7 Overlay Networks I – Wireless Systems

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. Since the communication range of the usually mobile nodes is limited, destination nodes are not always directly reachable. The data may have to be routed over intermediate nodes (multi-hop routing), and therefore every node has to have router capabilities. Such networks are called *ad hoc networks*. They impose higher requirements on routing algorithms such as adaptability to dynamic link changes and awareness of the limited energy in mobile nodes while maintaining high throughput and small delays. In this section we study the question of how to maintain wireless links between the nodes so that

- as a minimum requirement, every node is reachable from every other node (as long as this is in principle possible), i.e. the graph formed by the links is connected,
- or better, every node has a route along the links to every other node with a close to minimum possible energy consumption (over all possible sequences of hops).

Ideally, 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.

A naive approach for coming up with wireless links is that every node tries to establish wireless links with its k nearest neighbors for some integer k. However, Figure 1 demonstrates that it is easy to come up with examples for which the graph formed by the links would not be connected. So the naive approach does not work in general.

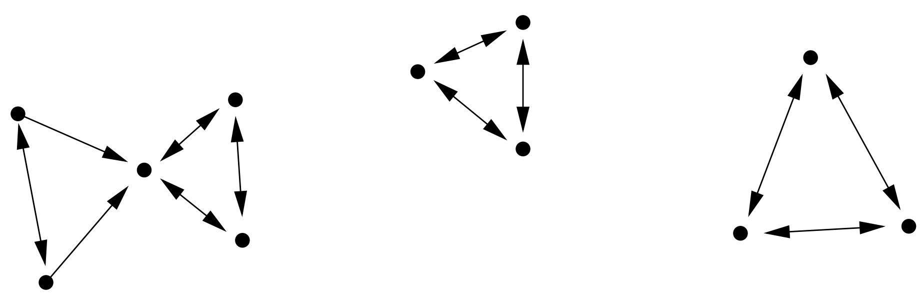

Figure 1: A counterexample for the naive approach with $k = 2$.

Fortunately, there are better approaches. To simplify the analysis, we assume that the wireless nodes are distributed in a perfect, 2-dimensional Euclidean space. Every node can transmit information over a potentially unlimited range, but the transmission of a packet over a range of r requires an energy consumption of r^d for some constant $d \geq 1$. In reality, d is usually in the range $[2, 5]$. Given two points $v = (x, y)$ and $w = (x', y')$ in a 2-dimensional Euclidean space, their *Euclidean distance* is defined as

$$|v, w| = \sqrt{(x - x')^2 + (y - y')^2} \, .$$

Furthermore, the *Euclidean length* of a path p traversing the nodes $v_0, v_1, \ldots, v_\ell$ is defined as

$$|p| = \sum_{i=0}^{\ell-1} |v_i, v_{i+1}| \ .$$

To avoid considering pathological cases, we sometimes assume that no two pairs of nodes have exactly the same Euclidean distance. A set of nodes with this property is called *non-degenerate*. The following definitions formalize the class of graphs we are searching for.

Definition 7.1 *Let V be a set of nodes in a Euclidean space. A graph $G = (V, E)$ is called a c-*spanner *if for all $u, v \in V$ there exists a (directed) path p from u to v in G with*

$$|p| \leq c \cdot |u, v| \ .$$

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

*G is a (c, d)-*power spanner *if for all $v, w \in V$ there is a path $p = (v = v_0, v_1, \ldots, v_\ell = w)$ from u to v in G with*

$$\sum_{i=0}^{\ell-1} |v_i, v_{i+1}|^d \leq c \cdot \min_{(v=u_0,u_1,\ldots,u_k=w)} \sum_{i=0}^{k-1} |u_i, u_{i+1}|^d \ .$$

If for all $d > 1$ there exists a constant c so that G is a (c, d)-power spanner, then we simply call G a power spanner.

Spanners belong to the class of *proximity graphs*, which are graphs that attempt to represent the spatial arrangement of nodes. We start with a general relationship between spanners and power spanners.

Theorem 7.2 ([2]) *Every c-spanner is a (c^d, d)-power spanner for every $d \geq 1$.*

Proof. Let $G = (V, E)$ be a c-spanner. Consider an arbitrary pair $v, w \in V$ and let $p = (v = v_0, v_1, \ldots, v_\ell = w)$ be an energy optimal path from v to w. Since G is a c-spanner, every pair of nodes (v_i, v_{i+1}), $i \in \{0, \ldots, \ell - 1\}$, has a path $q_i = (v_i = u_0, u_1, \ldots, u_k = v_{i+1})$ along edges of G with

$$\sum_{j=0}^{k-1} |u_j, u_{j+1}| \leq c \cdot |v_i, v_{i+1}| \ .$$

Hence, for every $d \geq 1$,

$$\sum_{j=0}^{k-1} |u_j, u_{j+1}|^d \leq \left(\sum_{j=0}^{k-1} |u_j, u_{j+1}| \right)^d \leq c^d \cdot |v_i, v_{i+1}|^d \ .$$

Connecting the paths q_i for all (v_i, v_{i+1}) gives a path q along edges of G with a power stretch factor of at most c^d, which proves the theorem. $\sqcap$

Hence, in order to prove that a graph is a power spanner, it suffices to prove that it is a spanner. We will look at three different ways of constructing spanners.

7.1 Gabriel graphs

We start with the class of Gabriel graphs.

Definition 7.3 *Let V be a set of nodes in a Euclidean space. The* Gabriel graph $GG(V)$ *of V consists of all edges $\{u, v\}$ with the property that the open sphere through u and v with diameter $|u, v|$ does not contain any other node from V.*

An example of a Gabriel graph is given in Figure 2.

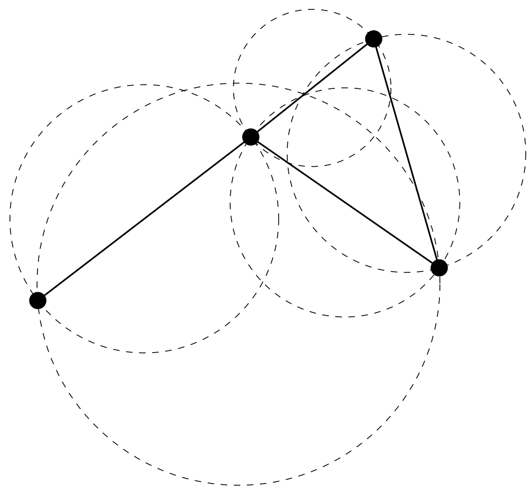

Figure 2: A Gabriel graph.

First, we may ask whether a Gabriel graph is always connected. This is shown by the following theorem.

Theorem 7.4 ([4]) *For every set of nodes V in the Euclidean space, $GG(V)$ is connected.*

Proof. For any pair of nodes $u, v \in V$, let $S_{u,v}$ denote the open sphere with diameter $|u, v|$ through u and v.

We prove the theorem by induction on the distances between the pairs of nodes. We start with a pair of lowest distance, say (u, w). There cannot be a node $v \in S_{u,w}$, because otherwise $|u, v| < |u, w|$, contradicting the minimality of $|u, w|$. Hence, $GG(V)$ contains the edge $\{u, w\}$, and therefore u and w are connected.

Now, suppose that we already know for the k pairs with lowest distance that each of them is connected. Then let (u, w) be the next pair to be considered. We distinguish between two cases:

1. There is some node $v \in S_{u,w}$: Since $|u, v| < |u, w|$ and $|v, w| < |u, w|$, it follows from our induction hypothesis that u and v and v and w must be connected. Hence, also u and w must be connected.

2. There is no node in $S_{u,w}$: Then $\{u, w\}$ is an edge in $GG(V)$, and therefore u and w must be connected.

Hence, all pairs of nodes in V must be connected, and therefore $GG(V)$ must be connected. $\sqcap$

The next two theorems demonstrate that the Gabriel graph is a bad spanner but a good power spanner.

Theorem 7.5 *The stretch factor of a Gabriel graph can be more than $n^{0.11}$.*

Proof. Consider the snowflake structure in Figure 3. As can be seen from the picture (see the nodes v and w), the stretch factor of the Gabriel graph is equal to

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

if $\epsilon > 0$ is sufficiently small.

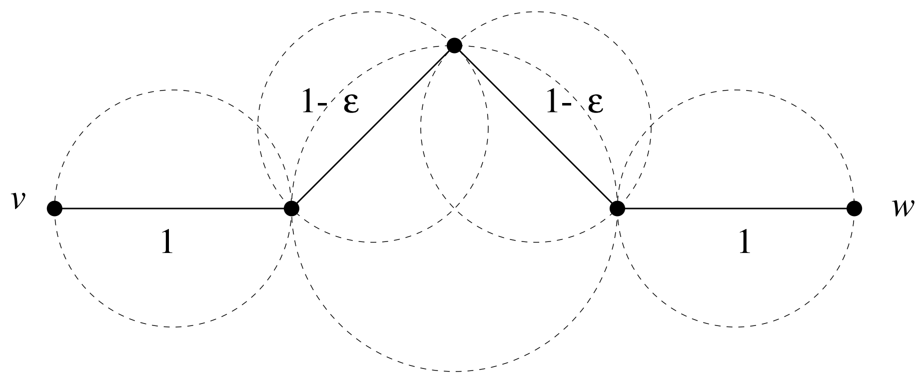

Figure 3: The snowflake structure.

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

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

$\sqcap$

With better techniques one can even create a counterexample with stretch factor $\Omega(\sqrt{n})$ [4].

Theorem 7.6 ([4]) *For every $d \geq 2$, the Gabriel graph is an optimal power spanner.*

Proof. Consider any pair of nodes $u, w \in V$ and let p be their optimal energy path. Consider an arbitrary edge $\{x, y\}$ in p. Suppose that there is a node $v \in S_{x,y}$. Then, by the Theorem of Thales, $|x,v|^d + |v,y|^d < |x,y|^d$ for every $d \geq 2$. Hence, replacing the edge $\{x,y\}$ in p by $\{x,v\}$ and $\{v,y\}$ would reduce its energy, which would contradict our assumption that p is an optimal energy path. Hence, there cannot be a node in $S_{x,y}$, and therefore $\{x,y\}$ must be an edge in the Gabriel graph. Thus, p must be a path in the Gabriel graph, which proves the theorem. $\sqcap$

However, Gabriel graphs have two problems:

- The degree may be very high (see Figure 4), and
- investigating open spheres between nodes in a wireless setting may not be easy.

Hence, we also look at other alternatives.

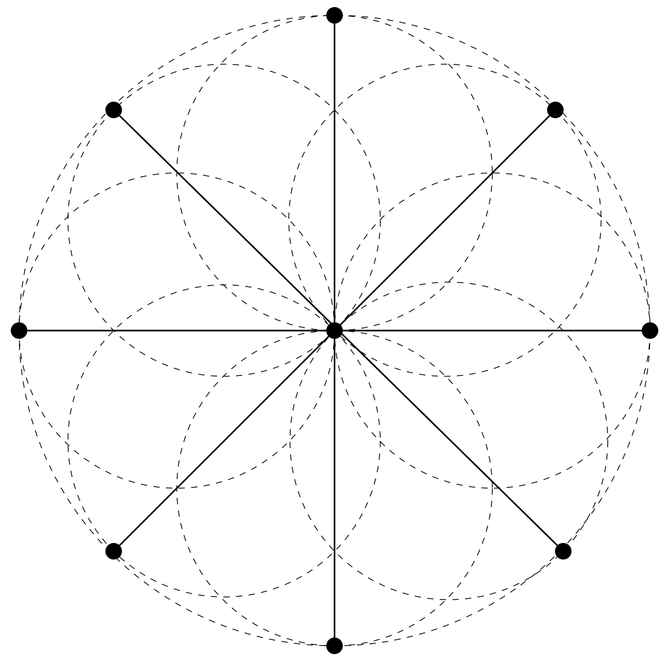

Figure 4: Gabriel graph for the unit sphere with one node in its center and all other nodes on its surface.

7.2 Yao graphs

The basic idea underlying the Yao graphs is to cut the space around each node into sectors of equal angle and to connect each node to the nearest neighbor in each of its sectors (see Figure 5). For any pair of nodes u, v, let $C_{u,v}$ denote the sector (or cone) of u containing v.

Definition 7.7 *Let V be a non-degenerate set of nodes in a 2-dimensional Euclidean space and let $k \in \mathbb{N}$. Suppose that the space around every node $v \in V$ is cut into k sectors with angle $\Theta = 2\pi/k$. Then the* Yao graph $YG_\Theta(V)$ *of V consists of the following set of edges:*

$$E = \{(u,v) \mid u,v \in V \textit{ and there is no } w \in V \textit{ with } w \in C_{u,v} \textit{ and } |u,w| < |u,v|\} .$$

The Yao graph is also known as Θ-graph. One can also define Yao graphs for higher dimensions, but we will restrict our attention to the 2-dimensional case.

Theorem 7.8 *If $\Theta = 2\pi/k$ with $k > 6$, then $YG_\Theta(V)$ is a spanner with stretch factor at most*

$$\frac{1}{1 - 2\sin(\Theta/2)} .$$

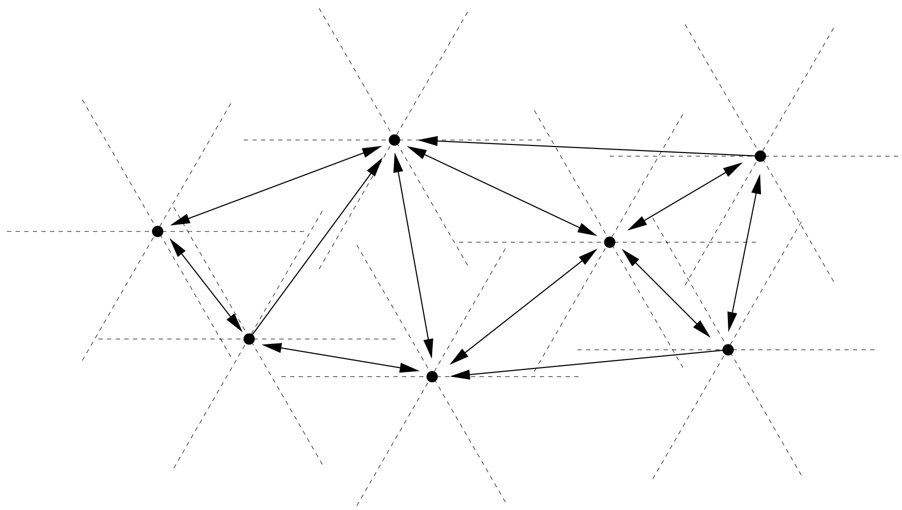

Figure 5: An example of a Yao graph.

Proof. We first need a lemma.

Lemma 7.9 *Let $p \in \mathbb{R}^2$ be a point and C be a sector originating at p. Furthermore, let q and r be two points in C with $|p,q| \leq |p,r|$. Then $|q,r| \leq |p,r| - (1 - 2\sin(\Theta/2))|p,q|$.*

Proof.

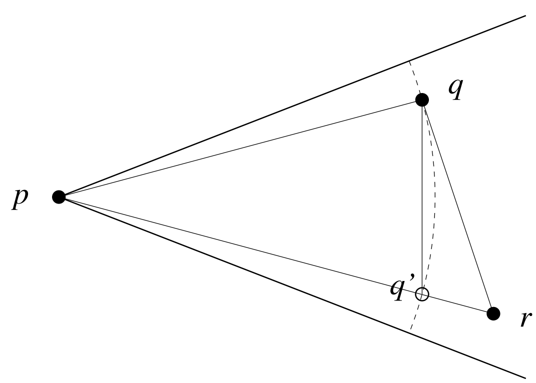

Figure 6: The sector of p that contains r.

Consider Figure 6. In this figure, q' represents the point on the line from p to r with the same distance to p as q. Applying the triangle inequality to q, q', and r, we get

$$|q,r| \leq |q,q'| + |q',r| \ . \tag{1}$$

$|q,q'|$ is certainly maximized if q and q' are on opposite sides of the sector. Hence,

$$|q,q'| \leq 2\sin(\Theta/2) \cdot |p,q| \ . \tag{2}$$

Moreover,

$$|q', r| = |p, r| - |p, q'| = |p, r| - |p, q| \ . \tag{3}$$

Plugging (2) and (3) into (1) yields

$$\begin{aligned} |q, r| &\leq 2\sin(\Theta/2) \cdot |p, q| + |p, r| - |p, q| \\ &= |p, r| - (1 - 2\sin(\Theta/2))|p, q| \ . \end{aligned}$$

$\sqcap$

Given a source-destination pair (s, t), consider the following strategy to get from s to t: Always take the edge whose other endpoint lies in the same sector as t.

Let the path obtained by this rule be $p = (s = v_0, v_1, \ldots, v_\ell = t)$. The path indeed ends at t, because for $k > 6$ we are guaranteed to have

$$|v_i, t| > |v_{i+1}, t|$$

for all i (see Figure 7).

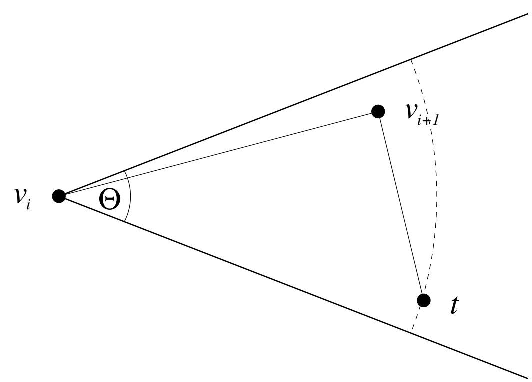

Figure 7: Figure illustrating that for $\Theta < \pi/3$, v_{i+1} is closer to t than v_i.

Using Lemma 7.9, it holds that

$$\sum_{i=0}^{\ell-1} |v_{i+1}, t| \leq \sum_{i=0}^{\ell-1} \left(|v_i, t| - (1 - 2\sin(\Theta/2))|v_i, v_{i+1}|\right) \ .$$

Rearranging the terms yields

$$\begin{aligned} \sum_{i=0}^{\ell-1} |v_i, v_{i+1}| &\leq \frac{1}{1 - 2\sin(\Theta/2)} \sum_{i=0}^{\ell-1} \left(|v_i, t| - |v_{i+1}, t|\right) \\ &= \frac{1}{1 - 2\sin(\Theta/2)} \cdot |s, t| \ , \end{aligned}$$

which proves the theorem. $\sqcap$

Combining this with Theorem 7.2 yields the following result.

Corollary 7.10 *If $\Theta = 2\pi/k$ with $k > 6$, then $YG_\Theta(V)$ is a (c,d)-power spanner for every $d \geq 1$ with*

$$c \leq \left(\frac{1}{1 - 2\sin(\Theta/2)}\right)^d .$$

A better result was shown by Li et al.

Theorem 7.11 ([4]) *If $\Theta = 2\pi/k$ with $k > 6$, then $YG_\Theta(V)$ is a (c,d)-power spanner for every $d \geq 1$ with*

$$c \leq \frac{1}{1 - (2\sin(\Theta/2))^d} .$$

The drawback of the Yao graph is that although its out-degree is at most k, its in-degree may be as high as $n-1$ (consider, for example, the unit sphere with one node in its center and all other nodes on its surface). Various sub-graphs of the Yao graph have been suggested to remove this drawback. We will present two of them here.

Definition 7.12 *The* sparsified Yao graph *$SpYG_\Theta(V)$ is a sub-graph of $YG_\Theta(V)$ with edge set*

$$E = \{(u,v) \in E(YG_\Theta(V)) \mid \textit{for all } w \in V \textit{ with } (w,v) \in E(YG_\Theta(V)) \textit{ and } w \in C_{v,u}\textit{: } |v,w| > |v,u|\} .$$

In words, the sparsified Yao graph only keeps the shortest of all incoming edges of v for every node v. Hence, the sparsified Yao graph has an in-degree of at most k and an outdegree of at most k, and therefore a degree of at most $2k$.

Definition 7.13 *The* symmetric Yao graph *$SyYG_\Theta(V)$ is a sub-graph of $YG_\Theta(V)$ with edge set*

$$E = \{(u,v) \in E(YG_\Theta(V)) \mid (v,u) \in E(YG_\Theta(V))\} .$$

In words, the symmetric Yao graph only keeps an edge (u,v) if not only v is the nearest neighbor of u in $C_{u,v}$ but also u is the nearest neighbor of v in $C_{v,u}$. Hence, the symmetric Yao graph has a degree of at most k. Obviously,

$$SyYG_\Theta(V) \subseteq SpYG_\Theta(V) \subseteq YG_\Theta(V) .$$

Thus, it suffices to prove connectivity for $SyYG_\Theta(V)$ in order to prove connectivity for both variants of the Yao graph.

Theorem 7.14 ([2]) *For all non-degenerate node sets V and $k > 6$, $SyYG_\Theta(V)$ is connected.*

Proof. The proof is very similar to the proof for the Gabriel graph. We prove the theorem by induction on the distances between the pairs of nodes, starting with the pair of lowest distance, say (u,w). In this case, w must be the nearest neighbor in a sector of u and u must be the nearest neighbor in a sector of$w$. Hence, $(u,w) \in E(SyYG_\Theta(V))$, and therefore u and w are connected.

Now, suppose that we already know for the k pairs with lowest distance that each of them is connected. Then let (u,w) be the next pair to be considered. We distinguish between three cases:

1. There is a node $v \in C_{u,w}$ that is closer to u then w: Since $k > 6$, it holds in this case that $|u, v| < |u, w|$ and $|v, w| < |u, w|$. Hence, according to our induction hypothesis, u and v and v and w must be connected. Thus, also u and w must be connected.

2. There is a node $v \in C_{w,u}$ that is closer to w then u: Using the same arguments as for the first case, it follows that also in this case u and w must be connected.

3. None of the two cases above hold: Then w must be the nearest neighbor of u in $C_{u,w}$ and u must be the nearest neighbor of w in $C_{w,u}$. Hence, $(u, w) \in E(SyYG_\Theta(V))$, and therefore u and w are connected.

Thus, all pairs of nodes in V must be connected, and therefore $SyYG_\Theta(V)$ must be connected. $\sqcap$

Unfortunately, the symmetric Yao graph is not a good power spanner.

Theorem 7.15 ([1]) *The symmetric Yao graph is not a (c, d)-power spanner for any constant c and any $d \geq 1$.*

However, the sparsified Yao graph is a power spanner.

Theorem 7.16 ([3]) *If k is a sufficiently large constant and $d \geq 2$, then the sparsified Yao graph is a (c, d)-power spanner for some constant c.*

The proofs of both theorems are quite involved, and therefore we do not present them here.

7.3 Hierarchical layer graphs

Although the sparsified Yao graph satisfies now all of our original requirements, it has the drawback that it can cause a large amount of update work if the node set V changes (which may happen in mobile ad hoc networks).

Theorem 7.17 ([2]) *There are node sets V where $\Theta(|V|)$ edges may have to be changed if an enter or leave operation happens in a Yao graph, sparsified Yao graph, or symmetric Yao graph.*

Proof. A bad situation for all Yao based graphs occurs, when two rows of nodes $U = \{u_1, \ldots, u_m\}$ and $V = \{v_1, \ldots, v_m\}$ are placed on two parallel lines so that the edge (u_i, v_i) is orthogonal to (u_1, u_n) and (v_1, v_n) and all nodes in U are in the same sector of a node in V and vice versa.

In this situation,we have $m = n/2$ edges for all Yao-based topologies, which all have to be erased if a node w pops up in the middle of the network. The inverse situation occurs if we remove this node. $\sqcap$

To remove this problem, *Hierarchical Layer Graphs* (HL-graphs) have been suggested [2]. An HL-graph consists of w layers $L_1, L_2, \ldots, L_w$ for some $w \geq 1$. The lowest layer contains all nodes in V. The node set of a higher layer is a subset of the node set of a lower layer, and the highest layer contains only a single node, i.e.

$$V = V(L_1) \supseteq V(L_2) \supseteq \ldots \supseteq V(L_w) = \{v_0\} .$$

The crucial property of these layers is that in each layer L_i, nodes obey a minimum distance:

$$\forall u, v \in V(L_i) : \ |u, v| \geq r_i \ .$$

Furthermore, all nodes in the next lower layer must be covered by this distance:

$$\forall u \in V(L_i) \ \exists v \in V(L_{i+1}) : \ |u, v| \leq r_{i+1} \ .$$

The construction uses parameters $\alpha \geq \beta > 1$, where starting with some $r_0 < \min_{u,v \in V} |u, v|$ we use radii $r_i = \beta^i \cdot r_0$ and we define in layer L_i the edge set $E(L_i)$ by

$$E(L_i) = \{(u, v) \mid u, v \in V(L_i) \ \wedge \ |u, v| \leq \alpha \cdot r_i\} \ .$$

If $(\max_{u,v} |u, v|)/(\min_{u,v} |u, v|) = n^{O(1)}$ (which is a reasonable assumption), then the number of layers we need is $O(\log n)$. Combining this with the fact that the degree of L_i is at most $(2\alpha + 1)^2$ [1], it follows that if α is constant, the degree of every node is $O(\log n)$. In addition to this, the following results can be shown.

Theorem 7.18 ([2]) *If $\alpha > 2\beta/(\beta - 1)$, the HL-graph is a c-spanner for*

$$c = \beta \cdot \frac{\alpha(\beta - 1) + 2\beta}{\alpha(\beta - 1) - 2\beta} \ .$$

Theorem 7.19 *For every set of nodes V with $(\max_{u,v} |u, v|)/(\min_{u,v} |u, v|) = n^{O(1)}$ in a 2-dimensional Euclidean space, every enter and leave operation only requires $O(\log n)$ changes in the edge set of the HL-graph.*

Hence, HL-graphs are much more adaptive to changes than Yao graphs.

References

[1] 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.

[2] 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.

[3] L. Jia, R. Rajaraman, and C. Scheideler. On local algorithms for topology control and routing in ad hoc networks. Technical report, Johns Hopkins University, 2002.

[4] X.-Y. Li, P.-J. Wan, and Y. Wang. Power efficient and sparse spanner for wireless ad hoc networks. In *IEEE International Conference on Computer Communications and Networks*, 2002.