

# Assessment of Urban-Scale Wireless Networks with a Small Number of Measurements

Joshua Robinson  
Rice University  
Houston, TX  
jpr@rice.edu

Ram Swaminathan  
HP Labs  
Palo Alto, CA  
ram.swaminathan@hp.com

Edward W. Knightly  
Rice University  
Houston, TX  
knightly@rice.edu \*

## ABSTRACT

In order to evaluate, improve, or expand a deployed, city-wide wireless mesh network, it is necessary to assess the network's spatial performance. In this paper, we present a general framework to accurately predict a network's well-served area, termed the *metric region*, via a small number of measurements. Assessment of deployed networks must address two key issues: non-uniform physical-layer propagation and high spatial variance in performance. Addressing non-uniformity, our framework estimates a mesh node's metric region via a data-driven sectorization of the region. We find each sector's boundary (radius) with a two-stage process of estimation and then measurement-driven "push-pull" refinement of the estimated boundary. To address high spatial variation, our coverage estimation couples signal strength measurements with terrain information from publicly available digital maps to estimate propagation characteristics between a wireless node and the client's location. To limit measurements and yield connected metric regions, we consider performance metrics (such as signal strength) to be monotonic with distance from the wireless node within each sector. We show that despite measured violations in coverage monotonicity, we obtain high accuracy with this assumption. We validate our estimation and refinement framework with measurements from 30,000 client locations obtained in each of two currently operational mesh networks, *Google-WiFi* and *TFA*. We study three illustrative metrics: *coverage*, *modulation rate*, and *redundancy*, and find that to achieve a given accuracy, our framework requires two to five times fewer measurements than grid sampling strategies. Finally, we use the framework to evaluate the two deployments and study the average size and location of their coverage holes as well as the impact of client association policies on load-balancing.

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

\*This research was supported by NSF grants CNS-0325971 and CNS-0721894 and by HP Labs.

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 '08*, September 14–19, 2008, San Francisco, California, USA.  
Copyright 2008 ACM 978-1-60558-096-8/08/09 ...\$5.00.

**General Terms:** Measurement, Performance, Reliability, Experimentation, Design, Algorithms

**Keywords:** Mesh Networks, Coverage, Measurement Assessment, Deployment, Wireless, Access Network, Signal Strength

## 1. INTRODUCTION

City-wide wireless mesh networks provide Internet connectivity through the deployment of outdoor 802.11-based mesh nodes [4]. To evaluate, expand, or improve performance in a deployed network, a network operator must first assess the current spatial performance of the network. We formulate the wireless network assessment problem as a problem of identifying *metric regions*, i.e., identifying locations in the network where the given performance metric meets a conformance threshold. Existing assessment strategies either require exhaustive measurements [19] or use detailed physical-layer object descriptions to precisely estimate propagation characteristics [20]. Unfortunately, these approaches are expensive and often impractical, especially for incremental network upgrades.

In this paper, we present a general framework to assess the spatial performance of a deployed mesh network using a constrained number of measurements. We estimate metric regions by coupling the use of coarse-grained terrain maps with the construction of virtual sectors of differing radii overlaid on the physical topology. We evaluate the framework itself and study two deployed urban mesh networks using measurement sets from approximately 30,000 client locations in each network. In particular, our contributions are as follows.

First, we devise a mesh network assessment framework that divides each node's metric region into a number of virtual *sectors*. We use a two-stage process to first estimate the metric sector boundaries (radii) and then to refine each boundary estimate through the selection of a small number of measurements. For coverage, the difficulty in estimating a metric region is due to complex and highly variable interactions with the physical environment, e.g., see [13]. Thus, we use the geometry of the terrain obtained from publicly available digital maps to account for differences in propagation characteristics among regions. Moreover, we refine the estimated sector boundaries with a small number of measurements guided by a push/pull heuristic that selects measurement locations and adjusts the estimated boundary. By assuming a monotonic decay in performance with distance, we model a metric region as a single, connected, multi-sector, multi-radii region. We show that despite some monotonicity violations, we obtain high characterization accuracy.

We then validate our framework and evaluate two oper-

ational networks by performing an extensive set of client measurements from two currently deployed wireless mesh networks, *GoogleWiFi* and *TFA*.<sup>1</sup> We show that for a given accuracy in describing metric regions, our framework requires two to five times fewer measurements than a grid sampling strategy. We compare our boundary refinement heuristic to a simplified ray-tracing approach, and show that our heuristic is more robust to monotonicity violations and obtains better accuracy. Using our framework, we find that the TFA network is deployed with sufficient density so that coverage holes occur only on the network edges. However, for the GoogleWiFi network, the *frequency* of coverage holes is much less dependent on deployment density, although the *size* of the holes decreases for higher node densities. This points to a key challenge of covering holes with additional nodes: a large number of additional nodes would be required to eliminate numerous small coverage holes in which half have radius of less than 10 meters. Lastly, we investigate client association policies and find a 20% loss in client throughput due to uneven spacing of mesh nodes.

The rest of this paper is organized as follows. Section 2 defines the assessment problem and our proposed framework. Section 3 presents estimation, sectorization, and measurement-based refinement algorithms. Section 4 validates the accuracy of our framework with real data sets and Section 5 evaluates two studied deployments. Section 6 discusses related work and finally, Section 7 concludes.

## 2. ASSESSMENT FRAMEWORK

In this section, we formally define the network assessment problem and then set up the required notations and definitions for our estimation and refinement framework.

### 2.1 Problem Definition

The problem we address is how to accurately characterize the coverage of a network with a small number of measurements. Generally, we characterize a network by identifying specific areas in the network where the measured value of a given metric (e.g., signal strength) exceeds a given threshold. Therefore, a good characterization scheme is one with a high accuracy for identifying unmeasured locations as above or below the given threshold. This formulation addresses scenarios such as: 1) a network operator wishing to identify dead-spots in order to add nodes and improve performance, or 2) a municipality wishing to determine if a deployed network conforms to contractual performance requirements. We next formally define the assessment problem with a constraint on the number of measurements allowed.

We consider a terrain  $\mathcal{T}$ , which consists of a continuous space of points,  $p \in \mathcal{T}$ , on a 2-d Cartesian plane. Similarly, we define the set of mesh nodes  $\mathcal{N}$ , where each node  $n \in \mathcal{N}$  is defined by a coordinate pair in 2-d space. Let  $M$  represent a specific performance metric; our study focuses on an SNR-based coverage metric, but also includes modulation rate and redundancy. For each point  $p$ , we define  $M(p)$  as the measurable value of metric  $M$  at point  $p$ . Measurement cost is assumed to be identical for all points  $p \in \mathcal{T}$ , but not for all metrics.

We begin by characterizing a single point with respect to a given metric  $M$  and given threshold,  $\theta_M$ , which represents the metric's performance cutoff. A point  $p$  satisfies metric  $M$

if  $M(p) \geq \theta_M$ . Characterization is then defined as correctly *predicting* if a location satisfies metric  $M$ . For metric  $M$ , a mesh node's metric region is the set of all points  $p \in \mathcal{T}$  such that  $M(p) \geq \theta_M$ . A mesh *network's* metric region is the union of all mesh node metric regions in set  $\mathcal{N}$ .

The role of measurements in the characterization of a network is to gain additional knowledge with which to increase the accuracy of predicting the value of  $M$  at an unknown location. In order to limit the measurement expense of the assessment study, we add a constraint that limits the total number of measurements. Stated as an optimization problem, we seek to maximize the characterization accuracy over a terrain  $\mathcal{T}$  subject to a constraint on the total number of measurements taken. We consider both versions of the problem where we characterize the metric region of a single node or of an entire network.

A key challenge for signal strength based metrics is that physical-layer transmissions do not propagate uniformly at all angles from a mesh node and signal strength does not monotonically decrease with distance. Further, there is no known practical way to *a priori* characterize the large changes in signal strength over short distances.

### 2.2 Metric Definitions

We first define three performance metrics for coverage, modulation rate, and redundancy. The coverage metric is based on the received signal-to-noise ratio (SNR), labeled  $P_{dB}(p, n)$ , at a client point  $p$  from node  $n$ . A conformance threshold,  $\theta_c$ , indicates the minimum acceptable SNR.

**Definition 1.** Consider a terrain  $\mathcal{T}$ , a location  $p$ , and a mesh node  $n$  in  $\mathcal{T}$ . The location  $p$  is **covered** by  $n$  if the received SNR at  $p$  with respect to  $n$ ,  $P_{dB}(p, n) \geq \theta_c$ . The **coverage region** of  $n$  is the set of all points in  $\mathcal{T}$  covered by  $n$ .

The second metric is modulation rate, which captures the expected value of the physical-layer modulation rate in use at a given location. This value is a function of SNR and the rate selection protocol used, e.g., Auto-Rate Fallback (ARF).

**Definition 2.** Let  $n$  be a mesh node and  $p$  be a client location in a terrain  $\mathcal{T}$ . The **modulation rate** of  $p$  with respect to  $n$  is the expected physical layer modulation rate in use. The **modulation-rate region** of  $n$  is the set of all points in  $\mathcal{T}$  with expected modulation rate at least threshold  $\theta_r$ .

We now define the coverage redundancy metric, which is based directly on the coverage metric and is the number of mesh nodes which cover a given point.

**Definition 3.** The **redundancy** of a location  $p$  in a terrain  $\mathcal{T}$  is the number of mesh nodes that cover  $p$ . The  **$k$ -redundancy region** of  $\mathcal{T}$  is the set of all points in  $\mathcal{T}$  with redundancy  $k$  or greater.

### 2.3 Metric Sector Framework

Our framework uses terrain information to divide the mesh node metric region into virtual sectors of varying angular widths and radii. To accurately characterize the network's diverse propagation environment, we independently estimate sector angles and boundaries

<sup>1</sup><http://wifi.google.com> and <http://tfa.rice.edu>.

More formally, a **metric sector** of mesh node  $n$  is a sector of the circle centered at  $n$  contained between angles  $\phi_1$  and  $\phi_2$ . We consider *monotonic* performance metrics defined as follows. Let the function  $\mathbf{d}(p_1, p_2)$  denote the distance between points  $p_1$  and  $p_2$  in a terrain, then:

**Definition 4.** Let  $\mathcal{T}$  be a terrain and  $M$  be a metric.  $M$  is **monotonic** in  $\mathcal{T}$  if for every mesh node  $n$  in  $\mathcal{T}$ , for any ray  $R$  emanating from  $n$  and for any two points  $p_1$  and  $p_2$  on  $R$ , if  $\mathbf{d}(p_1, n) < \mathbf{d}(p_2, n)$ , then  $M(p_1) \geq M(p_2)$ .

While we assume performance measures such as signal strength decay monotonically for each ray, the use of multiple sectors with different radii does *not* require monotonicity among rays nor among sectors. For example, a far away signal strength can be greater than that of a closer distance provided that the two points are on rays having different angle from the originating node.

We assume that this monotonicity property is satisfied for coverage and show later in Figure 8 that the coverage metric mostly satisfies this property. The modulation rate metric also satisfies monotonicity, whereas the redundancy metric does not.

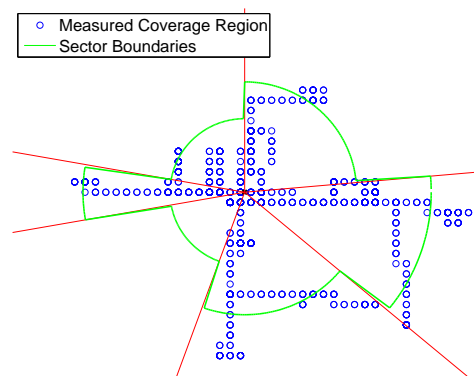
Let the *boundary* of a metric sector be the arc segment between angles  $\phi_1$  and  $\phi_2$ , which defines the sector’s border at radial distance  $r$  from the mesh node. With this definition, we characterize a monotonic metric at an unknown location based on whether it is inside the metric boundary or not. The disjoint union of all metric sectors and sector boundaries defines the metric region. Note that the region boundary is non-uniform as it depends on the environment specifics in the region, and is different for each performance metric.

Thus, our proposed framework overlays a sector-based structure on the assessment problem. The objective to maximize predictive accuracy translates to minimizing the difference between the estimated and true metric boundary. The framework provides three types of variables to optimize on a *per-node* basis: 1) the number of sectors, 2) each sector’s boundaries,  $\phi_1$  and  $\phi_2$ , and 3) the boundary distance  $r$  for each sector. The optimal solution is approached as the number of sectors goes to infinity, allowing the boundary to vary over smaller and smaller angles. In practice, we employ a small number of sectors because there is significant correlation over moderate angular distances, and the grouped boundary allows refinement with few measurements per sector, increasing overall accuracy.

Figure 1 depicts an example of our framework’s operation with a mesh node in the center of the figure, six virtual sectors displayed, and the estimated sector boundaries. In the next section, we present our techniques for choosing the sector borders  $\phi$  and boundary distances  $r$ , in order to heuristically improve the selection of boundaries.

### 3. ESTIMATION AND REFINEMENT

In this section, we describe our estimation techniques, including an estimator for coverage that exploits terrain information from digital maps. We then show how to use the estimates to drive sectorization so that chosen sectors have an approximately uniform boundary throughout. Lastly, we present an online heuristic to choose measurement locations in order to refine the metric sector boundary.



**Figure 1: Example metric sectorization and boundaries for an example GoogleWiFi mesh node’s coverage region.**

### 3.1 Performance Metric Estimation

We first present a coverage estimator which exploits terrain information to improve accuracy. We also introduce simple estimators for the modulation rate and redundancy metrics. Both coverage and modulation rate satisfy monotonicity, while redundancy does not, although it is calculated as a function of the coverage metric.

For coverage estimation, the environment has an average propagation environment (path loss) throughout. Yet, specific areas exhibit different propagation behavior due to different terrain (e.g., streets vs. buildings). Thus, an antenna’s transmission not only experiences different attenuation at each angle, but each ray also faces varying attenuation as it moves away from the source. To address this uncertainty, our key technique is to couple terrain maps with measurements in order to better estimate SNR at a point. We do this by calculating an average path loss for the entire network, and then for each measurement pair, we use the terrain information to estimate the shadowing, i.e., the deviation (in dB) from the average path loss. We next describe the terrain information and then our estimation equations.

*Terrain features* encompass any type of physical area of the input map, such as buildings, fields, or trees, all of which are approximated with polygons. Figure 2 shows the publicly available digital map that we use to extract the Google-WiFi terrain feature information. The TFA terrain map is not shown, but is similar. The number of different feature types and resolution of the terrain features determines the amount of information gained from the map, and is dependent on how the map processing algorithm groups similar features. Edge-detection image processing algorithms can be used to input satellite and city maps [9]. The maps used in our evaluation show zoning information and so a simple heuristic algorithm suffices to perfectly identify all terrain features. The output of the map processing algorithm is the set of polygons representing the terrain features. We then use training measurements to assign attenuation weights,  $C_f$ , to each feature type to indicate the feature’s impact on pathloss estimation. Note that our studied networks feature homogeneous antenna heights and we restrict to 2-d maps.

We estimate coverage using the standard log-normal path loss equation with shadowing [19]. Our key technique is to use terrain features to estimate the shadowing value for




Figure 2: Terrain map of Mountain View, CA, for the measurement study area.

each individual link. Shadowing accounts for the random variations in signal strength between node and client pairs at the same distance  $\mathbf{d}(n, p)$ , which are due to differences in the scattering and attenuation environment and is usually represented as a zero-mean Gaussian random variable [19].

Therefore, instead of estimating based only on average path loss, we also define a terrain-informed shadowing estimator,  $\beta(n, p)$ , to capture the specific path's deviation (higher or lower) from the average path loss. Recall that the received power  $P_{dB}$  is a function of the measured power,  $P_0$ , at reference distance  $d_0$ , and the average path loss exponent  $\alpha$ . Our estimate for the SNR is then:

$$P_{dB}(p, n) = P_0 - 10\alpha \log \frac{\mathbf{d}(n, p)}{d_0} + \beta(n, p) \quad (1)$$

The terrain-informed estimator,  $\beta(n, p)$ , depends on a) the terrain features in  $\mathcal{T}$  that lie along the ray between the mesh node  $n$  and point  $p$ , b) the width of this ray's intersection with each feature, and c) the feature type and weight,  $C_f$ . Specifically,  $\beta(n, p)$  is defined as the sum of each intervening feature's impact on pathloss:

$$\beta(n, p) = \sum_{f \in F} C_f \times \mathbf{w}(n, p, f) \quad (2)$$

where  $F$  is set of all features in the terrain  $\mathcal{T}$ ,  $C_f$  is the weight of a feature (attenuation in dB per unit distance), and  $\mathbf{w}(n, p, f)$  is the intersection width of the ray between  $n$  and  $p$  on the terrain feature  $f$ . In other words, each terrain feature that a link intersects either adds or subtracts from the value of the estimated pathloss, as a function of the feature weight  $C_f$ .

The  $\alpha$  and  $C_f$  terms above must be determined with some measurement overhead for each network. Training measurement locations are chosen randomly throughout the terrain, where each link intersects a subset of the terrain features in question. The training measurements must pass through a representative set of terrain features to capture each feature's effect on pathloss. In other words, we take measurements driving around the edges of terrain features, as opposed to measurements within each feature. The mea-

sured SNR values and measurement distances, in combination with Equations (1) and (2), then lead to a system of equations with the parameters as unknowns. We use minimum-mean squared error fitting to choose values of  $\alpha$  and  $C_f$  which best fit the measurements and equations. Section 4 studies the number of measurements needed per feature type, found to be between 10 and 20 for high accuracy estimation. Our approach for incorporating small-scale terrain features builds upon empirical models for outdoor path loss prediction [16] in macrocells with adjustments for terrain environments.

The modulation rate estimator builds upon the coverage metric as follows. The constant  $C_r$  maps SNR to an expected modulation rate choice,  $T(n, p)$ , from the set of possible physical layer modulation rates as:  $T(n, p) = C_r \times P_{dB}(p, n)$ , where  $C_r$  is dependent on the interface technology in use. Finally, estimation for redundancy derives directly from the coverage estimation discussed previously.

### 3.2 Estimating Monotonic Metric Regions


The objective of the estimation algorithm is two-fold: to choose sector locations (angles  $\phi_1$  and  $\phi_2$ ) and estimate the metric boundary distance of each sector. Because the total number of measurements needed is a function of how many sectors are used, the algorithm merges boundary sections to reduce the number of sectors considered to the desired number and to output sectors with approximately uniform propagation throughout. The only measurements required for this algorithm are training measurements for the metric estimation function parameters, e.g., values of pathloss exponent  $\alpha$ .

Algorithm **Estimate-Mono-Metric-Region** (terrain  $\mathcal{T}$ , mesh nodes  $\mathcal{N}$ , metric  $M$ )

1. For every mesh node  $n \in \mathcal{N}$ , do Steps 2 through 6.
2. Pick a set of rays at uniformly spaced angles from the mesh node. Call this set  $R$ , where the number of rays is chosen to be significantly larger than the desired number of final output sectors.
3. For each ray in set  $R$ , a) traverse the ray along the terrain map identifying terrain features and the respective type and attenuation; b) estimate the value of the metric  $M$  using a metric estimator to identify the boundary point  $x$  on the ray, where  $M(x) = \theta_M$ ; and c) connect the boundary points  $x$  on each ray to identify the estimated boundary of the metric sector.
4. Create a mapping from each ray's angular position to the estimated metric boundary distance,  $\mathbf{d}(n, x)$ .
5. Curve fit a step function to the above mapping, minimizing the mean-squared error between the estimated boundary distance and step function approximation. The number of steps corresponds to the number of allowed sectors, the height of each step is the boundary distance of each sector, and the cutoff points of each step are the sector border angles  $\phi_1$  and  $\phi_2$ .
6. Output set of sectors with borders defined by step function cutoff points.

Figure 3 shows an example of the estimation algorithm output. We divide the region surrounding the mesh node

into 360 sectors with equally spaced rays (one per sector) and estimate the boundary distance of each sector. Since estimation requires only a constant number of training measurements, the number of rays chosen in this step does not increase the measurement budget. We then merge the sectors to result in the sectorized ranges also plotted (step function). Also, included are the measured ranges for this mesh node for the angles where we have data points available.



**Figure 3: Example of sectorization process in algorithm Estimate-Mono-Metric-Region for a Google-WiFi mesh node. As a function of angle from the mesh node, we plot the estimated ranges, the sectorized estimate, and the measured ranges.**

### 3.3 Estimated Boundary Refinement

We now describe our algorithm to choose measurement locations in order to refine the boundary estimate of each sector. The refinement is challenging due to monotonicity violations and a noisy boundary. Therefore, we present a push/pull refinement heuristic, which is robust to these challenges by keeping little state in order to recover from anomalous measurements. Generally, the algorithm measures at the estimated boundary and then either pulls or pushes the estimated boundary nearer or farther from the mesh node based on the measurement result. The algorithm terminates when a boundary is found or the measurement budget per sector is exceeded. Algorithm **Refine-Estimate** (terrain  $\mathcal{T}$ , wireless nodes  $\mathcal{N}$ , sectors, metric  $M$ )

1. For each mesh node  $n \in \mathcal{N}$ , do Steps 2 through 6.
2. For each sector of mesh node  $n$ , do Steps 3 through 5.
3. From the location of  $n$ , draw one bisecting ray through each sector and identify point  $x$ , where the boundary intersects the ray.
4. Perform heuristic boundary refinement. While per-sector measurement budget not exceeded, take one measurement as close as possible to the estimated boundary  $x$ . If measurement is in metric region, move boundary point  $x$  away from mesh node by constant distance, and vice-versa if measurement is outside of metric region. Stop if measurement is within tolerance (e.g.,  $\pm 3$  dB) of threshold value. Label the resulting boundary point on the ray as  $z$ .

5. Draw arc through  $z$  to identify the refined boundary estimate for the metric sector.
6. Merge all the metric sectors with revised boundaries to get the refined estimate of the metric region of  $n$ .

**Algorithm Properties.** By limiting the number of measurements per sector in the algorithm Refine-Estimate and limiting the number of sectors in the algorithm Estimate-Mono-Metric-Region, we ensure an upper bound on the total number of measurements taken, which is the product of the number of mesh nodes, the number of sectors per mesh node, and the maximum number of measurements per sector. Measurements are required only for 1) training measurements to estimate parameters  $\alpha$  and  $C_f$  and 2) boundary refinement measurements in step 4 of algorithm Refine-Estimate. There are two reasons for the actual number of measurements to be less than this bound: a) boundary refinement requires fewer measurements, and b) overlapping mesh node regions allow a measurement to be taken for multiple mesh nodes at one time.

### 3.4 Non-Monotonic Metrics

Our proposed estimation algorithm assumes a connected metric region and thus a single metric boundary, derived from the monotonicity property. We extend our framework to non-connected regions by considering only pairwise estimation instead of boundaries. The key aspect that allows us to estimate a disconnected region is a positive (additive) value for  $C_f$  terms in Equation (2). We later show in Section 4 that pairwise estimation is more accurate, but does not gain as much from additional measurements.

A simple ray-tracing method for improving our estimates with measurements involves localized refinement of the  $\beta$  term in the estimator Equation (1). Instead of assigning values of  $C_f$  based on global terrain features, we estimate the  $C_f$  terms for only those features within a sector. We take uniformly distributed measurements per sector in order to refine  $C_f$  values. This simple ray-tracing algorithm has the benefit of not assuming a connected region.

## 4. FRAMEWORK VALIDATION

We now validate our proposed framework using measurements from the GoogleWiFi and TFA networks with three performance metrics: coverage, modulation rate, and redundancy. We first introduce the network architecture and our measurement methodology. For each metric, we evaluate the accuracy and measurement overhead of our framework algorithms. For the coverage metric, we also discuss the sources of inaccuracy and the frequency of monotonicity violations. The remaining metrics build upon coverage regions and we evaluate how accurately they can be estimated using our framework.

### 4.1 Framework Validation Methodology

The measurement study consists of measurements from approximately 35,000 locations in the GoogleWiFi network and 29,000 locations in the TFA network.<sup>2</sup> We apply and validate our framework using only small (100s of measurements) subsets of the data, whereas we use all data to evaluate the accuracy of our framework. Each coverage measurement corresponds with a GPS location reading. For our

<sup>2</sup>All measurement data available for download from [tfa.rice.edu](http://tfa.rice.edu)

modulation rate metric, measurement pairs consist of the current SNR and modulation rate, and we extrapolate these measurements to a full set (all locations) of modulation rate measurements based on the probability distribution at each SNR value. The coverage threshold is set at  $\theta_c = 25$  dB as this value allowed download throughputs of approximately 1 Mbps in both networks.

#### 4.1.1 Removing Measurement Bias

Each coverage measurement point is a 4-tuple with  $x, y$  GPS coordinates, a node identifier, and an SNR value. We measure locations throughout all streets in the studied area, i.e., wardriving, as well as measuring in parking lots and driveways where possible. As the studied networks are not intended for pervasive indoor access, we focus on outdoor measurements.

We calculate the coverage of a network as the fraction of locations with measured signal strength above threshold  $\theta_c$ . Because the error of the GPS positioning is 3 meters, we consider a single location as covered if there exist any measurements above  $\theta_c$  within a 3 meter radius. We then define *predictive accuracy* as the fraction of locations that our framework correctly predicts as covered or not.

To account for an uneven spatial distribution in our measurements, we generate a set of 2-d sample points from a spatial Poisson process with intensity of 10,000 samples per  $\text{km}^2$ . Our evaluation then considers only sample points that are within 3 meters (the GPS accuracy) of at least one measurement. If multiple co-located measurements exist from a single mesh node, we consider the median of the measured SNR values. To determine the predictive accuracy of our estimation of a node’s metric region, we weight the distribution of measurement distances,  $\mathbf{d}(n, p)$ , to give less weight to progressively longer distances. The weights are based on the empirically measured distribution of distance from a client location to the nearest three mesh nodes.

The received signal power cannot be measured below a minimum receive power as the wireless card is not able to distinguish between weak transmissions and noise. In order to distinguish coverage holes from locations that the measurement study did not visit, we infer a coverage hole (uncovered location) with respect to a mesh node only if there is no measured value for that node but there is for another node. All measurements are obtained from the client, and hence all predictions pertain to the client performance.


#### 4.1.2 Alternate Techniques

We compare our framework against two alternate approaches: uniform propagation estimation and grid-based sampling. Uniform propagation restricts the shape of a metric region to a circle with radius determined only by the path loss exponent  $\alpha$ . In comparisons, we use the most accurate value of the radius for each circular region, assuming sufficient measurements have been taken. Second, grid-based sampling provides a fair comparison for a given number of measurements. For a given density, measurements are taken on a grid, and unknown points are estimated via interpolation.

### 4.2 Measurement Study Background

At the time of this study, the GoogleWiFi network consists of 447 Tropos mesh nodes mounted mostly on city light posts and covering a total outdoor area of  $31 \text{ km}^2$  in Mountain View, California. The Tropos nodes consist of a 7.4

dBi antenna and a single 802.11g wireless interface. Figure 2 shows the digital map providing terrain feature information for GoogleWiFi, which is a publicly available economic zoning map of the city of Mountain View. We measured a  $12 \text{ km}^2$  region, encompassing the northwest quadrant of the network, as shown in Figure 4. The client measurement platform for the GoogleWiFi study was a laptop with an external 802.11g wireless adapter, 3 dBi antenna, and GPS receiver.



**Figure 4: Coverage map in the GoogleWiFi network with circles representing covered locations and x’s representing holes in the  $12 \text{ km}^2$  measured area with 168 mesh nodes.**


The TFA Network is an urban mesh network, deployed in southeast Houston by Rice University [3]. At the time of the measurement study [6], the network consisted of 17 mesh nodes, providing coverage to a  $3 \text{ km}^2$  area. Each mesh node features a high gain 15 dBi omnidirectional antenna placed approximately 10 meters above the ground, higher than most of the houses and some of the trees in the neighborhood. The coverage region for the TFA network is shown in Figure 5. The TFA measurement platform was a laptop inside a car with an 802.11b wireless interface, a 7 dBi external antenna, and a GPS receiver.

The two networks have several key structural differences. The antennas used in TFA are taller and have higher gain, indicating a larger coverage region. This is offset partially by the difference in terrain, as the TFA network is filled with larger, denser trees which act as attenuators. Moreover, the GoogleWiFi nodes are mounted on light poles along streets whereas most TFA nodes are mounted against houses in the interior of a residential block.

### 4.3 Coverage Metric

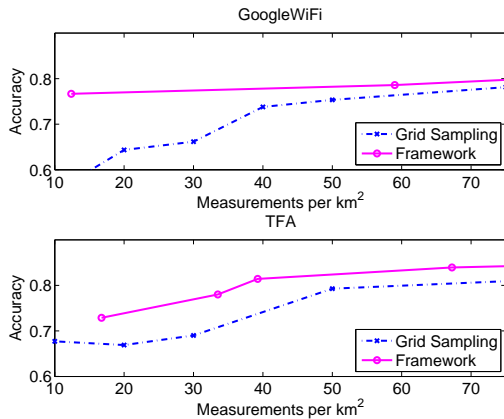
We focus on validating the predictive accuracy of our framework for the coverage metric region based on SNR measurements from both GoogleWiFi and TFA. We separately investigate our terrain-informed estimation technique and boundary refinement algorithm. Also, we evaluate the degree to which the coverage metric obeys the monotonicity property and the resulting impact on framework accuracy.

*For a given accuracy level, the framework reduces the required number of measurements by a factor of two to five as compared to a sampling and interpolation strategy.* Figure 6 presents the predictive accuracy of our framework and a



**Figure 5: Coverage region of the TFA network with circles representing covered locations and x's representing holes in 3 km<sup>2</sup> measured area, including the network edge where weak signal was measured.**

grid sampling strategy in both network scenarios as a function of the measurement budget per km<sup>2</sup>. These values do not reflect a one-time overhead of approximately 60 training measurements. Note that both approaches have a practical upper limit to accuracy, regardless of measurement budget, which we later show is due to monotonicity violations. Moreover, grid sampling accuracy does not exceed our framework's accuracy until approximately 10,000 measurements per km<sup>2</sup>, i.e., an exhaustive study.



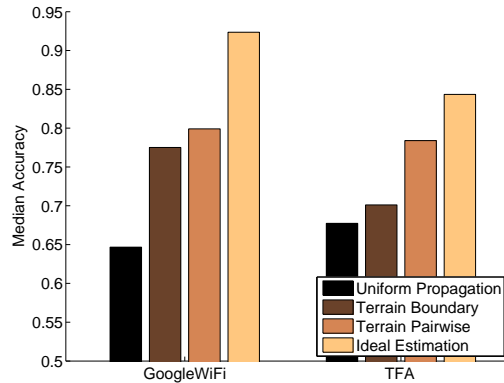
**Figure 6: Predictive accuracy of network coverage assessment with measurement budget, comparing our framework with grid sampling strategy.**

Figure 6 presents the total accuracy for covered and not covered locations. Though total accuracy improves, our accuracy for predicting *covered* locations is constant at 90% for all measurement budgets. As measurement budget increases, the ability of our framework to correctly predict coverage holes improves from 63% accuracy to 90% accuracy, leading to the improvements in Figure 6. Coverage holes are more difficult to correctly predict because *all* nearby mesh nodes' regions must be correctly predicted in order to correctly predict a hole.

### 4.3.1 Coverage Estimation Accuracy

We now focus on predicting a particular *node's* metric region, as opposed to the *network's* metric region by comparing four estimation strategies: 1) ideal estimation assuming all measurements are known, 2) per-location coverage estimation using Equation (1), 3) Algorithm Estimate-Mono-Metric-Region, and 4) the uniform propagation model. Note the difference between per-location estimation using only Equation (1) and estimation based on sectorization and boundary estimation is that the boundary-based estimation averages over a 20 degree sector before predicting locations.

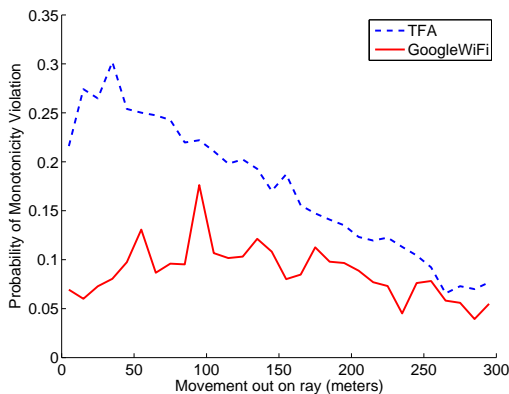
*Terrain-informed estimation halves the estimation errors compared to uniform propagation.* Figure 7 presents the median accuracy of the four described estimation strategies. Terrain-informed estimation eliminates more than half of the errors resulting from the uniform propagation model. Surprisingly, the sectorization and boundary averaging has little effect in the GoogleWiFi network, but significantly impacts the TFA network. Increasing the number of sectors from 10 per node to 36 per node eliminates this difference in the TFA network. These results indicate that the TFA metric boundary locations are correlated within approximately 10 degrees due to larger variations in tree foliage than in the GoogleWiFi terrain. Note that grid sampling requires approximately 100 measurements per km<sup>2</sup> in order for interpolation accuracy to be 80%. For assessment on a per-node basis (e.g., for network upgrades), this represents a relatively large cost increase as compared to estimation which requires no additional measurements.



**Figure 7: Predictive accuracy for the TFA and GoogleWiFi networks using uniform-propagation estimation, terrain-informed estimation, and ideal estimation.**

We now study the causes of errors in our estimation framework: monotonicity violations and the suboptimal choice of sector boundary. Monotonicity violations result in mispredictions even when the optimal sector boundary is chosen, which is a limitation of any framework that assumes a connected coverage region. The suboptimal boundary choice means that Equation (1) leads to a suboptimal choice of sector boundary location due to insufficient terrain feature information. For 31% of the measured sectors, the terrain-informed boundary estimation results in accuracy within 1% of the optimal boundary location, but for the remaining sectors, the terrain information is not sufficient for perfect estimation.

The probability of monotonicity violation is higher and has a stronger dependence on distance in the TFA network. Figure 8 depicts the probability that a measurement farther from the mesh node has a better signal strength than a nearer measurement within a sector width of 0.1 degrees. The difference between the two networks is in part due to fact that the measured TFA network is mostly residential blocks without line-of-sight, whereas the GoogleWiFi environment features more open space and line-of-sight along streets (where nodes are mounted on light posts). When considering optimal accuracy in 10 degree sectors, 20% of the measured sectors in the GoogleWiFi network feature zero monotonicity violations and perfect accuracy, whereas this is true for only 1% of TFA sectors. Overall, non-monotonicity contributes 10-15% average error, but with a large range 0-40% per sector. The average error is only slightly (3%) higher for the TFA network, despite the greater violation probability. The range in sector accuracy is significantly smaller for TFA though, indicating that the GoogleWiFi results in Figure 8 are averaged over a broader variety of propagation environments. Note that we factor out temporal fluctuations by disregarding an SNR increase of less than 3 dB, as measurements show 90% of co-located measurements vary less than 3 dB.



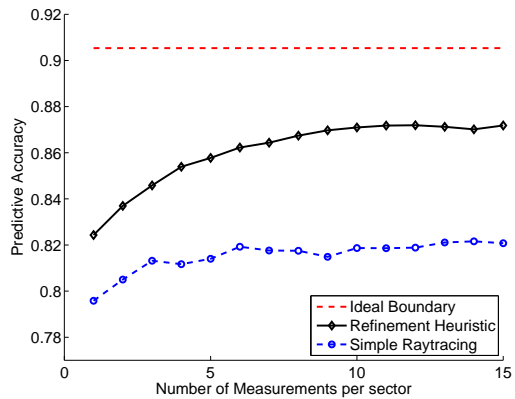
**Figure 8: Violations of the monotonicity property.** The probability that signal strength increases by more than 3 dB when increasing radius in a sector of width 0.1 degrees.

**Requiring metric regions to be connected incurs a minimal penalty in accuracy.** In particular, we found that adding a small number of disconnected metric regions has a minor benefit (1.5%), even with ideally chosen boundaries. That is, if we allow a metric sector to also have one disconnected section, we then need to find three boundaries. Our experiments show that even for boundaries that are found optimally, the mean accuracy increases from 89.8% to 91.4%. The remaining errors occur due to noisy metric boundaries.

#### 4.3.2 Refinement with Measurements

We now study the *refinement* phase of our evaluation framework using Algorithm Refine-Estimate. Our experiments consider 10 degrees wide sectors with more than 50 measurement points spread throughout the sector. The simple ray-tracing algorithm is also tested, with measurements taken within each sector to refine local values of  $C_f$  terms.

The boundary refinement heuristic outperforms simple ray-tracing for improving accuracy with measurements. Figure 9 depicts accuracy as a function of the measurement budget and indicates that the refinement stage improves average estimation accuracy from 82% to 88%. The ideal estimation line indicates the accuracy with all measurements known and the optimally chosen sector boundary. Localized refinement of  $C_f$  terms has small impact, indicating insufficient map granularity for simple ray-tracing techniques. The gain from the measurement refinement phase is less than the gain of exploiting terrain information for estimation, underscoring the importance of accurate  $\alpha$  and  $C_f$  parameters.



**Figure 9: Accuracy of boundary refinement algorithm for the coverage region of mesh nodes in GoogleWiFi.**

Our framework’s upper bound on the number of measurements is four times the actual needed number of measurements. There are two reasons Algorithm Refine-Estimate requires fewer measurements: boundary refinement stops early when a boundary is found and multiple mesh nodes are measured at a location with no extra cost. The first condition occurs because 30% of sectors require no more than three measurements to find the sector boundary. If the algorithm takes advantage of existing measurements within 15 meters of the estimated boundary instead of requiring a new measurement, the total required measurement budget is reduced by one-fourth. Note that we assume that taking an SNR measurement at point  $p$  from node  $n$  is approximately the same cost as taking an SNR measurement at  $p$  to all nodes  $n \in \mathcal{N}$ .

Terrain parameter  $C_f$  estimation requires a moderate (10-20) number of measurements per terrain feature type. Equation (1) requires an estimated value of the average path loss exponent  $\alpha$  and the weight,  $C_f$ , per terrain feature, to be determined as one-time overhead. We choose random measurements from our full set of measurements to estimate the average path loss  $\alpha$  and the  $C_f$  value of each terrain feature type. A moderately sized study of between 10 and 20 measurements achieves predictive accuracy within 2% of the best predictive accuracy. Note that there are six feature types in the terrain map used for GoogleWiFi.

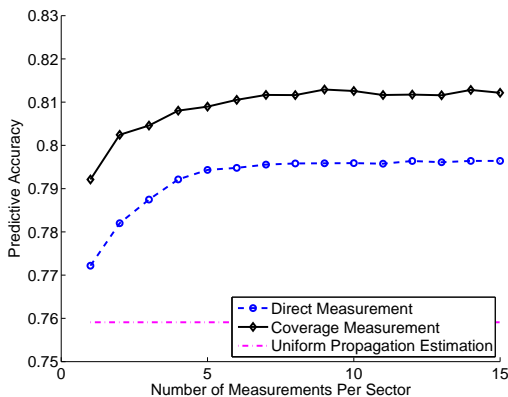
## 4.4 Modulation Rate Metric

The modulation rate metric captures the expected physical-layer modulation rate in use at a location, a value that depends on SNR and the rate selection protocol, e.g., ARF.

This section uses the GoogleWiFi network to compare the boundary refinement algorithm in two scenarios: 1) measurement of the modulation rate directly and 2) refinement using coverage measurements and then mapping coverage to modulation rate boundaries. The second approach involves empirically mapping the modulation rate threshold  $\theta_r$  to the coverage threshold  $\theta_c$  which corresponds to the desired modulation rate region. We next describe in more detail these two approaches.

We use our SNR measurement set for the first refinement approach, estimating the modulation rate using a piecewise linear function to map SNR value to expected modulation rate. We use this piecewise linear function to map the relationship for the specific network environment, though idealized values of this mapping can be obtained from the card manufacturer’s specifications. The constant  $C_r$  is a stepwise function with two cutoff SNR values,  $C_1$  and  $C_2$ . Below  $C_1$ , the expected rate is the minimum (1 Mbps), and above  $C_2$ , the expected rate is the maximum (54 Mbps). Between  $C_1$  and  $C_2$ , linear interpolation is used to find the expected rate. In our measurement study, we found  $C_1 = 14$  dB and  $C_2 = 32$  dB.

For the second refinement approach, the modulation rate measurements consist of an SNR value and the current physical-layer modulation rate, sampled after transmission of periodic ICMP packets. These measurements reflect the extent of cross-traffic and hidden terminal effects in the Google-WiFi network, which add noise to our linear mapping function. As expected, experiments show the modulation rate metric to be monotonic *on average* with respect to the coverage metric and therefore, to distance also.



**Figure 10: Comparison of boundary refinement strategies for modulation rate metric in Google-WiFi, either directly measuring modulation rate values or estimating from coverage measurements.**

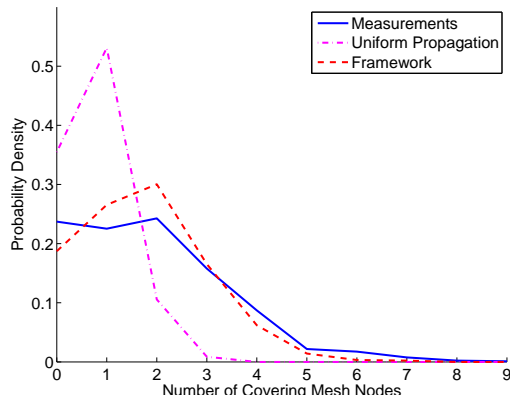
*Directly measuring modulation rate regions is less accurate than first using measurements to refine coverage estimates.* For metric boundary estimation, we compare two approaches: 1) directly measuring modulation rate for each sector and 2) measuring SNR and then estimating modulation rate sector boundaries from refined coverage boundaries. Figure 10 evaluates the improvement due to measurement refinement on the mean predictive accuracy of the modulation rate region of a metric sector of width 10 de-

grees. Direct measurement is worse because the extra variation in modulation rate at a given SNR value adds noise and makes it harder to find the refined boundary.

Note that MAC and network-layer metrics require associating with the mesh node in question and small data transmissions, both of which significantly increase the required measurement time and further motivate using coverage measurements to refine boundary estimates. The one-time overhead of experimentally characterizing the mapping between SNR and modulation rate is approximately 20 random measurements to achieve within 1% of the best estimation accuracy.

## 4.5 Coverage Redundancy Metric

*The estimation and refinement framework greatly increases accuracy at predicting 2-redundancy.* We estimate the metric redundancy region using the estimated coverage metric regions, as it requires only coverage information. Figure 11 compares the distribution of the  $k$ -redundancy metric in the GoogleWiFi network with the estimated values from our framework and the uniform propagation model. For predicting if a location is 2-redundant, this translates to a uniform propagation accuracy of 50% and an accuracy of 84% for our framework. The requirement to correctly predict multiple regions at each point accentuates the difference in accuracy between our framework and uniform propagation estimation. The relative accuracies in TFA (not shown) are approximately identical.



**Figure 11: Distribution of the number of covering mesh nodes for each location in the GoogleWiFi network, comparing measured values with estimations.**

## 5. DEPLOYMENT EVALUATION

Having validated the assessment framework, we next apply the framework and measurement study to investigate the two deployed networks, TFA and GoogleWiFi. Specifically, we evaluate the effect of deployment density on coverage holes and study client association policies to understand the load-balancing qualities of deployments.

### 5.1 Coverage Holes and Deployment Efficiency

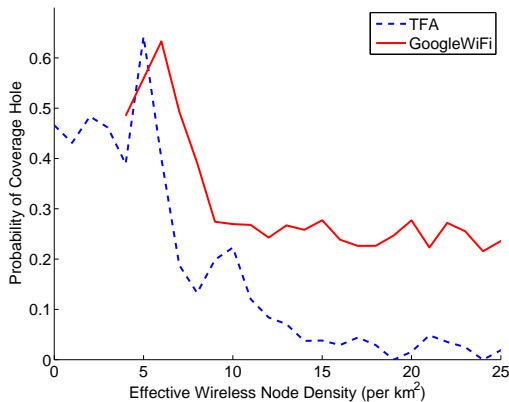
Here, we evaluate the efficiency of the two deployments in terms of the density of deployed nodes and their chosen locations. We first consider the sizes of the coverage holes in

both networks, and then focus on the effect of deployment density on the likelihood of coverage holes.

*Half of measured locations without coverage are small holes, within 10 meters from a covered location.* We consider a measured location to be a coverage hole if it is not within 3 meters of a location in the coverage region of any mesh node. The *size* of a coverage hole is the distance to the nearest covered location. Contrasting estimation techniques, grid sampling with 30 measurements per  $\text{km}^2$  predicts *three times* more coverage holes than exist, whereas our framework overpredicts by only 25%. This means that grid sampling approach would conclude that coverage holes are  $3\times$  more common and two to four times as large on average.

We next examine the distribution of deployment densities in each network with the goal of understanding how to best deploy a network in terms of the coverage metric region. We compare two deployment strategies: minimizing the maximum distance to the nearest mesh node and the looser restriction of deploying at a specific localized node density. We reverse engineer both the TFA and GoogleWiFi networks to determine how to improve their deployment strategy.

We calculate the *deployment density* at a client location as local node density per  $\text{km}^2$  within a circle around the client point with radius of 400 meters. Note that this is a client-centric definition of density, which allows us to evaluate locations of different density in a non-uniformly deployed mesh network. We focus on this specific localized region size because our data shows that less than 2% of client locations are covered only by a node farther than 400 meters away. Because of the different technology, antenna gains, heights, and propagation environment, the two networks are not expected to require the same node density. The mean density in TFA is 11.2 nodes per  $\text{km}^2$ , with standard deviation of 7.4. The mean density in GoogleWiFi is 17 nodes per  $\text{km}^2$ , with standard deviation of 5.9.



**Figure 12: Probability of a coverage hole as a function of local deployment density for both networks.**

*Deploying to minimize the distance from a client to the nearest mesh node leads to a  $3\times$  over-deployment in TFA.* In Figure 12, we plot the probability of a coverage hole as a function of deployment density at a client point. If we seek a deployment with 90% coverage, we must keep the maximum distance to a mesh node below 130 meters. This corresponds to a deployed density in a square grid of approximately 30 nodes per  $\text{km}^2$ . However, Figure 12 indicates that only 11

nodes per  $\text{km}^2$  are required to attain 90% coverage. This difference is due to the fact that coverage is often provided by the mesh node that is *not* closest to the client. For a deployment in progress, our assessment framework can be used to determine the best density, per the above analysis.

In the GoogleWiFi network, we find the deployment density has little impact on coverage hole probability above 8 nodes per  $\text{km}^2$ . Considering the nearest node deployment strategy, we find that in order to limit the probability of a coverage hole to less than 10% in GoogleWiFi, the distance to nearest node must be less than 80 meters, which translates to a deployment density of approximately 77 mesh nodes per  $\text{km}^2$ . This density is dramatically higher than the actual deployed density of 17 nodes per  $\text{km}^2$ , although the actual probability of a coverage hole is approximately 25%. In order to achieve 90% coverage, client-side solutions, such as use of higher-gain antennas, may be more cost effective than the very dense deployment.

*Coverage holes in GoogleWiFi are correlated between mesh nodes, leading to significantly higher ( $4\times$ ) node density requirements to decrease holes from 25% to 10%.* Surprisingly, the TFA network’s coverage-hole probability quickly approaches zero as density increases, whereas the GoogleWiFi network does not, indicating that further increasing node density has diminishing impact on coverage hole probability in GoogleWiFi. For all client locations, we examine each mesh node within 400 meters and the boundary distance for the sector that the location belongs to. For the coverage holes, we found that 75% of these sectors had boundary distance of at least 50 meters less than the average boundary distance of 178 meters. Therefore, the GoogleWiFi environments presents a significantly greater deployment challenge. *In general, improving the coverage of deployed mesh networks is challenging because the coverage holes are small and spread out, requiring many new node locations to remove them.*

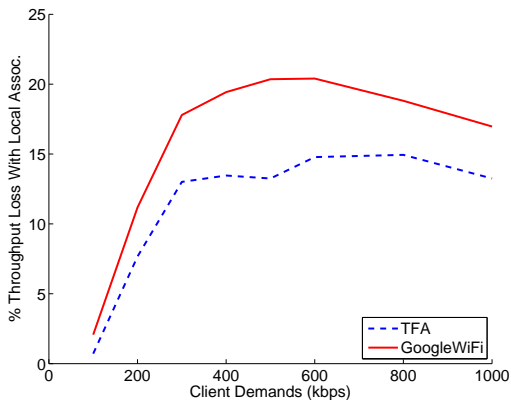
## 5.2 Load-Balanced Node Deployment

We next investigate the load-balancing qualities of a deployment through investigating association policies for clients in range of multiple mesh nodes. The objective is to show that given standard client association policies, there is significant imbalance in the number of client locations each mesh nodes serves. We compare association based on the client’s strongest signal strength with an ideal policy that jointly considers signal strength and load balancing.

To study this issue, we first define a method for calculating the load on each mesh node under a hypothetical population-based model. In particular, we consider the offered load of a location  $i$  to be a function of the client demand  $q_i$  (in kilobits) at  $i$  and the expected time required to serve a fixed sized packet,  $d_i$ , (measured in seconds per kilobit) at  $i$ . Our framework provides the estimate for  $d_i$  based on the predicted modulation rate at location  $i$ , which is in turn based on the SNR at the location. Let  $Q_n$  represent the set of clients associated with mesh node  $n$ , then the load of node  $n$  is  $L(n, Q_n) = \sum_{i \in Q_n} d_i q_i$ . As load is measured in time units, if  $L > 1$ , then the mesh node is saturated. In this case, the clients’ served load will be the offered load  $q_i$  divided by  $L$ . Otherwise, the full client demand,  $q_i$ , is served. We further assume that the access tier for each mesh node operates on an independent frequency channel, and hence there is no interference between nodes.

We now describe two client association policies. Most existing clients employ a policy in which they compare received SNR values of all APs within range and associate to the node with the strongest SNR. For a load-balanced policy, we formulate and solve the association problem as a maximum flow problem on bipartite graphs [17] with one set of nodes as the client locations, the other set as mesh node locations, and an edge in the graph when a client node is in the coverage region of a mesh node. The supply is  $d_i \times q_i$  for every client node  $i$ , the capacity of each edge emanating from  $i$  is also  $d_i \times q_i$ , and the demand is 1 for every mesh node.

For our experiment setup, we vary the client demand from 100 kbps to 1 Mbps, and use the mapping from SNR to modulation rate presented in Section 4. The term  $d_i$  is calculated as twice the inverse of the modulation rate to also account for overhead.



**Figure 13: Percentage gain in average client throughput from using a centralized, load-balancing association scheme versus realistic, local association with strongest SNR.**

*Max-flow load-balanced association improves total client throughput by 20% in GoogleWiFi as compared to strongest-SNR association.* Figure 13 shows 15-20% gain in average client throughput for the centralized, load-balancing policy versus local association. In a regular grid deployment, the local policy results in the same association as the load-balanced policy due to the fact that the same number and quality (SNR) of clients associate with each node. The loss in throughput of the strongest SNR policy is then also the loss due to the uneven deployment of mesh nodes, resulting in some overloaded nodes and some under-loaded nodes. Note that the most gain available is at moderate offered loads, where there is significant load imbalance on mesh nodes and some nodes have available capacity. The gain in the TFA network is less because a larger fraction of the client nodes have only one possible node to associate with.

## 6. RELATED WORK

**Measurement Strategies.** While many previous studies present measurements of the coverage of a wireless network [1, 3, 6, 11, 20], none have proposed a framework for choosing the number and location of measurements to characterize a metric region via a small number of measurements. Ray-tracing performs detailed simulation and prediction of physical-layer propagation in order to estimate

physical-layer propagation [11, 19, 20]. However, ray-tracing requires highly detailed information about the environment, such as building materials and thickness, to achieve high accuracy in outdoor environments. Other studies [1, 3] have used a small set of coverage measurements to estimate parameters such as path loss and shadowing, but assume a uniform propagation (circular) model.

**Physical Layer Models.** The Okumura model [16] and related models [5] for outdoor propagation are widely-used empirical models to predict signal propagation in urban environments. However, they apply to different carrier frequencies, coarse-grained terrain features, and clients located more than 1 km from the base station. Our terrain-informed estimation builds upon indoor propagation modeling techniques that use attenuation factors derived from building blueprints to show wall locations, thickness, and material [14]. A good review of such propagation modeling techniques is in [8]. In contrast, our terrain-informed estimation is a sectorized technique applied to outdoor environments via the “push-pull” boundary refinement technique. Moreover, we operate with finer-grained resolution of terrain features and utilize training measurements to obtain localized estimates of attenuation.

**Cellular Coverage.** Cellular networks feature different frequency bands, antenna heights, and propagation environments, but still encounter related coverage assessment problems. Previous studies have addressed how to characterize coverage in a cellular network with a small number of signal strength measurements [15] and in the presence of random shadowing (violations of coverage monotonicity) [2, 18]. However, prior work uses uniform propagation models and does not provide strategies for selecting measurement locations. Likewise, “diamond-shaped” coverage regions were studied for Manhattan cellular networks [7], with exhaustive measurements used for a downtown-only environment.

**WLAN and Mesh Measurement Studies.** For indoor WLAN deployments, exhaustive measurements have been compared to propagation modeling tools [21]. Other measurement studies in mesh and WLAN networks focused on client usage and mobility [6, 12] or protocol performance. Finally, a related problem is node placement, i.e., choosing locations to deploy wireless APs: this problem has been formulated as an optimization problem requiring as input exhaustive measurements [10]. In contrast, our framework can provide input to a placement algorithm with substantially smaller measurement overhead.

## 7. CONCLUSION

In this paper, we present a measurement framework to accurately characterize the metric regions of a mesh node using only a small number of measurements. The primary metric we consider is coverage (signal strength), and the technique of estimating and refining metric regions can be extended to other monotonic metrics. We also use our framework to study modulation rate and redundancy metrics.

We utilize publicly available terrain maps to improve coverage estimation, showing that coarse-grained maps significantly improve estimation accuracy. This improved estimation leads to fewer needed measurements to refine the estimated region boundaries. We further examine the sources of estimation error and find that coverage monotonicity violations account for an average of 10% error, although with much greater variation per sector in GoogleWiFi than in

TFA. With our framework, we then reverse-engineer the TFA and GoogleWiFi deployments and find that improving the performance of GoogleWiFi through the addition of nodes would be highly problematic (costly) due to the prevalence of numerous coverage holes having small area.

## Acknowledgments

We would like to thank Mustafa Uysal for his suggestions and feedback, Anastasios Giannoulis for providing TFA measurement data, and Karl Garcia at Google for generously providing information about the GoogleWiFi network.

## 8. REFERENCES

- [1] A. Arjona and S. Takala. The Google muni WiFi network—can it compete with cellular voice? In *The Third Advanced International Conference on Telecommunications*, May 2007.
- [2] P. Bernardin, M. F. Yee, and T. Ellis. Cell radius inaccuracy: a new measure of coverage reliability. *IEEE Transactions on Vehicular Technology*, 47(4):1215–1226, November 1998.
- [3] J. Camp, J. Robinson, C. Steger, and E. Knightly. Measurement driven deployment of a two-tier urban mesh access network. In *Proceedings of ACM MobiSys*, Uppsala, Sweden, June 2006.
- [4] S. Cherry. Wi-Fi nodes to talk amongst themselves. *IEEE Spectrum Magazine*, July 2006.
- [5] V. Erceg, L.J. Greenstein, S. Tjandra, S.R. Parkoff, A. Gupta, B. Kulic, A. Julius, and R. Jastrzab. An empirically-based path loss model for wireless channels in suburban environments. *Proceedings of IEEE GLOBECOM*, November 1998.
- [6] A. Giannoulis, M. Fiore, and E. Knightly. Supporting vehicular mobility in urban multi-hop wireless networks. In *Proceedings of ACM MobiSys*, Breckenridge, CO, June 2008.
- [7] A. Goldsmith and L. Greenstein. A measurement-based model for predicting coverage areas of urban microcells. *IEEE Journal on Selected Areas in Communications*, 11(7):1013–1023, September 1993.
- [8] M. Iskander and Z. Yun. Propagation prediction models for wireless communication systems. *IEEE Transactions on Microwave Theory and Techniques*, 50(3):662–673, Mar 2002.
- [9] X. Jin and C. H. Davis. Automated building extraction from high-resolution satellite imagery in urban areas using structural, contextual, and spectral information. *EURASIP Journal on Applied Signal Processing*, 2005(1):2196–2206, 2005.
- [10] M. Kamenetsky and M. Unbehaun. Coverage planning for outdoor wireless LAN systems. In *International Zurich Seminar on Broadband Communications*, Zurich, Switz., February 2002.
- [11] S.-C. Kim, B. J. Guarino Jr., T. M. Willis III, V. Erceg, S. J. Fortune, R. A. Valenzuela, L. W. Thomas, J. Ling, and J. D. Moore. Radio propagation measurements and prediction using three-dimensional ray tracing in urban environments at 908 MHz and 1.9 GHz. *IEEE Transactions on Vehicular Technology*, 48(3):931–946, May 1999.
- [12] D. Kotz and K. Essien. Analysis of a campus-wide wireless network. *Wireless Networks*, 11(1-2):115–133, January 2005.
- [13] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. In *Proceedings of the ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, October 2004.
- [14] M. Lott and I. Forkel. A multi-wall-and-floor model for indoor radio propagation. *Proceedings of IEEE Vehicular Technology Conference*, 1:464–468, 2001.
- [15] K. Manoj, P. Bernardin, and L. Tamil. Coverage prediction for cellular networks from limited signal strength measurements. *The Ninth IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, 3:1147–1151, September 1998.
- [16] T. Okumura, E. Ohmori, and K. Fukuda. Field strength and its variability in VHF and UHF land mobile service. *Review Electrical Communication Laboratory*, 16(9-10):825–873, 1968.
- [17] C. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Courier Dover Publications, 1998.
- [18] R. Pattuelli and V. Zingarelli. Precision of the estimation of area coverage by planning tools in cellular systems. *IEEE Personal Communications*, 7(3):50–53, Jun 2000.
- [19] T.S. Rappaport. *Wireless Communications, Principles & Practice*. Emerging Technologies Series. Prentice Hall, Upper Saddle River, New Jersey, 1996.
- [20] T. Rautiainen, G. Wolffe, and R. Hoppe. Verifying path loss and delay spread predictions of a 3d ray tracing propagation model in urban environment. *Proceedings of IEEE Vehicular Technology Conference*, 4:2470–2474, 2002.
- [21] S. Zvanovec, P. Pechac, and M. Klepal. Wireless LAN networks design: Site survey or propagation modeling? *Radio Engineering Prague*, 12(4):42–49, 2003.