Exploiting Overlap for Provisioning of Access Points in Wireless Networks

Purushottam Kulkarni and **Prashant Shenoy**
Department of Computer Science
University of Massachusetts
Amherst, MA 01003
{purukulk,shenoy}@cs.umass.edu

Abstract

802.11–based wireless LANs have become commonplace in home and office environments. In this paper, we address the problems of access point provisioning in congested hot spots and multi-hop mesh networks. We argue that both problems are closely related and show that they are NP–Complete. We propose a single greedy heuristic that exploits the notion of overlapping placement of access points to address both problems. A LP relaxation approach is also proposed to solve the provisioning problem in congested hotspots. An experimental evaluation using synthetic user distributions and a real office environment shows that exploiting overlap and the presence of multiple channels can be an effective approach for alleviating hotspots and extending coverage. With uniform and skewed distribution of end-hosts, a small addition to the minimum budget required gave benefits of upto 65% and 85% respectively.

1 Introduction

Wireless local area networks based on the 802.11 standard have become commonplace in office buildings, airports, hotels, shopping malls and homes. A typical wireless network consists of a one more access points that serve as gateways between the wired and wireless networks and provide network access to users within its transmission range. Deploying a wireless network involves two questions: how many access points to place and where? Since each access point has a fixed transmission range and a finite bandwidth that is shared among its users, the placement and number of access points is typically governed by the physical area that needs to be covered and the distribution of users in that area.

Several issues can complicate the placement of access points in a wireless network. First, the distribution of users in an environment may be skewed. For instance, the user density will be typically higher in conference rooms and class rooms than in a cluster of office cubicles. If a large number of users are closely clustered in a small physical area, all users will connect to the closest access point and share its bandwidth. If the number of users is large, the access point will overload and provide low bandwidth and poor performance to its users. Further, some access points impose limits on the number of connected users or on the number of DHCP addresses that are handed out, and may deny connectivity altogether when the number of users exceeds the maximum threshold. The wireless network must be designed to prevent such *hotspots* and the resulting congestion.

Second, it may not be possible to place wired access points—access points that are connected to a wired network—in each desirable location. For instance, existing cabling within a building may impose constraints that preclude network cables in a particular location where an access point needs to be placed. This may especially true when access points need to be placed outside a building for outdoors coverage. Another example is that of a community wireless network, also referred to as mesh networks, where a small number of users wirelessly share their broadband network connections among homes in their community. In this scenario, only a few access points are connected directly to the wired network. A multi-hop wireless network is created where each access point is within transmission range of other access points and has a path to a wired access point via some number of hops. In effect, access points act as relays on a multi-hop path to a wired access point in addition to providing wireless coverage. Given a fixed number of wired access points, the placement of relays is crucial for extending the range and reach of the wireless network to all users in the community.

In this work, we use the notion of *overlap* to address both problems: the placement of access points in congested hotspots and the placement of relays in multi-hop wireless networks. Congestion caused by hot spots can be alleviated by placing multiple access points with overlapping coverage in areas with high user density. Overlapping access points can operate on non-interfering frequencies and share the user load, thereby increasing the performance seen by each user. Similarly, by judiciously placing access points so that each access point overlaps with atleast one other access point, a multi-hop wireless network can be created to provide coverage in areas where deploying wired access points is infeasible.

In general, the placement of access points in a wireless network is well studied. Numerous studies have focused on the problem of placing base stations in a cellular network [10, 13]. Cellular networks are designed to minimize “dead zones” and minimize call blocking probability—the placement of multiple overlapping base stations is not an issue in these networks. Hotspot congestion has been studied explicitly in [5], where an admission control policy is used to guarantee a certain bandwidth to each user; the placement problem was not explicitly studied in this work. Commercial products are available, for instance from Cisco, for deploying multiple access points in a single location. However, the issue of how many overlapping access points to deploy and where still merits research attention. Finally, a recent work studied the placement of wired taps in a multi-hop wireless (mesh) network [14]. This work assumes that the placement of access points is *given* and proposes techniques for optimal placement of wired connections in this wireless network. We study the exact opposite problem—we assume that the the location of wired access points is given and focus on the placement of relays to extend the reach of the multi-hop wireless network.

In this paper, we address the problem of access point provisioning in congested hot spots and multi-hop mesh networks. We argue that both problems are closely related and show that the placement problem is NP-Complete. We propose a single greedy heuristic to address both problems—our heuristic exploits the notion of overlapping placement of access points to address these problems. A LP-relaxation technique is also proposed to address the load balancing problem. Our placement techniques employ a three step approach. In the first step, an optimal placement of access points to provide coverage in all desirable location is determined. In the second step, additional (overlapping) access points are placed in regions of high user density. The final step intelligently assigns operating channels (frequencies) to access points. For instance, overlapping access points in hotspots are assigned distinct frequencies to maximize the total available channel bandwidth, while non overlapping access points can be assigned the same channel. Similarly, all relays on a path to a wired access point are assigned the same frequency, while access points on overlapping paths are assigned different frequencies. We evaluate the effectiveness of our techniques under different user density distributions as well as for placement of access points in a real office building. Our results show that for uniform and skewed distribution of end-hosts, a small number of additional overlapping access points can reduce load by as much as 65% and 85%, respectively.

The rest of the paper is organized as follows: Section 2 presents some background and the problem formulation. The hardness results of the problems and the proposed approximate algorithms are presented in Section 3. We evaluate our solutions in Section 4. Section 5 presents related work and Section 6 presents our conclusions.

2 Background and Problem Formulation

A wireless network is assumed to consist of a certain number of access points (also referred to as base stations). Each access point (AP) operates on a certain frequency (or channel) and has a certain range. End-hosts within transmission range of the access point can *associate* with it and communicate with other end-hosts via the access point. If the node moves out of range of the AP, it must re-associate with another access point. Depending on its physical placement, the transmission range of an access point may or may not overlap with other access points. A end-host has the option of associating with any access point, when multiple access points provide coverage at a particular physical location. Typically, wireless nodes employ an association policy that connects them to an access point with the best signal strength or the best signal to noise ratio—both metrics denote a good channel. Other association policies are possible where a the load on the access point and the signal to noise ratio ratio are both used to determine an association Such policies attempt to balance the load across overlapping access points, while providing a good channel to the end-host. Hardware vendors [1] already implement association policies in their wireless cards that use information such as the load on access point (number of end-hosts associated), error characteristics, and signal strength to intelligently associate with an access point. Our work assumes that wireless nodes implement such a policy.

The wireless channel (frequency) used by each access point is especially important in the presence of overlap. If overlapping access points use the same channel, then they will share the bandwidth of that channel. In contrast, if overlapping access points are assigned distinct channels, then each will communicate with its associated users on a different frequency, thereby increasing the total capacity and bandwidth of the wireless network. In 802.11b, for instance, three non-overlapping channels are available, while 802.11a supports twelve non-overlapping channels; these channels must be judiciously assigned in regions of overlapping coverage. This is illustrated in Figure 1(a), where overlapping APs are assigned two distinct channels, while non-overlapping APs operate on the same channel.

Last, an access point may or may not be connected to a wired network. When an access point is connected to a wired network, it can directly relay data between nodes on the wired and the wireless network. If an access point is not wired, packets from an end-host can only be forwarded to an overlapping access point over the wireless channel. Data is forwarded

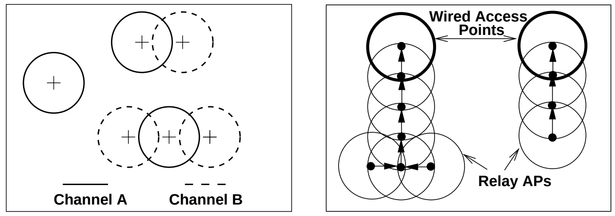

(a) Use of non-overlapping channels

(b) Extending wireless coverage

Figure 1: Use of non–interfering channels to increase throughput and overlapping access points to extend wireless coverage.

from one access point to another until it reaches a wired access point, which serves as the gateway to the wired network. In such networks, which are also referred to as mesh networks [2, 3], wired access points serve as gateways, while other access points serve as relays to a gateway access point; all access points, whether wired or not, provide wireless access to associated end-hosts. This is illustrated in Figure 1(b) where relays form an wireless overlay path to two gateway APs.

With this background, we formulate the two problems studied in this paper.

2.1 Provisioning for Alleviating Hot Spots

Consider a physical region where a wireless network needs to be deployed. We assume that an end-host density map for the region is available. The density map provides the distribution of end-hosts in the physical region and is used to guide the placement of access points in the region. Each density point in the map represents a certain number of users at that physical location. Each density point also has a mean bandwidth requirement associated with it, which represents the mean amount of traffic that the users will send and receive [15].

Given the dimensions of the physical region, the density map, and a budget k on the number of access points, the objective of the provisioning algorithm is to determine a placement—if one exists—for each of the k access points so that (i) all users in the region are covered, and (ii) the load on the most heavily loaded access point is minimized. Thus, the placement algorithm should maximize network coverage while minimizing load on access points in hotspot regions. Multiple overlapping access points may be deployed in hotspots (regions of high user density), which enables the user load to be distributed across these APs. Our goal is to minimize the maximize load seen by an access point in a hotspot, or put another way, maximize the minimum bandwidth seen by a user.

This placement problem can be formulated as an *Integer Linear Program* as follows.

Input:

A: An area to be covered with wireless access

N: All possible locations for access point placement

k: Budget on the number of access points

m: Number of available channels of each location

D: A density map of user distribution in area A

Output:

(i) Location of each access access points

(ii) The number of end-hosts covered and the average load that will be seen by each access point

Let x_i be a binary variable that denotes whether an access point is placed at location "i"; up to m access points may be placed at each physical location, $1 \leq i \leq m \times |N|$. Let f_{ij} be a binary variable that denotes if users at density point "j" connect to an access point at location "i". Also, $n = m \times |N|$ and d = $|D|$.

Objective:

$Minimize\ [\ Maximum \sum_{j=1}^{d} f_{ij}\ ]\ (\forall i; 1 \leq i \leq\ n)$

where $\sum_{j=1}^{d} f_{ij}$ denotes the load on an AP.

Constraints:

$$\sum_{i=1}^{n} x_i \quad \leq \quad k \tag{1}$$

$$\sum_{i=1}^{n} f_{ij} \quad = \quad 1\ \ \forall\, j \tag{2}$$

$$f_{ij} \leq x_i \quad (3)$$
$$x_i \in \{0,1\} \quad (4)$$
$$f_{ij} \in \{0,1\} \quad (5)$$

The first condition constrains the number of deployed access points to be no greater than the budget. The second constraint restricts each user to be associated with a single access point, and the third constraint ensures that a user only associates with an access point where one exists.
The above optimization is an instance of a Minimax optimization problem. The Minimax problem can be converted to an Integer Linear Program by modifying the objective function and adding to the constraints as follows:

Objective: Minimize M
Additional Constraints:

$$M >= \sum_{j=1}^{d} f_{ij} \quad \forall i \in n$$

The solution of the Integer Program is a set of locations where access points should be placed, denoted by x_i, and the expected load on each access point, denoted by $\sum_j f_{ij}$.

2.2 Provisioning in Multi-hop Wireless Networks

Consider a physical region where a multi-hop wireless network needs to be deployed. Like before, we assume that an user density map is available for the region. We assume that a certain number of wired access points are already deployed in this environment and as a result, a subset of the end-hosts are already covered. The location of existing wired access points is assumed to be given.
The objective is to construct a multi-hop wireless network using the remaining k access points. It is assumed that these k access points will not directly connect to a wired network, and instead will do so via one of the existing wired access points. The placement algorithm needs to determine a placement of these access points so that (i) all users are covered, (ii) each access point has a multi-hop path to a wired access point, and (iii) the load on the most heavily loaded access point is minimized. In addition to providing network coverage to users, each AP in this scenario also functions as a relay in a wireless overlay. Each relay forwards data from other APs in its transmission range along a path to a wired AP. Like before, a desirable placement will maximize coverage, minimize the load on heavily loaded access points, and ensure connectivity between wired APs and AP relays. Observe that overlap between access points for is necessary for relaying data as well as to balance load.
The following optimization problem results from this placement problem:

Objective:
$Minimize\ [\ Maximum \sum_{j=1}^{d} f_{ij}\ ] \ \ (\forall i; 1 \leq i \leq n)$
Constraints:

$$\sum_{i=1}^{n} x_i \leq k \quad (6)$$
$$\sum_{i=1}^{n} f_{ij} = 1 \ \ \forall j \quad (7)$$
$$f_{ij} \leq x_i \quad (8)$$
$$x_i \in \{0,1\} \quad (9)$$
$$f_{ij} \in \{0,1\} \quad (10)$$
$$If\ x_i = 1 \quad then \quad \exists\ a\ path\ to\ a\ wired\ AP\ location \quad (11)$$

As can be seen from the above formulations, there is considerable similarity between the two problems. In fact, the latter problem is identical to the first, except for an additional constraint that requires a wireless overlay path to a wired access point. Given these similarities, we present a unified approach for addressing both problems in the next section.

3 Provisioning Algorithms

In this section, we first show that the placement problem is NP-complete. We then present a greedy heuristic to address both the hotspot congestion and the multi-hop wireless placement problems. A LP relaxation of the Integer Program formulation is also presented for load balancing in congested wireless networks.

3.1 Hardness Properties

To show that the placement problem is NP-complete, we consider a simpler version of the problem where only wireless coverage is considered and load balancing across APs is ignored. The simpler version only attempts to place APs to provide coverage to all users.

Lemma 1 *The problem of fi nding the minimum number of access points to provide wireless coverage to all users is NP–Complete.*

Proof Sketch: The problem of finding the minimum number of access points to cover all density points can be stated as the *minimum dominating set* problem on a bipartite graph. Given a graph $G(V, E)$, where V is the set of vertexes and E the set of edges, a *dominating set* for G is a subset of vertices's V' such that: for all vertices's $u \in V - V'$ there is an edge (u, v) with vertex $v \in V'$. Since the minimum dominating set on bipartite graph is NP–Complete [9], the simpler version of provisioning access points for coverage is NP–Complete.
The formulation of the problem is as follows: the bipartite graph $G(U, V : E)$ has set U as the set of all access point locations and set V representing the set of all density points. E is the set of edges and an edge $(u, v) \in E$ $(u \in U \; and \, v \in V)$ if distance between location u and density point v is less than r, the communication range. The problem of finding the minimum number of access points in set U that cover all density points is same as solving the minimum dominating set problem on a bipartite graph. Since, the minimum dominating set on bipartite graph is NP–Complete [9], the simpler version of provisioning access points for coverage is NP–Complete.

Lemma 2 *The problem of fi nding the minimum number of access points for coverage in a multi– hop wireless network is NP– Complete.*

Proof Sketch: The above problem can be stated as the *minimum connected dominating set* problem on a bipartite graph. Since the minimum minimum dominating set on bipartite graph is NP–Complete [9], the simpler version of provisioning access points for multi–hop networks is NP–Complete.
To prove hardness of the problem of provisioning to extending wireless coverage, we make two additions to the graph construction. To sets U and V, we add special vertexes u' and v' respectively, such that u' is connected to all vertexes representing static access point locations in U and to vertex v'. Additionally, all vertexes in set U (except the special vertex u') have an edge connecting them if distance between them is less then r, the communication range. Solving the minimum connected dominating set on the bipartite graph G, will yield the minimum number of connected access points to be placed that cover all users. Since, minimum connected dominating set with bipartite graphs is NP–Complete [9], the problem of extending wireless coverage is also NP–Complete.

3.2 A Greedy Approach

Since the placement problems for alleviating load and extending coverage in a multi–hop network are NP-Complete, this section presents a greedy heuristic to determine a placement of access points. Table 1 lists the steps in our heuristic. Set A is the set of all locations where access points are to be placed, k is the budget, D the set of all density points and $C(i)$ the set of density points covered by access point at location "i". The greedy heuristic works in three phases. First, access points are placed to maximize coverage. Next, additional access points are placed in regions of high user density. Last, the available channels are judiciously assigned to the each access point.

Phase 1: Coverage Maximization
The first phase places access points to greedily cover all density points in a physical region. At each step, an access point is placed at the location that will cover the maximum number of yet-to-be-covered users. The number of users covered by an AP at a location is simply the sum of all users that are within its transmission range and are not already covered by an existing AP. The essential idea is to cover all users using a minimum number of access points, and to use the remaining access points in the second phase for distributing load across APs. Table 1, steps 1 through 4 depicts the first phase.
For multi-hop wireless placement, at each step, only those locations that will connect to a wired AP via one or more existing APs are considered. This constraint ensures that each placed access point forms an overlay path to route its traffic to a wired access point.
Phase 2: Load Balancing
Once all users have wireless coverage, the second phase attempts to improve the performance seen by each user. Assuming all users equally share the bandwidth of the wireless channel, users in high density areas will see a lower bandwidth per user than those in low density regions. Consequently placing additional access points in high density regions (and operating them

1. $A = \emptyset$ /* coverage maximization phase */
2. If $|A| = k$ goto Step 7
3. Choose access point 'i' that covers maximum density points in D
('i' should have a path to a wired access point in set A if extending coverage)
    $A = A + i$
    $D = D - C(i)$
4. If $D \neq \emptyset$ goto Step 2
5. If $|A| = k$ goto Step 7 /* load balancing phase */
6. Else
    Choose access point 'i' which
    Minimizes [Maximum n_i] over all $i \in A$
    (where n_i is load at access point i)
    $A = A + i$
    goto Step 5
7. If $D = \emptyset$ *All density points covered*
   Else *All density points not covered*

Table 1: The greedy heuristic for access point placement.

on a non-overlapping frequency) will improve the bandwidth seen by these users. The second phase incrementally places each remaining AP at the location that will see the maximum improvement in performance by this placement. Observe that the number of available channels m constrains the maximum number of overlapping APs at a physical location (since at most m non-overlapping frequencies are available for the APs). By greedily placing access points in congested regions that will benefit most by this placement, this step maximizes the minimum bandwidth seen by a user (or minimizes the load on the most heavily loaded AP). This placement process continues until the budget of k access points is exhausted. Table 1, steps 5 and 6 depict the second phase.

For multi-hop wireless placement, only those locations with an existing path to a wired APs are considered in each step. Such locations and their paths are formed in the first phase and the greedy heuristic aims to co–locate access points at these locations to alleviate the load on each additional path.

Phase 3: Channel Assignment

The third phase assigns operating channels to each access point. For the hotspot placement, the objective of channel assignment is to eliminate or minimize the number of interfering channels assigned to overlapping access points. This can be viewed as a graph coloring problem, where each node in the graph is an access point and an edge between two nodes denotes overlap between the corresponding access points. The objective is to colors nodes in the graph so that nodes connected by an edge have distinct colors (to the extent possible), and the number of available colors is determined by the number of channels. We use a greedy coloring heuristic that colors access points in the descending order of their degree of overlap; the degree of overlap is the number of APs that overlap with an access point. Colors are reused if the degree of overlap at a particular location(node degree) is greater than the number of available channels (colors).

Channel assignment for the multi-hop AP placement problem is done differently. Observe that all access points on an overlay path to a wired access point need to operate on the same channel in order to forward packets to one another. Hence, rather than coloring individual access points, entire paths (or the entire tree rooted at the wired access point) is colored in each step. In this case, each tree rooted at a wired AP is represented as a node for graph coloring. Two nodes are connected by an edge if any two access points in the corresponding trees overlap with one another. Given such a graph, the above coloring heuristic can be used to assign colors(channels) to nodes(access points)—colors are greedily assigned in decreasing order of the node degree.

3.3 A LP Relaxation Approach

Another solution for the problem of provisioning access points for alleviating hot spots as stated in Section 2.1 is based on the LP relaxation technique. Similar to the greedy approach, this solution also uses a three-phase procedure for coverage, placement of additional access points to alleviate load and channel assignment.

Phase 1: Relaxation to maximize coverage. The first step relaxes the integer constraint of the Integer Program. The Linear Program is setup with a budget equal to the total number of access point locations and a single channel at each location. The solution to the Linear Program is a set of fractional values for the x_i and f_{ij} variables, i.e.: the access point locations and the associations of density points. A rounding scheme converts the fractional x_i values to binary 0 or 1 values, such that the problem constraints on both the x_i and f_{ij} variables are satisfied. For this purpose, the x_i variables are sorted according to their fractional values in descending order. The rounding mechanism uses the smallest subset starting from the largest valued x_i variable to be placed such that all users are covered. An alternative approach is to use a *greedy* approach and select x_i variables which cover the most density points in each step. In this approach, a location x_i is chosen which covers the most density points; these density points are then removed from the rounding process and the next location x_j which covers most of the remaining density points is chosen. The process is repeated till all users are covered. Our technique, uses both the above rounding schemes and selects the smaller subset of x_i variables to cover all density points.

Objective: Minimize M
Constraints:

$$\begin{aligned}
M >= \quad & \textstyle\sum_{j=1}^{d} f_{ij} && \forall i \in n \\
\sum_{i=1}^{n} x_i \quad & <= && k \\
\sum_{i=1}^{n} f_{ij} \quad & = && 1 \quad \forall j \\
x_i \quad & \leq && 1 \texttt{ if } x_j \neq 1 \texttt{ in Step 1} \\
x_i \quad & = && 1 \texttt{ if } x_j = 1 \texttt{ in Step 1}
\end{aligned}$$

where: $i = i/m, n = m \times |N|$ and d = $|D|$

Table 2: Second step of Linear Program formulation.

Phase 2: Relaxation to alleviate hotspots. In the second phase, a Linear Program is formulated using the solution from Phase 1 to place additional access points. The Linear Program is initialized with access points placed in Phase 1 and extended with the additional budget and number of channels permitted at each location. Solution to the second–phase Linear Program (Table 2) yields a solution which minimizes the maximum load over all placed APs. The solution is represented with a subset x_i variables from Phase 1 set to 1 and a set of fractional valued x_i and f_{ij} variables, To round the fractional values of the x_i variables the following procedure is used: Locations of access points are ordered in descending values of x_i(without considering those placed in Phase 1). Next, assuming j APs are placed in Phase 1 and k is the available budget, the first '$k - j$' x_i variables and corresponding locations are selected and set to 1. The rest x_i variables are set to 0. The x_i variables set to value 1 are the access point locations for placement.
Phase 3: Channel Assignment. The third step used for channel assignment of access points to be placed is same as that described in Section 3.2.

4 Experimental Evaluation

In this section, we report empirical results of our techniques. The 802.11 parameters used in the evaluation are number of available channels, budget of access points and the communication range of each access point. In our experiments, if the communication range is r, we assume all points in a radius of Euclidean distance r to be within communication range of the access point. Our solutions can be extended to include other communication models based on access point location and radio propagation characteristics.

4.1 Experimental Setup

In order to evaluate the performance of our techniques, we use four different density point distributions. Three synthetic distributions representing different scenarios and are differentiated by the concentration of density points in an area. The three scenarios are: (i) *Uniform*, where users are uniformly distributed in an area (Figure 2(a)), (ii) *Some Skew*, which is a skewed distribution with concentration of density points in a few regions (Figure 2(b)) and (iii) *Heavy Skew*, where a majority of the density points distributed in a small region (Figure 2(c)). The above distributions are in an area of $200m \times 200m$ with 400 density points. The fourth distribution is obtained from the second floor of the Computer Science Building of the University of Massachusetts, Amherst. The density point distribution is constructed using locations of cubicles for students, locations of faculty offices and seating arrangements in conference rooms. The floor area is 77 meters$\times$35 meters and has 166 density points as shown in Figure 2(d). The dark dots in the figure are the current locations of the access points. Each scenario also has a set of potential locations where access points can be placed. These locations for the three synthetic scenarios are marked by 16 intersection points of the grid lines in Figure 2.
The metric used for comparison using the above scenarios is the *maximum load* over all access points. For multi-hop networks, as all non-wired access points communicate via a wired access point, the maximum load over the wired access points is used as the metric.

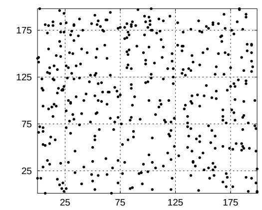

(a) Scenario 1 (Uniform)

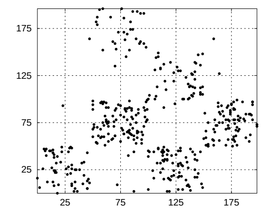

(b) Scenario 2 (Some Skew)

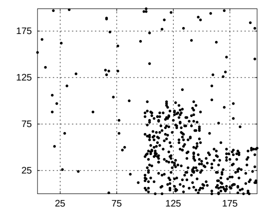

(c) Scenario 3 (Heavy Skew)

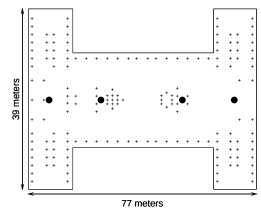

(d) CS Bldg. 2nd floor

Figure 2: Scenarios considered in experimental evaluation.

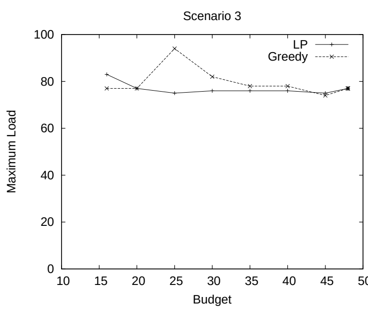

(a) Number of channels=3

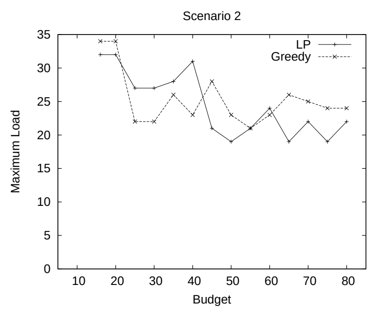

(b) Number of channels=9

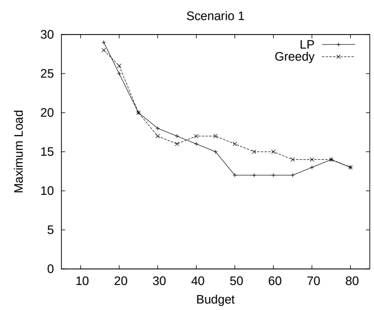

(c) Number of channels=12

Figure 3: Effect of varying the available budget for placement.

4.2 Provisioning for Load Balancing

This section compares the performance of the two heuristics, Greedy and LP–based, to provision access points for coverage and alleviating hot spots. The load on an access point is determined based on the number of end-hosts in communication range, operating channel, overlapping regions and association policy. The association policy used by the heuristics considers load at access points and end-hosts associate with the least loaded AP in range. The following experiments report the value of maximum load over the placed access points.

4.2.1 Exploiting Overlap

Budget	Scenario1 #channels			Scenario 3 #channels		
	2	3	12	2	3	12
16	47	36	28	111	83	32
20	41	40	25	111	77	27
25	42	30	20	111	75	25
MAX	47	29	10	113	77	21

Table 3: Maximum load after exploiting overlap.

Table 3 reports the benefits of exploiting overlap with excess budget and multiple channels to reduce load imbalance at access points. For Scenarios 1 and 3, 11 access points are needed to cover every single user without exploiting overlap. With a single channel, the maximum load in each case is 82 and 221, respectively. Any increase in budget beyond 11 enables overlapping placement to alleviate hotspots. The table reports the load on the most heavily loaded AP for budgets of 16 to infinity. As can be seen from the table, even with a small increase in budget from 11 to 16 coupled with multiple channels provides substantial gains. With a budget of 16 and 3 channels the reduction in load is 56% and 62% for the two scenarios respectively. With 12 channels the corresponding numbers are 65% and 85% respectively. The maximum benefits that can be obtained with maximum budget and 12 channels for the two scenarios is 87% and 90% respectively. The result shows that our approaches are able to exploit overlap effectively, even with a small increase in minimum budget required for coverage and in presence of multiple channels.

4.2.2 Effect of Budget

In this section, we study the effect of varying budget on maximum load. The scenarios have 16 potential access point locations, each access point has a range of 50m and the budget is varied for 16 to 80 access points. Figure 3(a) plots the maximum load with Scenario 3 and a limit of 3 co-located channels at each location, Figure 3(b) reports Scenario 2 with 9 channels per location and 3(c) reports Scenario 1 with 12 channels per location. As can be seen, with increase in the budget, both the LP-based and Greedy heuristics are successful in decreasing the maximum load. Clearly, an increased budget provides more opportunities to distribute load evenly and reduce hotspots. The interesting point to note is that with 9 and 12 channels at each location, the LP-based heuristic outperforms the Greedy heuristic in certain regions. In Fig 3(c), with a constrained budget of up to 35 APs and a higher budget of more than 75 APs both heuristics perform similarly. At intermediate budget values from 35 to 70 APs, the LP-based approach performs better than the Greedy approach. This result suggests that with constrained and surplus resources both heuristics perform similar and the LP-based heuristic performs better with intermediate resources.

4.2.3 Effect of Range

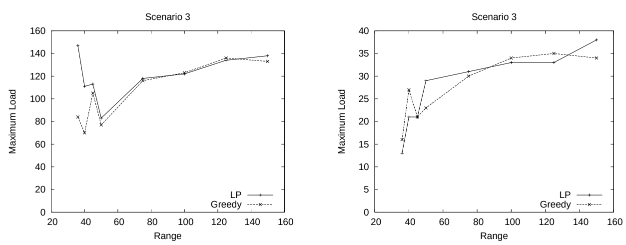

(a) Number of channels=3 and Budget=16

(b) Number of channels=12 and Budget=50

Figure 4: Effect of varying the communication range of each access point.

The communication range of an AP determines whether an end-host can associate with it. In this experiment, we use Scenario 3 to vary the communication range from 35m to 150m and report its effect on the maximum load. Figures 4(a) and (b) depict our results for 3 and 12 channels and with budgets of 16 and 50 access points, respectively. In general, we see that for a fixed number of channels and budget, as we increase range, the maximum load tends to increase. This is because as range increases, more APs overlap with one another, and with a fixed number of channels, more of them get assigned interfering channels. Both the LP and the Greedy solutions have comparable performance with increase in range.

4.2.4 Effect of Number of Channels

To study the effect of number of available channels, the three scenarios were used with budgets of 20, 30 and 75 APs and a range of 50m. The number channels at each location was systematically varied from 1 to 12. As seen from Figure 5, as the number of available channels increases, the imbalance and the maximum load on an AP decreases. A larger number of channels provides more opportunities to place APs in hot spots, reducing the congestion at hot spots. Increasing the number of channels beyond a point bring diminishing benefits, since the algorithms have already placed sufficient overlapping APs and any additional channels can not be put to use. Again, we find that the greedy and LP heuristics have comparable performance.

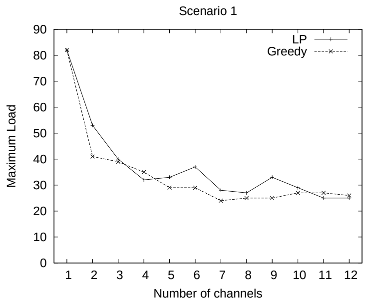

(a) Budget=20

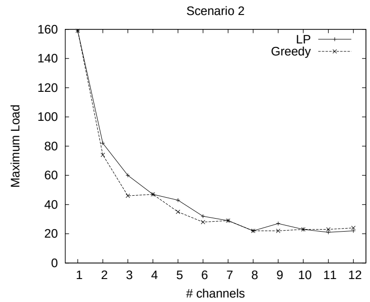

(b) Budget=30

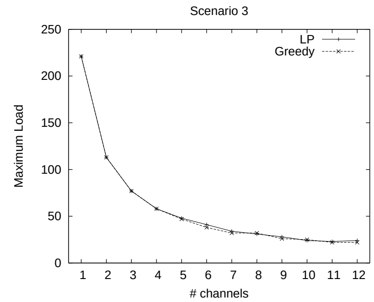

(c) Budget=75

Figure 5: Effect of varying the number of channels at each location.

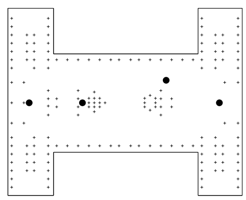

(a) Greedy (Range=25m)

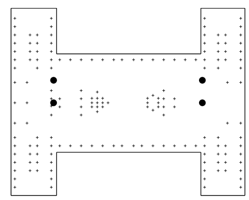

(b) LP (Range=25m)

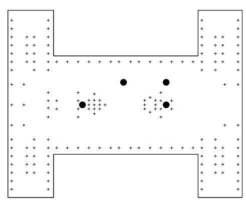

(c) Greedy (Range=30m)

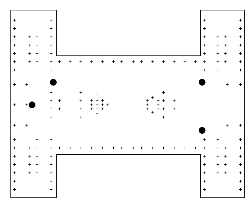

(d) LP (Range=30m)

Figure 6: Access Point placement in CS Bldg.

4.2.5 Case Study: Placement in a Real Building

(Budget, Range)	#channels=2		#channels=3	
	Current	Our Approach	Current	Our Approach
(4 , 25)	49	44	43	44
(4 , 30)	56	44	43	43
(4 , 35)	59	56	51	47

Table 4: Results of case study with different range values in meters.

We consider density distribution in the Computer Science building and determine the placement of access points. Figures 6(a)—(d) show the placement of access points using the Greedy and LP-based heuristics for range values of 25m and 30m respectively. As can be seen in both cases, the locations of access points is different from the current location of access points as shown in Figure 2(d). Table 4 reports the effect of varying range on maximum load with the current placement and our approaches. We only report the better of the LP and greedy placements in each case. A budget of 4 APs with 2 and 3 channels at each location were used. The table shows that our approach is able to improve upon the current placement when two channels are available; for three channels, the current and our placements are comparable, since there are plenty of non-overlapping frequencies available to both placements.

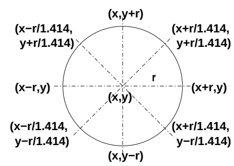

Figure 7: Potential neighbors to place access points to extend coverage.

4.3 Multi-Hop Wireless Networks

Next, we present experimental results for provisioning access points to form a multi-hop a wireless network. For each placed access point, the experimental setup allowed eight neighboring locations to place additional overlapping access points and extend coverage. Figure 7 shows the eight potential locations within a communication range of r for an access point at location x, y. For the experiments reported, the wired access point locations were $(50, 50)$, $(50, 150)$, $(100, 100)$, $(150, 50)$ and $(150, 150)$ for an area of $200m \times 200m$ with the three synthetic input scenarios.

Using a single channel at each access point location and a communication range of 40m, Figures 8(a), 8(b) and 8(c), show locations of access points and the wireless overlay path to reach a wired access point for the three scenarios. As can be seen pictorially, the placement of APs matches the distribution of users in each case. As shown in the figure, the Greedy heuristic requires 22, 14 and 17 access points to cover all users in Scenarios 1, 2 and 3, respectively.

4.4 Effect of Varying Resources

As non-wired APs communicate over a wireless overlay rooted at a wired AP, the *maximum load* is computed at the wired access points. Figure 9 reports experiments for extending wireless coverage with input Scenario 1.

Budget: Figure 9(a) plots the maximum load with varying budget from 22 to 228 APs and a range of 40m. With increase in budget there is increased opportunity to co-locate access points and form longer overlapping paths to distribute load. As can be seen, the maximum load decreases monotonically with increase in budget and the maximum benefits are obtained with 12 channels per location. Considering 6, 9 and 12 channels per location, a *knee* in the values of maximum load can be seen with increased budget. This suggests that a moderate budget with higher number of channels can provide significant benefits. Using 9 channels, the maximum load with a budget of 100 access points is 22 and that with a budget of 171 is 16.

Range: Next, we vary communication range of each access point from 25m to 100m with a fixed budget of 100 access points and report the results in Figure 9(b). Using a single channel, with increase in communication range the degree of overlap between access points increases and as result the load on each increases rapidly. With more than a single channel, the channels are reused to form non-overlapping paths to wired APs and provide benefits to minimize load. For this scenario, higher number of channels almost provide the same benefits. The least load using 9 channels is 9 density points, whereas with 12 channels is 7.

Number of channels: Figure 9(c) plots the maximum load with varying number of channels at each location and a communication range of 40m. Increasing number of channels per location, increases the potential for number of non-overlapping paths. As can be seen from figure, with a fixed budget increasing channels per location exploits overlap to co-locate APs and decrease load. In our scenario, 4 to 5 channels per location produced substantial gains and increasing the channels further produced diminishing benefits. Using a budget of 100 access points, the maximum load with 6 channels is 26 and with 12 channels is 22.

5 Related Work

The behavior of users and network performance parameters of wireless LANs have been studied in several scenarios. The different scenarios studied are: a Conference setting [6], a Classroom Area Network [4], a Corporate Wireless Local-Area Network [7] and a Campus Wide Area Network [11]. The studies measured parameters like user mobility, user and load distribution, network bandwidth, throughput at access points, application throughput, channel errors and time of day effects. All studies validate the point that user load and bandwidth requirements at access points are uneven and need intelligent association policies and access point provisioning techniques.

The *hotspot* problem for wireless LANs and an admission control and association policy to admit (associate) users based on their QoS requirements was studied in [5]. Their work focuses on admission control strategies and, unlike us, they are not concerned with provisioning issues. Further, they attempt to meet bandwidth needs of each user, which is assumed to be

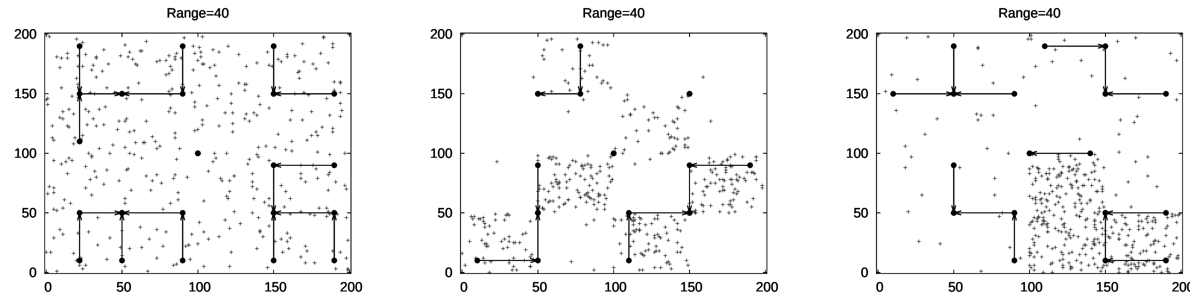

(a) Scenario 1 (Uniform) (b) Scenario 2 (Some Skew) (c) Scenario 3 (Heavy Skew)

Figure 8: Extending wireless coverage with Range=40m

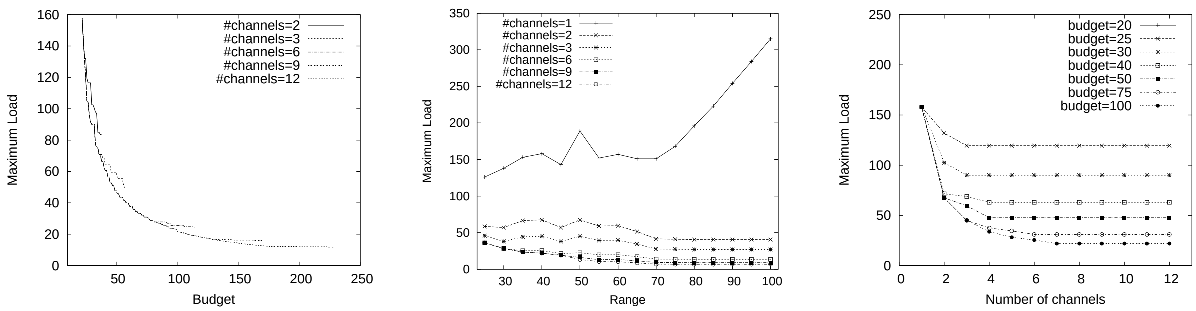

(a) Range=40m (b) Budget=100 (c) Range=40m

Figure 9: Effect of varying Budget, Communication Range and Number of channels at each location on Maximum Load.

given, while we attempt to evenly distribute the load. Intelligent association policies, also known as *association control* are studied in [8]. In this work, the problem of association with bandwidth constraints is formulated. Solutions presented aim at providing network-wide load balancing in terms of bandwidth allocation. Currently, our techniques use an association policy based on the load at APs and can be extended to include the above techniques.

The problem of placing minimum number of ITAPs (Internet Transit Access Points) in order to form a multi–hop wireless overlay for Internet access is addressed in [14]. In our work, we consider the reverse problem—given a set of wired access points, provisioning access points to extend coverage using a wireless overlay and also minimizing the load imbalance. Examples of practical instantiations of wireless mesh or overlay networks for Internet access and to as backbone networks are described in [2] and [3].

The concept of demand nodes is widely used to solve facility location problems and has been used in [10, 13, 15] to represent spatial distribution of nodes and their traffic demands. Solutions to various versions of the problem which maximize number of served nodes, minimize number of base-stations required, reduce number of interfering channels are proposed. None of these techniques address the issue of fair distribution of load while provisioning.

6 Conclusions

In this paper, we addressed the problem of access point provisioning for wireless networks in congested hot spots and multi-hop mesh networks. We argued that both problems are closely related and showed that the problems are NP-Complete. We proposed a greedy heuristic and a LP relaxation approach—both heuristics exploit the notion of overlapping placement of access points. An experimental evaluation shows that our approaches can effectively exploit overlap and presence of multiple channels to alleviate hotspots and extend coverage. With uniform and skewed distribution of end-hosts, a small addition to the minimum budget required gave benefits up to 65% and 85% respectively.

References

[1] Cisco Wireless Networking Products. http://www.cisco.com/en/US/products/hw/wireless/index.html.

[2] Self-Organizing Neighborhood Wireless Mesh Networks. http://research.microsoft.com/mesh/.

[3] MIT Roofnet Implementation, 2003. http://pdos.lcs.mit.edu/roofnet/.

[4] Guangwei Bai and Carey Williamson. Simulation Evaluation of Wireless Web Performance in an IEEE 802.11b Classroom Area Network. In *28th Annual IEEE Conference on Local Computer Networks (LCN 2003)*, 2003.

[5] A. Balachandran, G. Voelker, and P. Bahl. Hot-Spot Congestion Relief in Public-Area Wireless Networks. *IEEE Workshop on Mobile Computing Systems and Applications*, pages 70–80, June 2002.

[6] A. Balachandran, G. Voelker, P. Bahl, and P. Rangan. Characterizing User Behavior and Network Performance in a Public Wireless LAN. In *ACM SIGMETRICS*, June 2002.

[7] M. Balazinska and P. Castro. Charaterizing Mobility and Network Usage in a Corporate Wireless Local-Area Network. In *First International Conference on Mobile Systems, Applications, and Services (MobiSys)*, May 2003.

[8] Yigal Bejerano, Seung-Jae Han, and Li (Erran) Li. Fairness and Load Balancing in Wireless LANs Using Association Control. In *Tenth Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2004.

[9] Michael Garey and David Johnson. *Computers and Intractability*. W. H. Freeman and Company, San Francisco, CA, USA, 1979. ISBN 0-7167-1044-7.

[10] C. Glasser, S. Reith, and H. Vollmer. The Complexity of Base Station Positioning in Cellular Networks. In *ICALP Workshops*, 2000.

[11] D. Kotz and K. Essien. Analysis of a campus-wide wireless network. In *Eighth Annual International Conference on Mobile Computing and Networking (Mobicom)*, September 2002.

[12] Purushottam Kulkarni and Prashant Shenoy. Exploiting Overlap for Provisioning of Access Points in Wireless Networks. Technical Report TR04-XX, University of Massachusetts, Amherst, 2004.

[13] Rudolf Mathar and Thomas Niessen. Optimum positioning of base stations for cellular radio networks. *Wireless Networks*, 6:421–428, 2000.

[14] Lili Qiu, Ranveer Chandra, Kamal Jain, and Mohammad Mahdian. Optimizing the Placement of Integration Points in Multihop Wireless Networks, 2004. To Appear in Twelfth IEEE International Conference on Network Protocols.

[15] K. Tutschku, K. Leibnitz, and P. Tran-Gia. ICEPT - An Integrated Cellular Network Planning Tool, May 1997. IEEE/VTS 47th Vehicular Technology Conference.