

Optimizing the Placement of Internet TAPs in Wireless Neighborhood Networks

Ranveer Chandra[†], Lili Qiu^{*}, Kamal Jain^{*}, Mohammad Mahdian^{*}
Cornell University[†] Microsoft Research^{*}

Abstract

Efficient integration of a multi-hop wireless network with the Internet is an important research problem. In a wireless neighborhood network, a few Internet Transit Access Points (ITAPs), serving as gateways to the Internet, are deployed across the neighborhood; houses are equipped with low-cost antennas, and form a multi-hop wireless network among themselves to cooperatively route traffic to the Internet through the ITAPs. Furthermore, the placement of Internet TAPs is a critical determinant of system performance and resource usage. In this paper, we explore the placement problem under three wireless link models. For each link model, we develop algorithms to make informed placement decisions based on neighborhood layouts, user demands, and wireless link characteristics. We also extend our algorithms to provide fault tolerance and handle significant workload variation. We evaluate our placement algorithms and show that our algorithms yield close to optimal solutions over a wide range of scenarios we have considered.

1. Introduction

In wireless neighborhood networks, a few Internet Transit Access Points (ITAPs) are placed to relay data from the wireless multi-hop network to the Internet and vice versa. One of the key challenges in designing and deploying such networks is an efficient integration of multi-hop wireless networks with the Internet.

Wireless neighborhood networks are characterized by two important design constraints. They should be easy and cheap to deploy. Moreover, in order to be competitive to DSL or cable, such networks should provide Quality of Service (QoS) to end users. To achieve efficient utilization of network resources and good user performance, it is imperative to have a strategic placement of ITAPs. A good ITAP placement algorithm will have to (i) efficiently use wireless capacity, (ii) take into account the impact of wireless interference on network throughput, and (iii) be robust in face of

failures and changes in user demands. There has been little previous work on this subject.

In this paper, we propose a series of ITAP placement algorithms to achieve these goals. Our key contributions are:

- We formulate the ITAP placement problem under three wireless models. For each model, we develop algorithms to efficiently place ITAPs in the network. Our algorithms aim to minimize the number of required ITAPs while satisfying users' bandwidth requirements. We demonstrate the efficiency of the algorithms through analysis and simulation.
- To enhance robustness, we present a fault tolerance version of the placement algorithm that provides bandwidth guarantees in the presence of failures.
- We extend the algorithms to take into account variable traffic demands by developing an approximation algorithm to simultaneously optimize ITAP placement for demands over multiple periods. This algorithm is very useful in practice since user demands often exhibit periodic changes (e.g., diurnal patterns).

2. Related Work

There has been a recent surge of interest in building wireless neighborhood networks. [1] presents a scheme to build neighborhood networks using standard 802.11b Wi-Fi technology [20] by carefully positioning access points in the community. However, it requires a large number of access points. Moreover, it requires direct communication between machines and access points, which is difficult to meet in real terrains. Nokia's Rooftop technology, presented in [8], provides broadband access to households using a multi-hop approach that overcomes the shortcomings of [1]. The idea is to use a mesh network model with each house deploying a radio, as considered in this paper. This radio serves the dual purpose of connecting to the Internet and routing packets for neighboring houses [4]. The deployment and management cost of Internet TAPs in such networks is significant, and so it is crucial to minimize the required number of ITAPs to

provide QoS and fault tolerance guarantees. However, these problems are not addressed in [1, 8].

There have been a number of interesting studies on placing servers at strategic locations for better performance and efficient resource utilization in the Internet. For example, the authors in [19, 16, 22] examine placement of Web proxies or server replicas to optimize clients' performance; and Jamin *et al.* [15] examines the placement problem for Internet instrumentation. Facility location problems are also related to ITAP placement problem, and have been considered extensively in the fields of operation research (e.g., [18, 25]). Approximation algorithms with good worst case behavior have been proposed for different variants of this problem. The previous work on server placement or facility location cannot be applied to our context because they optimize locality in absence of link capacity constraints. This may be fine for the Internet, but is not sufficient for wireless networks since wireless bandwidth is much more limited. Moreover, the impact of wireless interference, and considerations of fault tolerance and workload variation make the ITAP placement problem very different from those studied earlier.

The work closest to ours is the pioneering work in [3]. It aims to minimize the number of ITAPs for multi-hop neighborhood networks based on the assumption that ITAPs use a Time Division Multiple Access (TDMA) scheme to provide Internet access to users. However, TDMA is difficult to implement in multi-hop networks due to synchronization and channel constraints [2]. Furthermore, the proposed slotted approach might not utilize all the available bandwidth due to unused slots. In comparison, in this paper we look at more general and efficient MAC schemes, such as IEEE 802.11. Removing the TDMA MAC assumption yields completely different designs, and increases applicability of the resulting algorithms. As we will show in performance evaluation, our placement algorithms significantly out-perform the algorithm in [3] under general MAC schemes.

In summary, placing ITAPs under the impacts of *link capacity constraints, wireless interference, fault tolerance, and variable traffic demands* is a unique challenge that we aim to address in this paper.

3. Problem Description and Network Model

The ITAP-placement problem, in its simplest form, is to place a minimum number of ITAPs that can serve a given set of nodes on a plane, which we call houses. A house h is said to be successfully served, if its demand, w_h , is satisfied by the ITAP placement. A house h is served by an ITAP i through a path between h and i . This path is allowed to pass through other houses, but any two consecutive points on this path must have wireless connectivity between them. We are usually interested in the fractional version of this problem. That is, we consider the flexibility that a house is allowed to

route its traffic over multiple paths to reach an ITAP.

This problem can be modeled using the following graph-theoretic approach. Let \mathcal{H} denote the set of houses, \mathcal{I} denote the set of possible ITAP positions, and w_h denote the demand from house h . We construct a graph G on the set of vertices $\mathcal{H} \cup \mathcal{I}$ by connecting two nodes if and only if there is wireless connectivity between them. The goal is to open the smallest number of ITAPs (denoted by the set \mathcal{I}'), such that in the graph $G[\mathcal{H} \cup \mathcal{I}']$, one can route w_h units of traffic from house h to points in \mathcal{I}' simultaneously, without violating capacity constraints on vertices and edges of the graph.

The edge capacity, Cap_e , in the graph denotes the capacity of a wireless link. In addition, each node also has an upper bound on how fast traffic can go through it. Therefore, we also assign each node with a capacity, Cap_h . Usually $\text{Cap}_h = \text{Cap}_e$, as both represent the capacity of a wireless link. (Our schemes work even when $\text{Cap}_h \neq \text{Cap}_e$, e.g., when a node's processing speed becomes the bottleneck.) Moreover, each ITAP also has a capacity limit, based on its connection to the Internet and its processing speed. We call this capacity, the ITAP capacity, Cap_i .

In addition, another input to the placement algorithms is a wireless connectivity graph (among houses). We can determine whether two houses have wireless connectivity using real measurements, and give the connectivity graph to our placement algorithms for deciding ITAP locations. In our performance evaluation, since we do not have wireless connectivity graphs based on real measurements, we instead derive connectivity graphs based on the *protocol model* [12]. In this model, two nodes i and j can communicate directly with each other if and only if their Euclidean distance is within a communication radius, CR . Given the position of all the nodes, we can easily construct a connectivity graph by connecting two nodes with an edge if their distance is within CR . However our placement algorithms can also work with other wireless connectivity models (e.g., physical model [12] or based on real measurements).

3.1. Incorporating Wireless Interference

There are several ways to model wireless interference. One approach is to use a fine-grained interference model based on the notion of a conflict graph, introduced in [14]. The main challenge of using the fine-grained interference model is high complexity (sometimes prohibitive) for even a moderate-sized network.

We use a coarse-grained interference model that estimates a relation between throughput and wireless interference. Since there is no single available function that captures the impact of interference on wireless throughput, we estimate the wireless throughput using two related functions. In our discussion, Throughput_l denotes the amount of throughput on a link along a path of length l , assuming

each wireless link capacity is 1. The other function, $g(l)$, denotes the amount of link capacity consumed if it is on a path of length l and the end-to-end throughput of the path is 1. Assuming the end-to-end throughput increases proportionally with the edge capacity, which is true in practice, we have $g(l) = \frac{1}{\text{throughput}_l}$.

In this paper, we study the following models separately:

1. **Ideal link model:** If $\text{throughput}_l = 1$ for all l , or equivalently, $g(l) = 1$, we get the basic version of the problem. This model is appropriate for the environment with very efficient use of spectrum. A number of technologies, such as directional antennas (e.g., [7, 11]), power control, multiple radios, and multiple channels, all strive to achieve close to this model by minimizing throughput degradation due to wireless interference.
2. **General link model:** A more general model is when throughput_l or $g(l)$ is a linear function of l . As we will show in Section 4.2, we can formulate the ITAP placement problem for the general link model as an integer linear program, and use a polynomial algorithm to approximate the solution. In addition, we also develop more efficient heuristics for two forms of $g(l)$. The first is the **bounded hop-count model**. If $\text{throughput}_l = 1$ for $l \leq k$ and $\text{throughput}_l = 0$ for $l > k$ (or equivalently, $g(l) = 1$ for $l \leq k$ and $g(l) = \infty$ for $l > k$), we get a variant in which flow cannot be routed through paths of length more than k . This approximates the case where we try to ensure each flow gets at least a threshold amount of throughput by avoiding paths that exceeds a hop-count threshold. The second is the **smooth throughput degradation model**, which corresponds to the case when $\text{throughput}_l = \frac{1}{l}$, where l is the number of hops in the path. This is equivalent to $g(l) = l$ for all l 's (i.e., the capacity consumed is equal to the flow times the number of hops). This represents a conservative estimate of throughput in a linear-chain network, and therefore this model is appropriate when tight bandwidth guarantees are desired.

Note that the above models capture wireless interference and contention among nodes whose paths to ITAPs share common links or nodes. A more accurate model will have to handle interference among nodes on independent paths (e.g., using the conflict graph [14]). However, in Section 5, we use packet-level wireless network simulations to show that an ITAP placement based on the above models gives satisfactory performance.

3.2. Generic Approach

In the following sections, we will investigate different variants of the placement problem. Our generic approach is

as follows. Given a set of potential ITAP locations, which may include all or a subset of points in the neighborhood, we first prune the search space by grouping points into equivalence class, where each equivalence class is represented by the set of houses that are reachable via a wireless link. For example, if points A and B have wireless connectivity to the same set of houses, then they are equivalent as far as ITAP placement is concerned. Therefore we only need to search through all the equivalence classes, instead of all points on the plane. (Refer to [21] for details). Then based on our choice of wireless link model, fault-tolerance requirements, and variability in user demands, we apply one of the placement algorithms described in Section 4, Section 7.1, and Section 7.2 to determine ITAP locations.

4. Placement Algorithms

In this section, we study how to place ITAPs under the impacts of link capacity constraints and wireless interference.

4.1. Ideal Link Model

First, we consider the placement problem for the ideal link model. We formulate the problem as a linear program, and present an approximation algorithm.

4.1.1. Problem Formulation: We formulate the placement problem for the ideal link model as an integer program shown in Figure 1. For each edge e and house h , we have a variable $x_{e,h}$ to indicate the amount of flow from h to ITAPs that is routed through e . For each ITAP i we have a variable y_i that indicates the number of ITAPs opened at the location i (More precisely, y_i is the number of ITAPs opened at locations in the equivalence class i , where the equivalence class is introduced in Section 3.) Cap_e , Cap_h , and Cap_i denote the capacity of the edge e , house h , and ITAP i , respectively; w_h denotes the traffic demand generated from house h .

Now we present a brief explanation of the above integer program. The first constraint, $\sum_{e=(v,h')} x_{e,h} = \sum_{e=(h',v)} x_{e,h}$, formulates the flow conservation constraint, i.e., for every house except the house originating the flow, the total amount of flow entering the house is equal to the total amount of flow exiting it. The inequality $\sum_{e=(h,v)} x_{e,h} \geq w_h$ formulates the constraint that each house has w_h amount of flow to send, and the third constraint indicates that a house does not receive flow sent by itself. The next three inequalities of the above integer program capture the capacity constraints on the edges, houses, and ITAPs. The inequality $\sum_{e=(v,i)} x_{e,h} \leq w_h y_i$ says that no house is allowed to send any traffic to an ITAP unless the ITAP is open. Notice that this inequality is redundant and follows from the ITAP capacity constraint and the assumption that y_i is an integer. However, if we want to relax

$$\begin{aligned}
& \text{minimize} && \sum_{i \in \mathcal{I}} y_i \\
& \text{subject to} && \sum_{e=(v,h')} x_{e,h} = \sum_{e=(h',v)} x_{e,h} \quad \forall h, h' \in \mathcal{H}, h' \neq h \\
& && \sum_{e=(h,v)} x_{e,h} \geq w_h \quad \forall h \in \mathcal{H} \\
& && \sum_{e=(v,h)} x_{e,h} = 0 \quad \forall h \in \mathcal{H} \\
& && \sum_h x_{e,h} \leq \text{Cap}_e \quad \forall e \in E(G) \\
& && \sum_{h',e=(v,h)} x_{e,h'} \leq \text{Cap}_h \quad \forall h \in \mathcal{H} \\
& && \sum_{h',e=(v,i)} x_{e,h'} \leq \text{Cap}_i y_i \quad \forall i \in \mathcal{I} \\
& && \sum_{e=(v,i)} x_{e,h} \leq w_h y_i \quad \forall i \in \mathcal{I}, h \in \mathcal{H} \\
& && x_{e,h} \geq 0 \quad \forall e \in E(G), h \in \mathcal{H} \\
& && y_i \in \{0, 1, 2, \dots\} \quad \forall i \in \mathcal{I}
\end{aligned}$$

Figure 1: Integer program formulation for the ideal link model

the integrality assumption on y_i 's in order to derive a lower bound using an LP solver (see Section 4.3.4 for example), then it is important to include this inequality in the linear program to get a tighter lower bound.

The following theorem shows that it is computationally hard to optimally solve the ITAP placement problem for the ideal link model. Refer to [21] for the proof.

Theorem 1 *It is NP-hard to find a minimum number of ITAPs required to cover a neighborhood in an ideal link model. Moreover, the problem has no polynomial approximation algorithm with an approximation ratio better than $\ln n$ unless $P = NP$.*

4.1.2. Our Approach – Greedy Placement We design the following greedy placement. We iteratively pick an ITAP that maximizes the total demands satisfied when opened in conjunction with the ITAPs chosen in the previous iterations. The major challenge is to determine how to make a greedy move in each iteration. This involves efficiently computing the total user demands that can be served by a given set of ITAPs. We make an important observation: computing the total satisfied demands can be formulated as finding a maximum flow problem. This is easy to see since our formulation for the ideal link model shown in Figure 1 satisfies the three properties, namely, capacity constraint, skew symmetry, and flow conservation. This suggests that we can apply the network flow algorithms [9] to efficiently determine the satisfied demands. A few transformations are required to make the network flow algorithm applicable. Figure 2 shows a skeleton of the algorithm, which finds a multiset S of ITAPs to open, where a multiset is the same as a set, except that it allows duplicate elements. Allowing duplicate elements in S indicates that we can open multiple

Input:	Set of houses \mathcal{H} , set of ITAPs \mathcal{I} , graph G on the set $\mathcal{H} \cup \mathcal{I}$ with capacities on its edges and vertices.
Output:	A multiset S of ITAPs to be opened.
begin	
	$S := \emptyset; Flow := 0;$
while $Flow$ is less than the total demand do	
	$max := 0;$
	for each $j \in \mathcal{I}$ do
	<ul style="list-style-type: none"> • Let G' be the subgraph of G induced on $\mathcal{H} \cup S \cup \{j\}$, with the same capacities as G. (If there are duplicates in $S \cup \{j\}$, we create one point for each duplicated element.) • For each house, transform its vertex capacity constraint to an edge capacity constraint by replacing the house h with two nodes, in_h and out_h; and connect in_h to out_h using a directed edge with capacity cap_h; all incoming edges toward the house go to in_h and all out-going edges from h come from out_h. • Add two vertices s and t to G', edges of capacity w_h from s to each $h \in \mathcal{H}$, and edges of capacity cap_i from each $i \in S \cup \{j\}$ to t. • Find the maximum flow from s to t in G'; Let f be the value of this flow. • if $f > max$, then $max := f; bestITAP := j;$
	endfor;
	$S := S \cup \{bestITAP\}; Flow := max;$
	endwhile;
	end.

Figure 2: Greedy placement algorithm in the ideal link model

ITAPs in the locations that belong to the same equivalence class (i.e., reachable from the same set of houses), which is certainly feasible.

The following theorem shows a worst-case bound on the performance of the above algorithm. An empirical performance analysis of this algorithm is presented in Section 6.1.

Theorem 2 *Consider the ITAP placement problem in the ideal link model with integral demands and integral house and link capacities, and let D denote the total demand of the houses. The approximation factor of the greedy algorithm for this problem is at most $\ln(D)$. In other words, if the optimal solution for the ITAP placement problem opens K ITAPs, the greedy algorithm opens at most $K \ln(D)$ ITAPs.*

We will need the following lemma to prove the above theorem. This lemma is non-trivial and uses the Ford-Fulkerson maximum flow-minimum cut theorem [9] in the proof.

Lemma 3 *Assume a multiset S of ITAPs are opened. Consider an optimal way of routing the maximum total demands from houses to the ITAPs in S , and let f_i denote the amount of traffic routed to ITAP i in this solution, where $i \in S$. Assume that at a later time, a multiset $S \cup T$ of ITAPs are opened. Then, there is an optimal way of routing the maximum total demands from houses to these ITAPs in which f_i units of traffic is routed to ITAP i for every $i \in S$.*

Refer to [21] for the proofs of the above theorem and lemma. Based on Theorem 2, we have the following corollary.

Corollary 4 *Let N be the number of houses. The approximation factor of the greedy algorithm in the ideal link model is $\ln(N)$ when the capacities of edges and vertices are integer-valued and every house has either zero or one unit of demand.*

Remark 1. Corollary 4 in combination with Theorem 1, shows that this algorithm achieves the best possible (worst-case) approximation ratio for the graph theoretic model when every house has either zero or one unit of demand. Furthermore, even though in our model we allow fractional routing of the flow, our greedy algorithm always finds an integer solution in this case, i.e., the demand from each house will be served through one path to an open ITAP. This is a consequence of the integrality theorem [9].

Remark 2. Note that $\ln(D)$ is the *worst-case* bound for heterogeneous demands. To make the worst-case bound tighter, we can normalize house demands, edge capacities, and node capacities before we apply the greedy placement algorithm. This yields a lower approximation factor, since D is reduced after normalization. Moreover, as we will show in Section 6.1, in practice the greedy algorithm performs quite close to the optimal, and much better than the worst-case bounds, $\ln(D)$ or $\ln(N)$.

4.2. General Link Model

The problem of efficient ITAP placement is more challenging when the throughput along a path varies with the path length. This corresponds to the general link model introduced in Section 3.1.

4.2.1. Problem Formulation: We formulate the placement problem for the general link model as an integer program shown in Figure 3. Here $x_{e,h,l,j}$ denotes the total amount of flow routed from house h to the ITAPs using a path of length l when edge e is the j 'th edge along the path. Variable y_i is an indicator of the number of ITAPs opened in the equivalence class i , and each house h has w_h units of traffic to send. The throughput degradation function for a path of length l is denoted by $g(l)$. L is an upper bound on the number of hops on a communication path, and if there is no such upper bound, we set $L = |\mathcal{H}|$. The other variables in Figure 3 are similar to the ones used in Figure 1.

The following theorem is an immediate consequence of Theorem 1, as the ideal link model is a special case of the general link model, when $g(l) = 1$.

Theorem 5 *It is NP-hard to find a minimum number of ITAPs to cover a neighborhood for a general link model.*

$$\begin{aligned}
& \text{minimize} && \sum_{i \in \mathcal{I}} y_i \\
& \text{subject to} && \sum_{e=(v,h')} x_{e,h,l,j} = \sum_{e=(h',v)} x_{e,h,l,j+1} && \forall h, h' \in H, h' \neq h, \\
& && && l, j \in \{1, \dots, L\}, j < l \\
& && \sum_{e=(h,v), l} x_{e,h,l,1} \geq w_h && \forall h \in H \\
& && \sum_{h,l,j \leq l} g(l) x_{e,h,l,j} \leq \text{Cap}_e && \forall e \in E(G) \\
& && \sum_{h',e=(v,h), l,j \leq l} g(l) x_{e,h',l,j} \leq \text{Cap}_h && \forall h \in H \\
& && \sum_{h',e=(v,i), l,j \leq l} g(l) x_{e,h',l,j} \leq \text{Cap}_i y_i && \forall i \in I \\
& && \sum_{e=(u,i), l,j \leq l} x_{e,h,l,j} \leq w_h y_i && \forall i \in I, h \in H \\
& && x_{e,h,l,j} \geq 0 && \forall e \in E(G), h \in H, \\
& && && l, j \in \{1, \dots, L\}, j \leq l \\
& && y_i \in \{0, 1, 2, \dots\} && \forall i \in I
\end{aligned}$$

Figure 3: Integer program formulation for the general link model, where $g(l)$ models throughput degradation with increasing hop-count.

4.2.2. Our Approach of Greedy Placement: The high-level idea of the greedy algorithm is similar to the one presented for the ideal link model. We iteratively select ITAPs to maximize the total user demands satisfied. The new challenge is to determine a greedy move in this model. Unlike in the ideal link model, we cannot compute the total satisfied demands by modeling it as a network flow problem since the amount of flow now depends on the path length. As we will describe below, this computation can be done by solving a linear program, or by using a heuristic.

Expensive algorithm for the general link model: Without making assumptions about $g(l)$, we can compute the total satisfied user demands, for a given set I' of ITAPs, by solving a slightly modified LP problem than the one in Figure 3. In this linear program, we replace the variable y_i by the number of occurrences of i in I' (This amounts to removing all the variables corresponding to edges ending in ITAP positions outside I' and removing inequalities containing these variables). The objective will be to maximize $\sum_h \sum_{e=(h,v), l} x_{e,h,l,1}$, which corresponds to maximizing the satisfied demands. We also modify the second constraint to be $\sum_{e=(h,v), l} x_{e,h,l,1} \leq w_h$ in order to limit the maximum flow from each house h .

In theory, solving a linear program takes polynomial time. However, in practice an LP solver, such as cplex [10], can only handle small-sized networks under this model due to the fast increase in the number of variables and constraints with the network size.

Below we develop more efficient algorithms for two special forms of $g(l)$: (i) bounded hop-count: $g(l) = 1$ for all $l \leq k$, and $g(l) = \infty$ for $l > k$, and (ii) smooth degradation:

$g_l = l$ for all l .

Efficient algorithm for the bounded hop-count model: We can use the following greedy algorithm to find the total demands satisfied by a given set of ITAPs. The hop-count constraint suggests we should favor short paths in the graph. Therefore, in each iteration, the algorithm finds the shortest path from demand points to opened ITAPs in the residual graph, routes one unit of flow along this path, and decreases the capacities of the edges on the path by one in the residual graph. This is continued until the shortest path found has length more than the hop-count bound. This algorithm is similar to the algorithm proposed in [13]. While this heuristic does not guarantee computing the maximum flow (so each greedy step is not local optimal), it works very well in practice as shown in Section 6.2.1.

Efficient algorithm for the smooth throughput degradation model: When $g(l) = l$ or $throughput_l = \frac{1}{l}$, the total demands satisfied by a set of ITAPs are given by the expression: $maximize \sum_{p_i \in P} \frac{1}{|p_i|}$ where P is a collection of edge-disjoint paths in the graph, and $|p_i|$ denotes the length of the path p_i . Therefore to maximize this objective function, our heuristic should prefer imbalance in path lengths, and this motivates the following algorithm.

As the heuristic for the bounded hop-count model, in the smooth throughput degradation model we compute the total satisfied demands by the selected ITAPs through iteratively removing shortest paths in the residual graph. However, we make the following modifications. First, since we no longer have bounds on hop-count, we continue picking paths until there is no path between any demand point and any open ITAP. Second, to ensure the throughput follows $throughput(l) = 1/l$, after we obtain all the paths, we re-adjust the demand satisfied along each path p , denoted as SD_p , according to the throughput function. The total satisfied demands are the sum of SD_p over all paths p . Although this algorithm does not always find the maximum flow (so each greedy step is not local optimal), it yields very good performance as shown in Section 6.2.2.

4.3. Alternative Algorithms

We compare our greedy placement algorithm to four alternative approaches.

4.3.1. Augmenting Placement: The idea of the augmenting placement algorithm is similar to the greedy algorithm. The main difference is that in the augmenting algorithm we do not make a greedy move; instead we are satisfied with any ITAP that increases the total amount of demand satisfied. More specifically, we search over the set of possible ITAP locations, and open the first ITAP we see that results in an increase in the amount of satisfied demand when opened together with the already opened ITAPs.

The augmenting placement algorithm can be applied to all three wireless link models with the following difference.

In the ideal link model, we compute the total amount of demand satisfied under a given set of ITAPs by finding the maximum flow in the graph; in the general link models, we use the heuristics described in Section 4.2.2 to derive the total amount of demand satisfied.

4.3.2. Clustering-based Placement: We compare our placement algorithms to the clustering-based scheme, proposed in [3]. The basic idea of the algorithm is to partition the network nodes into a minimum number of disjoint clusters, and place an ITAP in each cluster. We use the Greedy Dominating Independent Set (DIS) [3] heuristic to determine a set of clusterheads, which are used as possible ITAP locations. The nodes are then clustered to ensure that each node is associated with the closest clusterhead, and a shortest path tree rooted at the clusterhead is used for sending packets from and delivering packets to the cluster. The cluster is further divided into sub-clusters if either the weight or relay-load constraints are violated. The weight constraint specifies that an ITAP can serve nodes as long as the sum of their demands does not exceed the capacity of the ITAP, and the relay-load constraint specifies an upper bound on the maximum flow that can go through a node in the neighborhood cluster. We refer the reader to [3] for more details of this algorithm.

To apply the clustering-based algorithm for the ideal link model, in our simulations we use the ITAP capacity instead of wireless capacity when checking the weight constraint of placing an ITAP at a particular house; this is necessary since the ITAP capacity can be greater than the wireless capacity in our simulations. This ensures a fair comparison of the clustering algorithm with our placement schemes.

To apply the algorithm to the bounded-hop count model, we make the following modification. We divide a cluster into sub-clusters not only when the weight or relay-load constraints are violated, but also when the distance between any node and its clusterhead exceeds the hop-count threshold. The algorithm, however, does *not* apply to the smooth throughput degradation model.

4.3.3. Random Placement: This algorithm randomly places an ITAP at a house iteratively until all the user demands are satisfied. To avoid wasting resource, it ensures that each house has at most one ITAP. This approximates un-coordinated deployment of ITAPs in a neighborhood, and gives a baseline to evaluate the benefits of the more sophisticated algorithms presented above.

As the greedy and augmenting algorithms, there are three variants of random placement algorithms for different wireless link models. They differ in how we compute the total demand satisfied under a given set of ITAPs. We run the maximum flow algorithm to compute the satisfied demand under the ideal link model, and apply the heuristics in Section 4.2.2 to compute the satisfied demand under the general link models.

4.3.4. Lower Bound: It is useful to compare our algorithms with the optimal solution. However, our problem is NP-hard, and it is too expensive to derive an optimal solution. Therefore we compare our algorithms with the lower bounds. We derive the lower bound by relaxing the integer constraints on y_i in the formulation in Figure 1 and solving the relaxed LP problem using cplex [10]. The lower bound is a useful data point to compare with, as it gives an upper bound on the difference between a practical algorithm and the optimal.

We use the same scheme to derive lower bounds for all three link models. For the general link models, the derived lower bounds are loose since the throughput degradation with hop count is ignored. Relaxing the integrality constraint in Figure 3, and solving the relaxed linear program could give tighter bounds, but it is computationally very expensive. As we will show in Section 6.2.1 and Section 6.2.2, the results from our greedy and augmenting algorithms are still close to these loose lower bounds.

5. Validation

To validate the wireless link models used in this paper, we run simulations in Qualnet [23], a commercial network simulator as follows. Given a neighborhood layout, the placement algorithms determine the ITAP locations and the set of paths each house uses to reach the ITAPs. We use the same neighborhood layout and ITAP locations in the simulations. Every node in the simulation uses an omnidirectional antenna and 802.11b MAC, with the communication range and interference range being 195 meters and 376 meters, respectively. Every house sends CBR traffic to the ITAPs at the rate specified by the placement algorithms' output. To support multi-path routing, we implement a probabilistic source-routing scheme in Qualnet, where the paths used in source routing and the probability that each path is chosen are based on the placement algorithms' output.

As shown in Figure 4, the ITAPs, determined using the smooth degradation model, satisfy the user demands to a great extent: around 80% houses have their demands completely satisfied when houses are randomly placed in $1000 \times 1000 m^2$, and all houses receive their demands when houses are randomly placed in $1500 \times 1500 m^2$. The better performance in the latter scenario comes from the fact that the larger separation among houses lowers interference among cross traffic. Note that even for the former case, we can further improve the clients' throughput by over-provisioning. As shown in the same figure, with over-provisioning (assuming that each user's demand is 500 Kbps when the actual demand is 208 Kbps), most of the clients' demands are satisfied.

Since ideal link and bounded hop-count models are more optimistic about the impact of interference, they are more




Figure 4: Validation of general link models: CDF of clients' throughput, where $N = 50$, $WC = 5Mbps$, and $w_h = 208Kbps \forall h \in H$.

suitable for the environments with efficient spectral use (e.g., when directional antennas are used). As part of our future work, we plan to evaluate how well these two models capture the impact of wireless interference under such environments.

6. Performance Evaluation

In this section, we evaluate the performance of different placement algorithms using various network topologies, house demands, and link models.

6.1. Performance Under the Ideal Link Model

First, we look at the performance under the ideal link model under various scenarios. We use the following notations in our discussion.

- N : the number of houses
- WC : a wireless link's capacity
- IC : an ITAP's capacity
- CR : communication radius
- HR : average inter-house distance
- w_h : house h 's demand

We compare the performance of different algorithms by varying each of the above parameters. In our evaluation, we use both random topologies and a real neighborhood topology. The random topologies are generated by randomly placing houses in a region of size $N \times N$, and varying the communication radius. The real neighborhood topology contains 105 houses, spanning over a region of $1106m \times 1130m$. (We cannot reveal the source of the data for confidentiality.) The average inter-house distance in the real topology is 74 meters. Unless otherwise specified, for the same parameter setting, we run simulations three times, and report the average number of ITAPs required for each placement algorithm.

Effects of the communication radius: We start by examining the effect of communication radius (CR) on the




Figure 5: Ideal link model: varying communication radius, where $N = 100$, $WC = 6$, $IC = 100$, and $w_h = 1 \forall h \in H$.

placement algorithms. It is easy to see from the problem formulation that only the ratio, $\frac{CR}{HR}$, is important. Therefore in our evaluation, we vary the communication radius from 1 to 50, while fixing the inter-house distance by randomly placing 100 houses in an area of 100×100 , which yields an average inter-house distance of 4.5 - 6.

Figure 5 illustrates the number of ITAPs required with varying CR . We make the following observations. First, we see that an increase in CR results in a greater overlap of wireless coverage of the houses, and therefore fewer ITAPs are sufficient to satisfy the house demands. Second, comparing the performance across different algorithms, we observe that the greedy algorithm performs very close to the lower bound over all cases. Interestingly, the augmenting algorithm performs quite well, too. The good performance of the augmenting algorithm comes from the requirement that new ITAPs should lead to throughput improvement, which avoids wasting resource on the already covered region. This is especially useful after several ITAPs have been placed, since at this point only a few locations remain that can further increase the satisfied demands.

In contrast, the clustering and random-house placement schemes perform much worse. Compared with the greedy strategy, both schemes often require 2 to 10 times as many ITAPs. Note that when the communication radius is very large, the clustering algorithm yields worse performance than the random-house placement. This is because in the clustering algorithm data dissemination follows a shortest path tree, instead of maximizing the total amount of flow that can be pushed to the ITAPs. In comparison, the other algorithms, including the random-house placement, run the maximum flow algorithm to maximize the total satisfied demands.

Effects of network size: Next we study the impact of network size on the placement algorithms. We randomly place N houses in an $N \times N$ area while fixing the communication radius to 10. Figure 6 shows the number of required ITAPs using the different placement algorithms for various




Figure 6: Ideal link model: varying the number of nodes, where $CR = 10$, $WC = 6$, $IC = 100$, and $w_h = 1 \forall h \in H$.

network sizes. As we would expect, an increase in the number of houses leads to a larger number of ITAPs required to cover the neighborhood. Moreover, the greedy algorithm continues to perform very well, with its curve mostly overlapping with the lower bound. The augmenting algorithm performs slightly worse, whereas the clustering and random algorithms perform much worse – requiring up to 5 and 8 times as many ITAPs, respectively. In addition, the benefit of greedy algorithm increases as the network gets larger.

Effects of wireless link capacity: We also study the effects of wireless bandwidth on the placement algorithms. We observe that the relative ranking of the algorithms stays the same. The effect of bandwidth is only pronounced when it is very limited. For example, when the wireless bandwidth is equal to a single house's demand, the number of ITAPs required is considerably large. As the bandwidth increases and the wireless link is no longer the bottleneck, the number of required ITAPs remains the same with a further increase in the wireless link capacity.

Effects of the ITAP capacity: We compare the placement algorithms by varying the ITAP capacity. When ITAP capacity is small and hence is a bottleneck, the number of required ITAPs decreases proportionally with an increase in ITAP capacity. As the ITAP capacity is large enough and no longer the bottleneck, the number of required ITAPs is unaffected by a further increase in ITAP capacity. Moreover, the relative performance of different placement algorithms is consistent with the previous scenarios.

Effects of heterogeneous house demands: So far we have considered homogeneous house demands (i.e., each house generates one unit demand). A number of studies show that realistic user demands are very heterogeneous, and often exhibit Zipf-like distributions [5, 6]. Motivated by these findings, below we evaluate the placement algorithms when house demands follow a Zipf distribution. The results are qualitatively the same as those of using the homogeneous house demands. The greedy algorithm continues to out-perform the others significantly and yield nearly optimal solutions.




Figure 7: Ideal link model: real neighborhood topology with various communication radii; $N = 105$, $WC = 6$, $IC = 100$; the house demands follow a Zipf distribution.

Real neighborhood topology: Finally we evaluate the placement algorithms using a real neighborhood topology of 105 houses. We again use Zipf-distributed house demands. As shown in Figure 7, initially when the communication range is too small, most houses are unreachable from other houses, and therefore all the algorithms require close to 105 ITAPs. As the communication range increases, fewer ITAPs are needed to cover the neighborhood. At the extreme, when the communication range reaches 250 meters, the neighborhood forms a single connected component, and therefore most algorithms require only one ITAP. (When bandwidth is more limited or the impact of wireless interference is larger as shown in the next section, we may need multiple ITAPs even for a single connected component.) Moreover, the greedy algorithm performs close to optimal over all communication radii considered.

6.2. Performance Under the General Link Models

In this section, we evaluate the performance of placement algorithms under two general link models, namely bounded hop-count and smooth throughput degradation models.

6.2.1. Bounded Hop-count Model: We compare the placement algorithms for bounded-hop count model by varying the hop-count threshold, communication radius, and neighborhood topology.

Effects of hop-count threshold: First we compare the placement algorithms by varying the hop-count threshold. As shown in Figure 8, when the hop-count threshold increases, the effect of hop-count reduces, since all or most paths are within hop-count limit. Comparing the different placement algorithms, we see that the greedy placement performs very close to the lower bound, especially for large hop-count thresholds. When the hop-count threshold is small, the gap between the lower bound and greedy

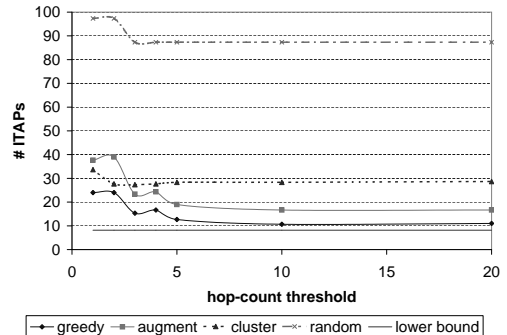


Figure 8: Bounded hop-count model: varying the hop-count threshold, where $N = 100$, $CR = 10$, $WC = 6$, $IC = 100$, and $w_h = 1 \forall h \in H$.




Figure 9: Bounded hop-count model: a real neighborhood topology for various communication radii, where $N = 105$, $WC = 6$, $IC = 100$, hop-count threshold = 3, and the house demands follow a Zipf distribution.

algorithm is slightly larger, since the lower bound ignores throughput degradation with the hop-count, and is not as tight. Compared with the greedy algorithm, the augmenting algorithm requires 50% more ITAPs; the clustering algorithm in [3], requires 2 - 3 times as many ITAPs; and the random algorithm requires 4 to 8 times as many ITAPs.

Effects of communication radius: Next we fix the hop-count threshold to 3, and vary the communication radius. As expected, an increase in communication radius reduces the number of ITAPs required to cover the neighborhood. Moreover the greedy continues to perform significantly better than the alternatives. We observe similar results for other hop-count thresholds.

Real neighborhood topology: We also evaluate the placement algorithms using the real neighborhood topology. As shown in Figure 9, the results are qualitatively the same as the random topologies. The greedy algorithm performs very close to the lower bound for all the communication radii considered.

6.2.2. Smooth Throughput Degradation Model: Next we empirically study the placement algorithms for the




Figure 10: Smooth throughput degradation model: varying the communication radius, where $N = 100$, $WC = 6$, $IC = 100$, and $w_h = 1 \forall h \in H$.

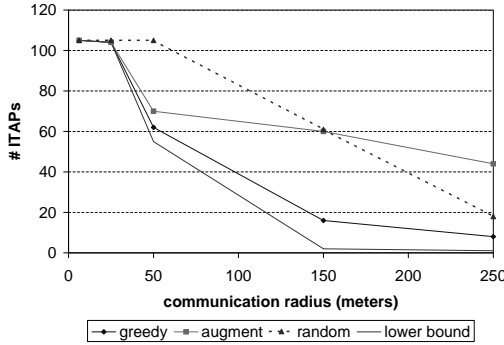


Figure 11: Smooth throughput degradation model: a real neighborhood topology for various communication radii, where $N = 105$, $WC = 6$, $IC = 100$, and the house demands follow a Zipf distribution.

smooth throughput degradation model.

Effects of communication radius: Figure 10 shows that the number of required ITAPs decreases with increasing communication radius. The gap between the performance of different algorithms is the largest when the communication radius is between 5 and 20 (The average inter-house is around 5.). This can be explained as follows. When the radius is very small, most houses are disconnected from one another. So the number of ITAPs required is nearly the number of houses. When the radius is very large, most houses are reachable from one another within one or few hops, and the number of ITAPs required is close to 1. However, for the practical scenario with medium communication radius, the gap between different algorithms is most significant. Note that the lower bound, which is derived by ignoring the impact of hop-count on throughput, is more loose for this scenario. Even then the greedy is still competitive when compared with these loose lower bounds.

Real neighborhood topology: Figure 11 shows the results from the real neighborhood topology. As we can see, the greedy placement continues to perform much better than the alternatives.

7. Practical Considerations

We now present algorithms for handling two practical requirements: providing fault tolerance and handling workload variation.

7.1. Providing Fault Tolerance

A practical solution to the ITAP placement problem should ensure Internet connectivity to all the houses in the neighborhood, even in the presence of a few ITAP and house failures. Here we present an enhancement to our algorithm by incorporating this fault tolerance constraint. Fault tolerance is achieved by providing multiple node disjoint paths from a house to ITAPs, and over-provisioning the delivery paths. Over-provisioning is a scheme that allocates more flow to a house than its demand, and therefore helps to provide QoS guarantees even when there are a few failures. Note that we consider node disjoint paths instead of edge disjoint paths to avoid correlated link failures caused by a problematic node. Also since the ultimate goal is to provide Internet connectivity irrespective of which ITAP is used, we allow a user's demand to route to different ITAPs.

7.1.1. Problem Formulation: Let each house have one unit of demand, and d independent paths to reach the ITAPs; the average failure probability of a path be p ; and the over-provisioning factor be f (i.e., each independent path allocates $\frac{f}{d}$ capacity to a house, and the total capacity allocated to a house by d independent paths is f).

Since for every house, there are d independent paths from this house to ITAPs and the probability of failure of each path is p , the probability that exactly i of these paths fail is $\binom{d}{i} p^i (1-p)^{d-i}$, assuming independent path failures. In this case, the amount of traffic that can be delivered is $\min(\frac{(d-i)f}{d}, 1)$. Therefore, the expected fraction of the traffic from a house that can reach an ITAP, $S(f, p, d)$, is given by the following formula.

$$S(f, p, d) = \sum_{i=0}^d \binom{d}{i} p^i (1-p)^{d-i} \min\left(\frac{(d-i)f}{d}, 1\right)$$

Given the expected guarantee desired by the home users, $S(f, p, d)$, we can use the above expression to derive the overprovision factor, f , based on path failure probability and the number of independent paths. We now provide fault tolerant LP formulations for the ideal and general link models.

Ideal Link Model with the Fault Tolerance Constraint : Figure 12 provides an LP formulation of the fault tolerant problem for the ideal case, i.e. when throughput is independent of the path length. For each edge e and each house h , the variable $x_{e,h}$ indicates the amount of flow from h to ITAPs that is routed through e . Also, for each ITAP

$$\begin{aligned}
& \text{minimize} && \sum_{i \in \mathcal{I}} y_i \\
& \text{subject to} && \sum_{e=(v,h')} x_{e,h} = \sum_{e=(h',v)} x_{e,h} \leq w_h \quad \forall h, h' \in \mathcal{H}, h' \neq h \\
& && \sum_{e=(h,v)} x_{e,h} - \sum_{e=(v,h)} x_{e,h} \geq w_h d \quad \forall h \in \mathcal{H} \\
& && \frac{f}{d} \sum_{e \in E(G)} x_{e,h} \leq \text{Cap}_e \quad \forall e \in E(G) \\
& && \frac{f}{d} \sum_{h', e=(v,h)} x_{e,h'} \leq \text{Cap}_h \quad \forall h \in \mathcal{H} \\
& && \frac{f}{d} \sum_{h', e=(v,i)} x_{e,h'} \leq \text{Cap}_i y_i \quad \forall i \in \mathcal{I} \\
& && \sum_{e=(v,i)} x_{e,h} \leq w_h y_i \quad \forall i \in \mathcal{I}, h \in \mathcal{H} \\
& && x_{e,h} \geq 0 \quad \forall e \in E(G), h \in \mathcal{H} \\
& && y_i \in \{0, 1, 2, \dots\} \quad \forall i \in \mathcal{I}
\end{aligned}$$

Figure 12: LP formulation for the ideal link model with fault tolerance constraints, where d is the number of independent paths, and f is the over-provision factor.

i , the variable y_i denotes the number of ITAPs opened in equivalence class i . The above integer LP is similar to the one in Section 4.1.1. The differences are as follows: (i) the constraint $\leq w_h$ added to the first inequality, (ii) a change in the second constraint from w_h to $w_h d$ in the amount of flow originating from each house, and (iii) a multiplicative factor of $\frac{f}{d}$ on the left-hand side of the capacity constraints (since the amount of capacity each path allocates to each house is $\frac{f}{d}$). The first modification ensures that the flow from each house is served by independent paths; (ii) and (iii) are for the over-provisioning purpose.

Similarly we can formulate integer LP for general link models with the fault tolerance constraint. Refer to [21] for details. For all link models, we have Theorem 6.

Theorem 6 *It is NP-hard to find a minimum number of ITAPs required to cover a neighborhood while providing fault tolerance.*

7.1.2. Placement Algorithms: The greedy, augmenting and random placement algorithms are based on the same idea described in the previous sections of this paper. However, they differ in the way they compute the total demands supported by a given set of ITAPs. For the ideal case, we compute the satisfied demands by slightly modifying the LP in Figure 12, and solving the resulting LP. The objective function is changed to be maximizing $\sum_h (\sum_{e=(h,v)} x_{e,h} - \sum_{e=(v,h)} x_{e,h})$, which corresponds to maximizing the supported demands. The variables y_i are replaced by the number of occurrences of i in I' . Furthermore, the second constraint is changed to $\sum_{e=(h,v)} x_{e,h} - \sum_{e=(v,h)} x_{e,h} \leq w_h d$ in order to limit the maximum flow from a node. For the general link model, we compute the satisfied demands by applying similar modifications to the corresponding integer linear program.

We compare the above algorithms to the lower bound, derived by relaxing the integrality constraint and solving the relaxed linear program. Our evaluation [21] shows that for all algorithms the number of ITAPs required increases linearly with the number of independent paths. Moreover, the results of the greedy algorithm are very close to the lower bound, and significantly better than the other two.

7.2. Handle Workload Variation

In practice, user demands change over time, and often exhibit diurnal patterns [6, 17, 24]. Since it is not easy to change ITAP locations once they are deployed, a good ITAP placement should handle demands over all periods. In this section, we describe and evaluate two approaches to handle variable workloads. While our discussion focuses on the non fault-tolerant version of the placement problems, the ideas carry over easily to the fault-tolerant version as well.

One approach to take into account workload change is to provision ITAPs based on the peak workload. That is, if $w[h][t]$ denotes the demand of house h at time t , we use $\max_t w[h][t]$ as the demand for house h , and feed this as an input to the placement algorithms described in the previous sections. We call this approach *peak load based placement*. This algorithm is simple, but sometimes inefficient, e.g., when different houses' demands peak at different times.

To improve efficiency without sacrificing user performance, we now explore how to optimize ITAP locations for demands over multiple time intervals. More formally, the problem can be stated as follows. Each house h has demand $w[h][t]$ at time t . Our goal is to place a set of ITAPs such that at any time t , they can serve all the demands generated at t , i.e., $w[h][t]$ for all h 's.

Here we describe a greedy heuristic with a logarithmic worst-case bound for the ideal link model. The same idea applies to other link models. The high-level idea is to iteratively place the ITAP such that together with the already opened ITAPs it maximizes the total demands served. Unlike in the previous section, here the total demands include demands over multiple time intervals. More specifically, we place an ITAP such that it maximizes $\sum_{t \in T} SD_t$, where SD_t is the total satisfied demands at time t . This can be computed by changing the greedy algorithm of Section 4.1.2 as follows. In every iteration, for every $j \in \mathcal{I}$ and $t \in T$, we construct the graph G' as in the algorithm of Section 4.1.2 based on the demands at time period t . Then we compute the maximum flow $f_{j,t}$ in this graph. After these computations, we pick the ITAP j that maximizes $\sum_t f_{j,t}$, and open it. We call this algorithm *multiple-demand-based greedy placement* (M-greedy, for short). In the following theorem, we show a worst-case bound on the M-greedy's performance in the ideal link model.

Theorem 7 *Consider the ITAP placement problem in the ideal link model with integral demands and capacities, and*

let D_t be the total demand in period t . The approximation factor of the M -greedy algorithm for this problem is at most $\ln(\sum_t D_t)$. In other words, if the optimal algorithm requires K ITAPs to serve demands over L time periods, then the M -greedy requires at most $K \ln(\sum_t D_t)$ ITAPs.

Refer to [21] for the proof. Based on the above theorem, we have the following corollary.

Corollary 8 Let L denote the total number of periods, and N denote the number of houses. The approximation factor of the M -greedy in the ideal link model is $\ln(LN)$, when the capacities of edges and vertices are integer-valued and every house has either zero or one unit of demand at any time t .

This is easy to see because $\sum_t D_t \leq LN$.

The approximation factor of the greedy placement using the peak load is at most a factor of $\sum_t (\ln D_t)$ (refer to [21] for details). When $D_t = D_m$ for all t 's, its cost is at most $L \ln(D_m)$. This is roughly L times the approximation factor of the M -greedy algorithm.

In addition to the worst-case analysis, we evaluate the effectiveness of the algorithms empirically, and observe that the number of ITAPs required to serve demands exhibiting diurnal patterns using M -greedy is only slightly higher than the maximum number of ITAPs required to serve either of the two periods. In the interest of space, we refer readers to [21] for details.

8. Conclusions

In this paper we look at the problem of efficient ITAP placement to provide Internet connectivity in multi-hop wireless networks. We make three major contributions in this paper. First, we formulate the ITAP placement problem under various wireless models, and design algorithms for each model. Second, we address several practical issues when using these algorithms. In particular, we extend the placement algorithms to provide fault tolerance and to handle variable user demands. These two enhancements improve robustness of our placement schemes in face of failures and demand changes. Third, we demonstrate the efficiency of our placement algorithms using analysis and simulations, and show that the greedy algorithms give close to optimal solutions over a variety of scenarios we have considered. To our knowledge this is the first paper that looks at the ITAP placement problem for general MAC schemes.

References

- [1] 802.11b community network list. <http://www.toaster.net/wireless/community.html>.
- [2] S. Basagni, M. Conti, S. Giordano, and I. Stojmenovi, editors. *Mobile Ad Hoc Networking*. Wiley-IEEE Press, Jul. 2004.
- [3] Y. Bejerano. Efficient integration of multi-hop wireless and wired networks with QoS constraints. In *ACM MOBICOM*, Sep. 2002.
- [4] D. A. Beyer, M. Vestrch, and J. J. Garcia-Luna-Aceyes. The rooftop community network: Free, High-Speed network access for communities. In *The First 100 Feet: Options for Internet and Broadband Access*. MIT Press: Cambridge, 1999.
- [5] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker. Web caching and Zipf-link distributions: Evidence and implications. In *Proc. of IEEE INFOCOM 99*, Mar. 1999.
- [6] M. Chesire, A. Wolman, G. M. Voelker, and H. M. Levy. Measurement and analysis of a streaming-media workload. In *Proc. of 3rd USENIX Symposium on Internet Technologies and Systems*, Mar. 2001.
- [7] R. R. Choudhury, X. Yang, R. Ramanathan, and N. H. Vaidya. Using directional antennas for medium access control in ad hoc networks. In *Proc. of ACM MOBICOM*, Sept. 2002.
- [8] Wireless networking reference - community wireless/rooftop systems. http://www.practically-networked.com/tools/wireless_articles_community.htm.
- [9] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms (Second Edition)*. MIT Press, Cambridge, 2001.
- [10] CPLEX version 7.5.0. <http://www.ilog.com/products/cplex/>.
- [11] T. ElBatt and B. Ryu. On the channel reservation schemes for ad-hoc networks: Utilizing directional antennas. In *5th IEEE Int'l Symposium on Wireless Personal Multimedia Communications*, 2002. http://www.wins.hrl.com/projects/adhoc/d/elbatt_wpmc02.pdf.
- [12] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2), Mar. 2000.
- [13] V. Guruswami, S. Khanna, R. Rajaraman, B. Shepherd, and M. Yannakakis. Near-optimal hardness results and approximation algorithms for edge-disjoint paths and related problems. In *Proc. of the 31st ACM Symposium on Theory of Computing (STOC)*, pages 19–28, May 1999.
- [14] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. In *Proc. ACM MOBICOM*, Sept. 2003.
- [15] S. Jamin, C. Jin, Y. Jin, D. Raz, Y. Shavitt, and L. Zhang. On the placement of Internet instrumentation. In *IEEE INFOCOM*, Mar. 2000.
- [16] S. Jamin, C. Jin, A. Kurc, D. Raz, and Y. Shavitt. Constrained mirror placement on the Internet. In *Proc. of IEEE Infocom*, Apr. 2001.
- [17] D. Kotz and K. Essien. Analysis of a campus-wide wireless network. In *Proc. of ACM MOBICOM*, Sept. 2002.
- [18] A. Kuehn and M. Hamburger. A heuristic program for locating warehouses. *Management Science*, 9:643–666, 1963.
- [19] B. Li et al. On the optimal placement of Web proxies in the Internet. In *Proc. of IEEE INFOCOM*, Mar. 1999.
- [20] L. M. S. C. of the IEEE Computer Society. Wireless LAN medium access control (MAC) and physical layer (PHY) specifications. *IEEE Standard 802.11*, 1999.
- [21] L. Qiu, R. Chandra, K. Jain, and M. Mahdian. On the placement of integration points in multi-hop wireless networks. *MSR Technical Report MSR-TR-2004-43*, 2004.
- [22] L. Qiu, V. N. Padmanabhan, and G. Voelker. On the placement of Web server replicas. In *Proc. of IEEE Infocom*, Apr. 2001.
- [23] Qualnet. <http://www.scalable-networks.com/>.
- [24] S. Saroiu, K. P. Gummadi, R. J. Dunn, S. D. Gribble, and H. M. Levy. An analysis of Internet content delivery systems. In *Proc. of the 5th Symposium on Operating Systems Design and Implementation (OSDI'2002)*, Dec. 2002.
- [25] V. V. Vazirani. *Approximation Algorithms*. Springer-Verlag, 2001.