Spanners, Weak Spanners, and Power Spanners for Wireless Networks

Christian Schindelhauer*, Klaus Volbert*, and Martin Ziegler*

Heinz Nixdorf Institute, Paderborn University
Institute of Computer Science,
{schindel, kvolbert, ziegler}@uni-paderborn.de

Abstract. For $c \in \mathbb{R}$, a *c-spanner* is a subgraph of a complete Euclidean graph satisfying that between any two vertices there exists a path of weighted length at most c times their geometric distance. Based on this property to approximate a complete weighted graph, sparse spanners have found many applications, e.g., in FPTAS, geometric searching, and radio networks. In a *weak* c-spanner, this path may be arbitrary long but must remain within a disk of radius c-times the Euclidean distance between the vertices. Finally in a *c-power* spanner, the total energy consumed on such a path, where the energy is given by the sum of the *squares* of the edge lengths on this path, must be at most c-times the square of the geometric distance of the direct link.

While it is known that any c-spanner is also both a weak C_1-spanner and a C_2-power spanner (for appropriate C_1, C_2 depending only on c but not on the graph under consideration), we show that the converse fails: There exists a family of c_1-power spanners that are no weak C-spanners and also a family of weak c_2-spanners that are no C-spanners for any fixed C (and thus no uniform spanners, either). However the deepest result of the present work reveals that any weak spanner *is* also a uniform power spanner. We further generalize the latter notion by considering (c, δ)-power spanners where the sum of the δ-th powers of the lengths has to be bounded; so $(\cdot, 2)$-power spanners coincide with the usual power spanners and $(\cdot, 1)$-power spanners are classical spanners. Interestingly, these $(\cdot, \delta)$-power spanners form a strict hierarchy where the above results still hold for any $\delta \geq 2$; some even hold for $\delta > 1$ while counterexamples exist for $\delta < 2$. We show that every self-similar curve of fractal dimension $d > \delta$ is no (C, δ)-power spanner for any fixed C, in general.

1 Motivation

Spanners have appeared in Computer Science with the advent of Computational Geometry [4, 18], raised further in interest as a tool for approximating NP-hard problems [13] and, quite recently, for routing and topology control in ad-hoc networks [1, 12, 8, 7, 11]. Roughly speaking, they approximate the complete Euclidean graph on a set of geometric vertices while having only linearly many edges. The formal condition for a c-spanner

* Partially supported by the DFG-Sonderforschungsbereich 376 and by the EU within 6th Framework Programme under contract 001907 "Dynamically Evolving, Large Scale Information Systems" (DELIS).

R. Fleischer and G. Trippen (Eds.): ISAAC 2004, LNCS 3341, pp. 805–821, 2004.

$G = (V, E)$ is that between any two $\mathbf{u}, \mathbf{v} \in V$, the edge $(\mathbf{u}, \mathbf{v})$ may be absent provided there exists a path in G from $\mathbf{u}$ to $\mathbf{v}$ of length at most c-times the Euclidean distance between $\mathbf{u}$ and $\mathbf{v}$; see Figure 1(a). In particular, this path remains within a circle around $\mathbf{u}$ of radius c. For applications in geometric searching [5], it has turned out that graphs with the latter, weaker condition suffice; see Figure 1(b). In [10] such weak spanners are used to approximate congestion-optimal radio networks. Several constructions yield both spanners [6] and weak spanners [5] with arbitrarily describable approximation ratio. Among them, some furthermore benefit from nice locality properties which led to successful applications in ad-hoc routing networks [7, 9, 17, 16]. However in order to restrict the power consumption during such a communication (which for physical reasons grows quadratically with the Euclidean length of each link), one is interested in routing paths, say between $\mathbf{u}$ and $\mathbf{v}$, whose sum of *squares* of lengths of the individual steps is bounded by c-times the square of the Euclidean distance between $\mathbf{u}$ and $\mathbf{v}$; see Figure 1(c). Such graphs are known as *c-power* spanners [7, 9].

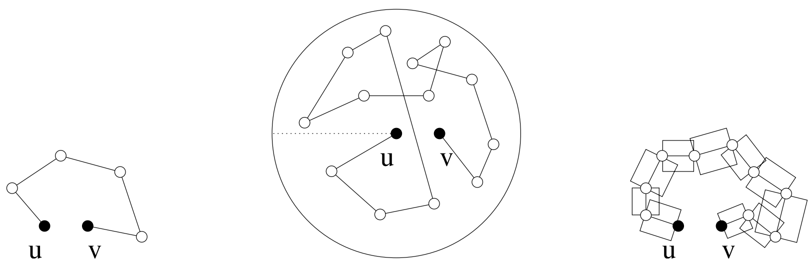

(a) Bounded length (b) Bounded radius (c) Bounded energy

Fig. 1. Spanner, Weak Spanner and Power Spanner

Finally, when power consumption is of minor interest but the routing time is dominated by the number of individual steps, sparse graphs are desired which between any vertices $\mathbf{u}$ and $\mathbf{v}$ provides a path containing at most c further vertices. These are the so-called *c-hop* spanners [1]. In this paper, we investigate the relations between these various types of spanners. Observe that any strongly connected finite geometric graph is a C-spanner for *some*[1] value C. Therefore the question on the relation between spanners and weak spanners rather asks whether any weak c-spanner is a C-spanner for some value C depending *only* on c. Based on a construction from [3], we answer this to the negative. For some weak c-spanners it is proved that they are also C-power spanners for some value C [7, 8] using involved constructions. One major contribution of our work generalizes and simplifies such results by showing that in the plane in fact *any* weak c-spanner is a C-power spanner with $C = \mathcal{O}(c^8)$. Moreover, we investigate the notion of a (c, δ)-power spanner [7] which

[1] Consider for any pair $\mathbf{u}, \mathbf{v}$ of vertices some path from $\mathbf{u}$ to $\mathbf{v}$ and the ratio of its length to the distance between $\mathbf{u}$ and $\mathbf{v}$; then taking for C the maximum over the (finitely many) pairs $\mathbf{u}, \mathbf{v}$ will do.

- for $\delta = 1$ coincides with c-spanners
- for $\delta = 2$ coincides with (usual) c-power spanners
- for $\delta = 0$ coincides with c-hop spanners, i.e. graphs with diameter c
- for $\delta > 2$ reflects transmission properties of radio networks (e.g., for δ up to 6 or even 8) .

We show that these form a strict hierarchy: For $\Delta > \delta > 0$, any (c, δ)-power spanner is also a (C, Δ)-power spanner with C depending only on c and Δ/δ; whereas we give examples of (C, Δ)-power spanners that are no (c, δ)-power spanner for any fixed c. Our main contribution is that any weak c-spanner is also a (C, δ)-power spanner for arbitrary $\delta \geq 2$ with C depending on c and δ only. We finally show that this claim is best possible by presenting, for arbitrary $\delta < 2$, weak c-spanners which are no (C, δ)-power spanner for any fixed C.

This paper is organized as follows: Section 2 formally defines the different types of spanners under consideration. Section 3 shows that, while any c-spanner is also a weak c-spanner, a weak c-spanner is in general no C-spanner for any C depending just on c. Section 4 similarly reveals the relations between spanners and power spanners. The central Section 5 of the this work investigates the relation between weak spanners and power spanners. Theorem 3 gives an example of a power spanner which is no weak spanner. Our major contributions then prove that, surprisingly, any weak c-spanner is also a C-power spanner with C depending only on c. For different values of δ, we obtain different upper bounds to C in terms of c: For $\delta = 2$ (power spanners in the original sense), we show $C \leq \mathcal{O}(c^8)$, see Theorem 5; for $\delta > 2$, we have $C \leq \mathcal{O}\big(c^{2+\delta}/(1 - 2^{2-\delta})\big)$, see Theorem 4. However for $\delta < 2$, we present counter-examples of unbounded C, that is, in this case provably *not* any weak c-spanner is a (C, δ)-power spanner. Further, we generalize our construction and analysis to self-similar fractal curves. Section 6 finally shows that for different δ, the respective classes of $(\cdot, \delta)$-power spanners form a strict hierarchy. In Section 7 we discuss extensions of our results to higher-dimensional cases, before we conclude this work presenting applications of our results concerning power-efficient wireless networks in Section 8.

2 Preliminaries

We focus on the two-dimensional case, that is, directed graphs $G = (V, E)$ with finite $V \subseteq \mathbb{R}^2$; extensions to higher dimensions are discussed in Section 7. Let $|\mathbf{u} - \mathbf{v}|$ denote the Euclidean distance between $\mathbf{u}, \mathbf{v} \in V$. A *path* from $\mathbf{u}$ to $\mathbf{v}$ in G is a finite sequence $P = (\mathbf{u} = \mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_\ell = \mathbf{v})$ of vertices $\mathbf{u}_i \in V$ such that $(\mathbf{u}_{i-1}, \mathbf{u}_i) \in E$ for all $i = 2, \ldots, \ell$. Occasionally, we also encounter the more general situation of a path from $\mathbf{u}$ to $\mathbf{v}$ not necessarily in G; this means that $\mathbf{u}_i \in V$ still holds but the requirement $(\mathbf{u}_i, \mathbf{u}_{i+1}) \in E$ is dropped. The *radius* of a path P is the real number $\max_{i=2}^{\ell} |\mathbf{u} - \mathbf{u}_i|$. The (Euclidean) *length* of P is given by $\sum_{i=2}^{\ell} |\mathbf{u}_i - \mathbf{u}_{i-1}|$; the *hop length* is $\ell - 1$; for $\delta \geq 0$,

$$\|P\|^\delta \quad := \quad \sum_{i=2}^{\ell} |\mathbf{u}_i - \mathbf{u}_{i-1}|^\delta \tag{1}$$

denotes the *δ-cost* of P. The length is just the 1-cost whereas the hop length coincides with the 0-cost.

Definition 1. *Let $G = (V, E)$ be a geometric directed graph with finite $V \subseteq \mathbb{R}^2$ and $c > 0$. G is a c-*spanner*, if for all $\mathbf{u}, \mathbf{v} \in V$ there is a path P from $\mathbf{u}$ to $\mathbf{v}$ in G of length $\|P\|^1$ at most $c \cdot |\mathbf{u} - \mathbf{v}|$. G is a* weak c-spanner*, if for all $\mathbf{u}, \mathbf{v} \in V$ there is a path P from $\mathbf{u}$ to $\mathbf{v}$ in G of radius at most $c \cdot |\mathbf{u} - \mathbf{v}|$. For $\delta \geq 0$, G is a (c, δ)-*power spanner *if for all $\mathbf{u}, \mathbf{v} \in V$ there is a path P from $\mathbf{u}$ to $\mathbf{v}$ in G of δ-cost $\|P\|^\delta$ at most $c \cdot |\mathbf{u} - \mathbf{v}|^\delta$. G is a c-*power spanner*, if G is a $(c, 2)$-power spanner. The factor c is called* length stretch factor, weak stretch factor *or* power stretch factor, *respectively.*

Informally (see Figure 1), in a c-spanner there exists between two arbitrary vertices a path of length at most c-times the Euclidean distance between these vertices (bounded length). In a weak c-spanner, this path may be arbitrary long but must remain within a disk of radius c-times the Euclidean distance between the vertices (bounded radius). Finally in a c-power spanner, the energy consumed on such a path (e.g., the sum of the *squares* of the lengths of its constituting edges) must be at most c-times the one consumed on a putative direct link (bounded cost). Sometimes we shorten the notion of spanner, weak spanner and power spanner and omit constant parameters. So, if we say that a family of graphs is a *spanner*, then there exists a constant c such that all its members are c-spanners.

The attentive reader might have observed that our Definition 1 does not exactly match that from [7]. The latter required that the 2-cost of some path P from $\mathbf{u}$ to $\mathbf{v}$ in G is bounded by c-times the 2-cost of *any* path Q (not necessarily in G) from $\mathbf{u}$ to $\mathbf{v}$. However, both approaches are in fact equivalent:

Lemma 1. *Let $G = (V, E)$ be a (c, δ)-power spanner, $\mathbf{u}, \mathbf{v} \in V$, and let Q denote some path Q (not necessarily in G) from $\mathbf{u}$ to $\mathbf{v}$ of minimum δ-cost. Then there is a path P in G from $\mathbf{u}$ to $\mathbf{v}$ of δ-cost $\|P\|^\delta$ at most $c \cdot \|Q\|^\delta$.*

Proof. Let $Q = (\mathbf{u} = \mathbf{q}_1, \ldots, \mathbf{q}_L = \mathbf{v})$. For each $i = 2, \ldots, L$ there exists by presumption a path P_i in G from $\mathbf{q}_{i-1}$ to $\mathbf{q}_i$ of δ-cost at most $c \cdot |\mathbf{q}_i - \mathbf{q}_{i-1}|^\delta$. The concatenation of all these paths yields a path P from $\mathbf{u}$ to $\mathbf{v}$ in G with δ-cost $\|P\|^\delta$ at most $c \cdot \|Q\|^\delta$.

3 Spanners Versus Weak Spanners

Every c-spanner is also a weak c-spanner. Our first result shows that the converse fails in general.

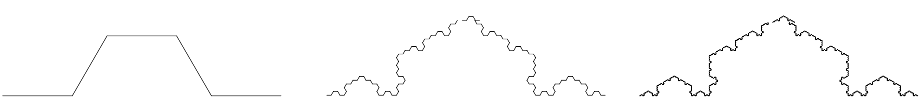

Fig. 2. Construction of EPPSTEIN provably yields no spanner but a weak spanner

Theorem 1. *There is a family of graphs $G = (V, E)$ with $V \subseteq \mathbb{R}^2$ all of which are weak $(\sqrt{3} + 1/2)$-spanners but no C-spanners for any fixed $C \in \mathbb{R}$.*

Proof. We show the claim using the fractal construction presented in [3] (see Figure 2). We briefly review its recursive definition which is similar to that of a KOCH Curve. At the beginning there are two vertices with distance 1. In the following steps we replace each edge by 5 new edges of equal length as follows: one horizontal, one at angle $\pi/4$, a second horizontal, another one at angle $-\pi/4$ and a third horizontal. After i steps we have a graph consisting of 5^i edges and $5^i + 1$ vertices. As shown in [3] this graph has unbounded length stretch factor. We argue that there exists a constant c such that it is a weak c-spanner. It is known that the area under the constructed curve is bounded by a constant and that the path between two vertices $u, v \in V$ lies completely in a disk around the midpoint of the segment between u and v with radius at most $(2 \cdot \sqrt{3}/2) = \sqrt{3}$ (see KOCH's Snowflake, Figure 6). Applying Observation 2 proves the claim.

The following observation says that, up to constants, it makes no difference in the definition of a weak spanner whether the radius is bounded with respect to center $\mathbf{u}$ (the starting one of the two points) or with respect to center $(\mathbf{u}+\mathbf{v})/2$ (the midpoint of the segment between the two points).

Observation 2. *Let $P = (\mathbf{u} = \mathbf{u}_1, \ldots, \mathbf{u}_\ell = \mathbf{v})$ be a path in the geometric graph $G = (V, E)$ such that $|\mathbf{u} - \mathbf{u}_i| \leq c \cdot |\mathbf{u} - \mathbf{v}|$ for all $i = 1, \ldots, \ell$. Then, $\mathbf{w} := (\mathbf{u}+\mathbf{v})/2$ satisfies by the triangle inequality*

$$|\mathbf{w}-\mathbf{u}_i| \quad = \quad |\mathbf{u}-\mathbf{u}_i+(\mathbf{v}-\mathbf{u})/2| \quad \leq \quad |\mathbf{u}-\mathbf{u}_i|+|\mathbf{v}-\mathbf{u}|/2 \quad \leq \quad (c+\tfrac{1}{2})\cdot|\mathbf{u}-\mathbf{v}| \ .$$

Conversely if P has $|\mathbf{w} - \mathbf{u}_i| \leq c \cdot |\mathbf{u} - \mathbf{v}|$ for all i, then

$$|\mathbf{u}-\mathbf{u}_i| \quad = \quad |\mathbf{w}-\mathbf{u}_i+(\mathbf{u}-\mathbf{v})/2| \quad \leq \quad |\mathbf{w}-\mathbf{u}_i|+|\mathbf{u}-\mathbf{v}|/2 \quad \leq \quad (c+\tfrac{1}{2})\cdot|\mathbf{u}-\mathbf{v}| \ .$$

4 Spanners Versus Power Spanners

In [7] it is shown that, for $\delta > 1$, every c-spanner is also a (c^δ, δ)-power spanner. However, conversely, for any $\delta > 1$, there are (c, δ)-power spanners which are no C-spanners for any fixed C: This follows from Theorem 3 as any C-spanner is a weak C-spanner as well.

5 Weak Spanners Versus Power Spanners

Now, we turn to the main contribution of the present paper and present our results concerning the relation between weak spanners and power spanners. Surprisingly, it turns out that any weak c-spanner is also a C-power spanner for some C depending only on c. But first observe that the converse in general fails:

Theorem 3. *In the plane and for any $\delta > 1$, there is a family of (c, δ)-power spanners which are no weak C-spanners for any fixed C.*

Proof. Let $V := \{\mathbf{u} = \mathbf{v}_1, \ldots, \mathbf{v} = \mathbf{v}_n\}$ be a set of n vertices placed on a circle scaled such that the Euclidean distance between $\mathbf{u}$ and $\mathbf{v}$ is 1 and $|\mathbf{v}_i - \mathbf{v}_{i+1}| = 1/i$ for all

$i = 1, \ldots, n-1$. Now, consider the graph $G = (V, E)$ with edges $(\mathbf{v}_i, \mathbf{v}_{i+1})$. First observe that G is a (c, δ)-power spanner with c independent of n. Indeed, its δ-power stretch factor is dominated by the δ-cost of the (unique) path P in G from $\mathbf{u}$ to $\mathbf{v}$ which amounts to

$$\|P\|^\delta \quad = \quad \sum_{i=1}^{n-1}(1/i)^\delta \quad \leq \quad \sum_{i=i}^{\infty}(1/i)^\delta \quad =: \quad c$$

a convergent series since $\delta > 1$. This is compared to the cost of the direct link from $\mathbf{u}$ to $\mathbf{v}$ of 1. On the other hand, the Euclidean length (that is, the 1-cost) of the polygonal chain from $\mathbf{u}$ to $\mathbf{v}$ is given by the unbounded harmonic series $\sum_{i=1}^{n-1}(1/i) = \Theta(\log n)$. Therefore also the radius of this polygonal chain cannot be bounded by any C independent of n, either.

In the sequel, we show that, conversely, any weak c-spanner is a (C, δ)-power spanner for both $\delta > 2$ (Section 5.1) and $\delta = 2$ (Section 5.2) with C depending only on c and δ. A counter-example in Section 5.3 reveals that this however does not hold for $\delta < 2$.

5.1 Weak Spanner Implies Power Spanner for $\delta > 2$

In this subsection, we show that any weak c-spanner is also a (C, δ)-power spanner for any $\delta > 2$ with C depending only on c and δ. By its definition, between vertices $\mathbf{u}, \mathbf{v}$, there exists a path P in G from $\mathbf{u}$ to $\mathbf{v}$ that remains within a disk around $\mathbf{u}$ of radius $c \cdot |\mathbf{u} - \mathbf{v}|$. However on the course of this path, two of its vertices $\mathbf{u}'$ and $\mathbf{v}'$ might come very close so that P, considered as a subgraph of G, in general is no weak c-spanner. On the other hand, G being a weak c-spanner, there exists also between $\mathbf{u}'$ and $\mathbf{v}'$ a path P' of small radius. Based on such repeated applications of the weak spanner property, we first assert the existence of a path which, considered as a subgraph of G, is weak $2c$-spanner.

Definition 2. *Let $G = (V, E)$ be a directed geometric graph and $e_1 := (\mathbf{u}_1, \mathbf{v}_1)$, $e_2 := (\mathbf{u}_2, \mathbf{v}_2)$ two of its edges. By their* distance *we mean the number*

$$\min\left\{|\mathbf{u}_1 - \mathbf{u}_2|, |\mathbf{v}_1 - \mathbf{v}_2|, |\mathbf{u}_1 - \mathbf{v}_2|, |\mathbf{v}_1 - \mathbf{u}_2|\right\} \ ,$$

that is, the Euclidean distance of a closest pair of their vertices (see Figure 3(b)).

Lemma 2. *Let $G = (V, E)$ be a weak c-spanner and $\mathbf{u}, \mathbf{v} \in V$. Then there is a path P from $\mathbf{u}$ to $\mathbf{v}$ in G which, as a subgraph of G, is a weak $2c$-spanner.*

Proof. The idea is to take the path P asserted by the weak spanner property for $\mathbf{u}$ and $\mathbf{v}$ and to, for any pair $\mathbf{u}', \mathbf{v}'$ of vertices on P for which P violates that property, locally replace that part of P by a path from $\mathbf{u}'$ to $\mathbf{v}'$ in G. However for these iterated improvements to eventually terminate, we perform them in decreasing order of the lengths of the edges involved.

W.l.o.g. we assume $|\mathbf{u} - \mathbf{v}| = 1$. Since G is a weak c-spanner there exists a path $P = (\mathbf{u} = \mathbf{u}_1, \ldots, \mathbf{u}_\ell = \mathbf{v})$ from $\mathbf{u}$ to $\mathbf{v}$ in G that lies completely within a disk

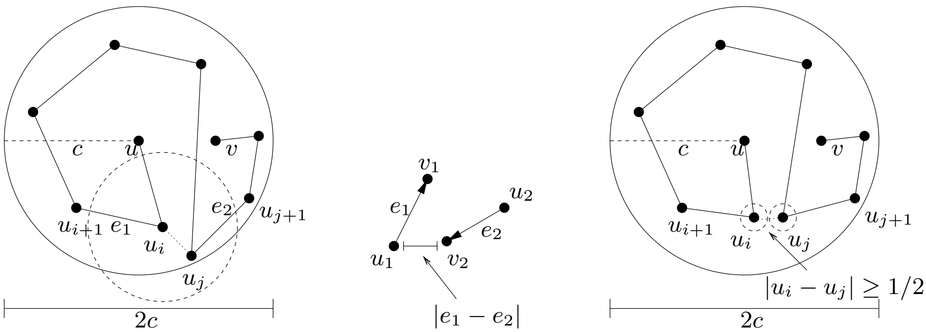

(a) Improving paths (b) Edge distance (c) Counting edges

Fig. 3. Construction and analysis of a path with low power stretch in a weak spanner

around $\mathbf{u}$ of radius c. In particular, any edge on this path has a length of at most $2c$, see Figure 3(a).

Now, consider all edges on this path of length between c and $2c$. For any pair $e_1 = (\mathbf{u}_i, \mathbf{u}_{i+1})$ and $e_2 = (\mathbf{u}_j, \mathbf{u}_{j+1})$ with $j > i$ closer than $\frac{1}{2}$ (Definition 2), improve that part of P by replacing it with a path according to the weak c-stretch property. Observe that, since the improvement is applied to vertices of distance at most $\frac{1}{2}$, this sub-path remains within a disk of radius $c/2$; in particular any edge introduced to P has length at most c and thus does not affect the edges of length between c and $2c$ currently considered. Moreover, after having performed such improvements to all edges of length between c and $2c$, the resulting path P', although it might now leave the disk around $\mathbf{u}$ of radius c, it does remain within radius $c + c/2$.

Next, we apply the same process to edges of length between c and $c/2$ and perform improvements on those closer than $\frac{1}{4}$. The thus obtained path P'' remains within a disk of radius $c+c/2+c/4$ while, for any pair of vertices $\mathbf{u}'$ and $\mathbf{v}'$ improved in the previous phase, the sub-path between them might increase in radius from $c \cdot |\mathbf{u}' - \mathbf{v}'|$ to at most $(c + c/2) \cdot |\mathbf{u}' - \mathbf{v}'|$.

As G is a finite graph, repeating this process for edges of length between $c/2$ and $c/4$ and so on, will eventually terminate and yield a path $\tilde{P}$ from $\mathbf{u}$ to $\mathbf{v}$ remaining within a disk of radius $c + c/2 + c/4 + \ldots = 2c$. Moreover, for *any* pair of vertices $\mathbf{u}', \mathbf{v}'$ in P, the sub-path between them has radius at most $(c+c/2+c/4+\ldots)\cdot|\mathbf{u}'-\mathbf{v}'|$ which proves that $\tilde{P}$ is indeed a weak $2c$-spanner.

Lemma 3. *Let $P = (\mathbf{u}_1, \ldots, \mathbf{u}_\ell)$ be a weak $2c$-spanner, $\mathbf{u}_i \in \mathbb{R}^2$, $|\mathbf{u}_1 - \mathbf{u}_\ell| = 1$. Then, P contains at most $(8c+1)^2$ edges of length greater than c; more generally, P contains at most $(8c + 1)^2 \cdot 4^k$ edges of length greater than $c/2^k$.*

Proof. Consider two edges $(\mathbf{u}_i, \mathbf{u}_{i+1})$ and $(\mathbf{u}_j, \mathbf{u}_{j+1})$ on P both of length at least c with $j > i$. P being a $2c$-weak spanner implies that, between vertices $\mathbf{u}_i$ and $\mathbf{u}_j$, the sub-path in P from $\mathbf{u}_i$ to $\mathbf{u}_j$ (which is unique and passes through $\mathbf{u}_{i+1}$), satisfies $c \leq |\mathbf{u}_i - \mathbf{u}_{i+1}| \leq 2c \cdot |\mathbf{u}_i - \mathbf{u}_j|$; hence, $|\mathbf{u}_i - \mathbf{u}_j| \geq \frac{1}{2}$, see Figure 3(c). In particular, placing an Euclidean disk B_i of radius $\frac{1}{4}$ around each starting vertex $\mathbf{u}_i$ of an edge of length at least c results in these disks being mutually disjoint. If m denotes the number of edges of length at least c, these disks thus cover a total area of $m\pi(\frac{1}{4})^2$. On the other hand, as all $\mathbf{u}_i$ lie within a single disk around $\mathbf{u}_1$ of radius $2c$, all disks B_i together cover an area of at most $\pi(2c+\frac{1}{4})^2$. Therefore,

$$m \quad \leq \quad \frac{\pi(2c+\frac{1}{4})^2}{\pi(\frac{1}{4})^2} \quad = \quad (8c+1)^2 \; .$$

For edges $(\mathbf{u}_i, \mathbf{u}_{i+1})$ and $(\mathbf{u}_j, \mathbf{u}_{j+1})$ on P longer than $c/2^k$, one similarly obtains $|\mathbf{u}_i - \mathbf{u}_j| \geq 2^{-k-1}$ so that, here, Euclidean disks of radius 2^{-k-2} can be placed mutually disjoint within the total area of $\pi(2c+2^{-k-2})^2$.

Theorem 4. *Let $G = (V, E)$ be a weak c-spanner with $V \subseteq \mathbb{R}^2$. Then G is a (C, δ)-power spanner for $\delta > 2$ where $C := (8c+1)^2 \cdot \dfrac{(2c)^\delta}{1-2^{2-\delta}}$.*

Proof. Fix $\mathbf{u}, \mathbf{v} \in V$, w.l.o.g. $|\mathbf{u}-\mathbf{v}| = 1$. In the following we analyze the δ-cost of the path P constructed in Lemma 2 for $\delta = 2+\epsilon$. We consider all edges on this path and divide them into classes depending on their lengths. By Lemma 3, there are at most $(8c+1)^2$ edges of length between c and $2c$, each one inducing δ-cost at most $(2c)^\delta$. More generally, we have at most $(8c+1)^2 \cdot 4^k$ edges of length between $c/2^k$ and $2c/2^k$ and the δ-cost of any such edge is at most $(2c/2^k)^\delta$. Summing up over all possible edges of P thus yields a total δ-cost of P of at most

$$\|P\|^\delta \quad \leq \quad \sum_{k=0}^{\infty} (8c+1)^2 \cdot 4^k \cdot \left(\frac{2c}{2^k}\right)^\delta \quad = \quad (8c+1)^2 \cdot \frac{(2c)^\delta}{1-2^{2-\delta}}$$

5.2 Weak Spanner Implies Power Spanner for $\boldsymbol{\delta = 2}$

The preceding section showed that, for fixed $\delta > 2$, any weak c-spanner is also a (C, δ)-power spanner. The present section yields the same for $\delta = 2$, a case which, however, turns out to be much more involved. Moreover, our bounds on C in terms of c become slightly worse. In fact, the deepest result of this work is the following:

Theorem 5. *Let $G = (V, E)$ be a weak c-spanner with $V \subseteq \mathbb{R}^2$. Then G is a $(C, 2)$-power spanner for $C := \mathcal{O}(c^8)$.*

Proof. First recall that between vertices $\mathbf{u}, \mathbf{v} \in V$ there is a path P in G from $\mathbf{u}$ to $\mathbf{v}$ which remains inside a square of length $\ell := 2c \cdot |\mathbf{u}-\mathbf{v}|$ and center $\mathbf{u}$. We denote such a square by $S_\mathbf{u}(\ell)$. By $\mathbf{s}$ we denote the starting point of the path and by $\mathbf{t}$ the end (target) point. We denote by $V(P)$ the vertex set of a path and by $E(P)$ the edge set of a path.

We give a constructive proof of the Lemma, i.e. given a path in G obeying the weak spanning property we construct a path which obeys the $\big(\mathcal{O}(c^8), 2\big)$-power spanner property. For this we iteratively apply a procedure called **clean-up** to a path, yielding paths with smaller and smaller costs. Besides the path P in G this procedure has parameters $L, d, D \in \mathbb{R}^+$. Hereby, L denotes the edge length of a square with middle point s containing the whole path. The parameters d, D are in the range $0 < 3(c\sqrt{3}+2)d \leq D \leq L$

```
Procedure clean-up (P, L, d, D)
  begin
    while three edges exist in P longer than 2√2cd
      starting or ending in the same cell of G
      do
        Let C be such a cell in G
        P ← contract(P, C)
      od
    while there exists a cell in G    where at least one vertex of P
      lies in each of its G  -sub-cells
      do
        Let C be such a cell of G
        Let rank   (u) denote the position of a vertex u in P
        Sort all cells Z1, . . . , Z(      )2 of G   in C
          according to minu∈   ∩  (  ){rank   (u)}
        Sort all cells Z'1, . . . , Z'(      )2 of G   in C
          according to maxu∈  '∩  (  ){rank   (u)}
        i ← 1
        while cell Z  is not horizontally neighbored
          to one of the cells {Z'1, . . . , Z'}
          nor Z' is horizontally neighbored
          to one of the cells {Z1, . . . , Z }
          do
            i ← i + 1
          od
        Let z and z' be the two neighbored cells
          from {Z1, . . . , Z } and {Z'1, . . . , Z'}
        P ← contract(P, z ∪ z')
      od
    return P
  end
```

```
Procedure contract (P = (v1, . . . , v   ) : path, A : area)
  begin
    Let v  be the first vertex of P in A
    Let v  be the last vertex of P in A
    Let P' = (w1, . . . , w  ) be a path between v  = w1
      and v  = w  satisfying the weak spanner property
    return (v1, . . . , v −1, w1, . . . , w  , v +1, . . . , v  )
  end
```

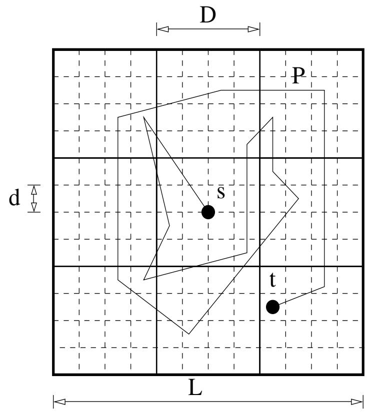

Fig. 4. The **clean-up** and the **contract** procedures and the idea for the Proof of Theorem 5

and can be chosen arbitrarily, yet fulfilling $D/d \in \mathbb{N}$ and $L/D \in \mathbb{N}$. These parameters define two edge-parallel grids G_d and G_D of grid size d and D such that boundaries of G_D are also edges of G_d. These grids fill out the square $S_{\mathbf{u}}(L)$, while the boundary edge of $S_{\mathbf{u}}(L)$ coincides with the boundary of G_d and G_D, see Fig. 4 The outcome of the procedure **clean-up** is a path $P' = \text{clean-up}(P, L, d, D)$ which reduces the cost of the path while obeying other constraints, as we show shortly.

In Figure 4 we describe the procedure **clean-up** which uses the procedure **contract** described in Figure 4. Let $D(A)$ denote the diameter of the area A.

Lemma 4. *The procedure* $P' = \textbf{contract}(P, A)$ *satisfies the following properties.*

- *Locality:* $\forall \mathbf{u} \in V\big(E(P) \setminus E(P')\big)$*:* $\min_{\mathbf{p} \in A} |\mathbf{u} - \mathbf{p}| \leq c \cdot D(A)$ *and* $\max_{\mathbf{p} \in A} |\mathbf{u} - \mathbf{p}| \leq (c+1) \cdot D(A)$.
- *Continuity of long edges:* $\forall e \in P' : |e| \geq 2c \cdot D(A) \Longrightarrow e \in P$.

Proof. The maximum distance between $\mathbf{v}_i$ and $\mathbf{v}_j$ is at most $D(A)$. The replacement path $(\mathbf{w}_1, \ldots, \mathbf{w}_k)$ is inside a disk of radius $c \cdot D(A)$. Hence for all vertices $\mathbf{u}$ of this replacement path we have $|\mathbf{u} - \mathbf{v}_i| \leq cD(A)$ and therefore $\min_{\mathbf{p} \in A} |\mathbf{u} - \mathbf{p}| \leq |\mathbf{u} - \mathbf{v}_i| \leq cD(A)$. From the triangle inequality it follows

$$\max_{\mathbf{p} \in A} |\mathbf{u} - \mathbf{p}| \leq D(A) + \min_{\mathbf{p} \in A} |\mathbf{u} - \mathbf{p}| \leq D(A) + |\mathbf{u} - \mathbf{v}_i| \leq (c+1)D(A) .$$

The second property follows from the fact that all new edges inserted in P' lie inside a disk of radius $D(A)$.

Lemma 5. *For $D \geq 3(c\sqrt{3}+2)d$ the procedure* **clean-up** *satisfies the four properties power efficiency, locality, empty space, and continuity of long edges.*

1. **Locality** *For all vertices $\mathbf{u} \in V(P')$ there exists $\mathbf{v} \in V(P)$ such that*
$$|\mathbf{u}-\mathbf{v}| \;\leq\; ((\sqrt{2}+\sqrt{3})c+2)\cdot d\,.$$
2. **Continuity of Long Edges** *For all edges e of P' with $|e| \geq 2\sqrt{3}cd$ it holds $e \in E(P)$.*
3. **Power Efficiency** *For all $k > 2\sqrt{3}c$:*
$$\sum_{e \in E(P'):2\sqrt{3}cd<|e|\leq kd} (|e|)^2 \;\leq\; k^2\, d^2\, \#F(P, G_d)\,,$$
where $\#F(P, G_d)$ denotes the number of grid cells of G_d where at least one vertex of P lies which is the end point of an edge of minimum length $2\sqrt{3}d$.
4. **Empty Space** *For all grid cells C of G_D we have at least one sub-cell of G_D within C without a vertex of P'.*

Proof. All cells of G_d are called sub-cells in this proof for distinguishing them from the cells of G_D.

Observe that the **clean-up** procedure uses only contract-operation to change the path. As parameters for this procedure we use either a grid sub-cell C of edge length d and diameter $D(C) = \sqrt{2}d$ or two horizontally neighbored grid sub-cells Z and Z' with edge lengths d with diameter $D(Z \cup Z')$.

Further note, that in the first loop each sub-cell C of the grid G_d will be treated by the contract-procedure once. The reason is that from the contract procedures edges with lengths less than $2\sqrt{2}d$ are produced, while each sub-cell will loose all but two edges of P with minimum length $2\sqrt{2}d$. This also proves that the first loop always halts.

Now consider the second while-loop and concentrate on the part inside the loop before the contract-operation takes place. Since in every sub-cell of C we have a vertex of P we can compute the ordering Z_i and Z'_i as described by the algorithm. The main observation is that until the first two neighbored sub-cells Z and Z' from these sets are found, no two sub-cells Z and Z' from $Z \in \{Z_j\}_{j\leq i}$ and $Z' \in \{Z'_j\}_{j\leq i}$ are horizontally neighbored. Hence, in the $2d$-surrounding of every point there is an empty sub-cell without any points of P.

The situation changes slightly if we apply the contract-operation. Then, an intermediate path will be added and possibly some of the empty sub-cells will start to contain vertices of the path. However, only sub-cells in a euclidean distance of $c\sqrt{3}d$ from sub-cells Z and Z' are affected by this operation. Now consider a square Q of $(c\sqrt{3}+2)d \times (c\sqrt{3}+2)d$ sub-cells in the middle of C. Then, at least two horizontally neighbored sub-cells will not be influenced by this contract-operation and thus remain empty.

One cannot completely neglect the influence of this operation to a neighbored grid cell of C. However, since $D \geq 3(c\sqrt{3}+2)d$ the inner square Q is not affected by contract-operation in neighbored grid cells of C because of the locality of the contract-operation.

This means if a cell C was object to the second while-loop, then an empty sub-cell will be produced which remains empty for the rest of the procedure. Hence, the second loop also terminates.

We now check the four required properties.

Locality. After the first loop the locality is satisfied even within a distance of $\sqrt{2}(c+1)d$. For this, observe that all treated cells contain end points of edges longer than $2\sqrt{2}cd$ which cannot be produced by contract-operations in this loop. Hence, if a cell is object to the contract-operation it was occupied by a vertex of P from the beginning. Then, from Lemma 4 it follows that for all new vertices of the path P there exists at least one old vertex in distance $\sqrt{2}(c+1)d$ after the first loop.

For the second loop we need to distinguish two cases. First, consider a cell C where in the inner square an empty sub-cell exists. In this case this cell will never be treated by this second loop. If new vertices are added to the path within this cell, then this will be caused by a contract-operation in a neighbored cell and will be considered in the second case.

Now consider all cells with preoccupied inner squares (preoccupation refers to the outcome of the first loop). These cells can be object to contract-operations of the second loop. However, they will add only vertices to their own sub-cells or to the outer sub-cells of neighbored cells. So, new vertices are added within a distance of $(\sqrt{3}c+1)d$ of vertices in the path at the beginning of the second loop. As we have seen above every such vertex is only $\sqrt{2}(c+1)d$ remote from an original vertex of a path. This gives a locality of distance $((\sqrt{3}+\sqrt{2})c+2)d$.

Continuity of Long Edges. Since the parametrized areas for the contract operation have a maximum diameter of $\sqrt{3}d$ this property follows directly from Lemma 4.

Power Efficiency. After the first loop the number of edges longer than $2\sqrt{2}cd$ is bounded by $\#F(P, G_d)$, because in every occupied sub-cell at most two edges start or end and each edge has two end points. Clearly, this number is an upper bound for edges longer than $2\sqrt{3}cd$. In the second loop no edges longer than $2\sqrt{3}cd$ will be added. This, directly implies the wanted bound.

Empty Space. As we have already pointed out the second loop always halts. Therefore the empty space property holds.

Lemma 6. *Given a path P_0 with source $\mathbf{s}$ and target $\mathbf{t}$ such that $\forall \mathbf{u} \in V(P_0) : |\mathbf{u} - \mathbf{s}| \leq c \cdot L$, where $L = c \cdot |\mathbf{u} - \mathbf{v}|$. Now iteratively apply $P_{i+1} = \text{clean-up}(P_i, L + \sum_{j=0}^{i} D_i, d_i, D_i)$ for $i = 1, 2, \ldots$ where $D_i = L\beta^{1-i}$, $d_i = L\beta^{-i}$ for $\beta = 20\sqrt{3}c$. Then, P_m for $m = \lceil \log_\beta \min_{\mathbf{u},\mathbf{v} \in V} |\mathbf{u} - \mathbf{v}| \rceil$ is path connecting $\mathbf{s}$ to $\mathbf{t}$ obeying the $\big(\mathcal{O}(c^8), 2\big)$ power spanner property.*

Proof. For this proof we make use of the four properties of the clean-up procedure. By assumption $c \geq \sqrt{3}$ we have $\beta \geq 60$.

First note that the the square of edge length L_i containing all vertices of path P_i can increase. However we can bound this effect by the locality property, giving $L_{i+1} \leq L_i + 2d_i$, where $d_i = L\cdot\beta^{-i}$. By assumption we have $c \geq \sqrt{3}$ and therefore $d_i \leq L4^{-i}$, which gives an upper bound of $L_i \leq 2L$ for all i.

Let $F_i = \#F(P_i, G_{d_i})$. Then $A_i = (d_i)^2 F_i$ denotes the area of all grid cells in G_{d_i} with a vertex of the path P_i which is the end point of an edge with length of at least $2\sqrt{3}d$. In the next iteration in each of this cells an empty space will be generated with an area of $(d_{i+1})^2$. Because of the locality property at least the following term of the edge length is subtracted

$$\sum_{j=1}^{\infty} 2\sqrt{3}d_{i+j} \leq \frac{1}{2}d_i \ .$$

Hence, an empty area of at least $\frac{1}{4}(d_i)^2$ remains after applying all clean-up-procedures. Let E_i be the sum of all these areas in this iteration. Therefore we have $A_i \leq 4\beta^2 E_i$. Clearly, these empty areas in this iteration do not intersect with empty areas in other areas (since they arise in areas which were not emptied before). Therefore all these spaces are inside the all-covering square of side length $2L$ yielding $\sum_{i=1}^{\infty} E_i \leq 4L^2$.

Because of the long edge continuity property, edges of minimum length $2\sqrt{3}d_i$ do not appear in rounds later than i. Therefore, the following sum S gives an upper bound on the power of the constructed path.

$$S = \sum_{i=1}^{\infty} \sum_{\substack{e\in E(P_i):\\ 2\sqrt{3}cd_i \leq |e| < 2\sqrt{3}c\beta d_i}} (|e|_2)^2 \ .$$

From the power efficiency property it now follows

$$\begin{aligned} S &\leq \sum_{i=1}^{\infty} 12c^2\beta^2(d_i)^2 \#F(P_i, G_{d_i}) = 12c^2\beta^2 \sum_{i=1}^{\infty}(d_i)^2 F_i = 12c^2\beta^2 \sum_{i=1}^{\infty} A_i \\ &\leq 48c^2\beta^4 \sum_{i=1}^{\infty} E_i \leq 192\, c^2\beta^4 L^2 \leq 192\, c^4\beta^4(|\mathbf{s}-\mathbf{t}|)^2 = O(c^8(|\mathbf{s}-\mathbf{t}|)^2) \end{aligned}$$

This lemma completes the proof of the theorem.

5.3 Weak Spanner Does Not Imply Power Spanner for $\delta < 2$

Theorem 6. *To any $\delta < 2$, there exists a family of geometric graphs $G = (V, E)$ with $V \subseteq \mathbb{R}^2$ which are weak c-spanners for a constant c but no (C, δ)-power spanners for any fixed C.*

Proof. As $\delta < 2$, there is a $k \in \mathbb{R}$ such that $2 < k < 4^{1/\delta}$. We present a recursive construction (see Figure 5). Fix $\mathbf{u}^1 = (1/2, 1/2) \in \mathbb{R}^2$. In each following recursion step j, we replace every existing vertex $\mathbf{u}^i = (u_x^i, u_y^i)$ by four new vertices $\mathbf{u}^{4i-3} = (u_x^i - d, u_y^i + d)$, $\mathbf{u}^{4i-2} = (u_x^i + d, u_y^i + d)$, $\mathbf{u}^{4i-1} = (u_x^i + d, u_y^i - d)$, and $\mathbf{u}^{4i} = (u_x^i - d, u_y^i - d)$ where $d := 1/(2k^j)$. Finally, we consider the graph $G_j := (V_j, E_j)$

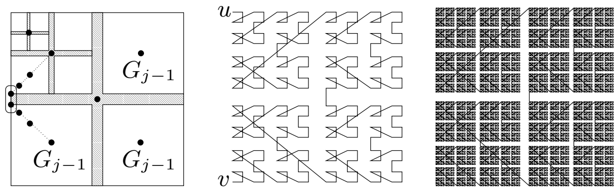

(a) Idea (b) After 4 steps (c) After 7 steps

Fig. 5. Recursive construction: The underlying idea and two examples for $k = 2.1$

with $V_j := \{\mathbf{u}^i \mid i \in \{1, \ldots, 4^j\}\}$ and $E_j := \{(\mathbf{u}^i, \mathbf{u}^{i+1}) \mid i \in \{1, \ldots, 4^j - 1\}\}$. The resulting graph after 4 recursion steps with $k = 2.1$ is given in Figure 5(b). Let $\mathbf{u} = \mathbf{u}^1$ and $\mathbf{v} = \mathbf{u}^{4^j}$.

Lemma 7. *The graph G_j is a weak c-spanner for $c := \frac{\sqrt{2}k(k-1)}{k-2}$ independent of j.*

Proof. We prove the claim by induction over j. For $j = 1$ the weak stretch factor is dominated by the path between $\mathbf{u}$ and $\mathbf{v}$. The distance between $\mathbf{u}$ and $\mathbf{v}$ is $1/k$. The farthest vertex on the path from $\mathbf{u}$ to $\mathbf{v}$ is $\mathbf{u}_3$. It holds that $|\mathbf{u} - \mathbf{u}_3| \le \sqrt{2}/k$. Hence, we get the weak stretch factor $\sqrt{2} \le \frac{\sqrt{2}k(k-1)}{k-2} = c$. Now, we consider G_j for any j. We can divide the graph G_j into four parts $G_j^1, \ldots, G_j^4$. By the definition of our recursive construction each part equals the graph G_{j-1}. For two vertices in one part the required weak c-spanner property holds by induction. We have to concentrate on two vertices which are chosen from two different parts. Since G_j^i is connected to G_j^{i+1} it is sufficient to consider a vertex from G_j^1 and a vertex from G_j^4. On the one hand, the weak stretch factor is affected by the shortest distance between such chosen vertices. On the other hand, this distance is given by (see also Figure 5(a))

$$\left(\frac{1}{2} \cdot \left(1 + \frac{1}{k} - \sum_{i=2}^{j} (\tfrac{1}{k})^i\right) - \frac{1}{2}\right) \cdot 2 \quad \ge \quad \frac{k-2}{k(k-1)}$$

The entire construction lies in a bounded square of side length 1, and hence we get a weak stretch factor of at most $\frac{\sqrt{2}k(k-1)}{k-2} = c$.

Lemma 8. *The graphs G_j are no (C, δ)-power spanners for any fixed C.*

Proof. It suffices to consider the δ-cost of the path from $\mathbf{u}$ to $\mathbf{v}$. The direct link from $\mathbf{u}$ to $\mathbf{v}$ has δ-cost at most 1. For any path P from $\mathbf{u}$ to $\mathbf{v}$ in G, it holds that

$$\|P\|^\delta \quad \ge \quad 3 \cdot 4^j \cdot \left((\tfrac{1}{k})^j\right)^\delta \quad = \quad 3 \cdot \left(\frac{4}{k^\delta}\right)^j$$

which goes to infinity if $j \to \infty$ for $k < 4^{1/\delta}$.

Combining Lemma 7 and Lemma 8 proves Theorem 6.

5.4 Fractal Dimension

The present section generalizes the construction and analysis used in Lemma 8. To this end, consider a self-similar polygonal fractal curve Γ as the result of repeated application of some generator K being a polygonal chain with starting point $\mathbf{u}$ and end point $\mathbf{v}$. This is illustrated in Fig. 2 showing a generator (left) and the resulting fractal curve (right part); see also [14, 2]. But other examples are plenty: the KOCH Snowflake or the space filling HILBERT Curve (Fig. 6). Recall that the fractal dimension of Γ is defined as

$$\frac{\log(\text{number of self-similar pieces})}{\log(\text{magnification factor})}$$

Theorem 7. *Let K be a polygonal chain, Γ_n the result of n-fold application of K, and Γ the final self-similar polygonal fractal curve with dimension d. Then, for all $\delta < d$, there is no fixed C such that Γ_n is a (C, δ)-power spanner for all n.*

Proof. Let p denote the number of self-similar pieces in Γ_n and m the magnification factor. Then by definition, we have $d = \log(p)/\log(m)$. Now consider the δ-cost of the (unique) path P in Γ_n from $\mathbf{u}$ to $\mathbf{v}$. Since Γ_n is constructed recursively we get in the n-th step:

$$\|P\|^\delta \quad = \quad p^n \cdot \left(\left(\tfrac{1}{m}\right)^n\right)^\delta \quad = \quad \left(\frac{p}{m^\delta}\right)^n$$

Note that $\|P\|^\delta$ is unbounded iff $p/m^\delta > 1$, that is, iff $\delta < \log(p)/\log(m) = d$.

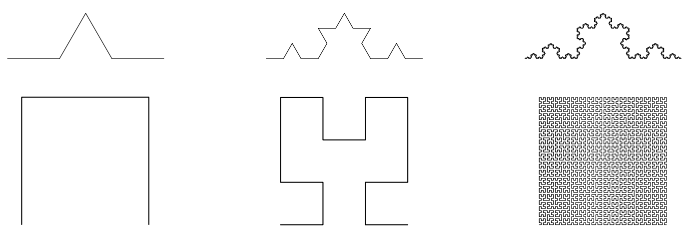

Fig. 6. Two Generators and the Fractal Curves they induce due to KOCH and HILBERT

The fractal dimensions of the KOCH and HILBERT Curves are well-known. Therefore by virtue of Theorem 7, the KOCH Curve is not a $(\cdot, \delta)$-power spanner for any $\delta < \log(4)/\log(3) \approx 1.26$; similarly, HILBERT's Curve is not a $(\cdot, \delta)$-power spanner for any $\delta < 2$. One can show that KOCH's Curve is a weak spanner (the proof is analogous to Theorem 1). However HILBERT's Curve is no weak spanner as its inner vertices come arbitrarily close to each other. Further examples for self-similar polygonal curves, e.g., SIERPINSKIS's triangle, can be found in [14, 2].

6 Power Spanner Hierarchy

In the following we show that for $\Delta > \delta > 0$, a (c, δ)-power spanner is also a (C, Δ)-power spanner with C depending only on c and Δ/δ. Then we show that the converse fails in general by presenting to each $\Delta > \delta > 0$ a family of graphs which are (c, Δ)-power spanners for some constant c but no (C, δ)-power spanners for any fixed C.

Theorem 8. *Let $G = (V, E)$ be a (c, δ)-power spanner with $V \subseteq \mathbb{R}^2$, $0 < \delta < \Delta$. Then G is also a (C, Δ)-power spanner for $C := c^{\Delta/\delta}$.*

Proof. Let $\mathbf{u}, \mathbf{v} \in V$ be two arbitrary vertices. Since G is a (c, δ)-power spanner there exists a path $P = (\mathbf{u} = \mathbf{u}_1, \ldots, \mathbf{u}_l = \mathbf{v})$ with $\|P\|^\delta = \sum_{i=1}^{l-1} |\mathbf{u}_i - \mathbf{u}_{i+1}|^\delta \leq c \cdot |\mathbf{u} - \mathbf{v}|_2^\delta$. The function $f(x) = x^{\Delta/\delta}$ being convex on $[0, \infty[$, one may apply JENSEN's Inequality:

$$\|P\|^\Delta = \sum_{i=1}^{l-1} |\mathbf{u}_i - \mathbf{u}_{i+1}|^\Delta = \sum_{i=1}^{l-1} \left(|\mathbf{u}_i - \mathbf{u}_{i+1}|^\delta\right)^{\Delta/\delta} \leq \left(\sum_{i=1}^{l-|!1} |\mathbf{u}_i - \mathbf{u}_{i+1}|^\delta\right)^{\Delta/\delta} \leq c^{\Delta/\delta} \cdot |\mathbf{u} - \mathbf{v}|^\Delta.$$

Theorem 9. *Let $0 < \delta < \Delta$. There is a family of geometric graphs which are (c, Δ)-power spanners but no (C, δ)-power spanners for any fixed C.*

Proof. We slightly modify the construction from the proof of Theorem 3 by placing n vertices $\mathbf{u} = \mathbf{u}_1, \ldots, \mathbf{u}_n = \mathbf{v}$ on an appropriately scaled circle such that the Euclidean distance between $\mathbf{u}$ and $\mathbf{v}$ is 1 and $|\mathbf{v}_i - \mathbf{v}_{i+1}| = (1/i)^{1/\delta}$ for all $i = 1, \ldots, n-1$. Now, in the graph $G = (V, E)$ with edges $(\mathbf{v}_i, \mathbf{v}_{i+1})$, the unique path P from $\mathbf{u}$ to $\mathbf{v}$ has Δ-cost

$$\|P\|^\Delta = \sum_{i=1}^{n-1} (1/i)^{\Delta/\delta} \leq \sum_{i=i}^{\infty} (1/i)^{\Delta/\delta} =: c$$

a convergent series since $\Delta/\delta > 1$. This is to be compared to the cost of the direct link from $\mathbf{u}$ to $\mathbf{v}$ which amounts to 1 both w.r.t. Δ and δ. On the other hand, the δ-cost of P is given by the harmonic series $\sum_{i=1}^{n-1} (1/i)^{\delta/\delta} = \Theta(\log n)$ and thus cannot be bounded by any C independent of n.

7 Higher-Dimensional Case

For simplicity, most results in this work have been formulated for the case of (not necessarily planar) geometric graphs in the plane. They immediately apply to higher dimensions as well, however with the exception of Section 5 (Weak versus Power Spanners). In fact, similar techniques yield that, for instance in 3D, each weak c-spanner is a (C, δ)-power spanner for $\delta \geq 3$ with C depending only on c and δ whereas to any $\delta < 3$, there are counter-examples of weak c-spanners that are not (C, δ)-power spanners for any fixed C; analogously in higher dimensions.

8 Conclusions

We investigate the relations between spanners, weak spanners, and power spanners. In the plane, for $\delta \geq 2$ it turns out that being a spanner is the strongest property, followed by being a weak spanner and finally being a $(\cdot, \delta)$-power spanner. For $1 < \delta < 2$, spanner is still strongest whereas weak spanner and $(\cdot, \delta)$-power spanner are not related to each other. For $0 < \delta < 1$ finally, $(\cdot, \delta)$-power spanner implies both spanner and weak spanner. For higher dimensions, similar relations/independencies hold. All stretch factors in these relations are constant and are pairwise polynomially bounded.

In [9, 7] a geometric graph called YY or SparsY-Graph was investigated as a topology for wireless networks. It is constructed by dividing the area around each vertex into $k \in \mathbb{N}$ non-overlapping sectors or cones of angle $\theta = 2\pi/k$ each. In each sector of a vertex, there is at most one outgoing edge and if there is one, then this goes to the nearest neighbor in this sector. A vertex accepts in each of its sectors only one ingoing edge and this must be the shortest one in this sector. For this graph the relation between weak and power spanner is exemplarily investigated in [9, 15] by performing experiments on uniformly and randomly distributed vertex sets. They conclude that the SparsY-Graph might be a spanner and also a power spanner. The first conjecture is still open, while the latter was independently proven in [8] and [7].

Observe that SparsY is well-known to yield a good weak spanner already for $k > 6$ [7]. Regarding that our Theorem 5 asserts *any* weak spanner to be also a $(\cdot, \delta)$-power spanner for $\delta \geq 2$, this includes the above result *and* weakens the presumption from $k \geq 120$ [8] to $k > 6$.

Although our results are exhaustive with respect to the different kinds of geometric graphs and in terms of δ, one might wonder about the optimality of the bounds obtained for C's dependence on c; for instance: Any c-spanner is a (C, δ)-power spanner for $C = c^\delta$, $\delta > 1$; and this bound *is* optimal. But is there some $C = o(c^4)$ such that any weak c-spanner is a (C, δ)-power spanner as long as $\delta > 2$? Is there some $C = o(c^8)$ such that any weak c-spanner is a $(C, 2)$-power spanner ?

References

1. K. Alzoubi, X.-Y. Li, Y. Wang, P.J. Wan, and O. Frieder. Geometric spanners for wireless ad hoc networks. *IEEE Transactions on Parallel and Distributed Systems*, 14(4):408–421, 2003.
2. D. Eppstein. The Geometry Junkyard: Fractals, www.ics.uci.edu/∼eppstein/junkyard/fractal.html.
3. D. Eppstein. Beta-skeletons have unbounded dilation. Technical Report ICS-TR-96-15, 1996.
4. D. Eppstein. Spanning trees and spanners. In *Handbook of Computational Geometry*, pages 425–461. 2000.
5. M. Fischer, T. Lukovszki, and M. Ziegler. Geometric searching in walkthrough animations with weak spanners in real time. In *6th Annual European Symposium on Algorithms (ESA '98)*, pages 163–174, 1998.
6. M. Fischer, F. Meyer auf der Heide, and W.-B. Strothmann. Dynamic data structures for realtime management of large geometric scenes. In *5th Annual European Symposium on Algorithms (ESA '97)*, pages 157–170, 1997.

7. M. Grünewald, T. Lukovszki, C. Schindelhauer, and K. Volbert. Distributed Maintenance of Resource Efficient Wireless Network Topologies (Ext. Abstract). In *8th EURO-PAR'02*, pages 935–946, 2002.
8. L. Jia, R. Rajaraman, and C. Scheideler. On local algorithms for topology control and routing in ad hoc networks. In *Proc. 15th ACM Symposium on Parallel Algorithms and Architectures (SPAA '03)*, pages 220–229, 2003.
9. 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 (ICCCN01)*, pages 564–567, 2001.
10. F. Meyer auf der Heide, C. Schindelhauer, K. Volbert, and M. Grünewald. Energy, congestion and dilation in radio networks. In *Proc. 14th Symposium on Parallel Algorithms and Architectures (SPAA'02)*, pages 230–237, 2002.
11. F. Meyer auf der Heide, C. Schindelhauer, K. Volbert, and M. Grünewald. Congestion, Dilation, and Energy in Radio Networks. *Theory of Computing Systems*, 37(3):343–370, 2004.
12. R. Rajaraman. Topology control and routing in ad hoc networks: a survey. *SIGACT News*, 33(2):60–73, 2002.
13. S. B. Rao and W. D. Smith. Approximating geometrical graphs via spanners and banyans. In *Proceedings of the 30th annual ACM symposium on Theory of computing*, pages 540–550, 1998.
14. C. Tricot. *Curves and Fractal Dimension*. Springer, 1995.
15. K. Volbert. Experimental Analysis of Adjustable Sectorized Topologies for Static Ad Hoc Networks, accepted for DIALM-POMC 2004.
16. Y. Wang and X.-Y. Li. Distributed Spanner with Bounded Degree for Wireless Ad Hoc Networks. In *Parallel and Distributed Computing Issues in Wireless networks and Mobile Computing*, page 120, 2002.
17. Y. Wang, X.-Y. Li, P.-J. Wan, and O. Frieder. Sparse power efficient topology for wireless networks. In *Proc. ACM Hawaii International Conference on System Sciences (HICSS'02)*, page 296, 2002.
18. A. C.-C. Yao. On Constructing Minimum Spanning Trees in k-dimensional space and related problems. *SIAM J. Comput.*, 11:721–736, 1982.