

Practical Conflict Graphs in the Wild

Xia Zhou, Zengbin Zhang, Gang Wang, Xiaoxiao Yu, Ben Y. Zhao, and Haitao Zheng, *Senior Member, IEEE*

Abstract—Today, most spectrum allocation algorithms use conflict graphs to capture interference conditions. The use of conflict graphs, however, is often questioned by the wireless community for two reasons. First, building accurate conflict graphs requires significant overhead, and hence does not scale to outdoor networks. Second, conflict graphs cannot properly capture accumulative interference. In this paper, we use large-scale measurement data as ground truth to understand how severe these problems are 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. Calibrated models are imperfect, and we study the impact of their errors on multiple steps in the process, from calibrating propagation models, predicting signal strengths, to 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. Our work produces several findings. Calibrated propagation models generate *location-dependent prediction errors*, ultimately producing conservative conflict graphs. While these “estimated conflict graphs” lower spectrum utilization, their conservative nature improves reliability by reducing the impact of accumulative interference. Finally, we propose a graph augmentation technique to address remaining accumulative interference.

Index Terms—Conflict graphs, dynamic spectrum access, interference.

I. INTRODUCTION

TO SUPPORT the rapid growth of today’s wireless technologies, current reforms in radio spectrum management target the use of on-demand auctions and secondary markets. These spectrum markets not only hold the promise of great profit for spectrum owners, but also allow spectrum users (e.g., small cell providers) to purchase exclusive spectrum usage on an on-demand basis. Enabling this type of spectrum reform requires two tightly coupled components: *a model of interference patterns* among spectrum users and *an allocation algorithm* that uses this interference model to distribute spectrum efficiently. The goal is to maximize spectrum utilization by parallelizing noninterfering transmissions whenever possible.

Manuscript received August 22, 2013; revised December 20, 2013; accepted February 11, 2014; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor A. Capone.

X. Zhou is with the Department of Computer Science, Dartmouth College, Hanover, NH 03755 USA (e-mail: xia@cs.dartmouth.edu).

Z. Zhang, G. Wang, B. Y. Zhao, and H. Zheng are with the Department of Computer Science, University of California, Santa Barbara, Santa Barbara, CA 93106 USA.

X. Yu is with the Department of Electrical and Computer Engineering, Tsinghua University, Beijing 100084, China.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2014.2306416

Most prior designs propose allocation algorithms using an interference model called “conflict graph” [1]. As the name suggests, a conflict graph is a simple graphical representation of the interference condition between any two spectrum users.¹ This model simplifies spectrum allocation design, leading to a series of highly efficient allocation algorithms with bounded performance and polynomial-time complexity [3]–[9]. In contrast, alternative physical interference models are complex and entail unbounded performance loss when used to develop allocation algorithms [10]–[12].

Despite their popularity, the practical value of conflict graphs is often questioned by the wireless community for two key reasons. *First*, building an accurate conflict graph is difficult. Given the complex nature of wireless propagation, it requires detailed measurements covering all combinations of sender/receiver locations. This type of per-link signal measurement is feasible for indoor WLANs [8], [13]–[18], but impractical for outdoor networks targeted by spectrum markets. Existing proposals build artificial conflict graphs using a simple distance-based criterion [3], [19]–[21] or using signal strengths generated by simple RF propagation models with rule-of-thumb parameters [6], [22]. These simplifications, however, produce incorrect interference estimates that lead to poor performance [23], [24].

Second, because conflict graphs only define interference conditions between any two spectrum users, they cannot capture the impact of interference accumulated from multiple concurrent transmissions on the same frequency band. Such “mismatch” leads to unpredicted, harmful interference at allocated users [25], breaking the exclusive usage promised by the spectrum market. Without guarantees that their transmissions operate without interference, users would have little incentive to purchase spectrum from the spectrum market.

In this paper, we use a data-driven approach to understand 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 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, addressing the community’s 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 collecting exhaustive measurements, this approach only measures a subset of locations. These sampled measurements are used to calibrate a propagation model, which can predict signal values for all locations in the area. These predictions are used in lieu of exhaustive measurements to build the conflict graph. This approach has

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) [2].




Fig. 1. Our high-level methodology. We build *estimated conflict graphs* using a small number of signal measurements at randomly selected locations. We use these samples to calibrate a propagation model, predict signal strength values, and construct conflict graphs. We examine the accuracy of estimated conflict graphs by comparing them to *measured conflict graphs* built from exhaustive signal measurements, using graph similarity and spectrum allocation benchmarks.

two advantages. First, prior works have shown that measurement-calibrated propagation models are much more accurate than those with rule-of-thumb parameters [26]–[29]. Second, because measurements are collected by sensors or trusted network subscribers, this approach incurs low overhead and can offer continuous measurements in real time. This allows conflict graphs to adapt to network changes. We recognize, however, that calibrated propagation models are imperfect and will introduce errors in predicted signal maps [30]–[33]. Thus, we must understand whether these errors propagate to resulting conflict graphs and if they affect the efficacy of spectrum allocations for users.

Our high-level methodology is as follows (Fig. 1).

- Use a relatively small number of signal measurements to calibrate RF propagation models.
- Use calibrated models to build predicted signal maps, and use these maps to produce “estimated conflict graphs.”
- Compare estimated graph to “measured conflict graph” built from exhaustive measurements on graph similarity.
- Evaluate end-to-end impact by running spectrum allocation on both conflict graphs and comparing allocation results.

To the best of our knowledge, our work is the first empirical study on the practical usability of conflict graphs for dynamic spectrum distribution. Our work differs from existing works on constructing conflict graphs. First, our work focuses on outdoor networks, unlike prior efforts [8], [13]–[18] that build indoor conflict graphs using exhaustive signal measurements. Second, our work targets dynamic spectrum markets where users are at unplanned places and the resulting conflict graph can be of arbitrary shape. This is fundamentally different from cellular networks [26], [27], [34]–[36] that 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 produce location-dependent prediction errors. They tend to underpredict signal strengths at short distances, and overpredict them at long distances. We observe this pattern consistently for multiple datasets.
- These prediction errors lead to conservative conflict graphs that rarely miss conflict edges, but commonly introduce extraneous conflict edges.
- This leads to spectrum utilization loss compared to measured conflict graphs. These extra edges, on the other hand, reduce the impact of accumulative interference, thus achieving more reliable spectrum usage.

TABLE I
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

- Our graph augmentation technique eliminates the artifact of accumulative interference and boosts the reliability of spectrum allocation to 96+%. Once augmented, estimated graphs achieve 85+% of the ideal allocation’s utilization.

II. METHODOLOGY

We aim to examine key issues of using conflict graphs for dynamic spectrum distribution. We consider conflict graphs built from measurement-calibrated propagation models because they require little measurement overhead and are much more accurate than those built with rule-of-thumb parameters.

Our approach (Fig. 1) consists of four steps: 1) collecting real signal maps via measurements and using them as ground truth; 2) using sampled measurements to calibrate propagation models and predicting network-wide signal maps; 3) building conflict graphs using 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. Next, we briefly describe our assumptions and present each step in detail.

Assumptions: Our study uses wardriving measurements of outdoor WiFi networks. We assume that WiFi channels have the same propagation properties. We use WiFi band as an example of distributing spectrum in outdoor networks, where we know base station locations. Our work can be extended to other frequencies by adjusting the propagation model to consider frequency differences [37], [38]. We do not consider the impact of MAC protocols because in the exclusive usage scenario, a market user can use any MAC protocol [9], [39].

A. Collecting Signal Maps

We use wardriving measurements at three outdoor municipal WiFi networks, one of which was collected by our group. Each dataset consists of beacon RSS values of WiFi access points (APs) measured in an outdoor area of size 3–7 km^2 , measurement GPS locations, and AP locations. We average multiple RSS readings per location to derive a map of average signal strengths for each AP. Table I summarizes the datasets.




Fig. 2. Measured area in the GoogleWiFi dataset. Triangles are the APs detected, and dots are measured locations on the streets.

GoogleWiFi: Collected by our research group in April 2010, this dataset covers a 7-km² residential area of the Google WiFi network in Mountain View, CA, USA. Fig. 2 shows measurement locations (as dots) and APs (as triangles). We used three co-located laptops equipped with customized WiFi cards that have higher receive sensitivity.² Thus, this dataset records detailed signal strength values of 78 APs at 11 447 distinct locations (with an average 5 m 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 [40] consists of RSS values in a 7-km² area of an 802.11x municipal network in Portland, OR, USA. It was collected by a research group from the University of Colorado, Boulder, CO, USA, 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, Houston, TX, USA, this measurement data covers 22 APs in a 3-km² area of the TFA network in Houston, TX [41]. It includes measurements from 27 855 locations.

To use these datasets, we treat each AP as the transmitter of a market user and measured locations in its coverage area as its receiver locations. Our measurements are for WiFi networks, but measured signal maps and resulting conflict graphs are independent of specific MAC protocols. This matches the exclusive usage scenario, where a market user is free to use any MAC protocol in its authorized spectrum range.

B. Calibrating Propagation Models

To build “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. These include the simple uniform path-loss model and complex models that support environmental features. We now describe our high-level approach to model calibration and signal map prediction. We leave details to Section III.

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

We begin with choosing samples from our exhaustive measurements and vary sample density between 1.4 and 100 samples per km². We then use the minimum mean square error (MMSE) fitting method to determine best-fit parameters for each propagation model. Once a given model is calibrated, we interpolate signal values at other locations to build the complete signal map.

C. Constructing Conflict Graphs

We now have two signal maps, one from our exhaustive signal measurements and the other interpolated by our calibrated propagation model. We use them respectively to build a *measured conflict graph*, i.e., ground truth, and an *estimated conflict graph*. In these conflict graphs, each node represents a spectrum market user, and each edge represents a conflict between two users. A user maps to a stationary transmitter, i.e., an AP in our signal maps, and its coverage area including its receiver locations. 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-interference-plus-noise ratio (SINR) above β [5]. Here, coverage area, 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 values and determine conflict conditions using the bottom $(1 - \gamma)$ -percentile value. That is, nodes 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 that a spectrum user is willing to tolerate to maximize capacity [42]. When $\gamma = 1$, (1) reduces to the minimal SINR-based criterion [10], [12], [21].

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

D. Evaluating Graph Accuracy

Finally, we examine the accuracy of 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. We use both graph similarity metrics and spectrum allocation benchmarks.

For graph similarity, we compare edges in the estimated and measured conflict graphs, using the measured graph as ground truth. The edge differences are classified as “extraneous edges” and “missing edges.” We analyze patterns of extraneous and missing edges and explain their appearance based on errors in signal map prediction.

To understand how graph errors affect spectrum users, we feed estimated and measured conflict graphs to two well-known spectrum allocation algorithms and compare allocation results. These end-to-end tests answer two questions: Will graph errors translate to significant loss in spectrum efficiency and reliability, and will the “uncaptured” accumulative interference lead to significant loss?

Next, we present detailed results from each step of our analysis. We start by the accuracy of signal prediction using calibrated models (Section III). Then, we build and compare measured and estimated graphs on graph similarity (Section IV) and spectrum allocation performance (Section V).

III. SIGNAL PREDICTION ACCURACY

We first examine errors introduced when we use incomplete measurements to calibrate propagation models and predict signal strengths. We aim to understand whether there are error patterns that are likely to manifest in resulting conflict graphs. We take several representative propagation models, calibrate them using controlled measurement samples, and evaluate their signal predictions using the full dataset as ground truth.

A. Propagation Models and Calibration

We choose four representative propagation models, which capture urban street environments that best match our datasets. They range from the simplest uniform path-loss model to sophisticated models that incorporate terrain features. We focus on understanding whether existing propagation models can offer accurate signal predictions, rather than refining existing propagation models or proposing new models (we leave these to future works).

Uniform Path-Loss Model (Uniform) [30]: As the simplest and most-used model, it captures signal attenuation using a single path-loss exponent. Calibration is straightforward: Use MMSE to determine the best-fit path-loss exponent.

Two-Ray Model (Two-Ray) [30]: This model uses two path-loss exponents to capture the dual-slope signal propagation in urban environments, i.e., signal attenuates faster after a certain distance [43]. To calibrate this model, we partition sample measurements into two sets using a distance threshold. To set the distance threshold, we apply exhaustive search to identify the optimal threshold that minimizes the overall MMSE. For each set, we use a separate MMSE fitting to determine the best-fit path-loss exponent.

Terrain-Based Model (Terrain) [29]: This model uses terrain information to capture the nonuniformity of radio propagation caused by different terrains. It divides a transmitter’s

coverage area into sectors and applies a terrain-specific shadowing factor in each sector. Our calibration follows the procedure in [29]. Since we do not have terrain information (street, buildings, etc.) for MetroFi and TFA datasets, we only provide results for GoogleWiFi dataset.

Street Model (Street) [44]: This model targets urban microcell networks and assumes that signals mostly propagate along the streets, with minor reflection and/or diffraction across streets. To calibrate this model, we categorize signal propagation into three types based on the number of reflections it encounters (no reflection, one reflection, and multiple reflections). We divide measurement samples into these categories and train parameters for each propagation type separately. Like Terrain, this model requires street information, and thus can only be calibrated using the GoogleWiFi dataset.

B. 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 -30 dB (underprediction) to 30 dB (overprediction). Our key findings are as follows.

Impact of Sensor Placement: We deploy sensors to collect measurement samples. The search for optimal sensor placement is still an open problem [45], so we compare two known strategies to place sensors: *random* [29] and *intersection-based* placement. The latter is recommended for the Street model because street intersections have direct paths to nearby APs [44]. Our results show that intersection-based placement only helps the Street model, achieving 0.2 dB reduction in average prediction error. Hence, for all later results, we apply intersection-based placement only for the Street model. In addition, we observe a small variance across prediction results obtained from different sensor placement instances. This indicates that a small number of measurements are sufficient to calibrate propagation models to a reasonable level of stability.

Impact of Sampling: We calibrate our models using measurement samples collected by sensors. We vary the density of samples/sensors from 1.4 to 100 samples per km^2 , or 10 to 700 total samples for an area of size 7 km^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 yields surprising results, e.g., we find that the calibrated path-loss 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 .

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 [30]–[33]. Specifically, prediction error varies across locations and can be approximated by a zero-mean Gaussian distribution. Fig. 3 shows the probability density function (PDF) of prediction errors and its Gaussian approximation using the GoogleWiFi dataset and the Street model. The same trend holds for other datasets and models [46], and we omit those results for brevity. Table II lists the standard deviation of the prediction error under each model and dataset.

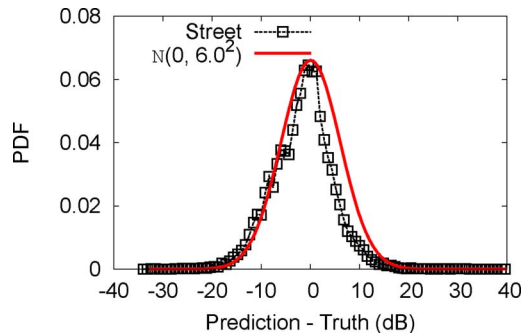


Fig. 3. Probability density distributions of prediction errors using calibrated Street model. The reference zero-mean Gaussian curve is also displayed for the calibrated model. Prediction errors approximately follow zero-mean Gaussian distributions with standard deviation of 6.

TABLE II
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

Impact of Receiver Location: When examining the correlation between prediction error and location, we observe that *all four propagation models tend to underpredict signal strength in area near the transmitter, and overpredict signal strength in area far from the transmitter.*

We illustrate this pattern in Fig. 4 by plotting the measured signal strength distribution of a randomly selected AP and levels of signal overprediction and underprediction at different locations. We define signal overprediction (underprediction) to be when the predicted signal strength is larger (smaller) than the actual value by more than 1 dB. 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 Fig. 5, we sort each AP’s measurement locations by their distances to the AP and group them into buckets of 0.05 km. For each interval, we calculate the occurrence of accurate signal prediction (absolute error ≤ 1 dB), overprediction, and underprediction. We observe that the trend is consistent across all settings.³

One possible explanation is that these propagation models still cannot fully capture how RF signals attenuate more quickly after traveling a certain distance in urban street environments [30], [43], [47]. Although both the Terrain and Street models seek to capture the nonuniform signal propagation, they still use the single-slope path-loss 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 [47]. Thus, we still observe errors for this model. This observation also motivates the need for new propagation

³While we show one propagation model for each dataset, the other models lead to the same trend [46]. We omit those results for brevity.

TABLE III
CONFLICT GRAPH STABILITY AT $\gamma = 0.9$

Propagation Model	Edge Differences				
	0	1	2	3	≥ 4
Uniform	100%	0%	0%	0%	0%
Two-Ray	61%	34%	4%	1%	0%
Terrain	37%	43%	20%	0%	0%
Street	27%	49%	23%	1%	0%

models that better capture the dual-slope effect. This is outside of the scope of this paper, and we leave it to future work.

Summary: We observe that propagation models, even after careful calibration, introduce visible but location-dependent errors in signal prediction. This naturally leads to the question: How will these signal prediction errors translate into errors in estimated conflict graphs? We explore this question next.

IV. CONFLICT GRAPH ACCURACY

We now examine the accuracy of “estimated conflict graphs” built from predicted signal maps using calibrated models. We ask the question: *What is the impact of imperfect signal strength predictions on the accuracy of 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.”

We use graph similarity 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.

A. Graph Stability Results

Before examining the accuracy of estimated conflict graphs, first we need to understand whether estimated graphs change significantly when we vary sensor locations. To examine the stability of estimated graphs against sensor placements, we use 50 sets of random sensor locations to calibrate propagation models and produce 50 graph instances. We then compare each pair of graphs and calculate the statistical distribution of edge differences across all graph pair comparison.

As shown in Table III, these graphs instances are highly similar (≤ 3 edge differences). Among different propagation models, Street model is more sensitive to changes of sensor locations because it has more dimensions and parameters in its configuration. Overall, our results demonstrate that given a sensor placement strategy, different sensor locations do not significantly affect resulting estimated graphs. Hence, we will show average results in all our later analysis.

B. 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:

- *Correct edges:* edges found in both estimated and measured conflict graphs;

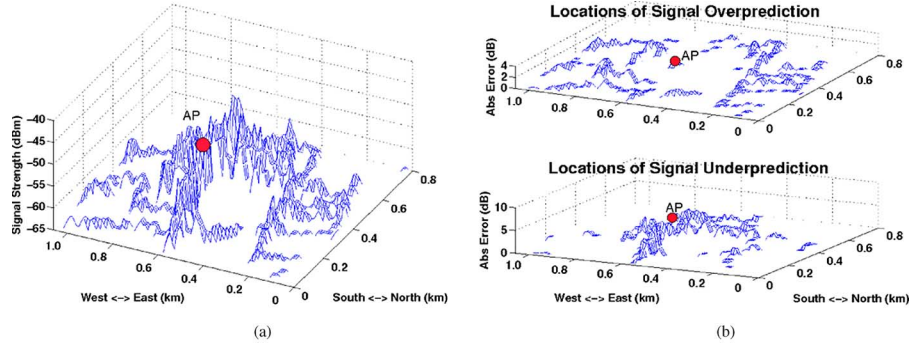


Fig. 4. (a) Measured RSS. Spatial distribution of measured signal strengths. Blank areas cannot be measured because of buildings or obstacles. (b) Prediction errors. Areas of signal overprediction and underprediction and absolute error values, using the Street model.

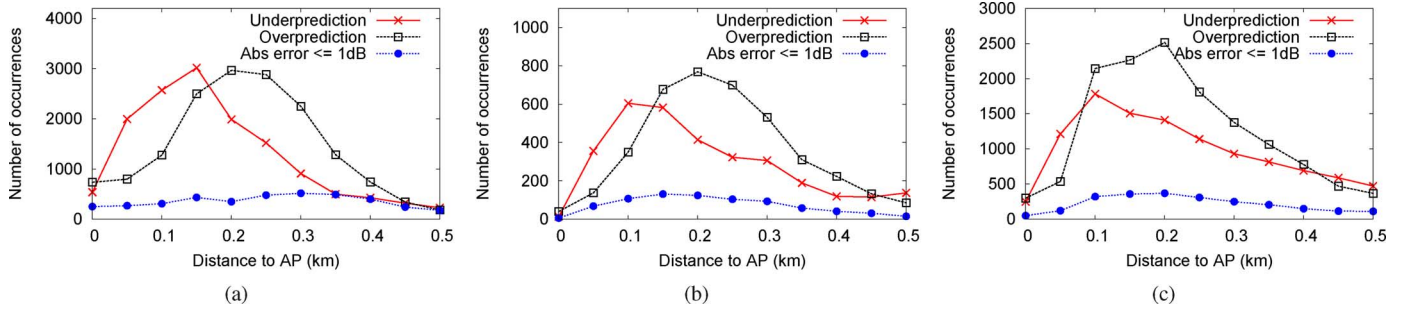


Fig. 5. As the distance to AP increases, the main prediction error gradually switches from underprediction to overprediction. Here, each data point summarizes prediction errors within a distance interval of 0.05 km. We show results with distances < 0.55 km because the number of data points above this range is insufficient to demonstrate any trends. (a) Street model, GoogleWiFi. (b) Two-Ray model, TFA. (c) Two-Ray model, MetroFi.

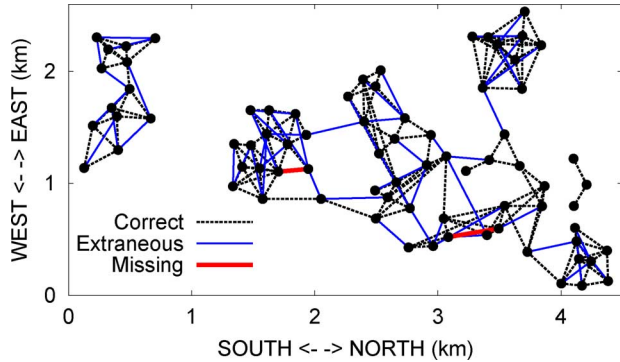


Fig. 6. Accuracy of an estimated conflict graph.

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

Fig. 6 shows an example estimated conflict graph 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 captures most edges.

We then compute *normalized edge errors* as the number of extraneous and missing edges normalized by the total number

of edges in the measured conflict graph. Fig. 7(a) shows normalized edge errors as the value of γ varies. We display normalized extraneous edges as positive values and normalized missing edges as negative values. We see that most 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 performs the best overall. This is likely because of the higher accuracy achieved by the Street model, which treats 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 fewer edge errors.

Fig. 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 more⁴ 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% ($\gamma = 1$).

We further examine how different sensor placement strategy affects graph accuracy. In Fig. 7(b), we compare random and intersection-based strategy using the Street model. We see that intersection-based placement generally provides moderate

⁴As γ grows from 0.8 to 1, the edge counts of the measured conflict graphs are 104, 132, 162, 243, 446, respectively.

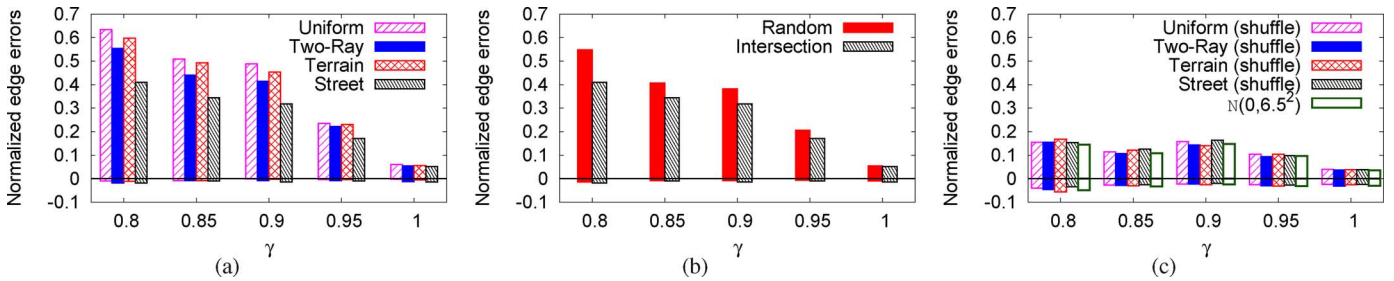


Fig. 7. (a) Varying propagation models. (b) Varying sensor placement strategies. (a) and (b) show edge errors of estimated conflict graphs, normalized by the number of edges in the measured conflict graph. Negative (positive) bars denote the normalized count of missing (extraneous) edges. (c) Modified estimated conflict graphs. (c) shows edge errors of “modified” estimated conflict graphs by removing the location dependency in prediction errors. The number of extraneous edges decreases significantly.

(15%–23%) reduction in extraneous edges. Thus overall, we conclude that the Street model with intersection-based sensor placement produces the most accurate estimated conflict graph.

C. Why Do Extraneous Edges Dominate?

The fact that extraneous edges dominate errors is attributed to two factors. The first is the location-dependent pattern of signal prediction errors (Section III). It causes underprediction of signal strength and overprediction 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 predicted signal maps, shuffles them randomly across different locations, and adds them back to the measured signal map. The second method uses synthetic prediction errors following a zero-mean Gaussian distribution with standard deviation of 6.5 and adds them to the measured signal map. Fig. 7(c) shows that these modified estimated graphs have much fewer extraneous edges than their unmodified counterparts, only 5%–15% versus 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 miss an edge than to add an edge. To miss an edge between i and j , both predicted ratios of conflict-free locations [q_i^j and q_j^i , defined by (1)] must exceed γ . In contrast, erroneously adding an edge between i and j requires only 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 prediction errors [Fig. 7(c)].

We note that these extraneous edges are not due to possible undermeasurement 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 identify those edges in the estimated conflict graph and correct them. After carefully examining our traces, we found no distinctive




Fig. 8. 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.

characteristics that distinguish extraneous edges from correct edges. For example, Fig. 8 plots the value of p_{ij} [defined by (1)] for each node pair i and j , calculated from the predicted signal strength distribution. We use different markers to separate correct and extraneous edges. We see that there is no clear distinction between the two sets.

D. Summary of Findings

Our graph accuracy analysis reveals two key findings.

- 1) *Estimated graphs are conservative.* Extraneous edge errors dominate; estimated graphs rarely miss edges (< 2%).
- 2) *Location-dependent signal prediction errors are the main cause of extraneous edges.* The location-dependent pattern in signal prediction errors causes 70+% of SINR values underpredicted, and is the main cause of extraneous edges.

V. IMPACT ON SPECTRUM ALLOCATION

After examining the accuracy of estimated conflict graphs at graph level, we now quantify the impact of these graph errors on spectrum market users. Also, given the lack of representation of accumulative interference in conflict graphs, we aim to understand how this artifact affects the quality of spectrum usage. For this, we distribute spectrum using 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.

A. Spectrum Allocation Benchmarks

To translate conflict graphs into actual spectrum allocations, we use two representative allocation algorithms: *single channel allocation* (SCA) and *multichannel allocation* (MCA). They efficiently distribute a given spectrum range across market users,




Fig. 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.

while ensuring that no conflicting users receive the same spectrum band. SCA allocates equal amount of continuous spectrum frequencies to each user. This problem reduces to the graph-coloring problem, which uses the minimal number of channels to ensure that each user receives a channel and does not conflict with another [7]. The fewer the required channels, the larger the channel bandwidth. In comparison, MCA divides the spectrum range into a number of channels. A user can receive multiple channels even if they are not continuously aligned in frequency [3]. MCA allocates channels to maximize a predefined system utility, e.g., proportional fairness [3].

We evaluate the resulting spectrum allocation based on two metrics. First, *spectrum efficiency* represents the average amount of spectrum received per market user, normalized by available spectrum. Next, we use *spectrum reliability* to capture whether each market user receives exclusive spectrum usage. A user receives exclusive usage if on each assigned channel, the percentage of receivers with actual SINRs $\geq \beta$ is no less than γ . For our purposes, we use *network-level spectrum reliability*, which is the percentage of users receiving exclusive spectrum usage.

Given a spectrum allocation, we use our full dataset to calculate actual SINRs at each user's receiver locations. The actual SINRs include interference accumulated from *all* other users on the same channel. This captures the true performance perceived by allocated users. While it is well known that conflict graphs fail to cover accumulative interference and this will affect spectrum allocations, we aim to understand the severity of this problem, and seek solutions to address it.

B. Spectrum Allocation Results

Spectrum Efficiency: Fig. 9 compares the spectrum efficiency of using the measured and estimated conflict graphs in allocating spectrum. Because extraneous conflict edges will prevent some nonconflicting 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 $\gamma = 0.8$). As γ increases, the efficiency loss goes down because extraneous edges decrease significantly. As before, the Street model outperforms the others because of its higher accuracy in signal prediction, which leads to less edge errors in the resulting 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 Fig. 10 that the reliability is not 100% for both allocation algorithms. Unless $\gamma = 1$, the measured conflict graph makes 10%–50% of market users unsatisfied. In comparison, the estimated graphs lead to more reliable spectrum usage because of extraneous edges. In this regard, extraneous edges improve spectrum reliability.

The results demonstrate that accumulative interference causes noticeable impact on spectrum usage. Because interference experienced by a receiver is the sum of signals from transmitters operating on the same frequency, 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 has lower spectrum reuse, and thus the level of accumulative interference. This effect motivates us to examine the conditions under which accumulative interference is 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 [48]. Their network is set up in an indoor environment, where rich obstacles limit signal propagation and hence the impact of accumulative interference. This begs the question: In an outdoor network, under what conditions will accumulative interference matter?

To answer this question, we first examine spatial locations of the market users with reliability violations in the GoogleWiFi dataset. We see that most of them are clustered in the area center with high density. This indicates that node density is a large contributing factor. To examine the impact of node density, we build new market configurations by sampling APs in the GoogleWiFi dataset uniformly, while keeping each AP's coverage area the same. For each new configuration, we build a conflict graph from exhaustive signal measurements and examine its reliability using the MCA allocation. Our results show that the percentage of users with reliability violations grows with 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. To use conflict graphs in practice, we must address this artifact.

VI. GRAPH AUGMENTATION

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

A. Challenges

To reduce the impact of accumulative interference, one intuitive method is to add edges into existing conflict graphs to make them more conservative. This essentially reduces the number of

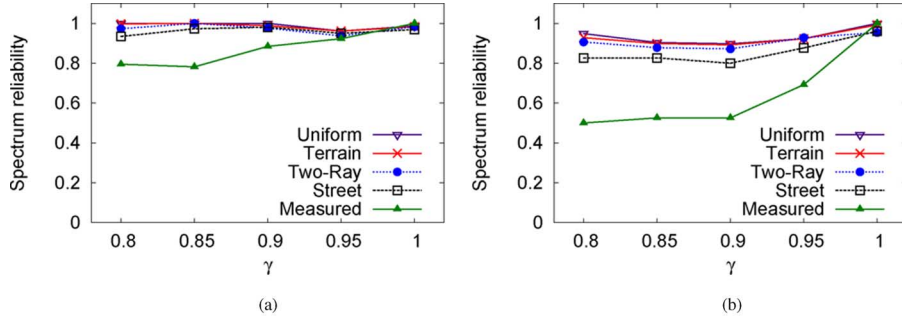


Fig. 10. Spectrum reliability when using measured and estimated conflict graphs to distribute spectrum. The reliability is between 80%–98% for the estimated graph and drops to 50% for the measured conflict graph. This indicates that the impact of accumulative interference is noticeable, and it is not captured by these conflict graphs. (a) SCA. (b) MCA.

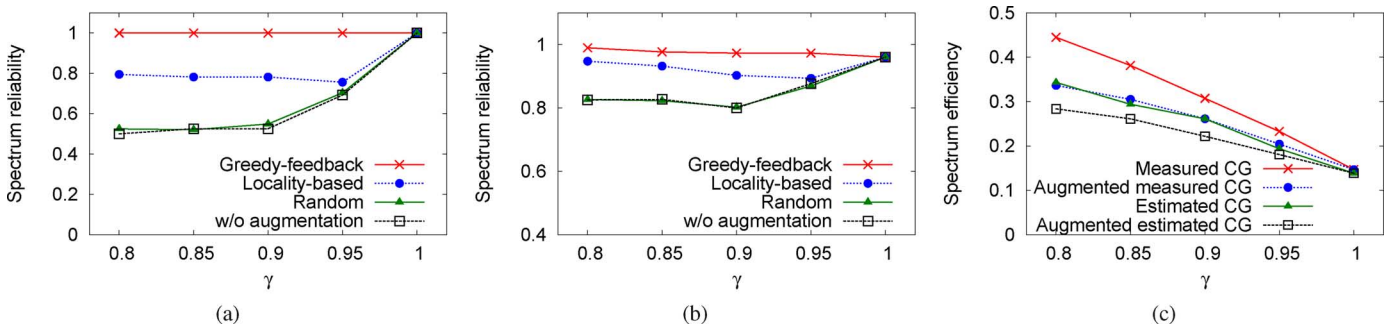


Fig. 11. Performance of graph augmentation. The estimated graphs are generated using the Street model. (a) Spectrum reliability (measured CG) and (b) spectrum reliability (estimated CG) before and after graph augmentation, using different augmentation algorithms. Greedy-Feedback outperforms the other two. (c) Spectrum efficiency before and after graph augmentation via Greedy-Feedback.

users allocated with the same spectrum channel, thus the amount of accumulative interference. However, adding 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 of this task, let us begin with two straw-man solutions. The first solution is to add edges randomly to unconnected node pairs (referred to as *Random*). A smarter alternative is to sort unconnected node pairs by their physical distances, 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) an added edge might not effectively reduce accumulative interference; and 2) it is difficult to determine the necessary number of edge additions.

B. Greedy-Feedback Graph Augmentation

We overcome the above challenges by proposing a greedy algorithm to add edges gradually and intelligently. Our 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 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 nodes i and j and repeat the

above process until all nodes have met the reliability requirement γ . The augmentation converges because as we add edges, the number of users allocated on each channel decreases monotonically. This reduces accumulative interference and improves each user’s reliability.

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 measured signal maps. In contrast, to augment estimated graphs, we estimate reliability using predicted signal maps.

C. Evaluation Results

We evaluate the effectiveness of our augmentation algorithm by comparing spectrum reliability and efficiency before and after augmentation. We use MCA allocation since it suffers more accumulative interference. We have two key observations.

Augmentation Boosts Spectrum Reliability: Our proposed greedy-feedback augmentation is highly effective to address accumulative interference. Fig. 11(a) and (b) compares three augmentation techniques on 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. We see that Greedy-Feedback significantly outperforms the others. It removes the impact of accumulative interference on the measured graph, and boosts the reliability of estimated graphs to 96+%. The reliability of estimated graphs is not 100% because its augmentation relies on imperfect signal predictions to estimate reliability. In contrast, Random has no visible improvement on reliability, while Locality-based augmentation is halfway between Random and Greedy-Feedback.

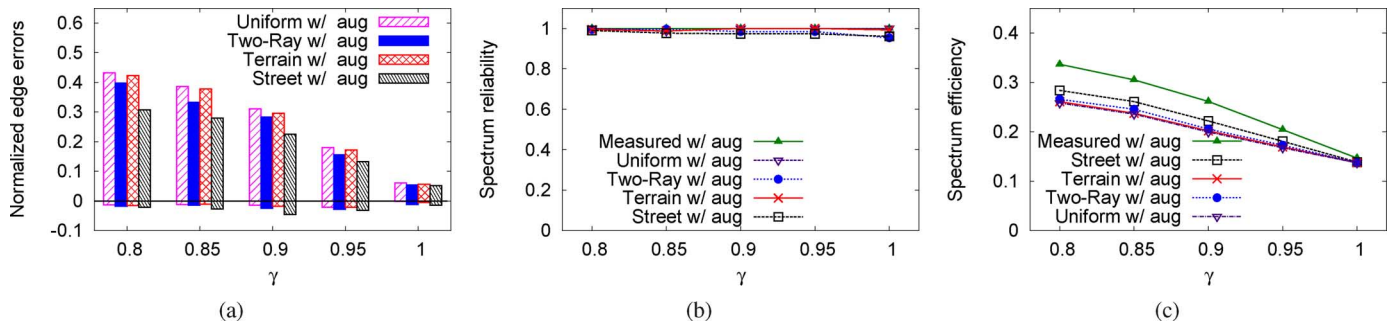


Fig. 12. Accuracy of augmented estimated conflict graphs using the four propagation models, compared to the ideal conflict graph. (a) Graph edge errors. Edge errors reduce considerably with graph augmentation. (b) Spectrum reliability. Reliability of using estimated graphs in spectrum allocation increases to 96+%. (c) Spectrum efficiency. Efficiency loss of using estimated graphs reduces to no more than 21%.

By adding edges, graph augmentation does lead to lower spectrum efficiency. From Fig. 11(c), we see that by using Greedy-Feedback, we get efficiency loss between 0%–25% for the measured graph and 0%–15% for estimated graphs. The loss for estimated graphs is lower because it adds fewer edges. We see that the proposed graph augmentation is effective against accumulative interference.

Augmentation Improves Graph Accuracy: Augmenting the measured conflict graph produces an *ideal* conflict graph, which captures real conflicts and the impact of accumulative interference. We now look at the accuracy of augmented, estimated conflict graphs relative to this ideal graph. Fig. 12(a) plots normalized edge errors of estimated graphs after augmentation. The ratios of extraneous edges reduce from 5%–60% [Fig. 7(a)] to 5%–40%, while the ratios of missing edges remain similar. This is because augmentation adds edges to the measured graph, and some added edges are already in estimated graphs. This reduces extra edge errors.

In addition, we see that augmentation also improves spectrum efficiency and reliability of estimated graphs relative to the ideal graph. Fig. 12(b) and (c) shows that the use of estimated conflict graphs achieves nearly 100% spectrum reliability and only leads to $\leq 21\%$ loss in spectrum efficiency. Overall, the Street model is the most efficient ($\leq 15\%$ efficiency loss, 96+% reliability) among the four propagation models.

VII. PRACTICAL CONSIDERATIONS

In this section, we discuss practical issues when building measurement-calibrated conflict graphs in outdoor networks.

Deploying Sensors: Building measurement-calibrated graphs relies on sensors for collecting signal measurements. Planning the sensor deployment will involve determining sensor locations, density, and sensing range, all while taking into account propagation environments, radio frequency, transmitter locations, coverage, and density. For our experiments, we conduct small-scale on-site tests to identify the saturation point where further increase in the sensor density leads to diminishing returns. In practice, optimizing sensor count and placement might become a location-specific optimization problem.

Collecting Signal Strength: Our current design assumes that each spectrum user’s transmitter periodically sends beacon signals on a specific frequency band that has identical propagation property as the other bands. Sensors extract received signal

strength values from these beacons, which are then used to calibrate propagation models and estimate conflict graphs. When frequency bands in the spectrum market have different propagation property, the system will need a beaconing scheme for each band without disrupting data transmissions.

Capturing Temporal Variations: To capture the impact of long-term temporal variations in signal strength, we need to adapt conflict graphs using periodical sensor measurements. For applications that must consider short-term signal fading, we can build the conflict graph based on the outage SINR, i.e., the bottom $x\%$ of SINR observed within a certain time period. This is an ongoing work for our study.

VIII. 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 conflict graphs that capture interference conditions between link pairs. They build conflict graphs using either a simple distance-based criterion [19], [20] or exhaustive per-link measurements in an active [8], [13]–[17], [24], [49] or passive [18], [50]–[52] manner. These link-based conflict graphs are for indoor WiFi networks where transmission links are known. They become impractical for outdoor mobile networks that spectrum markets target.

The second category builds coverage-based conflict graphs based on propagation models, either with rule-of-thumb parameters [6], [21], [22] or calibrated by on-site measurements [53]. 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, researchers have looked into other interference models such as the SINR model [10]–[12]. Most works target cellular networks and jointly optimize the placement (and transmission power) of base stations [26], [27], [34]–[36]. Our work differs in that we target spectrum markets where spectrum users come from independent entities with unplanned locations. We use conflict graphs as the interference model because they are simple and commonly used by existing spectrum allocation solutions.

While the limitations of conflict graphs have been recognized [25], [34], no study has used real-world, large-scale

measurements to examine the severity of these limitations in practice. Our work aims to bridge this gap and to further examine whether these limitations can be overcome/compensated by modifying the process of constructing conflict graphs. Recent work examines the accuracy of general interference models for small-scale networks using per-link measurements [23]. Our work was inspired by this work, yet we focus on large-scale outdoor networks where per-link measurement is infeasible.

Measurement-Calibrated Propagation Models: Measurement studies show that RF propagation models with rule-of-thumb parameters introduce large errors in signal prediction [30]–[32]. When calibrated using on-site measurements, however, these propagation models offer higher accuracy and have been used in cell planning [26], [27], interference management [33], and coverage prediction [28], [29]. Our work complements these prior works and is also inspired by prior work on measurement-calibrated models for social network graphs [54].

IX. 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 end-to-end impact on spectrum allocation. We found that 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 WiFi frequencies studied by this paper (and their nearby frequencies), our proposed techniques provide a scalable and accurate end-to-end solution for spectrum allocation.

REFERENCES

- [1] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu, “Impact of interference on multi-hop wireless network performance,” in *Proc. MobiCom*, 2003, pp. 66–80.
- [2] C. Peng, H. Zheng, and B. Y. Zhao, “Utilization and fairness in spectrum assignment for opportunistic spectrum access,” *Mobile Netw. Appl.*, vol. 11, no. 4, pp. 555–576, 2006.
- [3] L. Cao and H. Zheng, “Distributed spectrum allocation via local bargaining,” in *Proc. IEEE SECON*, 2005, pp. 475–486.
- [4] S. Gandhi, C. Buragohain, L. Cao, H. Zheng, and S. Suri, “A general framework for wireless spectrum auctions,” in *Proc. IEEE DySPAN*, 2007, pp. 22–33.
- [5] I. Katzela and M. Naghshineh, “Channel assignment schemes for cellular mobile telecommunication systems: A comprehensive survey,” *IEEE Pers. Commun.*, vol. 3, no. 3, pp. 10–31, Jun. 1996.
- [6] M. C. Necker, “Towards frequency reuse 1 cellular FDM/TDM systems,” in *Proc. MSWiM*, 2006, pp. 338–346.
- [7] S. Ramanathan, “A unified framework and algorithm for channel assignment in wireless networks,” *Wireless Netw.*, vol. 5, no. 2, pp. 81–94, 1999.
- [8] E. Rozner, Y. Mehta, A. Akella, and L. Qiu, “Traffic-aware channel assignment in enterprise wireless LANs,” in *Proc. IEEE ICNP*, 2007, pp. 133–143.
- [9] X. Zhou, S. Gandhi, S. Suri, and H. Zheng, “eBay in the sky: Strategy-proof wireless spectrum auctions,” in *Proc. MobiCom*, 2008, pp. 2–13.
- [10] L. Cao, L. Yang, X. Zhou, and Z. Zheng, “Optimus: SINR-driven spectrum distribution via constraint transformation,” in *Proc. IEEE DySPAN*, 2010, pp. 1–12.
- [11] A. Capone and M. Trubian, “Channel assignment problem in cellular systems: A new model and a tabu search algorithm,” *IEEE Trans. Veh. Technol.*, vol. 48, no. 4, pp. 1252–1260, Jul. 1999.
- [12] P. Subramanian, M. Al-Ayyoub, H. Gupta, S. R. Das, and M. M. Buddhikot, “Near-optimal dynamic spectrum allocation in cellular networks,” in *Proc. IEEE DySPAN*, 2008, pp. 1–11.
- [13] N. Ahmed, U. Ismail, and S. Keshav, “Online estimation of RF interference,” in *Proc. CoNEXT*, 2008, Art. no. 4.
- [14] N. Ahmed and S. Keshav, “SMARTA: A self-managing architecture for thin access points,” in *Proc. CoNEXT*, 2006, Art. no. 9.
- [15] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste, “DIRC: Increasing indoor wireless capacity using directional antennas,” in *Proc. SIGCOMM*, 2009, pp. 171–182.
- [16] D. Niculescu, “Interference map for 802.11 networks,” in *Proc. IMC*, 2007, pp. 339–350.
- [17] S. Rayanchu, V. Shrivastava, S. Banerjee, and R. Chandra, “FLUID: Improving throughputs in enterprise wireless LANs through flexible channelization,” in *Proc. MobiCom*, 2011, pp. 1–12.
- [18] V. Shrivastava, S. Rayanchu, S. Banerjee, and K. Papagiannaki, “PIE in the sky: Online passive interference estimation for enterprise WLANs,” in *Proc. NSDI*, 2011, p. 25.
- [19] M. Alicherry, R. Bhatia, and L. E. Li, “Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks,” in *Proc. MobiCom*, 2005, pp. 58–72.
- [20] H. Li, Y. Cheng, C. Zhou, and P. Wan, “Multi-dimensional conflict graph based computing for optimal capacity in MR-MC wireless networks,” in *Proc. IEEE ICDCS*, 2010, pp. 774–783.
- [21] L. Yang, L. Cao, and H. Zheng, “Physical interference driven dynamic spectrum management,” in *Proc. IEEE DySPAN*, 2008, pp. 1–12.
- [22] Z. Haas, J. Winters, and D. Johnson, “Simulation study of the capacity bounds in cellular systems,” in *Proc. IEEE PIMRC/WCN*, 1994, vol. 4, pp. 1114–1120.
- [23] R. Maheshwari, S. Jain, and S. R. Das, “A measurement study of interference modeling and scheduling in low-power wireless networks,” in *Proc. SenSys*, 2008, pp. 141–154.
- [24] J. Padhye, S. Agarwal, V. N. Padmanabhan, L. Qiu, A. Rao, and B. Zill, “Estimation of link interference in static multi-hop wireless networks,” in *Proc. IMC*, 2005, p. 28.
- [25] T. Moscibroda, R. Wattenhofer, and Y. Weber, “Protocol design beyond graph-based models,” in *Proc. HotNets*, 2006.
- [26] S. Hurley, “Planning effective cellular mobile radio networks,” *IEEE Trans. Veh. Technol.*, vol. 51, no. 2, pp. 243–253, Mar. 2002.
- [27] A. Mishra, *Fundamentals of Cellular Network Planning and Optimization*. New York, NY, USA: Wiley, 2004.
- [28] R. Murty, R. Chandra, T. Moscibroda, and P. Bahl, “SenseLess: A database-driven white spaces network,” in *Proc. IEEE DySPAN*, 2011, pp. 10–21.
- [29] J. Robinson, R. Swaminathan, and E. W. Knightly, “Assessment of urban-scale wireless networks with a small number of measurements,” in *Proc. MobiCom*, 2008, pp. 187–198.
- [30] A. Goldsmith, *Wireless Communications*. Cambridge, U.K.: Cambridge Univ. Press.
- [31] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott, “Experimental evaluation of wireless simulation assumptions,” in *Proc. MSWiM*, 2004, pp. 78–82.
- [32] C. Phillips, D. Sicker, and D. Grunwald, “Bounding the error of path loss models,” in *Proc. IEEE DySPAN*, 2011, pp. 71–82.
- [33] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, “Measurement-based models of delivery and interference in static wireless networks,” in *Proc. SIGCOMM*, 2006, pp. 51–62.
- [34] E. Amaldi, A. Capone, F. Malucelli, and C. Mannino, “Optimization problems and models for planning cellular networks,” in *Handbook of Optimization in Telecommunication*. New York, NY, USA: Springer, 2006.
- [35] L. Song and J. Shen, *Evolved Cellular Network Planning and Optimization for UMTS and LTE*. Boca Raton, FL, USA: CRC Press, 2010.
- [36] K. Tutschku, N. Gerlich, and P. Tran-Gia, “An integrated approach to cellular network planning,” in *Proc. Int. Netw. Planning Symp.*, 1996.
- [37] S. Deb, V. Srinivasan, and R. Maheshwari, “Dynamic spectrum access in DTV whitespaces: Design rules, architecture and algorithms,” in *Proc. MobiCom*, 2009, pp. 1–12.

- [38] M. Riback, J. Medbo, J.-E. Berg, F. Harrysson, and H. Asplund, "Carrier frequency effects on path loss," in *Proc. IEEE VTC*, 2006, pp. 2717–2721.
- [39] FCC, Washington, DC, USA, "Secondary markets in radio spectrum," 2000 [Online]. Available: <http://www.fcc.gov/events/secondary-markets-radio-spectrum>
- [40] R. Senior, "MetroFi," 2007 [Online]. Available: <http://crawdad.cs.dartmouth.edu/~crawdad/pdx/metrofi>
- [41] Rice University, Houston, TX, USA, "TFA," [Online]. Available: <http://tfa.rice.edu/measurements/>
- [42] Trapeze Networks, Pleasanton, CA, USA, "Coverage or capacity? Best use of 802.11n," Whitepaper, 2011.
- [43] E. Green, "Radio link design for microcellular systems," *Brit. Telecom Technol. J.*, vol. 8, no. 1, pp. 85–96, 1990.
- [44] A. Goldsmith and L. Greenstein, "A measurement-based model for predicting coverage areas of urban microcells," *IEEE J. Sel. Areas Commun.*, vol. 11, no. 7, pp. 1013–1023, Sep. 1993.
- [45] M. Younis and K. Akkaya, "Strategies and techniques for node placement in wireless sensor networks: A survey," *Ad Hoc Netw.*, vol. 6, no. 4, pp. 621–655, 2008.
- [46] X. Zhou, Z. Zhang, G. Wang, X. Yu, B. Y. Zhao, and H. Zheng, "Practical conflict graphs for dynamic spectrum distribution," in *Proc. SIGMETRICS*, 2013, pp. 5–16.
- [47] E. Green and M. Hata, "Microcellular propagation measurements in an urban environment," in *Proc. IEEE PIMRC*, 1991, pp. 324–328.
- [48] S. M. Das, D. Koutsonikolas, Y. C. Hu, and D. Peroulis, "Characterizing multi-way interference in wireless mesh networks," in *Proc. WiNTECH*, 2006, pp. 57–64.
- [49] L. Qiu, Y. Zhang, F. Wang, M. K. Han, and R. Mahajan, "A general model of wireless interference," in *Proc. MobiCom*, 2007, pp. 171–182.
- [50] K. Cai, M. Blackstock, M. J. Feeley, and C. Krasic, "Non-intrusive, dynamic interference detection for 802.11 networks," in *Proc. IMC*, 2009, pp. 377–383.
- [51] K.-Y. Jang, M. Carrera, K. Psounis, and R. Govindan, "Passive on-line in-band interference inference in centralized WLANs," USC, Los Angeles, CA, USA, Tech. Rep. 916, 2010.
- [52] J. Manweiler, N. Santhapuri, S. Sen, R. R. Choudhury, S. Nelakuditi, and K. Munagala, "Order matters: Transmission reordering in wireless networks," in *Proc. MobiCom*, 2009, pp. 61–72.
- [53] A. Kuurne, "Mobile measurement based frequency planning in GSM networks," M.S. thesis, Helsinki University of Technology, Espoo, Finland, 2001.
- [54] A. Sala, L. Cao, C. Wilson, R. Zablit, H. Zheng, and B. Y. Zhao, "Measurement-calibrated graph models for social network experiments," in *Proc. WWW*, 2010, pp. 861–870.



Xia Zhou received the B.S. degree from Wuhan University, Wuhan, China, in 2004, the M.S. degree from Peking University, Beijing, China, in 2007, and the Ph.D. degree from the University of California, Santa Barbara, CA, USA, in 2013, all in computer science.

She is an Assistant Professor with the Computer Science Department, Dartmouth College, Hanover, NH, USA. Her research interests include wireless networking systems and measurements, and mobile computing.

Dr. Zhou received the 2007–2011 Chancellor's Fellowship at the University of California, Santa Barbara.



Zengbin Zhang received the B.S. and M.S. degrees in electronic engineering from Tsinghua University, Beijing, China, in 2006 and 2009, respectively, and is currently pursuing the Ph.D. degree in computer science at the University of California, Santa Barbara, CA, USA.

His research interests include wireless networking, mobile computing, and social networks.



Gang Wang received the B.E. degree in electrical engineering from Tsinghua University, Beijing, China, in 2010, and is currently pursuing the Ph.D. degree in computer science at the University of California, Santa Barbara, CA, USA.

His research interests are security and privacy, mobile networks, online social networks, and crowd-sourcing systems.



Xiaoxiao Yu received the B.S. and Ph.D. degrees in electronic engineering from Tsinghua University, Beijing, China, in 2008 and 2013, respectively.

During her Ph.D. studies, she focused on data analysis and mining in wireless sensor networks under the supervision of Prof. Victor O.K. Li and Prof. Lin Zhang. She is now an Assistant Professor with Tsinghua University. Her recent research interests are transferred to elite education.



Ben Y. Zhao received the B.S. degree from Yale University, New Haven, CT, USA, in 1997, and the M.S. and Ph.D. degrees from the University of California (UC), Berkeley, CA, USA, in 2000 and 2004, respectively, all in computer science.

He is a Professor with the Computer Science Department, UC, Santa Barbara, CA, USA. He has published over 100 publications in areas of security and privacy, networked and distributed systems, wireless networks, and data-intensive computing.

Prof. Zhao has served as Program Chair for top conferences (WOSN, WWW 2013 OSN track, IPTPS, IEEE P2P) and is a co-founder and Steering Committee member of the ACM Conference on Online Social Networks (COSN). He is a recipient of the National Science Foundation's CAREER Award, *MIT Technology Review's* TR-35 Award (Young Innovators Under 35), and *Computerworld's* Top 40 Technology Innovators Award. His work has been covered by media outlets such as *The New York Times*, *The Boston Globe*, *MIT Technology Review*, and *Slashdot*.



Haitao Zheng (M'99–SM'09) received the B.S. degree (with highest honor) in electrical engineering from Xian Jiaotong University, Xian, China, in 1995, and the M.S.E.E. and Ph.D. degrees in electrical and computer engineering from the University of Maryland, College Park, MD, USA, in 1998 and 1999, respectively.

She joined the Wireless Research Lab, Bell-Labs, Lucent Technologies, Crawford Hill, NJ, USA, as a Member of Technical Staff in 1999 and moved to Microsoft Research Asia, Beijing, China, as a project lead and Researcher in 2004. Since 2005, she has been a faculty member with the Computer Science Department, University of California, Santa Barbara, where she is now a Professor. Her recent research interests include wireless systems and networking and social networks.

Dr. Zheng was named as the 2005 MIT Technology Review Top 35 Innovators under the age of 35 for her work on cognitive radios. Her work was selected by *MIT Technology Review* as one of the 10 Emerging Technologies in 2006. She also received the 2002 Bell Laboratories President's Gold Award from Lucent Bell-Labs and the 1998–1999 George Harhalakis Outstanding Graduate Student Award from the Institute of System Research, University of Maryland.