Journal of Interconnection Networks

(QUASI) SPANNERS FOR MOBILE AD HOC NETWORKS

PETRA BERENBRINK*
School of Computing Science, Simon Fraser University, Burnaby, BC, V5A 1S6 Canada
petra@sfu.ca

THOMAS FRIEDETZKY
Department of Computer Science, Science Laboratories, South Road, Durham, DH1 3LE UK
tom.friedetzky@durham.ac.uk

JÁN MAŇUCH†
School of Computing Science, Simon Fraser University, Burnaby, BC, V5A 1S6 Canada
jmanuch@sfu.ca

LADISLAV STACHO‡
Department of Mathematics, Simon Fraser University, Burnaby, BC, V5A 1S6 Canada
lstacho@sfu.ca

Received (received date)
Revised (revised date)

We introduce a notion of t-quasi-spanner as an alternative to classical t-spanners. Our motivation for quasi-spanner comes from a problem of computing a sparse backbone for ad hoc wireless networks with fixed transmission ranges. We study computational complexity of the problem of computing sparsest quasi-spanner. Then we concentrate on the case $t = 2$ and give a distributed algorithm for computing a 2-quasi-spanner with linear number of edges. Finally, we give a modification of our algorithm which performs well on a class of random ad hoc networks.

Keywords: mobile ad-hoc networks, spanners, hop distance, unit distance graph

1. Introduction

Research in ad hoc networks is very much concerned with practical applicability. Hence, it is not surprising that there is a variety of (theoretical) models for communication in such networks, defined by technological constraints. To name but a few, there are *single* vs *multiple frequencies*, *fixed* vs *variable transmission ranges*, *directed* or *undirected antennae*, *quickly moving* vs *slowly moving* vs *static nodes*,

*Research supported by NSERC of Canada Grant.
†Research supported by PIMS.
‡Research supported by NSERC of Canada Grant.

nodes with *uniform* or *non-uniform* transmission ranges. In the case of variable transmission ranges, a node can adjust the sending power to the minimum power needed to reach a given receiver. If the transmission range is fixed, nodes can only send with one fixed power (not necessarily the same for all nodes). Uniform nodes have all the same (maximum) transmission ranges, and in the case of non-uniform nodes, different nodes can have different transmission ranges.

In this work, we concentrate on the problem of establishing the communication substructure — *a spanner computation*. It is commonly accepted that a good communication structure should be sparse enough so that the routing algorithms run fast, and moreover, nodes which are connected by a path within the original graph should still be connected by a path (not much longer than the original path) in the communication structure. The later requirement has several implications. For example, it will guarantee that the communication delays in the substructure are bounded, and it also may help to minimize the transmission interferences.

We assume that nodes have a fixed and uniform transmission range. They can only send with one frequency, and the antennae are not directed. Furthermore, we assume that all nodes lie in the plane. We use the so-called *unit disk graph* as an underlying model.

Several people investigated the problem of constructing a good communication structure. One approach to a spanner computation for ad hoc networks assumes nodes with variable transmission ranges. The distance of two nodes is measured in terms of Euclidean distance instead of number of edges. The goal is to construct a *geometric spanner* (i.e., a spanner that preserves the path length in terms of Euclidean distance) with a constant stretch factor, a linear number of edges (linear in $|V|$, the number of nodes), and constant out-degree.

A lot of work has been done using the Yao graph[14] and its variations[6,12,8,5,9]. It can be shown that the Yao graph contains a constant factor geometric spanner, but some nodes can have degree of $\Omega(n)$. Some authors[6,13,8,5] convert the Yao graph into a constant degree graph, but they rely on global rankings and are not local. Wang et. al.[12] present the first local approach to compute a constant degree geometric spanner that has constant energy stretch for a special graph class. Jia et. al.[6] generalize the result to arbitrary unit disk graphs. Funke et. al.[4] use a different approach. They develop a data structure that supports so-called k-hop queries returning a path with a hop length at most k approximating the energy consumption of communication.

In all these constructions, the subgraph is sparse, however the communication delay is not bounded because the geometric spanner approximates the Euclidean distance and not the hop distance. Furthermore, these algorithms rely on knowing the geometric coordinates of all the nodes in the network. This usually implies that the construction of such a spanner must be preceded by a radio location based vertex discovery process relying on an available Geographic Positioning System (GPS) that will enable nodes to discover their coordinates.

Another possible approach is to compute so called t-spanner (i.e., a subgraph that "preserves" hop distances) of the original network for some small constant t. This subgraph would give a bounded communication delay (dependent on t). Spanners have been introduced by Peleg and Ullman[11] where they used them to synchronize asynchronous networks. A survey of results and constructibility of (sparse) spanners was given by Peleg and Schaffer[10]. The problem of finding a sparsest t-spanner was shown to be NP-complete[2] for most values of t. In case $t = 2$ (i.e., the length of a shortest path in a spanner is at most twice the length of a shortest path in the original graph) an approximation algorithm to construct a 2-spanner was given by Kortsarz and Peleg[7]. However, 2-spanners cannot have a linear number of edges in general networks[7], and moreover most known algorithms for computing a t-spanner are not distributed. Recently Dubhashi et. al.[3] presented a distributed non-constant delay $O(\log n)$-spanner algorithm.

In the context of ad hoc networks, we refer the reader to Alzoubi et. al.[1] for an overview over the work on spanners. In that paper, the authors also construct a 5-spanner in a localized manner. Their spanner is actually close to 3-spanner, since for every nodes at hop distance k in the unit distance graph, there is a path of length $3k + 2$ in the spanner. Their approach is very different from ours. First, they construct a connected dominating set in which every node has a constant degree. Then they connect nodes which are not part of the set to neighbors which are in the set. The final degrees of nodes are not bounded by any constant, and the constant presented in the paper bounding the partial degrees in the connected dominating set is huge. They do not give any bounds on the total number of edges.

A natural question arises: Does there exist a 2-spanner with a linear number of edges in *unit disk graphs*? This question seems rather difficult and we do not have any answer for it. However, we introduce a very reasonable relaxation of the 2-spanner, and show that this type of spanner has a linear number of edges.

The basic idea behind the relaxation is the following: For ad hoc networks that we consider, it completely suffices to route a message to a node that is in the transmission range of the destination node instead of routing it directly to the destination node. If a node is receiving a message that is destined for it, it can easily accept the message disregarding the fact that the sender is not its neighbor in the spanner. Indeed, consider the situation when a node u has a message for its neighbor v but v is not a neighbor of u in the spanner. Now, instead of sending the message to the first node on the path from u to v in the spanner, u can just send the message directly to v. Note that this can slightly decrease the communication load. For this to work, we need to assume that the node u does not only know the overlay network but also all nodes in its direct neighborhood. This is a reasonable assumption since in our construction, the nodes need to obtain their 2-hop neighborhoods to build the spanner.

This leads us to a definition of a *quasi-spanner* as a subgraph where for every pair (u, v) of nodes, a path from u to v in the spanner is allowed to end at an arbitrary neighbor of v (instead of v). We show that this small deviation from the original definition of spanner has fruitful consequences for constructing sparse communication topologies. Moreover, the only geometric information used in our algorithms is the adjacency relation: nodes do not need to know their geometric coordinates, neither their distances from each other, they only need to know which nodes they can communicate with. Since, in our algorithm the nodes have to pass information about their neighborhoods to the neighbors, it is necessary to assume that each node in the network has a unique id.

1.1. *The Model*

We assume that the nodes in our ad hoc network can communicate with all neighbors which are at Euclidean distance at most one, i.e., all nodes within Euclidean distance one are connected by an edge. Furthermore, nodes can only send messages on one fixed frequency. If nodes at Euclidean distance larger than one want to communicate with each other we use multi-hop routing. We use the unit disk graph model for our network.

Definition 1 *Consider a finite collection V of points in a finite region A of the 2-dimensional Euclidean space $\mathbb{R}^2$. A* unit disk graph on V *is the graph G with node set $V(G) = V$ and edge set*

$$E(G) = \{(u, v);\quad u, v \in V,\ |uv| \leq 1\}\,.$$

Definition 2 *Let $p = (u = u_1, u_2, \ldots, u_k = v)$ be a $u - v$ path in a graph G. Throughout this paper we assume that a path does not contain repeated nodes. The* length *of p is the number of edges on p, i.e., $k - 1$. The length of a shortest $u - v$ path is the* hop distance *of u and v, denoted by $\mathrm{d}_G(u, v)$.*

Note that whenever we use the term "hop distance" we count the number of edges of a shortest path.

Definition 3 *Let H be a spanning subgraph of a unit disk graph G, i.e., $V(H) = V$ and $E(H) \subseteq E(G)$. We say that two nodes $u, v \in V$ are* connected *in H if there exists a path $(u, u_1, \ldots, u_k, v)$ in H. Going a step further, we say that u and v are* quasi-connected *if there exists a path $(u, u_1, \ldots, u_k, v)$ in G with all but the last edge in $E(H)$. Such a path will be referred to as a* quasi-path *in H. Given two quasi-connected nodes u and v, their* quasi-distance $\mathrm{qd}_H(u, v)$ *is the length of a shortest $u - v$ quasi-path in H.*

Furthermore, we say that H is a spanner *(resp. a* quasi-spanner*) of G if every connected pair of nodes in G is also connected (resp. quasi-connected) in H. Finally, we say that H is a* t-spanner *of G if for every pair u, v of connected nodes in G,*

$$\mathrm{d}_H(u, v) \leq t \cdot \mathrm{d}_G(u, v)\,,$$

and that H is a t-quasi-spanner *of G if for every pair u, v of connected nodes in G,*

$$\mathrm{qd}_H(u, v) \leq t \cdot (\mathrm{d}_G(u, v) - 1) + 1\,.$$

1.2. *Our Results*

In Section 2 we will first show that the problem of determining for a given graph G and two integers t, m, whether or not G has a t-quasi-spanner with m edges is NP-complete. In Section 3 we give a distributed algorithm constructing 2-quasi-spanner of an arbitrary unit disk graph. In Subsection 3.2 we show that our quasi-spanner has linear number of edges, i.e., at most $36|V|$ edges. In Subsection 3.3 we analyze the communication complexity of our distributed algorithm, and show that overall running time is $\mathcal{O}(n)$, the time required to find a minimal dominating set of the whole unit disk graph. Finally, in Section 4 we show that in the case of randomly distributed nodes, the expected number of edges can be much smaller.

2. Complexity Results

Here we show that the problem of constructing a sparsest t-quasi-spanner in general graphs is intractable. As one would expect, the reduction is from the "sparsest t-spanner problem".

The problem of determining for a given graph G (not necessarily a unit disk graph) and two integers $t, m \geq 1$, whether G has a t-spanner with at most m edges is NP-complete[10]. In what follows we extend their result to t-quasi-spanners.

Theorem 1 *The problem of determining for a given graph G (not necessarily a unit disk graph) and two integers $t, m \geq 1$, whether G has a t-quasi-spanner with at most m edges is NP-complete for all $t \geq 2$.*

Proof. To show that the problem is in NP, we have to be able to verify that a given graph H is a t-quasi-spanner of G in polynomial time. Using Floyd-Warshall algorithm we can compute the lengths of shortest paths connecting any pair of nodes of G (resp. H) in polynomial time. For H, we have to perform an additional operation to correctly compute shortest quasi-distances: for each pair u, v, the length of a shortest quasi-path equals 1 plus the minimum hop distance from u to any neighbor of v in G. The operation takes $\mathcal{O}(|E(G)|)$ steps. Finally, for every pair of nodes we check whether the ratio between the hop distance in G and the quasi-distance in H is smaller than t. Obviously, all steps can be done in polynomial time.

For a given graph G, let G' be the graph obtained from G by adding a new node v' for every node $v \in V(G)$ and joining it to v by an edge. Let $E' = \{(v, v');\ v \in V(G)\}$.

First, we will show that for any t-quasi-spanner H' of G' with at most $m+|V(G)|$ edges, the graph $H = H'[V(G)]$ is a t-spanner of G with at most m edges. Let u and v be any two nodes in $V(G)$. Since H' is a t-quasi-spanner, there is a quasi-path from u to v' of length at most $t(\mathrm{d}_{G'}(u, v') - 1) + 1$ in H'. Since v is the unique neighbor of v' in G', the edge (v, v') is the only edge of the path not in G. Therefore, there is a path from u to v of length at most $t(\mathrm{d}_{G'}(u, v') - 1) = t \cdot \mathrm{d}_G(u, v)$ in H. This shows that H is a t-spanner of G. The existence of a quasi-path from u' to v shows that $E' \subseteq E(H')$. It follows that t-spanner H has at most m edges.

Second, we will show that for any t-spanner H of G with at most m edges, the graph $H' = H \cup E'$ is a t-quasi-spanner of G' with at most $m + |V(G)|$ edges. Consider a pair of nodes in G'. There are four cases depending on whether the first (resp. the second) node of the pair is in $V(G)$ or not.

1. The pair u, v': we have a path p from u to v of length at most $t \cdot \mathrm{d}_G(u,v)$ in H. Adding edge (v, v') to the end of p, we get a path from u to v' of length at most $t \cdot \mathrm{d}_G(u,v) + 1 = t(\mathrm{d}_{G'}(u,v') - 1) + 1$ in H'.

2-3. The pair u', v (resp. u', v'): the proof is similar to the first case.

4. The pair u, v: take a shortest path from u to v in G. Let z be the last but one node on the path. There is a path p from u to z of length at most $t \cdot \mathrm{d}_G(u,z) = t(\mathrm{d}_G(u,v) - 1)$ in H. Form a quasi-path from u to v of length at most $t(\mathrm{d}_{G'}(u,v) - 1) + 1$ in H' by appending the edge (z, v) to p.

Obviously, the t-quasi-spanner H' has at most $m + |V(G)|$ edges.

Hence, G has a t-spanner with at most m edges if and only if G' has a t-quasi-spanner with at most $m + |V(G)|$ edges. NP-completeness of the t-quasi-spanner problem follows by NP-completeness of the t-spanner problem. □

Similarly, as for spanners, the complexity of the problem of determining whether a unit disk graph has a t-quasi-spanner with at most m edges is unknown.

3. Computation of a 2-quasi-spanner

In this section, we first observe that 2-quasi-spanner cannot have all vertices of constant degree. Then we describe a distributed algorithm for constructing a spanning subgraph of a unit disk graph which has a linear number of edges.

Theorem 2 *Let u be a node of a unit disk graph G. Let d and D be the number of nodes at hop distance 1 and 2 from u in G, respectively. Let H be a 2-quasi-spanner of G. Then some node in $N = \mathrm{N}^1_G(u) \cup \mathrm{N}^2_G(u)$ has degree at least $\sqrt{\frac{D}{d}}$ in H.*

Proof. Let Δ be the maximum degree in H of nodes in N. For every node $x \in \mathrm{N}^2_G(u)$, there is a path in G of length 2 from x to u. Since H is a 2-quasi-spanner of G, for every such a node, there is a node $y \in \mathrm{N}^1_G(u)$ such that there exists an $x - y$ path in H of a length at most two. Since, the degree of every node in N is at most Δ in H, there is at most $d \cdot \Delta$ paths of length 1 and at most $d \cdot \Delta \cdot (\Delta - 1)$ paths of length 2 ending at a $y \in \mathrm{N}^1_G(u)$. Therefore, we must have

$$d \cdot \Delta + d \cdot \Delta \cdot (\Delta - 1) \geq D\,.$$

Hence $\Delta \geq \sqrt{D/d}$ as claimed. □

Note that it is easy to construct a unit disk graph containing a node for which the D/d ratio is as high as $\mathcal{O}(n)$ where n is the number of nodes.

3.1. *The Algorithm*

The main idea of the algorithm is the following: For every node we select a constant number of nodes within two-hops from it covering all other nodes in the circle with radius 2 centered at the node. As we will show the constructed graph will have at most $36|V|$ edges and will be a 2-quasi-spanner.

Definition 4 *The* k-neighborhood *of a node* v*, denoted as* $\mathrm{N}_G^k(v)$*, is the set of all nodes at hop distance* k *from* v*. Similarly, we denote the set of all nodes at hop distance at most* k *from* v *as* $\mathrm{N}_G^{\leqslant k}(v)$*.*

Definition 5 *A subset* C *of a set* S *is called a* covering *of* S *if for every point* $s \in S$ *there is a point* $c \in C$ *such that* $|sc| \leq 1$*.*

Given a unit disk graph G *and a node* $u \in V$*, a covering of* $\mathrm{N}_G^2(u)$ *will be denoted as* $T(u)$*. A covering* $T(u)$ *is* minimal *if no other covering of* $\mathrm{N}_G^2(u)$ *is a subset of* $T(u)$*.*

Note that a covering $T(u)$ is a dominating set of the subgraph of G induced by $\mathrm{N}_G^2(u)$. The following is a simple greedy procedure constructing a minimal covering.

Procedure `Min-Covering`(u, G)
Output: minimal covering $T(u)$ of $\mathrm{N}_G^2(u)$

$N := \mathrm{N}_G^2(u)$
$T(u) = \emptyset$
while $N \neq \emptyset$ **do**
  choose $v \in N$
  $T(u) := T(u) \cup \{v\}$
  for every $x \in \mathrm{N}_G^1(v) \cap N$ **do**
    $N := N - \{x\}$
  end for
end while

An example of a minimal covering $T(u)$ of the set $\mathrm{N}_G^2(u)$ is depicted in Figure 1. The size of $T(u)$ is 6 and 11 edges are needed to connect these nodes to the node u. In the next section, we will show that any minimal covering has at most 18 nodes, and hence, requires at most 36 edges to be added to the quasi-spanner. Furthermore, assuming that the collection of points V is "dense" in A, we will substantially improve both bounds.

We close this section with an algorithm for constructing a 2-quasi-spanner. The algorithm is very simple — every node will build a minimal covering around it. We will analyze the algorithm in the following section.

Algorithm 1 (2-quasi-spanner)
Input: *Unit disk graph* $G = (V, E(G))$
Output: *Spanning subgraph* $K = (V, E(K))$ *of* G

for *each node* u ***do***
  $E(u) := \emptyset$

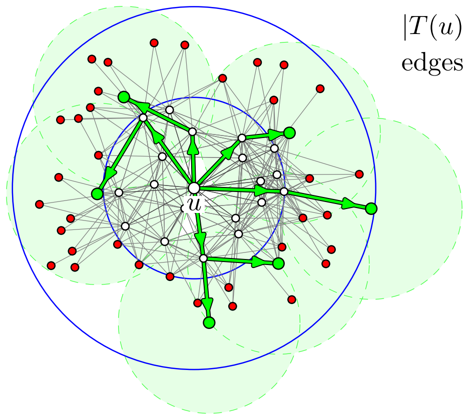

Figure 1: The set $T(u)$ (light gray/green big points) is a covering of the set $\mathrm{N}^2_G(u)$ (dark/red points).

$T(u) :=$ `Min-Covering`(u, G)
  for *every* $v \in T(u)$ ***do***
    find a $w \in \mathrm{N}^1_G(u) \cap \mathrm{N}^1_G(v)$
    $E(u) := E(u) \cup \{(u, w),\ (w, v)\}$
  end for
end for
$E(H) := \cup_{u \in V} E(u)$

An example of a unit disk graph and the corresponding 2-quasi-spanner constructed by our algorithm is shown in Figure 2.

3.2. *Analysis of Algorithm 1*

In this section, we show that Algorithm 1 indeed produces a sparse 2-quasi-spanner. This is stated in the following theorem.

Theorem 3 *The graph K produced by Algorithm 1 is a 2-quasi-spanner of G and has at most $36|V|$ edges, i.e., the average degree of K is at most 72.*

The proof of the theorem is split into the following four lemmas. In the first lemma, we show that if a subgraph S of a unit disk graph G has the property that for every node $u \in V$ there exists a covering of $\mathrm{N}^2_G(u)$ such that each node of this covering is at hop distance two from the node u in the graph S, then S is a 2-quasi-spanner of G. It is straightforward to see that Algorithm 1 produces a subgraph K with this property, hence K is indeed a 2-quasi-spanner of G.

Lemma 1 *If a subgraph S of a unit disk graph G has the property that for every node $u \in V$ there exists a covering $T(u)$ of $\mathrm{N}^2_G(u)$ contained in $\mathrm{N}^2_S(u)$, then S is a 2-quasi-spanner of G.*

Proof. Let $(p_1, \ldots, p_k)$ be a shortest path connecting p_1 and p_k in S. We will

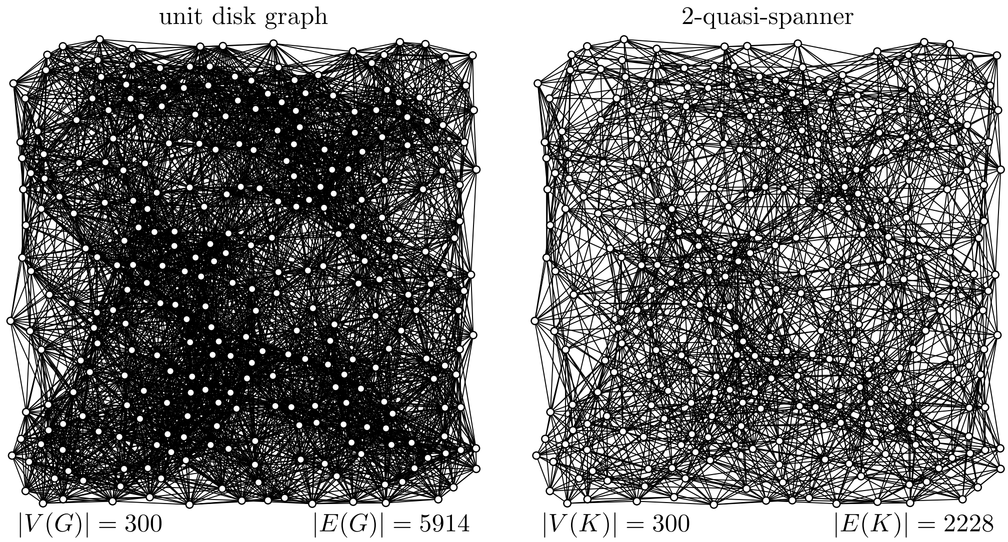

Figure 2: An example of a unit disk graph with 300 nodes and the corresponding 2-quasi-spanner K constructed using Algorithm 1.

construct a trail

$$(r_1, q_2, r_2, \ldots, r_{k-2}, q_{k-1}, r_{k-1}, p_k)$$

quasi-connecting $p_1 = r_1$ and p_k in S such that

$$r_i \text{ is at hop distance at most 1 from } p_{i+1} \text{ for all } i = 1, \ldots, k-1.$$

This property will guarantee that the last edge (r_{k-1}, p_k) is in the unit disk graph which is sufficient by the definition of quasi-path.

Also note that, as usual, in a trail both nodes and edges can repeat; in addition we allow two consecutive nodes in a trail to be the same (i.e., “virtual loops”). Obviously, existence of a trail connecting u and v implies existence of a path connecting u and v of length at most the length of the trail.

Such an alternative trail is constructed inductively. Since $r_1 = p_1$, the property is satisfied for $i = 1$. Assume that $|r_{i-1}p_i| \leq 1$. Then (r_{i-1}, p_i, p_{i+1}) is a path of length 2 in G, i.e., $p_{i+1} \in \mathrm{N}_G^{\leqslant 2}(r_{i-1})$. It follows from the definition of the covering that every node in $\mathrm{N}_G^{\leqslant 2}(r_{i-1})$ is within hop distance one either from r_{i-1} or from some node in $T(r_{i-1}) \subseteq \mathrm{N}_S^2(r_{i-1})$. Hence, there exists $w \in \mathrm{N}_S^2(r_{i-1}) \cup \{r_{i-1}\}$ such that $|p_{i+1}w| \leq 1$. If $w = r_{i-1}$ then we can set $r_i = q_i = r_{i-1}$ and we are done. Otherwise, we have a path (r_{i-1}, v, w) in S, and we set $q_i = v$ and $r_i = w$. See Figure 3.

Counting the nodes on the constructed trail shows that there is a $p_1 - p_k$ quasi-path of length at most $2k - 1$ as required. □

Next, we find an upper bound for the number of edges of 2-quasi-spanner K. We will show that this number is linear in the number of nodes. This will follow

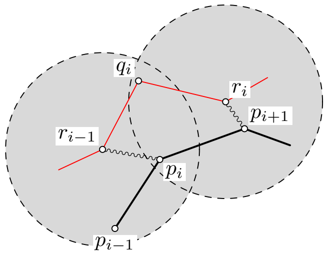

Figure 3: Thick edges indicate the original path $(\dots, p_i, p_{i+1}, \dots)$ in G and thin edges indicate the approximating trail $(\dots, r_{i-1}, q_i, r_i, \dots)$ in S. By inductive hypothesis, the node p_i is covered by r_{i-1}. Since $|p_{i+1}r_{i-1}| \leq 2$, there exists a node r_i covering p_{i+1}.

instantly if we show that for every node $u \in V$, the procedure `Min-Covering`(u, G) will produce $T(u)$ of size at most 18, and hence, Algorithm 1 will add at most 36 edges in step 6 for node u. This follows from the next lemma.

Lemma 2 *Let T be a set of points in the ring $R_u(1,2)$ such that the Euclidean distance between any two points in T is greater than 1. Then $|T| \leq 18$.*

The proof of the lemma is in Appendix A. Note that Lemma 2 cannot be improved as there exists a configuration of 18 points in $R_u(1,2)$ such that their mutual Euclidean distances are greater than one, cf. Figure 4. Hence all 18 points in the second neighborhood of the central node u have to be selected for the covering set. However, it is enough (and necessary) to add 30 edges to the graph K to get these points connected to u.

Proof of Theorem 3. It is enough to show that 2-quasi-spanner produced by Algorithm 1 will not have too many edges. Lemma 2 implies that any minimal covering $T(u)$ of $\mathrm{N}^2_G(u)$ has at most 18 points. In line 6 of the algorithm, at most two edges are added to the 2-quasi-spanner for each of the 18 points of the covering. Hence, the bound $36|V|$ on number of edges follows directly. □

Figure 2 shows an example of a unit disk graph and the corresponding 2-quasi-spanner constructed by Algorithm 1.

3.3. *Communication Complexity*

In this subsection we will analyze various aspects of communication complexity of our algorithm: the number and size of all messages and the total time needed to set up the 2-quasi-spanner. We will also compare complexities of our construction of 2-quasi-spanners and the construction of classical 5-spanners by Alzoubi et. al.[1].

Theorem 4 *Let G be a unit disk graph with n nodes and the maximal degree Δ. Assume that broadcasting of one message (of size $\mathcal{O}(\Delta \log n)$) takes time t. A 2-quasi-spanner can be constructed in time $\mathcal{O}(t)$ using $\mathcal{O}(n)$ messages with a total size $\mathcal{O}(\Delta n \log n)$ by Algorithm 1.*

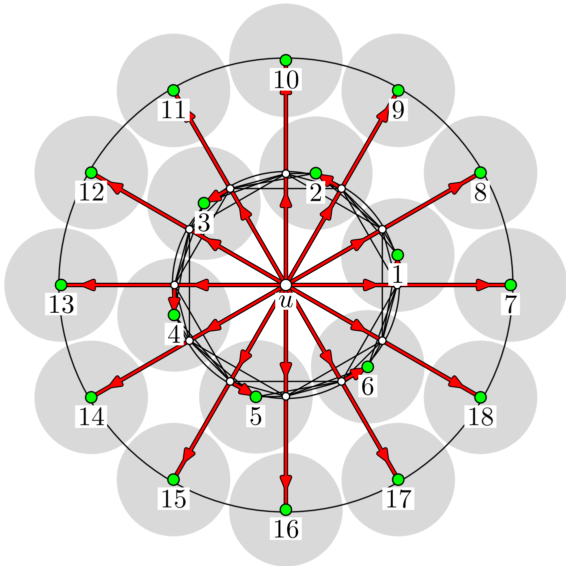

Figure 4: An example showing the lower bound 18 for the minimal cardinality of the covering set $T(u)$ is sharp. Nodes 1 to 6 are placed evenly on the circle of radius $1+\epsilon$ centered at u starting at angle 15. Similarly, nodes 7 to 18 are placed evenly on the circle of radius $2-\epsilon$ centered at u starting at angle 0. This placement guarantees that these 18 nodes are in the ring $R_u(1,2)$ and there are no edges in unit disk graph joining them. To complete the construction we add 12 new nodes (depicted as small white circles) placed evenly on the circle of radius 1 centered at u starting at angle 0. Note that each white node covers only one node 7 to 18. It follows that all 18 named vertices are in the cover and to connect them to u we need at least 30 edges.

Note that the total amount of information needed to be broadcast by the algorithm can be as high as $\mathcal{O}(n^2 \log n)$ and the whole process will take $\mathcal{O}(n)$ time in dense networks.

Proof. Each node u has to find a minimal covering of its second neighborhood $\mathrm{N}^2_G(u)$. The following protocol achieves the bounds on communication complexities required by Theorem 4:

1. Each node broadcasts its id. The size of each message is $\mathcal{O}(\log n)$. As a result each node knows its 1-neighborhood.

2. Each node broadcasts its 1-neighborhood. The size of each message is $\mathcal{O}(\Delta \log n)$. As a result each node knows its 2-neighborhood.

3. Now, each node is ready to execute procedure `Min-Covering`. However, to perform lines 6–8 of the procedure, more communication is required. Each time $v \in N$ is selected as a member of covering $T(u)$, u sends a request to a node connecting u and v for a 1-neighborhood of v. When the connecting

node sends this information back to u, u is able to perform lines 6–8 of the procedure. Each such request will require 2 messages of size $\mathcal{O}(\Delta \log n)$. By Lemma 2, there will be at most 18 such requests for node u.

Simple counting of the number and size of messages gives the required bounds. Note also that the computations of coverings for different nodes are independent and can run in parallel. Hence, the construction of a 2-quasi-spanner can be completed in time $\mathcal{O}(t)$. □

Finally, let us compare the complexities of Algorithm 1 and the procedure constructing a 5-spanner by Alzoubi et. al.[1]. The main step of their algorithm is to build a dominating set of the whole graph G. This requires $\mathcal{O}(n)$ messages, each of size $\mathcal{O}(\log n)$. However, selecting dominators cannot be always done in constant time. In the worst case scenario, each node on a path of G might have to wait until it is decided whether its neighbor on the path is in the dominating set. Hence, the computation can require $\Theta(d(G))$ time, where $d(G)$ is the diameter of G. In the worst case, for instance, a path with nodes numbered consecutively, the construction will take linear time.

4. A 2-quasi-spanner in a Random Unit Disk Graph

Theorem 3 gives an upper bound $36|V|$ for the number of edges of the 2-quasi-spanner constructed by Algorithm 1. In this section we improve the above upper bound for random mobile networks. We will consider random mobile networks which we model as random unit disk graphs in a certain region $\mathcal{A}$. We will show that if the node density of the graph is large enough, then we can improve the above upper bound by more than half with high probability.

Let A denote the area of the region $\mathcal{A}$. A random unit disk graph $G = G(\mathcal{A}, n)$ is a unit disk graph obtained by placing n nodes uniformly and independently at random into $\mathcal{A}$. Let $\rho(A, n) = n\pi/A$ be the *density* of this random unit disk graph, i.e., the expected number of nodes that fall into a fixed circle with radius 1. This density happens to be equal to the expected number of nodes a fixed node of G can reach.

Observe that the covering depicted in Figure 4 is the worst possible since the chosen nodes of the covering are close to the borders of the ring $R_u(1, 2)$. Hence, they cover only a small part of the ring. If we choose the covering nodes from the inside of the ring, each of them will cover substantially larger portion of the ring. Hence, a small number of such nodes will cover $R_u(1, 2)$. This, in turn, means that all nodes in $\mathrm{N}^2_G(u)$ are covered. In this section we assume that nodes are placed in $\mathcal{A}$ with a uniform random distribution with high density. For this reason it is reasonable to assume that we can find enough nodes inside of the ring. The following definition gives us a method how to find such an efficient covering.

Definition 6 *A δ-*wide configuration $\mathcal{C}_u(\delta)$ of *u is a collection of 15 circles C_1, C_2, …, C_{10} and $\bar{C}_1$, $\bar{C}_2$, …, $\bar{C}_5$ with radii δ. A* realization $(x_1, \ldots, x_{10}, y_1, \ldots, y_5)$ *of*

$\mathcal{C}_u(\delta)$ in $S \subseteq \mathbb{R}^2$ is a collection of 15 points in S such that $x_1 \in C_1$, $x_2 \in C_2$, …, $x_{10} \in C_{10}$ and $y_1 \in \bar{C}_1$, $y_2 \in \bar{C}_2$, …, $y_5 \in \bar{C}_5$. We say that $\mathcal{C}_u(\delta)$ is overcasting *if for any realization $(x_1, \ldots, x_{10}, y_1, \ldots, y_5)$ in $\mathbb{R}^2$,*

(a) points $x_1, x_2, \ldots, x_{10}$ constitute a covering of $R_u(1, 2)$,

(b) points $y_1, y_2, \ldots, y_5$ lie in the unit circle centered at u, and

(c) for all $i = 1, 2, \ldots, 5$, we have $|y_i x_{2i-1}| \leq 1$ and $|y_i x_{2i}| \leq 1$.

Note that any realization of a configuration $\mathcal{C}_u(\delta)$ is a covering of $\mathrm{N}_G^2(u)$ with 10 points (instead of 18) and uses 15 edges to connect them to u (instead of 36). We will prove the existence of an overcasting δ-wide configuration of u later. The following algorithm is an extension of Algorithm 1. For each node $u \in V$ it tries to find a realization of a configuration $\mathcal{C}_u(\delta)$. If such a realization of $\mathcal{C}_u(\delta)$ exists it will add at most 15 edges into the constructed spanner, otherwise it will cover u using the method of Algorithm 1 (adding at most 36 edges).

Algorithm 2 (Light 2-quasi-spanner)
Input: *Random unit disk graph $G = (V, E(G))$*
Output: *Spanning subgraph $R = (V, E(R))$ of G*

for *each node u* ***do***
  $E(u) := \emptyset$
  take any overcasting δ-configuration $\mathcal{C}_u(\delta)$
  if *exists a realization $(x_1, \ldots, x_{10}, y_1, \ldots, y_5)$ of $\mathcal{C}_u(\delta)$* ***then***
    for $i := 1$ *upto* 5 ***do***
      $E(u) := E(u) \cup \{(u, y_i), (y_i, x_{2i-1}), (y_i, x_{2i})\}$
    end for
  else
    $T(u) :=$ `Min-Covering`(u, G)
    for *every $v \in T(u)$* ***do***
      find a $w \in \mathrm{N}_G^1(u) \cap \mathrm{N}_G^1(v)$
      $E(u) := E(u) \cup \{(u, w),\ (w, v)\}$
    end for
  end if
end for
$E(R) := \cup_{u \in V} E(u)$

It is straightforward to see that the subgraph R produced by Algorithm 2 satisfies the assumption of Lemma 1, hence it is a 2-quasi-spanner of G.

In the following, we show the existence of an overcasting δ-wide configuration $\mathcal{C}_u(\delta)$ with $\delta = \frac{4}{25}$ which is done in Lemma 3.

Lemma 3 *For any point u, there exists an overcasting $\frac{4}{25}$-wide configuration $\mathcal{C}_u^*(\frac{4}{25})$.*

Proof. Let $\mathcal{C}^*_u(\frac{4}{25}) = (C_1, \ldots, C_{10}, \bar{C}_1, \ldots, \bar{C}_5)$, where the centers c_1 of C_1, …, c_{10} of C_{10}, and $\bar{c}_1$ of $\bar{C}_1$, …, $\bar{c}_5$ of $\bar{C}_5$ will be determined such that the conditions (a)–(c) of Definition 6 are satisfied, i.e., $\mathcal{C}^*_u(\frac{4}{25})$ will be an overcasting $\frac{4}{25}$-wide configuration.

We set up the centers as follows:

- Points $\bar{c}_1, \ldots, \bar{c}_5$ are placed evenly on a circle centered at u with radius $\frac{21}{25}$. This guarantees that the condition (b) is satisfied. Note that placing one of $\bar{c}_i$'s specifies exact positions of the others.

- Points $c_1, \ldots, c_{10}$ are placed evenly on a circle centered at u with radius a such that for every point $\bar{c}_i$, the Euclidean distances to c_{2i-1} and c_{2i} are the same and equal to $\frac{17}{25}$. Since the Euclidean distance between $\bar{c}_i$ and c_{2i-1} (resp. c_{2i}) is exactly $1 - 2 \cdot \frac{4}{25}$, the Euclidean distance between any point $y_i \in \bar{C}_i$ and any point $x_{2i-1} \in C_{2i-1}$ (resp. $x_{2i} \in C_{2i}$) is at most 1. Hence, the condition (c) is satisfied. Note that given positions of $\bar{c}_i$'s the radius a and the positions of all c_i's are fixed.

Now it suffices to verify the condition (a). For this we need to calculate the radius a, which can be done easily using Fact 1 on triangle $\triangle u\bar{c}_i c_{2i}$, cf. Figure 5. By our set-up,

$$|u\bar{c}_i| = \tfrac{21}{25}, \quad |\bar{c}_i c_{2i}| = \tfrac{17}{25}, \quad |c_{2i}u| = a \quad \text{and,} \quad |\measuredangle c_{2i} u \bar{c}_i| = \tfrac{2\pi}{20}$$

If we plug in the above value into Fact 1 we obtain a quadratic equation in a having only feasible solution (between 1 and 2):

$$a = \tfrac{21}{25}\cos(\tfrac{\pi}{10}) + \tfrac{1}{25} \cdot \sqrt{21^2\cos^2(\tfrac{\pi}{10}) - 152}\,.$$

Fix any realization $(x_1, \ldots, x_{10}, y_1, \ldots, y_5)$ of $\mathcal{C}^*_u(\frac{4}{25})$. We will prove that unit circles centered at x_j and x_{j+1} cover whole sector $S_j = \measuredangle x_j u x_{j+1} \cap R_u(1, 2)$, hence the condition (a) holds.

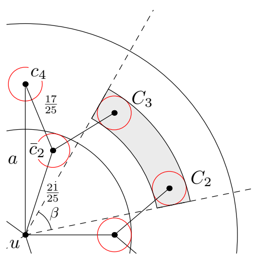

Figure 5: Construction of an overcasting $\frac{4}{25}$-wide configuration $\mathcal{C}^*_u(\frac{4}{25})$. The grey region depicts the sector $\measuredangle x_2 u x_3 \cap R_u(a - \frac{4}{25}, a + \frac{4}{25})$.

Clearly, the points x_j and x_{j+1} belong to the ring $R_u(a - \frac{4}{25}, a + \frac{4}{25})$. The angle $|\measuredangle x_j u c_j| \le \arcsin \frac{4}{25a}$, where the equality holds when x_j is a tangent point of C_j

from u, cf. Figure 5. It follows that the angle $|\measuredangle x_j u x_{j+1}| \leq \frac{\pi}{5} + 2\arcsin\frac{4}{25a} = \beta$. By Lemma A.2, any point z in sector S_j not covered by x_j or x_{j+1} must satisfies either

$$|uz| < \omega(a + \tfrac{4}{25}, \beta) < 1 \quad \text{or} \quad |uz| > \omega'(a - \tfrac{4}{25}, \beta) > 2$$

which is a contradiction. □

For every node $u \in V$, fix an overcasting $\frac{4}{25}$-wide configuration $\mathcal{C}^*_u(\frac{4}{25})$ constructed in Lemma 3. We say that a node u is *good* if there is a realization of $\mathcal{C}^*_u(\frac{4}{25})$ in V, otherwise the node u is *bad*.

In the following lemma we state results regarding the probability of the event that nodes $u \in V$ are good.

Clearly, almost all statements regarding "goodness" or "badness" of nodes depend on the *shape* of the region $\mathcal{A}$ that we are looking at, as well as the *position* of the node within $\mathcal{A}$. If, for instance, a node happens to be very close to the border, than this will certainly influence the probability for this node to have a certain number of neighbors.

To overcome this problem, we define *core nodes* and *outer nodes*. We define the *core* of the region to be the sub-region such that all points have a Euclidean distance of at least $a + 4/25$ from the boundary, with $a \approx 1.52$ as defined in the proof of Lemma 3. Note that the core may be empty, depending on the shape of $\mathcal{A}$. Now we say that all nodes within the core are *core nodes*.

Lemma 4

1. *Fix some core node u. For $n \geq 43$ and any $\alpha \geq 1$, if $\rho \geq 40\ln(15\alpha)$ then with probability at least $1 - 1/\alpha$, the node u is good.*

2. *For $n \geq 43$ and any $\beta \geq 1$, if $\rho \geq 40\ln(15\beta)$, then the expected number of bad core nodes is at most n/β.*

3. *If $n \geq 43$ and $\rho \geq 80\ln n + O(1)$, then there is no bad core node with probability at least $1 - 1/n$.*

Proof. For node u, let $C_{u,1}, C_{u,2}, \ldots, C_{u,15}$ denote the 15 circles of radius $4/25$, in $\mathcal{C}^*_u(\frac{4}{25})$. Let c denote the area of a circle of radius $4/25$, i.e., $c = (4/25)^2 \cdot \pi \approx 0.08$. Fix some $C_{u,i}$. Let

$$q = \text{Prob(no node falls into } C_{u,i}).$$

Note that q does not depend upon u or i. Clearly,

$$q \leq \left(1 - \frac{c}{A}\right)^{n-1} \leq e^{-c(n-1)/A}.$$

The probability that (for fixed node u) there *exists* an empty $C_{u,i}$ for $i = 1, \ldots, 15$ can now be upper-bounded by $15q \leq 15e^{-c(n-1)/A}$. We can now prove the first statement of the lemma. We require $15q < 1/\alpha$ with α from the statement of the lemma. Consider the following sequence of equivalent inequalities.

$$15e^{-c(n-1)/A} < 1/\alpha \Leftrightarrow e^{-c(n-1)/A} < 1/15\alpha \Leftrightarrow -c(n-1)/A < -\ln(15\alpha) \Leftrightarrow A < \frac{c(n-1)}{\ln(15\alpha)}.$$

Using $\rho = \pi n/A$ together with this upper bound on A, we obtain as sufficient condition

$$\rho = \frac{\pi n}{A} > \frac{\pi n \ln(15\alpha)}{c(n-1)} = \frac{\pi n \ln(15\alpha)}{(4/25)^2\pi(n-1)} = \frac{625n\ln(15\alpha)}{16(n-1)}.$$

Since $625n/16(n-1) \leq 40$ for $n \geq 43$, this condition is clearly satisfied whenever $\rho > 40\ln(15\alpha)$.

As for the second statement, let B denote the number of bad core nodes. By the above, $E[B] \leq 15ne^{-c(n-1)/A}$. We require this quantity to be at most n/β. Consider the following equivalent inequalities.

$$\begin{aligned}
& 15ne^{-c(n-1)/A} \leq n/\beta \\
\Leftrightarrow\ & 15e^{-c(n-1)/A} \leq 1/\beta \\
\Leftrightarrow\ & e^{-c(n-1)/A} \leq 1/15\beta \\
\Leftrightarrow\ & -c(n-1)/A \leq \ln(1/15\beta) = -\ln(15\beta) \\
\Leftrightarrow\ & c(n-1)/A \geq \ln(15\beta) \\
\Leftrightarrow\ & A \leq \frac{c(n-1)}{\ln(15\beta)}
\end{aligned}$$

Equivalently,

$$\rho = \frac{n\pi}{A} \geq \frac{\pi n \ln(15\beta)}{c(n-1)} = \frac{\pi n \ln(15\beta)}{(4/25)^2 \cdot \pi(n-1)} = \frac{625n\ln(15\beta)}{16(n-1)}.$$

Again using the fact that $n \geq 43$ and thus $625n/16(n-1) \leq 40$, we can see that choosing $\rho \geq 40\ln(15\beta)$ is sufficient.

The third statement is a simple corollary. □

Note 1 *Statement 1 says that we can already achieve constant failure probability for a fixed core node u if the density* ρ *is a (sufficiently large) constant, i.e., by requiring that the expected number of nodes that fall into a unit circle is constant. Furthermore, we can achieve polynomially (in n) small failure probability by choosing* $\rho = \Theta(\log n)$*. Furthermore, statement 2 says that choosing* $\rho = \Theta(\log n)$ *results in a constant expected number of bad core nodes.*

We now apply the previous lemma to obtain results for a few, selected special cases. The proof of the theorem is a simple application of Lemma 4, together with an estimation of the number of core nodes of the area.

Theorem 5 *If $n \geq 43$, and the region $\mathcal{A}$ is a circle of radius r with $4 \leq r \leq \sqrt{n/81 \ln n}$, then the graph R produced by Algorithm 2 is a 2-quasi-spanner of G, and with probability at least $1 - o(1)$, has at most $15n$ edges, i.e., the average degree of R is at most 30.*

Proof. *(i)* By definition, $r \leq \sqrt{n/81 \ln n}$ implies $\rho = n/r^2 \geq 81 \ln n$, and thus we can use the third statement of Lemma 3 directly, showing that w.h.p. there is no bad core node.

(ii) Furthermore, $r \geq 2$ implies that there is indeed a core (which was defined to be the inner region in which all points have a Euclidean distance of at least $a + 4/25 \approx 1.68$ from the boundary). For $1 \leq i \leq n$, let $X_i = 1$ if the i-th node is a core node, and let $X_i = 0$ otherwise. Clearly,

$$\begin{aligned} \text{Prob}(X_i = 1) &= n \cdot \frac{(r - (a + 4/25))^2}{r^2} = n \cdot \frac{r^2 - 2r(a + 4/25) + (a + 4/25)^2}{r^2} \\ &> n \cdot \left(1 - \frac{2(a + 4/25)}{r}\right) > n \cdot \left(1 - \frac{1}{r}\right). \end{aligned}$$

Let $X = \sum_{i=1}^{n} X_i$. Now,

$$E[X] = \sum_{i=1}^{n} \text{Prob}(X_i = 1) > n \cdot \left(1 - \frac{1}{r}\right) \geq n/4$$

whenever $r \geq 4$. A straightforward application of Chernov bounds yields that $X > n/2$ with extremely high probability in this case. *(i)* and *(ii)* together yield the statement. □

Conclusions

In this paper we introduced the new notion of t-quasi spanners as an alternative to classical t-spanners. A quasi-spanner is a subgraph where for every pair of nodes (u, v) there exists a path from u to a direct neighbor of v using only edges of the spanner. The basic idea behind the relaxation is the following: For ad hoc networks that we consider, it completely suffices to route a message to a node that is in the transmission range of the destination node instead of routing it directly to the destination node. If a node is receiving a message that is destined for it, it can easily accept the message disregarding the fact that the sender is not its neighbor in the spanner.

We first showed that the problem of determining for a given (general) graph G and two integers t, m, whether or not G has a t-quasi-spanner with m edges is NP-complete. Then we presented a distributed algorithm constructing 2-quasi-spanner of an arbitrary unit disk graph. Our quasi-spanner has linear number of edges, i.e., $36|V|$. Finally we showed that in the case of randomly and distributed nodes the expected number of edges can be much smaller.

References

1. K. Alzoubi, X-Y. Li, Y. Wang, P-J. Wan, O. Frieder, "Geometric spanners for wireless ad hoc networks," *IEEE Trans. on parallel and distributed systems* **14**(5):1–13, 2003.
2. L. Cai, "NP-Completeness of Minimum Spanner Problems," *Discrete Applied Mathematics* **48**:187-194, 1994.
3. D. Dubhashi, A. Mei, A. Panconesi, J. Radhakrishnan, and A. Srinivasan, "Fast Distributed Algorithms for (Weakly) Connected Dominating Sets and Linear-Size Skeletons," *Proc. of 14th ACM-SIAM Symposium on Discrete Algorithms (SODA 03)*, 2003, pp. 717–724.
4. S. Funke, D. Matijevic, P. Sanders, "Approximating Energy Efficient Paths in Wireless Multi-hop Networks," *Proc. of 11th Annual European Symposium on Algorithms (ESA)*, LNCS **2832**, 2003, pp. 230–241.
5. M.Grünewald, T. Lukovski, C. Schindelhauer, and K. Volbert, "Distributed Maintenance of Resource Efficient Wireless Network Topologies," *Proc. of 8th International Euro-Par*, 2002, pp. 935–946.
6. L. Jia, R. Rajaraman, and C. Scheideler, "On local Algorithms for Topology Control and Routing in Ad Hoc Networks," *Proc. of 15th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, 2003, pp. 220–229.
7. G. Kortsarz, and D. Peleg, "Generating sparse 2-spanners," *J. Alg.* **17**:222–236, 1994.
8. L.Li, J. Halpern, V. Bahl, Y.-W. Wang, R. Wattenhofer, "Analysis of a Cone-based Distributed Topology Control Algorithm for Wireless Multi-Hop Networks," *Proc. of 21st ACM Symposium on Principles of Distributed Computing (PODC)*, 2002, pp. 264–273.
9. X.-Y. Li, P.J. Wan, Y. Wang, and O. Frieder, "Sparse Power Efficient Topology for Wireless Networks," *Journal of Parallel and Distributed Computing*, to appear.
10. D. Peleg, and A.A. Schaffer, "Graph spanners," *J. Graph Theory* **13**:99–116, 1989.
11. D. Peleg, and J. D. Ullman, "An Optimal Synchronizer for the Hypercube," *SIAM J. Computing* **18**:740–747, 1989.
12. Y.-W. Wang, X.-Y. Li and O. Frieder, "Distributed Spanner with Bounded Degree for Wireless Ad Hoc Networks," *International Journal on Foundations of Computer Science* **14**(2):183–200, 2003.
13. R. Wattenhofer, L. Li, P.Bahl, and Y.M. Wang, "Distributed Topology Control for Power Efficient Operation in Multihop Wireless Ad Hoc Networks," *Proc. of 20th IEEE INFOCOM*, 2001, pp. 1388–1397.
14. A. Yao, "On Constructing Minimum Spanning Trees in k-Dimensional Spaces and related Problems," *SIAM Journal on Computing* **11**(4):721-736, 1982.

Appendix A

To prove Lemma 2 we will need the following several results. The first one is the fundamental Law of Cosines:

Fact 1 *Consider a triangle with sides a, b, c. Let α be the angle opposite to the side a. Then*

$$\cos \alpha = \frac{b^2 + c^2 - a^2}{2bc} .$$

Obviously, the covering $T(u)$ of u produced by procedure `Min-Covering`(u, G) is minimal, and the points in $T(u)$ are pairwise at Euclidean distance greater than one. If we knew the minimal angle between any two such points, then we would

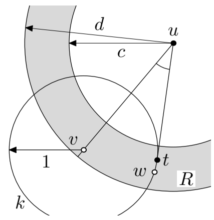

Figure A.1: Two points v and w with $|vw| = 1$ in the ring $R_u(c, d)$.

have an immediate upper bound on the size of $T(u)$, cf. Corollary A.1. Next lemma gives a lower bound on this angle.

Lemma A.1 *Let $R_u(c, d)$ be a ring between two circles with radii c and d, $1 \le c < d \le 2$, with a common center u. Let $\rho(c, d)$ be the minimal angle $|\measuredangle vuw|$ over all points $v, w \in R_u(c, d)$ with $|vw| \ge 1$. Then*

$$\rho(c,d) = \begin{cases} \arccos 1 - \frac{1}{2d^2}, & \text{if } c \ge d - \frac{1}{d}, \\ \arccos \frac{c^2+d^2-1}{2cd}, & \text{otherwise,} \end{cases} \tag{A.1}$$

and if $|vw| > 1$ then $|\measuredangle vuw| > \rho(c, d)$.
Furthermore, for a fixed c, the function $\rho(c, d)$ is decreasing on interval $\langle c, 2\rangle$.

Proof. Consider $v, w \in R_u(c, d)$ such that $|vw| \ge 1$ and the angle $|\measuredangle vuw| = \rho(c, d)$. Let k be the unit circle around v, see Figure A.1. If w does not lie on k then, by replacing w by the intersection point of k and the line vw, we would decrease the angle $|\measuredangle vuw|$, a contradiction. Hence, we may assume that w lies on k. Moreover, this also implies the second part of the lemma, i.e., if $|vw| > 1$ then $|\measuredangle vuw| > \rho(c, d)$.

Let t be the tangent point on k from u closer to w. By moving w along k away from t, we are decreasing the angle $|\measuredangle vuw|$. We can move w only until we hit either inner or outer circle of the ring. Therefore, necessarily w (and by symmetry also v) lies either on the inner or the outer circle. Hence, there are three possibilities:

- both v and w lie on the inner circle, i.e., by Fact 1, $\cos |\measuredangle vuw| = \frac{2c^2-1}{2c^2}$;
- v and w lie on different circles, i.e., $\cos |\measuredangle vuw| = \frac{c^2+d^2-1}{2cd}$; and
- both v and w lie on the outer circle, i.e., $\cos |\measuredangle vuw| = \frac{2d^2-1}{2d^2}$.

Thus, we obtain

$$\begin{aligned} \rho(c,d) &= \arccos \max \left\{ \frac{2c^2-1}{2c^2}, \frac{c^2+d^2-1}{2cd}, \frac{2d^2-1}{2d^2} \right\} \\ &= \arccos \max \left\{ \frac{c^2+d^2-1}{2cd}, 1 - \frac{1}{2d^2} \right\} = \begin{cases} \arccos 1 - \frac{1}{2d^2}, & \text{if } c \ge d - \frac{1}{d}, \\ \arccos \frac{c^2+d^2-1}{2cd}, & \text{otherwise.} \end{cases} \end{aligned}$$

The second part of the lemma is straightforward. □

Corollary A.1 *If T is a set of points in $R_u(c,d)$ such that the Euclidean distance of any two points in T is greater than 1 then $|T| < \frac{2\pi}{\rho(c,d)}$.*

Proof. Consider a set T of points in $R_u(c,d)$ such that the Euclidean distance of any two points in T is greater than 1. Let k be the cardinality of the set T. Pick a point $t_1 \in T$ and label the remaining points $t_2, \ldots, t_k$ of T so that

$$|\measuredangle t_1 u t_2| \leq |\measuredangle t_1 u t_3| \leq \cdots \leq |\measuredangle t_1 u t_k| .$$

Note that the sequence $t_1, \ldots, t_k$ has the order as it would be seen by an observer standing in the center of the ring and turning counterclockwise starting at t_1.

Now put $t_{k+1} = t_1$. We have $\sum_{i=1}^{k} |\measuredangle t_i u t_{i+1}| = 2\pi$, and, by Lemma A.1 , $|\measuredangle t_i u t_{i+1}| > \rho(c,d)$. Hence, $k < \frac{2\pi}{\rho(c,d)}$. □

Now, partition $T(u)$ into two subsets T_1 and T_2:

$$T_1 = \{v \in T(u); \quad |uv| < a\} \qquad \text{and}$$
$$T_2 = \{v \in T(u); \quad |uv| \geq a\},$$

where $a = \sqrt{2}$. By Corollary A.1, the ring $R_u(a,2)$ contains less than $\frac{2\pi}{\rho(a,2)} < 13$ points, while the ring $R_u(1,a)$ less than $\frac{2\pi}{\rho(1,a)} = 8$ points, i.e., $|T_1| \leq 7$ and $|T_2| \leq 12$. This gives an upper bound 19 for the size of $T(u)$. In what follows we improve this bound by one. As we show later, this cannot be further improved. We start with the following auxiliary lemma.

Lemma A.2 *Let c and d be constants such that $1 \leq c < d \leq 2$. Consider a ring $R_u(c,d)$ centered at u containing two points v, w such that $|vw| > 1$ and $|\measuredangle vuw| \leq \alpha < \frac{\pi}{3}$. For every point z in the sector $S = \measuredangle vuw \cap R_u(1,2)$ such that $|zv| > 1$ and $|zw| > 1$ then either*

$$|uz| < \omega(d,\alpha) = d \cdot \cos\tfrac{\alpha}{2} - \sqrt{d^2(\cos^2\tfrac{\alpha}{2} - 1) + 1}, \text{ or}$$
$$|uz| > \omega'(c,\alpha) = c \cdot \cos\tfrac{\alpha}{2} + \sqrt{c^2(\cos^2\tfrac{\alpha}{2} - 1) + 1}\,.$$

In particular, if $c \leq \frac{27}{14}$ then for $z \in \measuredangle vuw \cap R_u(1,c)$, the first inequality must occur.

Proof. Since $\alpha < \frac{\pi}{3}$ and $|uv|, |uw| \leq 2$, the unit circles C_v centered at v and C_w centered at w intersect at two points, say z_1 and z_2, where z_1 is closer to u, cf. Figure A.2. It can be easily seen that the point z must be in $\measuredangle vuw - C_v - C_w$, i.e., in the shadowed region in the figure. Hence, we have either $|uz| < |uz_1|$, or $|uz| > |uz_2|$. Now, when moving v (resp. w) toward u, both $|uz_1|$ and $|uz_2|$ are getting smaller. Let $\tilde{z}_1$ be the point at which $|uz_1|$ is maximized. By the above argument, this maximum is attained for $v = \tilde{v}$ and $w = \tilde{w}$ such that $|u\tilde{v}| = |u\tilde{w}| = d$. Similarly, let $\hat{z}_2$ be the point at which $|uz_2|$ is minimized. This minimum is attained for $v = \hat{v}$ and $w = \hat{w}$ such that $|u\hat{v}| = |u\hat{w}| = c$.

Since both triangles $\triangle \tilde{w}u\tilde{v}$ and $\triangle \tilde{w}\tilde{z}_1\tilde{v}$ are isosceles, the angle $|\measuredangle \tilde{z}_1 u \tilde{v}| = \frac{\alpha}{2}$. Using Fact 1 on $\triangle \tilde{z}_1\tilde{v}u$, we obtain $\cos\frac{\alpha}{2} = \dfrac{d^2 + |u\tilde{z}_1|^2 - 1^2}{2d|u\tilde{z}_1|}$. This quadratic equation

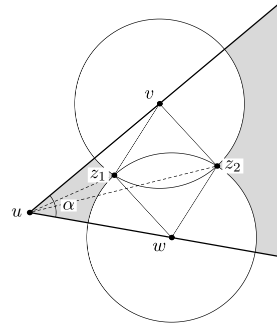

Figure A.2: A point in the sector $\measuredangle vuw$ not covered by v and w can occur only in the shadowed region).

has two solutions: Euclidean distances of intersection points of C_v and C_w from u. Since $\tilde{z}_1$ is the closer intersection point to u, its Euclidean distance from u is the smaller solution, i.e.,

$$|u\tilde{z}_1| = d \cdot \cos\tfrac{\alpha}{2} - \sqrt{d^2(\cos^2\tfrac{\alpha}{2} - 1) + 1} = \omega(d, \alpha)\,.$$

Similarly, we obtain

$$|u\hat{z}_2| = c \cdot \cos\tfrac{\alpha}{2} + \sqrt{c^2(\cos^2\tfrac{\alpha}{2} - 1) + 1} = \omega'(c, \alpha)\,.$$

Hence, we must have

$$|uz| < |u\tilde{z}_1| = \omega(d, \alpha) \quad \text{or} \quad |uz| > |u\hat{z}_2| = \omega'(c, \alpha)$$

as claimed in the first part of the lemma. In case $c \le \frac{27}{14}$, we have

$$|u\hat{z}_2| = \omega'(c, \alpha) > \tfrac{\sqrt{3}}{2}c + \sqrt{1 - \tfrac{c^2}{4}} \ge \tfrac{\sqrt{3}}{2}c + \tfrac{\sqrt{55}}{28} > c\,.$$

Since z is in $R_u(1, c)$, the case $|uz| > |u\hat{z}_2|$ is not possible, and the second part follows. □

Proof of Lemma 2. Partition T into the two subsets T_1 and T_2 defined above. As before, $|T_1| \le 7$ and $|T_2| \le 12$.

If $|T_1| < 7$ or $|T_2| < 12$, we are done. Hence, assume that $|T_1| = 7$ and $|T_2| = 12$. Let $t_1, \ldots, t_{12}$ be the elements of T_2 ordered in such a way that

$$|\measuredangle t_1ut_2| \le |\measuredangle t_1ut_3| \le \cdots \le |\measuredangle t_1ut_{12}|\,,$$

and $|\measuredangle t_1ut_2| \ge |\measuredangle t_iut_{t+1}|$, for all $i = 2, \ldots, 12$, where $t_{13} = t_1$. Consider sectors $S_i = \measuredangle t_iut_{i+1} \cap R_u(1, a)$. We start with proving two facts which will be used to obtain an upper bound on the number of points in T_1.

First, for every sector S_i, we determine the maximal Euclidean distance from the central point u to any point in $T_1 \cap S_i$.

Denote $\alpha_i = |\measuredangle t_i u t_{i+1}|$. By Lemma A.1, we have the lower bound $\alpha_i > \beta$, where $\beta = \rho(a,2) = \arccos \frac{5\sqrt{2}}{8}$. This yields the following upper bounds

$$\alpha_1 = 2\pi - \sum_{i=2}^{12} \alpha_i < 2\pi - 11\beta\,,$$
$$\alpha_j \le \tfrac{1}{2}\alpha_1 + \tfrac{1}{2}\alpha_j = \pi - \tfrac{1}{2} \sum_{i \ne 1, i \ne j} \alpha_i < \pi - \tfrac{1}{2}\cdot 10\beta = \pi - 5\beta\,, \quad \text{for all } j = 2, \ldots, 12.$$

Since $a \le \frac{27}{14}$ and $\pi - 5\beta < 2\pi - 11\beta < \frac{\pi}{3}$, by Lemma A.2, any point z in $T_1 \cap S_1$ satisfies $|uz| < \omega(2, 2\pi - 11\beta)$, and similarly, any point z in $T_1 \cap S_j$ satisfies $|uz| < \omega(2, \pi - 5\beta)$.

Second, we determine the maximum number of points in $T_1 \cap S_1$. Since all points in $T_1 \cap S_1$ belong to $R_u(1, \omega(2, 2\pi - 11\beta))$ and

$$1 < \omega(2, 2\pi - 11\beta) \le \sqrt{3} < 2\,, \tag{A.2}$$

we can apply the second part of Lemma A.1. Thus, we have the following upper bound on the number of points:

$$\left\lceil \frac{\alpha_1}{\rho(1, \omega(2, 2\pi - 11\beta))} \right\rceil \le \left\lceil \frac{2\pi - 11\beta}{\rho(1, \omega(2, 2\pi - 11\beta))} \right\rceil \le 2\,,$$

i.e., there are at most two points in $T_1 \cap S_1$.

To finish the proof, let $s_1, \ldots, s_7$ be the ordering of elements of T_1 such that

$$|\measuredangle s_1 u s_2| \le |\measuredangle s_1 u s_3| \le \cdots \le |\measuredangle s_1 u s_7|\,;$$

$|us_i| < \omega(2, 2\pi - 11\beta)$, for $i = 1, 2$; and $|us_i| < \omega(2, \pi - 5\beta)$, for all $i \ge 3$. Such an ordering is indeed possible, since there are at most 2 points in $T_1 \cap S_1$.

Finally, we lower bound angles between any two consecutive points in our ordering. Since, for points s_1 and s_2 we have a weaker upper bound $\omega(2, 2\pi - 11\beta)$ for their Euclidean distance from u, by Lemma A.1, the angles

$$|\measuredangle s_7 u s_1|, |\measuredangle s_1 u s_2|, |\measuredangle s_2 u s_3| \ge \rho(1, \omega(2, 2\pi - 11\beta))\,.$$

On the other hand, for points $s_3, \ldots, s_7$, the upper bound for their Euclidean distance from u is $\omega(2, \pi - 5\beta)$. Hence, Lemma A.1 gives a lower bound

$$|\measuredangle s_3 u s_4|, \ldots, |\measuredangle s_6 u s_7| \ge \rho(1, \omega(2, \pi - 5\beta))\,.$$

Since the seven angles must sum up to 2π, we have a contradiction with

$$|\measuredangle s_7 u s_1| + \cdots + |\measuredangle s_6 u s_7| > 3\rho(1, \omega(2, 2\pi - 11\beta)) + 4\rho(1, \omega(2, \pi - 5\beta)) > 2\pi\,.$$

□