solution such as SEISMIC. A video demonstration of SEISMIC is also available.⁶ Through this effort, we hope to enable more spectrum efficient white space networking.

Acknowledgements: We would like to thank the following people. Our shepherd, Theodoros Salonidis, for his valuable feedback, as well as Srinivasan Seshan, Brian Padalino, and Daniel Halperin. Jason Bilyeu for providing us with a large number of mics to test with, and for connecting us with the professional mic users community. Finally, James Stoffo for valuable feedback on the design of SEISMIC.

⁶KNOWS. <http://research.microsoft.com/en-us/projects/KNOWS/>

9. REFERENCES

- [1] Audio-Technica Tutorial on FM Mic Squelching, <http://www.audio-technica.com/cms/site/557bfecb9ab0cfee/index.html/>.
- [2] Connecting America: The Nation Broadband Plan (FCC), <http://download.broadband.gov/plan/national-broadband-plan.pdf>.
- [3] Private communication with J. Stoffo and J. Eskew of Professional Wireless Systems, <http://www.professionalwireless.com/>.
- [4] Ultra-wideband Technology for Short-Range, High-Rate Wireless Communications, <http://www.ieee.or.com/Archive/uwb.pdf>.
- [5] Unlicensed Operation in the TV Broadcast Bands. May 2004.
- [6] Demonstration of a Prototype Dynamic Spectrum Access System, Philips Research. In *DySPAN demo session*, 2008.
- [7] FCC press release, Evaluation of the Performance of Prototype TV-Band White Space Devices. November 2008.
- [8] FCC press release, FCC Adopts Rules for Unlicensed Use of Television White Spaces. November 2008.
- [9] Mitigating the Effects of Unlicensed Devices on Wireless Microphones, Ahren Hartman and Edgar Reihl SHURE Incorporated, SHURE Incorporated, 2009.
- [10] FCC press release, FCC Frees Up Vacant TV Airwaves For “Super Wi-Fi” Technologies and Other Technologies. September 2010.
- [11] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White Space Networking with Wi-Fi like Connectivity. In *SIGCOMM*, 2009.
- [12] R. E. D. Borth and B. Oberlie. Considerations for Successful Cognitive Radio Systems in US TV White Space. In *DySpan*, 2008.
- [13] G. Buchwald, S. Kuffner, M. Brown, and E. C. L. Ecklund. The Design and Operation of the IEEE 802.22.1 Disabling Beacon for the Protection of TV Whitespace Incumbents. In *DySpan*, 2008.
- [14] R.-R. Chen and X. Liu. Coexisting with csma-based reactive primary users. In *the proceedings of DySpan 2010*.
- [15] D. Halperin et al. Predictable 802.11 packet delivery from wireless channel measurements. In *SIGCOMM 2010*.
- [16] G. Judd, X. Wang, and P. Steenkiste. Efficient Channel-aware Rate Adaptation in Dynamic Environments. In *MobiSys*, 2008.
- [17] K. Leentvaar and J. Flint. The capture effect in fm receivers. In *IEEE Transactions on Communications*, 1976.
- [18] S. Mishra, R. Brodersen, S. Brink, and R. Mahadevappa. Detect and avoid: an ultra-wideband/wimax coexistence mechanism. *Communications Magazine, IEEE*, 45(6):68–75, 2007.
- [19] S. Mishra, R. Tandra, and A. Sahai. Coexistence with primary users of different scales. In *DySPAN 2007*, pages 158–167, 2007.
- [20] S. Narlanka, R. Chandra, P. Bahl, and I. Ferrell. A Hardware Platform for Utilizing the TV Bands with a Wi-Fi Radio. In *IEEE LANMAN*, June 2007.
- [21] G. Nychis, T. Hottelier, Z. Yang, S. Seshan, and P. Steenkiste. Enabling mac protocol implementations on software-defined radios. In *the proceedings of NSDI*, 2009.
- [22] J. Poston and W. Horne. Discontiguous ofdm considerations for dynamic spectrum access in idle tv channels. In *DySPAN*, 2005.
- [23] H. Rahul et al. Learning to share: Narrowband-friendly wideband networks. In *SIGCOMM 2008*.
- [24] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra. Perceptual Evaluation of Speech Quality (PESQ)-A New Method for Speech Quality Assessment of Telephone Networks and Codecs. In *IEEE Conf. on Acoustics, Speech, and Signal Processing*, 2001.
- [25] L. Yang et al. Supporting demanding wireless applications with frequency-agile radios. In *Proc. of NSDI (2010)*.
- [26] W. Yu-chun et al. Protection of Wireless Microphones in IEEE 802.22 Cognitive Radio Network. In *ICC Workshops*, 2009.