

Sybot: An Adaptive and Mobile Spectrum Survey System for WiFi Networks

Kyu-Han Kim¹, Alexander W. Min², Kang G. Shin²

¹Deutsche Telekom R&D Laboratories USA, Los Altos, CA 94022, U.S.A.

²Real-Time Computing Laboratory, The University of Michigan, Ann Arbor, MI 48109, U.S.A.

kyu-han.kim@telekom.com, {alexmin,kgshin}@eecs.umich.edu

ABSTRACT

Information on site-specific spectrum characteristics is essential to evaluate and improve the performance of wireless networks. However, it is usually very costly to obtain accurate spectrum-condition information in heterogeneous wireless environments. This paper presents a novel spectrum-survey system, called *Sybot* (Spectrum survey robot), that guides network engineers to efficiently monitor the spectrum condition (e.g., RSS) of WiFi networks. *Sybot* effectively controls mobility and employs three disparate monitoring techniques—complete, selective, and diagnostic—that help produce and maintain an accurate spectrum-condition map for challenging indoor WiFi networks. By adaptively triggering the most suitable of the three techniques, *Sybot* captures spatio-temporal changes in spectrum condition. Moreover, based on the monitoring results, *Sybot* automatically determines several key survey parameters, such as site-specific measurement time and space granularities. *Sybot* has been prototyped with a commodity IEEE 802.11 router and Linux OS, and experimentally evaluated, demonstrating its ability to generate accurate spectrum-condition maps while reducing the measurement effort (space, time) by more than 56%.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless Communication*

General Terms

Design, Experimentation, Measurement, Performance, Algorithms

Keywords

Spectrum site-survey, adaptive spectrum survey, spectrum map, wireless network, measurement

1. INTRODUCTION

To deal with the exponentially-increasing traffic volume of wireless local area networks [1, 2], a large number of access points

(APs) are being deployed, often redundantly, at homes, offices, campuses, and across cities [3–5]. However, such “chaotic” deployment causes wireless carriers or private network owners to often encounter difficulties in managing the performance and/or the spectrum usage of their networks [6–10].

To cope with such chaotic and complex wireless environments, spectrum site-surveys have been widely used to monitor comprehensive spectrum characteristics. In fact, the spectrum characteristics information is essential for many network services and mobile applications. For example, radio signal propagation from each AP within the deployment area is used as the basis for initial deployment and performance assessment of a wireless network [8, 11]. Accurate spectrum-condition information is key to accurate signal-based localization for mobile devices [12, 13]. Spatial spectrum footprints are useful for mobile users to pinpoint problematic areas for network troubleshooting [14].

Numerous spectrum site-monitoring techniques have been proposed, but they still suffer from several limitations as follows. First, commercial site-survey tools (e.g., [15–17]) can provide the information of Signal-to-Noise Ratio (SNR) from each AP by having a human engineer navigate through the network deployment area. However, such tools often require exhaustive measurements, and it is very challenging to determine site-specific measurement parameters such as measurement frequency, speed, or space granularity, given network or survey requirements. Next, a hybrid approach that uses both measurements and an empirical propagation model, has been proposed (e.g., [11]) for outdoor wireless networks. However, it does not provide the fine-grained survey results necessary for indoor localization systems [12]. Third, a sensor-based approach [13] can reduce manual measurement efforts by deploying spectrum sensors, but it requires the deployment of a large number of sensors, or provides limited survey accuracy due to fixed sensor locations. Finally, the use of a desktop infrastructure [18] has been proposed to avoid the deployment of expensive dedicated sensors. However, such an approach might not be able to produce a comprehensive spectrum map over the entire deployment area.

In this paper, we present an adaptive spectrum-survey system, called *Sybot*, that generates a spatial spectrum-condition map (e.g., Received Signal Strength (RSS) map) for indoor WiFi networks by addressing the following key challenges: (i) how to efficiently produce an accurate spectrum-condition map, (ii) how to maintain an up-to-date map in the presence of temporal and spatial variations in spectrum condition, and (iii) how to automatically determine key survey parameters to meet accuracy requirements on the spectrum map (e.g., <5dBm standard deviation). At its heart, *Sybot* is equipped with an adaptive survey algorithm that consists of three complementary monitoring techniques—*complete*, *selective*,

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiCom'10, September 20–24, 2010, Chicago, Illinois, USA.

Copyright 2010 ACM 978-1-4503-0181-7/10/09 ...\$10.00.

and *diagnostic*. These techniques are adaptively chosen to capture both repeatable and time-varying spectrum conditions, depending on changes in network usage or in physical environments. Briefly, when the network is not in use or lightly-used (e.g., night-time), Sybot invokes the complete monitoring to collect and build a “baseline” or repeatable network-condition map over the entire network deployment area. When the network is in active use (e.g., day-time), Sybot periodically triggers the selective monitoring through which the system can capture the time-variations in spectrum condition, while reducing the measurement overhead by measuring only part of the entire measurement space. When the variations in periodic survey results are unusual/abnormal in certain sub-areas, Sybot triggers the diagnostic monitoring to efficiently identify such “confined” sub-areas and updates the spectrum-condition map of only those sub-areas, as opposed to the entire deployment area.

We have implemented the above components of Sybot atop Linux OS running on a wireless router mounted on an iRobot [19]. We have deployed 12 APs equipped with IEEE 802.11 radios in the fourth floor of Computer Science and Engineering Building at The University of Michigan, and conducted an extensive measurement study using more than 10,000 measurement points for a period of four weeks. We used a robot to facilitate the measurement and evaluation process by automating the mobility each monitoring technique requires. On the other hand, our prototype demonstrates that the use of a mobile robot could be feasible to automate spectrum survey for certain settings such as large warehouses, office buildings, and airports during the night-time.

Our experimental results show that Sybot, indeed, generates a repeatable spectrum-condition map accurately reflecting physical characteristics, such as stationary obstacles. Next, the Sybot’s selective monitoring reduces the measurement space by more than 56 %, compared to the complete monitoring. Finally, the diagnostic monitoring effectively identifies unusual spectrum conditions and maintains the spectrum map to be up-to-date with 50 % less measurement effort than the traditional exhaustive spectrum-survey.

Our analysis of the measurement data confirms the effectiveness of Sybot’s design for site-specific spectrum monitoring as follows. First, Sybot can adaptively determine the granularity of survey parameters such as time interval, unit measurement size, or total measurement space, depending on a specific site (e.g., corridor or hall), distance to APs (close or far away), or unexpected events (obstacles or interferers). Second, Sybot can build and estimate a site-specific profile on the trade-off between accuracy and efficiency of a spectrum survey. For example, Sybot can reduce the measurement overhead by more than 65 % with a 3.5 dBm standard deviation in the survey results (see Section 5.3.3).

The rest of this paper is organized as follows. Section 2 describes the motivation behind this work. Section 3 presents the software architecture and algorithms of Sybot. Section 4 describes the prototype and implementation of Sybot. Section 5 describes our experimental evaluation of the Sybot prototype. Section 6 draws conclusions and discusses some of the remaining issues.

2. MOTIVATION

We first argue for the need of an efficient and accurate spectrum-survey system and then discuss the limitations of existing approaches to meeting this need.

2.1 Why Spectrum Site-Survey?

Despite their coverage benefits, the increasing number of WiFi networks face challenging coordination and performance problems, due mainly to site-specific (hence heterogeneous) spectrum propagation from each AP. To mitigate the problems associated with

spectrum heterogeneity, spectrum site-surveys are commonly used in wireless network services and applications, as can be seen in the following use-cases.

- *Deployment and assessment of wireless networks*: Spectrum site-survey results help network engineers determine the placement of network nodes [8]. When networks are incrementally expanded (as is often the case [2, 5]), one can easily determine where to place new nodes. Even after the deployment of network nodes, a site-survey is necessary to assess their performance [11].
- *Identification of sources of interference*: Spectrum surveys of deployed or neighboring networks allow the network operators to identify interference areas. Various wireless devices, such as cordless phones, may cause interference in certain areas, and measuring and using their spatial footprints is a common way to locate the interfering transmitters [14].
- *Supporting indoor localization systems*: Spectrum site-survey results or spectrum-propagation maps can improve the accuracy of location estimation by providing comprehensive RF signal signatures [12, 20, 21]. By comparing current signal strengths from multiple APs with those in maps, mobile users can estimate their current location.
- *Forecasting the connectivity of mobile users*: Spectrum-survey maps can help mobile users select the best AP to connect to, among those within their range [22]. Although mobile devices can build a connectivity profile of areas they visit, spectrum maps can instantly provide the profile even for new or dynamically-changing sites.

Motivated by the above and other potential use-cases, our goal in this paper is to develop an efficient spectrum site-survey system (equipped with techniques and tools) that accurately monitors the spatial characteristics of radio propagation and that reduces measurement efforts in terms of time and space.

2.2 Limitations of Existing Approaches

There has been a significant volume of work on characterizing spectrum propagation. We discuss below the pros and cons of using existing approaches for spectrum surveys.

Accuracy and repeatability: A spectrum site-survey system must collect information on spectrum-propagation characteristics that not only represents the actual network condition over time, but also shows a repeatable condition specific to each physical site. Here, “repeatability” is an important feature for the survey system to produce a baseline profile of spatial spectrum-condition. Many empirical radio propagation models [23] have been used to calculate the propagation path-loss. However, they are all based on extensive measurements in cellular network environments and have limitations in capturing site-specific spectrum conditions, especially in dynamic and heterogeneous indoor WiFi environments.

Ray-tracing techniques [24] and neural network models [25] have also been proposed to calculate path-loss. Ray-tracing can accurately predict the propagation of a signal by tracing rays from a transmitter at uniform angular intervals in all directions. However, this model requires information about the locations, thickness, and construction materials of walls, ceilings, and floors. The neural network models, such as a multi-layer perceptron algorithm, have been proposed for cellular networks, but they need an extensive training set of terrain information and SNR measurements.

Finally, navigation-based on-line measurement tools [11, 15, 17] have been proposed. These tools help network engineers collect spectrum-propagation information by traversing a path within the

network deployment area. Even though these tools allow for capturing a snapshot of each navigation, their accuracy often relies on several survey parameters, such as survey space granularities and frequency used for each site.

Efficiency and flexibility: A site-survey system must minimize the measurement time and must also be flexible to environment changes. Measurement-based surveys with portable tools [15, 17] are most popular at present, but it is tedious and time-consuming to repeat the same navigation through the deployment area to meet the various application requirements, including network deployment/assessment and RF-based localization systems. The use of empirical models may reduce the measurement overhead (e.g., [11]), while maintaining the site-survey accuracy. However, these models are suitable for outdoor wireless networks, and only provide limited (i.e., coarse-grained) information to indoor localization systems that require unique signal footprints every 1 meter. The deployment of sensors [13, 26] or the use of an inexpensive desktop infrastructure [18] can eliminate the need for physical navigation through the deployment area. However, the fixed locations of the sensors often cause network engineers to perform additional surveys to cover the entire network deployment area, or their locations need to be deliberately altered when the network deployment or physical environment changes. This is difficult to do in practice.

Adaptation and awareness: A site-survey system must be able to dynamically adjust the measurement granularity by recognizing site-specific spectrum characteristics. The navigation-based measurements rely on samples collected in a fixed unit space or at a fixed time interval (e.g., [8]). Such an approach can provide uniform measurement results, but cannot capture spatially heterogeneous spectrum-conditions. For example, if the variance of spectrum condition with respect to an AP in a room is larger than that in the corridor, the uniform spectrum survey is likely to have higher measurement error in the room than in the corridor.

3. THE SYBOT ARCHITECTURE

This section details the architecture of Sybot. First, we discuss the design rationale and overall operation of Sybot. We then present spectrum-survey metrics of interest, and finally, describe the Sybot’s adaptive spectrum-monitoring techniques.

3.1 Overview of Sybot

Sybot is a mobile spectrum-survey system that controls the mobility of a network engineer (or mobile robot) and measures the spatial spectrum conditions of already-deployed IEEE 802.11-based wireless networks using the following features.

- *Periodic and aperiodic monitoring:* Sybot triggers spectrum surveys at both pre-determined (e.g., every morning, afternoon, or evening) and requested (as needed) times to achieve monitoring accuracy. Being equipped with IEEE 802.11 NIC (Network Interface Card), Sybot monitors the spectrum condition of the deployed networks and, based on the monitored results, it guides a network engineer (or mobile robot [19]) to move and conduct spectrum survey at different time intervals and scales, depending on the specifics of each site.
- *Decomposition:* Sybot decomposes spectrum monitoring into three distinct but complementary types of survey, and includes specialized monitoring techniques for each type. This decomposition allows Sybot to selectively choose the best monitoring technique for improving efficiency and accuracy, as opposed to using only one technique during the entire survey period.

Algorithm 1 Sybot operations for the t_i -th survey

- (1) During the measurement period, t_m
 - 1: $\mathbb{L} \leftarrow$ list of APs visible from current grid, g_{cur} ;
 - 2: **for** $j=1$ to n **do** /* n is the size of \mathbb{L} */
 - 3: measure spectrum metrics to every $a \in \mathbb{L}$;
 - 4: move and randomize current location within g_{cur} ;
 - 5: **end for**
 - 6: derive conditions of g_{cur} using the measurements;
 - (2) During the navigation period, t_n
 - 7: $g_{next} \leftarrow$ determine the next grid under p_{cur} ;
 - 8: **if** $g_{next} == \text{NULL}$ **then** /* i -th survey is done */
 - 9: move to a start-point;
 - 10: enter the update period (3);
 - 11: **else** /* more space to survey */
 - 12: move to g_{next} ;
 - 13: enter the measurement period (1);
 - 14: **end if**
 - (3) During the update period, t_u
 - 15: update a spectrum-condition map under p_{cur} ;
 - 16: $c_i \leftarrow$ count grids whose condition deviates by σ ;
 - 17: **if** $c_i > 0$ and $p_{cur} == \text{SELECTIVE}$ **then**
 - 18: $p_{cur} = \text{DIAGNOSTIC}$; trigger p_{cur} ;
 - 19: **end if**
-

- *Use of spatio-temporal variance:* Sybot measures and computes spatial variations in spectrum condition over time, and extracts spatial locality of the condition. These spatial characteristics are then used to identify interference areas, or to minimize the measurement effort in both time and space.
- *Adaptive and controllable monitoring:* Sybot is designed to adapt to the site-specific survey requirements. Depending on the required level of spectrum-monitoring accuracy and site-specific characteristics, Sybot adaptively determines its monitoring granularities and type of monitoring technique to use.

Algorithm 1 describes the overall operation of Sybot, which consists of three sequential periods: (1) measurement period (t_m) during which Sybot directs an engineer (or robot) to navigate within a unit-space (or *grid*¹) determined by the monitoring technique currently in use (p_{cur}) and measures spectrum condition at multiple locations with respect to each accessible AP; (2) navigation period (t_n) during which Sybot determines the next grid to measure based on p_{cur} and moves to that grid. Sybot repeats the measure-then-navigate (for $t_m + t_n$ seconds) until it completely covers the measurement space under p_{cur} ; and (3) update period (t_u) after completing measurements and navigation, during which Sybot constructs or updates a spectrum map based on the measurement results. If Sybot identifies an unusual deviation (e.g., $>$ the normal standard deviation, σ) in the spectrum condition of a grid(s), it triggers on-demand monitoring, i.e., the diagnostic monitoring. Finally, Sybot waits for the next periodic monitoring instant.

3.2 Metrics of Interest

Sybot focuses on the characterization of radio signal propagation as spectrum-monitoring metrics. Specifically, to obtain the signal-propagation characteristics in grid i , Sybot measures the received signal strength (RSS) (γ) at m different locations within the grid during each measurement period (t_m), and uses their mean $\bar{\gamma}_i$ and

¹We assume and use a unit-grid (20 in \times 20 in), in which an engineer (or robot) carrying Sybot incrementally moves, and measures spectrum condition upon completion of each movement. Note that the “grid” can be of different shapes, such as a circle of 20 in radius.

standard deviation σ_i as metrics, which are defined as:

$$\bar{\gamma}_i = \frac{1}{m} \sum_{j=1}^m \gamma_i(j) \quad \sigma_i = \left(\frac{1}{m} \sum_{j=1}^m (\gamma_i(j) - \bar{\gamma}_i)^2 \right)^{1/2}, \quad (1)$$

where $\gamma_i(j)$ is the j -th RSS measurement in grid i . Sybot measures the RSS information by passively monitoring periodic beacon messages from the AP, which are transmitted once every 100 ms [27].² These measurement results for the area A at time t are used to construct a spectrum-condition map $M_A(t)$ for both γ and σ each— $M_A(\gamma, t)$ and $M_A(\sigma, t)$.

Sybot collects the above metrics to generate an accurate spectrum-condition map for the wireless network under consideration with a minimum number of measurements (thus reducing the survey or map-update time). Sybot is flexible enough to use other metrics (e.g., the packet-delivery ratio), but we focus on RSS-based metrics, not only because the RSS is a fundamental parameter to represent network performance, but also because many applications like localization systems use such metrics.

3.3 Adaptive Spectrum Monitoring

As described earlier, Sybot includes three spectrum monitoring techniques—complete, selective, and diagnostic—for efficiency and accuracy. Existing spectrum-survey systems rely mostly on a single measurement technique (e.g., trajectory-based scanning). Sybot decomposes a spectrum survey into three types, in order to reduce the measurement effort as well as to cope with the unpredictable spatio-temporal variations in spectrum condition.

Figure 1 depicts the Sybot’s adaptive monitoring approach. Briefly, when wireless networks are lightly used (e.g., the night-time), Sybot uses the *complete* monitoring to construct a baseline spectrum-condition map on a large-time scale, such as days (T_{day}). When networks are in active use and interfered with by the environment (e.g., moving obstacles or co-existing network activities during the daytime), Sybot uses the *selective* monitoring to capture temporal variations in spectrum condition on a time-scale of hours (T_{hour}). When Sybot detects large variations in spectrum condition in some areas based on complete/selective monitoring results, it triggers the *diagnostic* monitoring to locate such areas, quickly measure and update the spectrum map over a short period of time (T_{min}). Note that the values of the parameters (T_{day} , T_{hour} , T_{min}) are assumed to be provided by network operators, depending on network usage and monitoring requirements (e.g., variations in spectrum condition). Optimizing their values based on network utilization is beyond the scope of this paper.

In what follows, we will detail the above three monitoring techniques and the rationale of their use.

Complete Monitoring

The *complete* monitoring is designed to obtain a baseline radio-propagation characteristic of wireless networks and is triggered at coarse-grained time intervals. Numerous radio-propagation models [23] and computational methods [24, 25] have been proposed to acquire the propagation characteristics. However, they are suitable only for outdoor environments or require extensive human labor to tune the parameters of the models for each site (or even different corridors, as we will show in Section 5.3). Next, measurement-based approaches from stationary APs [7], desktop infrastructure [18], sensors [26], or mobile users [28] help detect performance problems in certain areas, but they do not provide fine-grained spectrum-condition information in the entire network coverage area.

²Sybot focuses on downlinks.




Figure 1: Adaptive approach: Sybot periodically triggers different monitoring techniques on different time scales ($T_{day}, T_{hour}, T_{min}$). Further, it triggers the corresponding techniques (e.g., diagnostic) on demand, whenever Sybot observes (or is informed of) an unusual deviation in spectrum condition.

The complete monitoring performs comprehensive spectrum surveys to collect repeatable spectrum-condition information governed mainly by the physical environments. Such repeatable spectrum-condition information is important, because it provides a baseline profile of spatial spectrum condition and is useful for the design of efficient monitoring techniques—selective and diagnostic. However, to incorporate the complete monitoring into Sybot, there are several challenges to overcome as follows.

- *Building a comprehensive map:* The complete monitoring has to provide a spatially-thorough spectrum map. The approach proposed in [11] estimates the coverage or boundary map of outdoor networks with a small number of measurements, but the boundary map is not good enough for such applications as indoor localization systems that rely on fine-grained spectrum-condition information (e.g., every square foot). The complete monitoring in Sybot virtually divides the network deployment area into small grids and measures the survey metrics of each grid within the deployment area. Note that we use a grid-shape unit space for ease of design. The grid can cover irregular survey areas (e.g., curves) and at the same time, the unit space can be of different shapes like a circle.
- *Selection of grid size:* The grid size is an important factor that balances the accuracy and the efficiency of a site survey. Using a fixed grid size might not accurately capture the heterogeneous spectrum condition of areas with large variances (needing to use a smaller grid size), or might waste time to measure the spectrum condition of open areas with small variances (needing to use a larger size). Thus, the complete monitoring adaptively determines the grid size based on the degree of heterogeneity in spectrum characteristics in a given measurement area. From a measured spectrum map, if the difference in measured signal strengths from neighboring grids is less than a predefined threshold, then Sybot linearly increases the grid size for the space and uses the size for the next survey. In Section 5.3.2, we will detail how to choose the grid size and the threshold by introducing the concept of *RSS distance* (or γ_{dist}).
- *Eliminating temporal variance:* One-time measurement with the complete monitoring can be biased due to unexpected events in certain areas (e.g., moving people or obstacles), and Sybot has to remove such temporal variance in constructing a baseline spectrum map. Sybot maintains a series of spectrum monitoring maps $\mathbb{M}(\gamma) = [M(\gamma, t - n + 1), \dots, M(\gamma, t)]$ and generates the baseline spectrum map $\mathbb{B}(\gamma) = \mathbb{E}[\mathbb{M}(\gamma)]$ based on n recent

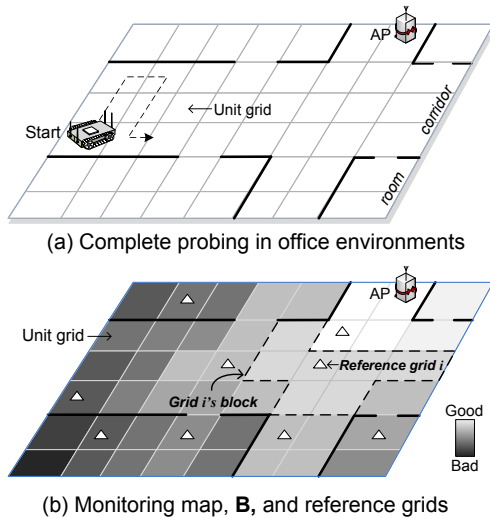


Figure 2: Example of complete and selective monitoring: (a) The complete monitoring progressively measures RSS over the target areas; (b) the complete monitoring result is then used to determine reference grids to capture temporal variations in spectrum condition.

spectrum maps to smooth such temporal variations. Note that although n is site-specific, a small number of spectrum maps is sufficient to produce a baseline map, as we will show in Section 5.3.1.

Figure 2 illustrates the complete monitoring. From the ‘start’ point in Figure 2(a), an engineer (or robot) carrying Sybot navigates through the measurement areas and measures spectrum condition in each unit grid of $20\text{ in} \times 20\text{ in}$, denoted as dots. After completing the measurement of every grid, Sybot accumulates current measurements with the previous spectrum maps and generates a new spectrum map, shown in Figure 2(b). Finally, for each corridor or room, if the variances of most grids are less than the threshold (e.g., 1 dBm), Sybot increases the unit grid size for the space to $40\text{ in} \times 40\text{ in}$ for the next survey.

Selective Monitoring

To reflect changes in the spectrum condition of wireless networks over time (e.g., in the order of hours) [4], Sybot uses the selective monitoring to capture such dynamics and updates the spectrum-condition map accordingly. Sensor-based network monitoring has been proposed to measure such dynamics, especially for the purpose of diagnosing network performance [26] or maintaining accuracy in localization systems [13]. However, they require the deployment of a large number of sensors (e.g., 8 sensors in $30\text{ m} \times 15\text{ m}$), and must painstakingly determine or adjust sensor locations over time. DAIR [18] proposed the use of an inexpensive desktop infrastructure for dense monitoring, but it often suffers from poor accuracy due to the static, unplanned placement of sensors (as we will show in Section 5.3.4).

The selective monitoring makes use of previous spectrum maps to reduce the monitoring space and maintain up-to-date spectrum-condition information. The selective monitoring measures spectrum condition for only a small set of reference grids and estimates spectrum condition over the entire network coverage area. However, to implement this idea in Sybot, there are several issues that must be addressed as follows.

- *Finding spatially-correlated and reference grids:* Sybot must find a group of grids that are spatially-correlated in spectrum condition. By using a complete-monitoring history, Sybot characterizes the site-specific spatial correlation among neighboring grids. Specifically, using the baseline spectrum map (\mathbb{B}), for each grid i , Sybot finds a set of neighboring grids (or block b) whose RSS (γ) is close to grid i 's within a given tolerance (π). Here, this set is called a *block b* , and the grid i is called a *reference grid* of the block.
- *Determining the smallest set:* After determining a block of each grid, because blocks of neighboring (reference) grids may overlap, there will be multiple combinations of reference grids to cover the entire deployment area. Sybot has to determine a combined set of reference grids whose size is minimum to reduce the measurement effort. Finding a globally optimal set of reference grids is an NP-hard problem, $O(2^n)$, where n is the total number of grids. Instead, Sybot uses a heuristic approach with which it iteratively includes the grid with the largest boundary set first in the reference-grid set. This algorithm performs reasonably well in minimizing the set size (see Section 5.3.3) with $O(n \log n + nm)$ complexity, where n is the total number of grids and m the average block size.
- *Controlling accuracy:* There exists an inherent trade-off between the efficiency and the accuracy of measurements, so the selective monitoring must have a knob to control the trade-off, depending on the network requirements. Sybot uses π as the control knob. The lower the π value, the higher the accuracy Sybot can achieve at the cost of more measurements, and vice versa. This trade-off profile can be built and used based on \mathbb{B} , as we show in Figure 12.

Let's consider the example in Figure 2. Using a spectrum map (shown in (b)) and a tolerance threshold (π) of 2 dBm , Sybot determines a set of reference grids, each represented by a triangle. Then, the selective monitoring measures the spectrum condition only for 9 grids, as opposed to the entire 48 grids, while still ensuring the monitoring variance within the tolerance π . Finally, Sybot updates the spectrum condition of correlated grids with the measurements at their reference grids. For the next triggering, Sybot can also rotate the reference grid within the block to opportunistically measure the spectrum condition of all grids within the block over time.

Diagnostic Monitoring

When wireless networks experience local environmental changes such as the appearance of new wireless interference sources or obstacles, Sybot uses the diagnostic monitoring to identify such areas and quickly update the spectrum-condition map of those areas. AP- or sensor-based network monitoring solutions [16, 18, 26] can indirectly detect changes in spectrum condition, and can be used for Sybot to trigger the diagnostic monitoring. However, they still require manual efforts to identify the problematic areas, or simply require engineers to conduct fine-grained a spectrum survey over the entire coverage area.

The diagnostic monitoring in Sybot detects abnormal spectrum condition changes and identifies the areas that need to be surveyed. By measuring only the spectrum condition of the thus-identified areas, Sybot can update the spectrum map very quickly and inexpensively. To implement this monitoring technique, there are, however, several challenges to overcome as discussed below.

- *Detecting abnormal changes:* Sybot must be able to detect the drastic changes in spectrum condition over certain areas, and the diagnostic monitoring makes use of both complete and se-

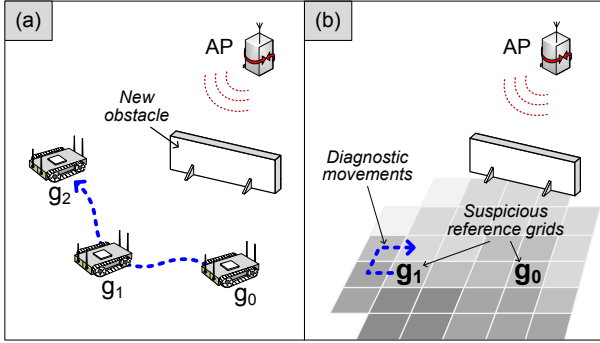


Figure 3: The diagnostic monitoring upon appearance of a new obstacle: (a) Sybot periodically performs the selective monitoring of the reference grids (g_i) to the AP. (b) The diagnostic monitoring identifies suspicious grids that experience unusually large deviations in spectrum condition, and incrementally measures their spectrum condition around the grids.

lective monitoring results for making such a decision. Whenever the selective monitoring is completed, Sybot calculates the difference between the most recent measurement and the baseline measurement ($diff_i = |\gamma_i - \bar{\gamma}_i|$) of each reference grid i . If the difference $diff_i$ is greater than a predefined threshold (e.g., an integer multiple of the grid’s standard deviation (σ_i)), then Sybot immediately triggers spectrum measurements on and around the suspected grids.

- *Speculating measurement areas:* On detecting deviations on the suspicious grids, the diagnostic monitoring has to estimate the spatial boundary of the deviation. By alternating incremental navigation and spectrum measurements, Sybot not only identifies the boundary, but also updates the spectrum map. Specifically, for each suspicious grid i , Sybot progressively navigates in a spiral trajectory and measures spectrum conditions of the grids, until their neighboring grids do not show large deviations. Because neighboring grids are likely to experience the same spatial deviation (i.e., spatial locality), Sybot explores the grids and updates their spectrum-condition information.
- *Exploiting external network monitoring information:* Sybot must be able to exploit network information on spectrum-condition changes. Network-monitoring infrastructures [16, 18] can provide information on network performance degradation at specific APs. Upon receiving such information, Sybot triggers the selective monitoring over the APs’ coverage areas, and then, if necessary, initiates the complete monitoring to update the areas’ spectrum maps.

Suppose that a new obstacle is placed at one location, as shown in Figure 3 (a). Sybot periodically performs the selective monitoring with respect to the AP. Once it finishes the survey, Sybot finds the reference grids g_0 and g_1 having large deviations from their baseline conditions (\mathbb{B}). Sybot then starts the complete monitoring from g_1 to its neighboring grids and applies the same method to g_0 . Finally, Sybot updates the spectrum map with the newly-measured results, as shown in Figure 3 (b).

4. SYSTEM PROTOTYPE OF SYBOT

We have implemented Sybot in Linux and built a proof-of-concept prototype for our evaluation and measurement study.

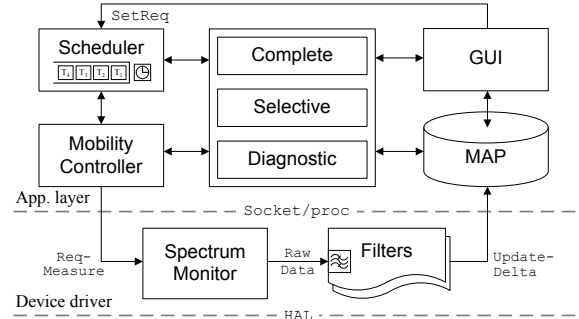


Figure 4: Software architecture of Sybot: The Sybot software design includes (1) a mobility control module in the application layer and (2) a spectrum monitoring module in the device driver (or link layer).

4.1 Software Implementation

Figure 4 shows the Sybot’s software architecture running in a wireless router, which consists of (1) a mobility control module at the application layer and (2) a spectrum monitoring module at the link layer (or device driver).

Mobility control module

This module is responsible for controlling and guiding movements of a network engineer (or a robot), and managing the measurement results. The module is implemented in the application layer and is composed of the following components. First, a graphic user interface (GUI) receives (sends) survey requirements (guidance) from (to) the network engineer. Based on the requirements, GUI initially schedules a spectrum survey. Then, the *scheduler* adaptively triggers complete, selective, or diagnostic monitoring using the algorithms in Section 3.3. During a survey, the *mobility controller* guides an engineer (or a robot) to a target location and triggers the monitoring module to take measurements.

Upon completion of measurements by the monitoring module, the mobility control module updates a spectrum-condition map and schedules the next monitoring technique, time, and areas. We have also implemented a positioning system based on the techniques in [29] for a robot to be used during our evaluation. The system provides high location accuracy (<10 cm error) without relying on any localization infrastructure. Since we merely use the positioning system which is not our claimed contribution, we omit its details.

Spectrum monitoring module

This module is responsible for measuring spectrum-survey metrics within a target space. Specifically, the module is implemented in an open MADWiFi device driver [30] and is composed of two components: spectrum monitor and filters. When the monitoring module receives a measurement request from the mobility controller (via a socket), the spectrum monitor starts collecting information on SNR from APs. Through a hardware abstraction layer (HAL) that Atheros-based chipset [31] provides, the monitor can acquire the above information available in the MAC layer.

Next, the filters (or survey metrics 1) process the collected raw data over the measurement space. Then, the processed information is reflected into MAP through a `/proc` interface.

4.2 Hardware Prototype

In addition to the software implementation of Sybot, we have used a mobile robot to automate our extensive evaluation and data




Figure 5: Sybot hardware prototype: A Sybot node is prototyped with an iRobot, a wireless router, and sonar sensors.

collection process. Although a mobile robot might not be able to provide sophisticated movements (getting around obstacles), basic driving capabilities (forward/backward/spin) are sufficient for navigating through an indoor environment. Moreover, the use of a mobile robot can reduce measurement errors, compared to a human, by taking same movement patterns (e.g., velocity, path, etc.).

Figure 5 depicts the hardware prototype that is composed of a mobile robot, a multi-radio wireless router, and sonar sensors. Specifically, the prototype (i) is built using a commodity robot, called iRobot Create [19] for mobility, which provides a well-defined API for movement control (e.g., a granularity of 1 cm movement) and is powerful enough to carry a wireless router as in [32, 33]; (ii) is equipped with an RB230 wireless router (233 MHz CPU, 128 MB memory) [34], and the router is installed with two IEEE 802.11 miniPCI NICs, each with a 5 dBi omni-directional antenna; and (iii) is equipped with an inexpensive sonar sensor on each side of the robot for estimating the current position of the robot.

5. PERFORMANCE EVALUATION


We have evaluated Sybot via extensive experimentation on our prototype and via thorough analysis of measurement results.

5.1 Testbed Setup


To evaluate Sybot in an indoor environment, we have deployed 12 IEEE 802.11-based APs in the 4th floor of Computer Science and Engineering building at The University of Michigan, with the topology shown in Figure 6(a). Each AP is deliberately placed in the ceiling or shelves to cover the entire 4th floor. All APs are equipped with an omni-directional antenna and operate at the IEEE 802.11a frequencies. Each AP is equipped with an Atheros-based miniPCI NIC and is tuned to use heterogeneous transmission power of 3–10 dBm so that every location in the given limited space may be covered by 3–4 APs.

5.2 Experiment Methodology

In the above testbed, we conducted extensive spectrum surveys using Sybot. Starting from ‘start-point’ in Figure 6(a), Sybot navigates through corridors A, B, C, D, and E and performs the complete, selective, and diagnostic monitoring with respect to each AP. We ran experiments during the early morning or evening hours when all corridors are accessible. During our experiments, people were allowed to walk through the survey areas. This might have caused a temporal variance in measurements, but the variance was negligible, since Sybot essentially used the average of multiple measurements with neighboring grids and measurements taken at different times (e.g., complete monitoring).



(a) Topology



(b) A measured spectrum map of a dotted box in (a)

Figure 6: Testbed topology (210 × 110 ft) and a spectrum-condition map: (a) 12 APs are deployed in our CSE building. (b) An example spectrum map constructed by Sybot for AP-9 over three corridors and a room (dotted box in Figure 6(a)).

During each spectrum survey, we used and tested various time-scales and experimental settings. First, for long-term spectrum measurements, we ran Sybot over the selected APs three times a day during late afternoon hours for 11 consecutive days. Next, for short-term measurements, we ran Sybot 5–10 times a day for every AP in our testbed. For each run, Sybot measures 2–3 corridors per AP and generates a spectrum map per AP. Finally, we used a 20 in × 30 in rectangle as the unit grid size to generate a high-resolution spectrum map and analysis. The size of grid i is said to be x if its area is x times larger than that of the unit grid.

5.3 Experimental Results and Their Analysis

Using the methodology and measurement data, we evaluated and analyzed Sybot’s performance.

5.3.1 Repeatability

We first study the repeatability of Sybot’s complete monitoring. The complete monitoring is said to be *repeatable* if the monitoring results (i.e., measured spectrum conditions) exhibit similar statistical behavior over time. Such repeatable information is important for both selective and diagnostic monitoring techniques.

Sybot periodically performs comprehensive spectrum surveys and each time constructs a baseline spectrum map that reflects the surrounding physical environment. To evaluate Sybot’s repeatability, we randomly selected several APs and analyzed a set of their spectrum maps generated by the complete monitoring over the corridors. Then, we plot the baseline spectrum map that consists of the average ($\bar{\gamma}$) and standard deviation (σ) of measured RSSs for each grid.

Figure 7 shows the constructed baseline spectrum map of $\bar{\gamma}$, where the baseline spectrum map accurately represents several signal-propagation characteristics. For example, $\bar{\gamma}$ in the figure gracefully diminishes as Sybot moves away from the AP located at (1000, 0), and shows the spatial RSS pattern (or gradient) over distance. This




Figure 7: An example of the complete monitoring result (i.e., baseline spectrum map) over a long-term period: The figure shows the average RSS $\bar{\gamma}$ of every grid over AP-3, measured via the complete monitoring, and reflects the real radio propagation over corridors.




Figure 8: Standard deviation σ of the complete monitoring results on Cor-B over a 11-day measurement period: The figure shows that σ is small and stable over time.




Figure 9: Histogram of standard deviation σ of the complete monitoring results: More than 87% of 894 measured grids show less than 4 dBm spread in the measured RSSs.

pattern is indeed repeatable in that the RSSs are stable (small σ) over the measured corridor.

Figures 8 and 9 show the histogram of σ in time and space, respectively. First, Figure 8 shows the average and the standard deviation (s.d.) of σ for Cor-B (112 grids) over a 11-day measurement period. The figure clearly indicates that Sybot provides accurate and stable measurement results (s.d. is less than 3 dBm) during the entire measurement period. This implies that a stable baseline spectrum map \mathbb{B} can be constructed based on a small number of spectrum maps \mathbb{M} . Next, Figure 9 shows the distribution of σ for 5 APs over 4 corridors (894 grids). As shown in the figure, more than 87% grids show a small standard deviation (< 4 dBm). We observed that some areas with high σ (> 6 dBm) actually experience physical changes (e.g., trash cans or doors), which is also captured in the map.

5.3.2 Impact of grid size

Next, we study the effect of grid size on measurement accuracy and efficiency. While a small (fine) grid size provides an accurate spectrum-condition map, it incurs a significant time overhead. Furthermore, determining the optimal grid size is also difficult due to the spatial heterogeneity in spectrum condition. To address these issues, we analyzed the spectrum maps of different corridors (Cor-B and Cor-C in Figure 6(a)), while varying the grid size. While increasing the grid size in multiples of the minimum size (i.e., 20 in \times 30 in), we analyzed the error introduced by the grid size in spectrum survey. For this analysis, we use the metric, called *RSS distance*, to quantify the measurement error, and the error on grid i , $\gamma_{dist}(i)$, is defined as:

$$\gamma_{dist}(i) = \max_k \{\gamma_i(k)\} - \min_k \{\gamma_i(k)\},$$

where $\gamma_i(k)$ is the measured RSS at point k within grid i . Given a set of $\gamma_i(k)$ measurements in grid i , the RSS distance is the difference between the maximum and the minimum RSS values, thus representing the degree of heterogeneity in RSS within the grid. Intuitively, the smaller γ_{dist} , the smaller deviation in the set of $\gamma_i(k)$.

We first compare the impact of grid size at different physical sites (corridors). Figures 10(a) and 10(b) show the empirical cumulative distribution functions (CDFs) of γ_{dist} for three different grid sizes—2, 4, and 6—over Cor-B and Cor-C. As shown in the figures, CDFs rise faster with small-sized grids, indicating better measurement accuracy. This confirms our expectation that a larger grid size introduces a greater measurement error. In addition, for each corridor, different grid sizes make different impacts on the measurement error.

Next, we compare the impact of grid size with different distances from an AP. Figure 10(c) shows the average RSS distance of 3 grid sizes at 3 different distance zones from AP-10. Given a grid size, the average RSS distance is shown to decrease as the distance from the AP increases. This is because that RSS changes (in dB) are more dynamic in a close proximity of the AP than the areas far away from it.

Therefore, the grid size for the complete monitoring should be carefully selected, depending on the physical site and distance to an AP. Using the complete monitoring results, Sybot can build a profile that estimates the impact of each grid size on site-specific spectrum characteristics. Furthermore, because of this non-uniformity of the characteristics, Sybot can apply the selective and diagnostic monitoring techniques to improve efficiency and accuracy.

5.3.3 Reducing the space to measure

Now, we evaluate the effectiveness of Sybot’s selective monitoring in reducing the measurement space during a spectrum survey.




Figure 10: Impact of grid size on the achievable accuracy of complete monitoring : (a-b) The accuracy is measured in terms of the RSS distance for grid sizes 2, 4 and 6 at two different corridors. The figures indicate that the grid size must be adaptively chosen for different measurement spaces or sites. (c) The average RSS distance decreases as the distance from the transmitter (AP) increases.

As discussed in the previous experiment, the spectrum condition is heterogeneous over space. This spatially-heterogeneous spectrum condition can be captured accurately via the complete monitoring with the minimum grid size. Capturing these spectrum characteristics is very useful for several purposes, such as learning/profiling spatial spectrum conditions during initial stages, finding the temporal variations of the areas, or designing a measurement strategy to save the survey resource/time by making only a small number of measurements. However, once Sybot acquires a stable spectrum map for a certain area, repeating the complete monitoring may degrade the spectrum-survey efficiency since such an exhaustive survey incurs excessive time and financial overheads. This can be problematic, especially when Sybot needs to cover a large area with limited time and financial budget.

The selective monitoring reduces the overheads by identifying areas with a similar spectrum condition and merging them together as a unit measurement block. Therefore, the main questions that the selective monitoring has to answer are: (1) how small or large the blocks are, (2) how to make a trade-off between efficiency and accuracy, and (3) how much of benefit the selective monitoring can provide over the complete monitoring.

First, to see the blocks formed by the selective monitoring, we ran Sybot over Cor-B with the tolerance threshold $\pi = 3.5$ dBm, and plotted the spectrum map constructed using the results of the complete monitoring, and the reference grids and measurement blocks chosen by the selective monitoring in Figure 11. Figure 11(a) shows comprehensive spectrum-propagation characteristics over distance. One can also observe that the spectrum conditions in close proximity of the AP, located at (0, 40), are diverse, whereas those far away from the AP are almost monotonic. This spectrum heterogeneity is exploited in selecting the reference grids by the selective monitoring (Figure 11(b)). The maps show that the reference grids (and the measurement blocks) are densely distributed near the AP, because of large spatial and temporal variations in RSSs, while they are sparsely placed where the signal is out of reach. Figure 11(c) shows the measurement blocks for each reference point. The measurement blocks are represented as a set of adjacent grids covered with the same color.

Next, we study the tradeoff between efficiency and accuracy of the selective monitoring. If the tolerance threshold (π) that determines the block size increases, Sybot reduces the number of measurements (or reference grids) at the cost of measurement accuracy. To show this tradeoff, we ran the selective monitoring for Cor-B with the spectrum map constructed based on the results of the com-




Figure 11: Selection of reference grids in selective monitoring (Cor-B): (a) The spectrum map is constructed based on the results from the complete monitoring; (b) Reference grids are so chosen as to cover the entire area with a minimum number of measurements; (c) A complete set of blocks that are generated by the selective monitoring algorithm. The same colored grid(s) represents one block.

plete monitoring and derives how much Sybot can reduce the measurement space. As shown in the Figure 12(a), as the threshold π (in dBm) increases, the selective monitoring becomes more *aggressive* in merging grids, thus reducing the number of reference grids. For example, when $\pi = 3.5$ (dBm), the number of reference grids can be reduced by 70 %, compared to the complete monitoring. Figure 12(b) plots the average and the standard deviation of measurement errors. It shows, on the other hand, that the measurement error increases as the threshold π increases. This is because large threshold π allows the selective monitoring to merge less similar grids, thus degrading the measurement accuracy. Therefore, Figure 12 shows a clear tradeoff between the measurement effort and




Figure 12: Reducing the measurement space and the resultant tradeoff: (a) The selective monitoring minimizes the measurement efforts by reducing the reference points, (b) at the cost of measurement accuracy.




Figure 13: Performance comparison of complete vs. selective monitoring: The selective monitoring reduces the measurement space by up to 72% with the grid size of 10.

accuracy, which can be used as a guideline for planning spectrum measurements.

Third, to study the advantage of the selective monitoring over the complete monitoring, we compare the average RSS distance (γ_{dist}) achieved by the complete and selective monitoring on the measurement data of Cor-B. Figure 13 shows γ_{dist} as a function of average grid size. The figure indicates that the average RSS distance increases almost linearly with the complete monitoring, while the distance remains below 1.5 dBm under the selective monitoring. This advantage comes from the dynamic size of a measurement block used in the selective monitoring. It merges grids with similar spectrum conditions as a unit measurement block, so its γ_{dist} is smaller than that measured by the complete monitoring, which uses a uniform grid size.

5.3.4 Gains from adaptive selection of reference grids

We also study the gains made with the adaptive selection of reference grids in the selective monitoring. To compare its performance, we use a sensor-based spectrum survey that relies on a fixed set of sensors or an existing desktop infrastructure (DAIR [18]). The sensor-based approach is simple and cost-effective since it uses existing wireless devices to capture the changes in spectrum condi-

tion, but such an approach depends highly on the availability and location of sensors. In this comparison, we assume that a baseline spectrum map is available from the complete monitoring, and measure/update the spectrum map over Cor-B, based on the results from the selective monitoring and a sensor-based approach. Finally, we evaluate the accuracy of updated maps by comparing them with a baseline monitoring result.

For fair comparison, we also implement and use an algorithm that effectively uses sensor-based measurements. Briefly, for each sensor, we measure the changes (δ) in the spectrum condition against the baseline spectrum map and then updates the spectrum map by applying δ to the previous condition of the grids around sensors. Furthermore, we consider the scenarios where different numbers of sensors—2, 4, and 6—are available and used for measuring spectrum condition. For each scenario, we perform the experiment multiple times by changing the positions of sensors. On the other hand, the selective monitoring updates the spectrum map by measuring the spectrum condition at the reference grids selected from a baseline spectrum map.

Table 1 shows the gains made with the selective monitoring over the sensor-based spectrum measurement. As shown in the table, even with an increased number of sensors in one corridor, the measurement error is still larger than that of the selective monitoring. In the worst case (i.e., skewed placement of the sensors), the measurement error (3.41) increases by more than 2 times of that (1.39) of the selective monitoring. Furthermore, the selective monitoring reduces the measurement error by up to an average of 51% compared to the sensor-based spectrum measurement, thanks to its ability to adaptively select the reference grids based on site-specific spectrum conditions.

Table 1: Performance comparison of selective monitoring vs. sensor-based approaches: The selective monitoring reduces the measurement error compared to the sensor-based method. Note that the numbers are in dBm.

| Method | Worst | Best | Mean | Std |
|-----------|-------|------|------|------|
| Sensor-2 | 3.41 | 2.01 | 2.73 | 1.84 |
| Sensor-4 | 2.98 | 2.22 | 2.65 | 1.86 |
| Sensor-6 | 3.77 | 2.08 | 2.57 | 1.87 |
| Selective | — | — | 1.39 | 1.60 |

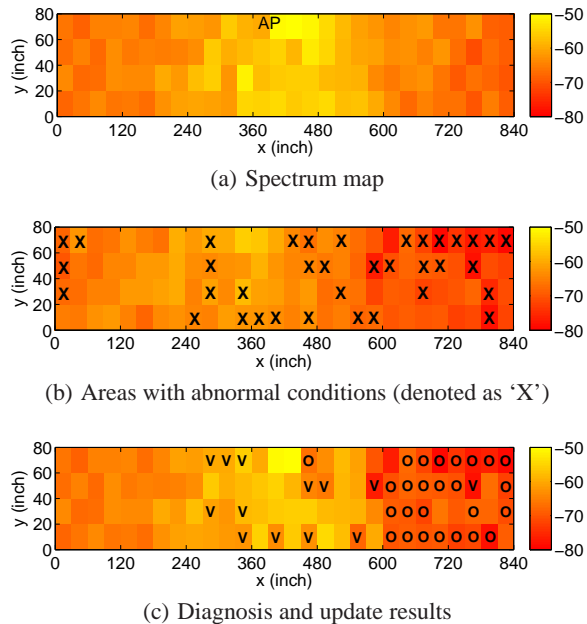


Figure 14: Example of the diagnostic monitoring with an obstacle: (a-b) The appearance of an obstacle (next to the AP) causes abnormal changes in spectrum condition (denoted as X in Figure 14(b)); (c) The diagnostic monitoring identifies the boundaries of areas with abnormal changes using fewer measurements.

5.3.5 Diagnosis of abnormal spectrum condition

We study the Sybot’s effectiveness in detecting and surveying unusual/abnormal changes in spectrum condition. By using both the selective monitoring results (for detection) and the baseline spectrum map (for range estimation), Sybot triggers the diagnostic monitoring to efficiently maintain an up-to-date spectrum map. To evaluate its efficiency, we placed an obstacle in the middle of Cor-C (480, 80) next to the AP (400, 80). Then, we ran the complete monitoring without the obstacle to obtain a baseline spectrum map and ran both the complete and diagnostic monitoring with the obstacle.

Figure 14 shows the above three measurement results, demonstrating the effectiveness of the diagnostic monitoring. The monitoring result without the obstacle (Figure 14(a)) appears as a regular radio propagation from the AP. However, the complete monitoring result with the obstacle (Figure 14(b)) clearly shows the effect of the obstacle and includes a large deviation in spectrum condition in the right side of the corridor. Grids with ‘X’ show larger deviations in their spectrum condition ($diff_i = |\gamma_i - \bar{\gamma}_i| > 2.5 \times \sigma_i$), compared to the baseline spectrum map (Figure 14(a)), due to the appearance of the obstacle. Using the complete/selective monitoring results, Sybot can quickly discover reference grids that experience large deviations (denoted as ‘V’ in Figure 14(c)). Then, Sybot incrementally updates the spectrum map by taking measurements only for those selected grids, which belong to the blocks of the reference grids identified by selective monitoring (‘O’). As shown in the figure, the diagnostic monitoring successfully estimates the problem areas, while reducing its survey space by 56%, compared to the complete monitoring.

6. CONCLUSION

We first discuss some of the remaining issues associated with Sybot and then make concluding remarks.

6.1 Discussion

Multiple APs: Although we presented examples of spectrum monitoring with respect to only one AP in Section 3 for ease of presentation, Sybot can simultaneously monitor the spectrum condition of multiple APs. Sybot employs a time-division measurement strategy for each AP at each grid. This can be extended to use multiple interfaces to perform measurements of multiple APs in parallel.

Multiple Sybots: Multiple Sybots may cooperate to conduct spectrum survey of a large coverage area. For instance, in a large building, an engineer/robot carrying a Sybot can conduct spectrum survey in a certain area (e.g., each floor) and combine his survey results with others’. However, it is difficult to effectively divide the space and merge separate spectrum maps into one. But cooperation among multiple Sybots will be able to reduce the measurement time and improve accuracy. This is part of our future inquiry.

6.2 Concluding Remarks

We presented Sybot, a novel spectrum site-survey system, for efficient and accurate spectrum monitoring in WiFi networks. Sybot adaptively controls mobility and also employs three complementary monitoring techniques that significantly reduce the measurement overhead and provide accurate spectrum-monitoring results under dynamic spectrum conditions. In addition, Sybot provides network engineers important control knobs to determine the trade-off between accuracy and efficiency in spectrum monitoring. Our experimental evaluation shows that Sybot reduces the measurement effort (e.g., the number of measurements) by more than 56%, compared to the conventional exhaustive survey. Moreover, our in-depth analysis of the measurement data has led to several useful guidelines for adjusting important survey parameters of Sybot.

Acknowledgement

This work is supported in part by the US National Science Foundation under Grant CNS-0721529 and NEC Laboratories America. We would like to thank our shepherd, Suman Banerjee, and anonymous reviewers for valuable comments and suggestions. We also thank Chong Chung for helping us collect measurement data through the Undergraduate Research Opportunity Program (UROP) at The University of Michigan.

7. REFERENCES

- [1] Wi-Fi hotspots are hot again, 2009. <http://www.muniwireless.com/2009/11/12/wifi-hotspots-are-hot-again/>.
- [2] AT&T sees big increase in WiFi connections, 2009. http://www.appleinsider.com/articles/09/07/30/att_sees_big_increase_in_wi-fi_connections_following_iphone_3_0.html.
- [3] J. Eriksson, S. Agarwal, P. Bahl, and J. Padhye. Feasibility study of mesh networks for all-wireless offices. In *Proceedings of ACM International Conference on Mobile Systems, Applications, and Systems (MobiSys)*, Uppsala, Sweden, June 2006.
- [4] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *Proceedings of ACM International Conference on Mobile Computing and Networking (MobiCom)*, Philadelphia, PA, September 2004.

- [5] T-Mobile HotSpot, US and Worldwide location, 2009. <https://selfcare.hotspot.t-mobile.com/locations/viewGlobalLocations.do>.
- [6] A. Akella, G. Judd, S. Seshan, and P. Steenkiste. Self-management in chaotic wireless deployments. In *Proceedings of ACM International Conference on Mobile Computing and Networking (MobiCom)*, Cologne, Germany, September 2005.
- [7] L. Qiu, P. Bahl, A. Rao, and L. Zhou. Troubleshooting multi-hop wireless networks. In *Proceedings of ACM SigMetrics (extended abstract)*, Alberta, Canada, June 2005.
- [8] A. Raniwala and T.-C. Chiueh. Deployment issues in enterprise wireless LANs. Technical Report TR-145, Technical Report, Experimental Computer Systems Lab, State University of New York, September 2003.
- [9] K.-H. Kim and K. G. Shin. On accurate measurement of link quality in multi-hop wireless mesh networks. In *Proceedings of ACM International Conference on Mobile Computing and Networking (MobiCom)*, Los Angeles, CA, September 2006.
- [10] K.-H. Kim and K. G. Shin. Extended abstract: Self-healing multi-radio wireless mesh networks. In *Proceedings of ACM International Conference on Mobile Computing and Networking (MobiCom)*, Quebec, Canada, September 2007.
- [11] J. Robinson, R. Swaminathan, and E. W. Knightly. Assessment of urban-scale wireless networks with a small number of measurements. In *Proceedings of ACM International Conference on Mobile Computing and Networking (MobiCom)*, San Francisco, CA, September 2008.
- [12] Paramvir Bahl and Venkata N. Padmanabhan. Radar: An in-building RF-based user location and tracking system. In *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)*, 2000.
- [13] J. Yin, Q. Yang, and L. M. Ni. Learning adaptive temporal radio maps for signal-strength-based location estimation. *IEEE Transactions on Mobile Computing*, July 2008.
- [14] W. Xu, T. Wood, W. Trappe, and Y. Zhang. Channel surfing and spatial retreats: Defenses against wireless denial of service. In *Proceedings of ACM Workshop on Wireless Security (WiSe)*, Philadelphia, PA, October 2004.
- [15] AirMagnet, Inc., 2002. <http://www.airmagnet.com>.
- [16] Newbury Networks, 2007. <http://www.newburynetworks.com>.
- [17] MetaGeek, LLC. <http://www.metageek.com>.
- [18] P. Bahl, J. Padhye, L. Ravindranath, M. Singh, A. Wolman, and B. Zill. DAIR: A framework for managing enterprise wireless networks using desktop infrastructure. In *Proceedings of ACM Workshop on Hot Topics in Networks (HotNets-IV)*, November 2005.
- [19] iRobot Corp. <http://www.irobot.com>.
- [20] Andrew M. Ladd, Kostas E. Bekris, Algis Rudys, Lydia E. Kavvaki, and Dan S. Wallach. Robotics-based location sensing using wireless Ethernet. *Wireless Networks*, 11(1–2):189–204, January 2005.
- [21] J. Yin, Q. Yang, and L. Ni. Adaptive temporal radio maps for indoor location estimation. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom)*, Kauai, Hawaii, March 2005.
- [22] A. J. Nicholson and B. D. Noble. Breadcrumbs: Forecasting mobile connectivity. In *Proceedings of the ACM International Conference on Mobile Computing and Networking (MobiCom)*, San Francisco, CA, September 2008.
- [23] T. S. Rappaport. *Wireless communications: Principles and practice*, second edition, January 2002.
- [24] K. R. Schaubach, N. J. Davis, and T. Rappaport. A ray tracing method for predicting path loss and delay spread in microcellular environments. In *Proceedings of IEEE Vehicular Technology Conference*, May 1992.
- [25] A. Nesovic, N. Neskovic, and D. Paunovic. Macrocell electric field strength prediction model based upon artificial neural networks. *IEEE Journal on Selected Areas in Communications*, August 2002.
- [26] Y.-C. Cheng, J. Bellardo, P. Benko, A. C. Snoeren, G. M. Voelker, and S. Savage. Jigsaw: Solving the puzzle of enterprise 802.11 analysis. In *Proceedings of ACM Special Interest Group on Data Communication (SIGCOMM)*, Pisa, Italy, September 2006.
- [27] H. Lee, S. Kim, O. Lee, S. Choi, and S. J. Lee. Available bandwidth-based association in IEEE 802.11 wireless LANs. In *Proceedings of ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, Vancouver, Canada, October 2008.
- [28] A. Adya, P. Bahl, R. Chandra, and L. Qiu. Architecture and techniques for diagnosing faults in IEEE 802.11 infrastructure networks. In *Proceedings of ACM International Conference on Mobile Computing and Networking (MobiCom)*, Philadelphia, PA, September 2004.
- [29] J. J. Leonard and H. F. Durrant-Whyte. Mobile robot localization by tracking geometric beacons. *IEEE Transactions on Robotics and Automation*, pages 376–382, June 1991.
- [30] MADWiFi. <http://www.madwifi.org>.
- [31] Atheros Communications. <http://www.atheros.com>.
- [32] P. De, R. Krishnan, A. Raniwala, K. Tatavarthi, N. A. Syed, J. Modi, and T. Chiueh. Mint-m: An autonomous mobile wireless experimentation platform. In *Proceedings of ACM International Conference on Mobile Systems, Applications, and Services (MobiSys)*, Uppsala, Sweden, June 2006.
- [33] J. Reich, V. Misra, and D. Rubenstein. Roomba MADNeT: A mobile ad-hoc delay tolerant network testbed. *SIGMOBILE Mobile Computing and Communications Review*, 12(1):68–70, 2008.
- [34] Router Board. <http://www.routerboard.com>.