Interference-Constrained Wireless Coverage in a Protocol Model

Prateek R. Kapadia*
Indian Institute of Technology, Bombay, India
prateek@it.iitb.ac.in

Om P. Damani
Indian Institute of Technology, Bombay, India
damani@it.iitb.ac.in

ABSTRACT

We present an $O(n \log n)$ algorithm to compute the coverage map of a given set of transmitters under interference constraints. That is, we compute the set of points that lie within the transmission range of one transmitter and lie outside the interference range of every other transmitter. To our knowledge, there is no existing satisfactory algorithm for this purpose. We assume that the transmission and interference ranges of each transmitter are circular disks.

We show that for an appropriate choice of 'distance measure', coverage at each point can be computed by considering only certain 'proximate' transmitters. Hence, we partition the plane into proximity regions and the coverage in these proximity regions is computed considering only proximate transmitters. We use an extension of Voronoi diagrams, called 'Power' diagrams, to represent the proximity regions.

Categories and Subject Descriptors: C.2.1 [Computer-Communication Networks]: Network Architecture and Design – *Wireless Communication, Network Topology*

General Terms: Algorithms

Keywords: Wireless Network Coverage and Interference, Computational Geometry

1. INTRODUCTION

The wireless network designer is often faced with the task of determining the coverage regions of a set of transmitters. The transmitters may correspond to base stations in an infrastructure network, wireless access points in a wireless local area network, or transmitter nodes in a wireless ad-hoc network. The inputs to a designer are the transmitter locations, their transmission power, and the channel model (modulation scheme, path loss, fading, etc.). Co-channel interference occurs at a receiver during simultaneous wireless communication between two transmitter-receiver pairs.

*Author is an employee of Reliance Communications, India, and a research scholar at IIT, Bombay

MSWiM'06, October 2–6, 2006, Torremolinos, Malaga, Spain.

Suppose the coverage map of a given set of transmitters is known. Then the designer can attempt to improve the coverage by varying the location of the transmitters, or their transmission power. Recently, Ahmed et al. ([1]) have proposed algorithms for transmission power assignment to access points (APs) under interference constraints. Their solutions also assume a protocol model similar to ours. They describe the computation of exact coverage to be a 'mathematically daunting' task. Our purpose in this work is to compute the coverage map efficiently using simple computational geometric primitives.

1.1 Problem Statement

We are given n transmitter locations (points) in the plane. We are also given the transmission and interference ranges of each transmitter. All transmitters simultaneously share the same wireless channel. We need to compute the set of points that lie within the transmission range of one transmitter, and outside the interference range of every other. We call this set the 'coverage region' of a transmitter. The union of all n coverage regions is the 'coverage map' of the network.

1.2 Model

Our solution makes the following assumptions:

1. The transmission and interference ranges of each transmitter is a circular disk centered at the transmitter location. Also, the transmission range is less than the interference range. (We use the terms "range" and "disk" interchangeably.)
2. Physical layer characteristics of the medium are not modeled. The model used is a "Protocol" Model, following Gupta et al. ([4]).
3. Each receiver has the same receive sensitivity. This means that the interference and transmission ranges are independent of the physical characteristics of the receiver.

This model idealizes the omni-directional transmitter in the plane.

The shaded region in Figure 1 shows the coverage region of one transmitter surrounded by 7 other transmitters. The interference ranges (disks) are shown in solid perimeter and the transmission disk of transmitter p is shown with a dotted perimeter.

1.3 Solution Approach

We show that for an appropriate choice of 'distance measure', coverage at each point can be computed by consider-

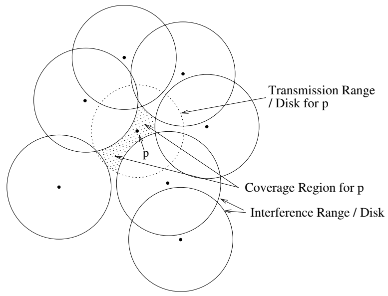

Figure 1: Coverage Region of a Transmitter

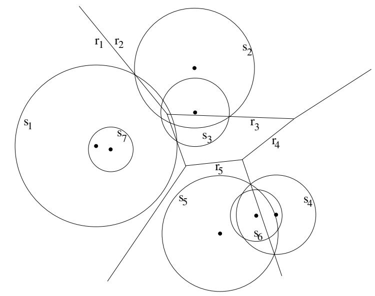

Figure 2: A Power Diagram: 7 disks, 5 regions (adapted from [2])

ing only certain 'proximate' transmitters. Hence, we partition the plane into 'proximity' regions and the coverage in these proximity regions is computed considering only proximate transmitters. We argue that the coverage region for a transmitter must necessarily lie in its proximity region. We further show that there exists a partition of this proximity region such that coverage in each partition can be decided by considering the interference from only one 'nearby' neighbor. Our algorithm uses this fact and computes the coverage map in $O(n \log n)$ time.

We use an extension of the well-known *Voronoi Diagram* ([9]) called the *Power Diagram* ([2]). Voronoi and power diagrams have been applied to wireless coverage problems in the context of sensor networks - by So et al.([10]), and before them by Meguerdichian et al.([8]). The work in [10] has been the main motivation and starting point for us. However, we are not aware of any published application of Voronoi or power diagrams to coverage under interference constraints.

2. COVERAGE REGIONS FROM A POWER DIAGRAM

The power diagram is based on a distance measure between a point and a disk - called the *power distance*. We denote the power distance by ρ. The power distance of a point p from a disk $\bigcirc$ of radius r and center c is defined by $\rho(p, \bigcirc) = d(p, c)^2 - r^2$; where $d(p, c)$ is the Euclidean distance between points p and c.

A power diagram of interference disks with equal radii is the Voronoi diagram of the disk centers ([2]). This observation led us to develop our approach first with equal ranges and the Voronoi diagram, and then later extend it to unequal ranges and the power diagram. The details of transition from Voronoi to power diagrams are available in our technical report ([5]).

Figure 2 shows a power diagram of 7 disks in the plane. The power diagram for a set of disks, Ψ, partitions the plane into convex polytopes. A point lies in the *power region* corresponding to disk $\bigcirc \in \Psi$ if its power distance from $\bigcirc$ is less than its power distance from every other disk in Ψ.

The next observation relates the distance measure ρ to interference disks.

OBSERVATION 2.1. *Each point on the interference disk for transmitter $\tilde{p}$ is also on the interference disk of every transmitter that is closer to the point than $\tilde{p}$, by distance measure ρ.*

PROOF. Let $\tilde{q} \in T \setminus \{\tilde{p}\}$. Let x be closer, by power distance measure ρ, to $\tilde{q}$ than $\tilde{p}$. Let $\bigcirc_{\tilde{p}}$ and $\bigcirc_{\tilde{q}}$ be the interference disks of $\tilde{p}$ and $\tilde{q}$, respectively.

Since $\rho(x, \bigcirc_{\tilde{q}}) < \rho(x, \bigcirc_{\tilde{p}})$, if x is in the interference range of $\tilde{p}$, $\rho(x, \tilde{p}) < 0$. Which means that $\rho(x, \tilde{q}) < 0$, i.e. x must also be in the interference range of $\tilde{q}$. □

We now state notation and expressions for a power region, its extreme points, a power neighborhood, and the power diagram.

The power region corresponding to a transmitter $\tilde{p} \in T$ is given by:
$\triangle(\tilde{p}, T) = \{x \mid \rho(x, \tilde{p}) < \rho(x, \tilde{q}), \forall \tilde{q} \in T \setminus \{\tilde{p}\}\}$
The extreme points of the power region for $\tilde{p}$ are given by:
$\partial(\tilde{p}) = \{x \mid \rho(x, \tilde{p}) \leq \rho(x, \tilde{q}), \forall \tilde{q} \in T \setminus \{\tilde{p}\}\} \setminus \triangle(\tilde{p})$
The power neighbors of transmitter $\tilde{p}$ are given by:
$\Gamma(\tilde{p}) = \{\tilde{q} \mid \partial(\tilde{p}) \cap \partial(\tilde{q}) \neq \phi\}$
The power diagram is given by:
$\mathbb{P}(T) = \bigcup_{\tilde{p} \in T} \partial(\tilde{p})$
When the set T is obvious from context, we denote the power region of $\tilde{p} \in T$ simply by $\triangle(\tilde{p})$.

We now prove an important property of the power region - it shows that the coverage region for a transmitter is confined to that transmitter's power region.

LEMMA 2.1 (COVERAGE IN POWER REGION). *Each point outside $\triangle(\tilde{p})$ that is on the transmission disk for $\tilde{p}$ is also on some interference disk other than that of $\tilde{p}$.*

PROOF. Let $\tilde{q} \in T \setminus \{\tilde{p}\}$, and $x \in \triangle(\tilde{q})$. Since the power diagram is a partition of the plane, such a $\tilde{q}$ exists. Thus, x is closer to $\tilde{q}$, in power distance, than it is to $\tilde{p}$. The transmission disk of $\tilde{p}$ is a subset of its interference disk; thus x lies on the interference disk of $\tilde{p}$. Hence, by Observation 2.1 x must also be on the interference disk of $\tilde{q}$. □

Lets assume that the power diagram for T has been built by some classical method ([2]). Suppose we remove the transmitter $\tilde{p}$ from T, and build $\mathbb{P}(T \setminus \{\tilde{p}\})$. For example, Figure 3 shows the effect of removing $\tilde{p}$:

1. Extension of some existing edges - for example, in Figure 3, solid edges extend dashed edges from the original power diagram
2. Deletion of some existing edges - for example, dashed edges enclosing r_3 in Figure 3
3. Addition of new vertices - for example, intersections of solid edge-extensions in Figure 3
4. Addition of new edges - for example, the solid edge between the two new vertices in Figure 3

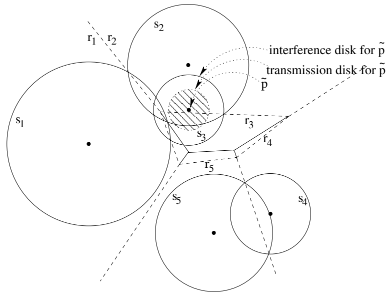

Figure 3: Power Diagram with Power Frame for $\tilde{p}$

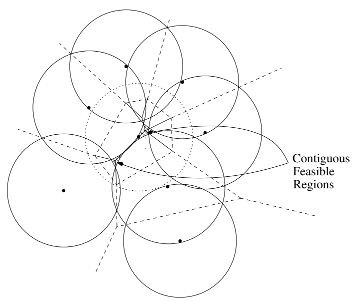

Figure 4: Feasible Coverage Area: Multiple Contiguous Feasible Regions

In Figure 3, s_3 is the interference disk of transmitter $\tilde{p}$. Note that only the power neighborhood of $\tilde{p}$ changes to form the new power diagram. We show, in Lemma 2.2, that this is true for every power diagram that does not have an empty power region. We call the new set of vertices, edges, and extension edges the *Power Frame* corresponding to the removed transmitter. Solid portions in Figure 3 represent the power frame.

Before we delve into our discussion, we define certain terms in the context of what has been said so far.

Feasible Coverage Area: For a given transmitter $\tilde{p}$, this is the set of points lying in the power region of $\tilde{p}$, but outside every interference disk other than that for $\tilde{p}$. We reiterate that due to Lemma 2.1, to compute coverage, we can restrict our attention only to the power region of $\tilde{p}$.

Actual Coverage Area: For a given transmitter $\tilde{p}$, this is the set of points that lie in its Feasible Coverage Area and also on its transmission disk. The Actual Coverage Area of $\tilde{p}$ is the intersection of the transmission disk of $\tilde{p}$ with the Feasible Coverage Area.

Contiguous Feasible Region: The Feasible Coverage Area may be composed of several disjoint maximal simply-connected subsets (see, for example, the two shaded sets in Figure 4). These subsets are called Contiguous Feasible Regions. In the subsequent text we use the term 'Feasible Region' when the word 'Contiguous' is obvious from the context.

Power Frame: For a given transmitter $\tilde{p}$, this is the set of points on extension edges, new edges and vertices in $\triangle(\tilde{p})$ obtained from deleting the point $\tilde{p}$ and adding new extreme points from the power diagram of $T \setminus \{\tilde{p}\}$. The power frame is thus the set of points in $\mathbb{P}(T \setminus \{\tilde{p}\}) \setminus \mathbb{P}(T)$. We next demonstrate how the Power Frame aids in reasoning about the Feasible and Actual Coverage Areas.

Power regions may be empty - for example, disks s_6 and s_7 in Figure 2 have empty power regions. Some of our formal results will require that no power region be empty. In Subsection 2.1, we justify that this assumption does not affect the coverage map in practical contexts. We also provide a method to remove transmitters with an empty power region. Thus the resulting power diagram will have no empty power regions.

Each edge in a power diagram corresponds to two power neighbors. We now show that each edge in the Power Frame for $\tilde{p}$ must correspond to two power neighbors of $\tilde{p}$, in $\mathbb{P}(T)$. We first observe a more general result - that all points in a power region are closer, by power distance ρ, to power neighbors than to any other transmitter.

In the ensuing discussion, we will implicitly assume the use of the distance measure ρ, the power distance; that is, we will say 'closest' or 'closer' (respectively, 'farthest' or 'farther') to mean closest or closer (farther, farthest, respectively) in power distance.

LEMMA 2.2. *Assume that every transmitter has a non-empty power region. Let x be a point in the power region of transmitter $\tilde{p}$, i.e. $\triangle(\tilde{p})$. If $\tilde{q}$ is the closest (by power distance) transmitter in $T \setminus \{\tilde{p}\}$ to x, then $\tilde{q}$ is a power neighbor of $\tilde{p}$.*

PROOF. $\triangle(\tilde{p}, T)$ is the power region of $\tilde{p}$ in the power diagram of the set of transmitters T. By definition, $x \in \triangle(\tilde{p}, T) \cap \triangle(\tilde{q}, T \setminus \{\tilde{p}\})$.

Assume that $\tilde{q}$ is not a power neighbor of $\tilde{p}$. We show that this leads to a contradiction.

Case 1 : Assume that an extreme point x_0 of $\triangle(\tilde{p}, T)$ lies in $\triangle(\tilde{q}, T \setminus \{\tilde{p}\})$. Let $\tilde{q_0}$ be a power neighbor of $\tilde{p}$ corresponding to the point x_0.

$\Rightarrow \rho(x_0, \tilde{p}) = \rho(x_0, \tilde{q_0})$

$\Rightarrow$ {Since $\tilde{q}$ is closer to x_0 than $\tilde{q_0}$,} $\rho(x_0, \tilde{p}) = \rho(x_0, q_0) > \rho(x_0, \tilde{q})$

But $x_0 \in \triangle(\tilde{p})$. Hence, this is a contradiction since no transmitter can be closer to x_0 than $\tilde{p}$. Thus, no extreme points of $\triangle(\tilde{p}, T)$ lie in $\triangle(\tilde{q}, T \setminus \{\tilde{p}\})$.

Case 2 : Assume that no extreme point of $\triangle(\tilde{p}, T)$ lies in $\triangle(\tilde{q}, T \setminus \{\tilde{p}\})$.

$\Rightarrow$ All edges in $\triangle(\tilde{p}, T)$ lie outside $\triangle(\tilde{q}, T \setminus \{\tilde{p}\})$.

$\Rightarrow$ Either $\triangle(\tilde{q}, T \setminus \{\tilde{p}\}) \subset \triangle(\tilde{p}, T)$, or $\triangle(\tilde{q}, T \setminus \{\tilde{p}\}) \cap \triangle(\tilde{p}) = \phi$.

(In other words, the power region of $\tilde{p}$ in $\mathbb{P}(T)$ either encloses that of $\tilde{q}$ in $\mathbb{P}(T \setminus \tilde{p})$, or the two power regions are disjoint.)

$\{x \in (\triangle(\tilde{p}, T) \cap \triangle(\tilde{q}, T \setminus \{\tilde{p}\})) \Rightarrow (\triangle(\tilde{q}, T \setminus \{\tilde{p}\}) \cap \triangle(\tilde{p})) \neq \phi\}$

$\Rightarrow \triangle(\tilde{q}, T \setminus \{\tilde{p}\}) \subset \triangle(\tilde{p}, T)$

$\Rightarrow$ each point in $\triangle(\tilde{q}, T \setminus \{\tilde{p}\})$ is closer to $\tilde{p}$ than $\tilde{q}$.

$\Rightarrow \triangle(\tilde{q}, T) = \phi$

This contradicts the assumption that no power region is empty. □

This result leads to important observations - Power Frame edges lie in the power region of the removed transmitter, and Power Frame edges correspond only to power neighbors.

COROLLARY 2.2.1. *If a point $x \in \perp(\tilde{p})$ is closest to transmitter $\tilde{q} \in T \setminus \{\tilde{p}\}$, then $\tilde{q}$ is a power neighbor of $\tilde{p}$.*

PROOF. Only points in $\triangle(\tilde{p})$ may be mapped to different partitions upon removal of $\tilde{p}$ from $\mathbb{P}(T)$. Thus, $\perp(\tilde{p}) \subset \triangle(\tilde{p})$. Lemma 2.2 states that the closest transmitter to x in $T \setminus \{\tilde{p}\}$ is a neighbor of $\tilde{p}$. □

Let $\nu(x, S)$ denote the set of transmitters in S closest to x.

$$\nu(x,S) = \{U \subset S | \; \forall \tilde{p} \in U, \tilde{q} \in S \setminus \{\tilde{p}\}, \rho(x, \bigcirc_{\tilde{p}}) \leq \rho(x, \bigcirc_{\tilde{q}})$$

We now give an expression for the Power Frame corresponding to $\tilde{p}$. Due to Corollary 2.2.1, the Power Frame lies in the power region and has edges only from power neighbors:

$$\perp(\tilde{p}) = \{x \in \triangle(\tilde{p}) \mid \exists \{\tilde{q_1}, \tilde{q_2}\} \subseteq \nu(x, \Gamma(\tilde{p}))\}$$

We compute the Actual Coverage Area for $\tilde{p}$ using its Power Frame. We first note a property of the Power Frame that shows its correlation with the Feasible Coverage Region.

OBSERVATION 2.2. *$\perp(\tilde{p})$ partitions $\triangle(\tilde{p})$, and each partition corresponds to exactly one neighbor in $\Gamma(\tilde{p})$.*

PROOF. $\mathbb{P}(T \setminus \{\tilde{p}\})$, by definition, partitions the plane. By definition, the Power Frame is the subset of this power diagram lying inside $\triangle(\tilde{p})$. Hence $\triangle(\tilde{p})$ is partitioned by the Power Frame. Each point in a partition belongs to some power region in $\mathbb{P}(T \setminus \{\tilde{p}\})$. Due to Lemma 2.2, a point in such a partition can be closest only to a neighbor of p in $\mathbb{P}(T)$. By definition, each point in a partition can be closest only to one point in $T \setminus \{\tilde{p}\}$. Hence, the partition corresponds to exactly one neighbor. □

Note that the edges in the Power Frame in $\triangle(\tilde{p})$ bounding its partition correspond to the edges contributed by a neighbor of $\tilde{p}$. Further, since $\tilde{p}$ is closer than $\tilde{q}$ to each point in this partition, the edge between $\tilde{p}$ and $\tilde{q}$ also bounds the partition.

This observation is illustrated in Figure 3, where the Power Frame for $\tilde{p}$ corresponds to the four partitions of r_3.

We denote by $\blacktriangle(\tilde{p}, \tilde{q})$ the partition of $\triangle(\tilde{p})$ by edges on the Power Frame corresponding to $\tilde{q}$. Formally,

$$\blacktriangle(\tilde{p}, \tilde{q}) = \triangle(\tilde{p}) \cap \{x \mid \tilde{q} \in \nu(x, \Gamma(\tilde{p}))\}$$

We now show that in order to find all points in $\blacktriangle(\tilde{p}, \tilde{q})$ that lie in the interference range of *any* transmitter, it is sufficient to consider points that lie in the interference range of $\tilde{q}$.

COROLLARY 2.2.2. *If $x \in \blacktriangle(\tilde{p}, \tilde{q})$ and x is in the interference range of some transmitter $\tilde{t} \neq \tilde{p}$, then x is in the interference range of $\tilde{q}$.*

PROOF. By definition, every point in $\blacktriangle(\tilde{p}, \tilde{q})$ is closer to $\tilde{q}$ than any other point in T. Since x is closer to $\tilde{q}$ than $\tilde{t}$, by Observation 2.1, x is in the interference range of $\tilde{q}$. □

We denote the Actual Coverage Area of a transmitter $\tilde{p}$ by $\chi(\tilde{p})$. We now prove the principal contribution of this paper - if no power region is empty, then the partition given by the Power Frame allows us to compute the coverage area by excluding interference from just one transmitter.

THEOREM 2.1. *Assume no power region is empty. Then,*

$$\chi(\tilde{p}) = \bigcup_{q \in \Gamma(\tilde{p})} (\bigcirc^t_{\tilde{p}} \cap \blacktriangle(\tilde{p}, \tilde{q})) \setminus \bigcirc^i_{\tilde{q}}$$

PROOF. Let $\bigcirc^i_{\tilde{p}}$ and $\bigcirc^t_{\tilde{p}}$ denote the interference and transmission disks, respectively, of transmitter $\tilde{p}$. Since no power region is empty, Lemma 2.1 implies that the Actual Coverage Area lies inside $\triangle(\tilde{p})$. Also, Observation 2.2 states that the Contiguous Feasible Region for $\tilde{p}$ is composed of contributions from each neighbor. Corollary 2.2.2 shows that to find points within $\blacktriangle(\tilde{p}, \tilde{q})$ that lie in the Feasible Coverage Area, it is sufficient only to exclude points on the interference disk of $\tilde{q}$. Thus the Actual Coverage Area can be computed from the individual regions contributed by each partition. □

The advantage of using the Power Frame is now clear - it yields a partition that allows us to deal with only one interfering transmitter at a time. Also, to compute the Actual Coverage Area, only two arc intersection computations are required for each neighbor $\tilde{q}$ - one for $\bigcirc^t_{\tilde{p}}$, and one for $\bigcirc^i_{\tilde{q}}$.

In order to compute $\chi(\tilde{p})$ we need a generalized polygon representation that allows circular arcs as edges. Berberich et al. ([3]) study such generalized polygons, and the boolean set operations on these generalized polygons. We employ these operations in our algorithm in Section 3.

2.1 Removing Redundant Transmitters

The power region corresponding to a disk may be empty - as is the case with s_6 and s_7 in Figure 2. We show a more general result - if a disk and its corresponding power region have no points in common, then that disk is included in the union of other disks.

LEMMA 2.3 (EMPTY POWER REGIONS). *Let $\tilde{q}$ be a transmitter with interference disk $\bigcirc_{\tilde{q}}$, such that $\bigcirc_{\tilde{q}} \cap \triangle(\tilde{q}) = \phi$. Then,*

$$\bigcirc_{\tilde{q}} \subseteq \bigcup_{\tilde{p} \in T \setminus \tilde{q}} \bigcirc_{\tilde{p}}$$

PROOF. We prove this result by contradiction.

Let $\bigcirc_{\tilde{q}} \nsubseteq \bigcup_{\tilde{p} \in T \setminus \tilde{q}} \bigcirc_{\tilde{p}}$

$\Rightarrow \exists x \in \bigcirc_{\tilde{q}}$ such that $\forall \tilde{p} \neq \tilde{q},\ x \notin \bigcirc_{\tilde{p}}$

$\Rightarrow (\rho(x, \bigcirc_{\tilde{q}}) < 0) \wedge (\forall \tilde{p} \neq \tilde{q},\ \rho(x, \bigcirc_{\tilde{q}}) > 0)$

$\Rightarrow x \in \triangle(\tilde{q})$

$\Rightarrow x \in \triangle(\tilde{q}) \cap \bigcirc_{\tilde{q}}$

This contradicts the assumption that $\triangle(\tilde{q}) \cap \bigcirc_{\tilde{q}}$ is empty. □

If an interference disk is included in the union of other interference disks, the corresponding transmitter has an empty coverage region. Lemma 2.3 implies that a transmitter with an empty power region has an empty coverage region. Thus, in a pratical setting, such a transmitter would cause interference without contributing to coverage. Hence, the removal of this transmitter is justified.

3. ALGORITHM

We collate the observations made in the preceding text into an algorithm. The inputs to the algorithm are: a set T of transmitters, their locations in the plane, and their transmission and interference radii. The algorithm outputs a coverage map for T, denoted by $\widehat{\chi}(T)$.

ALGORITHM 3.1 (COVERAGE MAP).

1. Initialize: $\widehat{\chi}(T) \leftarrow \phi$, $T' \leftarrow T$
2. Compute the power diagram $\mathbb{P}(T)$.
3. For each transmitter $\tilde{p} \in T$, do If $\triangle(\tilde{p}) = \phi$, $T' \leftarrow T' \setminus \{\tilde{p}\}$.
4. For each transmitter $\tilde{p} \in T'$, do
   (a) $\chi(\tilde{p}) \leftarrow \phi$
   (b) Find the power diagram of $\Gamma(\tilde{p})$, i.e.$\mathbb{P}(\Gamma(\tilde{p}))$.
   (c) For each region $\triangle(\tilde{q}, \Gamma(\tilde{p}))$, do

i. $\blacktriangle(\tilde{p},\tilde{q}) \leftarrow \triangle(\tilde{q},\Gamma(\tilde{p})) \cap \triangle(\tilde{p},T')$

ii. $\chi(\tilde{p}) \leftarrow \chi(\tilde{p}) \cup ((\blacktriangle(\tilde{p},\tilde{q}) \cap \bigcirc^t_{\tilde{p}}) \setminus \bigcirc^i_{\tilde{q}})$

5. For each transmitter $\tilde{p} \in T'$, do $\widehat{\chi}(T) \leftarrow \widehat{\chi}(T) \cup \chi(\tilde{p})$
□

3.1 Running Time Analysis

The sum of the number of neighbors over all transmitters is linear in n.

Observation 3.1 (Sum of Neighbors).

$$\sum_{\tilde{p}\in T} |\Gamma(\tilde{p})| = O(n)$$

Proof. The sum $\sum_{\tilde{p}\in T} |\Gamma(\tilde{p})|$ is also the number of ordered pairs $(\tilde{p},\tilde{q})$ such that $\tilde{p}$ and $\tilde{q}$ are neighbors in $\mathbb{P}(T)$. Since each power edge in $\mathbb{P}(T)$ corresponds to two transmitters, the latter is twice the number of power edges, which is $O(n)$ ([2]). □

Theorem 3.1 (Runtime). *The coverage map of n transmitters can be constructed in $O(n \log n)$ time.*

Proof. Step 2: A power diagram of n disks in the plane can be constructed in $O(n \log n)$ time ([2]).

Step 4b: The power diagram of $\Gamma(\tilde{p})$ can be constructed in $O(|\Gamma(\tilde{p})| \log |\Gamma(\tilde{p})|$ time. Now, $\log |\Gamma(\tilde{p})| \le \log n$

$\Rightarrow \sum_{\tilde{p}\in T} |\Gamma(\tilde{p})| \log |\Gamma(\tilde{p})| \le (\log n) \sum_{\tilde{p}\in T} |\Gamma(\tilde{p})|$

$\Rightarrow$ {By Observation 3.1} $(\log n) \sum_{\tilde{p}\in T} |\Gamma(\tilde{p})| = O(n \log n)$

The total time to compute the power diagrams for all transmitters is thus $O(n \log n)$.

Step 4(c)i: A well-known algorithm ([9]) for convex polygon intersection can be used to compute the partition. This algorithm is linear in the total number of edges, i.e. in our case $O(|\Gamma(\tilde{p})|)$. By Observation 3.1, the total time taken executing this step is $O(n)$.

Step 4(c)ii: The union and set difference operations can be performed by the sweep-line algorithm from [3]. This computation is also linear time in the number of line segments (edges) and arcs; i.e. in our case $O(|\Gamma(\tilde{p})|)$. By Observation 3.1, the total time taken executing this step is $O(n)$.

Step 4c: Each transmitter $\tilde{p}$ can contribute a partition only to a neighbor (Lemma 2.1). Thus the total number of partitions created by the algorithm is the sum of neighbors, which is $O(n)$ by Observation 3.1. Since each edge appears in at most two partitions, the total number of edges created in this step is also $O(n)$.

Hence, the coverage map can be computed in $O(n \log n)$ time. □

4. RELATED WORK

Voronoi and power diagrams have been applied by So et al. in [10] to solve a sensor network coverage problem. Their model does not include interference constraints. They determine whether a given set of n sensors covers all points in a field such that each point is covered by *at least* k sensors. However, even for $k = 1$, computing interference-limited coverage requires us to ensure that *exactly one disk* includes a point, which is a slightly harder problem.

Ahmed et al.([1]) study algorithms for transmission power assignment to access points (APs) under interference constraints. They select a random AP to increase power if more coverage is required, and decrease power *by an arbitrary amount* if it interferes with an existing power assignment. The model chosen there too is a protocol model. Our results can directly be applied to this problem - exactly computing how much power needs to be increased or decreased for an AP without the need for choosing arbitrary powers iteratively. They also compute a utility function that needs a coverage map. They use sampling in the absence of a known algorithm for computing the coverage map.

Kuhn et al.([7]) obtain a power assignment for each base station given a finite set of receiver locations. They solve a discrete optimization problem that minimizes the maximum number of base stations interfering with any receiver.

4.1 Future Directions

1. An immediate extension to this work is to allow up to k transmitters to interfere with a receiver. Higher-order power diagrams seem to be the natural tools to wield, as indicated by the generalizations in [10].
2. We assume constant signal strength at every point within a transmission and interference range. The decrease of signal strength with distance and time may be modeled with generalizations of Voronoi diagrams of point sets - Voronoi diagrams of convex objects ([6]), for example. An appropriate generalization would take us a step closer to modeling the physical layer - both in the contour of the coverage regions and received signal strength.

Acknowledgments: We thank Dr. Vinod Vasudevan (Reliance Communications), Dr. Amitabh Chaudhary (Univ. of Notre-Dame), and Dr. Abhay Karandikar (IIT, Bombay) for many fruitful discussions on our model and approach.

5. REFERENCES

[1] N. Ahmed and S. Keshav. A successive refinement approach to wireless infrastructure deployment. In *Wireless Comm. and Netw. Conf.*, 2006.

[2] F. Aurenhammer. Power diagrams: Properties, algorithms and applications. *SIAM J. Comput.*, 16:78–96, 1987.

[3] E. Berberich, A. Eigewillig, M. Hemmer, S. Hert, K. Melhorn, and E. Schömer. A computational basis for conic arcs and boolean operations on conic polygons. In *European Symp. on Algorithms*, pages 174–186, 2002.

[4] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans. Info. Theory*, pages 388–404, 2000.

[5] P. R. Kapadia and O. P. Damani. Interface–constrained wireless coverage in a protocol model. Tech. Rep. No. 21, 2006, KReSIT, IIT Bombay. http://www.it.iitb.ac.in/research/techreport/reports/21.pdf.

[6] M. I. Karavelas and M. Yvinec. The Voronoi diagram of planar convex objects. In *European Symp.on Algorithms*, pages 337–348, 2003.

[7] F. Kuhn, P. von Rickenbach, R. Wattenhofer, E. Welzl, and A. Zollinger. Interference in cellular networks: The minimum membership set cover problem. In *Computing and Combinatorics Conference*, pages 188–198, 2005.

[8] S. Meguerdichian, F. Kaushanfar, M. Potconjak, and M. B. Srivastava. Coverage problems in wireless ad-hoc sensor networks. In *IEEE Infocom*, pages 1380–1387, 2001.

[9] J. O'Rourke. *Computational Geometry in C - 2nd ed.* Cambridge University Press, 2001.

[10] M.-C. So and Y. Ye. On solving coverage problems in a wireless sensor network using Voronoi diagrams. In *Workshop on Internet and Network Economics*, 2005.