Minimizing Interference of a Wireless Ad-Hoc Network in a Plane

Magnús M. Halldórsson[1]* and Takeshi Tokuyama[4]**

[1] Dept. of Computer Science, Faculty of Engineering, University of Iceland, IS-107 Reykjavik, Iceland. Email: mmh@hi.is

[2] Graduate School of Information Sciences, Tohoku University, Sendai, 980-8579 Japan. Email: tokuyama@dais.is.tohoku.ac.jp

Abstract. We consider the problem of topology control of a wireless ad-hoc network on a given set of points in the plane, where we aim to minimize the maximum interference by assigning a suitable transmission radius to each point. By using computational geometric ideas and ϵ-net theory, we attain an $O(\sqrt{\Delta})$ bound for the maximum interference where Δ is the interference of a uniform-radius ad-hoc network. This generalizes a result given in [7] for the special case of highway model (i.e., one-dimensional problem) to the two-dimensional case. We also give a method based on quad-tree decomposition and bucketing that has another provable interference bound in terms of the ratio of the minimum distance to the radius of a uniform-radius ad-hoc network.

1 Introduction

Mobile wireless ad-hoc networks are an important subject in recent studies on communication networks. In a popular model, each mobile device is considered as a point (called *node*) in the Euclidean plane, and each node has a disk of a given transmission radius. Two nodes can communicate with each other if they are located within each other's disks.

The transmission radius is a monotone function of the electric power given to the node, which we assume to be a controllable parameter. Topology control involves assigning a suitable transmission radius to each node to form a connected network while minimizing some non-decreasing objective function of the radii. The most frequently studied objective is to minimize the power consumption, or the sum of the electric power given to the nodes. Making disks small has also another benefit, that is, to reduce the *interference*. Interference at a node is the number of disks containing it, and high interference increases the probability of packet collision of packets. Therefore, it is desirable to keep a low interference at every node.

Topology control for minimizing interference is bound to be a different task from that of minimizing energy. Traditionally, this has been addressed implicitly

* Partially supported by grants of the Icelandic Research Fund

** Partially supported by the project *New Horizons in Computing*, Grant-in-Aid for Scientific Research on Priority Areas, MEXT Japan.

by reducing the density of the communication graph. Burkhart *et al.* [2], however, showed that low interference is not implied by sparseness. Also, that networks constructed from nearest-neighbor connections can fail dismally to bound the interference. On the other hand, they gave experimental results that indicate that graph spanners help reduce interference in practice. Their work prompted the explicit study of interference minimization. Moscibroda and Wattenhofer [**?**] gave nearly tight approximation algorithms that bound the *average* interference of nodes.

The recent work of Rickenbach *et al.* [7] is the starting point of our study. They introduced the problem of bounding the maximum interference at a node, and gave algorithms for the special case where all the points are located on a line, called the *highway model*. Their algorithm constructs a network with an $O(\sqrt{\Delta})$ interference, where Δ is the interference of a uniform radius network, while it is shown that there exists an instance that requires $\Omega(\sqrt{n})$ interference. They also showed that the better one of a naive network and the above $O(\sqrt{\Delta})$ interference network attains a $O(\Delta^{1/4})$ approximation ratio.

For the two-dimensional problem, analogous results have not been reported yet (to the authors' knowledge). In this paper, we show that we can construct a network with an $O(\sqrt{\Delta})$ interference for any point set in the plane, extending the theory of [7] to the planar case (and even for any constant-dimensional space). We also give a network with an $O(\log(R/d))$ interference, where d is the minimum distance between points and R is the minimum radius of a uniform-radius network to attain connectivity. Our results rely on computational geometric tools such as local neighbor graphs, ϵ-nets, and quad-tree decompositions.

2 Mathematical formulation and terminology

We are given a set $V = \{\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n\}$ of points in a plane. For each $\mathbf{v}_i$, we assign a positive real number $r(\mathbf{v}_i)$ called the *transmission radius*. This can be considered as a *radius assignment* function

$$r : V \to \mathbb{R}_{>0}\,.$$

Consider the set $\mathcal{D} = \{D_1, D_2, \ldots, D_n\}$ of disks, where D_i has radius $r(\mathbf{v}_i)$ and its center at $\mathbf{v}_i$.

We define a wireless network on V, that is the graph $G(\mathcal{D}) = (V, E)$, where we have an undirected edge $(\mathbf{v}_i, \mathbf{v}_j)$ if and only if $\mathbf{v}_i \in D_j$ and $\mathbf{v}_j \in D_i$. In other words, $\mathbf{v}_i$ and $\mathbf{v}_j$ can directly communicate since they are within the transmission radius of each other. We say that the wireless network $G(\mathcal{D})$ is *feasible* iff it is connected.

The interference of $\mathcal{D}$ at a point $\mathbf{p}$ is the number of disks in $\mathcal{D}$ covering $\mathbf{p}$. That is,

$$I(\mathcal{D}, \mathbf{p}) = |\{i : \mathbf{p} \in D_i\}|.$$

The *interference* of a wireless network $G(\mathcal{D})$ is [1]

$$\max\{I(\mathcal{D}, \mathbf{p}) | \mathbf{p} \in \mathbb{R}^2\}.$$

The *interference minimization problem* is to find a radius assignment r to give a feasible network with the minimum interference.

One natural approach is to increase all radii uniformly until the the graph becomes connected. Let R_{min} be the infimum of the radius such that the network becomes connected, and refer to the network with all radii set to R_{min} as the *uniform-radius* network. Let Δ denote the interference of the uniform-radius network.

Although the problem is clearly an NP-optimization problem, it seems to be very difficult to find the optimal wireless network. Indeed, even the special case where all points V are located on a line (highway model) is considered to be difficult (although NP-hardness result is not known). Thus, we seek for a practical solution with some theoretical quality guarantee, in particular an upper bound of the interference or an approximation ratio to the optimal solution.

2.1 Review for the highway model

We briefly review some results for the highway model given by Rickenbach *et al.* [7]. Suppose that points of V are located on the x-axis in the sorted order with respect to their x-values.

Then, a naive method is to set $r(i) = \max(d(\mathbf{v}_i, \mathbf{v}_{i-1}), d(\mathbf{v}_i, \mathbf{v}_{i+1}))$ for $i = 1, 2, \ldots, n$, where we set $\mathbf{v}_0 = \mathbf{v}_1$ and $\mathbf{v}_{n+1} = \mathbf{v}_n$. It is easy to observe that $G(\mathcal{D})$ associated with this radius function is feasible: the network is called the *linear network*. Unfortunately, there is an example named *exponential chain* instance for which the linear network poorly performs. In the exponential chain, the points satisfies that $d(\mathbf{v}_i, \mathbf{v}_{i+1}) = 2^i$ for $i = 1, 2, \ldots n-1$, and it is easy that the interference of the point v_1 is $n-1$ in the linear network.

We can use a *hub-connected* network to reduce the worst-case interference. The general idea is as follows: We find a subset $W \subset V$ of points called *hubs* and first construct the linear network of hubs. Then, for each $\mathbf{v} \in V \setminus W$, we set

$$r(\mathbf{v}) = \min_{\mathbf{w} \in W} d(\mathbf{v}, \mathbf{w});$$

namely, $\mathbf{v}$ connects to its nearest hub. If we select every $\sqrt{n}$-th points in V as a hub, we have a set W of cardinality $\sqrt{n}$, and it is shown that $I(G(\mathcal{D})) = O(\sqrt{n})$ for this network. It has been shown that the interference is $\Omega(\sqrt{n})$ for the exponential chain, thus the hub-connected network is worst-case optimal. However, for each given instance, we can often design a network with a better interference. Indeed, there is a construction that $I(G(\mathcal{D})) = \sqrt{\Delta}$.

[1] We can also consider the version where we only consider interference at points of V, not all points in the plane. The results of this paper carry immediately over to that model.

2.2 Two-dimensional analogue of the linear network

Although the linear network performs poorly in the worst case, it is a basic structure that can also be constructed in a distributed fashion. That is, each point can connect to its right and left neighbors without the need of global information.

The first task is to extend this notion to the two-dimensional case, where we do not have clear definitions of the left and right neighbors. If we sort the points with respect to x-coordinate, and each point connects to the nearest neighbor with respect to the x-coordinate, we can obtain a feasible network. However, this ignores the y-coordinate, and usually gives a very bad network. Instead, we would like to use the Euclidean distance to measure the proximity of points.

Indeed, a network in which each node establishes (two-way) connection with its nearest neighbor is called a *nearest-neighbor forest*. The nearest-neighbor forest need not be connected, however, and we want give a connected network based on it. The minimum spanning tree $\mathrm{MST}(S)$ might be a direct two-dimensional analogue of the linear network. The wireless version is $\mathrm{WMST}(S)$ in which each node $\mathbf{p}_i$ has the radius $\max_{\mathbf{q}:(\mathbf{p}_i,q)\in \mathrm{MST}(S)} d(\mathbf{p}_i, q)$. Constructing a minimum spanning tree explicitly requires some global information; hence, we prefer graphs of a more local nature.

We briefly explain the *local neighborhood graph* (LNG) [8], since it inspires the construction of our hub-structure network given later.

For each point $\mathbf{p} \in \mathbb{R}^2$, we divide the plane into six cones $R_1(\mathbf{p}), R_2(\mathbf{p}), \ldots, R_6(\mathbf{p})$, where $R_k(\mathbf{p})$ is the region such that the argument angle about $\mathbf{p}$ is in the range $[\frac{(k-1)\pi}{3}, \frac{k\pi}{3})$.

Let $nb_k(\mathbf{p}, V)$ be the nearest point to $\mathbf{p}$ in $V \cap R_k(\mathbf{p})$. See Figure 1. The local neighbor graph $\mathrm{LNG}(V)$ is the graph connecting each $\mathbf{v} \in V$ to its six local neighbors.

The following elementary fact is important and will be used to show an interference bound for our network given later.

Lemma 1. *Suppose that* $\mathbf{u}$ *and* $\mathbf{v}$ *are in* $R_k(\mathbf{p})$ *and* $d(\mathbf{p}, \mathbf{u}) \le d(\mathbf{p}, \mathbf{v})$. *Then,* $d(\mathbf{u}, \mathbf{v}) < d(\mathbf{p}, \mathbf{v})$.

Proof. Straightforward from the fact that the diameter (distance between farthest pair of points) of a fan with the angle $\pi/3$ equals the radius of the circle.

The above lemma leads to the following fact [8], although we do not give a proof since we do not use this fact explicitly in the rest of the paper.

Lemma 2. $\mathrm{LNG}(V)$ *contains* $\mathrm{MST}(V)$. *Consequently, it is connected.*

Let $N_1(\mathbf{v}_i) = \{nb_k(\mathbf{v}_i, V) | 1 \le k \le 6\}$ and $N_2(\mathbf{v}_i) = \{\mathbf{w} \in V | \mathbf{v}_i \in N_1(\mathbf{w})\}$. If we set $r_i = \max\{d(\mathbf{v}_i, \mathbf{q}) | \mathbf{q} \in N_1(\mathbf{v}_i) \cup N_2(\mathbf{v}_i)\}$ for each $i = 1, 2, \ldots n$, we have a wireless network $\mathrm{WLNG}(V)$ that has $\mathrm{LNG}(V)$ as a subgraph.

We remark that $\mathrm{WLNG}(V)$ can be constructed locally: Each node increases its radius (up to a given limit) and sends a message until it receives acknowledgement from the local neighbor in each of six cones, and sends a connection

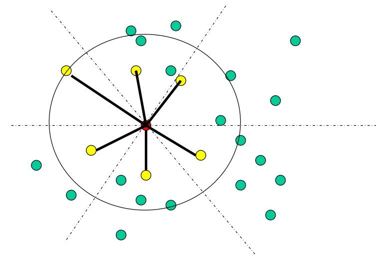

Fig. 1. Local neighbors of a point and a disk connecting them

request to each local neighbor. Then, each node that receives connection request increases the radius such that it can reach the sender. We remark that this method has the defect that we need to set the limit radius, since if there is an empty cone, we have to detect and ignore it to avoid increasing the radius to infinity.

2.3 Hub-connected network with $O(n^{1/2})$ interference

It is known that we can make a bad instance for which any network containing the nearest-neighbor forest has an $\Omega(n)$ interference while there exists a network with a constant interference for the instance [7]. Thus, if every node connects to its nearest neighbor, we can obtain neither a (nontrivial) absolute interference bound nor a rational bound to the optimal interference for each instance.

In order to attain a better interference bound, we consider a hub-connected network, where we select a subset W of V as the set of hubs. We construct WMST(W) as the core of the network, and propagate the connection around the core such that every vertex $\mathbf{v} \in V \setminus W$ is connected to the nearest hub to it. Note that we may use any connected network on W (e.g., $WLNG(W)$) as the core instead of WMST(W) in order to attain our main theoretical result: what is important is the choice of W.

Hub selection using an ϵ-net We apply ϵ-net theory to define the set W of hubs. Consider a family $\mathcal{R}$ of regions in the plane. Given a set V of n points, the pair $(V, \mathcal{R})$ is called a *range space*. An ϵ-net of the range space $(V, \mathcal{R})$ is a subset $S \subset V$ such that any region $R \in \mathcal{R}$ that contains at least ϵn points of V must contain at least one point of S. Intuitively, an ϵ-net is a uniformly distributed sample of V where the uniformity is measured by using the family $\mathcal{R}$ of regions.

The following theory (although readers need not be familiar with it) has many applications such as computational geometry [1] and learning theory: The *Vapnik-Chervonenkis*-dimension (VC dimension) of a range space is the largest size of a subset $A \in V$ such that all subsets of A are attained as an intersection of A and a region in $\mathcal{R}$. If VC dimension is low (say, a constant), we can always have a small ϵ-net (see [4] for example).

Here, we consider a range space associated with a family of sectors of disks. Consider a unit disk D, and divide it into six cone sectors P_k $(k = 1, 2, \ldots, 6)$ by the three diameter chords with argument angles 0, $\pi/3$, and $2\pi/3$. That is, $P_k = \{x \in D | (k-1)\pi/3 \leq arg(x) < k\pi/3\}$.

The family $\mathcal{P}_k$ is the set of all translated/scaled copies of P_k. We consider the family $\mathcal{P} = \cup_{1 \leq k \leq 6} \mathcal{P}_k$. Intuitively, it is the family of "1/6 piece of pies" in six rotated positions of any size located anywhere in the plane.

First, we give a weaker bound for the size of an ϵ-net of $\mathcal{P}$. Although this will be slightly improved later, the following result is useful since we do not need any complicated algorithm to find the ϵ-net.

Theorem 1. *A random sample of size $c\epsilon^{-1} \log \epsilon^{-1}$ becomes an ϵ-net for $\mathcal{P}$ with high probability if c is a sufficient large constant.*

Proof. If we construct an ϵ-net for each of $\mathcal{P}_k$, their union becomes an ϵ-net of $\mathcal{P}$. Thus, it suffices to show the existence of an ϵ-net of size $O(\epsilon^{-1} \log \epsilon^{-1})$ for each of $\mathcal{P}_k$. We have the theorem from the general theory of ϵ-nets [3, 4] of range spaces.

A family $\mathcal{R}$ of regions is said to be a family of *pseudo-disks* if for any non-collinear three points in the plane, there exists a unique $R \in \mathcal{R}$ such that those three points are on the boundary of R. The following better bound is known for a family of pseudo-disks.

Theorem 2. *[5] For any point set V, there is an ϵ-net of size $O(1/\epsilon)$ for a family of pseudo-disks.*

Consider the family $\mathcal{P}_k$ for $k = 1, 2, \ldots, 6$, say, $k = 1$. It is easy to see that for any noncollinear three points in the plane, there exists at most one $P \in \mathcal{P}_1$ such that the triple of points are on the boundary of P. Thus, $\mathcal{P}_1$ has a property that is very similar to pseudo-disks, but there may be triplets of points such that there is no $P \in \mathcal{P}_1$ such that the boundary of P goes through them. Nevertheless, we have the following theorem:

Theorem 3. *There exists an ϵ-net of V of size $O(1/\epsilon)$ for $\mathcal{P}$, and we can compute one in polynomial time.*

This theorem is of independent interest in the area of computational geometry. Since it probably requires too much geometric knowledge for a non-specialist to follow, we give (an outline of) the actual construction of such an ϵ-net later in a separate section.

The hub-connected network The construction is as follows: We first compute an $\sqrt{n^{-1}}$-net W of V such that the size of W is $O(\sqrt{n})$, which can be obtained by using Theorem 3 by setting $\epsilon = \sqrt{n^{-1}}$. Then, we form the wireless network WMST(W) (indeed, any connected network is fine for our purpose). Let $r_0(w)$ be the transmission radius of $w \in W$ in WMST(W).

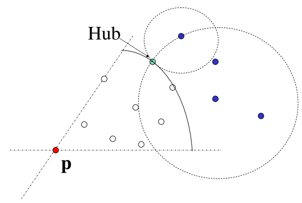

Fig. 2. No disk around a point outside the fan can reach $\mathbf{p}$

We call the elements of W *hubs*. Then, for each $v \in V \setminus W$, we find its nearest hub denoted by $hub(v)$. We set $r(v) = d(v, hub(v))$. For each hub $w \in W$, define the set $N(w) = \{v \in V \setminus W | hub(v) = w\}$. We set $r(w) = \max\{r_0(w), \max_{v \in N(w)} d(v, w)\}$ for each $w \in W$. We have determined r for each elements of v, and thus we obtain a wireless network GHUB(V).

Lemma 3. GHUB(V) *is connected.*

Proof. Since WMST(W) is connected, the induced subgraph of GHUB(V) by W is connected. Since other nodes are all connected to nodes in W, GHUB(V) is connected.

Theorem 4. *The interference of* GHUB(V) *is* $O(\sqrt{n})$.

Proof. Let c be a suitable constant such that $|W| < c\sqrt{n}$. We claim that any point $\mathbf{p} \in \mathbb{R}^2$ is covered by at most $(c+6)\sqrt{n}$ disks, or, more precisely, by $6\sqrt{n}$ disks except those around elements of W. Consider the cusp $R_1(\mathbf{p})$ whose argument angle interval is $[0, \pi/3]$. Because of symmetry, it suffices to show that at most $\sqrt{n}$ points in $R_1(\mathbf{p})$ can contain $\mathbf{p}$ in their disks. If there is no hub in $R_1(\mathbf{p})$, then $R_1(\mathbf{p})$ cannot contain more than $\sqrt{n}$ points because W is a $\sqrt{n^{-1}}$-net, and we are done. Otherwise, we can assume there is at least one hub in $R_1(\mathbf{p})$ (see Figure 2). Let $\mathbf{w}$ be the nearest hub to $\mathbf{p}$ in $R_1(\mathbf{p})$. We draw a circle C of radius $d(\mathbf{p}, \mathbf{w})$ around $\mathbf{p}$, and let P be the region (i.e., piece of pie) obtained as the intersection of the interior of C and $R_1(\mathbf{p})$. Since P does not contain a hub in its interior, P can contain at most $\sqrt{n}$ elements of V. Consider any point $\mathbf{x} \in V$ in $R_1(\mathbf{p}) \setminus P$. Then, it follows that $d(\mathbf{x}, \mathbf{w}) < d(\mathbf{x}, \mathbf{p})$ from Lemma 1 (here, it is crucial that the angle of a fan is $\pi/3$). Since $r(\mathbf{x})$ is the distance to its nearest hub, $r(\mathbf{x}) \leq d(\mathbf{x}, \mathbf{w}) < d(\mathbf{x}, \mathbf{p})$. Thus, $\mathbf{p}$ is not in the disk of $\mathbf{x}$. This completes the proof.

Note that if we use the weaker ϵ-net obtained by random sampling, we set $\epsilon = \sqrt{n^{-1} \log n}$ to have a network with an interference $O(\sqrt{n \log n})$.

3 A network with $O(\sqrt{\Delta})$ interference

Let us consider the uniform-radius network G_0 in which each disk has the same radius R_{min}. Recall that Δ is the interference of G_0. Although Δ can become as large as $\Omega(n)$, it can in practice be much smaller than n, or even $\sqrt{n}$. We show a construction of a network where the interference is $O(\sqrt{\Delta})$.

We use a standard localization method by bucketing. By scaling, we can assume that $R_{min} = 1$ to eliminate one parameter. We partition the plane into unit square buckets by an orthogonal grid. For simplicity of argument, we assume that there are no points on boundaries of buckets; this assumption is easy to remove. G_0 can connect a point $v \in B$ to points in bucket B or its eight neighbors. We say that two buckets B and B' are adjacent if there exists $v \in B$ and $v' \in B'$ such that the edge (v, v') is in G_0.

Lemma 4. *1. For each B, an adjacent bucket must be one of its eight neighbors in the grid.*
2. Each bucket contains $O(\Delta)$ points.

Proof. (i) is obvious, since the distance from any point in B to any bucket other than the eight neighbors is more than 1. For (ii), suppose that a bucket contains more than 4Δ points. We refine the buckets into four sub-buckets of size 0.5×0.5. One of the sub-bucket contains more than Δ points, and the center of the sub-bucket is covered by the unit disk about each point in its sub-bucket. This contradicts that the interference of G_0 is Δ.

Our construction is as follows: First, in each bucket B, we give a network with interference $O(\sqrt{\Delta})$ by using the construction given in the previous subsection, and set the radius of each point accordingly. Note that none of the disks in the construction has a radius larger than $\sqrt{2}$. Second, for each adjacent pair B, B' of buckets, select exactly one edge $(v, v') \in G_0$ connecting them. We call v and v' *connectors*. We enlarge the radius of each connector to 1 (if its current radius is less than 1) .

Now, we have defined all the radii, and accordingly we have a network LHUB(V).

Theorem 5. *The network* LHUB(V) *is connected, and its interference is $O(\sqrt{\Delta})$.*

Proof. The network is connected within each bucket, and the connection between buckets is same as G_0. Thus, it is connected. For each point p, it is interfered by points of at most 21 buckets (the neighbor buckets of Manhattan distance at most two), since the radius of the largest disk is $\sqrt{2}$. Each bucket contribute only $O(\sqrt{\Delta})$, excluding connectors. Also, there are only constant number of connectors in these buckets. Thus, we have the theorem.

4 A hierarchical construction

The GHUB network has two layers: hubs and others. LHUB network has three layers: connectors, hubs in buckets, and others. One may feel that we may have a better structure if we increase the number of layers. If we measure the worst-case interference by using the input size n or Δ, it is not possible to improve the worst case interference, since there is a lower bound of $\Omega(\sqrt{\Delta})$ even in the highway (i.e., one-dimensional) model. However, this can be advantageous in practice as we see if we measure the interference using a different parameter.

Let d be the minimum distance between two points in V. Below, we will give a network whose interference ratio is $O(\log(R_{min}/d))$, where R_{min} is the radius to give the uniform-radius network. As before, we scale the problem such that $R_{min} = 1$.

The same localization method works, and we can assume that all points are located in a unit square. Our approach is based on quad-tree decomposition. We adopt the convention that each square in the quad-tree decomposition includes its lower edge and its right edge, together with its lower two corner vertices.

We continue the following process from $k = 0$, where $U(S) = V$ if $k = 0$:

Quad-tree decomposition process Given a square S of size $2^{-k} \times 2^{-k}$ and a set $U(S) \subset V \cap S$ do the following.

1. If $U(S) = \emptyset$, terminate the process.
2. Otherwise, select a representative point $\mathbf{p}(S) \in V(S)$ arbitrarily, and remove $\mathbf{p}(S)$ from $U(S)$.
3. Partition S into four quadrants of size $2^{-(k+1)} \times 2^{-(k+1)}$. The point set $U(S)$ is partitioned accordingly. The at most four non-empty quadrants obtained are called *children* of S.
4. Apply the process iteratively to each child.

We call S' the parent of S if S is one of the children of S', and denote $S' = parent(S)$. We also say that $\mathbf{p}(S)$ is a child (resp. parent) of $\mathbf{p}(S')$ if S is a child (resp. parent) of S'. For the representative point $\mathbf{p}(S)$ of S, we set $r(\mathbf{p}(S)) = \max\{diag(S), d(\mathbf{p}(S), \mathbf{p}(parent(S))\}$, where $diag(S)$ is the length of the diagonal of the square S. Thus, we have assigned a radius to each point of V, and have a network QUAD(V).

Theorem 6. QUAD(V) *is connected, and its interference is* $O(\log d^{-1})$, *where* d *is the minimum distance between points of* V.

Proof. Since $r(\mathbf{p}(S)) \geq diag(S)$, the disk of $\mathbf{p}(S)$ contains all its children. Also, $r(\mathbf{p}(S)) \geq d(\mathbf{p}(S), \mathbf{p}(parent(S)))$ means that the disk also contains its parent. Thus, the points are connected via the tree structure of the parent-child relation.

Now, let us analyze the interference at a point $\mathbf{p}$. There are at most $O(\log d^{-1})$ different sizes of squares in the quad tree decompositions, since the diagonal length of the parent square of a smallest square must be at least d (otherwise, it can contain only one point). Consider a bucket size 2^{-k}, and analyze how many representative points of such buckets can interfere with $\mathbf{p}$. The radius $r(\mathbf{p}(S))$ of a representative point of a square S of this size is at most $2^{-k+1}\sqrt{2}$, since the distance from the representative point to any point in the parent square is at most $diag(parent(S)) = 2^{-k+1}\sqrt{2}$. Thus, S can interfere with $\mathbf{p}$ only if S intersects with the circle of radius $2^{-k+1}\sqrt{2}$ about $\mathbf{p}$. It is easy to see that there are only a constant number of such squares of this size. Thus, the interference at $\mathbf{p}$ is $O(\log d^{-1})$.

Note that in a practical implementation, we should apply a routine to shrink each disk as much as possible while keeping the connection to its parent and children.

5 Construction of a small-size ϵ-net

Here, we give an outline of a proof of Theorem 3. It suffices to show the following:

Theorem 7. *There exists an ϵ-net of V of size $O(1/\epsilon)$ for $\mathcal{P}_1$.*

We follow the argument of [5] with a (minor) modification. For simplicity, we assume that no two points of V lie on a horizontal line, a vertical line, or a line with the argument angle $\pi/3$. We call a member of $\mathcal{P}_1$ a *fan* in this section. For a fan P, we define $Int(P)$ and $cl(P)$ to be its interior and closure. Let $\partial(P) = cl(P) \setminus Int(P)$ be the boundary of its closure.

Given a subset $S \subset V$, a pair $(\mathbf{p}, \mathbf{p}')$ of points in S is extremal in S if, for any number $N > 0$, there is a fan P such that the area of P is larger than N, $Int(P) \cap S = \emptyset$ and $\{\mathbf{p}, \mathbf{p}'\} \in \partial(P)$.

We add a set X of three "extra" points $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3$ to V. Let ℓ_1 be a horizontal line that contains V in its lower halfplane. Let ℓ_2 be a line of argument angle $\pi/3$ that contains V in its upper halfplane. The points $\mathbf{q}_1$ and $\mathbf{q}_2$ are on the line ℓ_1, and the x-coordinate value of $\mathbf{q}_1$ (resp. $\mathbf{q}_2$) is sufficiently small (resp. large). The point $\mathbf{q}_3$ is on the line ℓ_2 and its y-coordinate value is sufficiently small. We can take these three points sufficiently far from V such that X satisfies the following conditions:

1. The triangle spanned by X contains all points of V.
2. For any fan P, we have another fan $P' \subseteq P$ such that $P' \cap V = P \cap V$, and $P' \cap X = \emptyset$.
3. For any pair of points in X, there is a fan P containing them on the boundary and containing no other points of V in it.
4. For any extremal pair $(\mathbf{p}, \mathbf{p}')$ of a subset S of V, we have a fan P with the largest size such that $Int(P) \cap (S \cup X) = \emptyset$ and $\{\mathbf{p}, \mathbf{p}'\} \in \partial P$. Note that one or more points of X lie on the boundary of P, and intutitively, X prevents $(\mathbf{p}, \mathbf{p}')$ to be an extremal pair in $S \cup X$.

Now, we fix $S \in V$ and consider $\tilde{S} = S \cup X$. We say a fan P an *empty fan* if it contains no point of $\tilde{S}$ in its interior. A pair of points $(\mathbf{p}, \mathbf{p}')$ is called a Voronoi pair if there exists an empty fan P containing $\mathbf{p}$ and $\mathbf{p}'$ in its boundary. Let $DT(S)$ be the graph whose node set is $\tilde{S}$ and edge set E is the set of all Voronoi pairs.

Lemma 5. *$DT(S)$ is connected, and gives a triangulation with the vertex set $\tilde{S}$ in the triangle spanned by X.*

Proof. It is easy to show that no pair of edges intersect each other using the fact that two fans intersect each other such that boundary curves intersect at most twice. Given an emtpy fan P containing a Voronoi pair $(\mathbf{p}, \mathbf{p}')$ on its boundary, we can grow P keeping the Voronoi pair on the boundary until we have another point $\mathbf{p}''$ in $\tilde{S}$ on its boundary. Then, we have a triangle $\mathbf{p}, \mathbf{p}', \mathbf{p}''$ in $DT(S)$ consisting of three edges. It can be shown (by case study) that if the Voronoi pair does not contain a point in X, we have exactly one such triangle in each side of the edge. Thus, we can show both the connectivity and triangulation property.

$DT(S)$ is called the generalized Delauney triangulation of S. For each triangle in $DT(S)$, the unique fan P containing three vertices of triangles on its boundary is called the *Voronoi fan* to the triangle. Note that a Voronoi fan contains no point of S in its interior.

The construction of [5] is as follows: Let $\delta = \epsilon/4$. We greedily find a maximal family of disjoint subsets $\{S_1, S_2, \ldots, S_k\}$ of V such that $|S_i| = \delta n$ and there exists a fan P_i such that $P_i \cap V = S_i$.

Let $S = \cup_{i=1}^k S_i$, and we make $DT(S)$. By definition, any fan P containing δn or more points of V must contain a point of S. Thus, for each triangle in $DT(S)$, there are at most δn points of V in the Voronoi fan. Moreover, the subgraph of $DT(S)$ induced by S_i is connected, and each Voronoi fan corresponding to a triangle in the induced subgraph contains no point of V in its interior.

We use $k+3$ colors to give a mutually different color to each set S_i and also each of three point of X. We give corresponding colors to vertices of $DT(S)$. For two colors (c_1, c_2), a triangle is called (c_1, c_2)-colored if its vertices use exactly the two colors.

For a fixed pair (c_1, c_2) of colors, the adjacency graph of the set of all (c_1, c_2)-colored triangles has neither a branching node (i.e., a node with degree three or more), nor a cycle: This can be shown by using the fact that S_i is intersection of a fan (a convex region) and V. Thus, the set of (c_1, c_2)-colored triangles is divided into maximal connected chains of triangles called *corridors*.

Lemma 6 ([5]). *There are $O(k)$ corridors.*

The corridors are refined into sub-corridors such that each sub-corridors has at most δn points of V in its triangles. The vertex set of subcorridors C consists of two monochromatic chains (possibly degenerated to points) in $D(S)$, and thus they have at most four endpoints.

Let Z be the set of all endpoints of all subcorridors in $DT(S)$.

Theorem 8. *Z is an ϵ-net of $V \cup X$, and its size is $O(1/\epsilon)$.*

Proof. Consider any fan P containing more than ϵn points of $V \cup X$. We assume that P contains no point of Z and derive contradiction. $P \cup V$ must be colored by at least three colors, since each monochromatic set (and the colorless set) has at most $\epsilon n/4$ points. P can contain no monochromatic chain in its iterior since it does not have a point in Z. If the fan P cuts both monochromatic chains of a subcorridor, $P \cup V$ must be bichromatic (by an argument given in [5], which we omit in this paper), and contradict to the assumption. Thus, for each subcorridor, P can only cut one of its mornochromatic chain. This implies that $P \cup V$ is monochromatic, and we have contradiction.

We finally show that $Z \setminus X$ is an ϵ'-net of V if $\epsilon < \epsilon' < 2\epsilon$. Indeed, suppose we have a fan P that contains $\epsilon' n$ points of V but no point in $Z \setminus X$. Thus, it must contain one or more points of X. We can shrink P such that only the points of X go outside of it. This new fan contains $\epsilon' n - 3$ points of V and contains no point in Z. Thus, this contradicts the fact that Z is an ϵ-net of $V \cup X$.

6 Concluding remarks

The theory can easily be generalized to any constant dimensional space, except that we only know a $O(\epsilon^{-1} \log^{-1} \epsilon^{-1})$ bound for ϵ-nets of the higher dimensional analogues of "the range space of pies".

Practically, we can improve the method in many ways. For example, in the construction of QUAD(V), we can stop the partitioning if $|U(S)| = 1$, and else partition $U(S)$ without selecting a representative point until there are at least two empty buckets. Also, we can mix the two methods: In each square S, we can replace the structure of QUAD(S) network within S by LHUB(S), if it gives a better interference.

There are several open problems: We can easily observe that an $\Omega(\sqrt{\log(R_{min}/d)})$ lower bound is attained by the "exponential chain instance" in the highway model. We conjecture that this lower bound is tight, although we currently only have an $O(\log(R_{min}/d)$ upper bound given in this paper. Moreover. for the highway model, the better one of linear network and hub network attains $O(\Delta^{1/4})$ approximation ratio to the optimal network. For the two-dimensional case, an analogous result has not been obtained yet.

References

1. P. K. Agarwal, Range Searching, Chapter 31 of *Handbook of Discrete and Computational Geometry* (ed. J. E. Goodman and J. O'Rourke), CRC Press 1997, pp. 575-598.
2. M. Burkhart, P. von Rickenbach, R. Wattenhofer, A. Zollinger: Does Topology Control Reduce Interference? Proc. MobiHoc 2004, 9-19 (2004).
3. D. Haussler and E. Welzl, ϵ-Nets and Simplex Range Queries, *Discrete & Compu. Geom.* **2** (1987), 237-256.
4. J. Matoušek, *Geometric Discrepancy, An Illustrated Guide*, Springer Verlag, 1999.
5. J. Matoušek, R. Seidel, E. Welzl, How to Net a Lot with Little: Small ϵ-Nets for Disks and Halfspaces, *Proc. 6th Symp. on Computational Geometry,*(1990), 16-22.
6. F. P. Preparata and M. I. Shamos, *Computational Geometry, an Introduction,* Texts and Monographs in Computer Science, Springer Verlag, 1988 (2nd ed).
7. P. von Rickenbach, S. Schmid, R. Wattenhofer, A. Zollinger: A Robust Interference Model for Wireless Ad-Hoc Networks, Proc. IPDPS 2005 (2005).

8. A. C.-C. Yao, On Constructing Minimum Spanning Trees in k-Dimensional Spaces and Related Problems, *SIAM J. Comput.*, **11** (1982), 721-736.