Efficient ϵ-net construction for an application to ad-hoc network

Takeshi Tokuyama

Abstract

The interference minimization in wireless ad-hoc networks is formulated as assigning a suitable transmission radius to each of the given points in the plane, so as to minimize the maximum number of transmission ranges overlapping any point. By using ideas from computational geometry and ϵ-net theory, we can attain an $O(\sqrt{\Delta})$ bound for the maximum interference where Δ is the interference of a uniform-radius ad-hoc network. In this talk, we focus on the construction algorithm of the ϵ-net required in the ad-hoc network application. This is an abstract of a conference presentation based on a joint work with Magnus Halldorsson.

Index Terms

Ad-Hoc Network, Computational Geometry, Algorithm.

I. Introduction

Wireless mobile ad-hoc networks (MANET) have received widespread interests as modern communication mod ls. In a common formulation, each mobile device is considered a point (or *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.

We assume that the transmission radius of nodes is a controllable parameter and a monotone function of the electric power given to the node. 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 collisions. 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 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 [6] gave nearly tight approximation algorithms that bound the *average* interference of nodes.

The recent work of Rickenbach *et al.* [8] 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 $O(\sqrt{n})$ interference, while it is shown that there exists an instance that requires $\Omega(\sqrt{n})$ interference.

For the two-dimensional problem, analogous results have been given by Halldorsson and Tokuyama [**?**], where construction of a network with an $O(\sqrt{n})$ interference is given for any point set in the plane, extending the theory of [8] to the planar case (and even for any constant-dimensional space). The method rely on ϵ-net of a family of regions that consists of $1/6$ pieces of disks (pies). In this paper, we show that we can use an ϵ-net of regular triangles instead of pies, so that the ϵ-net can be efficiently constructed. We note that we can combine the method with localization methods to obtain a practically good network.

II. 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 are to 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 center at $\mathbf{v}_i$.

We define a 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 network $G(\mathcal{D})$ is *feasible* iff it is connected.

Graduate School of Information Sciences, Tohoku University, Sendai, 980-8579 Japan. Email: tokuyama@dais.is.tohoku.ac.jp, Partially supported by the project New Horizons in Computing, Grant-in-Aid for Scientific Research on Priority Areas, MEXT Japan

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 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 that yields a feasible network with 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 appears 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 a practical solution with some theoretical quality guarantee, either an upper bound of the interference or a ratio of the approximation to the optimal solution.

A. Review for the highway model

We briefly review some results for the highway model given by Rickenbach *et al.* [8]. Suppose that points of V are located on the x-axis in the sorted order with respect to their x-values.

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*. The linear network has interference at most Δ and works well on typical practical instances, for example, on a randomly distributed point set. Unfortunately, there is an instance for which the linear network poorly performs. In the *exponential chain* forming this instance, 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 point in V as a hub, we have a set W of cardinality $\sqrt{n}$, and can show that $I(G(\mathcal{D})) = O(\sqrt{n})$ for this network. It has been shown that the minimum interference is $\Omega(\sqrt{n})$ for the exponential chain, thus the hub-connected network is worst-case optimal.

1) Hub selection using an ϵ-net.: For the two-dimensional problem, 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 (which the reader need not be familiar with) has numerous applications in 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 the VC dimension is low (say, a constant), we can always obtain a small ϵ-net (see [4] for example).

Here, we consider a range space associated with a family of regular triangles. Consider a regular triangle P_1 spanned by $(0,0)$, $(1,0)$, $(1/2, \sqrt{3}/2)$. Let P_2 be the reflected image of P_1 with respect to the x-axis. The family $\mathcal{P}_1$ (resp. $\mathcal{P}_2$) is the set of all translated and scaled copies of P_1 (resp. P_2). Let $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2$

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. In particular, this gives an easy local (fully distributed) algorithm.

Theorem II.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 $\mathcal{P}_k$ $(k = 1, 2)$, 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 $\mathcal{P}_k$. This follows from the general theory of ϵ-nets [3], [4] of range spaces. ■

Consider the family $\mathcal{P}_k$ for $k = 1, 2$, 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

[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.

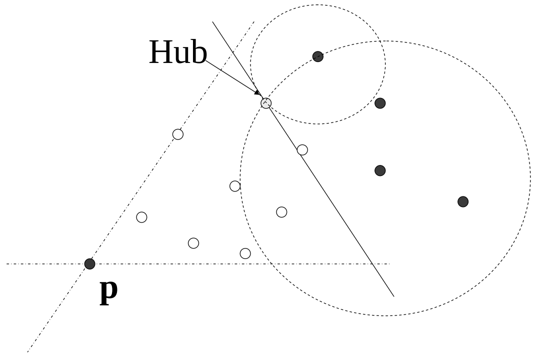

Fig. 1. No disk around a point outside the region $P(\mathbf{w})$ can reach $\mathbf{p}$

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 II.2. *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-expert to follow, we give (an outline of) the actual construction of such an ϵ-net later in a separate section.

2) 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 II.2 by setting $\epsilon = \sqrt{n^{-1}}$. Then, we form the network $\mathrm{WMST}(W)$ (or any connected network on W). Let $r_0(w)$ be the transmission radius of $w \in W$ in $\mathrm{WMST}(W)$.

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 network $\mathrm{GHUB}(V)$.

Lemma II.3. $\mathrm{GHUB}(V)$ *is connected.*

Proof: Since $\mathrm{WMST}(W)$ is connected, the subgraph of $\mathrm{GHUB}(V)$ induced by W is connected. Since other nodes are all connected to nodes in W, $\mathrm{GHUB}(V)$ is connected. ■

Theorem II.4. *The interference of* $\mathrm{GHUB}(V)$ *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 Fig. 1). Consider a hub $\mathbf{w} \neq \mathbf{p}$ in $R_1(\mathbf{p})$. We draw a line of argument angle $2\pi/3$ through $\mathbf{w}$ such that it makes a regular triangle $P(\mathbf{w}) \in \mathcal{P}_1$ together with the two boundary lines of $R_1(\mathbf{p})$. We select the hub $\mathbf{w}$ such that $P(\mathbf{w})$ is minimized. Then, $P(\mathbf{w})$ does not contain a hub in its interior, and hence $P(\mathbf{w})$ can contain at most $\sqrt{n}$ elements of V. Consider any point $\mathbf{x} \in V$ in $R_1(\mathbf{p}) \setminus P(\mathbf{w})$. Then, we can see that $d(\mathbf{x}, \mathbf{w}) < d(\mathbf{x}, \mathbf{p})$ analogously to Lemma **??**. 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 $\mathbf{x}$. We can deal with the other five cusps similarly. 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 $O(\sqrt{n \log n})$.

III. Construction of a small-size ϵ-net

Here, we give an outline of a proof of Theorem II.2. It suffices to show the following:

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

Although the above theorem can be generalized for a family of all translated/scaled copies of any given convex region, we only focus on $\mathcal{P}_1$ (i.e., the region of translated and scaled copies of a given regular triangle) in this paper to avoid unnecessary abstraction.

We follow the argument of [5] with a minor modification. Although there is little originality of the arguments given in this section, we would like to give the exact construction, since the general theory given in [5] seems to be difficult to follow for a non-expert in computational geometry. We remark that the published conference version of [5] has an error in its proof, and a corrected proof is in an unpublished manuscript[2].

For simplicity, we assume the non-degereracy condition that no two points of V lie on a horizontal line, a vertical line, or a line with the argument angle $\pi/3$. We often call a member of $\mathcal{P}_1$ a *range*, since we would like to use the term "triangle" for general triangles later, and also because the theory holds for more general range spaces. For a range $P \in \mathcal{P}_1$, we define $Int(P)$ and $cl(P)$ to be its interior and closure. The boundary of P is $\partial(P) = cl(P) \setminus Int(P)$.

The following lemma holds in a more general setting where P is a convex body and P' is its scaled and translated copy. It is easy to prove it for our ranges (i.e., isothetic regular triangles) by case study.

Lemma III.2. *For any pair P and P' of ranges, $P \setminus cl(P')$ is connected, and $\partial(P) \cap \partial(P')$ has at most two connected components.*

Corollary III.3. *Under the non-degeneracy condition , given any set A of three points of V, there is at most one range containing A on its boundary.*

Given a point set S, we call a range P an *empty range* (with respect to S) if it contains no point of S in its interior. A pair of points $(\mathbf{p}, \mathbf{p}')$ of S is called a *Delauney pair* if there exists an empty range P containing $\mathbf{p}$ and $\mathbf{p}'$ on its boundary. A Delauney pair is called *extremal* if, for any number $N > 0$, there is an empty range P containing them on its boundary such that the area of P is larger than N.

Let $DT(S)$ be the graph consisting of S (as the vertex set) and the set of all Delauney pairs (as the edge set). We draw each edge as the straight line segment between vertices.

Lemma III.4. *No two edges of E intersect each other at an interior point.*

Proof: Let e and f be edges intersecting at an interior point. Let P and P' be empty ranges containing $e = (\mathbf{p}, \mathbf{p}')$ and $f = (\mathbf{q}, \mathbf{q}')$, respectively. Because S is nondegenerate, we can shrink P (resp. P') if necessary such that they contain no point in $S \setminus \{\mathbf{p}, \mathbf{p}'\}$ (resp. $S \setminus \{\mathbf{p}, \mathbf{p}'\}$) in its closure. By the definition of empty ranges, $\mathbf{q}$ and $\mathbf{q}'$ (resp. $\mathbf{p}$ and $\mathbf{p}'$) are outside the interior of P (resp. P'). Let the edge e intersect $\partial(P \cap P')$ at points v_1 and v_2 and f intersect $\partial(P \cap P')$ at w_1 and w_2. If e and f intersect in the interior, these four points appears in a clockwise alternating order along $\partial(P \cap P')$, e.g. as v_1, w_2, v_2, w_2, since $P\mathcal{P}'$ is convex. Thus, $(P \cup P') \setminus Int(P \cap P')$ has four connected components. However, Lemma III.2 implies that $(P \cup P') \setminus Int(P \cap P')$ has at most two connected components. We have a contradiction. ■

Thus, $DT(S)$ gives a planar graph drawing. Indeed, it is the dual of the *generalized Voronoi diagram* [9]. We would like to claim that $DT(S)$ is a triangulation of the point set for a subset S of V. This is known to hold if we consider *pseudo disks* (assuming a suitable non-degeneracy condition) [5]. Here, a family of regions is called pseudo disks if for any three non-collinear points, there exists a unique region containing them on its boundary. Unfortunately, $\mathcal{P}_1$ does not satisfy the condition of pseudo disks, and $DT(S)$ is not always a triangulation. Indeed, the set S of three black points in Fig. 2 does not have a range containing it on the boundary, and $DT(S)$ has only two edges, thus is not a triangulation.

We slightly modify the set V and S to resolve the above problem. We add a set X of three "extra" points $\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3$ to V. Let ℓ be a horizontal line that contains V in its lower halfplane. The points $\mathbf{q}_1$ and $\mathbf{q}_2$ are on the line ℓ, the point $\mathbf{q}_3$ lies below ℓ, and the three form vertices of a regular triangle (actually, one in $\mathcal{P}_2$). We take these three points sufficiently far from V so that X satisfies the following conditions:

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

It is routine to show that we can find such an X. We write $\tilde{S}$ for $S \cup X$ for a subset S of V, and consider $DT(\tilde{S})$ instead of $DT(S)$. For the point set of Fig. 2, we obtain a triangulation by adding X (the three white points) as shown in Fig. 3.

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

Proof: Given an empty range P that has a Delauney pair $(\mathbf{p}, \mathbf{p}')$ on its boundary, we can grow P keeping the Delauney 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 $DT(\tilde{S})$ consisting of three edges. If the Delauney pair does not contain a point in X, we have one such triangle in each side of the edge. Indeed, we slide and scale the empty range towards two directions keeping the Delauney pair on the boundary. Since there is no Delauney pair parallel to boundary edges of the range, the size of the range grows to infinite if there is no triangle in $DT(\tilde{S})$ on one side. This contradicts the definition of X. Because of the non-crossing property of edges (Lemma III.4), there is exactly one such triangle on each side. Thus, each interior face of the planar graph $DT(\tilde{S})$ must be a triangle. Thus, $DT(\tilde{S})$ is a triangulation of the triangle spanned by X. ■

[2]This information, together with the address of the web page containing it, was given to the authors by J. Matoušek

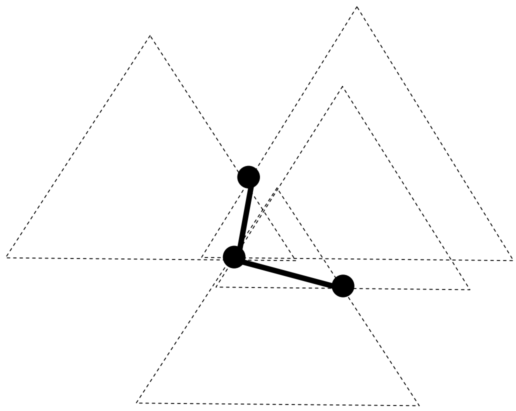

Fig. 2. A set of three points not contained on the boundary of any range.

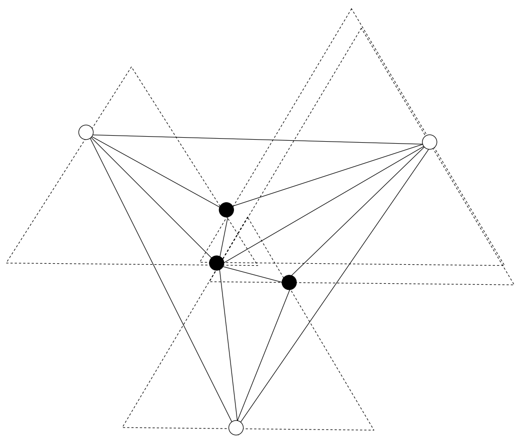

Fig. 3. Triangulation is obtained by adding the point set X.

We call $DT(\tilde{S})$ the generalized Delauney triangulation of S. For each triangle in $DT(\tilde{S})$, the unique range P containing three vertices of triangles on its boundary is called its *Voronoi range* [3]. Note that a Voronoi range contains no point of $\tilde{S}$ in its interior.

The construction of [5] is as follows: Let $\delta = \epsilon/5$. 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 range P_i such that $P_i \cap V = S_i$. Fig. 4 shows such a family where $\delta n = 5$.

Let $S = \cup_{i=1}^{k} S_i$, and we make $DT(\tilde{S})$. Any range P containing δn or more points of V must contain a point of S, since otherwise we can continue the greedy process to have a new subset S_{k+1} to contradict with the maximality of our family of subsets.

Thus, for each triangle in $DT(\tilde{S})$, there are at most δn points of V in its Voronoi range. Moreover, the subgraph D_i of $DT(\tilde{S})$ induced by S_i is connected: Otherwise, we can show that there is an empty range corresponding to an Delauney edge connecting two points in $S \setminus S_i$, and the intersection of the empty range and P_i violates Lemma III.2. Each triangle in D_i contains no point of V in its interior, and moreover, no cycle in D_i contains a point in $V \setminus S_i$ in its interior region (this comes from the fact that the range P_i defining S_i contains the convex hull of S_i).

We use $k+3$ colors to give a mutually different color to each set S_i and also to each of three point of X. The points in $V \setminus S$ are colorless. We give corresponding colors to vertices of $DT(\tilde{S})$. For two colors (c_1, c_2), a triangle is called (c_1, c_2)-colored if its vertices use exactly those two colors.

[3]This is analogue of Voronoi circle for an ordinary Voronoi diagram.

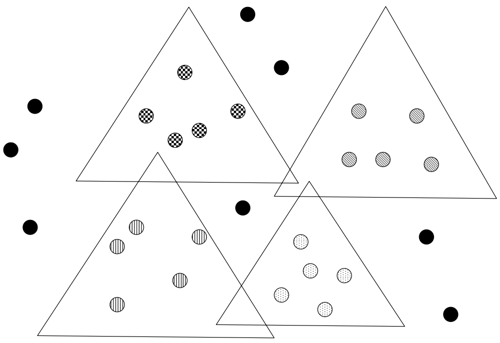

Fig. 4. Greedy procedure to find maximal family of disjoint subsets of size δn.

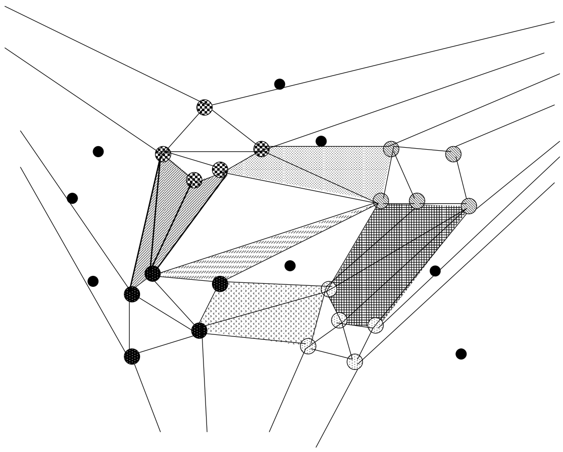

Fig. 5. Corridors in $DT(\tilde{S})$.

For a fixed pair (c_1, c_2) of colors, we divide the set of (c_1, c_2)-colored triangles into maximal connected chains of triangles such that each pair of consecutive triangles share a bicolored edge. Such a maximal chain is called a *corridor*.

Lemma III.6. *There are $O(k)$ corridors.*

Proof: Since any cycle in D_i only contains points of S_i, for each $i = 1, 2, \ldots, k$, we can contract each S_i of $DT(\tilde{S})$ into a point such that all bicolored edges in each corridor (say, corresponding colors of S_i and S_j) are replaced by an edge between S_i and S_j. This graph has k vertices, and each face has exactly three sides. Although this graph may have multiple edges, it has at most $3f/2$ edges (f is the number of faces), and the number of edges is $O(k)$. Thus, the number of corridors is $O(k)$. ■

The corridors are greedily refined into subcorridors containing at most δn points of V (indeed, they are colorless) in its triangles. Because the dual graph of each corridor is a tree of degree at most three, we can decompose it into $O(\frac{M}{\delta n})$ subcorridors if the corridor has M 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 $DT(\tilde{S})$.

Theorem III.7. *Z is an ϵ-net of $V \cup X$, and its size is $O(1/\epsilon)$.*

Proof: The number of sub-corridors is $O(k + \frac{1}{\delta}) = O(1/\epsilon)$. Thus $|Z| = O(1/\epsilon)$. Consider any range P containing more than ϵn points of $V \cup X$. We assume that P contains no point of Z and derive contradiction. Without loss of generality, we assume that P contains at least one point in S_1 (colored red, striped in Fig. 6). $P \setminus P_1$ has at most one connected component because of Lemma III.2, and let Y be the set of points in the component. If $Y = \emptyset$, P only contains red points, thus it can

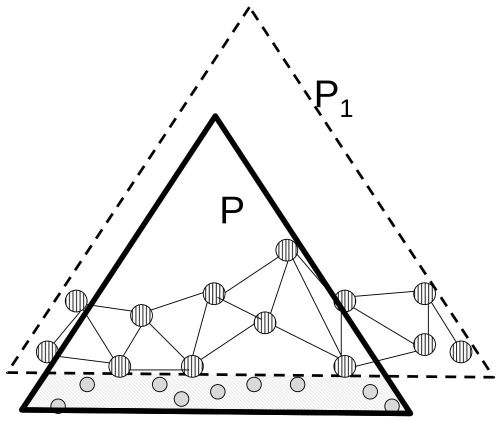

Fig. 6. Intersection of P and P_1, and the set Y (points in the shaded region).

have at most δn points; this is a contradiction. Thus, Y is nonempty.

Let C be the red monochromatic red chain in the union of corridors. Since $P \setminus P_1$ is connected, there is a unique connected component C_1 of $C \cap P$ such that the other side has at least one point of (thus, all points of) Y. C_1 must be a subchain of a red chain C_{red} of a subcorridor, since P contains no point of Z. Let C_{blue} be the partner chain of the subcorridor, colored blue, the color of the set S_2. If C_{blue} intersects P (including the case that P contains no blue points but only intersect edges), there is no non-blue point below C_{blue} (i.e., different side from C_{red}), since otherwise $P \setminus P_2$ must have two connected components, contradicting Lemma III.2. Let $e = (p_{red}, p_{blue})$ and $f = (q_{red}, q_{blue})$ be bi-colored edges at the two ends of the subcorridor. The subcorridor is bounded by C_{red}, C_{blue}, e and f as shown in Fig. 7.

If e intersects P, it must cut P into two pieces, since none of the endpoints of e is in P. Let R_e be the piece in the different side from the subcorridor. Similarly, we define R_f.

Let Q be the Voronoi range corresponding to the triangle containing e on the boundary (one of shaded triangles in Fig. 7) that is not in the corridor. If $R_e \setminus Q \neq emptyset$, then $P \setminus Q$ has two connected components, one on each side of the edge e, contradicting Lemma III.2. Thus, $R_e \in Q$, and thus R_e has at most δn points in it. R_f also has at most δn points.

The set of points in P consists of the five parts: The part above or on C_{red} only has red points, thus at most δn points. The part below or on C_{blue} only has blue points, thus at most δn points. Each of R_e and R_f has at most δn (colorless) points. Finally, the subcorridor has at most δn (colorless) points Thus, P has at most $5\delta n = \epsilon n$ points. ■

We finally show that $Z \setminus X$ is an ϵ'-net of V if $\epsilon n < \epsilon' n - 3$. Consequently, we have an ϵ-net of size $O(1/\epsilon)$ for the range space $(\mathcal{P}_1, V)$. Indeed, suppose we have a range 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 range 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 $V \cup X$.

It is almost trivial that Z can be computed in polynomial time. Indeed, it is known that a generalized Voronoi diagram can be constructed in $O(n \log n)$ time [9], and the most expensive process in the construction is to find the subsets $S_1, S_2, \ldots, S_k$ greedily.

Indeed, we have the following.

Theorem III.8. *The ϵ-net Z can be computed in $O(n^2 \log n)$ time if $\epsilon = \sqrt{n}$*

Proof: The number of regular triangles in $\mathcal{P}_1$ containing exactly $k = \delta n$ points is $O(kn)$, because of an analogous argument of the complexity of k-th Voronoi diagram. Since $k = O(\sqrt{n})$, there are $O(n^{1.5})$ such triangles, and let Y be the family of these triangles. For each $p \in V$, we can compute all the triangles in Y that contains p as a boundary point in $O(n \log n)$ time by using a sweep algorithm. Thus, Y can be constructed in $O(n^2 \log n)$ time. Next, we greedily select triangles in Y one by one, deleting all triangles in Y intersected by the selected triangle. Since we select at most $\delta^{-1} = \sqrt{n}/5$ triangles, this takes $O(\sqrt{n} n^{1.5}) = O(n^2)$ time. ■

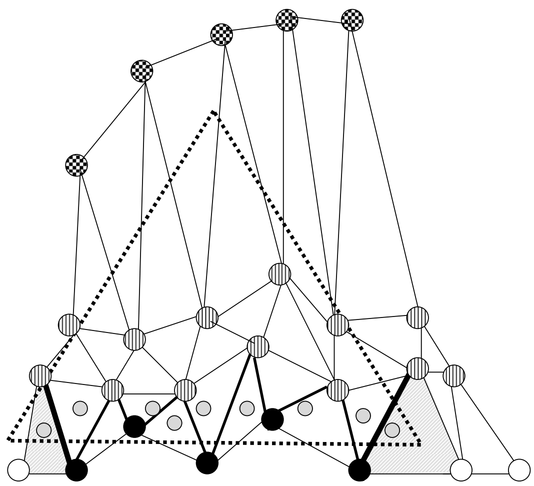

Fig. 7. A subcorridor intersecting P.

IV. 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 regular simplices".

Practically, we can improve the method in many ways. For example, in the construction of $\mathrm{QUAD}(V)$, 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 the $\mathrm{QUAD}(S)$ network within S by $\mathrm{LHUB}(S)$, if it gives a better interference.

There are several open problems: One may observe that the "exponential chain" instance attains a $\Omega(\sqrt{\log(R_{min}/d)})$ lower bound in the highway model. We conjecture that this lower bound is tight, although we currently only have the $O(\log(R_{min}/d))$ upper bound given in this paper. Moreover. for the highway model, the better one of a linear network and a 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. A corrected version is in J. Matoušek's web page, http://kam.mff.cuni.cz/~matousek/oldpaps.html.

[6] T. Moscibroda, R. Wattenhofer, Minimizing Interference in Ad Hoc and Sensor Networks, *Proc. DIALM-POMC Joint Workshop* 2005, 24–33.

[7] F. P. Preparata and M. I. Shamos, *Computational Geometry, an Introduction*, Texts and Monographs in Computer Science, Springer Verlag, 1988 (2nd ed).

[8] P. von Rickenbach, S. Schmid, R. Wattenhofer, A. Zollinger, A Robust Interference Model for Wireless Ad-Hoc Networks, Proc. IPDPS 2005 (2005).

[9] M. Sharir, Algorithmic Motion Planning, Chapter 40 of *Handbook of Discrete and Computational Geometry* ed. J. E. Goodman, J. O'Rourke, CRC Press, 1997, pp. 733–754.

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