Towards Scalable Planning of Wireless Networks

Mercy O. Jaiyeola
Dept. of Computer Science and Eng.
Mississippi State University
Mississippi State, MS, USA
moj31@msstate.edu

Maxwell Young
Dept. of Computer Science and Eng.
Mississippi State University
Mississippi State, MS, USA
my325@msstate.edu

Jin Xiao
IBM Thomas J. Watson Research Center
Yorktown Heights, NY, USA
jinoaix@gmail.com

Hugh Medal
Dept. of Industrial and Systems Eng.
University of Tennessee
Knoxville, TN, USA
hmedal@utk.edu

Greg Grimes
MS State Network Services
Mississippi State University
Mississippi State, MS, USA
gtg4@msstate.edu

David Schweitzer
Dept. of Industrial and Systems Eng.
Mississippi State University
Mississippi State, MS, USA
dts48@msstate.edu

***Abstract*—Current methods for planning wireless networks rely on a mixture of on-site measurements and predictive modeling. Unfortunately, such methods can be expensive and time-consuming when it comes to planning for venues of large dimension, or those with a vast number of wireless devices.**

In this short paper, we focus on an important aspect of scalable network planning: estimating the number of source devices (e.g., access points, base stations) needed to meet traffic demands. We propose a coarse-grained approach that models aggregate demands and interference sources. Minimizing the number of source devices is shown to be NP-complete. However, our coarse-grained approach permits an integer linear program that solves for the optimum while remaining tractable; an approximation result is also derived. Preliminary experiments using QualNet and Ekahau support our approach.

I. Introduction

Current methods for planning wireless networks are not scalable. Two obstacles to scalability are:

- *Measurement Cost.* Traditional approaches require on-site measurements which are both time-consuming and expensive. Even a limited number of measurements imposes a significant cost when dealing with venues that are vast in size and/or densely populated by wireless devices.
- *Computational Intractability.* Predictive planning (i.e., modeling) approaches often optimize using ***fine-grained features*** such as known demands at precise locations, exact dimensions of rooms, furniture placement, etc. However, this optimization is often time-consuming given that the underlying problem is combinatorially hard. For large venues, such techniques may be intractable.

For massive venues – such as ballgame stadiums and IoT networks – a natural question arises: ***Is there a scalable solution to wireless network planning?*** As with many other computational challenges, new approaches are needed as the problem size greatly increases.

Here, we make progress towards one aspect of the large-scale network design problem: *scalable estimation of the number of* **sources** *(i.e., access points, base stations, beacons, etc.) needed to satisfy demands per area.* Such an estimate is useful to a system designer who needs a fast and reliable cost estimate. This is key for efficient network provisioning, which impacts network management. Later, traditional methods can be used to place devices according to planning based on fine-grained features.

978-3-903176-15-7 © 2019 IFIP

We adopt an approach often used to model complex systems in physics [5] and biology [14] known as ***coarse-grained modeling***, which is a natural compromise between accuracy and tractability. The venue representation is simplified, and consists of only coarse-grained features: flexibly-sized (but typically large) areas, each with estimates for traffic demands, distances between areas, and maximum and minimum distances between devices. Using only this knowledge, we provide a predictive approach that is tractable while capturing the important interactions between devices at a large scale.

A. Our Contributions

This short paper reports on the following preliminary results:

- A course-grained modeling approach for estimating the required number of sources in general wireless networks.
- A proof that the problem is NP-complete and an ***integer linear program (ILP)*** for our coarse-grained approach.
- An algorithm that, subject to certain assumptions, guarantees an approximation ratio.
- Preliminary experimental results using the QualNet simulator [24] and Ekahau [11] on real-world data.

B. Related Work

Prior work on network planning is extensive and, given space constraints, we briefly survey only closely-related work.

Several commercial products exist: Ekahau [11], Aerohive Networks WiFi Planning Tool [1], TamoSoft Site Survey [26], and Netspot [21]. Yet, these are aimed at fine-grained predictive planning, and they apply only to WiFi networks.

Work from the research community often uses a combination of optimization and measurement points [3], [6], [10], [16], [18], [19], [25], [25]. However, again, these results are not focusing on the challenge of scalability.

Finally, genetic algorithms, simulated annealing, greedy heuristics exist [13], [15], [17], [20], [23]. While promising, these approaches are often proposed without a proof of NP-hardness, and they make no guarantees on solution quality.

II. The Wireless Planning Problem

Modeling Coarse-Grain Features. A venue is partitioned into disjoint ***areas***, each of which corresponds to a coarse-grained feature such as a speaking hall, group of offices, etc. where users are likely to congregate and expect wireless service. For each area $\boldsymbol{a_i}$, we assume that a rough estimate of aggregate demand, $\boldsymbol{D_i}$, is known. The set of all areas is denoted by $\boldsymbol{\mathcal{A}}$.

Areas a_i and a_j are ***neighbors*** if (i) any point in a_i is within $\boldsymbol{d}_{\text{sep}}$ meters from any point in a_j, and (ii) there is no intermediate area between them. Other neighbor definitions may be used without significantly changing things; our problem formulation is flexible, and ultimately any designation of areas/neighbors is at the discretion of the system designer.

Given $\mathcal{A}$, we define an ***area graph***, $\boldsymbol{G_{\mathcal{A}}}$, where each vertex v_i corresponds to a_i, and the weight of the vertex is D_i. Edge (v_i, v_j) has weight equal to the distance between a_i and a_j.

Signal Propagation. Communication between sources and clients may be degraded by signals originating in any other area (not only neighboring areas). To model the impact of such interference under our coarse-grained approach, we use the ***signal-to-interference-plus-noise ratio (SINR)*** [2], [28] for each area: $SINR(a_i) = \frac{P}{N+I}$ where $\boldsymbol{P}$ is the incoming signal power to a client in a_i, $\boldsymbol{N}$ is the floor noise, and $\boldsymbol{I}$ is the ***interference*** from other sources. More specifically, for a source s transmitting at P Watts, and for a ***path loss exponent*** $\boldsymbol{\alpha}$, the transmission is received by a client r if the following holds that $SINR(r)$:

$$\frac{P/(\delta_s^r)^\alpha}{N + \sum_{s' \neq s} P/(\delta_{s'}^r)^\alpha} \approx \frac{1/(\delta_s^r)^\alpha}{N + \sum_{s' \neq s} 1/(\delta_{s'}^r)^\alpha} \geq \beta$$

where $\boldsymbol{\beta} > 0$ is the ***SINR threshold***, and δ_s^r is the distance between a sender s and receiver r. In our coarse-grained modeling approach, we do not assume full knowledge of δ_s^rs, so our distance values are chosen conservatively as discussed in Section II-A.

In this short paper, we assume transmit power is homogenous, and that the floor noise is small relative to the total interference I. This is plausible in many settings given that N is often small (typically, -90 dBm to -100 dBm). We keep N in our ILP formulation, but this simplifies our model as the SINR power terms cancel in the numerator and denominator.

In practice, channel quality correlates with the degree to which β is exceeded. However, we may consider β to be set conservatively; thus, exceeding β implies high channel quality.

Overlapping Coverage. Clients can be arbitrarily located within an area a_i. Some will be in close proximity to a source in a_i, while others will be at the edges of a_i and closer to a source in a neighboring area. Given this, for each area a_i, the aggregate demand may be met in part by sources in a_i and in part by sources in a neighboring area (but not by sources in non-neighboring areas).

For any a_i, we assume that sources in a_i are on different channels. Note that this is plausible in WiFi networks, for example, given the use of channels in both the 2.4 GHz and 5 GHz bands. We emphasize that sources in any other areas (including neighbors) may share the same channel as any source in a_i, thus causing inteference.

Decision Problem. We formally define the following ***Wireless Planning (WIPLAN)*** decision problem. As input, we are given an area graph $G_{\mathcal{A}}$. The decision question is: *Can all demands be satisfied by some placement of s sources?*

A. ILP Formulation

Table I defines the parameters and variables for our mathematical formulation of WIPLAN.

Table I. Parameters and variables in our mathematical program.

Parameter	Definition
$\boldsymbol{d_j^i}$	*max.* distance from a source in a_j to a client in a_i.
$\boldsymbol{\ell_j^i}$	*min.* distance from a source in area a_j to a client in a_i.
$\boldsymbol{D_i}$	aggregate demand of an h-fraction of clients in area i (Mbps).
$\boldsymbol{C_i}$	channel capacity of a source in area a_i (Mbps).
$\boldsymbol{N}$	floor noise level (Watts).
$\boldsymbol{\beta}$	SINR threshold (unit-less).
$\boldsymbol{m}$	maximum number of sources in any area.
Variable	**Definition**
$\boldsymbol{s_i^j}$	excess fraction of service in area a_i that can be offered to a neighboring area a_j.
$\boldsymbol{x_i}$	number of sources in area a_i, where $0 \leq x_i \leq m$.
$\boldsymbol{y_i^j}$	has value 1 if source in area a_i can provide service to area a_j; that is, β is exceeded. Otherwise, the value is 0.

Parameter d_j^i (ℓ_j^i) denotes the max. (min.) distance from a source in a_j to a client in a_i. Note d_j^i appears in the numerator of the SINR constraints, while ℓ_j^i appears in the denominator (see Figure 1). This makes it more challenging to satisfy the SINR constraints, but is more conservative.

Initial Formulation. For clarity, we begin by describing a simple but non-linear mathematical program that captures the essentials of wireless coverage; this is presented in Figure 1.

For any area a_i, let $\boldsymbol{\mathcal{B}_i}$ denote the ***set of bordering areas*** to $a_i \in \mathcal{A}$. Constraint 2 enforces that the demand in each area $a_i \in A$ is met by a combination of the sources in a_i and those in neighboring areas $a_j \in \mathcal{B}_i$.

Constraint 3 ensures that a contribution $\sum_{a_j \in \mathcal{B}_i} s_i^j y_i^j D_j$ in Constraint 2 to neighbor a_j must come from sources in a_i.

Constraint 4 models whether the sources in a_i exceed the SINR threshold β. Observe that the equation can *always* be satisfied by setting $y_i^i = 0$; this interacts with the equations specified by Constraint 2 and implies that no demand is satisfied by sources in a_i, since $y_i^i x_a = 0$.

Constraint 5 is analogous to Constraint 4, but it addresses demand that can be met by each neighboring area.

Obtaining a Linear Program. The program in Figure 1 is non-linear. To linearize this program (see Figure 2), we create a proxy variable z_{jk}^i for the product $y_j^i x_k$; equality is enforced by Constraints 14 – 17. To see this, consider the case where $y_j^i = 1$, we must have $z_{jk}^i = x_k$; this is enforced by Constraint 14 and 16 since they force $z_{jk}^i \leq x_k$ and $z_{jk}^i \geq x_k - (1 - y_j^i)\, m = x_k$, respectively. Conversely, if $y_j^i = 0$, then $z_{jk}^i = 0$; this is enforced by Constraint 15 and 17 since they imply $z_{jk}^i \leq y_j^i\, m = 0$ and $z_{jk}^i \geq 0$, respectively.

Similarly, we create a proxy variable f_i^j for $s_i^j\, y_i^j$. Equality through Constraints 18-21 is proved similarly.

Upload and Download Demands. Our final ILP separates constraints for whether (i) a source in a_i can provide *download* service to clients in a_j, and (ii) a source in a_i can provide *upload* service to clients in a_j. This is accomplished via analogs to Constraints 10 – 13 where, for the SINR constraints, we use assume the necessary max. and min. distance information. Given space constraints, we omit these in our presentation.

III. NP-COMPLETENESS

We give a reduction from **DOMINATING SET (DOMSET)** where the input is a graph G and the descision question is whether there exists a set $\mathcal{DS}$ of b vertices such that all vertices either belong to $\mathcal{DS}$ or have a neighbor in $\mathcal{DS}$.

$$\min \sum_{a_i\in\mathcal{A}} x_i \quad (1)$$

s.t.

$$C_i x_i y_i^i - \sum_{a_j\in\mathcal{B}_i} s_i^j y_i^j D_j + \sum_{a_j\in\mathcal{B}_i} s_j^i y_j^i D_i \ge D_i \quad \forall a_i \in \mathcal{A} \quad (2)$$

$$C_i x_i y_i^i \ge \sum_{a_j\in\mathcal{B}_i} s_i^j y_i^j D_j \quad \forall a_i \in \mathcal{A} \quad (3)$$

$$\beta\, y_i^i \le \frac{(1/d_i^i)^\alpha}{N + \sum\limits_{a_k\in\mathcal{A}\setminus\{a_i\}} x_k/(\ell_k^i)^\alpha} \quad \forall a_i \in \mathcal{A} \quad (4)$$

$$\beta\, y_j^i \le \frac{(1/d_j^i)^\alpha}{N + \sum\limits_{a_k\in\mathcal{A}\setminus\{a_j\}} x_k/(\ell_k^i)^\alpha} \quad \forall a_i \in \mathcal{A} \text{ and } a_j \in \mathcal{B}_i \quad (5)$$

$$0 \le s_j^i \le 1 \quad (6)$$

$$y_i^j \in \{0,1\} \quad \forall a_i \in \mathcal{A} \text{ and } a_j \in \mathcal{B}_i \quad (7)$$

$$x_i \in \{0,\ldots,m\} \quad \forall a_i \in \mathcal{A} \quad (8)$$

Figure 1. A non-linear mathematical program for WIPLAN.

We consider a subset of WIPLAN instances with the following settings: download demand $D_i = 1$ (set upload demands to 0), $\ell_k^i = |A|\, m\, \beta$, $N = 0$, $\alpha \ge 2$, $d_j^i \le 1$ for all i, j, k. Now, the SINR constraints are met: each z_{jk}^i cannot exceed m, and there are less than $|A|$ terms in the SINR-equation sum; thus, the left-side sum is at most $|A|\, m\, \beta/(|A|\, m\, \beta)^\alpha \le 1/(|A|\, m\, \beta)$ while the right side is $1/(d_j^i)^\alpha \ge 1$.

We also set C_i for each $a_i \in \mathcal{A}$ equal to $\deg(v_i) + 1$ where $\deg(v_i)$ is the degree of vertex v_i in the area graph $G_\mathcal{A}$. This choice is motivated later in our proof of Theorem 1.

WIPLAN instances with these settings are referred to as ***relaxed***, and we will show that the set of such instances is NP-complete (implying that WIPLAN is NP-complete).

Theorem 1. WIPLAN *is NP-complete.*

Proof. Note: WIPLAN $\in$ NP because, given a solution for the n areas, it is verified in polytime through $O(n^2)$ constraints.

A DOMSET instance corresponds to a relaxed WIPLAN instance as follows. Each vertex $v_i \in G$ for DOMSET corresponds to $a_i \in \mathcal{A}$. Also, every neighbor v_j of v_i corresponds to area $a_j \in B_i$. From this, the resulting area graph $G_\mathcal{A}$ for WIPLAN is created with a topology that matches G. The parameter settings are those of a relaxed instance.

Consider a dominating set $\mathcal{DS}$ of size b. For each $v_i \in \mathcal{DS}$, set $x_i = 1$, $y_i^i = 1$ (so $z_{ii}^i = 1$), and $f_i^j = 1$ for $a_j \in B_i$. For every vertex $v_i \notin \mathcal{DS}$, set $z_{ii}^i = 0$ and $f_i^j = 0$. Since $C_i = \deg(v_i) + 1$ for each area $a_i \in \mathcal{A}$, the demand Constraints 10 and 11 are satisfied. To see this, note that if $v_i \in \mathcal{DS}$, this capacity is sufficient for contributions of at most $D_j = 1$ to all $\deg(v_i)$ neighbors while still leaving 1 to satisfy D_i. Else, if $v_i \notin \mathcal{DS}$, then a_i can meet its demand with help from its neighbors (each of which donates 1). Therefore, we have a solution to our relaxed WIPLAN instance that uses b sources.

Now, consider a solution to a relaxed WIPLAN instance with b sources. If $x_i \ge 1$ (note this implies $z_{ii}^i \ge 1$), put v_i in $\mathcal{DS}$. Thus, $|\mathcal{DS}| \le b$. To verify that $\mathcal{DS}$ dominates, assume otherwise; that is, there exists some v_i in G such that $v_i \notin \mathcal{DS}$ and none of its neighbors are either. Thus, the corresponding Constraint (10) cannot be satisfied since there is no source in a_i, nor in any of its neighbors; that is, D_i will not be met, and this contradicts the assumption of a WIPLAN solution. $\square$

$$\min \sum_{a_i\in\mathcal{A}} x_i \quad (9)$$

s.t.

$$C_i z_{ii}^i - \sum_{a_j\in\mathcal{B}_i} f_i^j D_j + \sum_{a_j\in\mathcal{B}_i} f_j^i D_i \ge D_i \quad \forall a_i \in A \quad (10)$$

$$C_i z_{ii}^i \ge \sum_{a_j\in\mathcal{B}_i} f_i^j D_j \quad \forall a_i \in A \quad (11)$$

$$N\beta\, y_i^i + \beta \sum_{a_k\in\mathcal{A}\setminus\{a_i\}} z_{ik}^i/(\ell_k^i)^\alpha \le \left(\frac{1}{d_i^i}\right)^\alpha \quad \forall a_i \in \mathcal{A} \quad (12)$$

$$N\beta\, y_j^i + \beta \sum_{a_k\in\mathcal{A}\setminus\{a_j\}} z_{jk}^i/(\ell_k^i)^\alpha \le \left(\frac{1}{d_j^i}\right)^\alpha \forall a_j \in \mathcal{B}_i \quad (13)$$

$$z_{jk}^i \le x_k \quad \forall a_i, a_k \in \mathcal{A} \text{ and } a_j \in \mathcal{B}_i \quad (14)$$

$$z_{jk}^i \le y_j^i\, m \quad \forall a_i, a_k \in \mathcal{A} \text{ and } a_j \in \mathcal{B}_i \quad (15)$$

$$z_{jk}^i \ge x_k - (1 - y_j^i) m \quad \forall a_i, a_k \in \mathcal{A} \text{ and } a_j \in \mathcal{B}_i \quad (16)$$

$$z_{jk}^i \ge 0 \quad \forall a_i, a_k \in \mathcal{A} \text{ and } a_j \in \mathcal{B}_i \quad (17)$$

$$f_i^j \le s_i^j \quad \forall a_i \in \mathcal{A} \text{ and } a_j \in \mathcal{B}_i \quad (18)$$

$$f_i^j \le y_i^j \quad \forall a_i \in \mathcal{A} \text{ and } a_j \in \mathcal{B}_i \quad (19)$$

$$f_i^j \ge s_i^j - (1 - y_i^j) \quad \forall a_i \in \mathcal{A} \text{ and } a_j \in \mathcal{B}_i \quad (20)$$

$$f_i^j \ge 0 \quad \forall a_i \in \mathcal{A} \text{ and } a_j \in \mathcal{B}_i \quad (21)$$

$$0 \le s_i^j \le 1 \quad (22)$$

$$y_j^i \in \{0,1\} \quad \forall a_i, a_j \in \mathcal{A} \quad (23)$$

$$x_i \in \{0,\ldots,m\} \quad \forall a_i \in \mathcal{A} \quad (24)$$

Figure 2. ILP capturing basic features of WIPLAN.

IV. APPROXIMATION ALGORITHM

While solving exactly using a tractable ILP is ideal, we report another preliminary result that may be useful for extremely large instances: an algorithm, GREEDYPLAN, that takes in $G_\mathcal{A}$ with max. degree Δ, and yields an approximation guarantee for WIPLAN if the assumptions below hold.

- ***A1.*** For any area a_i, sources in a_i and its neighboring areas are on different channels. Tentatively, this is plausible; for example, by using both the 2.4 GHz and 5 GHz bands for WiFi, and noting Δ will likely be small in practice.
- ***A2.*** Given that N is typically small relative to I, we formalize this as $N \le \epsilon\, I$ for a small constant $\epsilon > 0$.
- ***A3.*** The ratio $\boldsymbol{r} = \min\{d_{\text{sep}}/d_i^i, d_{\text{sep}}/d_j^i\} \ge ((1+\epsilon)(\beta(1+\Delta)m))^{1/\alpha}$, for all areas a_j that are not a neighbor of a_i.[1] Loosely, a client in a_i is closer by an r-factor to sources in its own area, or in a neighboring area than it is to a source in any other (non-neighboring) area.

  In many settings, **$\boldsymbol{\Delta}$ *will be small***. Consider $\alpha = 3.5$, $\epsilon = 0.1$, $m = 4$, $\Delta = 3$, and a "typical" $\beta = 10$ [9] (coincidentally, $\beta = 10$ translates to 10 dB $= 10\log_{10}(\beta)$); this yields a small $r \approx 4.4$.

Due to space limitations, we sketch the algorithm execution. A vertex whose demand (i) is not met is colored ***white***; (ii) is met, but has at least one (white) neighbor whose demand is not, is colored ***gray***; and (iii) is met, along with all neighbors, is colored ***black***. In each iteration of its execution, GREEDYPLAN finds a gray or white vertex v_i with the largest capacity (arbitrarily break any ties) and satisfies the demand for v_i and the demand of all of its neighbors by placing sources

[1]While it may appear that a large α helps us, note that N will swamp I for large α, and so the SINR will not exceed β. Hence, the assumption fails to hold if our earlier assumption $N \le \epsilon I$ does not hold.

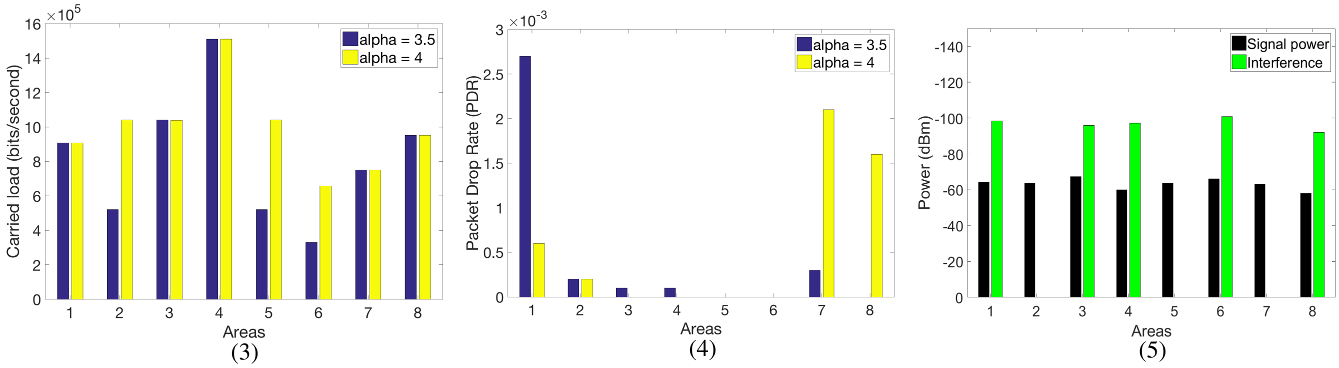

Figures 3-5. Average carried load, PDR, signal strength, and interference for clients, respectively.

in the corresponding area a_i. Then, v_i is colored black, and the neighbors' colors are updated.

Let OPT denote the optimum solution for WIPLAN. We omit the proof of our result below in this short paper.

Theorem 2. *Given assumptions A1 – A3, the number of sources assigned by* GREEDYPLAN *is* $O(\Delta^2) \cdot OPT$.

V. Preliminary Experimental Validation

We experiment with real-world data on a section of a football stadium from which distances and aggregate demands are inferred. There are 8 areas totaling roughly 3500 meters squared (37000 square feet). Each area has an aggregate demand between 1200 and 2300 Mbps corresponding to an average of 0.5 Mbps per client; these are unchanged in each experiment. Capacities differ per experiment, but are always between 300 and 400 Mbps in rough correspondence to the theoretical maximum of IEEE 802.11n using both the 2.4 GHz and 5 GHz bands. We use $N = -96$ dBm, and (a pessimistic) $\beta = 20$ (≈ 13 dB).

From these parameter settings, we construct our ILP and solve it using Gurobi [12]. Experiments use $\alpha = 3.5$ and $\alpha = 4$ to model the significant attenuation which arises in indoor environments such as office buildings and stadiums [22].

The network simulator QualNet [24] is used to evaluate our solutions. Sources are placed uniformly along the perimeter of each area, and clients send and receive at a constant bit rate via sources in a scenario lasting 2 (simulated) hours.

The automated AP placement and configuration (i.e., predictive planning) functionality of Ekahau [11] provides another point of comparison. Using an AutoCAD [4] layout, we set cinderblock walls, Cisco AP1040s [8] (using 2.4 GHz and 5 GHz with IEEE 802.11n) as sources (aligning with our capacities), and generic laptops as clients. Note that the AP capacities are fixed (unlike our ILP trials); therefore, Ekahau yields a single recommendation for our data set.

A. Our Initial Results

Due to space constraints, we summarize our preliminary findings. The ILPs were solved quickly; always less than 0.25 seconds. Solutions were stable, indicating 28 sources in total for almost all settings; and always in the range 26 to 31. By comparison, Ekahau recommended 25 sources for our data set.

Figures 3-5 illustrate a solution evaluation with QualNet for one run. The per-client average ***carried load*** (effective system throughput [7], [27]), ***packet drop rate (PDR)***, ***signal power***, and ***interference*** were reported by QualNet.

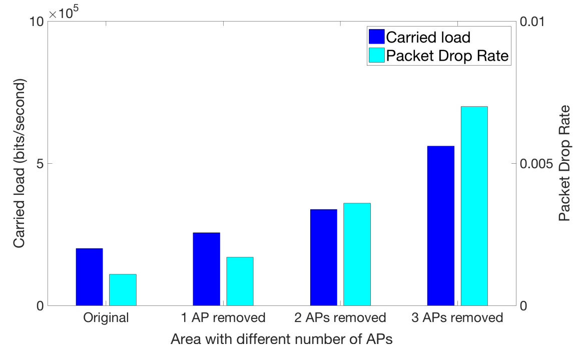

Figure 6. Carried load and PDR as sources are removed in an area.

Figure 3 illustrates an average carried load per client ranging from 0.4 to 1.5 Mbps for both $\alpha = 3.5$ and 4. This seems reasonable given that the average throughput per client to meet the specified download demand per area is 0.5 Mbps.

Figure 4 shows the average PDR for clients in each area. Notably, these PDRs are small, all less than 0.003. This is reassuring as it implies that the recommended number of sources from solving our WIPLAN instance allows for clients and sources to successfully communicate, and that the impact of interference (captured by SINR constraints) is tolerable. This is supported by Figure 5 which illustrates the average signal strength versus average interference values.

These results suggest that the number of sources is sufficient, but could we use fewer? To examine this, we experiment using less sources; Figure 6 depicts this for an area which was originally assigned 7 APs. We observe that PDR increases by a factor of ≈ 7, implying a performance degradation with less than the recommended number of sources. As a sanity check, we see that the carried load increases since fewer sources exist to handle the (unchanged) demand from clients.

VI. Future Work

While further research is needed, our coarse-grained modeling appears reasonable given the QualNet and Ekahau results. Our preliminary results suggest that the WIPLAN formulation does not lead to a grossly under or over-provisioned network.

We plan refine our coarse-grained approach on larger-scale scenarios, using more real-world data. Our approach is meant to be general; in addition to WiFi deployments, we plan to address IoT networks. Finally, we plan to evaluate the performance of our approximation algorithm.

References

[1] Aerohive. Wi-Fi planning tool, 2018. www.aerohive.com/planner/.

[2] J. G. Andrews, R. K. Ganti, M. Haenggi, N. Jindal, and S. Weber. A primer on spatial modeling and analysis in wireless networks. *IEEE Communications Magazine*, 48(11):156–163, 2010.

[3] Leena Arya, S.C. Sharma, and Millie Pant. Performance analysis of indoor positioning system. *International Journal of Advanced Computer Science and Applications*, 1(4), 2010.

[4] Autodesk. Autocad, 2018. www.autodesk.com/products/autocad/overview.

[5] Riccardo Baron, Daniel Trzesniak, Alex Vries, Andreas Elsener, Siewert Marrink, and Wilfred van Gunsteren. Comparison of thermodynamic properties of coarse-grained and atomic-level simulation models. 8:452–61, 02 2007.

[6] Roberto Battiti, Mauro Brunato, and Andrea Delai. Optimal wireless access point placement for location-dependent services. 2003.

[7] Cisco. Traffic analysis. www.cisco.com/c/en/us/td/docs/ios/solutions_docs/voip_solutions/TA_ISD.html, 2001.

[8] Cisco. Cisco aironet 1040 series access points data sheet, 2018. www.cisco.com/c/en/us/products/collateral/wireless/aironet-1140-series/data_sheet_c78-609338.html.

[9] N. Ehsan and R. L. Cruz. On the optimal sinr in random access networks with spatial reuse. In *2006 40th Annual Conference on Information Sciences and Systems*, pages 938–944, 2006.

[10] A. Eisenblatter, H.-F. Geerdes, and I. Siomina. Integrated access point placement and channel assignment for wireless LANs in an indoor office environment. In *2007 IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks, WOWMOM*, number 2007 IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks, WOWMOM, 2007.

[11] Ekahau. Wireless design, 2018. www.ekahau.com.

[12] Gurobi. http://www.gurobi.com/. Gurobi optimization: The fastest mathematical programming solver, 2018. http://www.gurobi.com/.

[13] Árpád Huszák, Győző Gódor, and Károly Farkas. Investigation of WLAN access point placement for indoor positioning. In Róbert Szabó and Attila Vidács, editors, *Information and Communication Technologies*, pages 350–361, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

[14] Helgi I. Ingólfsson, Cesar A. López, Jaakko J Uusitalo, Djurre H. de Jong, Srinivasa Murthy Gopal, Xavier Periole, and Siewert J. Marrink. The power of coarse graining in biomolecular simulations. In *Wiley interdisciplinary reviews. Computational molecular science*, 2014.

[15] Zhong Ji, T. K. Sarkar, and Bin-Hong Li. Methods for optimizing the location of base stations for indoor wireless communications. *IEEE Transactions on Antennas and Propagation*, 50(10):1481–1483, 2002.

[16] M. Kamenetsky and M. Unbehaun. Coverage Planning for outdoor wireless LAN systems. In *INTERNATIONAL ZURICH SEMINAR ON BROADBAND COMMUNICATIONS*, page 49, 2002.

[17] Seungseob Lee, SuKyoung Lee, Kyungsoo Kim, and Yoon Hyuk Kim. Base station placement algorithm for large-scale lte heterogeneous networks. *PLOS ONE*, 10:1–19, 10 2015.

[18] Y. Lin, W. Yu, and Y. Lostanlen. Optimization of wireless access point placement in realistic urban heterogeneous networks. In *2012 IEEE Global Communications Conference (GLOBECOM)*, pages 4963–4968, 2012.

[19] Y. Lin, W. Yu, and Y. Lostanlen. Optimization of wireless access point placement in realistic urban heterogeneous networks. In *2012 IEEE Global Communications Conference (GLOBECOM)*, pages 4963–4968, 2012.

[20] K. Maksuriwong, V. Varavithya, and N. Chaiyaratana. Wireless LAN access point placement using a multi-objective genetic algorithm. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, volume 2, pages 1944–1949, 2003.

[21] NetSpot. Wi-fi site surveys, analysis, troubleshooting, 2018. www.netspotapp.com.

[22] Charalampos Papamanthou, Franco P. Preparata, and Roberto Tamassia. Algorithms for location estimation based on rssi sampling. In Sándor P. Fekete, editor, *Algorithmic Aspects of Wireless Sensor Networks*, pages 72–86, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.

[23] N. F. Puspitasari, H. A. Fatta, and F. W. Wibowo. Implementation of greedy and simulated annealing algorithms for wireless access point placement. In *2015 3rd International Conference on Artificial Intelligence, Modelling and Simulation (AIMS)*, pages 165–170, 2015.

[24] Scalable Network Technologies. QualNet network simulation software. https://web.scalable-networks.com/qualnet-network-simulator-software, 2018.

[25] S. Sendra, J. Lloret, C. Turro, and J.M. Aguiar. IEEE 802.11a/b/g/n short-scale indoor wireless sensor placement. *International Journal of Ad Hoc and Ubiquitous Computing*, 15(1-3):68–82, 2014.

[26] Tamosoft. Tamograph site survey, 2018. www.tamos.com/products/wifi-site-survey/.

[27] David Tipper.http://www.pitt.edu/~dtipper/2720/2720_Slides4.pdf. Fundamentals of cellular networks.

[28] M. Z. Win, P. C. Pinto, and L. A. Shepp. A mathematical theory of network interference and its applications. *Proceedings of the IEEE*, 97(2):205–230, 2009.