Empirical Validation of Commodity Spectrum Monitoring

Ana Nika, Zhijing Li, Yanzi Zhu, Yibo Zhu, Ben Y. Zhao, Xia Zhou*, Haitao Zheng
University of California, Santa Barbara, USA, *Dartmouth College, USA
{anika, zhijing, yanzi, yibo, ravenben, htzheng}@cs.ucsb.edu, xia@cs.dartmouth.edu

ABSTRACT

We describe our efforts to empirically validate a distributed spectrum monitoring system built on commodity smartphones and embedded low-cost spectrum sensors. This system enables real-time spectrum sensing, identifies and locates active transmitters, and generates alarm events when detecting anomalous transmitters. To evaluate the feasibility of such a platform, we perform detailed experiments using a prototype hardware platform using smartphones and RTL dongles. We identify multiple sources of error in the sensing results and the end-user overhead (*i.e.* smartphone energy draw). We propose and implement a variety of techniques to identify and overcome errors and uncertainty in the data, and to reduce energy consumption. Our work demonstrates the basic viability of user-driven spectrum monitoring on commodity devices.

CCS Concepts

•Networks → Network design principles; Cognitive radios; Network monitoring; Mobile networks; Wireless access networks;

Keywords

Spectrum monitoring, Smartphones, Low-cost sensors

1. INTRODUCTION

Radio spectrum is a fixed and increasingly sought after resource. Licenses to existing spectrum bands are auctioned off by the FCC for billions to cellular carriers. To open up allocated spectrum for next generation wireless devices, the FCC is developing new tiered spectrum access models at multiple frequencies [6, 18]. Secondary devices can reuse "old" spectrum as long as they do not interfere with any primary (or legacy) users. Some secondary devices can become *protected entities* [18], who receive spectrum access free from interference by unlicensed secondary users.

The new model creates a contentious environment between different types of wireless devices, and adds further urgency to the development of spectrum monitoring tools to be used for transmitter detection, location and avoidance. Given the high cost in hardware and human resources for traditional spectrum monitoring, the FCC is partnering with industry to deploy online spectrum databases [1, 2] that maintain records for all primary transmitters, protected entities, and high power secondary transmitters. Secondary users can identify usable spectrum by querying the database by location and radio configuration. Spectrum databases are transparent, and easy to understand and utilize by secondary devices without paying for costly hardware to sense and detect primary users [24].

However, deploying spectrum databases does not address the difficult challenge of spectrum monitoring. Databases provide a simple way to catalog legacy wireless devices that are largely stationary, but do not simplify the task of sensing and locating new wireless devices that can be dynamic in both geographical and spectrum domains, *e.g.* portable access points for health and public safety entities, connectivity hubs for utility agencies, and interim cellular base stations to cover sudden traffic surges [6, 18]. As these devices continue to grow in number over time, the cost of spectrum databases will shift from the physical database to the cost of maintaining and updating entries to accurately reflect the frequency and physical location of active users. Exacerbating this challenge is the FCC's stringent location accuracy requirement of 50 meters [17].

A Case for Commodity Spectrum Monitoring. To date, spectrum monitoring is done by government agencies or cellular providers who perform measurements while driving around an area with specialized hardware such as spectrum analyzers. This method does not scale well for real-time, large-scale spectrum monitoring, given its costs in hardware and manpower [28]. As a result, measurement coverage is porous and sparse in many locations, making it simply impractical in less densely populated areas. Many transmitters would then evade detection, leading to large location errors in spectrum databases and undetected spectrum violations. One recent approach sought to address this problem by attaching spectrum analyzers to buses [47], but the system was severely constrained by bus routes and availability.

As flexible spectrum access policies grow in adoption around the world, it is clear that dedicated spectrum monitoring efforts will not achieve the scalability or coverage required. Instead, we believe such a system must include low-cost, commodity hardware, and leverage the growing population of active mobile devices, *e.g.* smartphones with embedded low-cost spectrum sensors. Such a distributed system, perhaps incentivized by network providers seeking to reduce spectrum monitoring costs, would have the key advantage of tying measurement density to user usage, where the system would generate the most dense measurement values and accurate sensing results in areas heavily frequented by users. Sensing results would be reported in real-time to a monitoring agency, which would process it to identify registered transmitters and locate usage anomalies.

Viability of Smartphone-based Spectrum Monitors. Deploy-

SenSys '16, November 14-16, 2016, Stanford, CA, USA
© 2016 ACM. ISBN 978-1-4503-4263-6/16/11. . . $15.00
DOI: http://dx.doi.org/10.1145/2994551.2994557

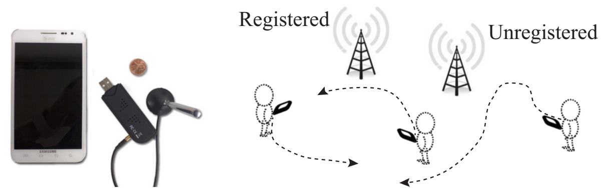

(a) Proof-of-concept platform (b) Collective spectrum measurements

Figure 1: Commodity spectrum monitoring using collective measurements on low-cost commodity devices.

ing a highly accurate commodity spectrum monitoring system is challenging. Low-cost spectrum sensors have limited sensitivity and bandwidth compared with specialized hardware [31]. User contributed data also suffers from unpredictable mobility and human error. These sources of noise and variance can be largely addressed by taking more samples. But doing so incurs additional costs in user participation (and possible incentive costs) and energy dissipation on user devices.

The goal of our work is to validate the viability of a future measurement platform based on smartphones and commodity sensors. We believe that as demand for wireless capacity continues to grow, next generation smartphones will come embedded with flexible radios that expose more low level RF data to the OS. To validate this approach to spectrum monitoring in the absence of prebuilt devices, we are using a prototype that combines commodity smartphones and an external RTL dongle interconnected via USB, monitoring a wide spectrum range of 52-2200MHz. Our platform, while impractical for today's cellular users, provides a lower-bound for analysis on the possible efficacy of spectrum monitoring using cognitive radio embedded smartphones. We hope this study and follow-ups can provide early validation for the viability of the smartphone-based spectrum monitoring platform. Finally, we also note that US government agencies have included ultra low-cost sensors like RTL dongles as potential candidates for spectrum monitoring [42].

We implement a "proof-of-concept" sensing platform by connecting $20 USB RTL dongles to today's smartphones (Figure 1), and collect 17 hours of outdoor spectrum measurements on TV bands from 48 volunteers. We identify in our monitoring data multiple sources of error, including RTL hardware noise, dynamic context, user mobility bias, and RF interference. Untreated, these artifacts lead to large errors (200 meters) in transmitter location estimation. To address this, we design multiple mechanisms to effectively remove artifacts of commodity measurements, producing accurate estimates of transmitter location (<40 meters error) from a few minutes of user measurements. Furthermore, we perform detailed energy analysis on our measurement platform, and identify ways to significantly lower energy dissipation on user devices at little or no cost on localization accuracy.

Most prior work on spectrum measurements focus on spectrum occupancy and not transmitter location, and rely on specialized, high-cost hardware [16, 24, 46, 47, 39]. We show that despite a variety of error sources and hardware limitations, low-cost commodity devices can be an effective approach to baseline spectrum monitoring. We believe our findings can generalize to systems using other low-cost sensor platforms (*e.g.* [33, 48]), which might experience different levels of hardware noise, but face the same challenges from user context, mobility bias, RF interference, and energy overhead.

2. BASIC DESIGN

We seek to study the basic viability of spectrum monitoring using collective user measurements with low-cost commodity devices. In this section, we set the context for our work by describing our "proof-of-concept" measurement hardware, and our basic monitoring system.

2.1 Measurement Hardware

While today's smartphones have multiple built-in radios, *e.g.* WiFi, Bluetooth, cellular, they only cover a very limited range of radio spectrum, *e.g.* 2.4GHz. To cover other frequencies, especially TV whitespaces (54-698MHz), our current platform leverages an external spectrum sensor.

Smartphone + $20 RTL. Our proof-of-concept platform consists of a commodity smartphone and an inexpensive Realtek dongle (RTL for brevity) [3] that connects to the smartphone via a USB cable. The RTL behaves as a spectrum sensor and collects raw spectrum usage signals; while the smartphone collects GPS data and acts as a "data processor", translating the raw data into a data stream that is more compact and meaningful for the monitoring system. We pick RTL because of its low cost (<$20), portability (<2oz weight), wide availability, and superior frequency coverage — it operates in 52–2200MHz with a sample rate up to 2.4MHz, and transfers raw I/Q samples to the connected host on the fly.

We built an Android app to run real-time spectrum measurements, by specifying frequency range, sampling rate and time duration. During a spectrum measurement, the smartphone obtains the I/Q samples from the RTL every 1ms and calculates the corresponding RSS value. To account for the impact of channel fading, the smartphone averages over all the RSS values gathered in a measurement cycle. For our basic design, we configure the measurement cycle to 1 second, compute 1000 RSS values (one per 1ms), and record the average. Thus the app generates a (time, GPS, RSS) tuple per second, amounting to 192KB of data per hour. In this case, the RTL is always on during a spectrum measurement, mapping to a 100% duty cycle.

Later we show that our design allows the RTL on time to be significantly reduced without affecting the monitoring result. For example, we can configure each cycle to be 10 seconds but within each cycle the RTL only performs measurements for the first 1 second, mapping to a 10% duty cycle. This generates 1000 RSS values which are then averaged to produce a (time, GPS, RSS) tuple once every 10 seconds.

Measurement Precision. Prior work [31] has shown that RTLs face two key disadvantages compared with conventional spectrum analyzers. *First*, RTLs have limited sensitivity and range, failing to detect weak signals. Our experiments on TV bands show that its sensing range is roughly 150m when detecting transmitters with 20dBm EIRP (100mW power) and 1500m for those with 1W power. *Second*, RTLs have a limited sensing bandwidth of 2.4MHz compared with USRP's 20MHz. To monitor a wideband, we can segment the target band into multiple 2.4MHz sections and allow RTLs to hop across the sections. Each hop faces a small frequency switching delay (up to 50ms [31]).

Energy Cost. Each RTL draws power from the smartphone via the USB connection. Detailed energy measurements [10] show that the total power draw depends on the specific tuner chip used. Between the two most popular RTL models, the Rafael Micro R820T dongle draws up to 1.2Watt while the FC0013 dongle draws about 0.6Watt [10]. These values are *on par* with the power draw of the LTE (1.5Watt) and WiFi (0.3Watt) radios on today's smartphones when operating in the receiving mode [23, 32]. For our study, we used the Rafael Micro R820T dongle model. Later in §6 we perform detailed analysis on RTL energy consumption, and discuss methods to minimize the per-user energy cost and the corresponding impact on the monitoring performance.

2.2 Spectrum Monitoring

Leveraging user measurements, we design our spectrum monitoring system to not only record the current utilization of each spectrum band, but also identify and locate active transmitters. Transmitter localization is a basic component of spectrum monitoring, and critical to the task of interference management and dynamic spectrum allocation.

Transmitter Identification. The first step to localizing a transmitter is to identify its signal. For TV whitespaces, the FCC requires that all (high power) devices transmit identifying information conforming to a standard, allowing observers to identify the device and its location [17]. Because the identification standard is not yet defined, in this paper we consider a simple and available solution: embedding a unique transmitter identifier (defined by the FCC) inside data transmissions as *cyclostationary features* [43].

Cyclostationary features are created when signals across some sequence of radio frequency segments are repeated, generating an easy to detect energy peak in the spectral correlation function (SCF) map. Prior work [43] developed a simple technique to achieve fine-grained control over positions of these signal peaks, encoding unique transmitter identifiers as cyclostationary features. The result is visible to any monitoring device that can sense signals on the transmitter's frequency, without decoding data.

In our system, the measurement devices can detect each feature by first capturing the RF signal on the transmitter's frequency and applying FFT to compute a normalized, discretized version of the SCF map, and then locating the feature peak using a correlation-based detection method [41]. This eliminates noise and random occurrences of cyclostationary property in the packet data itself. Later in §4.6, we show that our devices can also effectively identify registered wideband transmitters whose frequency bandwidth exceeds RTL's sensing bandwidth (2.4MHz).

Transmitter Localization Our basic design uses collective measurements from mobile users. While walking, a user uses her smartphone and RTL to collect spectrum measurements in the local area and submits the results in real-time to a monitoring agency, *e.g.* in snapshots of a few minutes each. The agency then analyzes these measurement snapshots to produce a complete view of the spectrum usage in a wide area, *e.g.*, the physical and frequency location of each detected transmitter. As users move (and start or stop their measurements), the system obtains a dynamic view of spectrum readings that scales with the number, density, and physical reach of users in the network. Our design does not require any specialized movement patterns for users.

Locating Registered Transmitters. When a registered transmitter embeds a valid ID (as cyclostationary features) in the monitored frequency, our devices can identify the feature location and extract its RSS traces. To locate a detected transmitter, we can apply a RSS-based transmitter localization algorithm on the collected data.

Locating Unregistered Transmitters. In the absence of any registration ID, the estimated transmitter location can be noisy, because the RSS can come from one or multiple transmitters. Assuming only a single transmitter is present, we can estimate its location based on RSS measurements. We leave the task of isolating and locating individual unregistered transmitters to future work.

2.3 Incentivizing User Participation

Our design assumes the agency can recruit mobile users at targeted monitoring areas. There are multiple forms of recruitment, including crowdsourcing and incentivizing in-network users [31]. Here a practical challenge is how to ensure adequate coverage. One potential solution is to leverage an ecosystem of network providers, where each provider leverages its own users (and their commodity mobile devices) to perform spectrum measurements. These service providers are active spectrum users who seek reliable spectrum usage to support/augment their services, and thus are incentivized to participate in spectrum monitoring and protect their own usage.

Spectrum measurements will come from two distinctive groups of users. First, passive measurements will be collected from each provider's own user population, by energy-efficient background app running on mobile devices. To incentivize participation, a network provider can reward participating users with small credits to network charges commensurate with actual measurements performed. Recent studies [14, 22] have shown that small monetary incentives will increase user participation in crowdsourcing tasks.

Second, the system can request on-demand measurements from users of other networks to augment passive data. Here a local network entity will predict the coverage from in-network users and trigger crowdsourcing requests from other networks' users in the target region. Providers pay non-network users for measurement tasks. All users, regardless of provider, run a crowdsourcing daemon and listens for locally broadcast measurement requests.

3. QUALITY OF USER-CONTRIBUTED MEASUREMENTS

For spectrum measurements contributed by end-users and low-cost commodity hardware, there exists obvious doubt on data quality and the impact on spectrum monitoring. In this work, we take a data-driven approach to study this concern. In the following, we first describe our efforts on collecting real world user measurements, then present our analysis on the quality of these measurements, and their impact on the accuracy of transmitter localization.

3.1 Real World Measurements

We recruited 48 volunteers via email announcements in our local area. They are between the ages of 20 and 40 and have different body shape and height. Each user was given a Galaxy SIII smartphone with our measurement app installed and an attached RTL device. In each experiment, the users walked (as they normally do) in a large neighborhood of the target transmitter, at least 200m×200m in size. No further instructions were given and the users had no knowledge of the transmitter location. For our measurements, we configure the RTLs to operate in 100% duty cycle. Participants used phones provided by us and walked along areas we specified, thus no personal information was leaked.

There is no active TV whitespace transmitter in our area with ground truth location information. Thus we set up our own transmitter using a USRP N210 radio, emitting OFDM signals on a 2.4MHz band or a 6MHz band (for wideband experiments). We place the transmitter roughly 4m above the ground in each experiment. We consider two available TV whitespace bands (569MHz and 653MHz)[1]. The majority of our experiments were on 569MHz. We configure our transmitter to emit at 100mW (20dBm), thus the signal detection range of a RTL is roughly 150m.

Measurement Environments. We performed extensive measurements at four outdoor environments, representing scenarios where

[1] We select the TV whitespace bands by querying two different whitespace databases, Spectrum Bridge [1], and Google Spectrum Database [2] to identify the available bands in our local area.

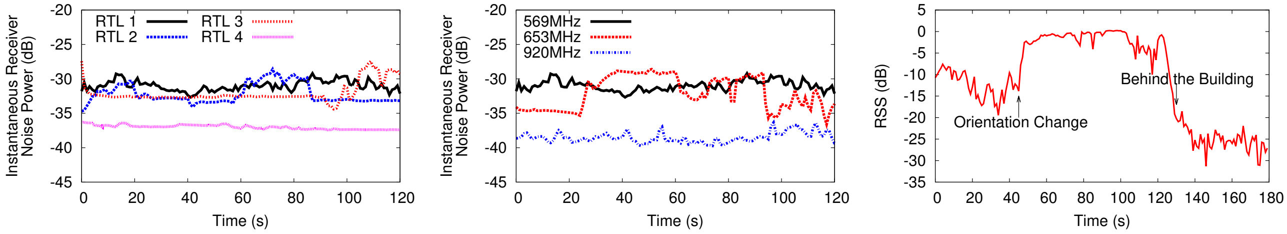

(a) Noise Power over Time (569MHz)

(b) Noise Power over Frequency

(c) Impact of Dynamic Context

Figure 2: Artifacts of commodity spectrum measurements. (a)-(b) RTL receiver noise varies over time and frequency and is device-dependent. (c) Dynamic user and environment context leads to sudden changes in RSS measurements.

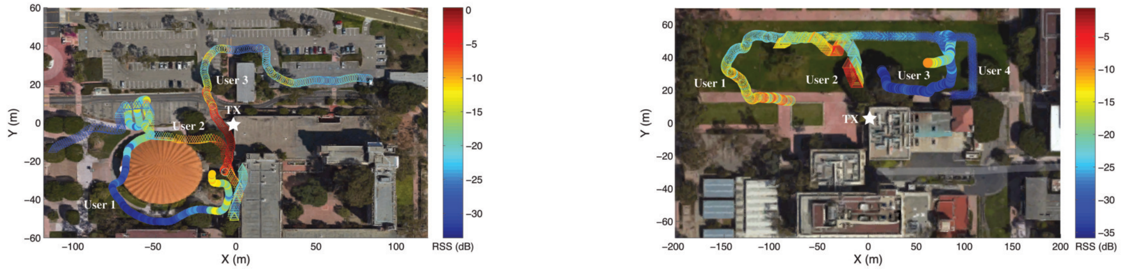

(a) Sample Scenario A

(b) Sample Scenario B

Figure 3: Sample user routes and RSS measurement results. The reported noise floor is in between -30 and -35 dB.

users perform spectrum measurements during daily (walking) activities. Each experiment round involved at most 6 randomly selected volunteers.

- *Open area* – An open lawn area with minimum obstacles beside our users. Thus, RSS measurements are mostly in line of sight to the transmitter, except that the user body can block the signal.
- *Areas around buildings* – A complex area where users walk in alleys between buildings and parking lots. Both static (buildings, parked cars, trees) and mobile obstacles (pedestrians, moving cars, and bikes) were present during the measurements.
- *Downtown sidewalk* – A complex area where users walk on sidewalks between outdoor shopping stores with various obstacles (*e.g.* buildings, trees, and pedestrians).
- *Outdoor plaza* – An outdoor food court area plus a pathway to the outdoor parking lot, where users walk around during busy lunch hours.

RSS Dataset. Our measurements took place between January and March 2015 and generated a dataset of more than 17 hours of measurement data or 360,000+ (time, GPS, RSS) tuples. Among them, 1 hour of measurement data was collected by setting up two transmitters, one as a registered user who embeds its registration ID inside the transmission and another as an unregistered (and interfering) transmitter.

3.2 Data Quality Analysis

Our analysis on the RSS dataset identified a large amount of noise and inconsistency across measurements. We also identified four key sources for these artifacts. The first three factors are responsible for creating inconsistent and noisy RSS values, while the last factor leads to uncontrolled and biased spatial coverage.

1. RTL Receiver Noise. The low-cost RTL devices are not calibrated and have a high receiver noise figure (also reported by other studies [4]). We found that the noise level is device-dependent, and varies over time and frequency, making it very hard to model and predict. To illustrate this, Figure 2(a)(b) plots the measured instantaneous noise power over time for four randomly-chosen RTL devices at 569MHz, and for one RTL device at three different frequencies (569MHz, 653MHz, 920MHz).

2. Dynamic Context. Unlike war-driving with a vehicle-mounted antenna, our measurements are carried out by walking users holding smartphones and RTLs. Changes in user movement pattern, body posture and local environment translate into random fluctuations in the data. Figures 2(c) shows an example where a user's reported RSS increases abruptly when her body orientation changes (from blocking the transmitter to unblocking), and later drops significantly when she walks behind a building.

3. RF Interference. When an interfering transmission is present, the measured RSS captures the aggregate power of the original signal and the interfering signal, and thus is very noisy. This is particularly harmful when unregistered transmitters "hide" behind a registered transmitter, *i.e.* using spectrum without registering as an authorized user. The corresponding RSS data will produce a wrong location of the registered transmitter.

4. Coverage Bias. Since we do not control user mobility patterns, user routes and measurement locations are uncontrolled and unpredictable. For example, Figure 3 plots the 3-minute routes taken by three RTLs for two scenarios, and the measured RSS values. For both routes, coverage around the transmitter is unbalanced. Such uncontrolled user routes create sampling bias, which is highly undesirable for transmitter localization.

3.3 Impact on Spectrum Monitoring

Together, these factors generate a considerable amount of noise and artifacts in RSS measurements. They are difficult to model and calibrate, leading to new challenges not found in conventional

	W.Centroid		W.Centroid P.		Gradient		Ecolocation	
	max	mean	max	mean	max	mean	max	mean
100% DC w/o int	82.6	43.1	81.7	35.4	197.1	47.8	117.3	41.2
100% DC w/ int	96.3	68.1	101.7	57.4	193.3	113.8	102.1	65.7
10% DC w/o Int	121.7	41.3	103.1	35.7	173.2	48.3	111.7	42.5
10% DC w/ int	136.2	71.2	113.6	67.5	183.4	143.8	121.9	70.1

Table 1: Localization error (in meters) obtained by applying conventional solutions on our RSS dataset. 100% DC references to 100% duty cycle, and w/o int refers to in absence of external interference.

spectrum monitoring, *i.e.* via war-driving with high-end spectrum analyzers [46, 47]. To quantify their impact on transmitter localization, we applied five popular transmitter localization methods to our data directly. We organize the dataset into 960 snapshots of 5 minutes each, and perform localization on each instance. We pick 5 minutes because it is roughly the time a user takes to walk 300 meters, *i.e.* twice of the RTL sensing range. Thus, it represents the duration that a RTL can capture the transmitter's signal.

We consider five popular localization algorithms: centroid, weighted centroid [13], weighted centroid with Gaussian prediction [30], gradient [20], and ecolocation [44]. We also examined other well-known solutions like trilateration [36] and calibrated propagation model, and found they perform much worse. Furthermore, we study the impact of RTL duty cycle. Since our measurements were taken by RTL being always on, *i.e.* 100% duty cycle, we emulate 10% duty cycle by subsampling the dataset by a factor of 10.

Table 1 lists the maximum and mean localization errors under three different conditions. We omit the Centroid result since it is worse than weighted Centroid. Overall, the maximum localization error can easily reach 100–200 meters, which is too coarse for common monitoring tasks, and clearly cannot meet the FCC requirement of accuracy within 50 meters [17]. Furthermore, the accuracy varies significantly across measurement instances, again confirming the large uncertainty on the data quality. Finally, we see that the localization error degrades largely under 10% duty cycle.

4. DEALING WITH NOISY DATA

Clearly the existing localization algorithms are unable to handle the noisy RTL measurements. To overcome this problem, we propose a robust spectrum monitoring system that combines *de-noising*, *interference removal* with *fidelity prediction*. These components allow us to remove the key noise and interference components from the RTL measurements, apply an existing transmitter localization on the cleaned data, and predict the accuracy of the localization result. As a result, our proposed solution provides three key benefits:

- **Improving localization accuracy** – By suppressing the noise and interference contribution, our solution effectively reduces the localization error.
- **Overcoming uncertainty** – By predicting the fidelity of each localization result, we enable effective decision-making in spectrum monitoring. The system can act based on fidelity levels to obtain quality results; precautionary measures may include skipping a particular measurement snapshot, or sending police devices with sophisticated hardware to do close-range verification.
- **Reducing RTL duty cycle** – By aggregating measurements across space, our solution reduces the amount of data required for accurate localization. This translates into significant reduction of RTL duty cycle, *e.g.* from 100% to 10%, with little impact on localization.

In the following, we describe the three components in detail. But to provide context, we start from briefly describing ecolocation [44], an existing transmitter localization algorithm used in our design, followed by a quick summary of our key contributions beyond ecolocation.

4.1 Background: Ecolocation

Ecolocation [44] is a widely-known algorithm for transmitter localization. The high-level idea is to capture the abstract relationship between RSS and link distance: the longer the link distance, the lower the RSS value. Unlike trilateration that represents the relationship via a propagation model, it applies a probabilistic approach to count, for each candidate transmitter location, how often the RSS-distance relationship is satisfied. The candidate location with the highest satisfaction rate is the final transmitter location.

Specifically, given a candidate transmitter location l, the algorithm calculates the distance between l and each measurement location i, referred to as $D_{l,i}$. Consider a pair of measurement locations i and j $(i \neq j)$. The pair satisfies the RSS-distance relationship if one of the three conditions is met: $(RSS_i > RSS_j)$ & $(D_{l,i} < D_{l,j})$, $(RSS_i < RSS_j)$ & $(D_{l,i} > D_{l,j})$, as well as $(RSS_i = RSS_j)$ & $(D_{l,i} = D_{l,j})$. The satisfaction rate $F(l)$ of the transmitter location l is the ratio between the number of measurement location pairs that meet one of the three conditions and the total number of distinct pairs. And the final transmitter location is $TX = \text{argmax}_l F(l)$.

As we will show below, our design leverages ecolocation as the underlying transmitter localization algorithm. We pick ecolocation over other candidates because it works well with small amounts of measurements and offers certain degree of robustness against noise.

4.2 Overview of Our Contributions

We make four new contributions beyond ecolocation.

- **Denoising & Localization** (§4.3) – To reduce the impact of noise, we partition the RTL data into context-based segments, apply ecolocation in each segment and aggregate the satisfaction rate across segments. Since the noise profile is much more consistent within each segment, this leads to a much more reliable estimate of the satisfaction rate, thus a more accurate localization result.
- **Predicting localization fidelity** (§4.4) – By comparing the measured satisfaction rate to the ideal value, we predict the fidelity of the localization result, and use it to aggregate localization results over time. This effectively reduces the uncertainty of the monitoring result.
- **Removing external interference** (§4.5) – By detecting and extracting the cyclostationary features of each registered transmitter, we can separate the RSS contribution of the registered transmitter and the interference, thus localizing them individually following the above two steps.
- **Wideband monitoring** (§4.6) – By scanning through multiple frequency ranges and combining probabilistic localization metrics across frequency (weighted by fidelity), we achieve accurate identification and location of wideband transmitters.

4.3 De-noising & Localization

After collecting a measurement snapshot (*i.e.*, x minutes of RSS measurements), we first partition the data into multiple segments (to isolate the noise), apply ecolocation on each segment S to obtain a per-segment satisfaction map $\{F_S(l)\}$, and then aggregate maps of multiple segments (and RTLs) into one ultimate satisfaction map $F(l) = \frac{1}{|S|} \sum_S F_S(l)$. We then determine the transmitter's location using $\{F(l)\}$.

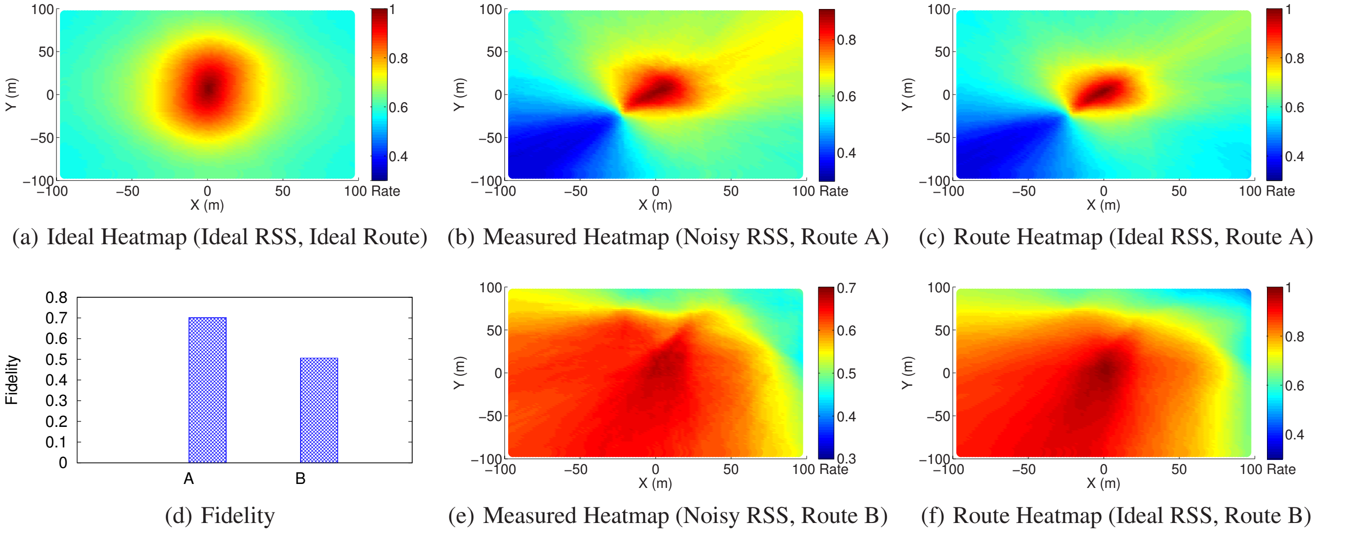

(a) Ideal Heatmap (Ideal RSS, Ideal Route) (b) Measured Heatmap (Noisy RSS, Route A) (c) Route Heatmap (Ideal RSS, Route A)

(d) Fidelity (e) Measured Heatmap (Noisy RSS, Route B) (f) Route Heatmap (Ideal RSS, Route B)

Figure 4: The satisfaction heatmaps used for transmitter localization and confidence computation.

10% duty cycle (DC)		100% DC w/ local avg		100% DC w/o avg	
max	mean	max	mean	max	mean
52.6	18.4	44.7	17.3	71.3	21.4

Table 2: Localization error (in meters) for 10% and 100% RTL duty cycle (DC).

Context based Data Segmentation. As a user walks around, the RSS value should vary smoothly over time unless the user moves behind a large obstacle or suddenly changes her body orientation (*e.g., from facing the transmitter to facing backwards and blocking the signal*). When these happen, the RSS value will experience a sudden change. With these issues in mind, we segment the data whenever two consecutive RSS observations differ by more than 8dB. We choose this threshold because our benchmark measurements show that human body blockage introduces at least 8dB loss in RSS (at least for the TV band). This value could be further optimized, which we leave to future work. Finally, we apply ecolocation on each segment S, producing a corresponding satisfaction map $F_S(l)$ per segment.

Reducing RTL duty cycle. By effectively combining data across space, we can reduce the amount of data required for accurate localization. Specifically, we find that building a good estimate of the satisfaction rate map $\{F_S(l)\}$ does not require RTL measurements at 100% duty cycle. This is because when computing the RSS and distance relationship for each location pair (i and j), i and j need to be sufficiently separated to minimize noise impact (as discussed in §4.1). Thus having fine-grained RSS data over time is only helpful when we take local average to reduce noise impact, *e.g.* averaging 10 seconds of measurements into 1 RSS value. But since each original RSS value is already an average over one second, the additional temporal average over a longer time window has limited benefits. For example, Table 2 lists the localization performance for 10% and 100% duty cycle values where 10% duty cycle refers to measuring every 1 second out of 10 seconds, while 100% refers to measuring in each of the 10 seconds. We see that 10% duty cycle performs similarly to 100% duty cycle with local averaging. As we will show in §6, such large reduction in the RTL duty cycle translates into significant reduction in energy consumption.

4.4 Predicting Localization Fidelity

After getting a localization result (from the above step), we wish to predict the fidelity (or the level of accuracy) of the result. For this we leverage the spatial distribution of the satisfaction rate $\{F(l)\}$.

Consider an ideal scenario where the RTL measurements are free of any noise, interference and dynamic context, *i.e.* the RSS value follows an ideal propagation model, and the measurements are evenly distributed around the target transmitter, *e.g.* a dense grid. The resulting satisfaction rate over space, hereby referred to the satisfaction heatmap, is shown in Figure 4(a). Here the target transmitter is located in the center (0,0), and $F(0,0) = 1$, and we assume a log-normal propagation model (for outdoor scenarios).

For comparison, we also plot in Figure 4(b) and (e) two (measured) satisfaction heatmaps built from two real RTL measurement sets on the same transmitter. For each, we center the heatmap at the estimated location. Clearly, we can observe a distinct difference between the two measured heatmaps, and their difference from the ideal heatmap (which is the same for both RTL measurements). These differences are caused by both the noise in the RTL measurements, but also the difference in user route (or coverage bias). To further illustrate these, we also plot in Figure 4(c) & (f) two *route* heatmaps, produced from using the actual RTL measurement locations but replacing each measured RSS with a model-generated value. By comparing these heatmaps, we can see that the difference between the ideal and route heatmaps is mostly on the heatmap structure, capturing the impact of user route (coverage bias). The difference between the route and measured heatmaps is the satisfaction value, reflecting the impact of measurement noise.

Motivated by these observations, we compute the location fidelity as the normalized cross-correlation between the ideal and measured heatmaps: $\lambda = \frac{1}{n}\Sigma_{x,y}\frac{(f(x,y)-\bar{f})(g(x,y)-\bar{g})}{\sigma_f\sigma_g}$, where $f(x,y)$ and $g(x,y)$ are the values of point (x,y) in the ideal and measured

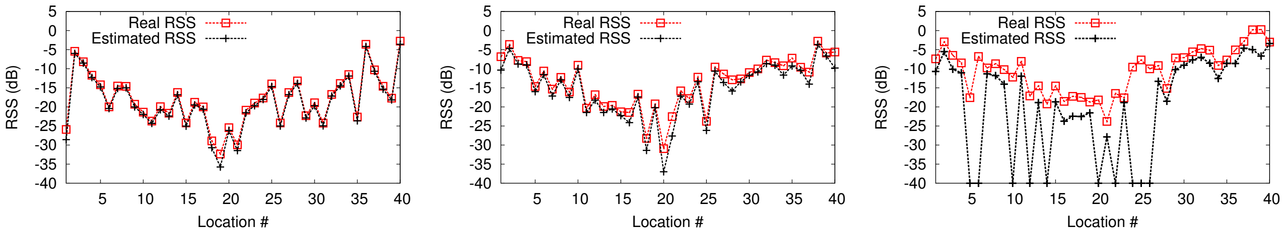

(a) No Interference (b) Weak Interference (c) Strong Interference

Figure 5: Raw received and estimated signal strength without and with interference.

heatmaps, n is the heatmap size, $\bar{f}$ is the average of f and σ_f is standard deviation of f. Figure 4(d) plots λ for both snapshots, 0.7 for snapshot A (13.3m location error) and 0.52 for snapshot B (51.7m location error).

Fidelity-Guided Temporal Combining. The above discussion shows that the monitoring performance depends on the coverage of user routes which varies over time. Thus we propose a temporal combining mechanism to further improve localization. For example, we can partition a 5-minute monitoring snapshot into 3 overlapping 3-minute slots, apply the above described method to estimate the transmitter location and its fidelity value, (l_i, λ_i), $i = 1, 2, 3$. We estimate the location as $\frac{l_1\lambda_1+l_2\lambda_2+l_3\lambda_3}{\lambda_1+\lambda_2+\lambda_3}$ and the new fidelity as $\frac{\lambda_1^2+\lambda_2^2+\lambda_3^2}{\lambda_1+\lambda_2+\lambda_3}$.

4.5 Removing Interference

Under interference, our RTL devices would observe a single transmission formed by the union of the registered and unregistered transmissions. The resulting RSS captures the sum of signal and interference power. Localization based on such data is obviously unreliable.

Feature based Interference Isolation. At each measurement location, if the amount of interference is moderate, our RTL devices can observe a valid cyclostationary feature (related to a registration ID). From the cyclostationary feature's peak strength, we can estimate the RSS of the registered transmission that actually carries the registration ID. This allows us to separate registered transmissions from those unregistered ones, *i.e.* the interference, thus locating the registered transmitter reliably in the presence of interference.

Specifically, the strength of the cyclostationary feature that carries the valid registration ID, is proportional to $\frac{SINR}{1+SINR}$ [43]. At each measurement location, we estimate the RSS of the registered transmitter S^* from its detected feature strength s and the raw RSS S_0: $S^* = \frac{s}{\rho} \cdot S_0$, where ρ is the maximum detectable feature strength that is hardware dependent; ρ=0.99 for the RTL radios in our experiments. Our approach can detect features using very few raw I/Q samples, thus reducing RTL duty cycle has no impact here.

As an example, Figure 5 shows the measurement results at 40 locations, comparing S_0 (raw RSS) and S^* (feature estimated RSS) under three scenarios: no interference, weak interference, and strong interference. The interfering transmitter is placed 90m away from the registered transmitter, whose power level is either the same as (weak interference) or 30dB higher than the target (strong interference).

In the absence of interference, S^* is almost identical to S_0. As the interference strength elevates, S_0 grows higher than S^*, and the difference increases gracefully with the interference strength. At locations where interference is strong, we are unable to extract features and mark S^* as the noise level -40dB. Overall, as long as we can detect the feature, S^* is a reasonable replacement of S_0, the raw RSS value, which we use to locate the registered transmitter. Our results in §5 also confirm that even under strong interference, our method can still locate the registered transmitter reliably.

Finally, we can estimate the total interference RSS at each location as $S_0 - S^*$ and use it to approximate the interferer location assuming only one interferer is present.

4.6 Wideband Monitoring

We now extend our design to wideband monitoring. Consider the task of monitoring a predefined frequency band, *e.g.* TV whitespace channels (6MHz each). We can split the band into 3 sections of 2MHz each; let each RTL hop across the sections sequentially and aggregate the results. This is feasible and efficient since RTLs have a small frequency switching delay (<50ms). An alternative is to divide the sections among users, but this requires higher user density. For simplicity, we focus on the frequency hopping method.

For robustness, wideband transmitters embed cyclostationary features over the entire band (6MHz). Thus to identify registered wideband transmitters, RTLs need to detect these wideband features by "stitching" multiple adjacent frequency observations together. Specifically, after monitoring each 2MHz section and build the corresponding SCF map, each RTL concatenates these maps in frequency to build a wideband SCF map for wideband feature detection. This requires the transmitter to transmit the same wideband feature for at least a time period long enough to complete a single scan. The stitching happens at each individual RTL and thus does not require tight user synchronization. We have validated this design using real measurements.

Fidelity Guided Frequency Combining. To identify registered wideband transmitters, our monitoring system needs to capture raw signals (and cyclostationary features) from multiple frequency sections. But to localize a detected transmitter, do we need to use data from all the frequency sections or just one? Ideally, one section should be enough. In practice, frequency selective fading or non-uniform interference profiles lead to performance fluctuations across frequency section and time. Thus we propose to combine localization results (in terms of the satisfaction heatmap as discussed in §4.3) across these frequency sections, weighted by their feature strength. This aggregation introduces frequency diversity into heatmap construction, further improving its reliability. Later our results show that it significantly improves localization accuracy.

4.7 Computation Complexity

For our localization design, the bulk of the computation lies in the satisfaction heatmap computation. Ideally we need to compute a fine-grained heatmap at locations near the target transmitter (which is unknown). For better efficiency, we first center the search area at the location reporting the highest RSS value among all the

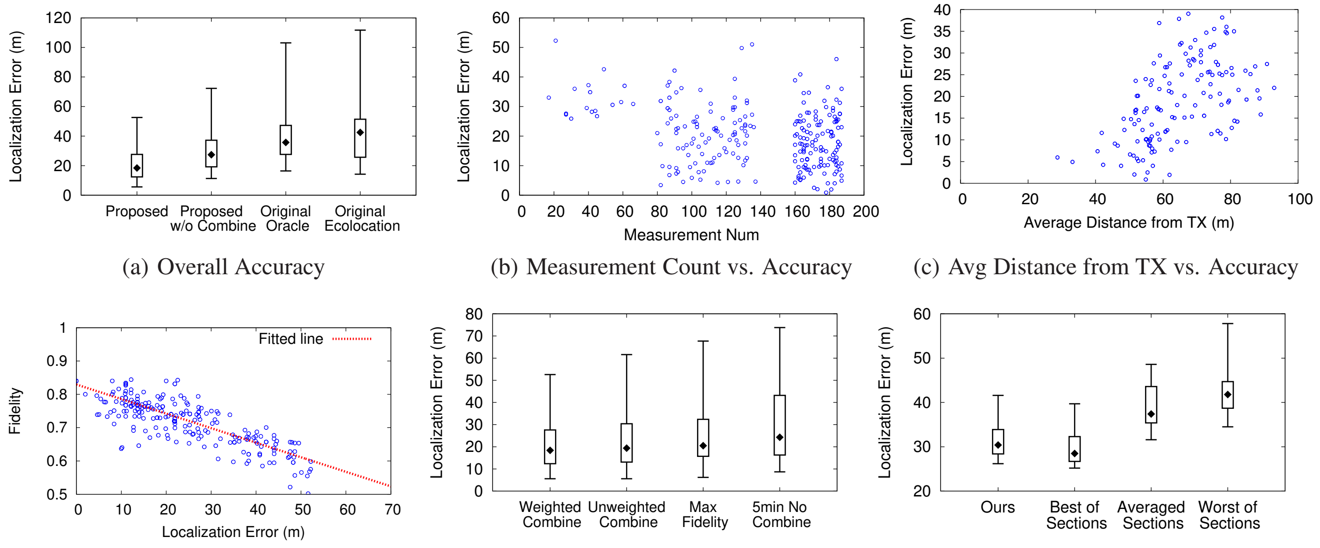

(a) Overall Accuracy

(b) Measurement Count vs. Accuracy

(c) Avg Distance from TX vs. Accuracy

(d) Fidelity Prediction

(e) Temporal Combining

(f) Frequency Combining

Figure 6: The localization accuracy of our proposed algorithm. (a) Quantiles (min, 25%, 50%, 75%, max) of the localization error across 960 snapshots. (b)-(c) Impact of measurement count and average distance from the transmitter on localization accuracy. (d) Our fidelity metric offers a good prediction of the accuracy level. (e) The effectiveness of our fidelity guided temporal combining. (f) The effectiveness of our fidelity guided frequency combining.

RTLs involved in the current snapshot. We then apply a multi-tiered search algorithm, first using a coarse granularity (1 per 10m) to identify regions of high satisfaction rates, followed by a fine-grained sampling (1 per 2m) at these regions. With this process, we limit the processing time per 5-min snapshot to 1 minute, using an non-optimized MATLAB code running on a standard MacBook Pro (2.2GHz CPU, 8GRAM). This can be further reduced using a good native C++ implementation running on a faster machine.

5. EVALUATION: LOCALIZATION ACCURACY

In this section, we evaluate the proposed spectrum monitoring system, focusing on the localization accuracy. We use our RTL RSS dataset described in §3, and 17 hours of measurements (16 hours without interference, and 1 hour with interference). We organize the dataset into 960 snapshots of 5 minutes each, performing localization on each snapshot. By default, we assume RTLs operate at 10% duty cycle, *i.e.* scan for 1s and then stay idle for 9s, which we emulate by subsampling our data by a factor of 10.

5.1 Accuracy in Absence of Interference

We start from the narrow band (2.4MHz) scenarios in absense of external interference. Figure 6(a) plots the quantile distribution of the localization error (min, 25%-tile, median, 75%-tile, max). Here we compare our proposed solution, our solution without temporal combining, the original ecolocation, and the best of the conventional localization methods, *i.e.* weighted centroid with Gaussian Prediction (as shown by Table 1). Compared to the two conventional localization methods, our proposed solution significantly reduces the localization error. The maximum error is bounded by 53m while the other two reach 112m and 103m (for 10% RTL duty cycle). When we increase RTL duty cycle to 100%, ours reduces to 44.8m while the best conventional method provides 82m.

Figure 6(a) also illustrates the breakdown of performance improvement by two components: denoising via segmentation and fidelity guided temporal combining. The difference between the original ecolocation and our proposed solution without temporal combining demonstrates the effectiveness of segmentation. The difference between our proposed solution w/ and w/o temporal combining shows the contribution of temporal combining. We can see that both components contribute to the accuracy improvement.

Performance Variance and Fidelity Prediction. We are also interested in understanding why the localization accuracy varies considerably across snapshots. For our solution, it varies between 5m to 53m, by a factor of 10. First we look at the number of measurements in the snapshot. Figure 6(b) shows that a snapshot with a smaller number of measurements (mostly because the number of RTLs is small) is likely to produce less accurate result, but the overall correlation is weak. A deeper analysis on the traces shows that the average distance to the transmitter is a more important factor (Figure 6(c)). As the user gets further from the transmitter, the impact of noise and sampling bias elevates, which degrades the localization performance.

We handle such variance by predicting the result fidelity. Figure 6(d) shows the predicted fidelity as a function of the localization error. We observe a good pattern between the two – higher confidence values (> 0.7) indicate more accurate localization ($< 30m$).

Effectiveness of Fidelity Guided Combining. We first consider temporal combining for narrowband monitoring. Figure 6(e) plots the quantile distribution of localization errors of the following four configurations for each 5-minute snapshot: fidelity (our proposed solution), (2) averaging the results of the 3 snapshots, (3) dividing the data into three 3-minute snapshots and selecting the localization result with the highest fidelity, and (4) no temporal combining. We see that weighted combining performs the best and significantly reduces the error tail. Compared with no combining, it reduces the maximum localization error from 75m to 52m. This result demonstrates the effectiveness of the fidelity guided temporal combining.

Next we study the proposed fidelity guided frequency combining used in wideband monitoring. For this we consider the scenario

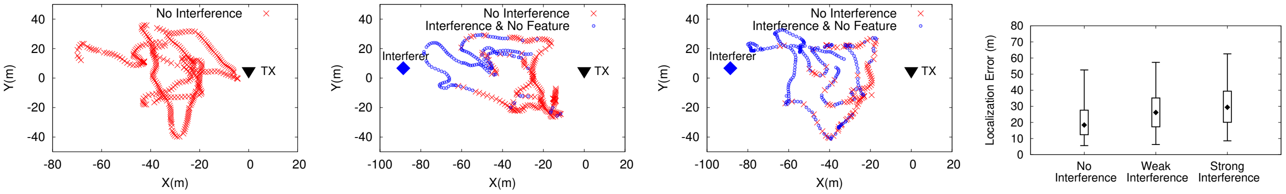

(a) No Interference (b) Weak Interference (c) Strong Interference (d) Localization Result

Figure 7: Locating transmitters when both registered and unregistered (interfering) users are present. (a)-(c) The RTL interference measurement results under no interference, weak and strong interference. (d) The localization error when using feature-estimated RSS to locate the registered transmitter.

of RTLs monitoring a 6MHz TV channel (566-572MHz) by hopping across 3 sections of 2MHz. To create non-uniformity among the sections[2], we configure the transmitter to emit 4MHz signals (567-571MHz). Thus the SNRs of the first and last sections are more than 3dB lower than the second section. Figure 6(f) plots the localization error over time (snapshots), using three approaches: weighted combining, averaging, random (in terms of the best and worst of the three sections). Our proposed weighted combining always outperforms averaging, and performs as good as or even better than picking the best localization result among all the sections.

5.2 Robustness to Interference

We now consider scenarios where both unregistered and registered transmitters are present. We setup a (registered) transmitter to emit OFDM signals with embedded features. After a few minutes, we turned on an interferer that is about 90m away. Users walked by the area (without the knowledge of the two transmitters) recording raw and feature-estimated signal strength. We repeated this experiment multiple times using different power levels at the interfering transmitter.

Figures 7 show 3-minute snapshots of two user routes in three scenarios: no interference, weak interference (the interferer has the same power level as the registered transmitter), and strong interference (the interferer's power is 30dB higher). We see that when there is no interference, the feature is always extracted (Figure 7(a)). Next, Figure 7(b) shows that under weak interference, at locations near the interferer the RTLs detect the difference between the raw RSS and the feature estimated RSS and mark the locations as "interference". Finally, as the interferer becomes stronger, the number of "interference" locations increases and they are located closer to the registered user (Figure 7(c)).

Finally, we use the feature-estimated signal strength to locate the registered transmitter. Figures 7(d) shows the localization results under weak and strong interference using each 5-min snapshot. The error tail increases by 5m when the interference level increases. This is because as the interferer becomes stronger, the number of locations where a feature can be detected reduces, thus providing less input to the localization algorithm. But overall, despite strong interference, our system can locate the registered user at an accuracy similar to that of the scenario without interference.

5.3 End User vs. Vehicle-based Monitoring

Prior works [46, 47] have proposed to use high-end measurement devices, *e.g.* spectrum analyzers ($3,500-$12,000), placed on top of selected city buses, to perform spectrum measurements. As the buses travel around, this approach enables detection and localization of high-power (*e.g.* 3.8W), static, and always-on TV whitespace transmitters. But it faces great challenges when detecting and localizing low-power (*e.g.* 100mw) transmitters with intermittent or dynamic transmissions. Being low-power and dynamic, these transmitters are often "out of sight" of the buses or vehicles (also shown by [47]). Yet they can be covered by nearby walking users within a few minutes.

3 RTLs Walking per 5 min	1 RF Explorer Walking per 5 min	1 RF Explorer Driving 10mph	1 RF Explorer Driving 20mph
25-28	50-120	29-77	42-189

Table 3: Comparing the localization error (in meters) of our RTL-based solution to those using high-end devices.

With this in mind, we compare our RTL based solution to two alternatives using a high-end measurement device. The first is a user walking near the transmitter while holding the high end device. For a fair comparison, we implemented this during our RTL measurements by a randomly-chosen user holding both a RTL and the high-end device to perform measurements. The second is a vehicle with the same high-end device driving by the transmitter. In this case, the transmitter is 90m away from the road. As the vehicle drives by the transmitter, it can obtain measurements over a short period of time. This approach is also used by prior work [47] to detect and localize low-power transmitters, *e.g.* 100mW. For the high-end device we chose the RF explorer because recent work [7] has shown that it has comparable performance to a professional spectrum analyzer (*i.e.,* Agilent N9344C), where the discrepancy in signal estimation in TV spectrum band is bounded by 2.8dB. The antenna attached to the RF explorer has the same gain factor (0dB) as that used by [47].

Table 3 lists the localization accuracy (in meters) of our RTL-based solution (with three users) and the RF-explorer based solutions (with 1 user/vehicle). We see that our solution, by providing more spatial coverage around the transmitter, outperforms the single user RF-explorer approach. The higher error in the vehicle-based approach comes from the fact that the vehicle spends only a short amount of time near the transmitter, which significantly limits the spatial coverage of the measurements. This observation aligns with those of the prior work [47].

6. EVALUATION: ACCURACY AND COST TRADEOFFS

Since an active RTL draws power from the smartphone, a key concern on our solution is whether the energy consumption can discourage users from participating. One potential solution is to reduce the RTL duty cycle to minimize energy consumption, but will this largely degrades the localization accuracy? In this section, we answer this question by performing a detailed study of RTL energy consumption, and exploring the tradeoff between accuracy,

[2] At this 6MHz TV band, the channel is frequency flat so that three sections have identical channel characteristics.

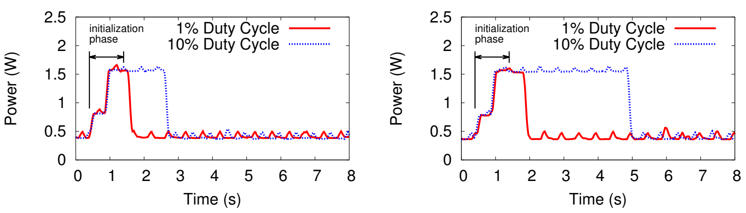

(a) RTL Power Draw: 2.4MHz

(b) RTL Power Draw: 6MHz

Figure 8: RTL power draw for monitoring 2.4MHz and 6MHz spectrum.

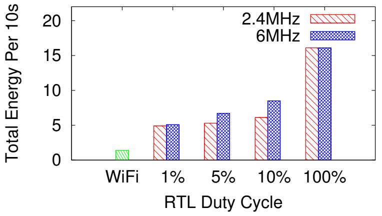

Figure 9: RTL energy consumption per 10s with different duty cycles.

energy consumption and user participation.

6.1 RTL Energy Analysis

Using the Monsoon Power Monitor [5], we measure the smartphone's power consumption every $0.2ms$. To compute the power consumption contributed by the attached RTL, we disable all background activities of the smartphone, turn off the smartphone screen, and measure the power draw when the RTL is unattached to the phone. We use this as a baseline and subtract it from the subsequent power measurements with the RTL attached. Furthermore, to study the impact of RTL duty cycle, we fix a 10s period and vary the RTL active scan duration (per 2.4MHz) between 0.1s and 10s, corresponding to 1% duty cycle and 100% duty cycle. We pick 10s because a longer cycle, *e.g.* 15s, will lead to insufficient monitoring data in each data segment at 1% duty cycle.

Power Draw of Narrowband (2.4MHz) Monitoring. Figure 8(a) plots the sample power draw over time using 0.1s and 1s RTL scan times, when monitoring a 2.4MHz band. We observe an extra "initialization" phase, which lasts for about 1s, and a 100ms "tail" phase. These do not contribute to the RSS measurements but consume power. At 100% duty cycle, these do not exist since the RTL is always on. Aside from the initialization and tail phases, the RTL has two states, idle and active. The idle state draws about 450mW power while the active state draws about 1.5W. These two numbers are slightly higher than those reported by [10], likely due to the difference in RTL manufacturers (although the devices use the same tuner type).

Figure 9 plots the total RTL energy consumption (J) over each 10s for different choices of RTL duty cycle. 100% duty cycle consumes 16.1J per 10s, but 10% duty cycle only consumes 6.1J, mapping to 62% of energy savings. Further reduction of duty cycles from 10% to 5% and 1% offers an additional energy saving of 13% and 19.7%, respectively. The energy reduction is not exactly proportional to the duty cycle reduction, because of the initialization phase that lasts 1s, and the fact that the RTL idle state also draws power. We believe that these overheads can be further optimized to reduce RTL energy consumption. Finally, our RTL devices use the Rafael R820T dongle which consumes 2x of power than another FC0013 model. Thus switching to the FC0013 model can potentially lead to more energy savings. We leave these to future works.

Power Draw of Wideband (6MHz) Monitoring. To monitor a TV band (of 6MHz), the RTL needs to hop across three bands of 2MHz each. Figure 8(b) plots the instantaneous power draw over time. We observe the same "initialization" phase, and the frequency switching is fast (<50ms). Thus the wideband monitoring just increases the RTL active period by a factor of 3. However, in terms of the total energy consumption (per 10s), Figure 9 shows that the 6MHz monitoring leads to less than 3x energy increase for duty cycles 5% - 10%. This happens because the idle state lasts much longer than the active state for these duty cycles and has more impact on the final energy.

Smartphone Battery Life. Having studied the RTL power consumption in detail, we now examine the amount of smartphone battery life when a user participates in our commodity monitoring. For this we need to consider both RTL and GPS energy consumption. For fairness, we do not duty cycle the GPS to match our RTL duty cycle. For the smartphones used in our study (and most smartphone models), GPS, when enabled, reports one reading every 1s [29]. Since constant location queries (via GPS) are becoming more commonplace (e.g., Pokemon Go), we can save energy by reusing cached GPS values. Finally, the energy cost to transmit data to the monitoring agency is negligible. For the 10% duty cycle, the app needs to transmit only 19KB of data per hour. This is implemented as a background activity and runs when other background activities run on the phone.

Figure 10 plots the battery life comparisons among different choices of duty cycle, for both narrowband (2.4MHz) and wideband (6MHz) monitoring. We see that at 10% duty cycle (1s scan time per 10s), the 2.4MHz and 6MHz monitoring can last about 7 hours and 5.8 hours, respectively. Further reduction of RTL duty cycle has marginal improvement because the GPS component still draws a considerable amount of power (423mW by measurements).

RTL vs. WiFi. A recent study has shown that WiFi sensing can be used to localize WiFi APs [50]. As a reference, we also measure the power draw of the WiFi scan and the corresponding battery life if we use it to perform sensing. For the Samsung Galaxy SIII phone (Android version 4.4.2) the WiFi scanning period is 3.5s per 10s. Our measurement shows that such WiFi scan consumes 1.4J energy over each 10s. With GPS on, the battery life is 10.8 hours, which is 3.8 hours longer than our RTL at 10% duty cycle.

Potential Energy Reduction. There are several potential directions to further reduce energy consumption, which we leave to future work. The *first* is to use more energy-efficient RTL hardware. For example, the FC0013 model offers more than 50% energy savings compared with our current RTL hardware [10]. Using this hardware model, we can potentially extend the smartphone battery life from 7 hours to 10.2 hours at 10% duty cycle. This closely matches that of the WiFi scan. *Second*, we can modify the default duty cycle of GPS to match that of the RTL radio. At 10% RTL duty cycle, this modification also increases the battery life to nearly 10 hours using our current RTL hardware, and 20 hours when we switch to the more energy-efficient hardware. *Third*, we can leverage user mobility context to schedule RTL measurements. For example, only when a user starts walking, which can be detected by the smartphone's accelerometer, we turn on the GPS and RTL to perform spectrum measurements.

6.2 Accuracy vs. Energy Consumption

While reducing the duty cycle lowers the energy consumption, it can also affect the localization accuracy. Figure 11 plots the lo-

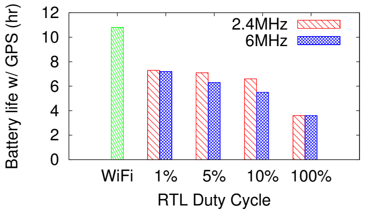

Figure 10: Smartphone battery life at different RTL duty cycles.

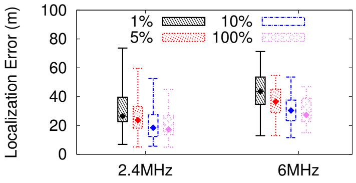

Figure 11: Localization accuracy at different RTL duty cycles.

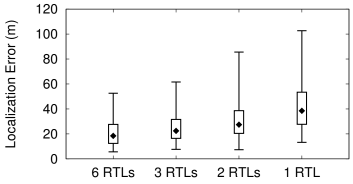

Figure 12: Localization accuracy with different user counts.

calization error for different RTL duty cycles. For both 2.4MHz and 6MHz monitoring, reducing RTL duty cycle does increase the localization error, especially the error tail. For example, the maximum error of 1% duty cycle is about 25m higher than that of 100% duty cycle. On the other hand, the localization performance of 10% duty cycle is on par with that of 100% duty cycle. By comparing the tradeoff between accuracy and energy cost, we find that 5-10% duty cycle is a sweet-spot . In this case, a participating user can expect 5.5-7 hours of battery life, while the monitoring system can bound the localization error by 60m.

6.3 Accuracy vs. User Participation

Finally, we study the tradeoff between user participation and localization accuracy. Figure 12 plots the quantile distribution of the localization error when each 5-min snapshot is gathered by 1, 2, 3, and 6 users. We see that to bound the localization error by 60m, our system does require more than 2 users to perform RTL measurements, *i.e.* they need to move and can capture the transmitter's signal. With one user, the localization error can reach 100m. On the other hand, our system does not require heavy user participation. Having three users actively walking near the transmitter can already achieve reasonable localization accuracy.

7. RELATED WORK

Spectrum Sensing & Measurements. Existing studies develop spectrum sensing techniques on narrowband [8, 34] and wideband signals [21, 45, 39], and improve robustness and scale using compressive sensing (*e.g.,* [26]) and collaborative sensing (*e.g.,* [16]). There are also multiple spectrum measurement platforms [16, 24, 46, 47] and some of them are used to refine TV propagation models [11, 46, 47]. Yet they all require specialized and costly spectrum analyzers (>$3500). Our work differs by using low-cost commodity radios (<$20) and collective user measurements.

Recent works implement low-cost "spectrum analyzers" using RTL and smartphone [10, 31], RTL and Raspberry Pi [33] or smartphone (WiFi) and frequency translator [48]. They also report the receiver noise caused by the low-cost radio. Another recent work [10] examines the energy consumption of RTLs when attached to smartphones. Our work differs by designing robust algorithms to deal with the noisy measurement data, sampling bias and RF interference, and by examining in detail the tradeoff between localization accuracy and energy and user cost.

Spectrum Misuse Detection. Existing studies examined spectrum misuse detection where secondary users interfere with an active primary user. They consider transmission characteristics such as the RSS distribution over space [25, 40], RSS variation [12] and physical channel features [9]. These studies either require dense sensor deployments or the availability of a sensor near each legitimate transmitter, infeasible for large-scale spectrum monitoring. These works also use data generated by propagation models. In contrast, our work uses collective measurements by low-cost RTLs to achieve real-time spectrum monitoring and transmitter location, and our evaluation is based on measurements from real-life scenarios.

Crowdsourcing Measurements. Recent efforts have leveraged crowdsourcing to collect large-scale wireless measurements, using them to characterize signal propagation and user mobility [15, 49], to understand network performance and coverage [19, 37, 38], and to improve indoor localization accuracy [35, 27]. Our work adopts a similar crowdsourcing approach but focuses on achieving real-time spectrum monitoring using low-cost commodity radios.

8. CONCLUSION AND FUTURE WORK

We propose real-time spectrum monitoring measurements using low-cost commodity devices where measurements scale naturally with the number, density and physical reach of mobile users in the network. We use a proof-of-concept platform, *i.e.* smartphone + RTL dongle, to perform empirical validation of the platform. We show that robust data analysis can help commodity measurements overcome a variety of error sources and produce meaningful results.

Moving forward, we plan to perform experiments and take a data-driven approach to multiple issues. *First*, we plan to expand our tests by locating and verifying existing TV whitespace transmitters beyond current measurements. This requires ground truth data on transmitter locations, which we hope to collect from industry partners. *Second*, we plan to expand our energy analysis using other RTL models, and optimize the measurement app to reduce energy consumption. *Third*, we will expand our work to account for false or incorrect data measurements from failures or malicious attackers, and develop mechanisms to identify and remove such anomalous reports.

9. ACKNOWLEDGMENTS

We thank our shepherd Prabal Dutta and the anonymous reviewers for their constructive feedback. This work was supported in part by NSF grants AST-1443956, AST-1443945 and CNS-1224100. Any opinions, findings, and conclusions or recommendations expressed in this material do not necessarily reflect the views of any funding agencies.

10. REFERENCES

[1] http://whitespaces.spectrumbridge.com/whitespaces/home.aspx.

[2] https://www.google.com/get/spectrumdatabase/.

[3] http://sdr.osmocom.org/trac/wiki/rtl-sdr.

[4] https://www.tablix.org/~avian/blog/archives/2015/03/noise_figure_measurements_of_rtl_sdr_dongles/.

[5] https://www.msoon.com/LabEquipment/PowerMonitor/.

[6] M. Altamaimi, M. B. Weiss, and M. McHenry. Enforcement and spectrum sharing: Case studies of federal-commercial sharing. In *TPRC*, 2013.

[7] A. Arcia-Moret, E. Pietrosemoli, and M. Zennaro. WhispPi: White space monitoring with Raspberry Pi. In *Global Information Infrastructure Symposium*, 2013.

[8] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with Wi-Fi like connectivity. In *SIGCOMM*, 2009.

[9] T. Bansal, B. Chen, and P. Sinha. FastProbe: Malicious user detection in cognitive radio networks through active transmissions. In *INFOCOM*, 2014.

[10] N. Brouwers and K. Langendoen. Will dynamic spectrum access drain my battery? *Embedded Software Report Series, ES-2014-01*, 2014.

[11] A. Chakraborty and S. R. Das. Measurement-augmented spectrum databases for white space spectrum. In *CoNEXT*, 2014.

[12] Z. Chen, T. Cooklev, C. Chen, and C. Pomalaza-Raez. Modeling primary user emulation attacks and defenses in cognitive radio networks. In *IPCCC*, 2009.

[13] Y.-C. Cheng, Y. Chawathe, A. LaMarca, and J. Krumm. Accuracy characterization for metropolitan-scale Wi-Fi localization. In *MobiSys*, 2005.

[14] J. M. Dyaberi, B. Parsons, V. S. Pai, K. Kannan, Y.-F. R. Chen, R. Jana, D. Stern, and A. Varshavsky. Managing cellular congestion using incentives. *IEEE Communications Magazine*, 50(11), 2012.

[15] A. Faggiani, E. Gregori, L. Lenzini, V. Luconi, and A. Vecchio. Network sensing through smartphone-based crowdsourcing. In *SenSys*, 2013.

[16] O. Fatemieh, R. Chandra, and C. Gunter. Secure collaborative sensing for crowdsourcing spectrum data in white space networks. In *DySPAN*, 2010.

[17] FCC. Second report and order and memorandum opinion and order. *FCC-08-260*, 2008.

[18] FCC. Report and order and second further notice of proposed rulemaking. *FCC-15-47*, 2015.

[19] A. Gember, A. Akella, J. Pang, A. Varshavsky, and R. Caceres. Obtaining in-context measurements of cellular network performance. In *IMC*, 2012.

[20] D. Han, D. G. Andersen, M. Kaminsky, K. Papagiannaki, and S. Seshan. Access point localization using local signal strength gradient. In *PAM*. 2009.

[21] H. Hassanieh, L. Shi, O. Abari, E. Hamed, and D. Katabi. GHz-wide sensing and decoding using the sparse fourier transform. In *INFOCOM*, 2014.

[22] G. Hsieh and R. Kocielnik. You get who you pay for: The impact of incentives on participation bias. In *CSCW*, 2016.

[23] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck. A close examination of performance and power characteristics of 4G LTE networks. In *MobiSys*, 2012.

[24] A. P. Iyer, K. Chintalapudi, V. Navda, R. Ramjee, V. N. Padmanabhan, and C. R. Murthy. SpecNet: Spectrum sensing sans frontières. In *NSDI*, 2011.

[25] P. Kaligineedi, M. Khabbazian, and V. Bhargava. Malicious user detection in a cognitive radio cooperative sensing system. *IEEE TWC*, 9(8):2488–2497, 2010.

[26] J. Laska, W. Bradley, T. Rondeau, K. Nolan, and B. Vigoda. Compressive sensing for dynamic spectrum access networks: Techniques and tradeoffs. In *DySPAN*, 2011.

[27] L. Li, G. Shen, C. Zhao, T. Moscibroda, J.-H. Lin, and F. Zhao. Experiencing and handling the diversity in data density and environmental locality in an indoor positioning service. In *MobiCom*, 2014.

[28] L. Littman and B. Revare. New times, new methods: Upgrading spectrum enforcement. Silicon Flatirons Roundtable Series on Entrepreneurship, Innovation, and Public Policy, Feb. 2014.

[29] J. Liu, B. Priyantha, T. Hart, H. S. Ramos, A. A. F. Loureiro, and Q. Wang. Energy efficient GPS sensing with cloud offloading. In *SenSys*, 2012.

[30] S. Liu, Y. Chen, W. Trappe, and L. J. Greenstein. Non-interactive localization of cognitive radios based on dynamic signal strength mapping. In *WONS*, 2009.

[31] A. Nika, Z. Zhang, X. Zhou, B. Y. Zhao, and H. Zheng. Towards commoditized real-time spectrum monitoring. In *HotWireless*, 2014.

[32] A. Nika, Y. Zhu, N. Ding, A. Jindal, Y. C. Hu, X. Zhou, B. Y. Zhao, and H. Zheng. Energy and performance of smartphone radio bundling in outdoor environments. In *WWW*, 2015.

[33] D. Pfammatter, D. Giustiniano, and V. Lenders. A software-defined sensor architecture for large-scale wideband spectrum monitoring. In *IPSN*, 2015.

[34] H. Rahul, N. Kushman, D. Katabi, C. Sodini, and F. Edalat. Learning to share: Narrowband-friendly wideband networks. In *SIGCOMM*, 2008.

[35] A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, and R. Sen. Zee: Zero-effort crowdsourcing for indoor localization. In *MobiCom*, 2012.

[36] A. Savvides, C.-C. Han, and M. B. Strivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *MobiCom*, 2001.

[37] S. Sen, J. Yoon, J. Hare, J. Ormont, and S. Banerjee. Can they hear me now?: A case for a client-assisted approach to monitoring wide-area wireless networks. In *IMC*, 2011.

[38] J. Shi, Z. Guan, C. Qiao, T. Melodia, D. Koutsonikolas, and G. Challen. Crowdsourcing access network spectrum allocation using smartphones. In *HotNets*, 2014.

[39] L. Shi, P. Bahl, and D. Katabi. Beyond sensing: Multi-GHz realtime spectrum analytics. In *NSDI*, 2015.

[40] L. Song, Y. Chen, W. Trappe, and L. Greenstein. ALDO: An anomaly detection framework for dynamic spectrum access networks. In *INFOCOM*, 2009.

[41] P. Sutton, K. Nolan, and L. Doyle. Cyclostationary signatures in practical cognitive radio applications. *IEEE JSAC*, 26(1):13–24, 2008.

[42] J. A. Wepman, B. L. Bedford, H. Ottke, and M. G. Cotton. RF sensors for spectrum monitoring applications: Fundamentals and RF performance test plan. *NTIA Report 15-519*, 2015.

[43] L. Yang, Z. Zhang, B. Y. Zhao, C. Kruegel, and H. Zheng. Enforcing dynamic spectrum access with spectrum permits. In *MobiHoc*, 2012.

[44] K. Yedavalli, B. Krishnamachari, S. Ravula, and B. Srinivasan. Ecolocation: a sequence based technique for RF localization in wireless sensor networks. In *IPSN*, 2005.

[45] S. Yoon, E. Li, S. C. Liew, R. R. Choudhury, I. Rhee, and K. Tan. QuickSense: Fast and energy-efficient channel sensing for dynamic spectrum access networks. In *INFOCOM*, 2013.

[46] T. Zhang and S. Banerjee. Inaccurate spectrum databases?:

Public transit to its rescue! In *HotNets*, 2013.

[47] T. Zhang, N. Leng, and S. Banerjee. A vehicle-based measurement framework for enhancing whitespace spectrum databases. In *MobiCom*, 2014.

[48] T. Zhang, A. Patro, N. Leng, and S. Banerjee. A wireless spectrum analyzer in your pocket. In *HotMobile*, 2015.

[49] Z. Zhang, L. Zhou, X. Zhao, G. Wang, Y. Su, M. Metzger, H. Zheng, and B. Y. Zhao. On the validity of geosocial mobility traces. In *HotNets*, 2013.

[50] Z. Zhang, X. Zhou, W. Zhang, Y. Zhang, G. Wang, B. Y. Zhao, and H. Zheng. I am the antenna: accurate outdoor AP location using smartphones. In *MobiCom*, 2011.