Bandwidth Provisioning in Infrastructure-based Wireless Networks Employing Directional Antennas

Shiva Prasad Kasiviswanathan[†], Sudarshan Vasudevan[‡], Bhuvan Urgaonkar[†]

[†]Dept. of Computer Science and Engineering
Pennsylvania State University
University Park, PA 16802
{kasivisw,bhuvan}@cse.psu.edu

[‡]Dept. of Computer Science
University of Massachusetts
Amherst, MA 01003
svasu@cs.umass.edu

Abstract—Wireless networks employing *directional* antennas are expected to proliferate. We study the problem of provisioning the subscribers of such networks in two fundamental settings: (i) *subscriber-centric*, where the objective is to fairly allocate bandwidth among the subscribers and (ii) *provider-centric*, where the objective is to maximize the revenue that the network provider generates from the subscribers. Using the well-regarded max-min fair allocation as the objective for a subscriber-centric network, we present an $O(n^4m)$ time dynamic programming algorithm for a network with m antennas mounted on a fixed access point and n subscribers. For a provider-centric network, we show that provisioning is NP-hard. For this problem, we present a linear-time algorithm that guarantees close to half of the optimum revenue. Finally, we discuss the complexities of the above problems in the presence of multiple access points.

I. INTRODUCTION

There has been a growing interest in the use of directional antennas for communication networks [6, 7, 22]. Unlike omni-directional antennas, which transmit in all directions simultaneously, directional antennas can transmit in a particular direction over a smaller area. For the same amount of power, directional antennas provide longer range and greater spatial reuse in comparison with omni-directional antennas. Several studies point out that directional antennas have the potential to improve the capacity of wireless networks [21, 24]. Furthermore, there have been several recent efforts [17] to build low-cost steerable directional antennas. It can therefore be imagined that directional antennas will be widely used in wireless networks. However, directional antennas introduce some unique challenges. In particular, each node has to determine the direction in which its antenna is positioned before communicating with other nodes.

There have been several proposals to handle the challenges introduced by usage of directional antennas in the context of neighbor discovery [23], medium access control [2, 18], and routing [6, 7] in multi-hop wireless networks. In this paper, we consider challenges introduced by using directional antennas in infrastructure-based wireless networks such as cellular networks and IEEE 802.11 WLANs. In these networks, wireless nodes affiliate themselves with *access points*, each of which typically has a high-speed connection to the rest of the network. There have been some recent proposals [25] which experimentally demonstrate the performance benefits

provided by using multi-radio access points. These works, however, assume that access points have omni-directional antennas. Our work, on the other hand, considers the case when each access point is equipped with multiple directional antennas, each using an orthogonal channel for communication with the subscribers affiliated to it. Our objective here is to design efficient algorithms that determine the positions of the antennas covering all the subscribers while meeting their performance requirements. It is easy to see that the problem of covering subscribers is trivial in the case of omni-directional antennas, but as we will see in this paper, the problem becomes significantly challenging when access points have directional antennas.

We consider two problem formulations in this paper. In the first problem formulation, that we call *subscriber-centric provisioning*, we are given a set of subscribers affiliated with an access point equipped with multiple directional antennas. Each antenna has an associated bandwidth that is to be shared among all the subscribers associated with that antenna. The problem here is to choose an assignment of subscribers to the antennas in a manner that ensures a max-min fair distribution of bandwidth to the affiliated subscribers i.e., the resulting assignment is such that the bandwidth allocated to any subscriber cannot be increased without reducing the bandwidth of a subscriber with an already smaller allocation [4]. Max-min fairness has been extensively used to implement QoS in both wired and wireless domains [14, 16]. This formulation captures the bandwidth provisioning problem that arises in wireless networks where the subscribers form a community or group that collectively pays the network provider for the aggregate bandwidth it provides. Furthermore, there exists an uniformity in how the quality of the network connectivity that the subscribers receive translates into their productivity, implying that it is desirable for the access point to align its antennas so as to achieve close to equal allocations of bandwidth among the subscribers.

In the second formulation, *provider-centric provisioning*, we consider the problem in which the service provider desires to maximize its revenue. In particular, each subscriber specifies a bandwidth requirement and if the requirement gets satisfied, pays the provider a cost proportional to the requirement. A subscriber does not pay the provider if its requirement is not satisfied. The problem here is to choose an assignment of

subscribers to the antennas that maximizes the revenue of the service provider. This formulation captures scenarios where the subscribers are independent entities, each interested in the quality of its own network connectivity and willing to pay for receiving it. In a recent paper [3] (some of) the authors considered a related problem of minimizing the number of antennas for satisfying all customer demands.

Both these problems are related to geometric covering problems, where we have a set of points in the plane and we want to cover them with (minimum) objects of some shape [8, 13]. In this paper, the shapes of sets are related to possible ranges of directional antennas. We make the following contributions in this paper:

- 1) We develop a polynomial-time algorithm that achieves the max-min fair allocation for the subscriber-centric model. The algorithm is based on dynamic programming and runs in $O(n^4 m)$ time.
- 2) We show that revenue-maximization for the provider-centric model is NP-hard. We develop an approximation algorithm that guarantees close to half of the optimal revenue. The algorithm is based on a greedy heuristic and runs in time linear in number of subscribers.

The rest of the paper is organized as follows. In Section II, we describe our network model and formulate the two bandwidth provisioning problems. In Section III, we describe a dynamic programming algorithm that achieves a max-min fair allocation of bandwidth among subscribers. In Section IV, we describe a greedy heuristic with a ≈ 2 -approximation ratio for the provider-centric provisioning problem. In Section V, we formulate the bandwidth provisioning problems in the presence of multiple access points and describe the complexities of these problems. Finally, we conclude with some future research directions in Section VI.

II. NETWORK MODEL AND PROBLEM DEFINITIONS

We begin this section by describing the wireless network model assumed in our research. The model we consider in this paper is similar to one developed in [3]. Next, we formally define two bandwidth provisioning problems that we study in the remainder of this paper under the assumptions of this model.

A. Background and Network Model

Our initial problem setting consists of an AP (access point) which has n wireless subscribers. In Section V, we consider a network with multiple APs. The AP is assumed to be connected to a wide-area network infrastructure. The AP functions as a gateway relaying traffic to and from the wireless subscribers. Each AP is equipped with m radios, each tuned to a separate orthogonal channel. Each channel has a bandwidth of b bits/second. Furthermore, the AP has m directional antennas (one for each radio) for communication with its subscribers. Each antenna has a range of R meters, a beamwidth (or angular range) $\rho (\leq \pi)$ and is capable of being steered (either mechanically or electronically) in a direction chosen by the AP. Similar to [3, 22, 24], we model the antenna pattern resulting from using such a directional antenna as a

Fig. 1. The figure to the left shows the model. A directional antenna with angular range ρ and reach R is mounted on the AP. The subscribers are spread on the plane around the AP. In the figure to the right, the shaded subscribers are in the set $Nbr_{right}(i)$ for a subscriber i .

circular sector with angle ρ . All subscribers are assumed to be at most R meters away from the AP and each subscriber is always assigned to at most one antenna. See Fig. 1.

In this paper, we assume that all the subscribers are static and their locations are known to the AP. Although these assumptions may be simplistic for networks with highly mobile subscribers, our research is a worthwhile exercise for the following two reasons. First, the theoretical bounds that our research provides for two fundamental provisioning problems apply equally well to networks with mobility. We believe that our algorithms are an essential first step towards developing provisioning algorithms for networks with mobile subscribers. Secondly, despite the popularity of networks with mobile subscribers such as cellular networks and ad hoc networks, infrastructure-based networks with little or no mobility are likely to continue being a significant class of networks. Examples of such networks include IEEE 802.11 WLANs in office buildings and universities providing connectivity to subscribers. In fact, there have been several recent measurement studies [1, 19] indicating that most of the wireless subscribers in these networks show very little or no mobility. Users in these workplaces typically connect to the network from their office desks, lab spaces, or cafeteria tables and are often seated in one location for prolonged periods with intermittent and limited episodes of mobility. The provisioning algorithms developed in our research are directly applicable to such networks.

We will deal with points in radial co-ordinates and with circular orderings. The origin of our co-ordinate system lies at the AP. The point (θ, r) is equivalent to the Cartesian point $(r \times \sin \theta, r \times \cos \theta)$. Equality of angles is understood modulo 2π . We denote the set of n subscribers by $U = \{1, \dots, n\}$ and the location of subscriber i by (θ_i, r_i) . A directional antenna is characterized by four parameters. The first three parameters, namely reach R , bandwidth b , and angular range ρ have already been introduced. The fourth parameter denotes the *direction* that the AP steers the antenna in, denoted by the radial angle α in our co-ordinate system. A directional antenna with parameters (R, b, ρ, α) can provide connectivity to subscribers with radial coordinates (θ, r) such that $\alpha \leq \theta \leq \alpha + \rho$ and $r \leq R$; we say that the antenna can *cover*

Fig. 2. The directed graph (U, E') for a set of subscribers. The graph captures spatial closeness.

these subscribers. Henceforth, we assume the bandwidth of each antenna to be 1 bandwidth unit¹; all variables denoting bandwidth are assumed to be scaled appropriately. For any subscriber i define the set of its *right neighbors* $Nbr_{right}(i)$ and the set of its *left neighbors* $Nbr_{left}(i)$ as follows.

$$\begin{aligned} Nbr_{right}(i) &= \{j \in U : \theta_j = \theta_i + \beta \text{ with } 0 \leq \beta \leq \rho\} \\ Nbr_{left}(i) &= \{j \in U : \theta_j = \theta_i - \beta \text{ with } 0 \leq \beta \leq \rho\} \end{aligned}$$

$Nbr_{right}(i)$ consists of all subscribers whose angular separation from i is at most ρ in the counter-clockwise direction and $Nbr_{left}(i)$ represents the same in the clockwise direction. Fig. 1 illustrates these concepts. We assume without loss of generality (w.l.o.g.) that θ_i is different for each $i \in U$. Our results are not affected by this simplification.

To capture spatial closeness of subscribers, we define a directed graph (U, E) where we have an edge $(i, j) \in E$ iff $Nbr_{right}(i) \cap Nbr_{left}(j) = \{i, j\}$. Therefore, the existence of the edge (i, j) implies that j follows i if in the counter-clockwise direction. Additionally, we say that subscribers i and j are *contiguous* if $(i, j) \in E$. If (U, E) has disconnected components, we introduce edges connecting each such component to the one next to it in the counter-clockwise direction. The resulting graph (U, E') , therefore, has a simple cycle. Let $path(a, b)$ denote the path from a to b in this graph. We use the function *successor* defined on E' as follows: for an element (i, j) of E' , $successor(i) = j$. See Fig. 2 for an example of this graph.

Also, we define for any set of subscribers S a *neighborhood* $NbrHood(S)$ as follows:

$$NbrHood(S) = Nbr_{right}(a) \cap Nbr_{left}(b)$$

if there exist $a, b \in S$ such that $S \subseteq Nbr_{right}(a) \cap Nbr_{left}(b)$. Intuitively, for a set S , its neighborhood $NbrHood(S)$ contains all the subscribers lying geographically between the first and last subscribers in S in an ordering of the members of S in the counter-clockwise direction. $NbrHood(S)$ is defined for a set S iff any two subscribers in S are within ρ of each other. For set S , we define $first_item(S)$ as the subscriber j such that for some subscriber k we have

¹ 1 bandwidth unit = b bits/second.

Fig. 3. The shaded subscribers are part of set S . The $NbrHood(S)$ contains all subscribers between $first_item(S)$ and $last_item(S)$.

$NbrHood(S) = Nbr_{right}(j) \cap Nbr_{left}(k)$. The subscriber k is defined as $last_item(S)$. See Fig. 3.

We define $first_item(S)$ as the subscriber a and $last_item(S)$ as the subscriber b .

Next, we formally define the two bandwidth provisioning problems that our research addresses.

B. Subscriber-centric Provisioning

The first setting that we consider is a wireless network where the network provider is expected to provide some form of *fair* bandwidth allocation when provisioning its subscribers. Typically, in networks that desire such provisioning, the subscribers form a community or group that collectively pays the network provider for the aggregate bandwidth it provides. The notion of fairness itself may vary based on the environment. As an example, consider an AP providing network connectivity to the faculty and students occupying offices and lab spaces in the university buildings surrounding it. Given the locations of these wireless subscribers and assuming uniformity in how the quality of their network connectivity translates into their productivity, it would be desirable for the AP to align its antennas in order to achieve equal allocations of bandwidth among the subscribers. Other notions of fairness may be appropriate in certain networks. For example, a network may wish to provide *proportional-share fairness* [12, 20], whereby subscribers have weights signifying their importance and the AP tries to distribute bandwidth in proportion to these weights. In this paper, we consider fairness to be synonymous with equal allocation of bandwidth. This requirement is captured by the well-regarded notion of *max-min fairness* [4]. We call the problem of achieving a max-min fair bandwidth allocation as the subscriber-centric provisioning and denote it by SUB-CENTRIC.

We begin by formally defining the notion of max-min fairness. Given two k -tuples, $X = (x_1, \dots, x_k)$ and $Y = (y_1, \dots, y_k)$, each in non-decreasing order, we say that X lexicographically dominates Y if $X = Y$, or there exists some index j such that $x_j > y_j$ and $x_i = y_i$ for all $i < j$. A *bandwidth allocation vector* is a n -tuple denoting the allocation of bandwidth to the n subscribers. Given two bandwidth allocation vectors $B(F)$ and $B(F')$, we say that $B(F)$

Fig. 4. NP-hardness of PROV-CENTRIC. The reduction from MAX-BIN-PACKING puts all subscribers on one straight line.

is as fair as $B(F')$ if $B(F)$ lexicographically dominates $B(F')$. In SUB-CENTRIC, we seek a disjoint collection of m sets $S_1, \dots, S_m \subseteq \mathcal{U}$ such that each set can be covered by one antenna and the resultant allocation vector is as fair as any other allocation vector in the lexicographic sense. The set S_j represents the subscribers *assigned* to antenna j , meaning the AP would provide connectivity to the subscribers in S_j using the antenna j . The AP would also decide how to partition the bandwidth of antenna j among the subscribers in S_j .

C. Provider-centric Provisioning

A fundamentally different provisioning problem arises in networks where the subscribers are independent entities, each interested in the quality of its own network connectivity. The subscribers are, therefore, willing to pay the provider for network bandwidth. We assume that the revenue obtained from a subscriber i is proportional to d_i if its requirement is satisfied and 0 otherwise. The objective of the provisioning, then, is to determine the bandwidth allocation that maximizes the revenue that the provider can generate from the subscribers. We assume that in such a provider-centric provisioning problem (denoted by PROV-CENTRIC), each subscriber has a bandwidth requirement $d_i > 0$. For a set of subscribers S , let $d(S) = \sum_{i \in S} d_i$. We define a set of subscribers S to be *valid* iff

- $d(S) \leq 1$;
- S can be covered with a single antenna.

A solution to PROV-CENTRIC consists of valid sets of subscribers $S_1, \dots, S_m$. The objective of maximizing revenue in this case is same as maximizing the total demand satisfied. That is we want valid sets $S_1, \dots, S_m$ maximizing $\sum_{i=1}^m d(S_i)$.

The problem of PROV-CENTRIC turns out to be NP-hard. The NP-hardness follows from a simple reduction from a variant of classical “bin packing” problem (see Garey and Johnson [11]). An instance of bin packing consists of a finite set $\mathcal{U}$ of items, a rational weight $w_i \in (0, 1]$ for each item u in $\mathcal{U}$, and a positive integer m . The question is whether there exists a partition of the items of $\mathcal{U}$ into disjoint subsets $S_1, \dots, S_m$, such that sums of sizes of items in each S_i is no more than 1.

We consider a variant of the classical bin packing problem (called MAX-BIN-PACKING [9]). In MAX-BIN-PACKING, the number of bins is fixed and the objective is to maximize the sum of weight of items packed. This problem is also NP-hard [9]. In order to reduce MAX-BIN-PACKING to PROV-CENTRIC, we start with an instance of MAX-BIN-PACKING and construct an instance of PROV-CENTRIC where all the subscribers are along one straight line drawn from the access point (actually it is enough if every pair of subscribers are within ρ of each other). See Fig. 4. The subscribers correspond to the items with d_i corresponding to w_i . The antennas correspond to the bins. Since all the subscribers are on a straight line, we simply align all the antennas along that line. The hardness for PROV-CENTRIC follows as the solutions to the instances of PROV-CENTRIC and MAX-BIN-PACKING are exactly the same.

III. SUBSCRIBER-CENTRIC BANDWIDTH PROVISIONING

In this section, we describe a dynamic programming formulation that provides a lexicographic fairest allocation of bandwidth vectors. The general idea is to show the existence of a solution in which the allocations are restricted to a very special form. In turn, the best solution in that special form can be found using dynamic programming. We start with any optimal solution and show that it can be transformed into another solution which has this special form. Note that this transformation is an existential proof (non-algorithmic) and the algorithm itself is described in the next subsection.

In this entire discussion we assume that all vectors are sorted in non-decreasing order. For any allocation F , let $B(F) = (b_1, \dots, b_n)$ be the bandwidth vector over n subscribers where b_i denotes the bandwidth that the subscriber i gets in the allocation. Let $B_j(F) \subseteq B(F)$ be the bandwidth vector for subscribers associated with antenna j . Let $\mathcal{F}$ denote any fairest allocation (optimum solution) with ordered sets $O_1, \dots, O_m$, so that $(\mathcal{U}, E')$ has a simple path with the subsequence $\text{first_item}(O_1), \text{first_item}(O_2), \dots, \text{first_item}(O_m)$ and $\text{NbrHood}(O_1)$ is not a subset of $\text{NbrHood}(O_i)$ for any $2 \leq i \leq m$. The set O_j denotes the set of all the subscribers allocated to antenna j ($1 \leq j \leq m$). We now establish some properties of the sets in the optimum solution. The following lemma notes the simple fact that in allocation $\mathcal{F}$ if $|O_j| = c$, then $B_j(\mathcal{F}) = (\frac{1}{c}, \dots, \frac{1}{c})$.

Lemma 3.1: Let $\mathcal{F}$ be any optimum solution for SUB-CENTRIC with sets $O_1, \dots, O_m$. For any set O_j with $|O_j| = c$, $B_j(\mathcal{F}) = (\frac{1}{c}, \dots, \frac{1}{c})$.

Proof. If $c = 1$, the proof is trivial. Else, if there exists even one subscriber $p \in O_j$ having higher than $1/c$ allocation in $B_j(\mathcal{F})$, then there exists at least one other subscriber $q \in O_j$ having an allocation of less than $1/c$. Assuming all other subscribers get the same allocation, assigning both p and q a bandwidth of $1/c$ leads to a more dominant vector than $B(\mathcal{F})$, leading to a contradiction that $B(\mathcal{F})$ is a fairest allocation. $\square$

As a consequence of the above Lemma, we have the following: when all the antennas are omni-directional, the fairest allocation vector allocates a bandwidth of $\frac{m}{n}$ to each subscriber.

Fig. 5. The figure to the left corresponds to the sets O_1 and O_2 . The figure to the right corresponds to the sets S_1 and S_2 obtained by running algorithm REPLACE.

```

x ← 0
y ← first_item(O1)
D ← |Debt(O1)|
Do forever
  If D ≥ 1
    D ← D - 1
    Sx ← Sx ∪ {y}
    y ← successor(y)
  Else
    If x = m
      terminate the loop
    x ← x + 1
    Sx ← ∅
    D ← |Ox|
  
```

Fig. 6. Algorithm REPLACE.

A set O of subscribers is said to be a *good* set if all the subscribers in it are contiguous and $\text{NbrHood}(O)$ is defined. We reassign subscribers in $O_1, \dots, O_m$ to create a new solution $\mathcal{F}'$ with sets $S_1, \dots, S_m$ such that $|O_j| = |S_j|$ and S_j is good. See Fig. 5. Note that since the cardinality of every set remains the same $B(\mathcal{F}') = B(\mathcal{F})$. For doing the reassignment we define

$$\text{debt}(O_j, O_k) = \begin{cases} \emptyset & \text{if } \text{first_item}(O_j) \notin \text{NbrHood}(O_k) \\ \text{NbrHood}(O_j) \cap O_k & \text{otherwise} \end{cases}$$

$$\text{Debt}(O_j) = \bigcup_{O_k \in \mathcal{F}} \text{debt}(O_j, O_k)$$

The process of replacing each O_j with S_j is described in algorithm REPLACE (Fig. 6). The following lemma proves the existence of solution $\mathcal{F}'$.

Lemma 3.2: Any solution for $\mathcal{F}$ for SUB-CENTRIC with sets $O_1, \dots, O_m$ can be replaced with a solution $\mathcal{F}'$ with sets $S_1, \dots, S_m$ such that for all $1 \leq j \leq m$ (i) $|S_j| = |O_j|$, (ii) subscribers in S_j are contiguous, and (iii) $\text{NbrHood}(S_j)$ is defined.

Sketch of Proof. Properties (i) and (ii) follow straightfor-

wardly from the algorithm REPLACE. We will divide the proof of property (iii) into three parts. The idea is to show that there exists at least one l such that $\text{NbrHood}(S_j)$ is a subset of $\text{NbrHood}(O_l)$ and $1 \leq l \leq j$.

Part 1: We first show that:

$$\bigcup_{i=1}^j S_i \subseteq \bigcup_{i=1}^j \text{NbrHood}(O_i)$$

To prove this we note that O_j and $\text{Debt}(O_j)$ are disjoint subsets of $\text{NbrHood}(O_j)$. Therefore, $|O_j| + |\text{Debt}(O_j)| \leq |\text{NbrHood}(O_j)|$. Now sets $S_1, \dots, S_j$ cover contiguous set of subscribers in $\text{path}(\text{first_item}(S_1), \text{last_item}(S_j))$. We finish the first part as follows:

- If $\text{first_item}(O_1) \in \text{NbrHood}(O_j)$, then the relation holds trivially as $\bigcup_{i=1}^j \text{NbrHood}(O_i) = \mathcal{U}$.
- Otherwise, $\bigcup_{i=1}^j \text{NbrHood}(O_i)$ contains at least $|\bigcup_{i=1}^j O_i \cup \text{Debt}(O_i)|$ subscribers. By construction we again guarantee that the relation holds.

As a consequence there also exists a set $s \in [1, j]$ with either $\text{last_item}(O_s) = \text{last_item}(S_j)$ or $\text{last_item}(O_s) \notin \bigcup_{i=1}^j S_i$.

Part 2: We show that:

$$|\text{path}(\text{first_item}(O_j), \text{first_item}(S_j))| \geq |\text{Debt}(O_j)|.$$

For $j = 1$, it immediately follows from the construction. For $j > 1$, we note that $|\text{path}(\text{first_item}(S_1), \text{first_item}(S_j))|$ equals to $|\bigcup_{i=1}^{j-1} O_i \cup \text{Debt}(O_j)| + 1$, but on the other hand $|\text{path}(\text{first_item}(S_1), \text{first_item}(O_j))|$ is at most $|\bigcup_{i=1}^{j-1} O_i \setminus \text{Debt}(O_j)| + 1$. Therefore, it is the case that,

$$|\text{path}(\text{first_item}(O_j), \text{first_item}(S_j))| \geq |\text{Debt}(O_j)|.$$

Part 3: From Part 2, we can infer that either $\text{NbrHood}(S_j) \subseteq \text{NbrHood}(O_j)$ or $\text{last_item}(O_j) \in \bigcup_{i=1}^j S_i$. If it was the first case, then definitely $\text{NbrHood}(S_j)$ is defined. For the second case, we condition on the location of $\text{last_item}(O_j)$. If $\text{last_item}(O_j) = \text{last_item}(S_j)$, then definitely $\text{NbrHood}(S_j) \subseteq \text{NbrHood}(O_j)$. If $\text{last_item}(O_j) \in S_k$ with $k \in [1, j]$, then from Part 1, $\text{NbrHood}(S_j) \subseteq \text{NbrHood}(O_s)$. Therefore in both these sub-cases, we get that $\text{NbrHood}(S_j)$ is defined. $\square$

A. Dynamic Programming Formulation

From Lemma 3.2, we know that we can convert any optimum solution $\mathcal{F}$ into a new optimum solution $\mathcal{F}'$ in which all sets in are good. We now show that, the solution $\mathcal{F}'$, i.e., sets $S_1, \dots, S_m$ can be found using dynamic programming.

For every $a, b \in \mathcal{U}$ the dynamic programming considers allocating exactly k ($\leq m$) antennas to subscribers in $\text{path}(a, b)$, a sub-problem denoted by $M[a, b, k]$. Remember that $\text{path}(a, b)$ defines a set of subscribers in the path between a and b in $(\mathcal{U}, E')$.

Recursive Formula in Dynamic Programming: Let Impossible be a special symbol, indicating that no feasible


```

j ← 1
i ← t
first_item(S1) ← i
While i ≤ n
  If di + d(Sj) ≤ 1 and
    Sj ∪ {i} ⊆ Nbrright(first_item(Sj))
    Sj ← Sj ∪ {i}
    i ← successor(i)
  Else
    D ← D ∪ d(Sj)
    j ← j + 1
    first_item(Sj) ← i
D ← D ∪ d(Sj)
Output the sum of m largest entries in D

```

Fig. 7. Algorithm GREEDY.

solution exists for the sub-problem. Now $M[a, b, k]$ for any $a, b \in \mathcal{U}$ is assigned to

$$\text{LMAX}_c\{M[a, c, k-1] \cup (\underbrace{\frac{1}{|\text{path}(d, b)|}, \dots, \frac{1}{|\text{path}(d, b)|}}_{\text{total of } |\text{path}(d, b)| \text{ times}})\}$$

if there exists $c \in \text{path}(a, b)$, $(c, d) \in E'$, and $d \in \text{path}(a, b)$ such that $M[a, c, k-1]$ is not assigned Impossible, and $\text{NbrHood}(\text{path}(d, b))$ is defined. The operator LMAX_c defines lexicographically the biggest vector over the range of c (we assume all vectors are sorted in non-decreasing order). If these conditions are not met, $M[a, b, k]$ is assigned Impossible, i.e., $\text{path}(a, b)$ cannot be covered with k antennas. For the base case we have $M[a, a, p] = (1)$, where $1 \leq p \leq m$. Also $M[a, b, 1]$ equals

$$(\underbrace{\frac{1}{|\text{path}(a, b)|}, \dots, \frac{1}{|\text{path}(a, b)|}}_{\text{total of } |\text{path}(a, b)| \text{ times}})$$

if $\text{NbrHood}(\text{path}(a, b))$ is defined, otherwise it is set to Impossible.

The correctness of our dynamic programming is straightforward (as for every $\text{path}(a, b)$ and every possible number of antennas that could be assigned to $\text{path}(a, b)$ we try all possible contiguous allocations, therefore finding $\mathcal{F}'$). The number of entries in the table M is of $O(n^2m)$ and each entry can be computed with $O(n^2)$ time, resulting in a total run-time of $O(n^4m)$. The space requirement of our algorithm is $O(n^3m)$. We conclude this section with the following.

Theorem 3.3: There exists an $O(n^4m)$ time algorithm for the problem of SUB-CENTRIC.

IV. PROVIDER-CENTRIC BANDWIDTH PROVISIONING

In this section, we describe a simple algorithm achieving ≈ 2 -approximation for the problem of PROV-CENTRIC. More precisely, we show that the revenue our algorithm generates is (almost) at least half of the revenue achievable by any optimal algorithm. For the simple case of omni-directional antennas, the problem of PROV-CENTRIC is exactly the same

Fig. 8. The figure shows operation of GREEDY. Starting from subscriber t , GREEDY uses two antennas to pack the five subscribers. We assume $\rho = 90$ and each subscriber has a demand of 0.2.

as the standard bin-packing problem and therefore admits an asymptotic polynomial approximation scheme [10, 15].

The algorithm GREEDY specified in Fig. 7 packs all the subscribers in $\mathcal{U}$ into valid sets (defined in Section II-C). To enable the packing, we define a linear ordering L on the set of subscribers $\mathcal{U}$. The linear ordering starts from a subscriber t satisfying

$$d_t = \min\{d(\text{Nbr}_{\text{left}}(i)) : i \in \mathcal{U}\}.$$

Two subscribers i, j are consecutive in this ordering if $(i, j) \in E'$, where E' is the edges of the graph constructed in Section II-A. We start greedy packing from the lowest unpacked subscriber (in the ordering) and close an antenna for packing purposes if either the antenna cannot accommodate a subscriber or if it reaches its maximum span (see Fig. 8).

The running time of the algorithm GREEDY is linear in the number of subscribers. We note that the sets produced by GREEDY are valid (follows from how sets are chosen for packing). Consider any optimal solution to the PROV-CENTRIC problem. Let $O_1, \dots, O_m$ be the sets of the optimal solution ordered by last subscriber in the linear ordering L , i.e., $\text{last_item}(O_j)$ appears before $\text{last_item}(O_{j+1})$ in the linear ordering L . Let $\text{OPT} = \sum_{j=1}^m d(O_j)$. For any $j \in [1, m]$, let $\bar{O}_j$ be the set derived from O_j by removing from it all the subscribers (if any) in $\text{Nbr}_{\text{left}}(t)$. Let $\bar{O} = \{\bar{O}_1, \dots, \bar{O}_m\}$ be the corresponding ordered set.

We define a graph G with its vertices as the sets in $\bar{O}$ and edges between $\bar{O}_i$ and $\bar{O}_j$ if the antennas representing $\bar{O}_i$ and $\bar{O}_j$ overlap, i.e., $\text{NbrHood}(\bar{O}_i) \cap \text{NbrHood}(\bar{O}_j) \neq \emptyset$. We next determine the approximation ratio of GREEDY.

Lemma 4.1: Let $S_1, \dots, S_g$ be the sets produced by the algorithm GREEDY. Let SOL be the output produced by the algorithm GREEDY. If $g \leq 2m$, then $\text{OPT} \leq 2 \cdot \text{SOL}$.

Proof. Note that $\sum_{j=1}^g d(S_j) = d(\mathcal{U})$ and $\text{OPT} \leq d(\mathcal{U})$. Now, if $g \leq 2m$, then the sum of the m largest entries in $\mathcal{D}$ is at least $d(\mathcal{U})/2$. Hence, the result follows. $\square$

For the sake of analysis, we partition the sets $\mathcal{S} = \{S_1, \dots, S_g\}$ constructed by GREEDY into three groups (set of sets).

Fig. 9. Algorithm REASSIGN

- Group S_a : sets $S_a \in \mathcal{S}$ satisfying $d(S_a) < 1/2$ and $d(\text{Nbr}_{\text{right}}(\text{first_item}(S_a))) > 1/2$.
- Group S_b : sets $S_b \in \mathcal{S}$ satisfying $d(S_b) < 1/2$ and $d(\text{Nbr}_{\text{right}}(\text{first_item}(S_b))) < 1/2$.
- Group S_c : sets $S_c \in \mathcal{S}$ satisfying $d(S_c) \geq 1/2$.

All the three groups are disjoint. The group S_a consists of antennas satisfying a total subscriber demand less than $1/2$ and are closed prematurely before spanning a sector of ρ . The group S_b consists of antennas which satisfy a total demand less than $1/2$ and are closed after satisfying all subscribers in a sector of ρ . Every set $S_b \in S_b$ spans a sector of ρ , therefore in every optimum solution no subscriber before $\text{first_item}(S_b)$ can be assigned to the same antenna as a subscriber following $\text{last_item}(S_b)$. The group S_c consists of antennas which satisfy a total demand at least $1/2$. The greedy packing strategy results in the following lemma.

Lemma 4.2: Let $S_a \in S_a$ with $\text{first_item}(S_a) \in U - \text{Nbr}_{\text{left}}(t)$. Then set $S_{a+1} \in S_c$.

Proof. Every such set S_a in S_a has demand less than $1/2$ and is closed because the next subscriber p considered for packing results in $d(S_a) + d_p > 1$. Therefore, $d_p > 1/2$. This implies that S_{a+1} which accommodates p has $d(S_{a+1}) > 1/2$ and S_{a+1} belongs to S_c . $\square$

We now transform the sets in $\mathcal{S}$ by eliminating sets in the group S_a . The algorithm REASSIGN is described in Fig. 9. Any set $\tilde{S}$ produced by the if part of the loop would satisfy $d(\tilde{S}) > 1$. The sets in $\tilde{\mathcal{S}}$ are also partitioned into groups (again set of sets) as follows:

- Group $\tilde{S}_a$: sets $\tilde{S}_a \in \tilde{\mathcal{S}}$ satisfying $d(\tilde{S}_a) > 1$.
- Group $\tilde{S}_b$: sets $\tilde{S}_b \in \tilde{\mathcal{S}}$ satisfying $d(\tilde{S}_b) < 1/2$ and $d(\text{Nbr}_{\text{right}}(\text{first_item}(\tilde{S}_b))) < 1/2$.
- Group $\tilde{S}_c$: sets $\tilde{S}_c \in \tilde{\mathcal{S}}$ satisfying $d(\tilde{S}_c) \geq 1/2$ and $d(\tilde{S}_c) \leq 1$.

Each set in $\tilde{S}_a$ is formed due to the algorithm REASSIGN and corresponds to two sets in $\mathcal{S}$. We now show that the total demand satisfied by GREEDY is very close to the demand satisfied by any optimal solution.

Lemma 4.3: Let SOL be the output produced by the algorithm GREEDY. Then, $\text{OPT} \leq 2 \cdot \text{SOL} + 1$.

Sketch of Proof. If $g \leq 2m$, Lemma 4.1 implies that $\text{OPT} \leq 2 \cdot \text{SOL} + 1$. So we next concentrate on the case when $g > 2m$.

Fig. 10. All the subscribers in the sector defined by $\text{Nbr}_{\text{left}}(t)$ are ignored in Case 2 of Lemma 4.3.

Note that $\text{OPT} \leq m$. We divide our analysis into two sub-cases:

Case 1: We first consider the case when the sum of demands satisfied by every consecutive pair of sets in $\mathcal{S}$ is at least 1. Since $g > 2m$, the top $2m$ entries in $\mathcal{D}$ sum to at least m . This implies that the top m entries in $\mathcal{D}$ sum to at least $m/2$. Therefore, we get $\text{OPT} \leq 2 \cdot \text{SOL}$.

Case 2: We next consider the case when there exists a pair of consecutive sets (say S_q and S_{q+1}) such that the sum of the demands satisfied by these sets is less than 1, i.e., $d(S_q) + d(S_{q+1}) < 1$, then $d(\text{Nbr}_{\text{left}}(t)) < 1$. This condition may arise when one of the antennas q and $q+1$ spans a sector of ρ and satisfies the demands of all the subscribers in that sector. Assume w.l.o.g. that it is the $(q+1)^{\text{st}}$ antenna. If the $(q+2)^{\text{nd}}$ antenna starts at some subscriber t' , then $d(\text{Nbr}_{\text{left}}(t')) < 1$, implying $d(\text{Nbr}_{\text{left}}(t)) < 1$.

We remove all the subscribers in the sector $\text{Nbr}_{\text{left}}(t)$ in the analysis (see Fig. 10). Since $d(\text{Nbr}_{\text{left}}(t)) < 1$, $\sum_{k=1}^m d(O_k) \leq \sum_{k=1}^m d(\tilde{O}_k) + 1$. We now construct $\hat{\mathcal{S}} \subseteq \tilde{\mathcal{S}}$ by assigning at most two sets from $\tilde{\mathcal{S}}$ to each set in $\tilde{\mathcal{O}}$. We scan the sets of $\tilde{\mathcal{O}}$ in order, doing the following step when we consider the k^{th} element (which is $\tilde{O}_k$) from $\tilde{\mathcal{O}}$.

If $\sum_{S \in \tilde{\mathcal{S}}} d(S) \geq \sum_{j=1}^k d(\tilde{O}_j)$, then we do not change the set $\hat{\mathcal{S}}$ (as it already satisfies more demands than what is satisfied by the first k sets in $\tilde{\mathcal{O}}$). Otherwise, consider the set of subscribers $\text{path}(t, \text{last_item}(O_k))$. For a set $\tilde{O}_j$, let $\ell(j)$ and $f(j)$ be the index number of the set in $\tilde{\mathcal{S}}$ containing $\text{last_item}(\tilde{O}_j)$ and $\text{first_item}(\tilde{O}_j)$ respectively, i.e., $\text{last_item}(\tilde{O}_j) \in \tilde{S}_{\ell(j)}$ and $\text{first_item}(\tilde{O}_j) \in \tilde{S}_{f(j)}$.

We scan $\tilde{S}_{\ell(k)}, \tilde{S}_{\ell(k)-1}, \dots, \tilde{S}_1$ in order until we encounter the first set $\tilde{S}_u \notin \hat{\mathcal{S}}$. If $d(\tilde{S}_u) \geq 1$, then assign $\tilde{S}_u$ to T_1 and $\emptyset$ to T_2 . Otherwise, we continue scanning until we reach the next set $\tilde{S}_v \notin \hat{\mathcal{S}}$. Again, if $d(\tilde{S}_v) \geq 1$, set T_1 to $\emptyset$ and T_2 to $\tilde{S}_v$. Otherwise, set T_1 to $\tilde{S}_u$ and T_2 to $\tilde{S}_v$ (in this case, both $d(\tilde{S}_u)$ and $d(\tilde{S}_v)$ are less than 1). Add T_1 and T_2 to $\hat{\mathcal{S}}$. Sets T_1 and T_2 are called the *representatives* of $\tilde{O}_k$. We only consider the case where both $\tilde{S}_u$ and $\tilde{S}_v$ exist. The cases when only one of $\tilde{S}_u$ or $\tilde{S}_v$ exists are straight forward and omitted.

The proof can be completed inductively by using a charging argument to show: $\sum_{S \in \hat{\mathcal{S}}} d(S) \geq \sum_{j=1}^k d(\tilde{O}_j)$. For the base

$$\begin{aligned}
&\ell(1) = f(1) + 1: \\
&d(\bar{O}_1) \leq d(S_{f(1)}) + d(S_{\ell(1)}) = d(T_1) + d(T_2) \\
&\ell(1) > f(1) + 1: \\
&d(\bar{O}_1) \leq 1 \leq d(S_{\ell(1)}) + d(S_{\ell(1)-1}) = d(T_1) + d(T_2)
\end{aligned}$$

Fig. 11. Inequalities for the basis of induction

case ($k = 1$), $d(T_1) = d(\bar{S}_{\ell(1)})$ and $d(T_2) = d(\bar{S}_{\ell(1)-1})$. If either of $d(\bar{S}_{\ell(1)})$ or $d(\bar{S}_{\ell(1)-1})$ is greater than 1, then the base case is true. Otherwise, we use the case distinctions in Fig. 11 to complete the proof.

By the inductive hypothesis, we assume that $\sum_{S \in \mathcal{S}} d(S) - d(T_1) - d(T_2) \geq \sum_{j=1}^{k-1} d(\bar{O}_j)$. For the cases when $d(T_1) \geq 1$, or $d(T_2) \geq 1$, or $d(T_1) + d(T_2) \geq 1$, the proof is straightforward. Therefore, the only remaining case is when at least one of $d(T_1)$ and $d(T_2)$ is less than $1/2$. Assume w.l.o.g. that $d(T_1) < 1/2$. Therefore, T_1 belongs to group $\bar{S}_b$.

From the construction of $\hat{\mathcal{S}}$, we know that no set in $\bar{O}$ has its `last_item` in $\bar{S}_u (= T_1)$. Furthermore, since $T_1 \in \bar{S}_b$, we infer that T_1 covers a sector of ρ . Suppose that the sets $\bar{O}_1, \bar{O}_2, \dots, \bar{O}_{\ell(k)}$ are divided into two parts. The first part $\bar{O}_1, \dots, \bar{O}_l$ contains the sets whose `last_item` are contained in sets before $\bar{S}_u$ (in $\bar{\mathcal{S}}$) and the second part consists of the remaining sets.

From our construction, whenever we consider a set (say $\mathcal{M}$) in the first part, we don't assign sets later than $\bar{S}_u$ as $\mathcal{M}$'s representatives. By inductive hypothesis, we get

$$\sum_{j=1}^l d(\bar{O}_j) \leq \sum_{w < u, \bar{S}_w \in \bar{\mathcal{S}}} d(\bar{S}_w).$$

If we consider the connected component of graph G containing $\bar{O}_k$, it does not contain any set from the first part. For the second part, we have:

$$\sum_{j=l+1}^k d(\bar{O}_j) \leq d\left(\bigcup_{j=l+1}^k \text{NbrHood}(\bar{O}_j)\right) \leq \sum_{w \geq u, \bar{S}_w \in \bar{\mathcal{S}}} d(\bar{S}_w).$$

This completes the induction. The fact that we use at most two sets from $\mathcal{S}$ as representatives for each set in $\bar{O}$ follows because either T_1 or T_2 each correspond to one set in $\mathcal{S}$ or one of them is $\emptyset$ and the other corresponds to two sets in $\mathcal{S}$.

To complete the proof, we note that for m sets in $\bar{O}$ we assign $2m$ sets from $\mathcal{S}$ to $\hat{\mathcal{S}}$ such that

$$\sum_{j=1}^m d(O_j) \leq \sum_{j=1}^m d(\bar{O}_j) + 1 \leq \sum_{S \in \mathcal{S}} d(S) + 1.$$

Since GREEDY picks m heaviest sets in $\mathcal{S}$, it gives the desired bound. $\square$

We conclude this section with the following theorem which follows from Lemma 4.3. The constant of proportionality arising in the theorem is what relates the revenue to subscriber requirement in the statement of PROV-CENTRIC.

Theorem 4.4: There exists a linear time algorithm for the problem of PROV-CENTRIC which produces a revenue which is at least twice the optimal revenue minus some constant of proportionality.

V. PROVISIONING IN A NETWORK WITH MULTIPLE ACCESS POINTS

In this section, we consider the generalizations of SUB-CENTRIC and PROV-CENTRIC to the case of multiple access points. Let K (considered part of the input) be the number of access points that can be established. Again there are n subscribers on the plane, who are to be serviced by directional antennas mounted on any one of the access points. Let m denote the number of antennas that can be mounted on each access point. Even though a subscriber could be in range of multiple antennas mounted on possibly different access points, he is assigned to only one antenna.

We reformulate the problems of SUB-CENTRIC and PROV-CENTRIC for this setting. For both problems, we want to know the location of the access points and for each access point i a set of subscribers $S_{(i,1)}, \dots, S_{(i,m)}$ such that each set can be covered by one antenna mounted on access point i . Also the sets need to satisfy the (disjointness) property that for $1 \leq p, q \leq K$ and $1 \leq r, s \leq m$, $S_{(p,r)} \cap S_{(q,s)} = \emptyset$ with $p = q$ and $r = s$ not being simultaneously true. For the problem of SUB-CENTRIC $_K$ the objective is to obtain a collection of sets $S_{(p,r)}$ such that the resultant allocation is the fairest. For the problem of PROV-CENTRIC $_K$ the objective is to obtain a collection of sets $S_{(p,r)}$ such that each set is valid and the resulting revenue is maximized.

The problem of PROV-CENTRIC $_K$ is definitely NP-hard here (as it was already the case for a single access point). But now so is the problem of SUB-CENTRIC $_K$. Actually, the problem is hard even for the simpler case of omni-directional antennas ($\rho = 360$). To show this, we argue that SUB-CENTRIC $_K$ on omni-directional antennas is at least as hard as the problem computing the maximum number of points in the plane that can be covered with K unit disks [13]. Consider any optimal solution for SUB-CENTRIC $_K$. It will have the property that number of subscribers within R of at least one access point is maximized (follows because every fair allocation vector will have as few preceding zeros as possible). Therefore disks (of radius R) placed on access point locations cover the maximum number of subscribers. Finally, covering maximum number of points (= subscribers) in the plane with K disks is a well known NP-hard problem [5, 13].

VI. CONCLUSIONS AND FUTURE RESEARCH DIRECTIONS

There has been an increasing interest in using directional antennas for wireless communications. Directional antennas offer several advantages such as increased spatial reuse and longer range as compared to omni-directional antennas. Consequently, directional antennas are expected to be more widely used in next-generation wireless networks. However, directional antennas also introduce some difficulties. In particular, each node has to position its antenna appropriately before communicating with other nodes. In this paper, we considered the use of directional antennas in infrastructure-based wireless networks in which each access point is equipped with multiple directional antennas for communication with its subscribers. The problem here is to determine how each access point aligns

its directional antennas so as to cover its subscribers and yet providing performance guarantees.

We formulated two bandwidth provisioning problems relevant in these networks: (i) subscriber-centric, where the objective was to achieve a fair allocation of bandwidth among the subscribers and (ii) provider-centric, where the objective was to maximize the revenue that the network provider generates from the subscribers. For a subscriber-centric network, under the notion of max-min fairness, we presented a dynamic programming algorithm that guaranteed fair allocation. For a provider-centric network, we presented a linear-time greedy algorithm that guaranteed close to half of the optimum revenue.

There are a number of interesting future directions that we are currently pursuing. We are investigating heuristics to the bandwidth provisioning problems for the case with multiple access points. The problem setting in this paper assumed that subscribers are static or move very infrequently. This allows the algorithms to be re-executed every time a subscriber moves. However, in highly mobile environments such an approach is clearly impractical. Efficient antenna positioning algorithms for these settings is an extremely interesting future direction currently under investigation.

VII. ACKNOWLEDGMENTS

We would like to thank Piotr Berman for many fruitful discussions on this work. Early work on this problem, particularly developing the model, was done jointly with Piotr.

REFERENCES

- [1] M. Balazinska and P. Castro. Characterizing mobility and network usage in a corporate wireless local-area network. In *Proceedings of ACM MobiSys '03*, pages 303–316, 2003.
- [2] L. Bao and J. Garcia-Luna-Aceves. Transmission scheduling in ad hoc networks with directional antennas. In *Proceedings of ACM MOBICOM '02*, pages 48–58, 2002.
- [3] P. Berman, J. Jeong, S. P. Kasiviswanathan, and B. Ugaonkar. Packing to angles and sectors. Technical report TR05-033, Electronic Colloquium on Computational Complexity, 2006.
- [4] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, 1987.
- [5] M. Charikar, S. Khuller, D. M. Mount, and G. Narasimhan. Algorithms for facility location problems with outliers. In *Proceedings of the ACM SODA '01*, pages 642–651, 2001.
- [6] R. Choudhury and N. Vaidya. Impact of directional antennas on ad hoc routing. In *Proceedings of the Personal and Wireless Communications '03*, pages 590–600, 2003.
- [7] R. Choudhury, X. Yang, R. Ramanathan, and N. Vaidya. Using directional antennas for medium access control in ad hoc networks. In *Proceedings of the ACM MOBICOM '02*, pages 59–70, 2002.
- [8] K. L. Clarkson and K. R. Varadarajan. Improved approximation algorithms for geometric set cover. In *Proceedings of the ACM SoCG '05*, pages 135–141, 2005.
- [9] E. G. Coffman, J. Y.-T. Leung, and D. W. Ting. Bin packing: Maximizing the number of pieces packed. *Acta Informatica*, 9:263–271, 1978.
- [10] W. F. de la Vega and G. S. Lueker. Bin packing can be solved within $1+\epsilon$ in linear time. *Combinatorica*, 1(4):349–355, 1981.
- [11] M. R. Garey and D. S. Johnson. *Computers and Intractability*. W.H. Freeman and Co., 1979.
- [12] P. Goyal, H. M. Vin, and H. Cheng. Start-time Fair Queuing: A Scheduling Algorithm for Integrated Services Packet Switching Networks. In *Proceedings of ACM SIGCOMM '96*, pages 157–168, 1996.
- [13] D. S. Hochbaum and W. Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *Journal of the ACM*, 32(1):130–136, Jan. 1985.
- [14] X. L. Huang and B. Bensaou. On max-min fairness and scheduling in wireless ad-hoc networks: analytical framework and implementation. In *Proceedings of the ACM MobiHoc '01*, pages 221–231. ACM, 2001.
- [15] N. Karmarkar and R. M. Karp. An efficient approximation scheme for the one-dimensional bin-packing problem. In *Proceedings of the IEEE FOCS '82*, 1982.
- [16] J. M. Kleinberg, Y. Rabani, and É. Tardos. Fairness in routing and load balancing. *J. Comput. Syst. Sci.*, 63(1):2–20, 2001.
- [17] Y. Kuga, J. Cha, J. A. Ritcey, and J. Kajiya. Mechanically steerable antennas using dielectric phase shifters. In *Proceedings of the IEEE-APS and URSI*, pages 161–164, 2004.
- [18] A. Nasipuri, S. Ye, J. You, and R. Hiromoto. A MAC protocol for mobile ad hoc networks using directional antennas. In *Proceedings of the IEEE WCNC '00*, pages 1214–1219, 2000.
- [19] T. Ojala, T. Hakanen, T. Mäkinen, and V. Rivinoja. Usage analysis of a large public wireless lan. In *Proceedings of IEEE WirelessComm '05*, pages 661–667, 2005.
- [20] A. K. Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: the multiple node case. *IEEE/ACM Transactions on Networking*, 2(2):137–150, 1994.
- [21] C. Peraki and S. Servetto. On the maximum stable throughput problem in random networks with directional antennas. In *Proceedings of the ACM MobiHoc '03*, pages 76–87, 2003.
- [22] R. Ramanathan. On the performance of beamforming antennas. In *Proceedings of the ACM MobiHoc '01*, pages 95–105, 2001.
- [23] S. Vasudevan, J. Kurose, and D. Towsley. On neighbor discovery in wireless networks with directional antennas. In *Proceedings of the IEEE INFOCOM '05*, pages 590–600, 2005.
- [24] S. Yi, Y. Pei, and S. Kalyanaraman. On the capacity improvement of ad hoc wireless networks using directional antennas. In *Proceedings of the ACM MobiHoc '03*, pages 108–116, 2003.
- [25] J. Zhu and S. Roy. 802.11 mesh networks with two-radio access points. In *Proceedings of the IEEE ICC '05*, pages 3609–3615, 2005.