

Practical Conflict Graphs for Dynamic Spectrum Distribution

Xia Zhou, Zengbin Zhang, Gang Wang, Xiaoxiao Yu[§], Ben Y. Zhao and Haitao Zheng
Department of Computer Science, U. C. Santa Barbara, USA
[§]Tsinghua University, Beijing, P. R. China
{xiazhou, zengbin, gangw, ravenben, htzheng}@cs.ucsb.edu, yuxiaox@gmail.com

ABSTRACT

Most spectrum distribution proposals today develop their allocation algorithms that use conflict graphs to capture interference relationships. The use of conflict graphs, however, is often questioned by the wireless community because of two issues. First, building conflict graphs requires significant overhead and hence generally does not scale to outdoor networks, and second, the resulting conflict graphs do not capture accumulative interference.

In this paper, we use large-scale measurement data as ground truth to understand just how severe these issues are in practice, and whether they can be overcome. We build “practical” conflict graphs using *measurement-calibrated propagation models*, which remove the need for exhaustive signal measurements by interpolating signal strengths using calibrated models. These propagation models are imperfect, and we study the impact of their errors by tracing the impact on multiple steps in the process, from calibrating propagation models to predicting signal strength and building conflict graphs. At each step, we analyze the introduction, propagation and final impact of errors, by comparing each intermediate result to its ground truth counterpart generated from measurements. Our work produces several findings. Calibrated propagation models generate *location-dependent* prediction errors, ultimately producing conservative conflict graphs. While these “estimated conflict graphs” lose some spectrum utilization, their conservative nature improves reliability by reducing the impact of accumulative interference. Finally, we propose a graph augmentation technique that addresses any remaining accumulative interference, the last missing piece in a practical spectrum distribution system using measurement-calibrated conflict graphs.

Categories and Subject Descriptors

C.2.3 [Network Operations]: Network Management

Keywords

Dynamic spectrum access; conflict graphs; interference

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.

SIGMETRICS'13, June 17-21, 2013, Pittsburgh, PA, USA.
Copyright 2013 ACM 978-1-4503-1900-3/13/06 ...\$15.00.

1. INTRODUCTION

Current reforms in radio spectrum management promise to spur rapid growth of wireless technologies, by using on-demand spectrum auctions and secondary markets [3, 50]. In an ideal scenario, these markets not only maximize profit for spectrum owners, but also allows spectrum users, *e.g.* small cell providers, to purchase spectrum on-the-fly and receive exclusive usage of allocated spectrum without the hassle of sharing.

Achieving this goal requires two tightly-coupled components: *an accurate model of interference patterns* among current spectrum users, and *an allocation algorithm* that uses this interference model to distribute spectrum efficiently. For spectrum, this means maximizing utilization by parallelizing non-interfering transmissions whenever possible.

The majority of prior works have chosen to develop allocation algorithms under an abstract interference model called “conflict graphs” [20]. As the name suggests, a conflict graph is a simple graphical representation of the interference condition between any two spectrum users¹. This simple interference structure greatly simplifies spectrum allocation design, leading to a series of highly efficient allocation algorithms with bounded performance and polynomial complexity [10, 13, 22, 32, 38, 43, 50]. In contrast, alternative physical interference models are highly complex, and have been shown in existing studies to produce unbounded performance loss when used to build allocation algorithms [9, 47].

As popular as they are, the practical value of conflict graphs is often questioned by the wireless community for two key reasons. *First*, building an accurate conflict graph for a specific physical area is very challenging. Given the complex nature of RF propagation, it requires detailed measurements covering all combinations of sender/receiver locations. This type of per-link signal measurement is feasible for indoor WLANs [5, 6, 26, 33, 39, 43, 45], but impractical for the outdoor networks targeted by spectrum markets. As a substitute, most current proposals build artificial conflict graphs using a simple distance-based criterion [7, 10, 48] or from signal strength values generated from simple RF propagation models with rule of thumb parameters [18, 32]. While these simplifications ease the process of designing and evaluating allocation algorithms, empirical studies have shown that they produce incorrect interference results that lead to poor performance [27, 34].

¹For a specific frequency band, if two users can operate concurrently without visible performance degradation, then they do not conflict. Otherwise they conflict and are connected with an edge (of the specific band) [35].




Figure 1: Our high level methodology. We build *estimated conflict graphs* by collecting limited signal measurements at a small number of randomly deployed sensors to calibrate a propagation model, and using the results to predict signal strength values and construct the conflict graph. We examine the accuracy of estimated conflict graphs by comparing them against *measured conflict graphs* built from exhaustive signal measurements, using graph similarity and spectrum allocation benchmarks.

Second, because conflict graphs only define interference conditions between any two spectrum users, they cannot capture the impact of interference accumulated from multiple concurrent transmitters in the same frequency band. Prior work [30] has shown that such “mismatch” leads to unpredicted and harmful interference at allocated users, breaking the exclusive usage guarantee offered by the spectrum market. Without guarantees that their transmissions would operate without interference, users would have little incentive to purchase from the spectrum market.

In this paper, we use a data-driven approach to gain a better understanding of the severity of these two issues. We use measurements as ground truth to quantify the severity of errors produced by building conflict graphs without exhaustive signal measurements, and to determine if these errors impact users in the form of poor spectrum allocations. We also seek to identify solutions to minimize these errors, and in doing so, addressing the community’s main concerns and promoting the continued use of conflict graphs in practice.

In our study, we build conflict graphs using *measurement-calibrated propagation models*. Instead of performing exhaustive measurements, this approach performs measurements on only a subset of locations. These results are used to calibrate a propagation model, which is used to make signal strength predictions for all locations in the area. These predictions are used in lieu of exhaustive measurements to build the conflict graph. This approach has two advantages. First, prior works have shown that measurement-calibrated propagation models are much more accurate than those built with rule-of-thumb parameters [19, 29, 31, 42]. Second, because measurements can be performed by sensors or even trusted network subscribers, this approach incurs low overhead, and can offer continuous measurements in real-time. This allows conflict graphs to adapt to constantly changing network environments and users. We recognize, however, that calibrated propagation models are imperfect, and will introduce errors in the predicted signal strength maps [14, 23, 36, 40]. So we must understand whether these errors carry through to become errors in conflict graphs, and if they impact the efficacy of spectrum allocations for users.

Our high level methodology is as follows:

- Use a relatively small number of signal measurements to calibrate RF propagation models;
- Use models to build predicted signal strength maps, and use those to produce “estimated conflict graphs”;
- Compare estimated conflict graph to “measured conflict graph” built from exhaustive signal measurements, in the form of missing or extraneous edges between the two;

- Evaluate end-to-end impact by running spectrum allocation on both conflict graphs and comparing them while considering the impact of accumulative interference.

To the best of our knowledge, our work is the first empirical study on the practical usability of conflict graph for dynamic spectrum distribution. Our work differs from existing works on constructing conflict graphs. First, focusing on outdoor environments, our work differs from prior work [5, 6, 26, 33, 39, 43, 45] that build indoor conflict graphs using exhaustive signal measurements. Second, our work targets dynamic spectrum markets where users requesting spectrum are located at unplanned places and the resulting conflict graph can be of arbitrary shape. This is fundamentally different from cellular networks [19, 29, 46] which optimize the placement (and transmit power) of base stations to produce conflict graphs of specific shapes.

Our measurement study leads to four key findings:

- Calibrated propagation models generate location-dependent signal prediction errors. They are more likely to underpredict signal strength at short distances, and overpredict them for long distance links. We consistently observe this pattern across multiple measurement datasets.
- These prediction errors lead to conservative conflict graphs that rarely miss actual conflict edges, but commonly introduce extraneous conflict edges.
- This leads to conservative spectrum allocations with utilization loss compared to measured conflict graphs. These extra edges, on the other hand, play a critical role in reducing the impact of accumulative interference, thus achieving more reliable links than allocations using only measured conflict graphs.
- A simple graph augmentation technique can effectively eliminate the artifact of accumulative interference from both conflict graphs, boosting the reliability of spectrum allocation to more than 96%. Once augmented, estimated conflict graphs also achieve utilization that is more than 85% of the ideal allocation.

2. METHODOLOGY

Using real data, we seek to understand key issues when using conflict graphs for dynamic spectrum distribution. We consider conflict graphs built from measurement-calibrated propagation models because they are practical, requiring little measurement overhead, and much more accurate than those built with rule-of-thumb parameters.

Our approach, shown in Figure 1, consists of four steps: 1) collecting real signal maps via measurements, and us-

ing them as ground truth; 2) using sampled subsets to calibrate propagation models, and predicting network-wide signal maps; 3) building conflict graphs from both measured and predicted signal maps, and 4) quantifying the accuracy of estimated conflict graphs using measured graphs as ground truth, via both graph similarity and spectrum allocation benchmarks. By examining both efficiency and reliability of the allocation, we examine the impact of accumulative interference. Next, we briefly describe our assumptions and present each step in detail.

Assumptions. The basis of our study comes from wardriving measurements of outdoor municipal WiFi networks. We assume that each spectrum user seeks to obtain one of the WiFi channels each with the same propagation properties. We use WiFi band as an example of distributing spectrum among outdoor networks, and also because this is the only outdoor network with known base station locations. Our work can easily be extended to other frequency bands by adjusting the propagation model to account for carrier frequency differences [12, 41]. We leave this to a future work.

2.1 Collecting Signal Maps

Our study uses wardriving measurements of outdoor municipal WiFi networks at three different cities, one of which was collected by our own group. Each dataset consists of beacon RSS values of WiFi access points (AP) measured in a large outdoor area of size $3\text{-}7\text{km}^2$, along with the location of each measurement and the locations of all APs. We average multiple RSS readings per location to derive a map of average signal strengths for each participating AP. Table 1 summarizes the datasets.

GoogleWiFi. Collected by our research group in April 2010, this dataset covers a 7km^2 residential area of the Google WiFi network in Mountain View, California. Figure 2 shows the measurement locations (as blue dots) and the APs (as red triangles). We used three co-located laptops equipped with customized WiFi cards² with higher receive sensitivity than normal cards. Thus this dataset records detailed signal strength values of 78 APs at 11,447 distinct locations (with an average 5m separation between nearby locations). More importantly, each location has signal strength values of 6+ APs in average, 2-3 times more than the other two datasets.

MetroFi. This dataset [2] consists of RSS values in a 7km^2 area of an 802.11x municipal network in Portland, Oregon. It was collected by a research group from University of Colorado in 2007. The dataset covers 30,991 distinct measured locations of 70 APs with known GPS locations. The average number of APs heard per location is only 2.3.

TFA. Collected by researchers from Rice University, this measurement data covers 22 APs in a 3km^2 area of the TFA network in Houston, Texas [4]. It includes measurements from 27,855 locations.

To use these datasets in our study, we treat each AP as the transmitter of a spectrum market user, and any measured location in its coverage area as the receiver positions of the market user. While our measurements are on WiFi

²We use WiFi cards from Wifly-City System Inc. Equipped with a 7dBi external omni antenna and a dual amplifier, they double the sensing range of standard WiFi cards. Following FCC rules, we only use the RX path of the card to receive beacons, with its TX path always turned off.

Table 1: Summary of the datasets used in our study

Dataset	Area size (km^2)	# of APs w/ GPS info	# of measured locations	Avg. # of APs heard per location
GoogleWiFi	7	78	11,447	6.2
MetroFi	7	70	30,991	2.3
TFA	3	22	27,855	2.7




Figure 2: Measured area in the GoogleWiFi dataset. Red triangles are the APs detected and blue dots are measured locations on the streets.

networks, both the measured signal maps and the resulting conflict graphs are independent of specific MAC protocols used. This is important, since it matches the exclusive usage scenario, where a spectrum market user is free to use any MAC protocol in its authorized spectrum range.

2.2 Calibrating Propagation Models

To generate “estimated conflict graphs,” we use samples of our measurements to calibrate existing propagation models. We select several well-known models designed specifically for urban street environments that match our datasets, including the simple uniform path loss model, and complex models that support specific environmental features like streets and building structures. We now describe our high-level approach to model calibration and signal map prediction. We leave the detailed discussion on each model and their calibration procedure to Section 3.

We begin by choosing sub-samples from the exhaustive measurement data. Since the search for optimal sampling methods is still an open problem [49], we randomly sample our data, and vary the density of the sample data between 1.4 and 100 samples per km^2 . We then use the Minimum Mean Squared Error (MMSE) fitting method to determine the best-fit parameters for each propagation model. Once the parameters have been calibrated for a given model, we then interpolate the signal values at other locations to build the complete signal strength map.

2.3 Constructing Conflict Graphs

We now have two signal strength maps, one from our exhaustive signal strength measurement data, and one interpolated from our calibrated signal propagation model. We use them respectively to build a *measured conflict graph*, *i.e.* ground truth, and an *estimated conflict graph*. These conflict graphs represent the interference patterns of spectrum market users, where each spectrum market user maps to a

stationary transmitter, *i.e.* an AP in our signal maps, and its coverage region corresponds to locations for its receivers.

The resulting conflict graph consists of a set of nodes, each mapping to a spectrum market user, and a set of edges, each representing a conflict between two nodes. To determine if two users conflict, we place their transmitters on the same spectrum channel and examine whether they both receive “exclusive spectrum usage.” A market user receives exclusive spectrum usage if γ -percentile of its qualified transmissions have signal to noise and interference ratio (SINR) above β [22]. Along with coverage area and transmit power, γ and β are operating parameters configured by spectrum market users in their spectrum purchase requests.

Consider two nodes i and j . Let $\text{SINR}_u^{i,j}$ represent the SINR value at location u in node i 's coverage area: $\text{SINR}_u^{i,j} = \frac{S_u^i}{I_u^j + N_0}$, $u \in U_i$, where S_u^i is the received signal strength at u from i 's transmitter, I_u^j is the interference strength from j 's transmitter, N_0 is the thermal noise, and U_i is the coverage area of i . We sort locations within each node's coverage area by their SINR, and determine conflict conditions using the bottom $(1 - \gamma)$ -percentile value. That is, node i and j conflict if and only if for either of the two coverage areas, the percentage of locations with $\text{SINR} \geq \beta$ is less than γ :

$$p_{ij} = \min(q_j^i, q_i^j) < \gamma, \quad (1)$$

where $q_i^j = \frac{|\{u|u \in U_i, \text{SINR}_u^{i,j} \geq \beta\}|}{|U_i|}$. Here $(1 - \gamma)$ represents the percentage of coverage holes a spectrum user is willing to tolerate to maximize capacity [1]. When $\gamma = 1$, eq. (1) reduces to the minimal SINR-based criterion [9, 47, 48].

Configuring Coverage Area and β . For simplicity, we assume each market user's coverage area includes all measurement locations whose SNR $\geq \beta$. If a single location falls into the coverage area of multiple users, we assume that it is associated with the user that maximizes its signal strength. We set $\beta=10\text{dB}$, which is the minimum SNR required to decode beacons in GoogleWiFi measurements. This allows us to use all measurement locations in our graph analysis. We have also experimented other β values (8–20dB). Since they lead to the same trend, we omit the results for brevity.

2.4 Evaluating Graph Accuracy

Finally, we examine the accuracy of the estimated conflict graphs and the artifact of accumulative interference not captured by these graphs. To do so, we compare the estimated (measurement-calibrated) conflict graph against the measured conflict graph built directly from measurements. Our analysis uses both graph similarity metrics and spectrum allocation benchmarks.

For graph similarity, we perform edge-based comparison of the two conflict maps, using the measured conflict graph as ground truth. This produces a set of “extraneous edges” and “missing edges” that capture the differences between the estimated graph and the measured graph. We analyze the patterns of extraneous and missing edges, and explain their appearance based on errors in signal map prediction.

To understand the impact of graph edge errors on spectrum users, we feed each type of conflict graphs to two well-known spectrum allocation benchmarks and compare the allocation results. These end-to-end tests provide answers to two questions: will the edge errors lead to significant loss in spectrum efficiency and reliability, and will the “uncaptured” accumulative interference also lead to significant loss?

In the following, we present detailed results of each of our analysis steps. We begin by examining the accuracy of signal map prediction using calibrated propagation models (Section 3). Then we build and compare measured and estimated conflict graphs in terms of graph similarity (Section 4) and spectrum allocation performance (Section 5).

3. SIGNAL PREDICTION ACCURACY

To examine the concern on the accuracy of measurement-calibrated conflict graphs, we begin with understanding the types of errors introduced when we use incomplete measurements to calibrate propagation models and predict signal strength values. More specifically, are there patterns in prediction errors likely to manifest later as errors in conflict graphs. Our goal is to answer the question: *how accurate are signal strength predictions made by measurement-calibrated propagation models, and does receiver location play a role in prediction accuracy?* We take several representative propagation models, calibrate them using controlled samples of our measurement data, and evaluate their signal strength predictions for locations missing from the sample, using the full dataset as ground truth.

3.1 Propagation Models and Calibration

We choose four representative propagation models for our work, because they capture urban street environments that best match our datasets. They range from the simplest uniform pathloss model to sophisticated models that incorporate features like streets and changes in terrain.

Uniform Pathloss Model (Uniform) [14]. The simplest and most-used model, this captures signal attenuation over distance using a single pathloss exponent. Calibration is straightforward: use Minimum Mean Square Error (MMSE) to determine the best-fit pathloss exponent.

Two-Ray Model (Two-Ray) [14]. This model uses two pathloss exponents to capture the dual slope feature of signal propagation in urban environments, *i.e.* signal attenuates faster after a certain distance. It offers higher accuracy than the uniform pathloss model in urban street environments [16]. To calibrate this model, we partition the sample measurements into two sets using a distance threshold, and for each set we use a separate MMSE fitting to determine the best-fit pathloss exponent. We also optimize the partition to minimize the overall MMSE.

Terrain-based Model (Terrain) [42]. This model leverages terrain information to capture the non-uniformity of radio propagation caused by different terrains. It divides the transmitter's coverage area into sectors, and applies a terrain-specific shadowing factor in each sector. We follow the procedure in [42] to calibrate this model. Since we do not have terrain information (street, buildings, *etc.*) for the MetroFi and TFA datasets, we only provide results for our GoogleWiFi dataset.

Street Model (Street) [15]. This model targets urban microcell networks, and assumes that signals are constrained to propagate along the streets along line-of-sight, with minor reflection and/or diffraction to cross streets. To calibrate this model, we categorize signal propagation into three types based on the number of reflections it encounters. These include those without any reflection, *i.e.* line-of-sight, those with one reflection, and those with multiple reflections. We divide measurement samples into these categories and train

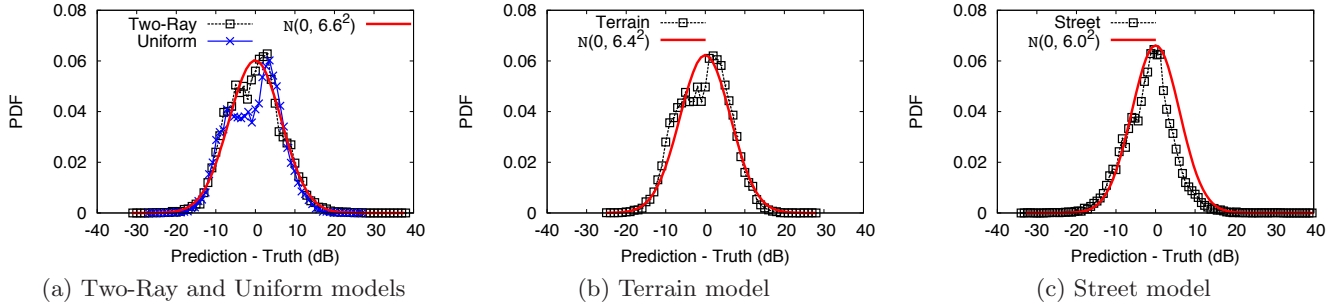


Figure 3: Probability density distributions of prediction errors using four calibrated propagation models. The reference zero-mean Gaussian curves are also displayed for each calibrated model. Prediction errors approximately follow zero-mean Gaussian distributions with standard deviation in $[6, 6.6]$.

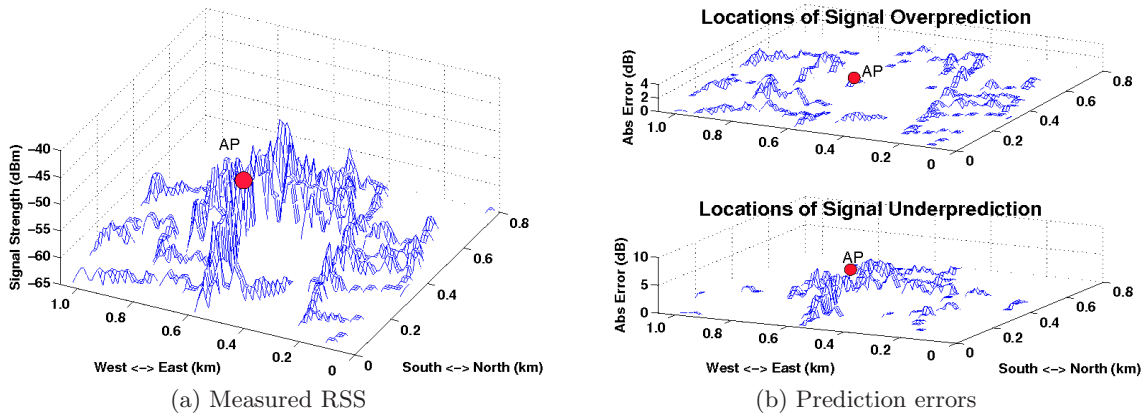


Figure 4: (a) Spatial distribution of measured signal strength values. Signal strength generally follows the power law, but varies significantly over space. Measurements could not be performed in areas with buildings or obstacles, and they show up as blank areas in plots. (b) Areas of signal overprediction and underprediction and absolute error values. Predictions use the Street model. Locations close to the AP tend to be underpredicted while those further away tend to be overpredicted.

the parameters for each propagation type separately. Like Terrain, this model requires street information, and thus can only be calibrated using the GoogleWiFi dataset.

3.2 Signal Prediction Results

We quantify signal prediction errors as the difference between the predicted signal strength (in dBm) and the measured signal strength (in dBm). We observe prediction errors that range from -30dB (under-prediction) to 30dB (over-prediction). We make three key observations.

Observation 1: Impact of Sampling. To calibrate our models, we randomly select sub-samples from the exhaustive measurement data. We vary the density of these samples from 1.4 to 100 samples per km^2 , or 10 to 700 total samples for an area of size $7km^2$. For all four models, we observe that increasing density beyond 34 samples per km^2 (239 total samples) leads to negligible gain in performance. Thus we use this sampling density for all our later tests. We also observe that calibration often yields surprising results, *e.g.* we find that the calibrated pathloss exponent for the Uniform model varies between 1.15 and 2.20 for our three datasets, while typical rule of thumb suggests 2 or 3.

Table 2: Standard deviation of prediction error

Dataset	Standard deviation			
	Uniform	Two-Ray	Terrain	Street
GoogleWiFi	6.6	6.6	6.4	6.0
MetroFi	8.4	8.1	N/A	N/A
TFA	7.6	7.4	N/A	N/A

Observation 2: Impact of Models. We observe that prediction errors are visible, but they do not vary significantly across models (the street model performs slightly better). This matches prior work [14, 23, 36, 40]. Specifically, prediction error varies across locations, and can be approximated by a zero-mean Gaussian distribution. Figure 3 shows the probability density function (PDF) of the prediction errors and its Gaussian approximation using the GoogleWiFi dataset. The same trend holds for MetroFi and TFA, and we omit those results for brevity. Table 2 lists the standard deviation of the prediction error under each model and dataset.

Observation 3: Impact of Receiver Location. When examining the correlation between prediction error and location, we observe that *all four propagation models tend to un-*

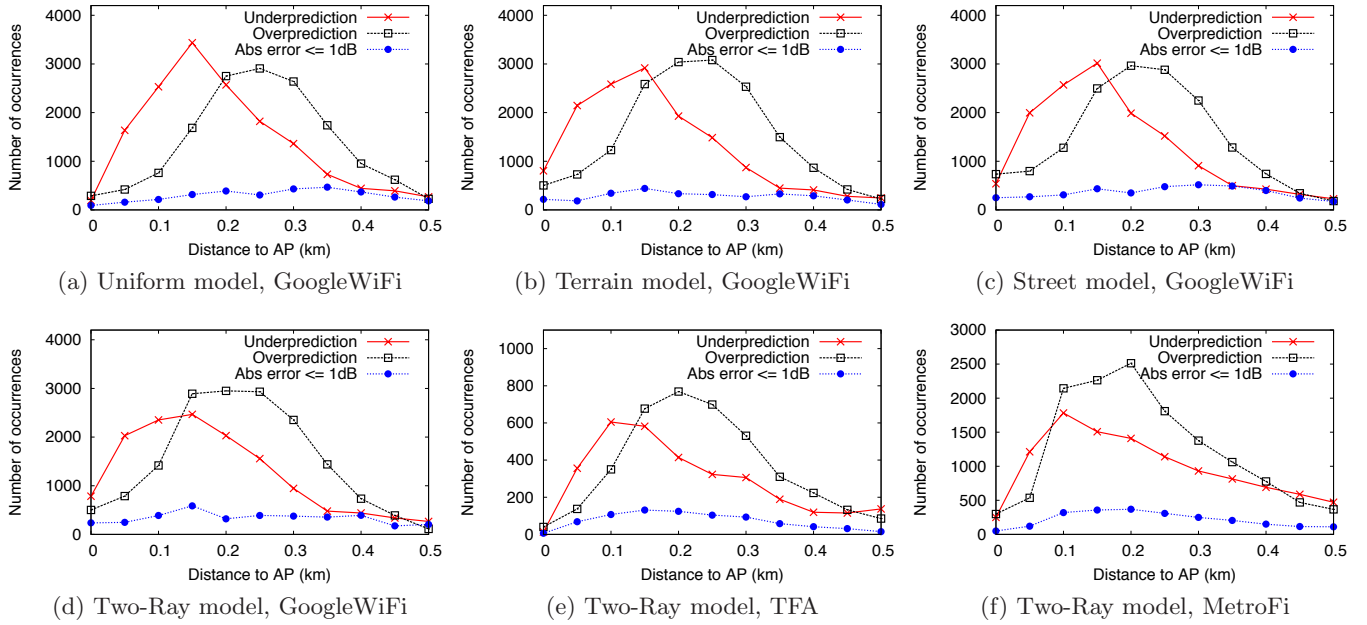


Figure 5: As the distance to AP increases, the main prediction error gradually switches from underprediction to overprediction. Here each data point summarizes the prediction errors within a distance interval of 0.05km. We show results with distances < 0.55 km because the number of data points above this range is insufficient to demonstrate any trends.

derpredict signal strength in area near the transmitter, and overpredict signal strength in areas far from the transmitter.

We illustrate this pattern graphically in Figure 4 by plotting the measured signal strength distribution of a randomly selected AP, and the levels of signal overprediction and underprediction at different locations. We define signal overprediction (under-prediction) to be when the predicted signal strength is larger (smaller) than the actual value by more than 1dB. This result shows a strong correlation between locations (in terms of their distance to the AP) and the type of expected error.

A closer look shows that this effect is consistent across all four propagation models and all three datasets. In Figure 5, we sort each AP’s measurement locations by their distances to the AP, and group them into buckets of 0.05km. For each interval, we calculate the occurrence of locations with accurate signal prediction (absolute error ≤ 1 dB) and those with overprediction and underprediction. We observe that the trend is consistent across all settings.

One possible explanation is that these propagation models still cannot fully capture how RF signals in urban street environments experience faster attenuation after traveling a certain distance [16, 17, 14]. Although both the Terrain and Street models seek to capture the impact of non-uniform signal propagation, they still use the single-slope pathloss model, and cannot fully reflect the dual slope feature of signal propagation. Thus the dual slope effect is more evident on these models. The Two-Ray model considers this feature, but is limited by the use of a uniform breakpoint distance that does not exist in practice [17]. Thus we still observe some errors for this model.

Summary of Findings. Our accuracy analysis shows that propagation models, even after careful calibration, in-

roduce visible but location-dependent errors in signal map prediction. This naturally leads to the question: how will the signal prediction errors translate into errors in estimated conflict graphs? We explore this question next.

4. CONFLICT GRAPH ACCURACY

Having analyzed the errors introduced by predictions from calibrated signal propagation models, we now examine the actual accuracy of the “estimated conflict graphs” produced using these results. More specifically, we ask the question: *what is the impact of imperfect signal strength predictions on the accuracy of their resulting conflict graphs?* As before, we use our measurement data as ground truth to produce “measured conflict graphs,” and use them to gauge the accuracy of “estimated conflict graphs.”

Here, we use similarity between the conflict graphs as a measure of the accuracy of estimated conflict graphs. Since both types of conflict graphs share the same vertices, graph similarity in this context reduces to a measure of overlap in the set of edges between graphs. In addition, our analysis is limited to the GoogleWiFi dataset, because lower receiver sensitivity in the other datasets resulted in extremely sparse conflict graphs with less than 20 edges. For each estimated conflict graph, we use 50 random sampling sets to produce 50 graph instances. Since these graphs across different samples are highly similar (≤ 3 edge difference), we only show average results here for brevity.

4.1 Graph Similarity Results

We compare each estimated conflict graph against the measured conflict graph, and classify each edge in the estimated graph as correct, extraneous, or missing:

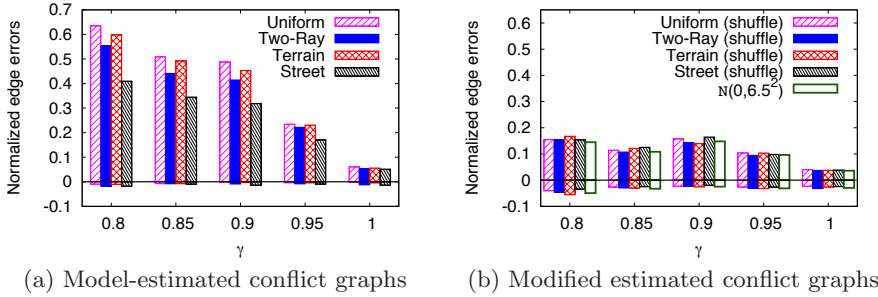


Figure 7: (a) Edge errors in the estimated conflict graph, normalized by the number of edges in the measured conflict graph. Negative (positive) bars denote the normalized count of missing (extraneous) edges. (b) Edge errors in the “modified” estimated conflict graph by removing the location dependency of the prediction errors. The amount of extraneous edges reduces significantly.

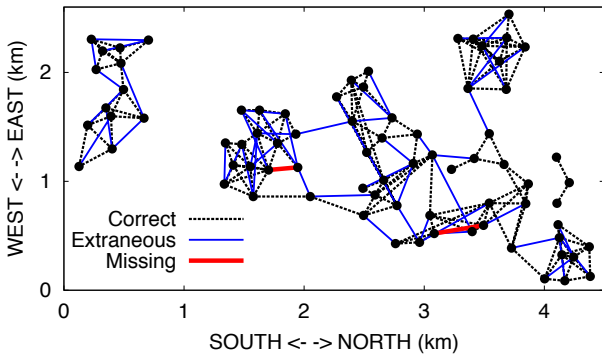


Figure 6: Accuracy of an estimated conflict graph.

- *Correct edges*: edges found in both estimated and measured conflict graphs.
- *Extraneous edges*: edges in the estimated conflict graph but not in the measured conflict graph; these edge errors make the estimated conflict graph more conservative, reducing spectrum utilization.
- *Missing edges*: edges in the measured conflict graph but missing in the estimated conflict graph; these errors are more harmful than extraneous edges, because they reduce the reliability of the estimated conflict graph and lead to harmful interference when conflicting nodes are assigned to the same channel.

Figure 6 shows a sample of estimated conflict graph generated using the Street model. Distances between nodes are shown to scale. Compared to the measured conflict graph with 162 edges, the estimated graph misses only 2 edges (thick red lines) and introduces 51 extraneous edges (blue lines). While slightly conservative, the estimated conflict graph is able to capture most of the edges.

We then compute the *normalized edge errors* as the number of extraneous and missing edges normalized by the total number of edges in the measured conflict graph. Figure 7(a) shows the normalized edge errors as the value of γ varies. We display normalized extraneous edges as positive values and the normalized missing edges as negative values. The




Figure 8: The values of p_{ij} for both correct and extraneous edges in the estimated conflict graph, using the Street model and $\gamma = 0.9$. We denote the two kinds of edges with different markers, and spread them out vertically using random Y values.

results show that the majority of edge errors are extraneous edges. Missing edges account for less than 2% of the edges of the measured graph. This pattern holds across different propagation models and for different values of γ .

Comparing across propagation models, we see that the choice of propagation models has only minor impact on the accuracy of the estimated conflict graph. The Uniform model is the most conservative and generates a slightly higher ratio of extraneous edges, and the Street model provides the best overall performance. This is likely because of the higher accuracy achieved by the Street model, which treats the reflected paths as the main components in non-line-of-sight (NLOS) scenarios. As a result, it is more accurate for urban street environments such as Mountain View, and leads to less edge errors in estimated graphs.

Figure 7(a) also shows that the normalized occurrence of extraneous edges decreases as γ increases. This is because increasing γ lowers the bar for two nodes to conflict with each other, thus producing more³ edges in the measured conflict graph, and shrinking the pool of potential extraneous edges for the estimated graph. Thus the ratio of extraneous edges decreases from 40-60% ($\gamma = 0.8$) to 5-8% ($\gamma = 1$).

4.2 Why Do Extraneous Edges Dominate?

The fact that extraneous edges dominate the errors can be attributed to two factors. The first is the location-dependent pattern of signal prediction errors described in Section 3. It causes under-prediction of signal strength and over-prediction of interference strength. Hence the majority (70+%) of pairwise SINR values are underpredicted, leading to many extraneous edges.

To verify this hypothesis, we build a new set of *modified* estimated conflict graphs using the same model-predicted signal maps, but make the prediction error randomly distributed across locations. We use two methods to remove the location dependency. The first method gathers the prediction errors of the model-generated signal maps, shuffles them randomly across different measurement locations, and adds them back to the measured signal map. The second method produces a synthetic pattern of prediction errors from a zero-mean Gaussian distribution with standard de-

³As γ grows from 0.8 to 1, the edge counts of the measured conflict graphs are 104, 132, 162, 243, 446, respectively.

violation of 6.5, and adds them to the measured signal map. Figure 7(b) shows that normalized edge errors in these modified estimated graphs have much fewer extraneous edges than their unmodified counterparts, only 5–15% vs. 5–60%. This confirms our hypothesis.

The second factor contributing to more extraneous edge errors is the fact that missing edge errors occur under more stringent conditions, *i.e.* it takes more signal errors to remove an edge than to add an edge. To erroneously remove an edge between i and j , both predicted ratios of conflict-free locations (q_i^j and q_j^i , defined by Eq. (1)) must exceed γ . In contrast, erroneously adding an edge between i and j only requires one of these two estimates to fall below γ . This factor explains why extraneous edges still outrun missing edges even after removing the location-dependency in the prediction errors (Figure 7(b)).

We note that these extraneous edges are not due to possible under-measurement of interference in our dataset, *i.e.* some weak interference signals may not be captured by our measurement receivers. This is because when computing SINR values used to build estimated graphs, we ignore interferers whose signals are not captured by the dataset.

Can We Identify Extraneous Edges? Since extraneous edges make up most of our observed edge errors, it is tempting to try to identify those edges in the estimated conflict graph and correct them. After carefully examining our traces, we found no distinctive characteristics that distinguish extraneous edges from correct edges. For example, Figure 8 plots the value of p_{ij} (defined by Eq. (1)) for each node pair i and j , calculated from the predicted signal strength distribution. We use different markers to separate the correct and extraneous edges. We see that there is no clear distinction between the two sets.

4.3 Summary of Findings

Our graph accuracy analysis reveals two key findings:

Estimated conflict graphs are conservative. The large majority of errors are extraneous edges; estimated graphs rarely miss edges (<2%).

Location-dependent signal prediction errors are the main cause of extraneous edges. The location-dependent error pattern in signal prediction triggers under-prediction of more than 70+% of SINR values, and is the main cause of extraneous edge errors.

5. IMPACT ON SPECTRUM ALLOCATION

After examining the accuracy of estimated conflict graphs at graph level, we now quantify the end-to-end impact of such graph errors on spectrum market users. Also, given the lack of representation of accumulative interference in conflict graphs, we aim to understand how such artifact affects the quality of the allocated spectrum. For this, we distribute spectrum using both measured and estimated conflict graphs, and evaluate end-to-end performance in terms of the efficiency of spectrum utilization and link reliability of allocated spectrum.

5.1 Spectrum Allocation Benchmarks

To translate conflict graphs into actual spectrum allocations, we use two representative allocation algorithms. Both seek to efficiently distribute a given spectrum range across

the market users, while ensuring that no conflicting users receive the same spectrum band.

- *Single channel allocation (SCA)* seeks to allocate equal amount of continuous spectrum frequencies to each spectrum market user. This problem reduces to the well-known graph-coloring problem, which uses the minimal number of channels to ensure that each market user receives one channel and does not conflict with another [38]. The fewer the number of channels required, the greater the bandwidth assigned to each channel.
- *Multi-channel allocation (MCA)* divides the spectrum range into a large number of channels. Each market user can receive multiple channels even if they are not continuously aligned in frequency [10]. MCA distributes the channels across market users to maximize a predefined system utility, *e.g.* proportional fairness [10].

We evaluate the resulting spectrum allocation based on:

- *Spectrum efficiency:* the average amount of spectrum received per market user, normalized by the available spectrum. We calculate it by counting the number of channels the market user receives and the frequency bandwidth of each channel.
- *Spectrum reliability:* whether each market user actually receives exclusive spectrum usage. A market user has received exclusive usage if on each assigned channel, the percentage of receivers whose actual SINRs exceed β is no less than γ . This is to evaluate the actual outcome of spectrum usage in terms of end-user coverage. For our purposes, we use *network-level spectrum reliability*, which is defined as the percentage of market users receiving exclusive spectrum usage.

Given a spectrum allocation, we use our full dataset to calculate the actual SINR observed at each market user’s receiver locations. The actual SINRs of each user include interference accumulated from *all* other market users operating on the same channel, which captures the true performance perceived by any allocated market user. While it is well known that the lack of coverage on accumulative interference will affect spectrum allocations under conflict graphs, our goal is to understand the severity of such degradation, and seek solutions to effectively suppress them.

5.2 Spectrum Allocation Results

Spectrum Efficiency. Figure 9 compares the spectrum efficiency of using the measured and estimated conflict graphs in allocating spectrum. Because extraneous conflict edges will prevent some non-conflicting users from reusing spectrum, spectrum efficiency using estimated graphs is lower than that of the measured graph. The efficiency loss, however, never exceeds 30%, even for cases where estimated graphs introduce 40%–60% extra edges (those with $\gamma = 0.8$). The loss of efficiency also reduces as γ increases, because the amount of extraneous edges reduces significantly. As before, the Street model outperforms the other models modestly because of its higher accuracy in signal prediction, which leads to less edge errors in the resulted conflict graph.

Spectrum Reliability. Using the actual SINR as the metric, we now examine the link reliability of the allocated spectrum. As expected, we see from Figure 10 that the reliability is not 100% for both allocation algorithms. Unless $\gamma = 1$, the measured conflict graph makes 10–50% of market






Figure 9: Spectrum efficiency using measured and estimated conflict graphs to distribute spectrum. The use of estimated conflict graphs leads to spectrum efficiency loss, which is bounded by 30% and becomes negligible as γ approaches 1.



(a) SCA



(b) MCA

Figure 10: Spectrum reliability when using measured and estimated conflict graphs to distribute spectrum. The reliability is between 80%-98% for the estimated graph and drops to 50% for the measured conflict graph. This indicates that the impact of accumulative interference is noticeable and it is not captured by these conflict graphs.

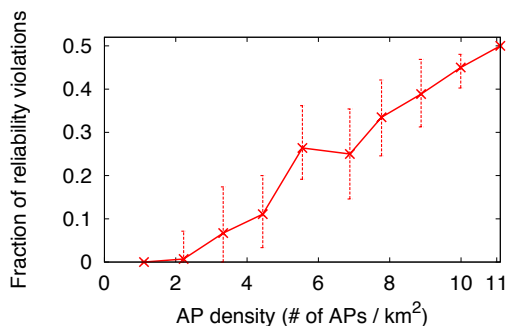


Figure 11: Reliability violations increase with the AP density. $\gamma = 0.9$, using measured conflict graphs. For the GoogleWiFi dataset, the AP density is 11 APs per km^2 .

users unsatisfactory. In comparison, the estimated graphs actually lead to more reliable spectrum usage for market users because of having extraneous edges. In this regard, the extraneous-edge errors help ensure spectrum reliability.

The results demonstrate that accumulative interference does cause noticeable impact on the spectrum usage. Because interference experienced by a receiver is the accumulative sum of signals from transmitters operating on the same frequency band, the higher the spectrum reuse in the neighborhood, the higher the level of accumulative interference. When using the measured conflict graph, the spectrum reuse level is very high, *e.g.* 30 market users per channel for $\gamma = 0.8$. Therefore the effect of accumulative interference is significant. As γ increases, the reuse level decreases, and so does the effect of accumulative interference. For estimated conflict graphs, their conservative allocation from extraneous conflict edges reduces spectrum reuse, and thus the level of accumulative interference. This effect also motivates us to further examine the conditions under which accumulative interference would be a prevalent effect.

How Prevalent Is Accumulative Interference? In contrast to our results, prior work on a 32-node network reports that accumulative interference has negligible effect on wireless transmissions [11]. This begs the question: under what conditions will accumulative interference matter? To answer this, we first examine the spatial locations of the

market users with reliability violations in our GoogleWiFi dataset. We see that most of them are clustered in the center of the physical area with high market user density. This indicates that node density is a large contributing factor.

To examine the impact of node density, we build a set of new market configurations by sampling the APs in GoogleWiFi dataset uniformly, while keeping each AP's coverage area unchanged. For each new configuration, we build a conflict graph from the exhaustive signal measurements, and examine its reliability using the MCA allocation. Figure 11 shows the percentage of market users with reliability violations, which grows with the AP density. For the current GoogleWiFi network, the average density is 11 APs per km^2 , which is common for municipal wireless networks. Thus we conclude that accumulative interference does matter in many current and future wireless deployments. We must address such artifact in order to use conflict graphs in practice.

6. GRAPH AUGMENTATION

The spectrum reliability violations we found for both measured and estimated conflict graphs are clearly undesirable for the practical deployment of spectrum markets. In this section, we seek for solutions to eliminate the artifact of accumulative interference for both conflict graphs. This ensures exclusive spectrum usage with reasonable level of reliability, addressing the key concerns on conflict graphs and promoting their practical usage.

6.1 Challenges

To reduce the impact of accumulative interference, one intuitive method is to augment the existing conflict graphs by adding more edges and making them more conservative. This essentially reduces the number of users who get allocated with the same spectrum channel, thus the amount of accumulative interference. However, adding more edges inevitably leads to loss in spectrum efficiency. Hence the key challenge is to minimize the number of edge additions while eliminating the artifact of accumulative interference.

To show the level of difficulty in this task, let us begin with two straw-man solutions. The first solution is to randomly add edges to unconnected node pairs (referred to as *Random*). A smarter alternative is to first sort unconnected node pairs by the physical distance between their transmit-




Figure 12: The performance of graph augmentation. The estimated graphs are generated from the Street model. (a)-(b) Spectrum reliability results before and after graph augmentation, using different augmentation algorithms (Greedy-Feedback, Locality-based, and Random). Greedy-Feedback is highly effective, and outperforms the other two. (c) Spectrum efficiency before and after graph augmentation via Greedy-Feedback. The improvement in reliability is at the cost of slightly degraded spectrum efficiency. The gap between measured and estimated graphs reduces to less than 15% after augmentation.

ters, and only add edges to the top- K closest node pairs. We refer to this approach as *Locality-based augmentation*. While simple, these two solutions face two drawbacks: 1) each added edge might not effectively reduce accumulative interference; and 2) it is difficult to determine the correct number of edges to add.

6.2 Greedy-Feedback Graph Augmentation

We overcome the above challenge by proposing a greedy algorithm to gradually and intelligently add edges. This algorithm stops adding edges when the (estimated) reliability reaches 100%, assuming wireless interference is the only source of reliability loss. Because the level of accumulative interference depends on the spectrum allocation algorithm, we integrate graph augmentation with spectrum allocation.

More specifically, the augmentation procedure works as follows. After allocating spectrum using the current conflict graph, we examine the reliability performance of each node, and identify the node i with the lowest reliability and its worst channel m . Next, we find node j , who is currently allocated with channel m , and whose removal will lead to the largest reliability improvement at i . We then add an edge between node i and j , and repeat the above process until all nodes have met the reliability requirement γ .

We use this approach to augment both measured and estimated conflict graphs. The only difference is that when augmenting a measured graph, we compute reliability using the real signal strength map. In contrast, we augment estimated graphs by estimating reliability from the predicted signal strength map.

6.3 Evaluation Results

We evaluate the effectiveness of our augmentation algorithm by comparing spectrum reliability and efficiency before and after augmentation. Our evaluation uses the MCA allocation since it suffers more accumulative interference.

Effectiveness of Graph Augmentation. In Figure 12(a)-(b), we first compare the three augmentation techniques in terms of spectrum reliability. For a fair comparison, we apply Random and Locality-based augmentation to add the same number of edges as that of Greedy-Feedback.

Greedy-Feedback graph augmentation is highly effective and significantly outperforms the other two techniques. It

completely removes the impact of accumulative interference on the measured conflict graph, and boosts the reliability of the estimated graphs to 96+%. The reliability of estimated graphs is not always 100%, because the augmentation algorithm relies on reliability predictions from signal strength estimates. In contrast, Random graph augmentation leads to no visible improvement on reliability while Locality-based augmentation is half way between Random and Greedy-Feedback.

By adding edges, graph augmentation does lead to lower spectrum efficiency. From Figure 12(c), we see that by using Greedy-Feedback, we get relative efficiency loss between 0–25% for the measured graph and 0–15% for the estimated graphs. The loss for the estimated graphs is lower because it adds less number of edges. We see that the proposed graph augmentation is effective against accumulative interference.

We also observe that graph augmentation has no effect when $\gamma = 1$. This is because both types of conflict graphs already contain a large number of edges (440+). The resulting spectrum allocation has very limited reuse across users, and the impact of accumulative interference is negligible.

Accuracy of Augmented Conflict Graphs. Applying graph augmentation to our measured conflict graph produces an *ideal* conflict graph, one that captures real conflicts and the impact of accumulative interference. We now look at how close the augmented, estimated conflict graph is relative to this ideal graph. Figure 13(a) plots the normalized edge errors of the estimated graphs after graph augmentation. The ratios of extraneous edges reduce from 5%-60% (Figure 7(a)) to 5%-40%, while the ratios of missing edges remain similar. This is because the augmentation on the measured graph adds more edges, and some of these edges already appear in the estimated graph.

Finally, we look at the spectrum efficiency and reliability of the estimated graph after augmentation, and see that both have improved relative to the ideal graph. Figure 13(b)-(c) show that the use of estimated conflict graphs achieves nearly 100% spectrum reliability and only leads to at most 21% loss in spectrum efficiency. Overall, we see that the Street model is the most efficient ($\leq 15\%$ efficiency loss, 96+% reliability) among the four propagation models.

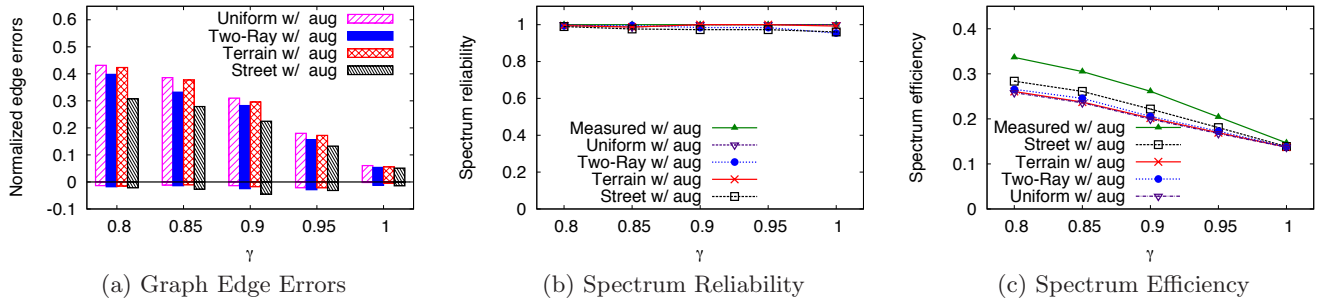


Figure 13: Accuracy of augmented estimated conflict graphs using the four propagation models, compared to the ideal conflict graph. (a) The edge errors reduce considerably with graph augmentation. (b) The reliability of using estimated graphs in spectrum allocation increases to 96+%. (c) The efficiency loss of using estimated graphs reduces to no more than 21%.

Key Findings. We summarize the key results on graph augmentation as the following two findings:

- **Graph augmentation is effective against accumulative interference.** Proper graph augmentation effectively boosts spectrum reliability to 96-100%, while maintaining spectrum efficiency.
- **Augmentation improves the accuracy of estimated conflict graphs.** Augmentation reduces the difference between measured and estimated conflict graphs. Using estimated graphs in spectrum allocation results in only 15% or less loss in spectrum efficiency.

7. DISCUSSION

We discuss possible extensions of our methodology beyond the scenarios covered by our study.

Temporal Signal Variations. Our work focuses on the long-term impact of interference by considering the average signal values. To understand the impact of temporal signal variations, we can adapt the conflict graphs based on periodical sensor measurements. For applications that must consider fast signal fading, the conflict edge can be determined using the outage SINR, the bottom $x\%$ of SINR observed within a certain time period. Examining the accuracy of such conflict graphs is an interesting future study.

Incorporating MAC Protocols. Our analysis does not consider the impact of MAC protocols because in the exclusive usage scenario, a spectrum market user can use any MAC protocol [3, 50]. However, for scenarios where all the users adopt the same MAC protocol, one can integrate a traffic-driven model [25, 43] into the conflict graph.

8. RELATED WORK

Conflict Graphs and Interference Models. We divide existing works into two categories based on the type of conflict graphs they use. The first category uses per-link signal measurements to capture interference conditions among individual links, using either active measurements [5, 6, 26, 33, 34, 37, 39, 43], or passive measurements [8, 21, 28, 45]. These link-based conflict graphs are for indoor WiFi networks where transmission links are known a priori. They are impractical for outdoor networks with mobile users that spectrum markets target.

The second category of works builds coverage-based conflict graphs based on propagation models, either with rule-of-thumb parameters [18, 32, 48], or calibrated by on-site measurements [24]. However, no one has used real-world measurements to evaluate the conflict graph accuracy. Our work is the first measurement study on this problem. We use both graph and spectrum allocation analysis to understand the feasibility of building accurate coverage-based conflict graphs for dynamic spectrum distribution.

Aside from conflict graphs, recent work examines the accuracy of general interference models for small-scale networks using per-link measurements [27]. Our work was inspired by this work, yet focuses on large-scale outdoor networks where per-link measurement is infeasible. We use conflict graphs as interference models because they are widely used by spectrum allocation solutions. Our methodology can be extended to other interference models such as SINR [9, 47].

Measurement-calibrated Propagation Models. Measurement studies show that RF propagation models with rule-of-thumb parameters introduce large errors in signal strength estimation [14, 23, 36]. When calibrated using on-site measurements, however, these propagation models offer higher accuracy, and have been used in cell planning [19, 29], interference management [40] and coverage prediction [31, 42]. Our work complements these prior works, and is also inspired by prior work on measurement-calibrated models for social network graphs [44].

9. CONCLUSION

Using large-scale signal measurements, we examined the severity of two key concerns on using conflict graphs for dynamic spectrum distribution. We focused on conflict graphs built from measurement-calibrated propagation models, and studied their accuracy and the end-to-end impact on spectrum allocation. We found that the resulting “estimated conflict graphs” are conservative compared to precise conflict graphs built from exhaustive signal measurements. Yet surprisingly, these extraneous edges improve link reliability by alleviating the impact of accumulative interference, an artifact not captured by conflict graphs. We proposed a graph augmentation technique to suppress the impact of accumulative interference. With this new technique, estimated conflict graphs can produce spectrum allocations that provide near-perfect link reliability, with spectrum efficiency

less than 15% away from the ideal allocation. We believe that for the WiFi frequencies studied by this paper, (and their nearby frequencies), our proposed techniques address existing concerns on conflict graphs, and provide a scalable and accurate end-to-end solution for spectrum allocation.

10. ACKNOWLEDGMENTS

We thank the anonymous reviewers for their helpful feedback, and Lei Yang for his help with the measurements. This work was supported in part by NSF grants CNS-0905667 and CNS-0915699.

11. REFERENCES

- [1] Coverage or capacity? best use of 802.11n. Trapeze Networks Whitepaper.
- [2] MetroFi. <http://crawdad.cs.dartmouth.edu/pdx/metrofi/2007/coverage>.
- [3] Secondary markets in radio spectrum. <http://www.fcc.gov/events/secondary-markets-radio-spectrum>.
- [4] TFA. <http://tfa.rice.edu/measurements/>.
- [5] AHMED, N., ISMAIL, U., AND KESHAV, S. Online estimation of RF interference. In *Proc. of CoNEXT* (2008).
- [6] AHMED, N., AND KESHAV, S. SMARTA: a self-managing architecture for thin access points. In *Proc. of CoNEXT* (2006).
- [7] ALICHERY, M., BHATIA, R., AND LI, L. E. Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks. In *Proc. of MobiCom* (2005).
- [8] CAI, K., ET AL. Non-intrusive, dynamic interference detection for 802.11 networks. In *Proc. of IMC* (2009).
- [9] CAO, L., ET AL. Optimus: SINR-Driven Spectrum Distribution via Constraint Transformation. In *Proc. of DySPAN* (2010).
- [10] CAO, L., AND ZHENG, H. Distributed spectrum allocation via local bargaining. In *Proc. of SECON* (2005).
- [11] DAS, S. M., ET AL. Characterizing multi-way interference in wireless mesh networks. In *Proc. of WiNTECH* (2006).
- [12] DEB, S., SRINIVASAN, V., AND MAHESHWARI, R. Dynamic spectrum access in DTV whitespaces: design rules, architecture and algorithms. In *Proc. of MobiCom* (2009).
- [13] GANDHI, S., BURAGOHAJ, C., CAO, L., ZHENG, H., AND SURI, S. A general framework for wireless spectrum auctions. In *Proc. of DySPAN* (2007).
- [14] GOLDSMITH, A. *Wireless communications*. Cambridge Univ Pr.
- [15] GOLDSMITH, A., AND GREENSTEIN, L. A measurement-based model for predicting coverage areas of urban microcells. *IEEE JSAC* 11, 7 (1993).
- [16] GREEN, E. Radio link design for microcellular systems. *British Telecom technology journal* 8, 1 (1990), 85–96.
- [17] GREEN, E., AND HATA, M. Microcellular propagation measurements in an urban environment. In *Proc. of PIMRC* (1991).
- [18] HAAS, Z., WINTERS, J., AND JOHNSON, D. Simulation study of the capacity bounds in cellular systems. In *PIMRC/WCN* (1994).
- [19] HURLEY, S. Planning effective cellular mobile radio networks. *IEEE TVT* 51, 2 (2002), 243–253.
- [20] JAIN, K., PADHYE, J., PADMANABHAN, V. N., AND QIU, L. Impact of interference on multi-hop wireless network performance. In *Proc. of MobiCom* (2003).
- [21] JANG, K.-Y., ET AL. Passive on-line in-band interference inference in centralized WLANs. In *Tech. Rep. 916, USC* (2010).
- [22] KATZELA, I., AND NAGHSHINEH, M. Channel assignment schemes for cellular mobile telecommunication systems: a comprehensive survey. *Personal Communications, IEEE* 3, 3 (1996), 10–31.
- [23] KOTZ, D., ET AL. Experimental evaluation of wireless simulation assumptions. In *Proc. of MSWiM* (2004).
- [24] KUURNE, A. Mobile measurement based frequency planning in GSM networks. *MS thesis, Helsinki University of Technology* (2001).
- [25] LI, Y., ET AL. Predictable performance optimization for wireless networks. In *Proc. of SIGCOMM* (2008).
- [26] LIU, X., ET AL. DIRC: Increasing indoor wireless capacity using directional antennas. In *Proc. of SIGCOMM* (2009).
- [27] MAHESHWARI, R., JAIN, S., AND DAS, S. R. A measurement study of interference modeling and scheduling in low-power wireless networks. In *Proc. of SenSys* (2008).
- [28] MANWEILER, J., ET AL. Order matters: Interference-aware transmission reordering in wireless networks. In *Proc. of MobiCom* (2009).
- [29] MISHRA, A. *Fundamentals of cellular network planning and optimisation*. Wiley Online Library, 2004.
- [30] MOSCIBRODA, T., WATTENHOFER, R., AND WEBER, Y. Protocol design beyond graph-based models. In *Proc. of HotNets* (2006).
- [31] MURTY, R., ET AL. SenseLess: A database-driven white spaces network. In *Proc. of DySPAN* (2011).
- [32] NECKER, M. C. Towards frequency reuse 1 cellular FDM/TDM systems. In *Proc. of MSWiM* (2006).
- [33] NICULESCU, D. Interference map for 802.11 networks. In *Proc. of IMC* (2007).
- [34] PADHYE, J., ET AL. Estimation of link interference in static multi-hop wireless networks. In *Proc. of IMC* (2005).
- [35] PENG, C., ZHENG, H., AND ZHAO, B. Y. Utilization and fairness in spectrum assignment for opportunistic spectrum access. *MONET* 11, 4 (2006), 555–576.
- [36] PHILLIPS, C., SICKER, D., AND GRUNWALD, D. Bounding the error of path loss models. In *Proc. of DySPAN* (2011).
- [37] QIU, L., ET AL. A general model of wireless interference. In *Proc. of MobiCom* (2007).
- [38] RAMANATHAN, S. A unified framework and algorithm for channel assignment in wireless networks. *Wirel. Netw.* 5, 2 (1999), 81–94.
- [39] RAYANCHU, S., ET AL. FLUID: improving throughputs in enterprise wireless LANs through flexible channelization. In *Proc. of MobiCom* (2011).
- [40] REIS, C., ET AL. Measurement-based models of delivery and interference in static wireless networks. In *Proc. of SIGCOMM* (2006).
- [41] RIBACK, M., ET AL. Carrier frequency effects on path loss. In *Proc. of VTC* (2006).
- [42] ROBINSON, J., SWAMINATHAN, R., AND KNIGHTLY, E. W. Assessment of urban-scale wireless networks with a small number of measurements. In *Proc. of MobiCom* (2008).
- [43] ROZNER, E., ET AL. Traffic-aware channel assignment in enterprise wireless networks. In *Proc. of ICNP* (2007).
- [44] SALA, A., CAO, L., WILSON, C., ZABLIT, R., ZHENG, H., AND ZHAO, B. Y. Measurement-calibrated graph models for social network experiments. In *Proc. of WWW* (2010).
- [45] SHRIVASTAVA, V., ET AL. PIE in the sky: Online passive interference estimation for enterprise WLANs. In *Proc. of NSDI* (2011).
- [46] SONG, L., AND SHEN, J. *Evolved cellular network planning and optimization for UMTS and LTE*. CRC Press, 2010.
- [47] SUBRAMANIAN, P., ET AL. Near-optimal dynamic spectrum allocation in cellular networks. In *Proc. of DySPAN* (2008).
- [48] YANG, L., CAO, L., AND ZHENG, H. Physical interference driven dynamic spectrum management. In *Proc. of DySPAN* (2008).
- [49] YOUNIS, M., AND AKKAYA, K. Strategies and techniques for node placement in wireless sensor networks: A survey. *Ad Hoc Networks* 6, 4 (2008), 621–655.
- [50] ZHOU, X., GANDHI, S., SURI, S., AND ZHENG, H. eBay in the sky: Strategy-proof wireless spectrum auctions. In *Proc. of MobiCom* (2008).