Bandwidth Provisioning in Infrastructure-based Wireless Networks Employing Directional Antennas[1]

Shiva Prasad Kasiviswanathan[a], Bo Zhao[c], Sudarashan Vasudevan[b], Bhuvan Urgaonkar[c]

[a] *IBM T.J. Watson Research Center*
[b] *Dept of Computer Science, University of Massachusetts Amherst*
[c] *Department of Computer Science and Engineering, Pennsylvania State University*

Abstract

Motivated by the widespread proliferation of wireless networks employing directional antennas, we study the problem of provisioning bandwidth in such networks. Given a set of subscribers and one or more access points possessing directional antennas, we formalize the problem of orienting these antennas 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 generated by satisfying the bandwidth requirements of subscribers. For both the problems, we first design algorithms for a network with only one access point working under the assumption that the number of antennas does not exceed the number of non-interfering channels. Using the well-regarded lexicographic max-min fair allocation as the objective for a subscriber-centric network, we present a dynamic programming algorithm that achieves the fairest allocation. For a provider-centric network, the allocation problem turns out to be NP-hard. We present a greedy heuristic based algorithm that guarantees almost half of the optimum revenue. We later enhance both these algorithms to operate in more general networks with multiple access points and no restrictions on the relative numbers of antennas and channels. A simulation-based evaluation using OPNET demonstrates the efficacy of our approaches and provides us further in insights into these problems.

1. Introduction

There has been a growing interest in the use of directional antennas in communication networks [1, 2, 3, 4]. 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 re-use in comparison with omni-directional antennas. Several studies point out that directional antennas have the potential to improve the capacity

[1] A preliminary version of this paper appeared in 11th International Conference on Distributed Computing and Networking.

Email addresses: kasivisw@gmail.com (Shiva Kasiviswanathan), bzhao@cse.psu.edu (Bo Zhao), svasu@cs.umass.edu (Sudarashan Vasudevan), bhuvan@cse.psu.edu (Bhuvan Urgaonkar)

Preprint submitted to Elsevier *March 20, 2011*

of wireless networks [3, 5, 6]. Furthermore, there have been several recent efforts [3, 4] to build low-cost steerable directional antennas. It can therefore be imagined that directional antennas will be widely used in a variety of wireless networks. However, directional antennas introduce some unique challenges in the design and operation of wireless networks. In particular, unlike networks with omni-directional antennas, such networks require mechanisms for orienting the antenna.

Our focus in this paper is on bandwidth allocation[1] in infrastructure-based wireless networks such as cellular networks and IEEE 802.11 WLANs. In these networks, wireless nodes (called subscribers) affiliate themselves with *base stations* (also called *access points (APs)* in this paper), each of which typically has a high-speed connection to the rest of the network. We assume that each subscriber can be only assigned to one antenna (and correspondingly, a single channel). A directional antenna can only provide connectivity to a subscriber if it geometrically covers it. The bandwidth allocated to a subscriber depends on the number of other subscribers covered by a directional antenna. For example, if a directional antenna on an AP transmits on a channel with bandwidth b and covers n subscribers sharing the bandwidth equally, each subscriber receives a bandwidth of b/n. We are interested in algorithms for orienting the antennas and assigning them to the subscribers such that the resulting allocation of bandwidth to the subscribers has some desirable properties (such as those concerning fairness or efficiency of allocation; we describe two concrete examples below.)

There have been several proposals to handle the challenges introduced by the usage of directional antennas in the context of neighbor discovery [7], medium access control [8, 9], and routing [10] in multi-hop wireless networks. Whereas existing research has experimentally demonstrated the performance benefits offered by using multi-radio access points (APs) equipped with *omni-directional* antennas [11], bandwidth allocation with directional antennas has not been studied. As we will see in this paper, the problem becomes significantly more challenging with directional antennas.

We consider two fundamental bandwidth allocation problems in this paper. In both the settings, we are given: (i) the locations of a set of wireless subscribers, and (ii) the locations of the APs present in the network. Each AP is equipped with one or more directional antennas. Each antenna has an associated bandwidth that is to be shared among all the subscribers associated with that antenna.

In the first problem formulation, which we call the *subscriber-centric provisioning*, the goal of the network is to choose an assignment of subscribers to the antennas that ensures a max-min fair distribution [12] of bandwidth among the affiliated subscribers. Max-min fairness is a simple, well-recognized egalitarian objective that has gained wide acceptance for defining fairness in both wired and wireless networks [12, 13]. This formulation captures the bandwidth provisioning problem that arises in a network where subscribers are homogeneous and do not pay differently for the service received. As an example, consider an AP providing network connectivity to users in offices and university buildings. Given the locations of these wireless users 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.

[1]We use the term *bandwidth* as synonymous with *data rate*, unless otherwise specified.

The second formulation we consider is called the *provider-centric provisioning*. Unlike the subscriber-centric problem, each subscriber specifies a bandwidth requirement to the network provider. If a subscriber's requirement can be satisfied, it pays the provider a cost proportional to the requirement. A subscriber does not pay the provider if its requirement is not satisfied. The goal of the provider is to choose orientations of the antennas and assignment of subscribers to the antennas such that the network provider's revenue is maximized. 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.

Both these problems are related to geometric covering problems, where we have a set of points in a two-dimensional plane and we want to cover them with the minimum possible number of objects of some shape [14, 15]. The shapes of objects correspond to the patterns of directional antennas. Similar to [5, 16, 7], we model a directional antenna as a circular sector centered at the AP. In a recent paper [17], the authors considered the problem of minimizing the number of directional antennas for satisfying all subscriber demands. Although, the revenue maximization problem considered in this paper is related to the problem considered in [17], it requires fundamentally different algorithmic ideas and techniques.

We make the following key contributions in this paper.

① We start by considering the restricted version of the problems where we have a single AP and number of antennas on the AP does not exceed the number of non-interfering channels (we refer to this set of restrictions as *RestAntChan*). Under these restrictions we,
a) Develop a dynamic programming algorithm that achieves an optimum max-min fair allocation for the subscriber-centric model.
b) Show that revenue-maximization for the provider-centric model is NP-hard. We develop a greedy algorithm that guarantees close to half of the optimal revenue in the worst case and runs in time linear in number of subscribers.

② Building upon the algorithms above, we devise heuristics for bandwidth allocation in networks with multiple APs and without any restrictions on the relative numbers of antennas per AP.

③ We evaluate our algorithms by conducting an extensive simulation study using the OPNET [18] simulator, both in a single-AP and the more general multiple AP settings. Our simulations show that our algorithms achieve near-optimal allocations over a range of parameter settings.

1.1. Related Work

Choudhury *et al.* [10] considered two MAC protocols for ad hoc networks using directional antennas. Ramanathan *et al.* [16] studied the capacity of wireless networks employing directional antennas, focusing on exploiting the longer ranges as well as the reduced interference that beamforming antennas can provide. Raman *et al.* [19] explored how to make 802.11 performance effective with directional antennas. Different combination of antenna modes has been studied by Yi *et al.* [5]. Directional antennas also have certain drawbacks. For example, in wireless MAC the directional antennas introduce the problem of *deafness*, caused when a node B is unable to communicate with another node

A because A is beamformed in a direction away from B. New MAC protocols have been designed to solve this problem [20, 21].

Research in fair allocation of the bandwidth resource has been active for decades. Lazar *et al.* [22] investigated the virtual path capacity allocation problem in a noncooperative setting. Heyman *et al.* [23] considered the performance of a bottleneck link shared by a fixed number of homogeneous sources alternately emitting documents and remaining inactive during a random think-time. Biswas *et al.* [24] examined the dynamic bandwidth allocation problem for real-time variable bit rate traffic in the context of ATM cell multiplexing. Fred *et al.* [25] studied the statistics of the realized throughput of elastic document transfers, accounting for the way network bandwidth is shared dynamically between the randomly varying number of concurrent flows. Duan *et al.* [26] studied the bandwidth provisioning problem for service overlay networks. They considered both the static and dynamic bandwidth provisioning models and their study takes into account various factors such as QoS, traffic demand distributions, and bandwidth costs. Allalouf *et al.* [27] presented centralized and distributed algorithms for finding an optimal global per-commodity max-min fair rate vector in a network with multiple source-demand pairs.

Organization. The rest of the paper is organized as follows. In Section 2, we describe our network model and formulate the two bandwidth provisioning problems in the context of a network with one AP. In Sections 3 and 4, we design and analyze algorithms for the restricted version (as described above) of the subscriber-centric and the provider-centric problems. In Section 5, we develop heuristics that relax these restrictions. In Section 6, we present a simulation study to evaluate our algorithms. Finally, we conclude with some future research directions in Section 7.

2. Network Model and Problem Definitions

We begin this section by describing the wireless network model assumed in our research. Next, we formally define the two bandwidth provisioning problems that we study in this paper. For sake of simplicity we define the model and problems with respect to a single AP. The definitions naturally generalize to the case of multiple APs.

2.1. Background and Network Model

We consider a wireless network consisting of one AP with m directional antennas, each capable of being steered in a direction chosen by the AP. Each antenna has a bandwidth of b bits/second, a transmission range of R meters, and a span of ρ degrees. Each antenna has to be associated with one of the available channels. The network has n wireless subscribers whose locations are fixed and are known to the AP. All subscribers are assumed to be at most R meters away from the AP.

We will deal with points in radial coordinates and with circular orderings. The origin of our coordinate system lies at the AP. A point (θ, r) is equivalent to a Cartesian point $(r \times \sin\theta, r \times \cos\theta)$. Equality of angles is understood modulo 2π, i.e., $\alpha = \beta$ means that for a certain integer j we have $\alpha = \beta + 2j\pi$. We denote the set of n subscribers by $U = \{1, \ldots, n\}$ and the location of subscriber i by (θ_i, r_i). A directional antenna is characterized by four parameters. The first three parameters, namely range R, bandwidth

b, and 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 coordinate 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* these subscribers. Henceforth, we assume the bandwidth of each antenna to be 1 bandwidth unit (1 bandwidth unit = b bits/second); all variables denoting bandwidth are assumed to be scaled appropriately. For subscriber i, we define the set of its *right neighbors* $\mathsf{Nbr}_{\mathsf{right}}(i)$ and the set of its *left neighbors* $\mathsf{Nbr}_{\mathsf{left}}(i)$ as follows.

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

$\mathsf{Nbr}_{\mathsf{right}}(i)$ consists of all subscribers whose angular separation from i is at most ρ in the counter-clockwise direction and $\mathsf{Nbr}_{\mathsf{left}}(i)$ represents the same in the clockwise direction.

2.2. Subscriber-centric Provisioning

The first setting that we consider is a wireless network where the network provider is expected to provide max-min fair bandwidth allocation when provisioning its subscribers. We call the problem of achieving a max-min fair bandwidth allocation as the subscriber-centric provisioning.

We begin by formally defining the notion of max-min fairness. It is defined via a type of equilibrium. Given two k-tuples, $X = (x_1, \ldots, x_k)$ and $Y = (y_1, \ldots, y_k)$, each in non-decreasing order, we say that X lexicographically dominates Y if $X = Y$, or there exists some index p such that $x_p > y_p$ and $x_q = y_q$ for all $q < p$. If X lexicographically dominates Y, we also say that X is *as fair as* Y. A *bandwidth allocation vector* is a n-tuple denoting the allocation of bandwidth to the n subscribers. A solution F to a subscriber-centric instance consists of a disjoint collection of m sets $S_1, \ldots, S_m \subseteq U$ such that each set can be covered by one antenna, and an allocation vector $B(F) = (b_1, \ldots, b_n)$ such that, for $j \in \{1, \ldots, m\}$,

- $\sum_{i \in S_j} b_i \leq 1$;
- $\bigcup_j S_j = U$.

Here, 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) and b_i denotes the bandwidth that the subscriber i gets in the allocation. Finally, $\forall i, j$, we require that the antennas covering S_i and S_j be associated with different channels if the sectors corresponding to the antennas intersect. Max-min fairness enforces the following *max-min equilibrium* condition on the allocation vector: For any collection of sets $\{S_1, \ldots, S_m\}$ and allocation vector $B(F)$, the vector $B(F)$ must be the fairest allocation.

2.3. Provider-centric Provisioning

A fundamentally different provisioning problem arises in networks where the subscribers are interested only in the quality of their own network connectivity. The subscribers are, therefore, willing to pay the provider for the network bandwidth. We assume that the revenue obtained from a subscriber i is $\Lambda \cdot d_i$ if its requirement d_i is satisfied (Λ

is a network-wide constant), and 0 otherwise. Here, $d_i \in [0, 1]$ (due to normalization of the antenna bandwidth mentioned above). The objective of the provisioning is to determine a bandwidth allocation that maximizes the revenue that the provider can generate from the subscribers. We assume that in such a provider-centric provisioning problem, each subscriber has a bandwidth requirement $d_i > 0$. For a set of subscribers S, define *weight* $d(S)$ as $\sum_{i \in S} d_i$. A set of subscribers S is said to be *valid* iff

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

A solution to a provider-centric instance consists of valid disjoint sets of subscribers $S_1, \ldots, S_m$. Again, $\forall i, j$, we require that the antennas covering S_i and S_j be associated with different channels if the sectors corresponding to the antennas intersect. The objective of maximizing revenue in this case is the same as maximizing the sum of weights. That is, we want valid sets $S_1, \ldots, S_m$ maximizing $\sum_{i=1}^{m} d(S_i)$. This problem is NP-hard. The NP-hardness follows from a simple reduction from a variant of classical "bin packing" problem and is described in Appendix Appendix A.

2.4. Notations

Recall from Section 2.1 that $\mathsf{Nbr}_{\mathsf{left}}(i)$ denotes the set of all subscribers whose angular separation from i is at most ρ in the counter-clockwise direction and $\mathsf{Nbr}_{\mathsf{right}}(i)$ represents the same in the clockwise direction. Without loss of generality, we assume that the direction θ_i of a subscriber i with respect to the AP is different for each $i \in U$.

To capture spatial closeness of subscribers, we define a directed graph (which is a simple directed cycle) (U, E). To define this graph consider the unit radius circle centered at the origin and passing though the set of radial coordinates $(1, \theta_1), (1, \theta_2), \ldots, (1, \theta_n)$. Now we walk along the circle in counter-clockwise direction and add an edge directed (i, j) to E if the point $(1, \theta_j)$ is encountered immediately after $(1, \theta_i)$. Additionally, we say that subscribers i and j are *contiguous*. We define the e function $\mathsf{successor}$ as follows: for an element (i, j) of E, $\mathsf{successor}(i) = j$. We use $\mathsf{path}(a, b)$ to denote the set of subscribers in the path from a to b in this graph.

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

$$\mathsf{NbrHood}(S) = \mathsf{Nbr}_{\mathsf{left}}(a) \cap \mathsf{Nbr}_{\mathsf{right}}(b),$$

if there exist $a, b \in S$ such that $S \subseteq \mathsf{Nbr}_{\mathsf{left}}(a) \cap \mathsf{Nbr}_{\mathsf{right}}(b)$. Intuitively, for a set S, its neighborhood $\mathsf{NbrHood}(S)$ 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. $\mathsf{NbrHood}(S)$ is defined for a set S iff any two subscribers in S are within ρ of each other. For set S, we define $\mathsf{first\text{-}item}(S)$ as the subscriber j such that for some subscriber k we have $\mathsf{NbrHood}(S) = \mathsf{Nbr}_{\mathsf{left}}(j) \cap \mathsf{Nbr}_{\mathsf{right}}(k)$. The subscriber k is defined as $\mathsf{last\text{-}item}(S)$. See Figure 1.

3. Subscriber-centric Bandwidth Provisioning

In this section, we describe a dynamic programming formulation that provides a max-min fair allocation in a network with a single AP possessing fewer antennas (m) than

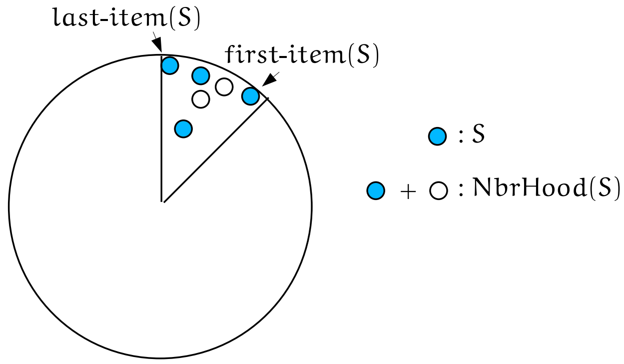

Figure 1: The subscribers denoted by the shaded circles are part of set S. The NbrHood(S) contains all subscribers between first-item(S) and last-item(S).

the number of non-interfering channels (i.e., under the restrictions of *RestAntChan*). 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 that has this special form. Note that this transformation is an existential proof (non-algorithmic) and the algorithm itself is described in the next subsection.

Let $\mathcal{F}$ denote any optimal solution with ordered sets $O_1, \ldots, O_m$, so that (U, E) has a simple path with the subsequence $\mathsf{first\text{-}item}(O_1), \mathsf{first\text{-}item}(O_2), \ldots, \mathsf{first\text{-}item}(O_m)$. We now establish some properties of the sets in the optimal solution. Let $B_j(\mathcal{F})$ be the bandwidth vector for subscribers associated with antenna j. The following lemma notes the simple fact that in the solution $\mathcal{F}$ if $|O_j| = c$, then $B_j(\mathcal{F}) = (\frac{1}{c}, \ldots, \frac{1}{c})$.

Lemma 3.1. *Let $\mathcal{F}$ be any optimal solution of a subscriber-centric instance with sets $O_1, \ldots, O_m$. For any set O_j with $|O_j| = c$, $B_j(\mathcal{F}) = \left(\frac{1}{c}, \ldots, \frac{1}{c}\right)$.*

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 fairer solution (contradiction). □

As a consequence of the above lemma, we have the following max-min equilibrium condition: when all the antennas are omni-directional ($\rho = 360$), the fairest allocation vector allocates a bandwidth of m/n to each subscriber.

A set S of subscribers is said to be a *desirable* set if: (i) all the subscribers in S are contiguous, i.e., there exists $a, b \in U$ such that $\mathsf{path}(a, b) = S$, and (ii) $\mathsf{NbrHood}(S)$ is defined. We reassign subscribers in $O_1, \ldots, O_m$ to create a new solution $\mathcal{F}'$ with sets $S_1, \ldots, S_m$ such that $|O_j| = |S_j|$ and S_j is desirable. Note that since the cardinality of every set remains the fairness is not affected (from Lemma 3.1).

The algorithm for replacing each O_j with S_j is described in Lemma 3.2. As shown in Figure 2, this algorithm replaces subscriber sets that cover geographically overlap-

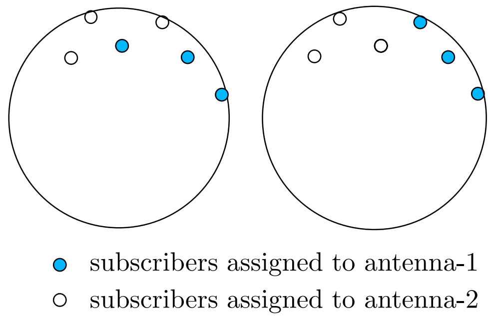

Figure 2: The figure on the left corresponds to the sets O_1 (shaded set) and O_2 (unshaded set). The figure to the right corresponds to the sets S_1 (shaded set) and S_2 (unshaded set) obtained by running the algorithm described in Lemma 3.2.

ping areas into sets that cover non-overlapping areas. Recall that all the subscribers in $\mathsf{NbrHood}(O_j)$ (neighborhood of O_j) are within angular distance ρ from each other. $\mathsf{NbrHood}(O_j)$ includes all the members of O_j and possibly members of other sets. The following lemma proves the existence of solution $\mathcal{F}'$.

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

Proof. Let O_i be a set such that for no other set O_k $(k \neq i)$, $\mathsf{NbrHood}(O_k) \subseteq \mathsf{NbrHood}(O_i)$ (there always such a set O_i). Consider $\mathsf{first\text{-}item}(O_i)$ and $\mathsf{last\text{-}item}(O_i)$. Let

$$\mathsf{path}(\mathsf{first\text{-}item}(O_i), \mathsf{last\text{-}item}(O_i)) = (u_1, \ldots, u_a)$$

where $u_1 = \mathsf{first\text{-}item}(O_i)$ and $u_a = \mathsf{last\text{-}item}(O_i)$. Now if

$$\mathsf{path}(\mathsf{first\text{-}item}(O_i), \mathsf{last\text{-}item}(O_i)) = O_i$$

then this means that the subscribers in O_i are contiguous and we set $S_i = O_i$. Otherwise, we run the algorithm Replace (Figure 3) to create S_i.

Let $\Gamma \subseteq \{O_1, \ldots, O_n\}$ be those sets for which the corresponding desirable set have not been created. So initially, $\Gamma = \{O_1, \ldots, O_n\}$.

Consider the sets O_i and O_p before and after the execution a single iteration of the "for" loop of the algorithm Replace. For simplicity, let us denote the sets O_i and O_p before the execution by $\hat{O}_i$ and $\hat{O}_p$, and the sets O_i and O_p after the execution by O_i^* and O_p^*.

There are three cases that arise. Note that the case $\mathsf{first\text{-}item}(O_p) \in \mathsf{NbrHood}(O_i)$ and $\mathsf{last\text{-}item}(O_p) \in \mathsf{NbrHood}(O_i)$ cannot arise by the way we chose O_i.
Case 1: If $\mathsf{first\text{-}item}(O_p) \in \mathsf{NbrHood}(O_i)$ and $\mathsf{last\text{-}item}(O_p) \notin \mathsf{NbrHood}(O_i)$ then due to the "if" part of the algorithm,

$$\mathsf{path}(\mathsf{first\text{-}item}(O_i^*), \mathsf{last\text{-}item}(O_i^*)) \subseteq \mathsf{path}(\mathsf{first\text{-}item}(\hat{O}_i), \mathsf{last\text{-}item}(\hat{O}_i)), \text{ and}$$
$$\mathsf{path}(\mathsf{first\text{-}item}(O_p^*), \mathsf{last\text{-}item}(O_p^*)) \subseteq \mathsf{path}(\mathsf{first\text{-}item}(\hat{O}_p), \mathsf{last\text{-}item}(\hat{O}_p)).$$

```
for l = 1 to a
    if u_l ∉ O_i then
        let p be such that u_l ∈ O_p and O_p ∈ Γ (if no such p exists then continue)
        if first-item(O_p) ∈ NbrHood(O_i) then
            temp ← last-item(O_i)
            O_i ← O_i \ {last-item(O_i)} ∪ {u_l}
            O_p ← O_p \ {u_l} ∪ {temp}
        else
            temp ← first-item(O_i)
            O_i ← O_i \ {first-item(O_i)} ∪ {u_l}
            O_p ← O_p \ {u_l} ∪ {temp}
S_i = O_i
```

Figure 3: Algorithm Replace.

Case 2: If $\mathsf{first\text{-}item}(O_p) \notin \mathsf{NbrHood}(O_i)$ and $\mathsf{last\text{-}item}(O_p) \in \mathsf{NbrHood}(O_i)$ then due to the "else" part of the algorithm,

$$\mathsf{path}(\mathsf{first\text{-}item}(O_i^*), \mathsf{last\text{-}item}(O_i^*)) \subseteq \mathsf{path}(\mathsf{first\text{-}item}(\hat{O}_i), \mathsf{last\text{-}item}(\hat{O}_i)), \text{ and}$$
$$\mathsf{path}(\mathsf{first\text{-}item}(O_p^*), \mathsf{last\text{-}item}(O_p^*)) \subseteq \mathsf{path}(\mathsf{first\text{-}item}(\hat{O}_p), \mathsf{last\text{-}item}(\hat{O}_p)).$$

Case 3: If $\mathsf{first\text{-}item}(O_p) \notin \mathsf{NbrHood}(O_i)$ and $\mathsf{last\text{-}item}(O_p) \notin \mathsf{NbrHood}(O_i)$ then due to the "else" part of the algorithm,

$$\mathsf{path}(\mathsf{first\text{-}item}(O_i^*), \mathsf{last\text{-}item}(O_i^*)) \subseteq \mathsf{path}(\mathsf{first\text{-}item}(\hat{O}_i), \mathsf{last\text{-}item}(\hat{O}_i)), \text{ and}$$
$$\mathsf{path}(\mathsf{first\text{-}item}(O_p^*), \mathsf{last\text{-}item}(O_p^*)) = \mathsf{path}(\mathsf{first\text{-}item}(\hat{O}_p), \mathsf{last\text{-}item}(\hat{O}_p)).$$

Therefore, at the end of the execution of the algorithm Replace, we have a desirable set S_i, and each O_k $(k \neq i)$ can be covered by an antenna (because as shown above the angular span of the O_k's doesn't increase). We now remove O_i from the Γ and repeat the above construction on the new $\Gamma = \{O_1, \ldots, O_{i-1}, O_{i+1}, \ldots, O_m\}$. The sets $S_1, \ldots, S_m$ we obtain at the end will be all desirable. □

3.1. Dynamic Programming Formulation

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

For every $a, b \in U$ the dynamic programming considers allocating exactly k $(\leq m)$ antennas to subscribers in $\mathsf{path}(a, b)$, a sub-problem whose optimal solution is denoted by $M[a, b, k]$. Intuitively, $M[a, b, k]$ denotes the optimal bandwidth allocation vector obtained by allocating k antennas to cover all the nodes in $\mathsf{path}(a, b)$. Remember that $\mathsf{path}(a, b)$ defines a set of subscribers in the path between a and b in (U, E).

Recursive Formula in Dynamic Programming: Let ϕ be a special symbol, indicating that no feasible solution exists for the sub-problem. For any $a, b \in U$, $M[a, b, k]$ is set as

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

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

$$\underbrace{\left(\frac{1}{|\mathsf{path}(a,b)|}, \ldots, \frac{1}{|\mathsf{path}(a,b)|}\right)}_{\text{total of } |\mathsf{path}(a,b)| \text{ times}}$$

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

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

Theorem 3.3. *The above algorithm finds the optimum solution for the subscriber-centric provisioning problem (under the restrictions of RestAntChan) in* $O(n^4 m)$ *time.*

4. Provider-centric Bandwidth Provisioning

As in the previous section, we consider a network with a single AP possessing fewer antennas (m) than the number of non-interfering channels (i.e., under the restrictions of *RestAntChan*). We describe a simple algorithm achieving ≈ 2-approximation for this restricted version of provider-centric provisioning. More precisely, we show that the revenue our algorithm generates is (almost) at least half of the revenue achievable by any optimal algorithm.

The algorithm Greedy specified in Figure 4 packs all the subscribers in U into valid sets (defined in Section 2.3). To enable the packing, we construct an ordered list L from the set of subscribers U. The list L starts with a subscriber t satisfying

$$d(\mathsf{Nbr}_{\mathsf{right}}(t)) = \min\{d(\mathsf{Nbr}_{\mathsf{right}}(i)) : i \in U\}.$$

Two subscribers i, j are consecutive in L if $(i,j) \in E$, where E is the edges of the graph constructed in Section 2.4. We start greedy packing from the lowest unpacked subscriber (in the list L) and close an antenna for packing purposes if either the antenna cannot accommodate an subscriber or if it reaches its maximum range (ρ). Therefore all the sets constructed by Greedy are valid.

```
j ← 1
i ← t
D ← ∅
first-item(S_1) ← i
while i ≤ n
    if d_i + d(S_j) ≤ 1 and
            S_j ∪ {i} ⊆ Nbr_left(first-item(S_j))
        S_j ← S_j ∪ {i}
        i ← successor(i)
    else
        D ← D ∪ d(S_j)
        j ← j + 1
        first-item(S_j) ← i
D ← D ∪ d(S_j)
output the sum of m largest entries in D
```

Figure 4: Algorithm GREEDY.

We now analyze the performance of GREEDY. Consider any optimal solution to a provider-centric instance. Let $O_1, \ldots, O_m$ be the sets of an optimal solution ordered by the position of their last subscriber in L, i.e., $\mathsf{last\text{-}item}(O_j)$ appears before $\mathsf{last\text{-}item}(O_{j+1})$ in L. Let $\mathsf{OPT} = \sum_{j=1}^{m} d(O_j)$. For any $j \in \{1, \ldots, m\}$, let $\bar{O}_j$ be the set derived from O_j by removing from it all the subscribers (if any) in $\mathsf{Nbr}_{\mathsf{right}}(t)$. Let $\bar{\mathcal{O}} = \{\bar{O}_1, \ldots, \bar{O}_m\}$ be the corresponding ordered set.

We define a graph G with its vertices as the sets in $\bar{\mathcal{O}}$ and edges between $\bar{O}_i$ and $\bar{O}_j$ if the antennas representing $\bar{O}_i$ and $\bar{O}_j$ overlap, i.e., $\mathsf{NbrHood}(\bar{O}_i) \cap \mathsf{NbrHood}(\bar{O}_j) \neq \emptyset$. We start by proving the approximation ratio of GREEDY for a special case.

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

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

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

- Group $\mathcal{P}_a$: sets $S \in \mathcal{P}$ satisfying $d(S) < 1/2$ and $d(\mathsf{Nbr}_{\mathsf{left}}(\mathsf{first\text{-}item}(S))) \geq 1/2$.
- Group $\mathcal{P}_b$: sets $S \in \mathcal{P}$ satisfying $d(S) < 1/2$ and $d(\mathsf{Nbr}_{\mathsf{left}}(\mathsf{first\text{-}item}(S))) < 1/2$.
- Group $\mathcal{P}_c$: sets $S \in \mathcal{P}$ satisfying $d(S) \geq 1/2$.

All the three groups are disjoint. The group $\mathcal{P}_a$ consists of sets with weight less than $1/2$ formed by antennas that are closed prematurely before covering a sector of ρ. The group $\mathcal{P}_b$ consists of sets with weight less than $1/2$ formed by antennas that are closed after satisfying all subscribers in a sector of ρ. Every set $S_b \in \mathcal{P}_b$ is formed by an antenna covering a sector of ρ, and therefore in every optimal solution no subscriber appearing before $\mathsf{first\text{-}item}(S_b)$ in L can be assigned to the same antenna as a subscriber following

```
i ← 1
j ← 1
while j ≤ g
    if S_j ∈ P_a and first-item(S_j) ∈ U \ Nbr_right(t)
        S̄_i ← S_j ∪ S_{j+1}
        j ← j + 1
    else
        S̄_i ← S_j
    P̄ ← P̄ ∪ S̄_i
    j ← j + 1
    i ← i + 1
```

Figure 5: Algorithm REASSIGN.

$\mathsf{last\text{-}item}(S_b)$ in L. The group $\mathcal{P}_c$ consists of sets with weight at least $1/2$. The greedy packing strategy results in the following lemma.

Lemma 4.2. *Let* $S_a \in \mathcal{P}_a$ *with* $\mathsf{first\text{-}item}(S_a) \in U \setminus \mathsf{Nbr}_{\mathsf{right}}(t)$. *Then set* $S_{a+1} \in \mathcal{P}_c$.

Proof. The set S_a is closed because the next subscriber (say p) considered for packing has $d_p > 1 - d(S_a)$. By definition $d(S_a) < 1/2$. Therefore $d_p > 1/2$. This implies that S_{a+1} which accommodates p has $d(S_{a+1}) > 1/2$ and therefore S_{a+1} belongs to $\mathcal{P}_c$. □

We now transform the sets in $\mathcal{P}$ by eliminating sets in the group $\mathcal{P}_a$. This is done using the algorithm REASSIGN is described in Figure 5. Any set $\bar{S}$ constructed in the "if" part of the loop would satisfy $d(\bar{S}) > 1$. The sets in $\bar{\mathcal{P}}$ are also partitioned into groups as follows:

- Group $\bar{\mathcal{P}}_a$: sets $\bar{S} \in \bar{\mathcal{P}}$ satisfying $d(\bar{S}) > 1$.
- Group $\bar{\mathcal{P}}_b$: sets $\bar{S} \in \bar{\mathcal{P}}$ satisfying $d(\bar{S}) < 1/2$.
- Group $\bar{\mathcal{P}}_c$: sets $\bar{S} \in \bar{\mathcal{P}}$ satisfying $1/2 \leq d(\bar{S}) \leq 1$.

Each set in $\bar{\mathcal{P}}_a$ is formed due to the algorithm REASSIGN and corresponds to two sets in $\mathcal{P}$. We now show that the total requirement satisfied by GREEDY is close to half of what could be satisfied by any optimal solution.

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

Proof. If $g \leq 2m$, Lemma 4.1 implies that $\mathsf{OPT} \leq 2 \cdot \mathsf{SOL}$. So we next concentrate on the case where $g > 2m$. Note that $\mathsf{OPT} \leq m$. We divide our analysis as:
Case 1: We first consider the case where the sum of weight of every consecutive pair of sets (S_i, S_{i+1}) in $\mathcal{P}$ 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 $\mathsf{OPT} \leq 2 \cdot \mathsf{SOL}$.
Case 2: We next consider the case where there exists a pair of consecutive sets (say S_q and S_{q+1}) such that the sum of weight of these sets is less than 1, i.e., $d(S_q) + d(S_{q+1}) <$

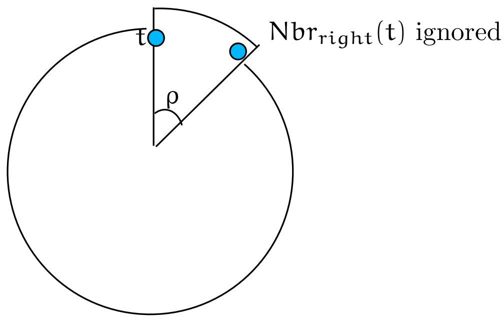

Figure 6: All the subscribers in the sector defined by $\mathsf{Nbr}_{\mathsf{right}}(t)$ are ignored in Case 2 of Lemma 4.3.

1. If there exists an unique pair of $(q, q+1)$ then again sum of top $2m$ entries in $\mathcal{D}$ is at least m. Lets us now assume $d(S_q) + d(S_{q+1}) < 1$ and $d(S_r) + d(S_{r+1}) < 1$ for $q + 1 \leq r$. If $d(S_q) \leq 1/2$, then $d(\mathsf{Nbr}_{\mathsf{right}}(t)) \leq d(S_q) \leq 1/2$ (as q^{th} antenna spans a sector of ρ). Otherwise if $d(S_q) > 1/2$, then $d(S_{q+1}) < 1/2$. Also because of the existence of pair $(r, r+1)$, the $q + 1^{st}$ antenna spans a sector of ρ, implying $d(\mathsf{Nbr}_{\mathsf{right}}(t)) \leq d(S_{q+1}) < 1/2$.

The remainder of the proof can be viewed as a charging scheme in which we show that for every set O in the optimal solution we need to assign (at most) two sets S and T generated by GREEDY such that $d(S) + d(T) \geq d(O)$. We remove all the subscribers in the sector $\mathsf{Nbr}_{\mathsf{right}}(t)$ in the analysis. Since $d(\mathsf{Nbr}_{\mathsf{right}}(t)) < 1/2$, $\sum_{k=1}^{m} d(O_k) \leq \sum_{k=1}^{m} d(\bar{O}_k) + 1/2$. We now construct $\widehat{\mathcal{P}} \subseteq \bar{\mathcal{P}}$ by assigning sets from $\bar{\mathcal{P}}$ to sets in $\bar{\mathcal{O}}$. We scan the sets of $\bar{\mathcal{O}}$ in order, doing the following when we consider the k^{th} element $(\bar{O}_k)$ from $\bar{\mathcal{O}}$.

If $\sum_{S \in \widehat{\mathcal{P}}} d(S) \geq \sum_{j=1}^{k} d(\bar{O}_j)$, then we do not change $\widehat{\mathcal{P}}$. Otherwise, consider the set $\mathsf{path}(t, \mathsf{last\text{-}item}(O_k))$. For a set $\bar{O}_j$, let $\ell(j)$ and $f(j)$ be the index number of the set in $\bar{\mathcal{P}}$ containing $\mathsf{last\text{-}item}(\bar{O}_j)$ and $\mathsf{first\text{-}item}(\bar{O}_j)$ respectively, i.e., $\mathsf{last\text{-}item}(\bar{O}_j) \in \bar{S}_{\ell(j)}$ and $\mathsf{first\text{-}item}(\bar{O}_j) \in \bar{S}_{f(j)}$.

We scan $\bar{S}_{\ell(k)}, \bar{S}_{\ell(k)-1}, \ldots, \bar{S}_1$ in order until we encounter the first set $\bar{S}_u \notin \widehat{\mathcal{P}}$. If $d(\bar{S}_u) \geq 1$, then assign $T_1 \leftarrow \bar{S}_u$ and $T_2 \leftarrow \emptyset$. Otherwise, we continue scanning until we reach the next set $\bar{S}_v \notin \widehat{\mathcal{P}}$. Again, if $d(\bar{S}_v) \geq 1$, set $T_1 \leftarrow \emptyset$ and $T_2 \leftarrow \bar{S}_v$. Otherwise, assign $T_1 \leftarrow \bar{S}_u$ and $T_2 \leftarrow \bar{S}_v$ (in this case, both $d(\bar{S}_u)$ and $d(\bar{S}_v)$ are less than 1). Add T_1 and T_2 to $\widehat{\mathcal{P}}$. Sets T_1 and T_2 are called the *representatives* of $\bar{O}_k$.

We only consider the case where both $\bar{S}_u$ and $\bar{S}_v$ exist (not null). The cases where only one of $\bar{S}_u$ or $\bar{S}_v$ exists are straight forward and omitted. The proof can be completed inductively by showing:

$$\sum_{S \in \widehat{\mathcal{P}}} d(S) \geq \sum_{j=1}^{k} d(\bar{O}_j).$$

For the base 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 trivially true. Otherwise, we use the case distinctions in Figure 7 to complete the proof.

By the inductive hypothesis, we assume that $\sum_{S \in \widehat{\mathcal{P}}} d(S) - d(T_1) - d(T_2) \geq \sum_{j=1}^{k-1} d(\bar{O}_j)$.

If $\ell(1) = f(1) + 1$:
$d(\bar{O}_1) \leq d(\bar{S}_{f(1)}) + d(\bar{S}_{\ell(1)}) = d(T_1) + d(T_2)$.
Else if $\ell(1) > f(1) + 1$:
$d(\bar{O}_1) \leq 1 \leq d(\bar{S}_{\ell(1)}) + d(\bar{S}_{\ell(1)-1}) = d(T_1) + d(T_2)$.

Figure 7: Inequalities for the basis of induction

For the cases where $d(T_1) \geq 1$, or $d(T_2) \geq 1$, or $d(T_1) + d(T_2) \geq 1$, the proof is straightforward. Therefore, in the only remaining case at least one of $d(T_1)$ or $d(T_2)$ is less than $1/2$. Assume w.l.o.g. that $d(T_1) < 1/2$. Therefore, T_1 belongs to group $\bar{\mathcal{P}}_b$.

From the construction of $\widehat{\mathcal{P}}$, we know that no set in $\bar{O}_1, \ldots, \bar{O}_{k-1}$ has its $\mathsf{last\text{-}item}$ in $\bar{S}_u (= T_1)$. Furthermore, since $T_1 \in \bar{\mathcal{P}}_b$, we infer that antenna corresponding to T_1 covers a sector of ρ. We divide the sets $\bar{O}_1, \bar{O}_2, \ldots, \bar{O}_k$ into two parts. The first part $\bar{O}_1, \ldots, \bar{O}_l$ contains the sets whose $\mathsf{last\text{-}item}$ are contained in sets before $\bar{S}_u$ (in ordered set $\bar{\mathcal{P}}$) and the second part consists of the remaining sets.

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

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

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:

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

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

Finally, we note that for the m sets in $\bar{\mathcal{O}}$ we assign $2m$ sets from $\mathcal{P}$ to $\widehat{\mathcal{P}}$ such that

$$\sum_{j=1}^{m} d(O_j) \leq \sum_{j=1}^{m} d(\bar{O}_j) + 1/2 \leq \sum_{S \in \widehat{\mathcal{P}}} d(S) + 1/2.$$

This completes the proof of the charging scheme. Note that since GREEDY picks m heaviest sets in $\mathcal{P}$, it gives the claimed bound. □

The running time of the algorithm GREEDY is linear in the number of subscribers. The constant of proportionality arising in the theorem is what relates the revenue to subscriber requirement in the statement of the provider-centric provisioning (see Section 2.3).

Theorem 4.4. *There exists a linear time algorithm for the problem of provider-centric provisioning (under the restrictions of* RestAntChan*) which produces a revenue which is at least half of the optimal revenue minus* $\Lambda/2$.

5. Practical Considerations

In this section, we attempt to relax two restrictions imposed on the bandwidth provisioning problems so far: (i) the restriction that number of antennas on an AP does not exceed the number of non-interfering channels, and (ii) the presence of only one AP in the network. We find these problems to be significantly less amenable to a theoretical study similar to that presented in the previous sections. We develop heuristics that build upon the algorithms developed for the restricted models.

5.1. Channel Assignment

In practice, wireless networks divide available bandwidth[2] into (possibly overlapping) channels. For example, in 802.11 networks there are 11 overlapping channels. Only 3 channels (say channels 1, 6, and 11) can be considered non-interfering. Consider a single AP network where the number of antennas exceeds the number of available non-interfering channels. (An example could be a single AP employing 802.11b and possessing 4 directional antennas.) At least two antennas would have to share the same channel in such a network. This has two important design implications that we must consider.

- Bandwidth Sharing: All antennas using the same channel share the aggregate bandwidth available to the channel.
- Interference: Orienting multiple antennas using the same channel such that the regions covered by them overlap may cause unacceptable degradation of signal quality and throughput.

Theoretical properties. These phenomena change the nature of our allocation problems. In particular, instead of treating each antenna as capable of sustaining a data rate of b bits/sec, *we must treat each non-interfering channel as the entity with a data rate of* b; all antennas using this channel must share the data rate b amongst them. We suspect that the subscriber-centric problem continues to be easy (i.e., not NP-hard). The provider-centric problem however continues to be NP-hard.

Heuristics. Label the channels $\mathsf{C1}, \mathsf{C2}, \ldots$. Also, denote the antennas as $\mathsf{A1}, \mathsf{A2}, \ldots$. Our heuristic for subscriber-centric provisioning, SUB-HEUR, works as follows. SUB-HEUR first derives antenna orientations and assignments of subscribers to antennas given by the dynamic programming based algorithm presented in Section 3.1. It then associates a channel to each antenna in a round-robin fashion (antennas $\mathsf{A1}$ gets $\mathsf{C1}$; $\mathsf{A2}$ gets $\mathsf{C2}$, etc.) Whenever multiple antennas covering intersecting sectors get associated with the same channel (this can happen in regions of high subscriber density), SUB-HEUR replaces these with a single antenna. It enlarges the angular span of this antenna to cover the total area covered by the antennas being replaced.

Our heuristic for provider-centric provisioning, PROV-HEUR, operates in a similar manner. It first derives antenna orientations and assignments of subscribers to antennas given by the approximation algorithm GREEDY (Section 4). It then associates channels with the antennas using the round-robin scheme described for SUB-HEUR. Whenever multiple antennas covering intersecting sectors get associated with the same channel

[2] In this section, we use the term bandwidth in its usual sense and distinguish it from data rate.

PROV-HEUR replaces these with a single antenna. It then considers the set of subscribers covered by the antennas being replaced in a decreasing order of their data rate requirements (and hence in a decreasing order of the revenue they would generate.) The subset whose requirements can be met by one antenna are retained; the rest are discarded (meaning not assigned to any antenna.) Finally, PROV-HEUR enlarges the angular span of this antenna to cover all the retained subscribers.

5.2. Multiple Access Points

The chief difficulty when allocating bandwidth in a network with multiple APs arises when dealing with a subscriber that can potentially be covered by more than one APs. For such subscribers, we need to decide association with which AP would best serve the objective of the network provider.

Theoretical properties. The hardness of the subscriber-centric problem in a network with multiple APs (with or without the restriction *RestAntChan*) is open. The provider-centric problem clearly is NP-hard (since it was NP-hard even with one AP). However, unlike for the single AP case, we do not yet have an algorithm with provable approximation ratio. We develop intuitively appealing heuristics for these problems and evaluate them using simulation in Section 6.

Heuristics. Each AP first determines antenna orientations and subscriber assignments *independently* using the heuristics developed above. For sets of antennas with intersecting sectors, the heuristics then attempt to modify the association of channels to antennas. The heuristic for subscriber-centric allocation operates by attempting to equally divide subscribers in intersecting areas among available antennas. The heuristic for provider-centric allocation attempts to include subscribers with highest revenue and drops subscribers that can not be accommodated. Remaining antennas using the same channel and covering intersecting sectors are merged and their angular spans adjusted appropriately.

6. Experimental Evaluation

We first demonstrate the efficacy of our algorithms in a single AP network. We then evaluate our heuristics in more practical settings when there are multiple APs.

6.1. Evaluation Methodology

We use the OPNET [18] network simulator version 12.0 to evaluate our algorithms. The transmission range of the AP is 50 meters. The APs employ the 802.11b MAC protocol that has 3 non-interfering channels (we will use channels 1, 6, and 11.) The peak data rate per channel is 11 Mbps. In a simulated AP, one MAC layer component can support multiple transmitter components, each of which connects to a directional antenna. Each MAC layer has one receiver component. The subscribers are assumed to employ directional antennas with very narrow angular spans. We assume bi-directional UDP traffic in our experiments. We emphasize here that the conclusions we make from our simulation results should apply even when subscribers do not have directional antennas for transmission/reception. The bandwidth allocation algorithms proposed in this paper do not necessitate directional antennas at the subscribers and only impact the bandwidth allocated for communication from the AP to the subscribers. Each simulation is performed for a duration of 600 seconds. We assume that each AP has enough antennas to cover the entire circle (with radius 50 meters) centered at it.

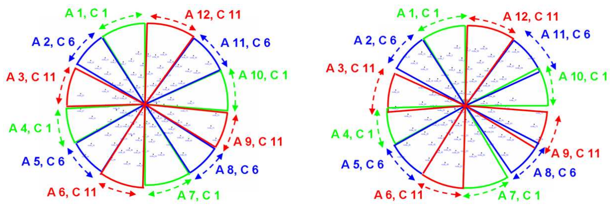

(a) Baseline (b) SUB-HEUR

Figure 8: Subscriber-centric provisioning: Antenna orientations for the Uniform subscriber distribution. In this case, our heuristic (SUB-HEUR) derives antenna orientations and subscriber assignments very similar to those derived by the algorithm oblivious to subscriber locations (baseline). *Note: In all the figures the pair Ax, Cy denotes antenna x using channel y.*

Throughout this evaluation, we compare the performance of our algorithms with a *baseline* algorithm that operates as follows: each AP orients its antennas so that they cover the entire area and the sectors covered by any two antennas do not intersect. The baseline algorithm represents a simple strategy that is oblivious of subscriber locations and provides an useful point of comparison.

6.2. Subscriber-centric Provisioning In Single-AP Networks

We begin by evaluating our heuristics for a single-AP network. We fix the number of subscribers to 100. We consider a variety of geographical distributions of subscribers. We present results for two distributions for which we use the following labels: (i) *Uniform*, in which subscriber locations are chosen uniformly at random over the given area and (ii) *Non-Uniform*, in which 80 subscribers are uniformly distributed within a small region (12% of the total area) and 20 subscribers are uniformly distributed in the remaining area. We assume 12 directional antennas on the AP, each with a span of 30°.

Figure 8 compares the antenna orientations derived by baseline and SUB-HEUR for the Uniform subscriber distribution. Note that for this distribution, baseline is the (statistically) optimal bandwidth allocation scheme. We observe that by deriving an antenna orientation close to that derived by baseline, SUB-HEUR is able to provide a fair allocation.

Figure 9 compares the antenna orientations derived by baseline and SUB-HEUR for the Non-Uniform subscriber distribution. We observe that the solution provided by SUB-HEUR exhibits the intuitively desirable feature of orienting more antennas where more subscribers reside. In Figure 10, we present the data rates available to the subscribers. Since there are 3 channels and subscribers using a channel share its bandwidth equally, there are 3 groups of subscriber throughputs. SUB-HEUR improves fairness significantly by ensuring that each channel was shared by an almost equal number of subscribers.

6.3. Provider-centric Provisioning In Single-AP Networks

We present results for an AP with 3 directional antennas each with a default span of 120°. (We chose a smaller number of antennas (3 as opposed to 12 as in the last section) to enable the evaluation of optimal revenue via an exhaustive search of all feasible antenna orientations. We present the performance of PROV-HEUR for the Non-Uniform subscriber

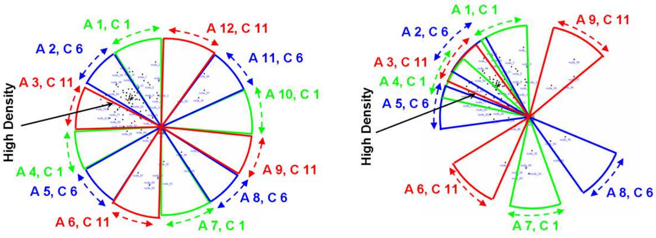

(a) Baseline (b) Sub-heur

Figure 9: Subscriber-centric provisioning: Antenna orientations for the Non-Uniform subscriber distribution. In this case, our heuristic (Sub-heur) derives antenna orientations and subscriber assignments very different from those derived by the baseline algorithm. In particular, more antennas are assigned to the region with higher concentration of subscribers.

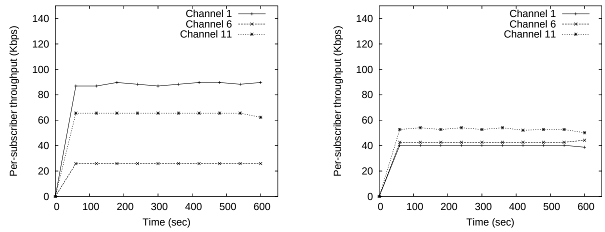

(a) Baseline (b) Sub-heur

Figure 10: Subscriber-centric provisioning: Comparison of the performance of the subscriber location oblivious (baseline) algorithm and our algorithm (Sub-heur). The baseline algorithm provides poor fairness - subscribers in the densely populated region receive much lower data rates than others. Sub-heur results in substantially improved fairness.

distribution in Figure 11. Figures 11(a) and (b) compare the antenna orientations for baseline and Prov-heur when we assume that each subscriber specifies a data rate requirement of 10 Kbps and provides a revenue of 10 units if its requirement is met. We assume homogeneous data requirements/revenue numbers because 802.11b does not provide support for differentially splitting bandwidth among subscribers assigned to an antenna. The emerging 802.11e protocol provides such support but is not implemented in OPNET yet. Notice that Prov-heur might leave some subscribers uncovered, as it is not trying to maximize the number of subscribers covered, bur rather it tries to maximize the revenue generated from the allocation.

In Figure 11 (part (c)), we compare the revenues generated by the baseline algorithm, Prov-heur, and the optimal solution for a variety of subscriber throughput demands. We see that the Prov-heur greatly improves upon the revenue obtained using the baseline algorithm. Even more encouragingly, we see that the performance of our Prov-heur is very close to that of the optimal solution.

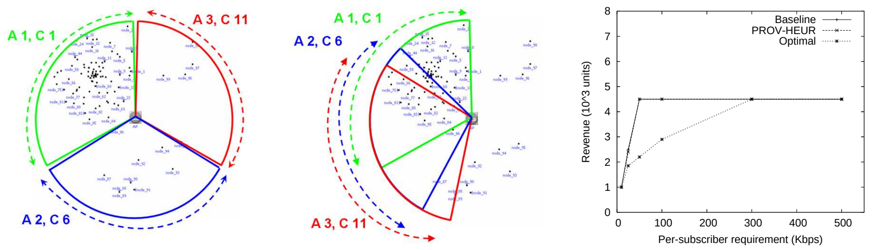

(a) Baseline (b) Prov-heur (c) Revenue Comparison

Figure 11: Provider-centric provisioning: Figures (a) and (b) compares the antenna orientations provided by baseline and Prov-heur. The figure (c) on the right compares the revenue provided by baseline and Prov-heur. The latter is found to closely match the optimal revenue (determined using an exhaustive search.

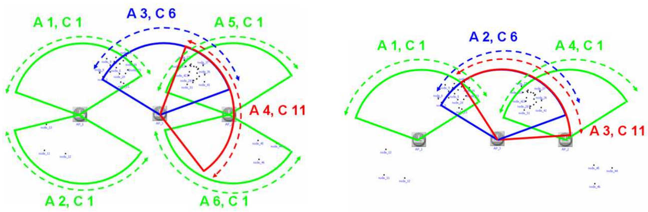

(a) Subscriber-centric (b) Provider-centric

Figure 12: Multiple APs: Antenna orientations for the subscriber-centric and the provider-centric provisioning.

6.4. Networks With Multiple APs

We depict the outcome of experiments with 3 APs in Figures 12 and **??**. We assume two antennas on each AP. The APs are placed on a line, with adjacent APs 50 meters apart. We choose two regions with high subscriber densities. Figures 12(a) and (b) depict the antenna orientations derived by the heuristics described in Section 6.4. We would like to highlight two key observations. First, our heuristics successfully derive channel assignments for antennas so as to avoid interference. Second, and more interestingly, notice the difference in the antenna orientations in the two cases. Whereas in the subscriber-centric setting (Figure 12(a)), the heuristic covers *all* the subscribers (as a max-min fair allocation should), in the provider-centric setting (Figure 12(b)), the heuristic chooses to cover the subscribers in the dense regions and leaves out some subscribers. The aggregate data rate provided is the same in both the cases, despite fewer antennas being used the provider-centric setting. This is because the antennas of an AP that share a channel also share its bandwidth. For example, the data rate provided by antenna A1 in Figure 12(b) is equal to that provided together by antennas A1 and A2 in Figure 12(a).

7. Conclusions and Future Research Directions

We studied the problem of bandwidth provisioning in wireless networks employing directional antennas in two contrasting settings: (i) subscriber-centric, where the objective is to max-min fairly allocate bandwidth among the subscribers and (ii) provider-centric, where the objective is to maximize the revenue generated by satisfying the bandwidth requirements of subscribers. In the subscriber-centric setting, we provide a dynamic programming algorithm that achieves a max-min fair allocation. In the provider-centric setting, the allocation problem becomes NP-hard and we provide a 2-approximation algorithm for it. We also support our theoretical results using extensive simulations.

Our research opens up a number of directions for future research. First, the problem of joint antenna orientation and channel assignment in networks with multiple APs is an interesting future direction. Second, in this paper, we considered a setting where the subscribers are static or move very infrequently. This allows our 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 direction.

8. Acknowledgements

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.

[1] S. Sen, B. Raman, Long distance wireless mesh network planning: Problem formulation and solution, in: WWW, ACM, 2007, pp. 893–902.

[2] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, S. R. Das, Mobisteer: Using steerable beam directional antenna for vehicular network access, in: ACM MobiSys, ACM, 2007, pp. 192–205.

[3] S. Bandyopadhyay, S. Roy, T. Ueda, Enhancing the performance of ad hoc wireless networks with smart antennas, CRC Press, 2006.

[4] Y. Kuga, J. Cha, J. A. Ritcey, J. Kajiya, Mechanically steerable antennas using dielectric phase shifters, in: Proceedings of the IEEE- APS and URSI, 2004, pp. 161–164.

[5] S. Yi, Y. Pei, S. Kalyanaraman, On the capacity improvement of ad hoc wireless networks using directional antennas, in: Proceedings of the ACM MobiHoc '03, 2003, pp. 108–116.

[6] C. Peraki, S. Servetto, On the maximum stable throughput problem in random networks with directional antennas, in: Proceedings of the ACM MobiHoc '03, 2003, pp. 76–87.

[7] S. Vasudevan, J. Kurose, D. Towsley, On neighbor discovery in wireless networks with directional antennas, in: Proceedings of the IEEE INFOCOM '05, 2005, pp. 590–600.

[8] A. Nasipuri, S. Ye, J. You, R. Hiromoto, A MAC protocol for mobile ad hoc networks using directional antennas, in: Proceedings of the IEEE WCNC '00, 2000, pp. 1214–1219.

[9] L. Bao, J. Garcia-Luna-Aceves, Transmission scheduling in ad hoc networks with directional antennas, in: Proceedings of the ACM MOBICOM '02, 2002, pp. 48–58.

[10] R. Choudhury, X. Yang, R. Ramanathan, N. Vaidya, Using directional antennas for medium access control in ad hoc networks, in: Proceedings of the ACM MOBICOM '02, 2002, pp. 59–70.

[11] J. Zhu, S. Roy, 802.11 mesh networks with two-radio access points, in: Proceedings of the IEEE ICC '05, 2005, pp. 3609–3615.

[12] D. Bertsekas, R. Gallager, Data Networks, Prentice Hall, 1987.

[13] X. L. Huang, B. Bensaou, On max-min fairness and scheduling in wireless ad-hoc networks: analytical framework and implementation, in: Proceedings of the ACM MobiHoc '01, ACM, 2001, pp. 221–231.
URL `http://doi.acm.org/10.1145/501416.501447`

[14] H. Brönnimann, M. T. Goodrich, Almost optimal set covers in finite vc-dimension, Discrete & Computational Geometry 14 (4) (1995) 463–479.

[15] K. L. Clarkson, K. R. Varadarajan, Improved approximation algorithms for geometric set cover, in: Proceedings of the ACM SoCG '05, 2005, pp. 135–141.

[16] R. Ramanathan, On the performance of beamforming antennas, in: Proceedings of the ACM MobiHoc '01, 2001, pp. 95–105.

[17] P. Berman, J. Jeong, S. P. Kasiviswanathan, B. Urgaonkar, Packing to angles and sectors, in: Proceedings of the ACM SPAA '07, 2007, pp. 171–180.

[18] Opnet modeler version 12.0, 2007, `http://www.opnet.com` (2007).

[19] B. Raman, K. Chebrolu, Design and evaluation of a new mac protocol for long-distance 802.11 mesh networks, in: MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking, ACM, New York, NY, USA, 2005, pp. 156–169. doi:http://doi.acm.org/10.1145/1080829.1080847.

[20] T. Korakis, G. Jakllari, L. Tassiulas, A mac protocol for full exploitation of directional antennas in ad-hoc wireless networks, in: MobiHoc '03: Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing, ACM, New York, NY, USA, 2003, pp. 98–107. doi:http://doi.acm.org/10.1145/778415.778428.

[21] R. R. Choudhury, N. H. Vaidya, Deafness: A mac problem in ad hoc networks when using directional antennas, in: ICNP '04: Proceedings of the 12th IEEE International Conference on Network Protocols, IEEE Computer Society, Washington, DC, USA, 2004, pp. 283–292.

[22] A. A. Lazar, A. Orda, D. E. Pendarakis, Virtual path bandwidth allocation in multiuser networks, IEEE/ACM Trans. Netw. 5 (6) (1997) 861–871. doi:http://dx.doi.org/10.1109/90.650145.

[23] D. P. Heyman, T. V. Lakshman, A. L. Neidhardt, A new method for analysing feedback-based protocols with applications to engineering web traffic over the internet, in: SIGMETRICS '97: Proceedings of the 1997 ACM SIGMETRICS international conference on Measurement and modeling of computer systems, ACM, New York, NY, USA, 1997, pp. 24–38. doi:http://doi.acm.org/10.1145/258612.258670.

[24] S. K. Biswas, R. Izmailov, Design of a fair bandwidth allocation policy for vbr traffic in atm networks, IEEE/ACM Trans. Netw. 8 (2) (2000) 212–223. doi:http://dx.doi.org/10.1109/90.842143.

[25] S. B. Fred, T. Bonald, A. Proutiere, G. Régnié, J. W. Roberts, Statistical bandwidth sharing: a study of congestion at flow level, SIGCOMM Comput. Commun. Rev. 31 (4) (2001) 111–122. doi:http://doi.acm.org/10.1145/964723.383068.

[26] Z. Duan, Z.-L. Zhang, Y. T. Hou, Service overlay networks: Slas, qos, and bandwidth provisioning, IEEE/ACM Trans. Netw. 11 (6) (2003) 870–883. doi:http://dx.doi.org/10.1109/TNET.2003.820436.

[27] M. Allalouf, Y. Shavitt, Centralized and distributed algorithms for routing and weighted max-min fair bandwidth allocation, IEEE/ACM Trans. Netw. 16 (5) (2008) 1015–1024. doi:http://dx.doi.org/10.1109/TNET.2007.905605.

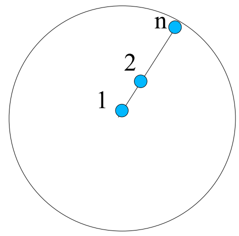

Figure A.13: NP-hardness of provider-centric provisioning. The reduction from bin-packing puts all subscribers on one straight line.

Appendix A. NP-Hardness of Provider-centric Provisioning

An instance of bin packing consists of a finite set U of items, a rational weight $w_i \in (0, 1]$ for each item u in U, and a positive integer m. The question is whether there exists a partition of the items of U into disjoint subsets $S_1, \ldots, S_m$, such that sums of sizes of items in each S_i is no more than 1.