Bounding Interference in Wireless Ad Hoc Networks with Nodes in Random Position*

Majid Khabbazian[1], Stephane Durocher[2], and Alireza Haghnegahdar[3]

[1] University of Winnipeg, Winnipeg, Canada, m.khabbazian@uwinnipeg.ca
[2] University of Manitoba, Winnipeg, Canada, durocher@cs.umanitoba.ca
[3] University of British Columbia, Vancouver, Canada, alirezah@ece.ubc.ca

Abstract. Given a set of positions for wireless nodes, the interference minimization problem is to assign a transmission radius (equivalently, a power level) to each node such that the resulting communication graph is connected, while minimizing the maximum interference. We consider the model introduced by von Rickenbach et al. (2005), in which each transmission range is represented by a ball and edges in the communication graph are symmetric. The problem is NP-complete in two dimensions (Buchin 2008) and no polynomial-time approximation algorithm is known. In this paper we show how to solve the problem efficiently in settings typical for wireless ad hoc networks. We show that if node positions are represented by a set P of n points selected uniformly and independently at random over a d-dimensional rectangular region, for any fixed d, then the topology given by the closure of the Euclidean minimum spanning tree of P has maximum interference $O(\log n)$ with high probability. We extend this bound to a general class of communication graphs over a broad set of probability distributions. We present a local algorithm that constructs a graph from this class; this is the first local algorithm to provide an upper bound on the expected maximum interference. Finally, we analyze an empirical evaluation of our algorithm by simulation.

1 Introduction

1.1 Motivation

Establishing connectivity in a wireless network can be a complex task for which various (sometimes conflicting) objectives must be optimized. To permit a packet to be routed from any origin node to any destination node in the network, the corresponding communication graph must be connected. In addition to requiring connectivity, various properties can be imposed on the network, including low power consumption [20, 27], bounded average traffic load [10, 12], small average

* This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC). The full version of this paper was posted online by the authors in November 2011 [15]. Proofs omitted here due to space constraints can found in the full version.

hop distance between sender-receiver pairs [1], low dilation (t-spanner) [1,3,6,7, 14,21,25], and minimal interference; this latter objective, minimizing interference (and, consequently, minimizing the required bandwidth), is the focus of much recent research [1,2,5,9,11,17–19,22–24,27–31] and of this paper.

We adopt the interference model introduced by von Rickenbach et al. [30] (see Section 1.2). We model transmission in a wireless network by assigning to each wireless node p a radius of transmission $r(p)$, such that every node within distance $r(p)$ of p can receive a transmission from p, whereas no node a greater distance from p can. The interference at node p is the number of nodes that have p within their respective radii of transmission. Given a set of wireless nodes whose positions are represented by a set of points P, we consider the problem of identifying a connected network on P that minimizes the maximum interference. The problem of constructing the network is equivalent to that of assigning a transmission radius to each node; once the transmission radius of each node is fixed, the corresponding communication graph and its associated maximum interference are also determined. Conversely, once a graph is fixed, each node's transmission radius is determined by the distance to its furthest neighbour.

Given a set of points P in the plane, finding a connected graph on P that minimizes the maximum interference is NP-complete [5]. A polynomial-time algorithm exists that returns a solution with maximum interference $O(\sqrt{n})$, where $n = |P|$ [11]. Even in one dimension, for every n there exists a set of n points P such that any graph on P has maximum interference $\Omega(\sqrt{n})$ [30]. All such known examples involve specific constructions (i.e., exponential chains). We are interested in investigating a more realistic class of wireless networks: those whose node positions observe common random distributions that better model actual wireless ad hoc networks.

When nodes are positioned on a line (often called the *highway model*), a simple heuristic is to assign to each node a radius of transmission that corresponds to the maximum of the distances to its respective nearest neighbours to the left and right. In the worst case, such a strategy can result in $\Theta(n)$ maximum interference when an optimal solution has only $\Theta(\sqrt{n})$ maximum interference [30]. Recently, Kranakis et al. [19] showed that if n nodes are positioned uniformly at random on an interval, then the maximum interference provided by this heuristic is $\Theta(\sqrt{\log n})$ with high probability.

In this paper, we examine the corresponding problem in two and higher dimensions. We generalize the nearest-neighbour path used in the highway model to the Euclidean minimum spanning tree (MST), and show that with high probability, the maximum interference of the MST of a set of n points selected uniformly at random over a d-dimensional region $[0,1]^d$ is $O(\log n)$, for any fixed $d \geq 1$. Our techniques differ significantly from those used by Kranakis et al. to achieve their results in one dimension. As we show in Section 3, our results also apply to a broad class of random distributions, denoted $\mathcal{D}$, that includes both the uniform random distribution and realistic distributions for modelling random motion in mobile wireless networks, as well as to a large class of connected spanning graphs that includes the MST.

In Section 3.4 we present a local algorithm that constructs a topology whose maximum interference is $O(\log n)$ with high probability when node positions are selected according to a distribution in $\mathcal{D}$. Previous local algorithms for topology control (e.g., the cone-based local algorithm (CBTC) [20]) attempt to reduce transmission radii (i.e., power consumption), but not necessarily the maximum interference. Although reducing transmission radii at many nodes is often necessary to reduce the maximum interference, the two objectives differ; specifically, some nodes may require large transmission radii to minimize the maximum interference. Ours is the first local algorithm to provide a non-trivial upper bound on maximum interference. Our algorithm can be applied to any existing topology to refine it and further reduce its maximum interference. Consequently, our solution can be used either independently, or paired with another topology control strategy. Finally, we discuss an empirical evaluation of our algorithm with a suite of simulation results in Section 4.

1.2 Model and Definitions

We represent the position of a wireless node as a point in Euclidean space, $\mathbb{R}^d$, for some fixed $d \geq 1$. For simplicity, we refer to each node by its corresponding point. Similarly, we represent a wireless network by its communication graph, a geometric graph whose vertices are a set of points $P \subseteq \mathbb{R}^d$. Given a (simple and undirected) graph G, we employ standard graph-theoretic notation, where $V(G)$ denotes the vertex set of G and $E(G)$ denotes its edge set. We say vertices u and v are *k-hop neighbours* if there is a simple path of length k from u to v in G. When $k = 1$ we say u and v are neighbours.

We assume a uniform range of communication for each node and consider bidirectional communication links, each of which is represented by an undirected graph edge connecting two nodes. Specifically, each node p has some *radius of transmission*, denoted by the function $r : P \to \mathbb{R}^+$, such that a node q receives a transmission from p if and only if $\text{dist}(p, q) \leq r(p)$, where $\text{dist}(p, q) = \|p - q\|_2$ denotes the Euclidean distance between points p and q in $\mathbb{R}^d$. For simplicity, suppose each node has an infinite radius of reception, regardless of its radius of transmission.

Definition 1 (Communication Graph). *A graph G is a* communication graph *with respect to a point set $P \subseteq \mathbb{R}^d$ and a function $r : P \to \mathbb{R}^+$ if (i) $V(G) = P$, and*

$$\text{(ii) } \forall \{p, q\} \subseteq V(G),\ \{p, q\} \in E(G) \Leftrightarrow \text{dist}(p, q) \leq \min\{r(p), r(q)\}. \quad (1)$$

Together, set P and function r uniquely determine the corresponding communication graph G. Alternatively, a communication graph can be defined as the closure of a given embedded graph. Specifically, if instead of being given P and r, we are given an arbitrary graph H embedded in $\mathbb{R}^d$, then the set P is trivially determined by $V(H)$ and a transmission radius for each node $p \in V(H)$ can be assigned to satisfy (1) by

$$r(p) = \max_{q \in \text{Adj}(p)} \text{dist}(p, q), \quad (2)$$

where $\text{Adj}(p) = \{q \mid \{q, p\} \in E(H)\}$ denotes the set of vertices adjacent to p in H. The communication graph determined by H is the unique edge-minimal supergraph of H that satisfies Definition 1. We denote this graph by H' and refer to it as the *closure* of graph H. Therefore, a communication graph G can be defined either as a function of a set of points P and an associated mapping of transmission radii $r : P \to \mathbb{R}^+$, or as the closure of a given embedded graph H (where $G = H'$).

Definition 2 (Interference). *Given a communication graph G, the* interference *at node p in $V(G)$ is*

$$\text{inter}_G(p) = |\{q \mid q \in V(G) \setminus \{p\} \text{ and } \text{dist}(q, p) \leq r(q)\}|$$

and the maximum interference *of G is* $\text{inter}(G) = \max_{p \in V(G)} \text{inter}_G(p)$.

In other words, the interference at node p, denoted $\text{inter}_G(p)$, is the number of nodes q such that p lies within q's radius of transmission. This does not imply the existence of the edge $\{p, q\}$ in the corresponding communication graph; such an edges exists if and only if the relationship is reciprocal, i.e., q also lies within p's radius of transmission.

Given a point set P, let $\mathcal{G}(P)$ denote the set of connected communication graphs on P. Let $\text{OPT}(P)$ denote the optimal maximum interference attainable over graphs in $\mathcal{G}(P)$. That is,

$$\text{OPT}(P) = \min_{G \in \mathcal{G}(P)} \text{inter}(G) = \min_{G \in \mathcal{G}(P)} \max_{p \in V(G)} \text{inter}_G(p).$$

Thus, given a set of points P representing the positions of wireless nodes, the *interference minimization problem* is to find a connected communication graph G on P that spans P such that the maximum interference is minimized (i.e., its maximum interference is $\text{OPT}(P)$). In this paper we examine the maximum interference of the communication graph determined by the closure of $\text{MST}(P)$, where $\text{MST}(P)$ denotes the Euclidean minimum spanning tree of the point set P. Our results apply with high probability, which refers to probability at least $1 - n^{-c}$, where $n = |P|$ denotes the number of networks nodes and $c \geq 1$ is fixed.

2 Related Work

Bidirectional Interference Model. In this paper we consider the bidirectional interference model (defined in Section 1.2). This model was introduced by von Rickenbach et al. [30], who gave a polynomial-time approximation algorithm that finds a solution with maximum interference $O(n^{1/4} \cdot \text{OPT}(P))$ for any given set of points P on a line, and a one-dimensional construction showing that $\text{OPT}(P) \in \Omega(\sqrt{n})$ in the worst case, where $n = |P|$. Halldórsson and Tokuyama [11] gave a polynomial-time algorithm that returns a solution with maximum interference $O(\sqrt{n})$ for any given set of n points in the plane. Buchin [5] showed that finding an optimal solution (one whose maximum interference is exactly $\mathrm{OPT}(P)$) is NP-complete in the plane. Tan et al. [29] gave an $O(n^3 n^{O(\mathrm{OPT}(P))})$-time algorithm for finding an optimal solution for any given set of points P on a line. Kranakis et al. [19] showed that for any set of n points P selected uniformly at random from the unit interval, the maximum interference of the nearest-neighbour path ($\mathrm{MST}(P)'$) has maximum interference $\Theta(\sqrt{\log n})$ with high probability. Sharma et al. [28] consider heuristic solutions to the two-dimensional problem. Finally, recent results by Devroye and Morin [9] extend some of the results presented in this paper and answer a number of open questions definitively to show that with high probability, when P is a set of n points in $\mathbb{R}^d$ selected uniformly at random from $[0,1]^d$, $\mathrm{inter}(\mathrm{MST}(P)') \in \Theta((\log n)^{1/2})$, $\mathrm{OPT}(P) \in O((\log n)^{1/3})$, and $\mathrm{OPT}(P) \in \omega((\log n)^{1/4})$.

Unidirectional Interference Model. If communication links are not bidirectional (i.e., edges are directed) and the communication graph is required to be strongly connected, then the worst-case maximum interference decreases. Under this model, von Rickenbach et al. [31] and Korman [17] give polynomial-time algorithms that return solutions with maximum interference $O(\log n)$ for any given set of points in the plane, and a one-dimensional construction showing that in the worst case $\mathrm{OPT}(P) \in \Omega(\log n)$.

Minimizing Average Interference. In addition to results that examine the problem of minimizing the maximum interference, some work has addressed the problem of minimizing the average interference, e.g., Tan et al. [29] and Moscibroda and Wattenhofer [24].

3 Bounds

3.1 Generalizing One-Dimensional Solutions

Before presenting our results on random sets of points, we begin with a brief discussion regarding the possibility of generalizing existing algorithms that provide approximate solutions for one-dimensional instances of the interference minimization problem (in an adversarial deterministic input setting).

Since the problem of identifying a graph that achieves the optimal (minimum) interference is NP-hard in two or more dimensions [5], it is natural to ask whether one can design a polynomial-time algorithm to return a good approximate solution. Although Rickenbach et al. [30] give a $\Theta(n^{1/4})$-approximate algorithm in one dimension [30], the current best polynomial-time algorithm in two (or more) dimensions by Halldórsson and Tokuyama [11] returns a solution whose maximum interference is $O(\sqrt{n})$; as noted by Halldórsson and Tokuyama, this algorithm is not known to guarantee any approximation factor better than the immediate bound of $O(\sqrt{n})$. The algorithm of Rickenbach et al. uses two strategies for constructing respective communication graphs, and returns the graph with the lower maximum interference; an elegant argument that depends on Lemma 1 bounds the resulting worst-case maximum interference by

$\Theta(n^{1/4} \cdot \mathrm{OPT}(P))$. The two strategies correspond roughly to a) $\mathrm{MST}(P)'$ and b) classifying every $\sqrt{n}$th node as a hub, joining each hub to its left and right neighbouring hubs to form a network backbone, and connecting each remaining node to its closest hub. The algorithm of Halldórsson and Tokuyama applies ϵ-nets, resulting in a strategy that is loosely analogous to a generalization of the hub strategy of Rickenbach et al. to higher dimensions. One might wonder whether the hybrid approach of Rickenbach et al. might be applicable in higher dimensions by returning $\mathrm{MST}(P)'$ or the communication graph constructed by the algorithm of Halldórsson and Tokuyama, whichever has lower maximum interference. To apply this idea directly would require generalizing the following property established by von Rickenbach et al. to higher dimensions:

Lemma 1 (von Rickenbach et al. [30] (2005)). *For any set of points $P \subseteq \mathbb{R}$,*

$$\mathrm{OPT}(P) \in \Omega\left(\sqrt{\mathrm{inter}(\mathrm{MST}(P)')}\right).$$

However, von Rickenbach et al. also show that for any n, there exists a set of n points $P \subseteq \mathbb{R}^2$ such that $\mathrm{OPT}(P) \in O(1)$ and $\mathrm{inter}(\mathrm{MST}(P)') \in \Theta(n)$, which implies that Lemma 1 does not hold in higher dimensions. Consequently, techniques such as those used by von Rickenbach et al. do not immediately generalize to higher dimensions.

3.2 Randomized Point Sets

Although using the hybrid approach of von Rickenbach et al. [30] directly may not be possible, Kranakis et al. [19] recently showed that if a set P of n points is selected uniformly at random from an interval, then the maximum interference of the communication graph determined by $\mathrm{MST}(P)'$ is $\Theta(\sqrt{\log n})$ with high probability. Throughout this section, we assume general position of points; specifically, we assume that the distance between each pair of nodes is unique.

We begin by introducing some definitions. An edge $\{p, q\} \in E(G)$ in a communication graph G is *primitive* if $\min\{r(p),\ r(q)\} = \mathrm{dist}(p, q)$. An edge $\{p, q\} \in E(G)$ in a communication graph G is *bridged* if there is a path joining p and q in G consisting of at most three edges, each of which is of length less than $\mathrm{dist}(p, q)$. Given a set of points P in $\mathbb{R}^d$, let $\mathcal{T}(P)$ denote the set of all communication graphs G with $V(G) = P$ such that no primitive edge in $E(G)$ is bridged.

Halldórsson and Tokuyama [11] and Maheshwari et al. [23] give respective centralized algorithms for constructing graphs G, each with interference $O(\log(d_{\max}(G)/d_{\min}(G)))$, where $d_{\max}(G)$ and $d_{\min}(G)$ are defined as in Theorem 1. As we show in Theorem 1, this bound holds for any graph G in the class $\mathcal{T}(P)$. In Section 3.4 we give a local algorithm for constructing a connected graph in $\mathcal{T}(P)$ on any given point set P.

Theorem 1. *Let P be a set of points in $\mathbb{R}^d$. For any graph $G \in \mathcal{T}(P)$,*

$$\mathrm{inter}(G) \in O\left(\log\left(\frac{d_{\max}(G)}{d_{\min}(G)}\right)\right),$$

where $d_{\max}(G) = \max_{\{s,t\}\in E(G)} \text{dist}(s,t)$ *and* $d_{\min}(G) = \min_{\{s,t\}\in E(G)} \text{dist}(s,t)$.

The proof is omitted due to space constraints. In the next lemma we show that $\text{MST}(P)'$ is in $\mathcal{T}(P)$. Consequently, $\mathcal{T}(P)$ is always non-empty.

Lemma 2. *For any set of points* $P \subseteq \mathbb{R}^d$, $\text{MST}(P)' \in \mathcal{T}(P)$.

Proof. The transmission range of each node $p \in P$ is determined by the length of the longest edge adjacent to p in $\text{MST}(P)$. Suppose there is a primitive edge $\{p_1, p_2\} \in E(\text{MST}(P))$ that is bridged. Therefore, there is a path T from p_1 to p_2 in $\text{MST}(P)'$ that contains at most three edges, each of which is of length less than $\text{dist}(p_1, p_2)$. Removing the edge $\{p_1, p_2\}$ partitions $\text{MST}(P)$ into two connected components, where p_1 and p_2 are in different components. By definition, T contains an edge that spans the two components. The two components can be joined using this edge (of length less than $\text{dist}(p_1, p_2)$) to obtain a new spanning tree whose weight is less than that of $\text{MST}(P)$, deriving a contradiction. Therefore, no primitive edge $\{p_1, p_2\} \in \text{MST}(P)$ can be bridged, implying $\text{MST}(P)' \in \mathcal{T}(P)$. □

Theorem 1 implies that the interference of any graph G in $\mathcal{T}(P)$ is bounded asymptotically by the logarithm of the ratio of the longest and shortest edges in G. While this ratio can be arbitrarily large in the worst case, we show that the ratio is bounded for many typical distributions of points. Specifically, if the ratio is $O(n^c)$ for some constant c, then the maximum interference is $O(\log n)$.

Definition 3 ($\mathcal{D}$). *Let $\mathcal{D}$ denote the class of distributions over $[0,1]^d$ such that for any $D \in \mathcal{D}$ and any set P of $n \geq 2$ points selected independently at random according to D, the minimum distance between any two points in P is greater than n^{-c} with high probability, for some constant c (independent of n).*

Theorem 2. *For any integers $d \geq 1$ and $n \geq 2$, any distribution $D \in \mathcal{D}$, and any set P of n points, each of which is selected independently at random over $[0,1]^d$ according to distribution D, with high probability, for all graphs $G \in \mathcal{T}(P)$,* $\text{inter}(G) \in O(\log n)$.

Proof. Let $d_{\min}(G) = \min_{\{s,t\}\in E(G)} \text{dist}(s,t)$ and $d_{\max}(G) = \max_{\{s,t\}\in E(G)} \text{dist}(s,t)$. Since points are contained in $[0,1]^d$, $d_{\max}(G) \leq \sqrt{d}$. Points in P are distributed according to a distribution $D \in \mathcal{D}$. By Definition 3, with high probability, $d_{\min}(G) \geq n^{-c}$ for some constant c. Thus, with high probability, we have

$$\log\left(\frac{d_{\max}(G)}{d_{\min}(G)}\right) \leq \log\left(\frac{\sqrt{d}}{n^{-c}}\right). \tag{3}$$

The result follows from (3), Theorem 1, and the fact that $\log(n^c\sqrt{d}) \in O(\log n)$ when d and c are constant. □

Lemma 3. *Let D be a distribution with domain $[0,1]^d$, for which there is a constant c' such that for any point $x \in [0,1]^d$, we have $D(x) \leq c'$, where $D(x)$ denotes the probability density function of D at $x \in [0,1]^d$. Then $D \in \mathcal{D}$.*

The proof is omitted due to space constraints.

Corollary 1. *The uniform distribution with domain* $[0,1]^d$ *is in* $\mathcal{D}$*.*

By Corollary 1 and Theorem 2, we can conclude that if a set P of $n \geq 2$ points is distributed uniformly in $[0,1]^d$, then with high probability, any communication graph in $G \in \mathcal{T}(P)$ will have maximum interference $O(\log n)$. This is expressed formally in the following corollary:

Corollary 2. *Choose any integers* $d \geq 1$ *and* $n \geq 2$*. Let* P *be a set of* n *points, each of which is selected independently and uniformly at random over* $[0,1]^d$*. With high probability, for all graphs* $G \in \mathcal{T}(P)$*,* $\text{inter}(G) \in O(\log n)$*.*

3.3 Mobility

Our results apply to the setting of mobility (e.g., mobile ad hoc wireless networks). Each node in a mobile network must periodically exchange information with its neighbours to update its local data storing positions and transmission radii of nodes within its local neighbourhood. The distribution of mobile nodes depends on the mobility model, which is not necessarily uniform. For example, when the network is distributed over a disc or a box-shaped region, the probability distribution associated with the random waypoint model achieves its maximum at the centre of the region, whereas the probability of finding a node close to the region's boundary approaches zero [12]. Since the maximum value of the probability distribution associated with the random waypoint model is constant [12], by Lemma 3 and Theorem 2, we can conclude that at any point in time, the maximum interference of the network is $O(\log n)$ with high probability. In general, this holds for any random mobility model whose corresponding probability distribution has a constant maximum value.

3.4 Local Algorithm

As discussed in Section 1.1, existing local algorithms for topology control attempt to reduce transmission radii, but not necessarily the maximum interference. By Lemma 2 and Theorem 2, if P is a set of n points selected according to a distribution in $\mathcal{D}$, then with high probability $\text{inter}(\text{MST}(P)') \in O(\log n)$. Unfortunately, a minimum spanning tree cannot be generated using only local information [16]. Thus, an interesting question is whether each node can assign itself a transmission radius using only local information such that the resulting communication graph belongs to $\mathcal{T}(P)$ while remaining connected. We answer this question affirmatively and present the first local algorithm (LocalRadiusReduction), that assigns a transmission radius to each node such that if the initial communication graph $G_{\max}$ is connected, then the resulting communication graph is a connected spanning subgraph of $G_{\max}$ that belongs to $\mathcal{T}(P)$. Consequently, the resulting topology has maximum interference $O(\log n)$ with high probability when nodes are selected according to any distribution in $\mathcal{D}$. Our algorithm can be applied to any existing topology to refine it and further reduce its maximum interference. Thus, our solution can be used either independently, or paired with another topology control strategy. The algorithm consists of three phases, which we now describe.

Let P be a set of $n \geq 2$ points in $\mathbb{R}^d$ and let $r_{\max} : P \rightarrow \mathbb{R}^+$ be a function that returns the maximum transmission radius allowable at each node. Let $G_{\max}$ denote the communication graph determined by P and $r_{\max}$. Suppose $G_{\max}$ is connected. Algorithm LocalRadiusReduction assumes that each node is initially aware of its maximum transmission radius, its spatial coordinates, and its unique identifier.

The algorithm begins with a local data acquisition phase, during which every node broadcasts its identity, maximum transmission radius, and coordinates in a node data message. Each message also specifies whether the data is associated with the sender or whether it is forwarded from a neighbour. Every node records the node data it receives and retransmits those messages that were not previously forwarded. Upon completing this phase, each node is aware of the corresponding data for all nodes within its 2-hop neighbourhood. The algorithm then proceeds to an asynchronous transmission radius reduction phase.

Consider a node u and let f denote its furthest neighbour. If u and f are bridged in $G_{\max}$, then u reduces its transmission radius to correspond to that of its next-furthest neighbour f', where $\text{dist}(u, f') < \text{dist}(u, f)$. This process iterates until u is not bridged with its furthest neighbour within its reduced transmission radius. We formalize the local transmission radius reduction algorithm in the pseudocode in Table 1 that computes the new transmission radius $r'(u)$ at node u.

```
radiusReductionComplete ← false
r'(u) ← r_max(u)
f ← u
for each v ∈ Adj(u)
   if dist(u, v) > dist(u, f)
      f ← v // furthest neighbour
while ¬radiusReductionComplete
   radiusModified ← false
   if BRIDGED(u, f)
      radiusModified ← true
      f ← u // identify next neighbour within distance r'(u)
      for each v ∈ Adj(u)
         if dist(u, v) < r'(u) and dist(u, v) > dist(u, f)
            f ← v
      r'(u) ← dist(u, f)
   radiusReductionComplete ← ¬radiusModified
return r'(u)
```

Table 1. pseudocode for Algorithm LocalRadiusReduction(u)

Algorithm LocalRadiusReduction is 2-local. Since transmission radii are decreased monotonically (and never increased), the while loop iterates $O(\Delta)$ times, where Δ denotes the maximum vertex degree in $G_{\max}$. Since each call to the subroutine Bridged terminates in $O(\Delta^2)$ time, each node determines its reduced transmission radius $r'(u)$ in $O(\Delta^3)$ time.

After completing the transmission radius reduction phase, the algorithm concludes with one final adjustment in the transmission radius to remove asymmetric edges. In this third and final phase, each node u broadcasts its reduced transmission radius $r'(u)$. Consider the set of nodes $\{v_1, \ldots, v_k\} \subseteq \mathrm{Adj}(u)$ such that $\mathrm{dist}(u, v_i) = r'(u)$ for all i (when points are in general position, $k = 1$, and there is a unique such node v_1). If $r'(v_i) < r'(u)$ for all i, then u can reduce its transmission radius to that of its furthest neighbour with which bidirectional communication is possible. Specifically,

$$r'(u) \leftarrow \max_{\substack{v \in \mathrm{Adj}(u) \\ \mathrm{dist}(u,v) \leq \min\{r'(u), r'(v)\}}} \mathrm{dist}(u, v). \tag{4}$$

The value of $r'(u)$ as defined in (4) is straightforward to compute in $O(\Delta)$ time.

Lemma 4. *The communication graph constructed by Algorithm* LocalRadiusReduction *is in $\mathcal{T}(P)$ and is connected if the initial communication graph $G_{\max}$ is connected.*

The proof is omitted due to space constraints. More generally, since transmission radii are only decreased, it can be shown that $G_{\min}$ and $G_{\max}$ have the same number of connected components by applying Lemma 4 on every connected component of $G_{\max}$.

4 Simulation

We simulated Algorithm LocalRadiusReduction to evaluate its performance in static and mobile wireless networks. In both settings, each node collects the list of its 2-hop neighbours in two rounds, applies the algorithm to reduce its transmission radius, and then broadcasts its computed transmission radius so neighbouring nodes can eliminate asymmetric edges and possibly further reduce their transmission radii. By the end of this stage, all asymmetric edges are removed and no new asymmetric edges are generated. Consequently, a node need not broadcast its transmission radius again after it has been further reduced.

We applied two mobility models to simulate mobile networks: random walk and random waypoint [13]. In both models each node's initial position is a point selected uniformly at random over the simulation region. In the random walk model, each node selects a new speed and direction uniformly at random over $[v_{\min}, v_{\max}]$ and $[0, 2\pi)$, respectively, at regular intervals. When a node encounters the simulation region's boundary, its direction is reversed (a rotation of π) to remain within the simulation region with the same speed. In the random waypoint model, each node moves along a straight trajectory with constant speed toward a destination point selected uniformly at random over $[v_{\min}, v_{\max}]$ and the simulation region, respectively. Upon reaching its destination, the node stops for a random pause time, after which it selects a new random destination and speed, and the process repeats.

We set the simulation region's dimensions to 1000 metres $\times$ 1000 metres. For both static and dynamic networks, we varied the number of nodes n from 50 to 1000 in increments of 50. We fixed the maximum transmission radius $r_{\max}$ for each network to 100, 200, or 300 metres. To compute the average maximum interference for static networks, for each n and $r_{\max}$ we generated 100,000 static networks, each with n nodes and maximum transmission radius $r_{\max}$, distributed uniformly at random in the simulation region. To compute the average maximum interference for mobile networks, for each n and $r_{\max}$ we generated 100,000 snapshots for each mobility model, each with n nodes and maximum transmission radius $r_{\max}$. We set the speed interval to $[0.2, 10]$ metres per second, and the pause time interval to $[0, 10]$ seconds (in the waypoint model). A snapshot of the network was recorded once every second over a simulation of 100,000 seconds.

We compared the average maximum interference of the topology constructed by the algorithm LocalRadiusReduction against the corresponding average maximum interference achieved respectively by two local topology control algorithms: i) the local computation of the intersection of the Gabriel graph and the unit disc graph (with unit radius $r_{\max}$) [4], and ii) the cone-based local topology control (CBTC) algorithm [20]. In addition, we evaluated the maximum interference achieved when each node uses a fixed radius of communication, i.e., the communication graph is a unit disc graph of radius $r_{\max}$ (100, 200, or 300 metres, respectively). See the full version [15] for figures displaying simulation results.

As shown, the average maximum interference of unit disc graph topologies increases linearly with n. Many of the unit disc graphs generated were disconnected when the transmission radius was set to 100 metres for small n. Since we require connectivity, we only considered values of n and $r_{\max}$ for which at least half of the networks generated were connected. When $r_{\max} = 100$ metres, a higher average maximum interference was measured at $n = 300$ than at $n = 400$. This is because many networks generated for $n = 300$ were discarded due to being disconnected. Consequently, the density of networks simulated for $n = 300$ was higher than the average density of a random network with $n = 300$ nodes, resulting in higher maximum interference.

Although both the local Gabriel and CBTC algorithms performed significantly better than the unit disc graphs, the lowest average maximum interference was achieved by the LocalRadiusReduction algorithm. Note that the LocalRadiusReduction algorithm reduces the maximum interference to $O(\log n)$ with high probability, irrespective of the initial maximum transmission radius $r_{\max}$.

Simulation results obtained using the random walk model closely match those obtained on a static network because the distribution of nodes at any time during a random walk is nearly uniform [8]. The average maximum interference increases slightly but remains logarithmic when the random waypoint model is used. The spatial distribution of nodes moving according to a random waypoint model is not uniform, and is maximized at the centre of the simulation region [12]. Consequently, the density of nodes is high near the centre, resulting in greater interference at these nodes.

Finally, we evaluated the algorithm LocalRadiusReduction using actual mobility trace data of Piorkowski et al. [26], consisting of GPS coordinates for trajectories of 537 taxi vehicles recorded over one month in 2008, driving throughout the San Fransisco Bay area. We selected the 500 largest traces, each of which has over 8000 sample points. To implement our algorithm, we selected n taxis among the 500 uniformly at random, ranging from $n = 50$ to $n = 500$ in increments of 50. The resulting average maximum interference is similar to that measured in our simulation results.

5 Discussion

Using Algorithm LocalRadiusReduction, each node determines its transmission radius as a function of its 2-hop neighbourhood. Alternatively, suppose each node could select its transmission radius at random using a suitable distribution over $[d_{\min}(G), d_{\max}(G)]$. Can such a strategy for assigning transmission radii ensure connectivity and low maximum interference with high probability? Similarly, additional topologies and local algorithms for constructing them might achieve $O(\log n)$ expected maximum interference. For example, our experimental results suggest that both the Gabriel graph and CBTC local topology control algorithms may provide $O(\log n)$ expected maximum interference. Since neither the Gabriel graph nor the CBTC topology of a set of points P is in $\mathcal{T}(P)$ in general, whether these bounds hold remains to be proved.

As mentioned in Section 2, multiple open questions related to interference on random sets of points were resolved recently by Devroye and Morin [9]. Several questions remain open related to the algorithmic problem of finding an *optimal* solution (one whose maximum interference is exactly $\mathrm{OPT}(P)$) when node positions may be selected adversarially. The complexity of the interference minimization in one dimension remains open; at present, it is unknown whether the problem is polynomial-time solvable or NP-hard [29]. While the problem is known to be NP-complete in two dimensions [5], no polynomial-time approximation algorithm nor any inapproximability hardness results are known.

Acknowledgements

Stephane Durocher thanks Csaba Tóth for insightful discussions related to the interference minimization problem in one dimension.

References

1. M. Benkert, J. Gudmundsson, H. Haverkort, and A. Wolff. Constructing minimum-interference networks. *Comp. Geom.: Theory & App.*, 40(3):179–194, 2008.

2. D. Bilò and G. Proietti. On the complexity of minimizing interference in ad-hoc and sensor networks. *Theor. Comp. Sci.*, 402(1):42–55, 2008.
3. P. Bose, J. Gudmundsson, and M. Smid. Constructing plane spanners of bounded degree and low weight. *Algorithmica*, 42(3–4):249–264, 2005.
4. P. Bose, P. Morin, I. Stojmenović, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Net.*, 7(6):609–616, 2001.
5. K. Buchin. Minimizing the maximum interference is hard. *CoRR*, abs/0802.2134, 2008.
6. M. Burkhart, P. von Rickenbach, R. Wattenhofer, and A. Zollinger. Does topology control reduce interference? In *Proc. ACM MobiHoc*, pages 9–19, 2004.
7. M. Damian, S. Pandit, and S. V. Pemmaraju. Local approximation schemes for topology control. In *Proc. ACM PODC*, pages 208–218, 2006.
8. A. Das Sarma, D. Nanongkai, and G. Pandurangan. Fast distributed random walks. In *Proc. ACM PODC*, pages 161–170, 2009.
9. L. Devroye and P. Morin. A note on interference in random point sets. *CoRR*, 1202.5945, 2012.
10. S. Durocher, E. Kranakis, D. Krizanc, and L. Narayanan. Balancing traffic load using one-turn rectilinear routing. *J. Interconn. Net.*, 10(1–2):93–120, 2009.
11. 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.
12. E. Hyytiä, P. Lassila, and J. Virtamo. Spatial node distribution of the random waypoint mobility model with applications. *IEEE Trans. Mob. Comp.*, 6(5):680–694, 2006.
13. D. B. Johnson and D. A. Maltz. Dynamic source routing in ad hoc wireless networks. In T. Imielinski and H. Korth, editors, *Mobile Computing*, volume 353. Kluwer Academic Publishers, 1996.
14. I. Kanj, L. Perković, and G. Xia. Computing lightweight spanners locally. In *Proc. DISC*, volume 5218 of *LNCS*, pages 365–378. Springer, 2008.
15. M. Khabbazian, S. Durocher, and A. Haghnegahdar. Bounding interference in wireless ad hoc networks with nodes in random position. *CoRR*, 1111.6689, 2011.
16. M. Khan, G. Pandurangan, and V. S. Anil Kumar. Distributed algorithms for constructing approximate minimum spanning trees in wireless sensor networks. *IEEE Trans. Parallel & Dist. Sys.*, 20(1):124–139, 2009.
17. M. Korman. Minimizing interference in ad-hoc networks with bounded communication radius. In *Proc. ISAAC*, LNCS. Springer, 2011. To appear.
18. D. Kowalski and M. Rokicki. Connectivity problem in wireless networks. In *Proc. DISC*, volume 6343 of *LNCS*, pages 344–358. Springer, 2010.
19. 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.
20. L. Li, J. Y. Halpern, P. Bahl, Y.-M. Wang, and R. Wattenhofer. A cone-based distributed topology-control algorithm for wireless multi-hop networks. *IEEE/ACM Trans. Net.*, 13(1):147–159, 2005.
21. X.-Y. Li, G. Calinescu, and P.-J. Wan. Distributed construction of a planar spanner and routing for ad hoc wireless networks. In *Proc. IEEE INFOCOM*, pages 1268–1277, 2002.
22. T. Locher, P. von Rickenbach, and R. Wattenhofer. Sensor networks continue to puzzle: Selected open problems. In *Proc. ICDCN*, volume 4904 of *LNCS*, pages 25–38. Springer, 2008.

23. A. Maheshwari, M. Smid, and N. Zeh. Low-interference networks in metric spaces with bounded doubling dimension. *Information Processing Letters*, 111(23–24):1120–1123, 2011.
24. T. Moscibroda and R. Wattenhofer. Minimizing interference in ad hoc and sensor networks. In *Proc. ACM DIALM-POMC*, pages 24–33, 2005.
25. G. Narasimhan and M. Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007.
26. M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser. CRAWDAD data set epfl/mobility (v. 2009-02-24). http://crawdad.cs.dartmouth.edu/epfl/mobility, 2009.
27. P. Santi. Topology control in wireless ad hoc and sensor networks. *ACM Comp. Surv.*, 37(2):164–194, 2005.
28. A. Sharma, N. Thakral, S. Udgata, and A. Pujari. Heuristics for minimizing interference in sensor networks. In *Proc. ICDCN*, volume 5408 of *LNCS*, pages 49–54. Springer, 2009.
29. H. Tan, T. Lou, F. Lau, Y. Wang, and S. Chen. Minimizing interference for the highway model in wireless ad-hoc and sensor networks. In *Proc. SOFSEM*, volume 6543 of *LNCS*, pages 520–532. Springer, 2011.
30. P. von Rickenbach, S. Schmid, R. Wattenhofer, and A. Zollinger. A robust interference model for wireless ad hoc networks. In *Proc. IEEE IPDPS*, pages 1–8, 2005.
31. P. von Rickenbach, R. Wattenhofer, and A. Zollinger. Algorithmic models of interference in wireless ad hoc and sensor networks. *IEEE/ACM Trans. Net.*, 17(1):172–185, 2009.