

A Vehicle-based Measurement Framework for Enhancing Whitespace Spectrum Databases

Tan Zhang[#], Ning Leng[§], Suman Banerjee[#]
University of Wisconsin-Madison
{tzhang, suman}@cs.wisc.edu[#], leng@stat.wisc.edu[§]

ABSTRACT

The present TV whitespace networks rely on spectrum occupancy databases to determine their operating channels. In this paper, we show that such databases cause non-negligible wastage of whitespace spectrum. We also report that whitespace channels can have very different quality due to interference from secondary devices and the leakage from TV broadcasts. Such disparity in channel quality is not captured by existing databases. We propose the use of spectrum measurements to overcome the above limitations of databases. In particular, we describe a system called V-Scope that leverages spectrum sensors on public vehicles to collect and report measurements from the road. These measurements are used as “anchor points” to construct various models to better determine whitespace spectrum, estimate its channel quality, and validate locations of primary and secondary devices.

We have deployed our system on a single metro bus traveling across a mid-sized US city. Based on measurements collected at above 1 million locations over 120 square-km area, we find that a commercial database causes under-utilization of certain whitespace channels over a large area (up to 71% measured locations). Our system can reclaim this spectrum wastage at up to 59% locations, correctly selecting all the suitable whitespace channels at 72 – 83% locations, and achieving a localization accuracy between 16 – 27m.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Design studies, Measurement techniques

Keywords

TV Whitespace; Spectrum Database; Vehicular Sensing; Propagation Model; Channel Quality; Localization

1. INTRODUCTION

Through recent rulings, various spectrum regulatory agencies across the world have opened up a wide swath of spectrum in the UHF television band (512 MHz - 698 MHz), commonly referred

to as TV whitespaces, for unlicensed use. As demand for mobile and wireless connectivity continues to grow, such spectrum is going to be particularly attractive to address continued spectrum crunch. One of the fundamental issues in utilizing TV whitespace spectrum is to accurately determine vacant and high-quality channels. As per FCC’s ruling [5], the secondary, whitespace devices must not interfere with licensed transmissions from primary spectrum incumbents, i.e., TV broadcasts and licensed wireless microphones. A current FCC-preferred mechanism for determining vacant channels is to have secondary devices to query *spectrum occupancy databases*. These databases are operated by some third-party companies following the FCC’s guideline. They rely on a *same*, widely-used propagation model (R6602 [3]) to predict the coverage contour of TV broadcasts. In addition, they set up a fixed (2km) protection contour around licensed microphones and reserve two channels nationwide for their exclusive usage. A channel is concluded to be whitespace if a secondary device is outside the predicted contours of all the primary devices.

Being based solely on a propagation model, such databases are likely to have errors in predicting the availability of whitespace spectrum. Further, the databases have no attempt to distinguish the quality of whitespace channels, nor are they responsible for validating the operating locations of primary and secondary devices. In this work, we explore the use of spectrum measurements combined with propagation models for augmenting databases to i) better determine whitespace spectrum, ii) estimate the quality of whitespace channels, and iii) localize primary and secondary devices.

In particular, we present *V-Scope* (Vehicular Spectrum Scope), a measurement system that leverages public transit buses to carry spectrum sensors and collect measurements opportunistically as they travel¹. The unique advantage of this approach is that each mobile sensor can add a proportional volume of useful measurements over a relatively *large* area. In addition, these measurements are likely to remain useful for some time, especially for TV broadcasts with a stable signal strength (e.g., < 5dB variation across a 6-week period). We therefore utilize these opportunistic measurements to refine various signal strength models, which can be used by existing databases to efficiently manage outdoor whitespaces.

Whitespace database limitations and V-Scope approach: Existing databases have inevitable inaccuracy in predicting the signal strength of primary devices. This is because the underlying propagation model is unable to capture the environment-induced variation, e.g., shadowing and multipath fading of specific contours, objects, and topologies. Since protecting primary users is at the top

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiCom’14, September 7-11, 2014, Maui, Hawaii, USA.
Copyright 2014 ACM 978-1-4503-2783-1/14/09 ...\$15.00.

¹Of course, public transit buses are just one of many possibilities for vehicles that can carry spectrum sensors; other potential examples are mail delivery trucks, taxicabs, and many other third party services that scour different city roads.

priority, the databases have to use a very *conservative* model configuration for predicting the coverage of primary devices, which causes under-utilization of some whitespace channels over large area (up to 71%). V-Scope uses two related techniques to improve the accuracy of databases with spectrum measurements. (i) It uses an enhanced version of feature detection to accurately detect primary signals while measuring their power up to the FCC-mandated sensing threshold (-114dBm [5]). (ii) Based on the measured power of primary devices, V-Scope refines the parameters of a propagation model, which can be used by databases to better predict the signal strength of primary devices. In our design, we have developed a *region specific model fitting* procedure to construct a fine-tuned propagation model for each local environment, and adopted a weighted regression algorithm to tolerate the non-uniform distribution of vehicular measurements in model fitting.

Additional enhancements to databases: V-Scope can provide databases with two additional functions enabled by these measurements. First, it allows databases to estimate the quality of whitespace channels. Our measurements show that the noise power in whitespace channels can differ by 40dB at a given location, which can significantly affect the performance of a whitespace link. The large variation in channel quality is contributed by the *co-channel interference* from unlicensed devices (whitespace devices and unlicensed microphones), and *adjacent-channel* leakage from TV broadcasts. Existing databases can not predict this co-channel interference because the precise power information of unlicensed devices is not available. Similarly, they are not able to estimate the leakage of TV broadcasts that depends on specific hardware attributes of a TV transmitter, e.g., sharpness of the front-end filter. To predict co-channel interference, V-Scope uses the same model fitting procedure described above to construct signal strength models for each unlicensed device. For adjacent-channel leakage, V-Scope constructs a leakage model based on an inherent power relationship between a TV signal and its leakage, which is identified by measurements. This model takes the predicted power of a TV broadcast to accurately estimate its leakage power at a same location.

The second function offered by V-Scope is to localize primary and secondary devices. This location information is necessary for constructing propagation models to predict the availability and quality of whitespace channels. However, existing databases do not store such information for most of unlicensed devices². It may also have inaccurate information for primary devices due to various reasons such as incorrect entries to databases and transmitter reallocation. V-Scope utilizes the measured signal strength of these devices to pinpoint their location, thereby providing database operators additional means to obtain and validate such information. We leverage a RSSI modeling technique commonly used in indoor localization systems [1, 16]. An unique challenge in applying this technique to our outdoor scenario is to correctly capture the *large-scale path loss* in front of environment-induced variation, e.g., terrain elevation, shadowing (§ 3.4). Existing model based techniques use all the measurements *indiscriminately* to construct a propagation model. This model is consequently biased by a noisy propagation trend leading to large error (up to 100m). We have enhanced the standard approach by carefully selecting measurements from *certain radiation sectors* and constructing a *sector-specific* model based on the chosen sectors to improve localization accuracy.

While our opportunistic wardriving approach has some unique advantages in collecting wide area measurements, a similar goal can certainly be achieved by other approaches such as local spectrum sensing and crowd-sourcing mobile whitespace devices. Re-

²Only those fixed, high-power whitespace devices can register their operating locations in the databases [5].

gardless of how measurements are collected, most of our proposed techniques remain useful for enhancing databases based on these measurements. In addition, our techniques can be beneficial to indoor whitespace determination systems like WISER [23].

Key contributions: We have deployed V-Scope on a single public transit bus traveling in a mid-sized US city for a 6-week period. Since our bus operator tends to rotate their buses through multiple routes in the course of each day, we have been able to collect spectrum measurements at more than one million distinct locations (with different pairs of latitude and longitude reported by a GPS device) over a 120 sq. km. area in and around our city.

Based on these measurements and the combination of our proposed techniques, we make the following contributions: (i) we show that commercial databases that are based solely on propagation models cause under-utilization of whitespace spectrum over a wide area (up to 71% measured locations for protecting TV and 70% locations for protecting microphones); (ii) V-Scope can reduce spectrum waste in protecting TV broadcasts by up to 59% locations; (iii) It can identify all the suitable whitespace channels at 72 – 83% locations under different channel quality constraints; and (iv) V-Scope can pinpoint TV-band devices at various locations with low error (16 – 27m), improving state-of-the-art localization techniques by 1.2 – 3.5 \times in outdoor scenario.

2. LIMITATION OF EXISTING SPECTRUM OCCUPANCY DATABASES

In this section, we motivate the design of V-Scope by demonstrating the two major limitations in existing spectrum databases. First, to tolerate the inaccuracy of propagation models, the databases assign very conservative protection contours for primary devices, which cause non-negligible wastage in whitespace utilization. Second, they are agnostic to the quality of whitespace channels, which can differ significantly at a given location, and thus can have a large impact on the performance of a whitespace network. We start by describing our measurement setup and datasets.

2.1 Measurement setup

Our current deployment consists of a server and a client as shown in Figure 1. The client has been deployed on a metro bus traveling in and around a mid-sized US city for a 6-week period. It uses a laptop to configure a portable spectrum analyzer (WSA4000 [20]) for collecting spectrum samples (FFTs) in all the UHF channels. An omni-directional TV antenna with 0dB gain is used to capture signals from all the directions. Based on these spectrum samples of a channel, the laptop performs some real-time analysis (§ 3.1), i.e., primary detection and power estimation, while obtaining the measured location from a GPS module. It uploads the GPS reading and measurement results to our server over cellular networks. The server is situated in our laboratory, with an Ethernet to receive the measurement results. It queries a commercial database [19] at measured locations to evaluate the database’s accuracy.

Datasets: Using this measurement setup, we have collected three datasets from different vehicles. Table 1 summarizes their salient features. *Dataset A* is our main dataset collected from the metro bus, and covers about one million distinct locations over a 120 sq. km. area. The data consists of signal type, power and the database’s prediction for all (30) UHF channels at each measured location. We will use this dataset to evaluate the performance of the database and V-Scope as well as studying whitespace channel quality. *Dataset B* contains measurements in a whitespace channel where our whitespace radio was transmitting 12Mbps UDP packets using the 802.11 baseband technology. This whitespace device

Dataset	Measured Locations	Coverage (sq. km.)	Channels	Vehicle	Target Device
A	1 million	120	30 UHF channels	A metro bus	All TV-band devices
B	3.5K	3	Channel 42	A personal vehicle	A 3.8W whitespace transmitter atop a 8-floor building
C	2.5K	2	Channel 28	A personal vehicle	A 100mW microphone in 5 buildings

Table 1: Summary of datasets.

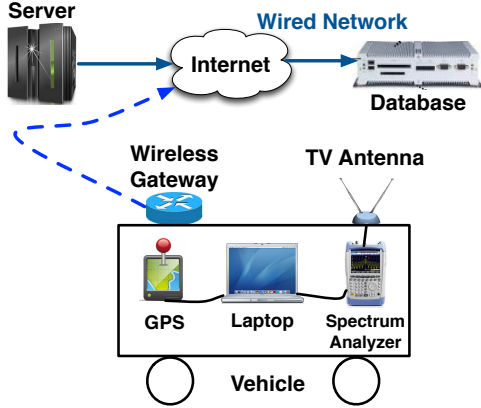


Figure 1: System architecture of V-Scope.

was mounted atop a 8-floor building with a transmit power of 3.8W. *Dataset C* contains measurements in a whitespace channel for a microphone device transmitting at 100mW in 5 different buildings at a downtown area. Dataset B and C were collected from a personal vehicle driving along multiple road stretches, and comprise the signal type and channel power at each measured location. We will use them to benchmark our measurement clustering algorithm (§ 3.2) and localization technique (§ 3.4).

2.2 Inefficiency in protecting primary users

We study the performance of a commercial database (SpectrumBridge [19]) in predicting TV and licensed microphones respectively. Note that other FCC-approved databases would produce the same prediction results. Our focus is to understand its performance tradeoff between the safety in primary protection and efficiency in whitespace utilization. All the evaluation in this section is based on dataset A unless otherwise mentioned.

Performance in protecting TV broadcasts: Borrowing the definition from prior work [10], we divide the errors in whitespace determination into two types, i.e., *false positive* and *false negative* as summarized in Table 2. A false positive is a location where the database mis-predicts an occupied channel (with measured power $\geq -114\text{dBm}$) as whitespaces, whereas the opposite being a false negative. We find negligible false positives ($<0.29\%$) of database in predicting all the TV broadcasts, which is similar to a prior report [10]. Thus, existing databases can reliably protect TV broadcasts. However, Figure 2(a) shows a 13–71% false negative rate in half of these channels, indicating these channels being unnecessarily blocked for unlicensed usage over a wide area. Figure 2(b) shows a typical spatial distribution of false negatives in one such channel, along with those locations with this channel correctly predicted to be whitespace (true negatives). We note most of these false negatives are at the north-east side, which is at the *closer* side to the TV tower. Thus, we believe these false negatives are caused by over-provisioning the protection contour of TV broadcasts.

Performance in protecting licensed microphones: To quantify false negatives in protecting licensed microphones, we cal-

	Ground truth	Prediction
False Positive	Occupied	Vacant
False Negative	Vacant	Occupied

Table 2: Two types of errors in whitespace determination.

culate the power of about 200k measurements collected in their reserved channels. Figure 2(c) shows that about 70% measurements have very low power, suggesting that microphone signals are unlikely to be present in many reserved locations. While these wardriving measurements might not capture every instance of microphone transmissions, they are statistically sufficient to suggest that licensed microphones are not making efficient use of their reserved channel, possibly due to coarse-grained reservation periods (24 hours) and over-sized protection contours (2km). On the other hand, we find no other type of signals in these channels, suggesting the database to be sufficient in protecting licensed microphones.

2.3 Variation in whitespace channel quality

We now study the variation in whitespace channel quality, which are caused by *co-channel interference* from secondary devices and *adjacent-channel leakage* from TV broadcasts.

Magnitude of quality variation: Figure 3 shows the CDF of absolute differences in noise power between the best channel and the worst channel at each measured location. *Overall*, *TV-adjacent*, and *Non-adjacent* choose the worst channel from all the whitespace channels, adjacent channels to TV broadcasts, and non-adjacent channels, respectively. All of them pick the best channel from all the whitespace channels. We observe that the worst channel in *Overall* has a 8dB higher power in median and 17dB higher power in 90-quartile compared to the best channel. Such variation is currently contributed by unlicensed microphone transmission and TV broadcast leakage. We also note the non-adjacent channels generally have better quality than adjacent channels due to the absence of broadcast leakage. However, its worst channel still has 8dB higher power than the best channel at 10% locations due to unlicensed microphones. We envision this quality variation to become more prominent with the future proliferation of whitespace devices.

Prevalence of quality variation: Figure 4 shows the CDF of the amount of whitespace spectrum below different channel power based on our measurements and the database. We observe that for about 80% locations, the amount of whitespace spectrum identified by measurements is similar across different power thresholds, which is about 18 – 24MHz more than that predicted by the database³. For the rest 20% locations, this measured amount differs by 12 – 48MHz (2 – 8 channels) between the highest and the lowest power threshold (-65dBm and -80dBm). Thus, whitespace devices are likely to operate in a low-quality channel at these locations, if they randomly pick a channel suggested by the database.

³Here we did not consider spectrum wastage in protecting licensed microphones.

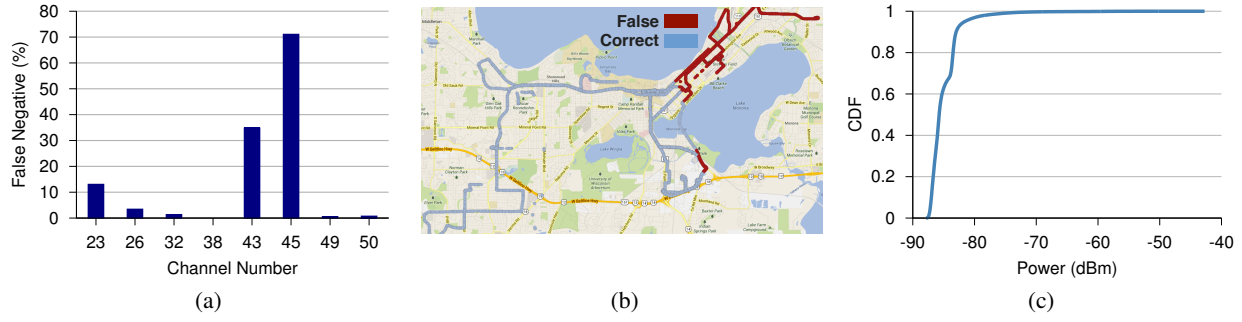


Figure 2: Wastage of whitespace spectrum by existing databases for protecting primary incumbents. (a) False negative rates in predicting TV broadcasts. (b) Spatial distribution of false negatives and true negatives in channel 43. (c) CDF of power measurements collected in reserved channels for licensed microphones.




Figure 3: CDF of the maximum power differences of whitespace channels at different measured locations.

3. V-SCOPE DESIGN

V-Scope is based on our measurement architecture described in § 2.1, and consists of following components to augment databases with measurements — i) primary detection, ii) measurement clustering, iii) propagation model refinement, iv) localization, and v) leakage model construction. Figure 5 shows its operation flow.

Overview: Our proposed system leverages a few clients mounted on public vehicles to detect primary and secondary devices based on spectrum samples collected in different UHF channels (§ 3.1). The detection results that include the device type and power of each channel are forwarded to our server. The server groups these measurements according to different device types in each UHF channel, and segregates each group for individual devices (§ 3.2). It then localizes those (secondary) devices if their location is not available in databases (§ 3.4). Using each device’s location and its associated measurements, the server constructs a propagation model that is tailored to local environment to predict its signal strength (§ 3.3). It also builds a model to predict the power of adjacent-channel leakage from TV broadcasts (§ 3.5). To predict whitespace channels in the vicinity of measurements, the databases can use our refined propagation models to better estimate the coverage of TV broadcasts. Similarly, the databases may use our models constructed for secondary users to predict their in-band interference, while leveraging the leakage model to estimate the adjacent-channel leakage from TV broadcasts. The sum of the predicted power of these interference signals is the quality of a whitespace channel. To predict for those locations far from measurements, the databases may switch back to use its default model (R6602). Finally, we envision the databases to use its current mechanism for protecting licensed microphones due to the limitations of our wardriving approach as will be discussed in § 3.1. We now explain each of these components in detail.

Figure 4: CDF of the available whitespace spectrum at different locations based on measurements and database prediction.

3.1 Zoom-in pilot tracking algorithm for primary detection

Our measurement module aims to utilize spectrum samples for detecting different types of primary transmitters and unlicensed devices (i.e., whitespace devices and unlicensed microphones). We accomplish this task by leveraging feature detection algorithms [8, 11] that identify different signals based on their spectral features. While such a technique performs well for detecting unlicensed signals, we find non-trivial challenges in satisfying the FCC’s requirement for detecting primary signals at up to -114dBm in our measurement setup. Such a stringent detection threshold is to tolerate the sensing inaccuracy caused by fading and shadowing, and to take account the transmission range of a whitespace device. We use TV detection to illustrate this challenge and our specific solution.

Existing algorithms detect a TV signal by tracking its pilot in the spectral domain. A pilot is a set of preambles in a TV packet, which produces a predominant peak at a fixed frequency. Unfortunately, for a TV signal close to the detection threshold, even this peak can be overwhelmed by noise and thus unable to be captured. To illustrate this, we attenuated a TV signal to be about -114dBm, while collecting its spectrum over a 6MHz TV channel with the largest number of FFTs (32768) available in our high-end spectrum analyzer [20]. As shown in Figure 6(a), the pilot of this TV spectrum is hard to be distinguished from the noise floor. This is because the noise spectral density is -159dBm/Hz in our spectrum analyzer. When using 32768 FFTs to represent a 6MHz band, the bandwidth per FFT bin is 183Hz (6MHz/32768) and thus the noise power per bin is about -135dBm. As will be discussed later, the pilot power is usually about 10 - 15dB lower than the total power of a TV signal, and thus can be as low as -129dBm. Comparing to the -135dBm noise floor, the 6dB stronger pilot is not robust enough for toler-

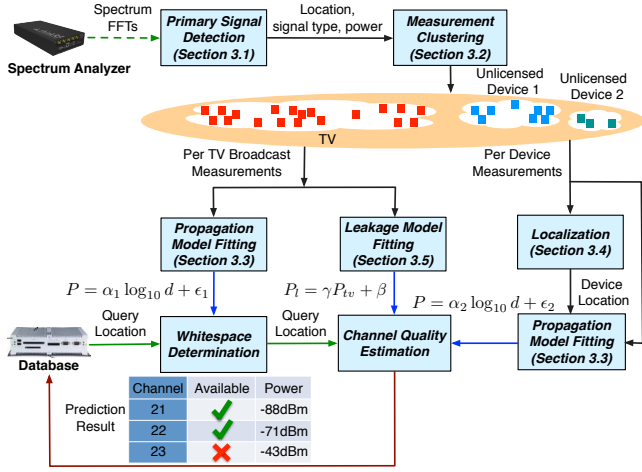


Figure 5: Flow of operations in V-Scope.

ating the signal strength fluctuation due to fading and shadowing. Prior systems [11] have overcome this challenge by leveraging low-noise amplifiers to amplify a weak signal by a fixed amount (e.g., 20dB). But such an approach is observed to cause *saturation* of a spectrum analyzer when capturing strong signals ($> -35\text{dBm}$), leading to distorted spectrum that can significantly degrade detection performance. The alternative solution of using tunable amplifiers would largely increase design complexity and hardware cost.

Zoom-in pilot tracking algorithm: Instead of resorting to signal amplification, V-Scope improves the sensitivity of feature detection by capturing the most prominent portion of spectrum with the available FFTs. To detect a TV signal, for example, we configure a spectrum analyzer to capture at a narrow band (488KHz) around its pilot frequency. This can effectively improve spectral resolution while reducing the noise floor, thereby producing a clear pilot peak as shown in Figure 6(b). Since this peak is well distinguishable at the detection threshold after zoom-in, V-Scope uses it as an unique feature to detect TV signals.

The detected pilot is, however, insufficient to determine whether the power of a TV signal to be above or below the -114dBm threshold. We also need the precise power information of a primary signal for constructing its signal strength models (§ 3.3). To achieve this goal, V-Scope leverages the power of a TV pilot to estimate its total power. According to the digital and analog TV standards [2], there is a fixed power offset between a TV signal and its pilot. For example, the pilot of a digital TV is required to be 11.3dB lower than its total power. Figure 7 shows this relationship indeed holds for a DTV signal at a wide range of power, albeit with some variation (10 – 15dB). Thus, V-Scope computes the total power of a TV signal by adding to its pilot power a constant offset η (20dB in our implementation). The estimated power is then compared with the -114dBm detection threshold to determine TV whitespaces.

Our proposed technique can be applied to microphone detection in a similar way. Briefly, we start by examining the 6MHz wide spectrum of each TV channel for detecting potential microphone tones that are narrow spikes carrying audio signals. Since these spikes can also be a result of noise fluctuation, we perform a narrow-band capture around each spike, extracting various features from its zoom-in spectrum for further validation. Once a microphone signal is confirmed, we use the power of its tones in the zoom-in spectrum as its total power because these tones contain most of the power ($> 95\%$) of a microphone signal. Despite




Figure 6: Different spectrum captures of a -114dBm digital TV signal. (a) Full-channel capture; (b) Zoom-in capture at the first 488KHz band. Both captures consist of 32768 FFTs.

Detected Ground truth	Digital TV	Analog TV	Microphone
Digital TV	94.9%	0.7%	4.4%
Analog TV	0.5%	97.4%	2.1%
Microphone	1.2%	0.7%	98.1%

Table 3: Accuracy of primary detection algorithm.

the effectiveness of this detection technique, V-Scope might not be able to sense all the intermittent microphone transmissions because a client visits each location only for a short amount of time. To guarantee sufficient protection, we envision the databases to protect licensed microphones as is. As our future work, we intend to enhance our system with static sensors deployed in the proximity of microphone devices to reliably capture their transmission.

To benchmark the accuracy of our detection algorithms, we collected spectrum data from 30 UHF channels at multiple locations. We established the ground truth results by using a TV receiver and our microphone transmitters. The identified primary signals were further attenuated for constructing spectrum traces at a wide range of power (-40dBm to -114dBm). A standard cross validation was then performed by randomly choosing 90% spectrum traces to detect the rest 10%. Table 3 shows the accuracy of our detection algorithm. We observe reasonably low error rates ($< 5\%$) in detecting different types of signals.

Summarizing, our measurement procedure works as follows. (a) A client captures spectrum fragments around pilot frequencies in each UHF channel. (b) It extracts a potential pilot by searching for the maximum FFT bin and including all the surrounding bins with power above a threshold. (c) From the obtained FFT bins, several features (e.g., power, center frequency) are extracted and fed to a classifier for detecting TV pilots and broadcast type (analog or digital). (d) If a pilot is detected, the client estimates its total power by adding a specific power offset η to the pilot power. (e) The estimated power is then compared with -114dBm threshold to determine the presence of TV signals. (f) If no TV signal is detected in a given channel, the client further captures a full-channel spectrum to detect microphones and unlicensed devices while measuring their in-band power. (g) If the full-channel spectrum contain potential microphone tones, a narrow-band capture is performed around each spike for validation. The entire procedure takes less than 1 second for processing all 30 UHF channels at a location, and produces results of *signal type* and *power* in each channel.

3.2 Measurement clustering

Since a V-Scope client has classified measurements to different device types (i.e., TV, microphone, whitespace devices), our clustering module further segregates measurements among devices of a




Figure 7: CDF of differences between the power of a digital TV signal and that of its pilot.




Figure 8: Accuracy of a fitted region model and a global model in predicting the power of TV signals.




Figure 9: Region models fitted by linear regression and weighted regression in predicting a TV broadcast.

same type. We developed a preliminary algorithm that leverages k -medoids clustering based on the *power* and *inter-distance* of measurements. k -medoids clustering aims to minimize the sum of pairwise dissimilarities of measurements in each cluster. We define the dissimilarity metric between pair of measurements using gower dissimilarity score [6]. A unique feature of gower dissimilarity score is its capability of assigning different weights to observations. By assigning higher weights to stronger measurements, our algorithm is robust to noise and outliers in choosing the center of clusters. Specially, we take only measurements with power $P_n > -90$ dB to reduce computational overhead. With N selected measurements, we first calculate distance between each pair of measures. This gives us a N -by- N symmetric matrix, in which the (i, n) th entry d_{in} indicates physical distance between measure i and n . We call columns of this matrix as variables, gower dissimilarity score between the i th and j th measure (row) is defined as a weighted sum of dissimilarities for each variable $\frac{\sum_{n=1, \dots, N} w_n D_{ij}^n}{\sum_{n=1, \dots, N} w_n}$, in which D_{ij}^n is defined as $|d_{in} - d_{jn}|$ and the weight of variable n is defined as $w_n = \log(P_n + 90)$. We iteratively apply this algorithm with different number of clusters k between $1 - k_{max}$, and determine the optimal k that can maximize the *silhouette* distance [17]. k_{max} is set to be 100 in our implementation and a R package [9] is used to perform clustering. To evaluate this algorithm, we merged groups of measurements corresponding to different transmission locations of our microphone device in dataset C. We applied our algorithm on the aggregated measurements, and found only 6% measurements mis-classified to a different location. One limitation of our current approach is that mis-classification may happen when devices are close to each other. However, the mis-classified measurements are generally far from a device location and have relatively lower power, thus imposing limited impact on the overall prediction accuracy of V-Scope.

3.3 Region model for predicting signal strength of different devices

Using the clustered measurements of each device, V-Scope refines the parameters of a propagation model that can be used to better predict its signal strength in the vicinity of measurements. To choose an appropriate model, we note that most of the UHF-band models can be generalized in a form of $P = \alpha \log_{10}(d) + \epsilon$, where P is the power of a device at a reception location, d is the distance between this location to the device, α is the rate at which the signal power attenuates over an increasing distance, and ϵ captures both the transmission power of a device and the fixed attenuation of environmental shadowing. The difference of these models lies in how the parameters α, ϵ are determined, e.g., based on an-




Figure 10: Illustration of region models in V-Scope.

tenna height (in Egli, Hata) and various environmental factors (in Longley-Rice, R6602). Since the measured signal strength can best reflect a propagation environment, V-Scope uses this general model and calculates α, ϵ based on measurements.

A standard model fitting procedure [13] solves α, ϵ by plugging measurements P_i, d_i into a given propagation model. This forms a set of linear equations $P_i = \alpha \log_{10} d_i + \epsilon$. It uses least-squares linear regression to calculate α, ϵ , with the objective of minimizing the squared sum of fitting errors $\sum_i (P_i - \alpha \log_{10} d_i - \epsilon)^2$. V-Scope improves this approach by (i) fitting an individual set of parameters for each local area to better model its propagation environment, and (ii) performing a weighted regression to avoid fitting bias caused by non-uniform distribution of vehicular measurements.

Region-specific model: V-Scope groups measurements into road segments and fits a different set of parameters for each segment as shown in Figure 10. The motivation is that these regions are likely to have different propagation characteristics, especially in an urban environment. This region-specific variation can hardly be captured by a global propagation model using a single set of parameters. To demonstrate this, we fit different α and ϵ based on measured signal strength of TV broadcasts in each 100m road segment. We compare this model with a global model fitted with all the measurements in predicting the strength of TV signals. Figure 8 shows that our region model achieves a median error of 1.4dB and 75 quartile error of 2.6dB, which are $3 \times$ and $2.9 \times$ lower than a global model. Thus, V-Scope fits a different set of model parameters (α, ϵ) to better capture each local propagation environment.

Weighted regression fitting: In fitting a region model, we note the non-evenly spaced vehicular measurements can degrade the performance of linear regression. Since a public vehicle drives at a varying speed and stops quite often, V-Scope collects measurements at non-uniform density. This causes linear regression to produce a biased model that favors *densely* measured area, and has




Figure 11: Measured signal strength vs. propagation range (in log scale) of a whitespace transmitter.




Figure 12: Measurement locations for the whitespace transmitter. A darker color indicates higher power. Two radiation sectors in Figure 11 are also marked.




Figure 13: Absolute correlation of all the sectors partitioned from the centroid location with different angles.




Figure 14: Power of a TV broadcast vs. power of its adjacent-channel leakage at different locations.

large errors at *sparingly* measured area. The underlying reason is that linear regression aims to minimize the squared sum of fitting error; the area with sparse measurements contribute less to this squared sum, thus being under-fitted. Figure 9 shows an example where we fit a model for a 100m region with most of training measurements collocated at a bus top. We observe that the model fitted by linear regression has up to 36dB error in predicting TV signal strength at the testing locations.

V-Scope uses weighted regression to compensate the effect of non-uniform measurement density in model fitting. The algorithm assigns a weight W_i to each measurement i , with the objective of minimizing the weighted squared sum of fitting errors $\sum_i W_i (P_i - \alpha \log_{10} d_i - \epsilon)$. A higher weight can indicate greater importance in fitting a measurement. Therefore, we assign *higher* weights to *sparse* measurements to compensate for the difference in measurement density in model fitting. To accurately capture this measurement sparsity, we calculate the weight of a measurement based on its total distance to other measurements $W_i = \sum_j \text{dist}(i, j)$. Figure 9 shows our model fitted by weighted regression achieves high prediction accuracy at all the testing locations.

To recapitulate, our model fitting procedure first bins measurements of each device into road segments. For each segment, it computes a weight for each measurement based on its distance to other measurements. It then takes these measurements and their weights as the input to weighted least-squares regression for constructing region models.

The fitted model will then be used to predict for a m -by- m square-shape region centered around the measured road segment. The region size m is the length of the road segment. We will explore the performance tradeoff between the prediction accuracy and stor-

age overhead of different region sizes in § 4.1. Finally, beyond the coverage of these regions, the database may use its default model (R6602) for predicting TV whitespaces.

3.4 Sector based localization

The above model fitting procedure requires the location of a transmitting device to calculate the distance d_i for each measurement. V-Scope leverages popular *RSSI modeling* techniques [1, 16] to localize a device if its location is unknown or awaits validation. Our motivation of using this technique comes from its flexibility in localizing a device from arbitrarily measured locations, and the simplicity of using signal strength (RSSI) information that is readily available from spectrum measurements. Such an approach works similarly to model fitting. It uses signal strength measurements and a propagation model to form a set of equations $P_i = \alpha \log_{10} d_i + \epsilon$. Here d_i is replaced with $\sqrt{(x_t - x_i)^2 + (y_t - y_i)^2}$, where (x_t, y_t) is the transmitter location and (x_i, y_i) is the location of each measurement. The algorithm solves $\alpha, \epsilon, x_t, y_t$ to estimate a device's location. V-Scope adapts this approach to outdoor scenario by (i) carefully selecting measurements in certain radiation sectors of a device and (ii) constructing a *sector-specific* model based on the chosen sectors.

While a RSSI modeling approach is reported to achieve a high accuracy within few meters in an indoor scenario, directly applying this technique to our outdoor area leads to large localization error (up to 100m in § 4.2.2). The reason is that its underlying propagation model matches poorly with the signal strength pattern of a transmitting device in an outdoor scenario. We use the wardriving measurements collected for our whitespace device in dataset B to demonstrate this. Figure 11(a) shows the received signal strength over different transmission ranges for all the measurements. A

propagation model expects a *linear* trend of the measured signal strength P_i over an increasing distance $\log_{10} d_i$, but we observe many measurements deviate from this trend (the fitted line). We use Pearson correlation coefficient to quantify the linearity of these measurements, with 1 and -1 being an extract positive and negative linear trend and 0 implying no correlation. We find the correlation value to be merely -0.57.

Such a poor linear trend is caused by surrounding environment, e.g., terrain elevation, obstacles, etc. To illustrate this, we decompose measurements according to different radiation sectors of the device with 10-degree angle as shown in Figure 12. Figure 11(b) is a scatter plot for measurements from one such sector. We observe a sharp drop of signal strength at a distance between 350 – 500m due to a building blocking the transmission, but a drastic increase in signal strength beyond 500m due to the rising terrain elevation coupled with diminishing blocking effect of the building at faraway locations. This environmental-induced variation can largely *perturb* a propagation model in capturing the large-scale path loss, leading to large localization error. Fortunately, we find measurements in some other radiation sectors present a less noisy propagation trend as show in Figure 11(c). This is owing to the environmental shadowing being unlikely to affect the signal along all the directions. Thus, V-Scope only uses measurements in a few sectors that present a good propagation trend for localization.

Localization procedure: Our proposed algorithm proceeds in five steps. (a) We use the *centroid* location of the *strongest* measurements (top 5dB) as the partitioning center, while grouping measurements into radiation sectors at a narrow angle (10 degree). The intuition behind this heuristic approach is that the centroid location is usually not far from a device’s actual location (e.g., 100m in Figure 12); and by using a narrow angle, some radiation sectors are likely to “fall through” the gaps of a blocking environment, thus having measurements following a good propagation trend. To validate this intuition, Figure 13 shows the *absolute* correlation of measurements in different sectors partitioned at various angles for the whitespace device. We observe that the narrowest angle used in our implementation indeed leads to best correlation (>0.9) for some sectors. (c) Post partitioning, we calculate the Pearson correlation value between the measured signal strength and the transmission range for those sectors that have at least some measurements present. We then select the sectors with a high correlation value (top 0.1 bin) for localization. (d) Given these candidate sectors, we use a *sector-specific* model with different α_j, ϵ_j for each sector j . This is because these chosen sectors can still have a different propagation trend in slope and intercept due to environmental variation. We omit this result for the sake of brevity. (e) We construct linear equations based on measurements from all the chosen sectors, and solve x_t, y_t and all α_j, ϵ_j to estimate a device’s location using an optimization function from a statistics package [15].

3.5 Adjacent-channel model for predicting TV broadcast leakage

V-Scope builds a model to predict the leakage power of a TV broadcast into its adjacent channels based on its in-band power⁴. It first identifies those locations with adjacent-channel measurements classified to be noise because the leakage is simply high-power noise. It then leverages a *linear* relationship between the measured power of this TV broadcast and that of its adjacent channel to construct the model. To demonstrate this power relationship, Figure 14 shows the power of a TV broadcast in channel 26 and that of its leakage in channel 27 at each chosen location. We observe

⁴Different TV broadcasts are usually allocated in far apart channels, and their adjacent channels therefore do not overlap.

that the leakage power increases approximately along a 45-degree line with TV power, and can be as high as -65dBm. The Pearson correlation between these datasets is 0.91. This prominent linear relationship comes from both a TV signal and its leakage being transmitted from a *same* location and at a power of *constant* difference. After traversing along similar paths to any location, their power still differ approximately by this amount leading to a linear trend between the in-band power and adjacent-channel power.

Since the specific leakage trend is different among TV transmitters, V-Scope constructs an individual model for each TV broadcast. The model takes the form $P_l = \gamma P_{tv} + \beta$, where P_{tv} and P_l are the power of a TV signal and its leakage at a given location, and γ, β are tunable parameters to capture the power relationship between these signals. V-Scope calculates γ, β by applying linear regression on the power measurements. To estimate the leakage of a TV broadcast at a given location, the databases can first use our region model to predict its in-band power P_{tv} , then leveraging the leakage model to predict P_l .

4. EVALUATION

In this section, we evaluate the performance of V-Scope based on three datasets collected over a 120 square-km area in and around a US city as described in § 2.1. We start by evaluating the overall performance of V-Scope in § 4.1. We then benchmark the performance of individual components in § 4.2. Overall, we find V-Scope can reduce false negative rates for protecting TV broadcasts by up to 59%, identify all the suitable whitespace channels at 72 – 83% locations under different channel quality constraints, and localize unlicensed devices at various locations with an error of 16 – 27m.

Implementation: We have implemented the measurement collection module and all the data processing modules in 7500 lines of Python, and built a database query utility in 650 lines of C++.

Methodology: We apply a standard five-fold cross-validation by using 80% randomly selected measurements to construct V-Scope models. The fitted models are used to predict the power of TV broadcasts, unlicensed devices and adjacent-channel leakage at remaining measured (testing) locations. We compare the predicted power of TV broadcasts with the -114dBm threshold to determine TV whitespace, while combining the predicted power of unlicensed devices and adjacent-channel leakage to estimate the noise power of whitespace channels. We compared our predictions with ground truth results based on measurements at these testing locations. To evaluate localization, we use measurements for our whitespace device and microphones operating at various locations, with ground truth locations determined by a GPS device.

Evaluation metrics: We use a variety of metrics to evaluate the performance of V-Scope such as false positive and false negative rates in predicting the availability of TV whitespaces, number of inappropriate whitespace channels selected under different channel quality constraints, absolute error in predicted signal strength, localization error between the actual location and predicted location of different TV-band devices, etc.

4.1 Overall performance of V-Scope

We use dataset A to evaluate overall performance of V-Scope in predicting whitespace channels and estimating their quality. We quantify the gain of V-Scope over a FCC-approved database [19] that predicts whitespace channels based solely on a propagation model without distinguishing their quality. We also evaluate the performance of V-Scope models fitted under different region sizes. This can provide database operators with insights about the tradeoff between prediction accuracy and storage overhead in choosing an appropriate region size for different spectrum management tasks.




Figure 15: Fraction of locations under-utilized by V-Scope and the database for predicting different TV broadcasts.




Figure 16: CDF of number of whitespace channels wrongly selected at different locations under various quality constraints.




Figure 17: Fraction of locations where all the suitable channels are correctly selected by different V-Scope models.

Approach	Database	VS-100m	VS-200m	VS-400m	VS-800m
False Positive Rate	0.29%	0.29%	0.22%	0.17%	0.17%

Table 4: False positive rates of the database and V-Scope models fitted at different region sizes in predicting TV broadcasts.

4.1.1 Predicting TV whitespace spectrum

We start by quantifying the false negative rates of the database and V-Scope models fitted at different region sizes for predicting TV broadcast channels that are accessible at least in part of our measured area. Figure 15 shows that different V-Scope models can reclaim the spectrum wastage of the databases by up to 59% locations. We then present the false positive rates averaged across different broadcast channels for these approaches in Table 4. Compared to the conservative database, we find the V-Scope models are able to offer same or even better protection to primary users. One important observation is that the V-Scope models fitted under different region sizes present little difference in the accuracy of whitespace determination, e.g., $\leq 2\%$ in false negative rate and $\leq 0.1\%$ in false positive rate. This is because most of their differences in predicted signal strengths are masked when comparing them with the detection threshold. We may therefore use a large region size (e.g., 800m) to fit V-Scope models for TV broadcasts, thereby reducing the overhead of storing model parameters.

4.1.2 Selecting suitable whitespace channels

Using the predicted whitespace channels, we evaluate V-Scope and the database in selecting appropriate channels under different power constraints. A whitespace channel is deemed to be suitable if its interference power is below a given constraint. Such a quality constraint can be estimated by network operators based on parameters such as the distance of wireless links, transmission power, and the minimum signal-to-noise ratio (SNR) for decoding a received signal (under desired modulations).

Figure 16 shows the CDF of the number of wrongly selected channels at different locations for a 100m V-Scope model and the database under different quality constraints. Note that a 5dB increase in two consecutive power constraints can lead to a 15 – 30Mbps drop in the achievable PHY rates for the 802.11n technology [21]. Without attempting to distinguish channel quality, the database can select all the appropriate channels at less than 2% of the locations, and have 3 – 4 channels wrongly selected for 50% of the locations. In contrast, V-Scope correctly selects all the qualified whitespace channels at 72% – 83% locations, and mis-predicts at most 1 channel for 92% – 97% locations. The much higher accu-

Signal Type	25th	Median	75th	95th	99th
TV broadcast	0.2dB	0.5dB	1.1dB	3.2dB	7.7dB
Unlicensed signal	0.2dB	0.6dB	1.3dB	3.4dB	6.9dB
TV leakage	0.3dB	0.7dB	1.4dB	3.6dB	6.9dB

Table 5: Absolute error in predicting the power of different types of signals by a 100m V-Scope model.




Figure 18: Accuracy in predicting the power of TV signals.

racy suggests that V-Scope can help avoid most of the performance penalty on a whitespace network due to channel mis-selection.

We break down the accuracy of V-Scope in predicting the power of different types of signals in Table 5. We note for all the signal types, the median error is below 0.7dB and a 95 quartile error below 3.6dB, thus explaining its high accuracy in channel selection.

We now study the impact of region sizes on the accuracy of V-Scope models for selecting suitable whitespace channels. Figure 17 shows the fraction of locations where all the appropriate channels are correctly identified by these models. We observe that the accuracy degrades at a larger region size. For example, a 800m region model correctly selects all the suitable channels at 8% – 11% fewer locations than a 100m region model, and a global model has 27% – 31% fewer such locations than a 100m model. This is because a model fitted for a larger region is not fine-tuned enough to capture local environment. The database operators may use this information to choose an appropriate region size for channel quality estimation based on *available storage* and *desired accuracy*.

4.2 Microbenchmarks of V-Scope

We now benchmark the performance of individual components in V-Scope, which are region-specific model in predicting TV broadcasts, sector-based localization algorithm, region-specific model in predicting unlicensed signals, and adjacent-channel models in predicting the leakage of TV broadcasts.




Figure 19: Accuracy in localizing a whitespace device (left) and microphone device (right).





Figure 20: Performance of different partitioning angles used by our approach for localizing the whitespace device.



(a) Whitespace transmitter.

(b) Unlicensed microphones.

Figure 21: Accuracy in predicting the power of unlicensed devices.




Figure 22: Accuracy in predicting the leakage power from TV broadcasts.

4.2.1 Accuracy in predicting TV broadcasts

We compare the V-Scope model with those models fitted by two alternative approaches. *Global* is a single model fitted for the entire measured area. *Local* is a region model fitted by linear regression instead of weighted regression. We quantify the gain of V-Scope models based on two datasets — dataset A collected during a 6-week period and a subset of it collected in an initial week.

Figure 18(a) shows the 99th quartile error of different models fitted based on the 1-week dataset. We first observe that *Global* has the highest prediction error since a global model can hardly be tailored to different local propagation environments. We then note a 19% – 40% reduction in prediction error achieved by *Local* over *Global* because *Local* tunes an individual model to each small region. V-Scope outperforms *Local* by 8% – 13% due to the use of weighted regression to compensate the non-uniform measurement density in model fitting.

As more measurements were subsequently collected, the measurement density tends to become uniform in each road segment. As a result, Figure 18(b) shows *Local* achieves a similar accuracy to V-Scope based on the 6-week data. This suggests that both versions of regression apply well on uniformly distributed measurements, but our weighted regression has its unique advantage in dealing with *non-uniform* measurements collected during a *short* wardriving period. Finally, we observe the accuracy in predicting the TV power improves at a smaller region size from both datasets.

4.2.2 Accuracy in localizing unlicensed devices

We use measurements collected for our whitespace transmitter and a microphone device in dataset B and C for evaluating our localization technique. The ground truth locations of these devices were obtained by a GPS device as mentioned before.

Overall accuracy: We compare following localization techniques. *Single-deter* and *Single-prob* are two popular RSSI modeling techniques used in EZ system [1] and WiFiNet system [16] respectively.

Single-deter uses a deterministic propagation model as described in § 3.4, whereas *Single-prob* uses a probability model based on the same propagation trend. *Sector-deter* and *Sector-prob* are our sector based versions of these common techniques. *Centroid* is the center location used by our technique to partition sectors, which is the geometric center of the strongest measurements (top 5dB).

Figure 19(left) shows the error of different algorithms for localizing our whitespace device. The measurements were selected based on different maximum power thresholds to emulate a variety of localization environment. We first observe that *Single-deter* achieves a low error of 26.9m using all the measurements (-45dBm threshold). Under different power thresholds, our sector based techniques improve *Single-deter* and *Single-prob* by 1.2 – 3 \times and 1.5 – 3.5 \times , because they carefully choose a few sectors and develop an individual propagation model for them. *Sector-deter* also refines the partitioning center *Centroid* by 1.2 – 4.1 \times . The error of *Sector-deter* and *Sector-prob* increases by up to 2.8 \times and 1.8 \times when using measurements at a lower power (e.g., -80dBm threshold). This is because the low-power measurements present a less distinct path loss trend and the partitioning center (*Centroid*) deviates more from the device’s actual location.

Figure 19(right) shows the accuracy of *Sector-deter* and *Single-deter* for localizing our microphone device in 5 different buildings based on clustered measurements as described in § 3.2. The maximum measured power ranges from -70dBm to -60dBm for different operating locations of the device. We observe that our technique *Sector-deter* achieves a low error between 16 – 27m, which are 2 – 2.8 \times lower than *Single-deter*.

Accuracy over different partitioning angles: Figure 20(a) shows the accuracy of *Sector-deter* with different partitioning angles for localizing the whitespace device. We observe that a larger angle leads to worse performance, e.g., the error increases by 1.5 \times from the 10-degree angle to the 30-degree angle, and by 2.06 \times to a 360-degree angle that is the baseline algorithm – *Single-deter*. To ex-

plain this performance degradation, we present the absolute correlation ratio averaged over the selected sectors for localization. Note that these sectors have a correlation ratio within the top 0.1 bin of all the sectors (§ 3.4). Figure 20(b) shows the absolute correlation decreases over an increasing angle, with a 360-degree sector having 0.36 lower correlation than a 10-degree. The highest correlation of 10-degree sectors indicates a best match of the selected measurements to the linear path loss trend, which in turn leads to highest accuracy. This justifies our use of the narrowest partitioning angle in the final algorithm.

4.2.3 Accuracy in predicting unlicensed signals

Using the device locations predicted by *Sector-deter*, we construct region-specific models for predicting the signal strength of the whitespace device and all (5) the microphone instances. Figure 21 shows the CDF of absolute errors in predicting the signal strength of different types of devices. We observe a 100m V-Scope model can achieve a median error of 1.3dB and 0.9dB, and a 75-quartile error of 2.8dB and 2.7dB for predicting the whitespace and microphone devices respectively. The prediction error increases at a larger region size, with a global model having about $3\times$ higher median error than a 100m region model for both types of devices.

4.2.4 Accuracy in predicting TV broadcast leakage

We compare our leakage model tuned to individual TV broadcasts with an alternative model (*All-Broadcast*) comprising a single set of parameters fitted for all the TV broadcasts. Figure 22 shows the CDF of absolute error for predicting the TV leakage power in all the adjacent channels based on dataset A. We observe that V-Scope achieves a 0.7dB median error and a 1.4dB 75-quartile error. *All-Broadcast*, albeit with slightly lower storage overhead, has a $1.6\times$ higher error in median and $1.4\times$ higher error at 75-quartile. The worse performance in *All-Broadcast* is because a single model fitted for all the TV broadcasts cannot accurately capture the specific leakage characteristics of each TV transmitter. Since the overhead of storing a separate set of model parameters for a few TV broadcasts is low, we choose this *broadcast-specific* model in our final design to improve accuracy.

5. DISCUSSION

Impact on FCC ruling: To harvest the benefits of V-Scope, we envision several changes to be made to the current FCC ruling. To ensure the validity of spectrum measurements, the measurement hardware needs to be rigorously tested to ensure the required detection accuracy, perhaps following the same testing procedure for the sensing-based whitespace devices. To utilize the opportunistic measurements, the database may choose an appropriate region size to fit the V-Scope models based on the desired accuracy. It may fall back to use its default model (R6602) when predicting those areas beyond the regions covered by measurements. Finally, while we find non-negligible temporal and spatial under-utilization of whitespace spectrum in microphone protection, our system has *inherent limitations* in reliably protecting licensed microphones with the use of vehicular measurements. To achieve this goal, perhaps some static sensors can be deployed in the proximity of microphone reserved locations for continuously monitoring their activity, or a proactive protocol can be adopted by microphones to alert interference from whitespace devices as suggested in prior work [12].

Coping with deployment cost: Our opportunistic wardriving approach has costs and overheads in deploying and managing whitespace sensors. Perhaps this measurement infrastructure is most useful in urban areas where seizing additional spectrum of good quality can be particularly beneficial to users. In such scenarios, spec-

trum database providers may contract with public vehicle operators to deploy this infrastructure, and recoup their costs by charging additional fees to use the services enabled by them. In this paper, we do not explore the economic aspects of opportunistic wardriving, but focus on the technical aspects in collecting and utilizing spectrum measurements to augment databases.

Addressing temporal variations and storage overhead: The collected measurements can become invalid with the change of environment and transmission behavior of (secondary) devices. While such temporal variation is observed to be small ($<5\text{dB}$) at most of our measured locations, it can increase with the future proliferation of whitespace devices. We envision to use time-based filtering, backed up by statistical algorithms such as Dixon’s test [4], to identify up-to-date measurements. To mitigate storage overhead, the databases may only store the fitted model parameters and device locations while discarding the measurements after using them.

Impact of measurement volume and coverage: Since our measurements are collected on a public transit bus opportunistically, the number of measurements in each region vary to a large extent. For example, different 100m regions can have measurements ranging from two to hundreds. We do not observe any clear trend between the number of measurements and the accuracy of the fitted model. Instead, the location of these measurements can play a more significant role. For example, using densely distributed measurements to fit a model can have lower accuracy than using a few measurements collected on sparse locations, which is the very motivation for our weighted regression model. Finally, the evaluation of our region model is not applicable to those locations beyond road segments. However, for those neighborhoods surrounded by the roads, we expect similar fading and shadowing effects introduced by the same environment, hence rendering our models to be relatively useful in predicting these areas.

6. RELATED WORK

Whitespace spectrum database: In its recent ruling, the FCC mandates commercial databases to use a well-known set of R6602 propagation curves [3] to estimate the coverage of TV broadcasts. We have found non-negligible spectrum wastage in these databases, possibly due to their inability to capture shadowing and fading in an urban environment.

Senseless [10] proposes an improved database design that is based on the Longley-Rice model incorporating terrain data and reported to incur low loss of whitespace spectrum. While V-Scope and Senseless share the same essence of using additional environmental information to calibrate a propagation model, their difference lies in the trade-off between the richness of the environmental data and the associated overhead in collecting such information. Senseless has the apparent advantage of obviating any measurement overhead. But the terrain information alone is somewhat insufficient to thoroughly reveal a propagation environment, e.g., the shadowing effect of buildings. In addition, Senseless is not designed to predict the quality of whitespace channels due to the lack of knowledge about the transmission power of secondary devices and the leakage characteristics of TV broadcasts. In contrast, V-Scope directly uses spectrum measurements in a local environment to tune a propagation model, which can therefore take account most of the environmental effects along with the transmission characteristics of TV-band devices. Hence V-Scope is likely to not only achieve better performance in predicting whitespace channels, but also estimate their quality. As the downside, V-Scope incurs the measurement and storage overhead and is only able to predict confined regions where measurements are available. Given these pros and cons, we believe the selection of either technique depends on the

desired functionality and the availability of training information, i.e., terrain data and spectrum measurements.

The design of V-Scope is also motivated by a recent indoor whitespace database (WISER [23]) that is purely based on measurements from sensors placed at strategic indoor locations. Targeting at a different scenario, V-Scope leverages opportunistic spectrum measurements to refine a propagation model for the vast outdoor area.

Propagation model enhancement: Some recent work [13, 14] aims to improve propagation models using a few measurements. Caleb et.al [13] proposes an adaptive path loss model for a 2.5GHz WiMax network based on linear regression. A geo-statistical approach has been subsequently proposed in [14], which interpolates systematically sampled measurements to build radio environment maps. Most of these approaches require measurement locations to be carefully chosen, which can hardly be guaranteed in vehicular sensing. V-Scope leverages weighted regression to remove fitting bias caused by the non-uniform distribution of measurements.

Localization: Existing RSSI modeling systems [1, 16] build a propagation model based on a few indoor measurements to pinpoint a device with an error of few meters. V-Scope adapts the RSSI modeling approach to outdoor scenario by carefully selecting measurements in certain radiation sectors and constructing a sector-specific model. Recent AoA based localization systems [7, 18, 22] are reported to achieve even higher accuracy within a meter, by measuring the arriving angle of a signal with an antenna array. Our current implementation leverages SISO based commercial spectrum analyzers, and thus cannot implement these techniques. We intend to explore their advantage in our future deployment when MIMO based spectrum analyzers become available.

Primary detection: Energy detection is the most straightforward algorithm for primary detection. However, it fails to detect a primary signal below a certain SNR threshold [8]. Subsequent work in feature detection [8, 11] proposes to leverage spectral features of a primary signals to improve detection sensitivity. V-Scope enhances the sensitivity of this technique with zoom-in capturing, while leveraging the detected features to estimate the strength of weak primary signals.

Finally, the concept of opportunistic wardriving on public vehicles was first reported in a position paper [24] that explored the accuracy of existing databases and the potential of a measurement-enhanced propagation model. In this paper, we significantly extend this prior concept by presenting a complete measurement framework that can enhance databases to accurately determine whitespace spectrum, predict its channel quality, and validate device location, backed up by extensive evaluation based on a large volume of measurement data.

7. CONCLUSION

We have presented V-Scope, a measurement framework for TV whitespaces. Our system leverages spectrum sensors mounted on public vehicles to collect spectrum measurements during the drive. These measurements are aggregated to refine various propagation models, which can augment spectrum databases to better identify whitespace spectrum, predict its channel quality, and validate the location of primary and secondary devices. We believe the concept of opportunistic wardriving and our techniques developed to utilize such measurements, can have broader applications in dynamic spectrum access beyond TV whitespaces.

Acknowledgments

We would like to thank the Madison Metro Transit for letting us collect measurements on the bus. We are also grateful to the anonymous reviewers whose feedback helped bring the paper to its final form. Tan Zhang and Suman Banerjee are supported in part by the US National Science Foundation through awards CNS-1040648, CNS-0916955, CNS-0855201, CNS-0747177, CNS-1064944, CNS-1059306, CNS-1345293, CNS-1343363, and CNS-1258290.

8. REFERENCES

- [1] K. Chintalapudi, A. Padmanabha Iyer, and V. N. Padmanabhan. Indoor localization without the pain. In *MobiCom*, 2010.
- [2] A. T. S. Committee. Atsc digital television standard - part 2: Rf transmission system characteristics, December 2011.
- [3] J. Damelin, W. Daniel, H. Fine, and G. Waldo. Development of VHF and UHF propagation curves for TV and FM broadcasting, 1966.
- [4] R. Dean and W. Dixon. Simplified statistics for small numbers of observations. *Analytical Chemistry*, 1951.
- [5] FCC. Unlicensed operation in the tv broadcast bands, second memorandum opinion and order, 2010.
- [6] J. C. Gower. A general coefficient of similarity and some of its properties. *Biometrics*, 1971.
- [7] K. Joshi, S. Hong, and S. Katti. Pinpoint: Localizing interfering radios. In *NSDI*, 2013.
- [8] H. Kim and K. G. Shin. In-band spectrum sensing in cognitive radio networks: energy detection or feature detection? In *MobiCom*, 2008.
- [9] M. Maechler, P. Rousseeuw, A. Struyf, M. Hubert, H. Kurt, M. Studer, and P. Roudier. cluster: Cluster analysis extended rousseeuw et al, 2014.
- [10] R. Murty, R. Chandra, T. Moscibroda, and P. V. Bahl. Senseless: A database-driven white spaces network. *IEEE Transactions on Mobile Computing*, 2012.
- [11] S. Narlanka, R. Chandra, P. Bahl, and J. Ferrell. A hardware platform for utilizing tv bands with a wi-fi radio. In *LANMAN*, 2007.
- [12] G. Nychis, R. Chandra, T. Moscibroda, I. Tashev, and P. Steenkiste. Reclaiming the white spaces: Spectrum efficient coexistence with primary users. In *CoNEXT*, 2011.
- [13] C. Phillips, D. Sicker, and D. Grunwald. Bounding the error of path loss models. In *DySPAN*, 2011.
- [14] C. Phillips, M. Ton, D. Sicker, and D. Grunwald. Practical radio environment mapping with geostatistics. In *DySpan*, 2012.
- [15] R Core Team. The R Stats Package. <http://stat.ethz.ch/R-manual/R-devel/library/stats/html/optim.html>.
- [16] S. Rayanchu, A. Patro, and S. Banerjee. Catching whales and minnows using wifinet: Deconstructing non-wifi interference using wifi hardware. In *NSDI*, 2012.
- [17] P. J. Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 1987.
- [18] S. Sen, J. Lee, K.-H. Kim, and P. Congdon. Avoiding multipath to revive inbuilding wifi localization. In *MobiSys*, 2013.
- [19] Spectrum Bridge Inc. Spectrum bridge database. <http://www.spectrumbridge.com/Home.aspx>.
- [20] ThinkRf. The thinkrf wireless signals intelligence platform. <http://www.thinkrf.com/>.
- [21] Ubiquiti Inc. Datasheet of ubiquiti sr71-a mimo cards. http://dl.ubnt.com/sr71a_datasheet.pdf.
- [22] J. Xiong and K. Jamieson. Arraytrack: A fine-grained indoor location system. In *NSDI*, 2013.
- [23] X. Ying, J. Zhang, L. Yan, G. Zhang, M. Chen, and R. Chandra. Exploring indoor white spaces in metropolises. In *MobiCom*, 2013.
- [24] T. Zhang and S. Banerjee. Inaccurate spectrum databases?: Public transit to its rescue! In *HotNets*, 2013.