Interference Minimization in k-Connected Wireless Networks*

Stephane Durocher† Sahar Mehrpour‡

Abstract

Given a set of positions for wireless nodes, the k-*connected interference minimization* problem seeks to assign a transmission radius to each node such that the resulting network is k-connected and the maximum interference is minimized. We show there exist sets of n points on the line for which any k-connected network has maximum interference $\Omega(\sqrt{kn})$. We present polynomial-time algorithms that assign transmission radii to any given set of n nodes to produce a k-connected network with maximum interference $O(\sqrt{kn})$ in one dimension and $O(\min\{k\sqrt{n}, k\log\lambda\})$ in two dimensions, where λ denotes the ratio of the longest to shortest distances between any pair of nodes.

1 Introduction

1.1 Interference Minimization and k-Connectivity

A network must be connected if a multi-hop communication channel is required between every pair of nodes. Various secondary objectives can be considered in addition to the connectivity requirement, often resulting in an optimization problem to construct a network that meets both criteria. Common additional objectives include minimizing the maximum or average power consumption, sender-receiver route length, node degree, ratio of route length to Euclidean distance, and, of particular relevance to wireless networks, interference [11]. By increasing or decreasing its transmission power, a wireless node increases or decreases its transmission range. If wireless signal strength is assumed to fade uniformly in all directions, then the range within which transmission exceeds a given minimum threshold corresponds to a disk centred at the point of transmission; we refer to the disk's radius as the transmitting node's *transmission radius*. Under the *receiver-based interference model* [16], two nodes p_1 and p_2 can communicate if they lie mutually in each other's transmission ranges, and any node q_1 that lies in the transmission range of a node q_2 receives interference from q_2, regardless of whether q_1 can communicate with q_2. Given a set of node positions as input, the objective of the *interference minimization problem* is to assign a transmission radius to each node to produce a connected network that minimizes the maximum interference among all nodes. The interference minimization problem has been examined extensively under the receiver-based interference model over the past decade (e.g., [2, 3, 5, 7, 11, 13, 16–18]).

Maintaining network connectivity is critical to preserving multi-hop communication channels between all pairs of nodes. Connectivity alone is insufficient to preserve communication in case of node failure: a connected network can become disconnected when even a single node fails. Guaranteeing network connectivity in the presence of node failure requires multiple disjoint routes joining every pair of nodes, i.e., redundancy in the network's connectivity. A network is k-*connected* if it remains connected whenever fewer than k nodes are removed. The factor k parameterizes the network's degree of connectivity. In this work, we examine interference minimization on k-connected networks. Given a set of node positions, the k-*connected interference minimization problem* is to assign a transmission radius to each node to produce a k-connected network while minimizing the maximum interference at any node. To the authors' knowledge, this is the first work to examine interference minimization in k-connected networks.

1.2 Definitions

We represent the position of a wireless node by a point $p_i \in \mathbb{R}^d$. The set $P = \{p_1, \ldots, p_n\} \subseteq \mathbb{R}^d$ represents positions for a set of n nodes, along with a corresponding function, $r : P \to \mathbb{R}^+$, that associates a positive real transmission radius with each node. Communication in a wireless network is often modelled by a symmetric disk graph (SDG); the *symmetric disk graph* of P with respect to r is an undirected graph with vertex set P and edge set $\{(p, q) \mid \{p, q\} \subseteq P \wedge r(p) \geq \mathrm{dist}(p, q) \wedge r(q) \geq \mathrm{dist}(p, q)\}$, where $\mathrm{dist}(u, v)$ denotes the Euclidean distance between the points u and v in $\mathbb{R}^d$ [1]. In this paper we focus on point sets in one or two dimensions ($d \in \{1, 2\}$).

von Rickenbach et al. [16] introduced the receiver-centric interference model. In this model, the *interference* at the node $p \in P$, denoted $I(p)$, is the number of nodes in P whose transmission range covers p. That is, $I(p) = |\{q \mid q \in P \wedge \mathrm{dist}(p, q) \leq r(q)\}|$. The *maximum interference* for the set of points P with transmission radii given by r is the maximum $I(p)$ over all

*This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC).

†University of Manitoba, Canada, durocher@cs.umanitoba.ca
‡University of Utah, USA, mehrpour@cs.utah.edu

$p \in P$. For a given graph G on the point set P, let $I(G) = \max_{p \in P} I(p)$. The *interference minimization problem* is to assign transmission radii (i.e., to define the function r) for a given set of points $P \subseteq \mathbb{R}^d$ such that the corresponding symmetric disk graph G is connected and $I(G)$ is minimized.

A graph G is *connected* if there is a path (a sequence of adjacent vertices) joining every pair of vertices in G. A graph G is *k-connected* if there are k disjoint paths between every pair of vertices in G or, equivalently, if the removal of any j vertices does not disconnect G, for all $j < k$. The *k-connected interference minimization problem* is to assign transmission radii (i.e., define the function r) for a given set of points $P \subseteq \mathbb{R}^d$ such that the corresponding symmetric disk graph G is k-connected and $I(G)$ is minimized. Let $\mathrm{OPT}_k(P)$ denote the minimum maximum interference among all k-connected networks on P.

Given a set $P \subseteq \mathbb{R}^d$, let $\mathrm{MST}(P)$ denote its Euclidean minimum spanning tree, $\mathrm{DT}(P)$ its generalized Delaunay triangulation, and $\lambda = d_{\max}/d_{\min}$ the ratio of the maximum and minimum distances between any two points in P, i.e., $d_{\max} = \max_{\{p,q\} \subseteq P} \mathrm{dist}(p,q)$ and $d_{\min} = \min_{\{p,q\} \subseteq P} \mathrm{dist}(p,q)$. A set $P = \{p_1, \ldots, p_n\}$ of n points in $\mathbb{R}$ ordered such that $p_i < p_j$ for all $i < j$ contains an *exponential chain* of size m if there exist m integers $1 \le a_1 < a_2 < \cdots < a_m \le n$ (or $n \ge a_1 > a_2 > \cdots > a_m \ge 1$) such that $\mathrm{dist}(p_{a_i}, p_{a_i+1}) \ge \mathrm{dist}(p_{a_1}, p_{a_i})$ for all $i \in \{1, \ldots, m\}$. That is, the transmission range of p_{a_i} in $\mathrm{MST}(P)$ covers $\{p_{a_{i+1}}, \ldots, p_{a_m}\}$. For example, the set $\{2^i \mid i \in \{0, \ldots, m\}\}$ forms an exponential chain of size m. See Figure 1. Given a set $P \subseteq \mathbb{R}$ of n node positions and an assignment of transmission radii corresponding to the symmetric disk graph G on P, von Rickenbach et al. [16] define a *hub* node as any vertex of G that has at least one neighbour to its right; a non-hub node in P has all of its neighbours to its left. For networks in $\mathbb{R}^2$, a subset $H \subseteq P$ may be identified as a set of hubs, where these hub nodes provide a connected or k-connected backbone to which non-hub nodes connect.

Recall the definition of an ϵ-net [8]. Given a set P of points in $\mathbb{R}^2$ and a family $\mathcal{R}$ of regions (*ranges*) in $\mathbb{R}^2$, the pair $(P, \mathcal{R})$ is a *range space*. For any given $\epsilon \in (0,1)$, an *ϵ-net* of the range space $(P, \mathcal{R})$ is a subset $S \subseteq P$ such that for any region $R \in \mathcal{R}$, if $|R \cap P| \ge \epsilon n$, then $R \cap S \ne \varnothing$. As do Halldórsson and Tokuyama [7], our algorithm uses the set $\mathcal{R}$ of ranges consisting of all equilateral triangles with one edge parallel to the x-axis.

1.3 Overview of Results

We begin with a discussion of related work in Section 2. In Section 3 we establish a lower bound of $\Omega(\sqrt{kn})$ on the worst-case maximum interference among all k-connected networks on a given set of n points in $\mathbb{R}$. This bound applies to point sets in $\mathbb{R}^d$ for any $d \ge 1$ and any $1 \le k < n$, and improves on the lower bounds of $\Omega(k)$ due to k-connectivity and $\Omega(\sqrt{n})$ for maximum interference in a connected network [16]. In Section 4 we generalize a technique introduced by von Rickenbach et al. [16] and apply it to give an $O(n \log(n/k))$-time algorithm that assigns transmission radii to any set of n nodes in $\mathbb{R}$ to give a k-connected network with maximum interference $O(\sqrt{kn})$ for any $1 \le k < n$, asymptotically matching our lower bound; interestingly, the dependence on k is $O(\sqrt{k})$, as opposed to being linear in k. In Section 5 we generalize techniques introduced by Halldórsson and Tokuyama [7] and apply them to develop two algorithms that assign transmission radii to any set P of n nodes in $\mathbb{R}^2$ to give k-connected networks with maximum interference $O(k \log \lambda)$ and $O(k\sqrt{n})$, respectively, in $O(n \log \lambda)$ and $O(nk + n \log n + k^3 \sqrt{n} \log n)$ time, respectively. We conclude with a discussion and directions for future research in Section 6.

2 Related Work

Buchin [3] showed that finding an optimal solution to the interference minimization problem is NP-complete in two dimensions. At present, the problem's complexity remains open in one dimension.

Several studies examine the interference minimization problem in one dimension, also known as the *highway model*. von Rickenbach et al. [16] gave an $O(n^2)$-time $O(n^{1/4})$-approximation algorithm and showed a tight asymptotic bound of $\Theta(\sqrt{n})$ on the worst-case minimum maximum interference of any set P of n points in $\mathbb{R}$. Their approximation algorithm applies one of two strategies, $\mathrm{MST}(P)$ or a hub backbone, whichever has lower interference. $\mathrm{MST}(P)$ provides low interference when P is near to being uniformly distributed. If P contains an exponential chain of size m, then $I(\mathrm{MST}(P)) \in \Omega(m)$ [16]. The hub strategy of von Rickenbach et al. [16] selects every $\sqrt{n}$th node as a hub according to their ordering on the line, forms a connected backbone network on the hubs (e.g., their MST), and connects each non-hub node to its nearest hub, giving a network with maximum interference $O(\sqrt{n})$ for any set of n points in $\mathbb{R}$. Tan et al. [18] gave an algorithm that finds an optimal solution for any set P of n points in $\mathbb{R}$ in $O(n^{3+\mathrm{OPT}_1(P)})$ time.

The interference problem has also been examined extensively in two dimensions. Halldórsson and Tokuyama [7] used ϵ-nets to define a backbone of $O(\sqrt{n})$ hub nodes, resulting in a network with maximum interference $O(\sqrt{n})$ for any set of n points in $\mathbb{R}^2$. See Section 2.1 for a detailed description. Halldórsson and Tokuyama [7] present a second algorithm using a quadtree decomposition that guarantees maximum interference $O(\log \lambda)$ for any set of points P in $\mathbb{R}^2$. As the quadtree is constructed, each non-empty square B_i of width w_i contains some set $P_i \subseteq P$. A representative point $p \in P_i$ is selected arbitrarily and its transmission radius is set to $\max\{\sqrt{2}w_i, \text{dist}(p,q)\}$, where q is the representative of the parent square to B_i. The square B_i is divided into four squares of width $w_i/2$ and $P_i \setminus \{p\}$ is partitioned accordingly. The recursion terminates when $P_i = \varnothing$. Still in $\mathbb{R}^2$, Holec [9] used linear programming to give an algorithm with maximum interference $O(\text{OPT}_1(P)^2 \log n)$. Aslanyan and Rolim [2] also proposed an algorithm that finds a connected network by applying an approximation algorithm for a variant of the minimum membership set cover problem.

In addition to the worst-case results described above, the interference minimization problem has been examined in the randomized setting. Kranakis et al. [13] proved that $\text{MST}(P)$ has maximum interference $\Theta((\log n)^{1/2})$ with high probability for any set P of n points selected uniformly at random in $[0,1]$. Khabbazian et al. [11] showed that $\text{MST}(P)$ has maximum interference $O(\log n)$ with high probability for any set P of n points selected uniformly at random in $[0,1]^d$; Devroye and Morin [5] improved these results to show that $\text{MST}(P)$ has maximum interference $\Theta((\log n)^{1/2})$ with high probability and, furthermore, that $\text{OPT}_1(P) \in O((\log n)^{1/3})$ and $\text{OPT}_1(P) \in \Omega((\log n)^{1/4})$ with high probability, showing that for nearly all point sets P, $\text{MST}(P)$ does not minimize interference.

2.1 $O(\sqrt{n})$ Interference in $\mathbb{R}^2$

We include a detailed overview of the algorithm of Halldórsson and Tokuyama [7] using ϵ-nets, which will be important to describe our algorithm presented in Section 5.2. Given a set P of n points in $\mathbb{R}^2$, the algorithm selects an ϵ-net $H \subseteq P$ of size $O(\epsilon^{-1})$ to serve as a set of hubs. Hubs are connected by $\text{MST}(H)$, and each non-hub node (the set $P \setminus H$) connects to its nearest hub in H. Each node receives interference from at most $|H| \in O(\epsilon^{-1})$ hubs and $O(n\epsilon)$ non-hub nodes. Consequently, the resulting network has maximum interference $O(\epsilon n + \epsilon^{-1})$, which corresponds to maximum interference $O(\sqrt{n})$ when $\epsilon = n^{-1/2}$.

Halldórsson and Tokuyama [7] describe the following algorithm to find an ϵ-net $H \subseteq P$ of size $O(\epsilon^{-1})$. The algorithm begins by greedily constructing a maximal family of disjoint subsets $\{P_1, \ldots, P_l\}$ such that for each i, $P_i \subseteq P$, $|P_i| = \epsilon n/5$, and there exists a range $R \in \mathcal{R}$ such that $R \cap P = P_i$. Select any range $R_0 \in \mathcal{R}$ such that $P \subseteq R_0$, and let $V(R_0)$ denote the set of three vertices on its boundary. Let $\widetilde{P} = V(R_0) \cup \bigcup_{i=1}^{l} P_i$. Two nodes $\{p,q\} \subseteq \widetilde{P}$ form a *generalized Delaunay pair* with respect to $\mathcal{R}$ if there exists a range $R \in \mathcal{R}$ such that p and q are on the boundary of R and $R \cap \widetilde{P} = \{p,q\}$. Construct $\text{DT}(\widetilde{P})$ by adding an edge between all generalized Delaunay pairs in $\widetilde{P}$. Consider a set of colours $\{c_1, \ldots, c_{l+3}\}$. For each i, assign each $p \in P_i$ the colour c_i, and colour the points in $V(R_0)$ distinctly using the three remaining colours. A *corridor* refers to a maximal chain of 2-coloured triangles in $\text{DT}(\widetilde{P})$. Each corridor is greedily partitioned into *subcorridors* such that the union of the Delaunay triangles in each subcorridor contains $\epsilon n/5$ nodes of P. The set of endpoints of subcorridors corresponds to the set H of hubs. Since each corridor contains $O(\epsilon n)$ points of P, the number of subcorridors and, therefore, $|H|$ are $O(\epsilon^{-1})$.

3 Lower Bounds

We show the following lower bound:

Theorem 1 *For every n and every k, $1 \le k \le n$, there exists a set of n points $P \subseteq \mathbb{R}$ such that every k-connected network on P has maximum interference $\Omega(\sqrt{kn})$.*

Proof. Consider the set $P = \{p \mid p = 2^i,\ i \in \{0, \ldots, n-1\}\}$ that forms an exponential chain of size n on the line. Consider any k-connected network on P. Let H denote the set of hub vertices and let S denote the set of non-hub vertices, where $|H| + |S| = n$. Since the network is k-connected, all vertices have between k and Δ neighbours, where Δ denotes the maximum vertex degree. Consequently, the first k vertices on the left of the chain are hubs and, furthermore, these k vertices form a clique. Every hub interferes with the leftmost node in the exponential chain. Therefore, the interference at the first node (and, therefore, the maximum interference) is at least $|H| - 1$. Similarly, the maximum interference is at least Δ. That is,

$$I(G) \ge \max\{|H| - 1, \Delta\}. \tag{1}$$

Let $E_{S \to H}$ denote the set of edges that join a non-hub vertex to a hub vertex. Similarly, let $E_{H \to H}$ denote the set of edges joining pairs of hubs. This gives,

$$k|S| \le |E_{S \to H}|. \tag{2}$$

Since the first k hubs form a clique, there are $\binom{k}{2}$ edges among these. So we have,

$$\binom{k}{2} \le |E_{H \to H}|. \tag{3}$$

The number of edge endpoints at a hub is bounded by

$$\begin{aligned} & |E_{S \to H}| + 2|E_{H \to H}| \le |H|\Delta \\ \Rightarrow\quad & k|S| + 2\binom{k}{2} \le |H| \cdot I(G) \text{ (by (1), (2) and (3))} \\ \Rightarrow\quad & k(n - |H|) + k(k-1) \le |H| \cdot I(G) \\ \Rightarrow\quad & k(n + k - 1) \le |H|(I(G) + k) \end{aligned}$$

$$\leq (I(G)+1)(I(G)+k) \text{ (by (1))}$$

$$\Rightarrow \quad I(G) \geq \frac{\sqrt{(4n-6)k+5k^2+1}-(k+1)}{2}. \quad (4)$$

Next we show that $I(G) \in \Omega(\sqrt{nk})$ for all $n \geq 5$. The result holds trivially for $n \in O(1)$ and, specifically, for $n < 5$. Assume

$$\begin{aligned}
& n \geq 5 \quad (5) \\
\Rightarrow \quad & 3n+k \geq 14 \text{ (since } k \geq 1) \\
\Rightarrow \quad & 3nk+k^2 \geq 14k \\
\Rightarrow \quad & 4nk+k^2 \geq 14k+3 \text{ (by (5) since } k \geq 1) \\
\Rightarrow \quad & 4nk-6k+5k^2+1 \geq 4k^2+8k+4 \\
\Rightarrow \quad & \sqrt{(4n-6)k+5k^2+1} \geq 2(k+1) \\
\Rightarrow \quad & -(k+1) \geq -\frac{\sqrt{(4n-6)k+5k^2+1}}{2} \\
\Rightarrow \quad & I(G) \geq \frac{\sqrt{(4n-6)k+5k^2+1}}{4} \text{ (by (4))} \\
& \geq \frac{\sqrt{2nk+5k^2}}{4} \text{ (by (5))} \\
& \in \Omega(\sqrt{nk}). \ \square
\end{aligned}$$

As we show in Theorem 2, the lower bound of Theorem 1 is asymptotically tight.

4 k-Connected Networks in One Dimension

In this section, we present an algorithm that constructs a k-connected network on any set P of n points in $\mathbb{R}$. Our algorithm generalizes the hub technique applied in the algorithm of von Rickenbach et al. [16] to construct a connected network with maximum interference $O(\sqrt{n})$, as discussed in Section 2.

Instead of every $\sqrt{n}$th node as in [16], we select every $\sqrt{n/(2k+1)}$th node as a hub, resulting in $\lceil\sqrt{n(2k+1)}\rceil$ hubs. Specifically, select the ith node as a hub if $i = \lfloor j\sqrt{n/(2k+1)}\rfloor$ for some $j \in \mathbb{Z}$ (where nodes are numbered $i = 0, \ldots, n-1$). Set each hub node's transmission radius to its furthest point in P (forming a clique on the hubs). Finally, set each non-hub node's transmission radius to the further of the kth hub to its left and the kth hub to its right.

Theorem 2 *Given any set P of n points in $\mathbb{R}$ and any $k < n$, transmission radii corresponding to a k-connected network on P with maximum interference $O(\sqrt{kn})$ can be found in $O(n\log(n/k))$ time.*

Proof. First we show that the network produced is k-connected.

$$\begin{aligned}
& n > k \\
\Rightarrow \quad & n > \frac{k}{2+1/k} \\
\Rightarrow \quad & \sqrt{n(2k+1)} > k. \\
\Rightarrow \quad & \left\lceil\sqrt{n(2k+1)}\right\rceil > k.
\end{aligned}$$

Therefore, there are at least k hubs. Since the hubs form a clique and each non-hub node is connected to k hubs, the network is k-connected.

Next we bound the maximum interference. Choose any point $p \in P$. The interference at p, denoted $I(p)$, is the sum of the interference it receives from hub and non-hub nodes. Hub nodes define a partition of non-hub nodes into $\lceil\sqrt{n(2k+1)}\rceil$ intervals. Suppose the hub at the left end of each interval belongs to that interval. Let I_i denote the interval that contains p, where intervals are numbered in order from the left. Let h_l and h_r denote the respective hubs at the left and right extremities of I_i. Three types of non-hub nodes interfere with p: nodes in I_i, nodes in I_j for $j < i$ that are connected to h_r, and nodes in I_j for $j > i$ that are connected to h_l. Since each non-hub node connects to its k nearest hubs, p may receive interference from non-hub nodes in k intervals on each side, or $2k$ total intervals, corresponding to at most $\lceil 2k\sqrt{n/(2k+1)}\rceil$ non-hub nodes in other intervals. In addition, p may receive interference from non-hub nodes within its own interval. Finally, p receives interference from at most $\lceil\sqrt{n(2k+1)}\rceil$ hubs. Summing these gives

$$\begin{aligned}
I(p) &\leq \left\lceil\sqrt{n(2k+1)}\right\rceil + \left\lceil 2k\sqrt{\frac{n}{2k+1}}\right\rceil + \left\lceil\sqrt{\frac{n}{2k+1}}\right\rceil \\
&< \sqrt{n(2k+1)} + (2k+1)\sqrt{\frac{n}{2k+1}} + 3 \\
&= 2\sqrt{n(2k+1)} + 3 \\
&\in O(\sqrt{kn}).
\end{aligned}$$

The hubs can be identified in $O(n\log(n/k))$ time by near-sorting P, e.g., by a partial execution of deterministic quicksort to partition P into blocks of size $\sqrt{n/(2k+1)}$ that returns the partition pivots in sorted order. The list of hubs is traversed in $O(\sqrt{n/k})$ time to assign a transmission radius to each hub, corresponding to the further of the leftmost or rightmost points in P. Non-hub nodes are examined in block sequence, in arbitrary order within a given block. Each non-hub's transmission radius is set to the maximum distance of its kth hub to the left and its kth hub to the right in $O(n)$ total time, achieved by simultaneously traversing the list of hubs and referring to the $(i-k)$th and $(i+k)$th hubs, where i denotes the block index. The total time is dominated by near-sorting, resulting in $O(n\log(n/k))$ time in the worst case. $\square$

This guaranteed $O(\sqrt{kn})$ maximum interference matches the lower bound of $\Omega(\sqrt{kn})$ established in Theorem 1, showing that our algorithm is asymptotically optimal in the worst case. Previously, we knew $I(G) \in \Omega(\sqrt{n})$ in the worst case, implied by $k = 1$ [16], and $I(G) \in \Omega(k)$, since every node in a k-connected graph has at least k neighbours. Furthermore, $I(G) \to n - 1$ as $k \to n-1$. The interesting implication of Theorem 2, however, is for values of k between these two extrema: that the worst-case maximum interference's dependence on k is sublinear for all values of k.

5 k-Connected Networks in Two Dimensions

In this section we present two algorithms that generalize techniques applied in algorithms of Halldórsson and Tokuyama [7] described in Section 2. Given a set P of n points in $\mathbb{R}^2$, our algorithms construct respective k-connected networks on P with maximum interference $O(k \log \lambda)$ and $O(k\sqrt{n})$, for any k.

5.1 Quadtree Decomposition

Theorem 3 *Given any set P of n points in $\mathbb{R}^2$ and any $k < n$, transmission radii corresponding to a k-connected network on P with maximum interference $O(k \log \lambda)$ can be found in $O(n \log \lambda)$ time, where $\lambda = d_{\max}/d_{\min}$ is the ratio of the maximum and minimum distances between any two points in P.*

Proof. Let B_0 be an axis-parallel square of minimum width $w_0 \leq d_{\max}$ that contains P. Select any set of k points $R_0 \subseteq P$ as representatives for B_0 and set their transmission radii to $\sqrt{2}w_0$. Divide B_0 into four subsquares of width $w_0/2$ and partition $P \setminus R_0$ accordingly. This procedure is applied recursively as follows. Each non-empty square B_i of width w_i contains some set $P_i \subseteq P$. Select a representative set $R_i \subseteq P_i$ arbitrarily, where $|R_i| = \min\{k, |P_i|\}$. Set the transmission radius of each $p \in R_i$ to $\max_{q \in B_j} \operatorname{dist}(p, q)$, where B_j is the parent square to B_i (i.e., q is one of the corners of B_j). The square B_i is divided into four squares of width $w_i/2$ and $P_i \setminus R_i$ is partitioned accordingly. The recursion terminates when $|P_i| \leq k$.

The first k representatives form a k-clique. Each remaining node is connected to the k representatives of its parent square. Consequently, any node forms a k-connected graph with its ancestors in the quadtree. Therefore, the entire network is k-connected.

The width of the root square is at most $d_{\max}$. The width of the lowest leaf square in the quadtree is at least $d_{\min}/(2\sqrt{2})$. Therefore, the height of the quadtree is at most $\lceil \log(2\sqrt{2}\lambda) \rceil = \lceil 3/2 + \log \lambda \rceil$. Each representative interferes with at most 32 cells at its level in the quadtree; see Figure 2. Therefore, each node $p \in P$ receives interference from at most $32k$ nodes at each level of the tree, for a total interference of at most $32k\lceil 3/2 + \log \lambda \rceil \in O(k \log \lambda)$.

At each node of the quadtree, k representatives are selected and have their transmission radii assigned, and the set P_i is partitioned into four subsets in $O(|P_i|)$ time. Since the quadtree's height is $O(\log \lambda)$, the total time is $O(n \log \lambda)$. □

5.2 $O(k\sqrt{n})$ Interference

In this section we describe an algorithm that constructs a k-connected network with maximum interference $O(k\sqrt{n})$ for any given set P of n points in $\mathbb{R}^2$. We assume a non-degeneracy condition on points, specifically, that no two points lie on the same line forming an angle of 0, $\pi/3$, or $2\pi/3$ with the x-axis.

This algorithm first selects a set H of $O(k\sqrt{n})$ hubs by finding an $((k\sqrt{n})^{-1})$-net of size $O(k\sqrt{n})$ on P as in the algorithm of Halldórsson and Tokuyama [7] described in Section 2.1. Consequently, any range containing at least $O(\sqrt{n}/k)$ points of P must contain a hub. Next, a k-connected backbone is built on the hubs. Finally, each non-hub node is connected to its k nearest hubs.

It suffices to k-connect the hubs by forming a clique on the hubs. Although the hubs could be k-connected by applying the algorithm recursively, this does not lead to any asymptotic reduction in the maximum interference. Connecting hubs by a tree, such as the MST or the local neighbourhood graph, does not guarantee k-connectivity after non-hubs connect to their k nearest hubs. For small k (e.g., $k \leq 3$) the Delaunay triangulation provides a good strategy for k-connecting hubs, but a more general strategy is required for larger k.

We analyze the maximum interference of the resulting network. Consider an arbitrary point $p \in P$. Divide the plane around p into six cones $R_1(P), \ldots, R_6(p)$ such that for each i, $R_i(p)$ is the cone consisting of all rays with apex p and angle in $[(i-1)\pi/3, i\pi/3]$. Without loss of generality, we consider the cone $R_1(p)$; analogous results apply to the remaining cones. Let $h_1, \ldots, h_k$ denote the k hubs nearest to p in $R_1(p)$ ordered by increasing distance to p. Let $l_\alpha(p)$ denote the line through p with angle α.

Lemma 4 *No point in $R_1(p) \cap (P \setminus H)$ lies on the right of $l_{2\pi/3}(h_k)$ and interferes with p.*

Proof. For the sake of contradiction, assume such a point q exists. Consequently, the transmission radius of q is at least $\operatorname{dist}(p, q)$, and so, q is connected to some hub $h \in H$ where $\operatorname{dist}(p, q) < \operatorname{dist}(q, h)$. However, $\operatorname{dist}(q, h_i) < \operatorname{dist}(p, q) < \operatorname{dist}(q, h)$ for all $i \in \{1, \ldots, k\}$, contradicting the fact that q is connected to its k nearest hubs. □

Lemma 5 *There are $O(k\sqrt{n})$ nodes in the area enclosed by $l_0(p)$, $l_{\pi/3}(p)$, and $l_{2\pi/3}(h_k)$.*

Proof. We decompose the range enclosed by $l_0(p)$, $l_{\pi/3}(p)$, and $l_{2\pi/3}(h_k)$ into smaller regions and count the vertices in each region. The first region is the range enclosed by $l_0(p)$, $l_{\pi/3}(p)$, and $l_{2\pi/3}(h_1)$. As this range contains no hub, it contains at most $c\sqrt{n}/k$ nodes of P, for some fixed $c \in \mathbb{R}^+$.

For each $i \in \{1, \ldots, k-1\}$, let Q_i denote the isosceles trapezoidal region enclosed by $l_0(p)$, $l_{\pi/3}(p)$, $l_{2\pi/3}(h_i)$, and $l_{2\pi/3}(h_{i+1})$. We identify ranges in $\mathcal{R}$ that contain no hub whose union covers Q_i. Let H'_1 be a list of the i nearest hubs to p in descending order according to their distance to $l_{\pi/3}(p)$. For each j, let h'_j denote the first hub in the list H'_j. Let A_1 be the range enclosed by $l_0(p)$, $l_{\pi/3}(h'_1)$, and $l_{2\pi/3}(h_k)$. For $j \geq 2$, let $H'_j = H'_{j-1} \setminus \{h'_{j-1}$ and all hubs below $l_0(h'_{j-1})\}$. If $H'_j \neq \varnothing$, let A_j be the range enclosed by $l_0(h'_{j-1})$, $l_{\pi/3}(h'_j)$, and $l_{2\pi/3}(h_i)$. Otherwise, A_{j-1} is the final range necessary to cover Q_i, and we let A_{j-1} be the range enclosed by $l_0(h'_{j-1})$, $l_{\pi/3}(p)$, and $l_{2\pi/3}(h_k)$. This procedure selects at most $i+1$ ranges whose union covers Q_i, each of which contains no hub in its interior. See Figure 3.

Along with the first range, the region $\bigcup_{i=1}^{k-1} Q_i$ is exactly the entire region enclosed by $l_0(p)$, $l_{\pi/3}(p)$, and $l_{2\pi/3}(h_k)$. Since each Q_i can be covered by $i+1$ ranges, each of which contains no hub in its interior, the entire region can be covered by $3k/2+k^2/2$ ranges. Since each empty range contains at most $c\sqrt{n}/k$ nodes of P, the region enclosed by $l_0(p)$, $l_{\pi/3}(p)$, and $l_{2\pi/3}(h_k)$ contains at most $ck\sqrt{n} \in O(k\sqrt{n})$ nodes of P. ☐

Theorem 6 *Given any set P of n points in $\mathbb{R}^2$ and any $k < n$, transmission radii corresponding to a k-connected network on P with maximum interference $O(k\sqrt{n})$ can be found in $O(nk + n\log n + k^3\sqrt{n}\log n)$ time.*

Proof. We first argue that the resulting network is k-connected. The clique of hubs is k-connected. Each non-hub node is connected to k hubs. Therefore, the entire network is k-connected.

Next we bound the maximum interference. By Lemmas 4 and 5, for any node $p \in P$, $O(k\sqrt{n})$ non-hub nodes interfere with p in each of the six cones around p. There are $O(k\sqrt{n})$ hubs, each of which may interfere with p. Therefore, $I(p) \in O(k\sqrt{n})$.

Finally we analyze the algorithm's running time. Since this algorithm require running part of the algorithm of Halldórsson and Tokuyama [7] described in Section 2.1, we begin by analyzing the time it takes to build the ϵ-net.

Greedily constructing the maximal family of disjoint subsets can be achieved in $O(n\log n)$ time. Similarly, the generalized Delaunay triangulation can be constructed in $O(n\log n)$ time [6] after constructing the Θ-graph (e.g., see [4, 10, 15]). Finding corridors, subcorridors, and their endpoints can done greedily in $O(n)$ time.

In our algorithm we form a clique on the set H of hubs, which can be done in $O(|H|\log|H|)$ time by finding the convex hull of the hubs and setting the transmission radius of each hub to the distance to its furthest hub in $O(\log|H|)$ time per hub using binary search on the boundary of the convex hull, or $O(|H|\log|H|)$ total time. In the final step, we set the transmission radius of each non-hub node to the distance to its kth nearest hub. To do so we can compute a k-nearest neighbour Voronoi diagram of the set H of hubs in $O(k^2|H|\log|H|)$ time [14], upon which a point location data structure (e.g., [12]) is constructed in $O(k|H|(\log k+\log|H|))$ time and applied in $O(k+\log|H|)$ time per non-hub node, or $O(nk+n\log|H|)$ total time. Thus, the running time is dominated by the larger of $O(nk)$, $O(n\log n)$, and $O(k^2|H|\log|H|)$. Since $|H| \in O(k\sqrt{n})$, this gives a total running time of $O(nk+n\log n+k^3\sqrt{n}\log n)$. ☐

6 Discussion and Directions for Future Research

We showed asymptotically tight upper and lower bounds of $\Theta(\sqrt{kn})$ on the worst-case maximum interference for k-connected networks in one dimension. The lower bound $\Omega(\sqrt{kn})$ applies in two dimensions, where we showed an upper bound of $O(k\sqrt{n})$, leaving open the question of whether a k-connected network with lower maximum interference can be found. In particular, is maximum interference $O(\sqrt{kn})$ always achievable in two dimensions?

von Rickenbach et al. [16] gave a polynomial-time algorithm that builds a connected network with interference at most $O(n^{1/4} \cdot \mathrm{OPT}_1(P))$ for any set P of n points on the line. Their algorithm constructs a network either by applying the hub strategy or returning $\mathrm{MST}(P)$, whichever has lower maximum interference. To bound the approximation factor they rely on a pair of lemmas showing that $\mathrm{OPT}_1(P) \in O(\sqrt{n})$ and $\mathrm{OPT}_1(P) \in \Omega(\sqrt{I(\mathrm{MST}(P))})$. A natural direction for future research is to determine whether this approximation algorithm can be generalized to build a k-connected network in one dimension. Instead of connecting to the nearest neighbours to the left and right as in a one-dimensional MST, we can consider the graph $\mathrm{MST}_k(P)$, in which each point connects to its k nearest neighbours to the left and k nearest neighbours to the right. In Theorem 2 we showed the generalization of the first lemma, i.e., that $\mathrm{OPT}_k(P) \in O(\sqrt{nk})$. It remains open whether the second lemma generalizes. I.e., is $\mathrm{OPT}_k(P) \in \Omega(\sqrt{I(\mathrm{MST}_k(P))})$ for any set $P \subseteq \mathbb{R}$?

Finally, Buchin [3] showed that the problem of finding a connected network that minimizes maximum interference for a given set of n points in two dimensions is NP-complete. The complexity of the interference minimization problem in one dimension remains an important open question.

References

[1] A. K. Abu-Affash, R. Aschner, P. Carmi, and M. J. Katz. The MST of symmetric disk graphs is light. *Comp. Geom.: Theory & Appl.*, 5(1–2):54–61, 2012.

[2] H. Aslanyan and J. D. P. Rolim. Interference minimization in wireless networks. In *Proc. EUC*, pages 444–449, 2010.

[3] K. Buchin. Minimizing the maximum interference is hard. *CoRR*, abs/0802.2134, 2008.

[4] K. Clarkson. Approximation algorithms for shortest path motion planning. In *Proc. ACM STOC*, pages 56–65, 1987.

[5] L. Devroye and P. Morin. A note on interference in random point sets. In *Proc. CCCG*, volume 24, pages 201–206, 2012.

[6] S. Fortune. Voronoi diagrams and Delaunay triangulations. *Comp. in Euclidean Geom.*, 1:193–233, 1992.

[7] M. M. Halldórsson and T. Tokuyama. Minimizing interference of a wireless ad-hoc network in a plane. *Theor. Comp. Sci.*, 402(1):29–42, 2008. Algorithmic Aspects of Wireless Sensor Networks.

[8] D. Haussler and E. Welzl. ϵ-nets and simplex range queries. *Disc. & Comp. Geom.*, 2:127–151, 1987.

[9] E. Holec. Using linear programming to minimize interference in wireless networks. Master's thesis, Univ. Minnesota, 2013.

[10] J. M. Keil. Approximating the complete Euclidean graph. In *Proc. SWAT*, pages 208–213. Springer, 1988.

[11] M. Khabbazian, S. Durocher, A. Haghnegahdar, and F. Kuhn. Bounding interference in wireless ad hoc networks with nodes in random position. *IEEE/ACM Trans. Net.*, 23(4):1078–1091, 2015.

[12] D. Kirkpatrick. Optimal search in planar subdivisions. *SIAM J. Comp.*, 12(1):28–35, 1983.

[13] E. Kranakis, D. Krizanc, P. Morin, L. Narayanan, and L. Stacho. A tight bound on the maximum interference of random sensors in the highway model. *CoRR*, abs/1007.2120, 2010.

[14] D.-T. Lee. On k-nearest neighbor Voronoi diagrams in the plane. *IEEE Trans. Comp.*, 31(6):478–487, 1982.

[15] G. Narasimhan and M. Smid. *Geometric spanner networks.* Cambridge University Press, 2007.

[16] P. v. Rickenbach, S. Schmid, R. Wattenhofer, and A. Zollinger. A robust interference model for wireless ad-hoc networks. In *Proc. IPDPS*, pages 239.1–8, 2005.

[17] H. Tan. *Minimizing interference in wireless sensor networks.* PhD thesis, Univ. Hong Kong, 2011.

[18] H. Tan, T. Lou, F. C. M. Lau, Y. Wang, and S. Chen. Minimizing interference for the highway model in wireless ad-hoc and sensor networks. In *Proc. SOFSEM*, pages 520–532, 2011.

A Appendix: Figures

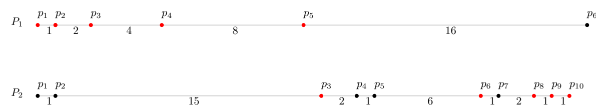

Figure 1: The first five points of P_1 form an exponential chain of size 5, where $a_1 = 5, a_2 = 4, \ldots, a_5 = 1$. An exponential chain need not be a perfect geometric sequence, nor need its points be consecutive. For example, $a_1 = 3$, $a_2 = 6$, $a_3 = 8$, $a_4 = 9$, $a_5 = 10$ (red points) is an exponential chain of size 5 in P_2. The exponential chain property holds for $i = 1$ in P_2 since $\mathrm{dist}(p_{a_1}, p_{a_1 - 1}) = \mathrm{dist}(p_3, p_2) = 15 \geq \max_{j=2}^{5} \mathrm{dist}(p_{a_1}, p_{a_j}) = \mathrm{dist}(p_3, p_{10}) = 14$; it also holds for all $i \in \{2, 3, 4\}$.

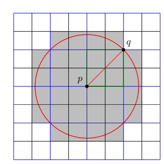

Figure 2: A point p is selected as a representative for a quadtree cell, denoted by the smaller bold green square of width w_i. In the worst case, the furthest representative, q, of the parent square of p, denoted by the larger bold green square, is a distance $2\sqrt{2}w_i$ from p. Consequently, p's transmission range interferes with at most 32 cells of the quadtree at its level.

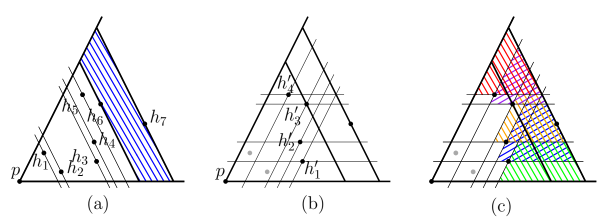

Figure 3: (a) The shaded region is the trapezoid Q_6. (b) Four hubs that determine the ranges used to cover Q_6. (c) The five empty ranges whose union covers Q_6.