13 Wireless Overlay Networks I

Radio networks are widely used today. People access voice and data services via mobile phones, Bluetooth technology replaces unhandy cables by wireless links, and wireless networking is possible via IEEE 802.11 compatible network equipment. Nodes in such networks exchange their data packets usually with fixed base stations that connect them with a wired backbone. However, in applications such as search and rescue missions or environmental monitoring, no explicit communication infrastructure may be available. In this case, the wireless hosts have to organize in a so-called wireless ad hoc network. As long as all of the hosts are within transmission range of each other, the problem of exchanging information in such a network basically boils down to designing suitable medium access control protocols, but if not all hosts can directly communicate with each other, we also need suitable routing algorithms. Designing routing algorithms for wireless ad hoc networks is an extremely challenging task and still research in progress. We mostly focus on the simpler question of how to maintain an overlay network of wireless links between the hosts so that, as long as this is in principle possible,

1. every node is reachable from every other node, i.e. the graph formed by the links is connected, and

2. for every pair of nodes (v, w) there is a route from v to w with a close to minimum possible hop distance or energy consumption.

The graph formed by the wireless links should also have a low degree to ensure a low maintenance cost, it should allow to find routes for the messages that have a low congestion, and it should be easy to update in case of arrivals or departures of nodes or changes in their positions. We will present various local-control algorithms for reaching these goals.

13.1 First approaches

The problem of designing overlay networks for wireless ad hoc networks has recently attracted a lot of attention. A basic requirement for these overlay network designs is that they maintain connectivity among the hosts, as long as this is possible. The most straightforward approach to achieve connectivity is to maintain a link between every pair of wireless hosts that are within their transmission range. However, this may require a high maintenance and update cost since the corresponding overlay network may have a high degree. Also, some links may have a high energy cost, and so a natural question would be whether these can be dropped without endangering connectivity.

An alternative approach would be to maintain connections only to the k nearest neighbors. However, Figure 1 demonstrates that it is easy to come up with examples in which the graph formed by the links would not be connected. So this approach does not work in general. As was shown by Xue and Kumar [11], it only works in specific cases. For example, if n hosts are distributed uniformly at random in a unit square and every host connects to more than $5.1774 \log n$ of its nearest neighbors, then the network formed by these links is connected with a probability that tends to 1 as n increases. But connecting to less than $0.074 \log n$ nearest neighbors results in almost sure disconnectivity.

Another possible approach is that every host maintains connections to k hosts chosen uniformly at random among all hosts within its transmission range. This also does not guarantee connectivity in general but works well in certain cases. For example, Dubhashi et al. [4] recently showed that if every

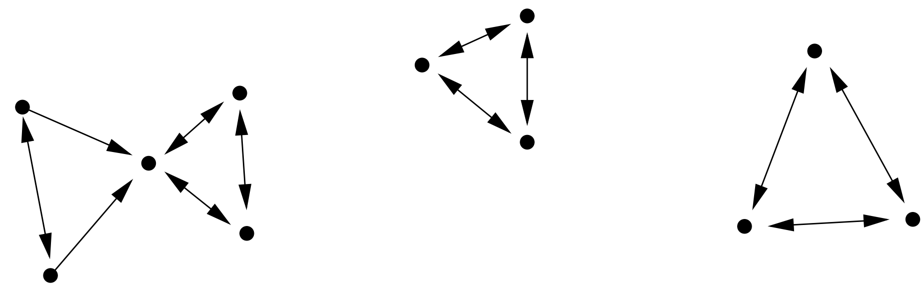

Figure 1: A counterexample for the naive approach with $k = 2$.

node has at least $\Theta(\log n)$ nodes within its transmission range, then choosing just 2 random nodes to connect to will establish connectivity almost surely.

We will only focus on approaches here that *guarantee* connectivity, no matter how the hosts are distributed, as long as this is in principle possible. Most of these approaches are based on so-called *spanners*, which are properly selected subgraphs of the graph of all possible connections between the wireless hosts so that the hosts are not only connected but their (hop or Euclidean) distance in that graph is closely related to their minimum (hop or Euclidean) distance when considering all possible connections. Spanners first appeared in computational geometry [5, 12], were then discovered as an interesting tool for approximating NP-hard problems [9], and have recently attracted a lot of attention in the context of routing and topology control in wireless ad hoc networks [1, 6, 7, 2, 8].

13.2 Notation

In the following, the wireless hosts are simply called *nodes*. To simplify our presentation, we assume that the nodes are distributed in a perfect 2-dimensional Euclidean space, or formally, the nodes represent a set of points $V \subset \mathbb{R}^2$, but all of the approaches presented here can also be extended to higher dimensions. Given any pair of nodes $u = (u_x, u_y), v = (v_x, v_y) \in \mathbb{R}^2$,

$$||uv|| = \sqrt{(u_x - v_x)^2 + (u_y - v_y)^2}$$

denotes the *Euclidean distance* between u and v, and given any sequence of nodes $s = (u_1, u_2, \ldots, u_k)$ and any $\delta \geq 0$,

$$||s||^\delta = \sum_{i=1}^{k-1} ||u_i u_{i+1}||^\delta$$

denotes the δ-*cost* of s. For any graph $G = (V, E)$, a node sequence $s = (u_1, u_2, \ldots, u_k)$ is called a *path* in G if $(u_i, u_{i+1}) \in E$ for all $1 \leq i < k$.

Given any directed graph $G = (V, E)$ and any two nodes $u, v \in V$, the δ-*distance* $d_G^\delta(u, v)$ of u and v in G is the minimum δ-cost $||p||^\delta$ over all paths p from u to v in G. If $\delta = 0$, then $d_G^\delta(u, v)$ gives the *topological (or hop) distance* of u and v in G, and if $\delta = 1$, $d_G^\delta(u, v)$ gives the *Euclidean distance* of u and v in G. Also cases with $\delta > 1$ are interesting for us because the transmission of a packet over a distance of r usually has an energy consumption that scales with r^δ for some $\delta > 1$. In reality, δ is usually in the range $[2, 5]$, where it is closer to 2 outdoors and closer to 5 indoors.

We assume that every node has a maximum transmission range of 1, i.e., every node $u \in V$ can send messages only to nodes $v \in V$ with $||uv|| \leq 1$. From this assumption it follows that every overlay network connecting these nodes can only be a subgraph of the following graph.

Definition 13.1 *For any point set $V \subset \mathbb{R}^2$, the* unit disk graph *of V, called $UDG(V)$, is a directed graph that contains all edges (u, v) with $||uv|| \leq 1$.*

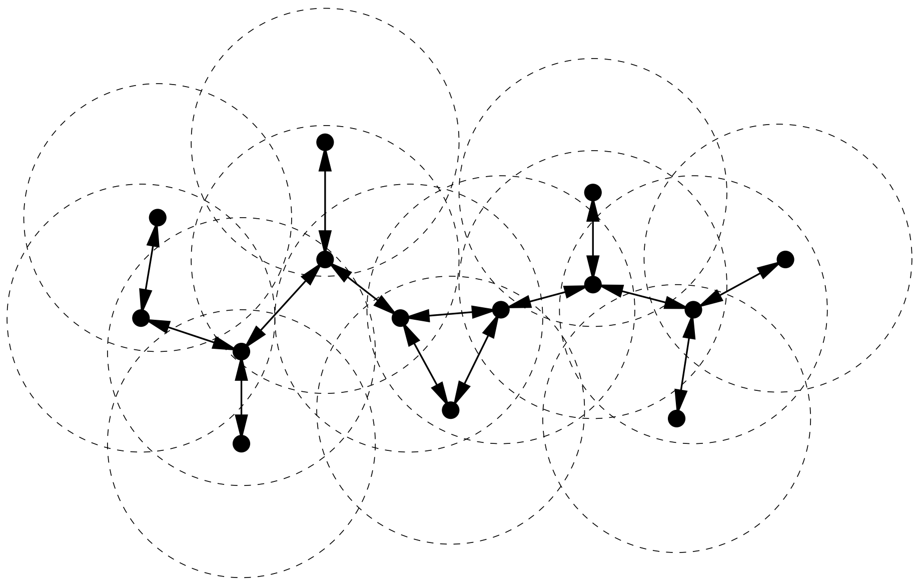

Figure 2: A connected unit disk graph.

In the following, we will always assume that V is chosen so that its UDG is connected and non-degenerate, i.e., there is a path in $UDG(V)$ between every pair of nodes and no two pairs of nodes have exactly the same Euclidean distance (see also Figure 2). The connectivity assumption is a prerequisite for our strategies below to establish a connected network among the nodes and the non-degenerateness property will simplify the proofs. When G is the UDG of V, we simply use $d^{\delta}(u, v)$ instead of $d^{\delta}_G(u, v)$.

Next, we introduce graph spanners. First, we define spanners in which arbitrary pairs of nodes can, in principle, be connected by an edge (i.e., we do not limit the transmission range of nodes).

Definition 13.2 *Consider any finite set of nodes $V \subset \mathbb{R}^2$, and let $c \geq 1$ be any constant.*

- *A graph $G = (V, E)$ is called a* geometric c-spanner *of V if for all $u, v \in V$ there exists a path p from u to v in G with*

$$||p|| \leq c \cdot ||uv|| \; .$$

  If G is a geometric c-spanner, c is called its stretch factor.

- *G is a (c, δ)-*power spanner *of V if for all $u, v \in V$ there is a path p from u to v in G with*

$$||p||^{\delta} \leq c \cdot ||uv||^{\delta} \; .$$

  If for all $\delta \geq 2$ there exists a constant c so that G is a (c, δ)-power spanner, then we simply call G a power spanner.

- *G is a* weak c-spanner *of V if for all $u, v \in V$ there is a path p from u to v in G that is within a disk of diameter at most*

$$c \cdot ||uv||$$

- *A graph $G = (V, E)$ is called a* constrained *(geometric, power, or weak) spanner of V if for every pair of nodes $u, v \in V$ there is a path p that, in addition to the specific requirement for the spanner type, also satisfies the condition that for every edge e in p,*

$$||e|| \leq ||uv||$$

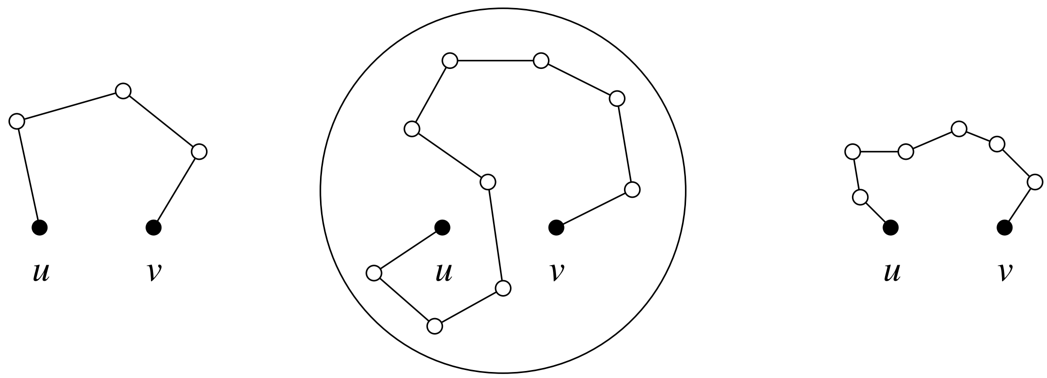

Figure 3: Examples of a spanner, weak spanner, and power spanner.

Since wireless nodes have a limited transmission range, the following spanner definitions are more relevant for ad hoc networks.

Definition 13.3 *Let $V \subset \mathbb{R}^2$ be any finite set of nodes with a connected UDG.*

- *A graph $G = (V, E)$ is called a* geometric c-spanner *of $UDG(V)$ if for all $u, v \in V$ there exists a path p from u to v in G with*

$$||p|| \leq c \cdot d(u, v) .$$

- *G is a* (c, δ)-power spanner *of $UDG(V)$ if for all $u, v \in V$ there is a path p from u to v in G with*

$$||p||^\delta \leq c \cdot d^\delta(u, v) .$$

- *G is a* weak c-spanner *of $UDG(V)$ if for all $u, v \in V$ there is a path p from u to v in G that is within a disk of diameter at most*

$$c \cdot d(u, v)$$

Interestingly, any constrained spanner of V in which all edges of length more than 1 are removed is also a spanner of the UDG of V, as shown in the next theorem.

Theorem 13.4 *Any constrained geometric c-spanner / (c, δ)-power spanner / weak c-spanner G of V restricted to edges of length at most 1 is also a geometric c-spanner / (c, δ)-power spanner / weak c-spanner of the UDG of V.*

Proof. Let U be the UDG of V. Suppose that G is a (c, δ)-power spanner of V for some $\delta \geq 0$. Then it holds for every pair of nodes $u, v \in V$ with $||uv|| \leq 1$ that there is a path p in $G \cap U$ with $||p||^\delta \leq c||uv||^\delta$. Now, consider an arbitrary pair $u, w \in V$, and let $p = (v_0, v_1, v_2, \ldots, v_k)$ be any path in U with $v_0 = u$ and $v_k = w$ that has a δ-cost of $d^\delta(u, w)$. Since $||v_i v_{i+1}|| \leq 1$ for all i, there is a path p_i from v_i to v_{i+1} in $G \cap U$ with $||p_i||^\delta \leq c||v_i v_{i+1}||^\delta$. Concatenating these paths, we end up with a path p' with

$$||p'||^\delta = \sum_{i=0}^{k-1} ||p_i||^\delta \leq \sum_{i=0}^{k-1} c||v_i v_{i+1}||^\delta = c \cdot d^\delta(u, w) \ .$$

Hence, $G \cap U$ is also a (c, δ)-power spanner of U. Since a geometric c-spanner is just a $(c, 1)$-power spanner, this also proves the theorem for constrained geometric spanners.

Finally, consider the case that G is a constrained weak c-spanner. Then it holds for every pair of nodes $u, v \in V$ with $||uv|| \leq 1$ that there is a path p in $G \cap U$ that is within a disk of diameter at most $c||uv||$. Consider now an arbitrary pair $u, w \in V$, and let $p = (v_0, v_1, v_2, \ldots, v_k)$ be any path in U with $v_0 = u$ and $v_k = w$ that has a Euclidean length of $d(u, w)$. Since $||v_i v_{i+1}|| \leq 1$ for all i, there is a path p_i from v_i to v_{i+1} in $G \cap U$ that is within a disk of diameter at most $c \cdot ||v_i v_{i+1}||$. Concatenating these paths, we end up with a path p' that is within a disk of diameter at most $c \cdot d(u, w)$. To prove this, we need the following straightforward fact.

Fact 13.5 *Any two disks of diameter d_1 and d_2 with a non-empty intersection are contained in a disk of diameter at most $d_1 + d_2$.*

Using this fact in an inductive manner on the length of p, it follows that when replacing the paths p_i in p' by their disks, p' is contained in a disk of radius at most

$$\sum_{i=0}^{k-1} c \cdot ||v_i v_{i+1}|| \leq c \cdot d(u, w)$$

$\square$

Hence, it suffices to present and analyze algorithms for constrained spanners in order to obtain overlay networks that are also spanners of UDGs.

13.3 Geometric spanners, power spanners, and weak spanners

Next, we study general relationships between the different kinds of spanners. All of these relationships hold for general spanners as well as constrained spanners. However, to simplify the presentation, we only prove the statements for general spanners. The reader may verify that they also hold for constrained spanners.

Theorem 13.6 *Every graph $G = (V, E)$ that is a (constrained) geometric c-spanner is also a (constrained) weak c-spanner.*

Proof. Consider any pair of nodes $u, w \in V$. Since G is a geometric c-spanner, there is a path $p = (v_0, v_1, v_2, \ldots, v_k)$ with $v_0 = u$ and $v_k = w$ that has a length of at most $c \cdot ||uw||$. Replacing each edge (v_i, v_{i+1}) by the disk of diameter $||v_i v_{i+1}||$ containing v_i and v_{i+1} and using Fact 13.5 in an inductive manner implies that p is contained in a disk of diameter at most

$$\sum_{i=0}^{k-1} ||v_i v_{i+1}|| \leq c \cdot ||uw||$$

□

However, the theorem does not hold any more when considering power spanners [10].

Theorem 13.7 *For any $\delta > 1$ there is a family of (constrained) (c, δ)-power spanners which are not a (constrained) weak C-spanner for any constant C.*

Proof. Let $V = \{v_1, v_2, \ldots, v_n\}$ be a set of n nodes placed on a circle scaled so that $||v_1 v_n|| = 1$ and $||v_i v_{i+1}|| = 1/i$ for all $1 \leq i < n$. Now, consider the graph $G = (V, E)$ with edges (v_i, v_{i+1}) for all $1 \leq i < n$ (see also Figure 4). First, we show that G is not a weak C-spanner for any constant C, and then we show that G is a (c, δ)-power spanner for all $\delta > 1$.

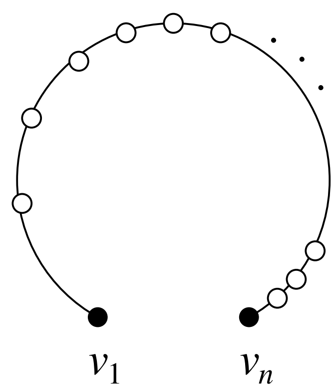

Figure 4: Example graph for Theorem 13.7.

It is easy to see that for every $n \geq 2$, the circumference of the circle through the n nodes is at least

$$\sum_{i=1}^{n-1} \frac{1}{i} \geq \int_{x=1}^{n} \frac{1}{x}\, dx = \ln n$$

which implies that its diameter is at least $(\ln n)/\pi$. Since the distance between v_1 and v_n is just 1 but any path from v_1 to v_n would have to traverse all nodes v_i along the circle, there cannot be a constant C so that G is a weak C-spanner.

On the other hand, if we look at the δ-cost of the unique path p from v_1 to v_n in G, we see that for $\delta = 1 + \epsilon > 1$,

$$||p||^\delta = \sum_{i=1}^{n-1} \left(\frac{1}{i}\right)^\delta \leq 2 \int_{x=1}^{n} -\frac{1}{\epsilon} \left(\frac{1}{x}\right)^\epsilon dx \leq \frac{2}{\epsilon} = \frac{2}{\delta - 1} \ .$$

Since p uses all edges of G, this is an upper bound for the δ-cost of any path connecting any other pair of nodes (v_i, v_j) in G. Hence, for all pairs of nodes (v_i, v_j) with $||v_i v_j|| \geq 1$, the δ-cost is at most

$2/(\delta - 1)$. Moreover, it is not hard to check that every pair of nodes (v_i, v_j) with $||v_i v_j|| < 1$ has a path p from v_i to v_j of length at most $1 + O(1/\ln n)$ and therefore of δ-cost at most $(1 + O(1/\ln n))^\delta$. Hence, G is a (c, δ)-power spanner for some constant c (depending on δ) for any constant $\delta > 1$. □

Also, the reverse direction of Theorem 13.6 is not true, i.e., the fact that a graph is a weak spanner does not imply in general that it is also a geometric spanner (see also [10]).

Theorem 13.8 *There exists a family of graphs $G = (V, E)$ with $V \subset \mathbb{R}^2$ all of which are (constrained) weak $2(\sqrt{2} + 1)$-spanners but not a (constrained) geometric c-spanner for any constant c.*

Proof. Consider the snowflake structure in Figure 5. As can be seen from the picture (see the nodes v and w), the stretch factor of the snowflake structure is equal to

$$\frac{4}{2 + \sqrt{2}} \geq 1.17$$

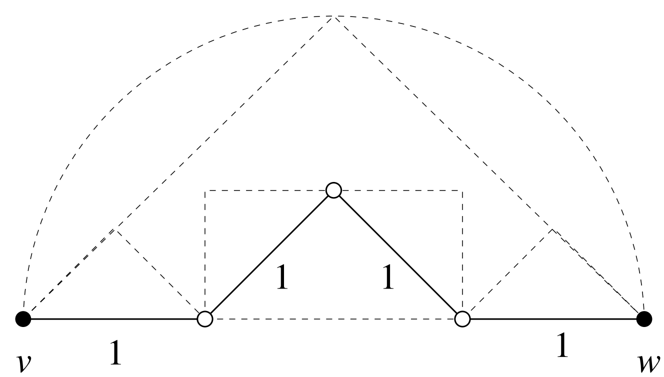

Figure 5: The basic snowflake structure.

Recursively replacing each edge by a snowflake structure over d levels increases the stretch factor to at least 1.17^d. Suppose now that we have n nodes, where n is a multiple of 4. Then we can use them to construct a snowflake structure with $d = \log_4 n$ levels. This results in a stretch factor of at least

$$1.17^{\log_4 n} = n^{(\log_2 1.17)/(\log_2 4)} > n^{0.11} .$$

Next we show that the recursive snowflake structure is a weak spanner. Consider the dashed triangle through v and w in Figure 5. This triangle certainly contains all the other triangles in the picture. Using this observation inductively, starting with the lowest level, it follows that for any recursion depth, the recursive snowflake structure with endpoints v and w is completely inside the triangle through v and w.

Now, let $G = (V, E)$ be any recursive snowflake structure of depth d. Consider any two nodes $v', w' \in V$. Let G' be the snowflake structure of largest depth within G that contains v' and w', and let S be the basic snowflake structure (i.e., we ignore further recursions) associated with G' and v and w be its endpoints (like in Figure 5). At this point we distinguish between two cases.

If v' and w' are associated with two non-adjacent edges in S, then it follows from the fact that any two points in non-adjacent triangles in Figure 5 have a distance of at least 1 that v' and w' must have a distance of at least 1. On the other hand, G' is completely contained in the disk of diameter $||vw||$ through v and w. Hence, G is a weak $2+\sqrt{2}$-spanner for these (v', w') pairs.

If v' and w' are associated with two adjacent edges in S, then we recurse further on the snowflake structures of v' and w' until, for the first time, v' and w' are not associated with adjacent edges of adjacent snowflake structures, or we reached the lowest level. In this case, v' and w' are contained in triangles that are at least two edges away from each other. Going through all cases for these triangles, one can easily check that they must have a distance of at least $\sqrt{2}$ (relative to an edge length of 1 in the snowflake structures of v' and w') while the snowflake structures of v' and w' are contained in a disk of diameter $2(2+\sqrt{2})$. Hence, G is a weak $2(\sqrt{2}+1)$-spanner for these (v', w')-pairs.

Combining the two cases proves the theorem. □

The next theorem studies the relationship between geometric spanners and power spanners.

Theorem 13.9 *Every (constrained) geometric c-spanner is a (constrained) (c^δ, δ)-power spanner for every $\delta \geq 1$.*

Proof. Let $G = (V, E)$ be a geometric c-spanner. Then it holds that for every pair of nodes $u, w \in V$ there is a path $p = (v_0, v_1, \ldots, v_\ell)$ in G with $v_0 = u$ and $v_\ell = w$ and $||p|| = \sum_{i=0}^{\ell-1} ||v_i v_{i+1}|| \leq c \cdot ||uw||$. Hence, for every $\delta \geq 1$,

$$||p||^\delta = \sum_{i=0}^{\ell-1} ||v_i v_{i+1}||^\delta \leq \left(\sum_{i=0}^{\ell-1} ||v_i v_{i+1}|| \right)^\delta \leq c^\delta \cdot ||uw||^\delta .$$

Therefore, G is also a power spanner for all $\delta \geq 1$, which proves the theorem. □

Hence, in order to prove that a graph is a power spanner, it suffices to prove that it is a geometric spanner. Interestingly, for $\delta \geq 2$, it even suffices to show that a graph is a weak spanner in order to prove that it is a power spanner. We only prove this fact for $\delta > 2$. The proof for $\delta = 2$ is involved and can be found (as well as the proof for $\delta > 2$) in [10].

Theorem 13.10 *Let $G = (V, E)$ be a (constrained) weak c-spanner. Then G is also a (constrained) (C, δ)-power spanner for $\delta > 2$ where $C = (4c+1)^2 \cdot \frac{c^\delta}{1-2^{2-\delta}}$.*

Proof. Consider any pair of nodes $v, w \in V$ and let p be any path from v to w that is in a disk $D(v, w)$ of diameter at most $c||vw||$. Suppose first there is no pair of nodes in p with a distance of at most $||vw||/2$. In this case, the disks of radius $||vw||/4$ around each of these nodes are disjoint. Since a disk of diameter d has a surface of $\pi(d/2)^2$, it follows that there can be at most

$$\frac{\pi((c+1/4)||vw||/2)^2}{\pi(||vw||/4)^2} \leq (4c+1)^2$$

nodes in $D(v, w)$ that are used by p. Since the δ-cost of any edge in $D(v, w)$ is at most $(c||vw||)^\delta$, it follows that the δ-cost of p is at most $(4c+1)^2 c^\delta ||vw||^\delta$.

Suppose now that there is a (not necessarily adjacent) pair of nodes in p with a distance of at most $||vw||/2$ (see Figure 6). Let v' be the first node reached when walking along p from v to w that has a node w' on p with $||v'w'|| \leq ||vw||/2$. Then we replace the part of p from v' to w' by the path p'

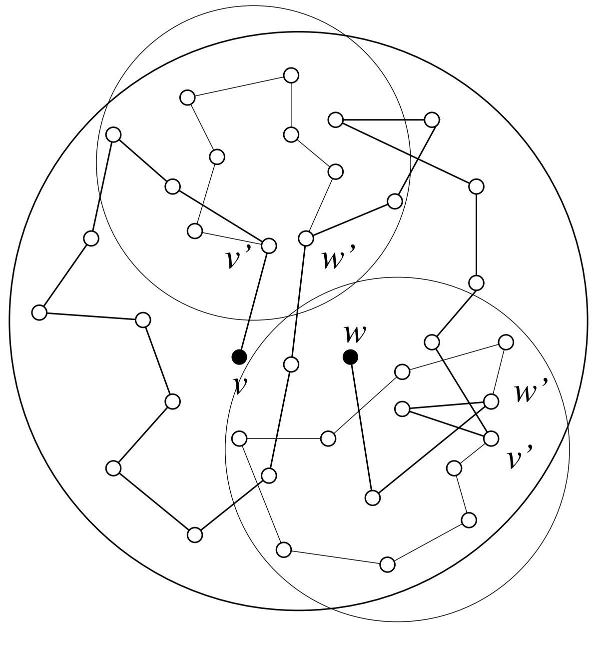

Figure 6: A possible path for the nodes v, w with two pairs v', w' that are too close.

from v' to w' that stays within a disk $D(v', w')$ of diameter at most $c||v'w'|| \leq c||vw||/2$. We continue walking along p at w' until we find the next node v' on p that has a node w' on p (*ignoring* the nodes in p') of distance at most $||vw||/2$. We do the same transformation for this pair and continue until the entire path p has been traversed.

After this transformation, all pairs of nodes (v', w') in p of distance at most $||vw||/2$ have been replaced by paths using edges of length at most $c||vw||/2$. Suppose that the nodes of these paths have a pairwise distance of more than $||vw||/4$. Since $D(v, w)$ can have at most $4(4c+1)^2$ such nodes and the δ-cost of any edge connecting these nodes is at most $(c||vw||/2)^\delta$, it follows that the δ-cost of all replaced parts of p is at most $4(4c+1)^2(c/2)^\delta||vw||^\delta$.

However, there may still be nodes in the new parts of p that have a pairwise distance of at most $||vw||/4$. Then we continue to reorganize the new parts of p as we did with the path p above. Each level i of reorganization creates an additional δ-cost of at most

$$2^{2i}(4c+1)^2(c/2^i)^\delta||vw||^\delta \ .$$

Since V is finite, and therefore the minimum distance between any two nodes is finite, this reorganization eventually terminates. At the end, the total δ-cost with $\delta > 2$ is at most

$$\sum_{i\geq 0}\left(\frac{2^2}{2^\delta}\right)^i(4c+1)^2c^\delta||vw||^\delta \ = \ (4c+1)^2c^\delta||vw||^\delta\sum_{i\geq 0}\frac{1}{2^{\delta-2}} \ = \ (4c+1)^2c^\delta||vw||^\delta\frac{1}{1-2^{2-\delta}}$$

□

However, a weak c-spanner may not be a (C, δ)-power spanner for any constant C if $\delta < 2$ [10].

Theorem 13.11 *For any $\delta < 2$ there exists a family of graphs $G = (V, E)$ with $V \subset \mathbb{R}^2$ which are (constrained) weak c-spanners for a constant c but not a (constrained) (C, δ)-power spanner for any constant C.*

Summing up Theorems 13.6, 13.7, 13.8, 13.9, and 13.10, we obtain the following interesting relationship between the class of all geometric spanners, weak spanners, and power spanners with $\delta \geq 2$:

$$\text{Geometric spanners} \quad \subset \quad \text{Weak spanners} \quad \subset \quad \text{Power spanners}$$

13.4 Proximity graphs

From our insights on spanners above it follows that it would often be sufficient to design protocols that guarantee a constrained weak c-spanner as long as this is possible because weak spanners are guaranteed to have energy-efficient paths. But how can such protocols be designed in a distributed way? Let us first focus on the weak spanner property. Consider the following definition:

Definition 13.12 *For any node set $V \subset \mathbb{R}^2$, the graph $G = (V, E)$ is called a* proximity graph *of V if and only if for all $u, w \in V$ it holds that*

- *$(u, w) \in E$ or*
- *there is a $v \in V$ with $(u, v) \in E$ and $||vw|| < ||uw||$.*

For an example of a node v satisfying the proximity conditions, see Figure 7. It is known that there are proximity graphs with a stretch factor as bad as $|V| - 1$ [3] but proximity graphs are always good weak spanners.

Theorem 13.13 *For any finite $V \subset \mathbb{R}^2$, every proximity graph of V is a weak 2-spanner.*

Proof. Let $G = (V, E)$ be any proximity graph of V. First we prove that G is connected. Certainly, a graph G is connected if and only if for every pair of nodes in G there is a path connecting these two nodes. So consider any pair of nodes $u, w \in V$. We distinguish between two cases:

1. $(u, w) \in E$: Then u and w are connected, and we are done.
2. There is a $v \in V$ with $(u, v) \in E$ and $||vw|| < ||uw||$: Then we use the edge (u, v) and get closer to w then we were before.

Since V is finite, we only have to apply case 2 a finite number of times until case 1 holds. Hence, G is connected.

Besides G being connected, it follows from the observation above that for any pair of nodes $u, w \in V$ there is a path p that monotonically converges against w. Hence, p is contained in a disk of diameter at most $2||uw||$, which proves the theorem. □

Hence, every proximity graph is also a power spanner of V for every $\delta \geq 2$. To make proximity graphs useful for ad hoc networks, we consider a constrained form of proximity graphs which are also known as relative neighborhood graphs [3].

Definition 13.14 *For any node set $V \subset \mathbb{R}^2$, the graph $G = (V, E)$ is called a* relative neighborhood graph *(RNG) of V if and only if for all $u, w \in V$ it holds that*

- *$(u, w) \in E$ or*

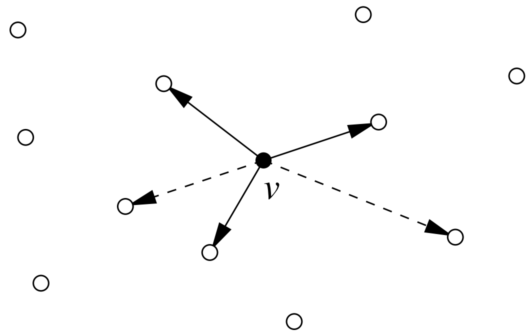

Figure 7: Connections satisfying the RNG condition for v. (Removing the dashed connections gives a minimum set of connections satisfying the RNG condition.)

- *there is a* $v \in V$ *with* $(u, v) \in E$, $||uv|| < ||uw||$, *and* $||vw|| < ||uw||$.

It is easy to verify that relative neighborhood graphs satisfy the condition on constrained graphs we formulated for spanners in Definition 13.3. Hence, Theorems 13.4, 13.10, and 13.13 imply that relative neighborhood graphs are weak and power spanners of the UDG of V for every $\delta \geq 2$.

Figure 8 shows a simple distributed protocol for minimal relative neighborhood graphs. In this protocol we assume that every node $u \in V$ knows its neighborhood

$$N(u) = \{v \in V \mid ||uv|| \leq 1\}$$

and the current positions of the nodes in $N(u)$. Node u also keeps track of three sets:

- $E(u)$: set of edges that u currently has to nodes in $N(u)$.
- $B(u)$: set of nodes $w \in N(u) \setminus E(u)$ that have a node $v \in E(u)$ with $||uv|| < ||uw||$ and $||vw|| < ||uw||$.
- $U(u)$: nodes in $N(u)$ that are not in $E(u)$ or $B(u)$ (for example, nodes that newly entered $N(u)$).

Theorem 13.15 *The RNG protocol self-stabilizes in at most 5 rounds. In the stable state, the outdegree of every node is at most 5.*

Proof. First we prove that the protocol self-stabilizes in at most 5 rounds. Suppose that node u is in an arbitrary state at the beginning of the first round. Only $N(u)$ and the positions of the nodes in $N(u)$ are assumed to be correct. Then it holds:

- After step 1) of round 1, $E(u)$ is minimal, and the nodes taken out of $E(u)$ were moved to $B(u)$.
- After step 2) of round 1, $B(u)$ is minimal, and the nodes taken out of $B(u)$ were moved to $U(u)$.
- After step 3) of round 1, $U(u)$ is empty and $E(u)$ and $B(u)$ satisfy the definitions of a RNG.

Protocol RNG:

For every node $u \in V$ repeatedly do:
1) for every node $w \in E(u)$:
if there is a node $v \in E(u)$ with $||uv|| < ||uw||$ and $||vw|| < ||uw||$ then
move w to $B(u)$ (i.e., remove edge (u, w))
2) for every node $w \in B(u)$:
if there is no $v \in E(u)$ with $||uv|| < ||uw||$ and $||vw|| < ||uw||$ then
move w to $U(u)$
3) for every node $w \in U(u)$:
if there is a node $v \in E(u)$ with $||uv|| < ||uw||$ and $||vw|| < ||uw||$ then
move w to $B(u)$
else move w to $E(u)$

Figure 8: A self-stabilizing protocol for relative neighborhood graphs.

If $E(u)$ is minimal after step 3, $E(u)$ and $B(u)$ will not be changed in later rounds as long as $N(u)$ and the positions of the nodes do not change. Hence, in this case, the protocol has stabilized.

$E(u)$ may not be minimal after step 3 of the first round, but $E(u)$ is guaranteed to contain the node $v_1 \in N(u)$ of minimum distance to u, no matter whether initially $v_1 \in E(u)$ or $v_1 \neq E(u)$. Let

$$B_1(u) = \{w \in N(u) \mid ||uv_1|| < ||uw|| \text{ and } ||v_1w|| < ||uw||\}$$

and let $v_2 \in N(u) \setminus (\{v_1\} \cup B_1(u))$ be the remaining node of closest distance to u (if it exists). Notice that v_2 can only be prevented to join $E(u)$ if a node in $B_1(u)$ is in $E(u)$. Hence, it holds:

- After step 1) of round 2, $v_1 \in E(u)$ and $B_1(u) \subseteq B(u)$ and $v_2 \in E(u) \cup B(u)$.
- After step 2) of round 2, $v_1 \in E(u)$ and $B_1(u) \subseteq B(u)$ and $v_2 \in E(u) \cup U(u)$.
- After step 3) of round 2, $v_1, v_2 \in E(u)$ and $B_1(u) \subseteq B(u)$.

Now, let

$$B_2(u) = \{w \in N(u) \setminus B_1(u) \mid ||uv_1|| < ||uw|| \text{ and } ||v_1w|| < ||uw||\}$$

and $v_3 = N(u) \setminus (\{v_1, v_2\} \cup B_1(u) \cup B_2(u))$ be the remaining node of closest distance to u (if it exists). Following the arguments for v_2, it is guaranteed that $v_3 \in E(u)$ and $B_2(u) \subseteq B(u)$ after round 3. This is continued until there is no node v_i. At this point, $E(u)$ and $B(u)$ are stable.

It remains to bound the number of rounds the protocol needs to stabilize. For this we need the following lemma, which implies that in the final $E(u)$ there can be at most 5 nodes, and therefore the protocol needs at most 5 rounds to stabilize. In the following, $\angle(v, u, w)$ denotes the angle between the lines $\overline{uv}$ and $\overline{uw}$.

Lemma 13.16 *In a minimal set $E(u)$ there cannot be two nodes $v, w \in E(u)$ with $\angle(v, u, w) \leq \pi/3$.*

Proof. Suppose that there are two nodes $v, w \in E(u)$ with $\angle(v, u, w) \leq \pi/3$. Let v be the closer of the two nodes (which is unique because we only consider non-degenerate sets V). Then $||uv|| < ||uw||$ and also $||vw|| < ||uw||$. □

If $|E(u)| \geq 6$, then there must be at least two nodes v and w in $E(u)$ with $\angle(v, u, w) \leq \pi/3$. But in this case, the lemma above implies that $E(u)$ cannot be minimal, which completes the proof of the theorem. $\square$

Though relative neighborhood graphs may be good weak spanners, they may not be geometric spanners or power spanners with a low cost. Here, two basic approaches have been pursued in the literature to obtain geometric spanners and/or power spanners with low cost:

- The nodes cut the space around them into sectors of equal angle θ, where θ is sufficiently small. Such graphs are also known as θ-graphs or Yao graphs.
- The nodes triangulate the space to form Delaunay-like graphs.

In the next section, we first consider Yao graphs and their variants, which we also call *sector-based spanners*, and afterwards we study Delaunay graphs and their variants, which we also call *planar spanners*.

References

[1] K. Alzoubi, X.-Y. Li, Y. Wang, P. Wan, and O. Frieder. Geometric spanners for wireless ad hoc networks. 14(4):408–421, 2003.

[2] F. M. auf der Heide, C. Schindelhauer, K. Volbert, and M. Grünewald. Energy, congestion and dilation in radio networks. In *Proc. of the 14th ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 230–237, 2002.

[3] P. Bose, L. Devroye, W. Evans, and D. Kirkpatrick. On the spanning ratio of Gabriel graphs and beta-skeletons. In *Proc. of Latin American Theoretical Informatics Conference (LATIN)*, 2002.

[4] D. Dubhashi, O. Häggström, A. Panconesi, and M. Sozio. Irrigating ad hoc networks in constant time. In *Proc. of the 17th ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 106–115, 2005.

[5] D. Eppstein. *Handbook of Computational Geometry*, chapter Spanning trees and spanners, pages 425–461. Elsevier, 2000.

[6] M. Grünewald, T. Lukovszki, C. Schindelhauer, and K. Volbert. Distributed maintenance of resource efficient wireless network topologies. In *European Conference on Parallel Computing (EUROPAR)*, pages 935–946, 2002.

[7] L. Jia, R. Rajaraman, and C. Scheideler. On local algorithms for topology control and routing in ad hoc networks. In *Proc. of the 15th ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, pages 220–229, 2003.

[8] R. Rajaraman. Topology control and routing in ad hoc networks: a survey. *SIGACT News*, 33(2):60–73, 2002.

[9] S. Rao and W. Smith. Approiximating geometrical graphs via spanners and banyans. In *Proc. of the 30th ACM Symp. on Theory of Computing (STOC)*, pages 540–550, 1998.

[10] C. Schindelhauer, K. Volbert, and M. Ziegler. Spanners, weak spanners, and power spanners for wireless networks. In *Proc. of 15th Annual International Symposium on Algorithms and Computation (ISAAC '04)*, pages 805–821, 2004.

[11] F. Xue and P. Kumar. The number of neighbors needed for connectivity of wireless networks. *Wireless Networks*, 10(2):169–181, 2004.

[12] A. C.-C. Yao. On constructing minimum spanning trees in k-dimensional space and related problems. *SIAM Journal on Computing*, 11:721–736, 1982.