Finding the Best Shortcut in a Geometric Network

Mohammad Farshi[*]
Dept. of Computer Science
TU Eindhoven
The Netherlands
m.farshi@tue.nl

Panos Giannopoulos[†]
Institut für Informatik
Freie Universität Berlin
Germany
panos@inf.fu-berlin.de

Joachim Gudmundsson
NICTA[‡]
Sydney, Australia
joachim.gudmundsson
@nicta.com.au

ABSTRACT

Given a Euclidean graph G in $\mathbb{R}^d$ with n vertices and m edges we consider the problem of adding a shortcut such that the stretch factor of the resulting graph is minimized. Currently, the fastest algorithm for computing the stretch factor of a Euclidean graph runs in $\mathcal{O}(mn + n^2 \log n)$ time, resulting in a trivial $\mathcal{O}(mn^3 + n^4 \log n)$ time algorithm for computing the optimal shortcut. First, we show that a simple modification yields the optimal solution in $\mathcal{O}(n^4)$ time using $\mathcal{O}(n^2)$ space. To reduce the running times we consider several approximation algorithms. Our main result is a $(2+\varepsilon)$-approximation algorithm with running time $\mathcal{O}(nm + n^2(\log n + 1/\varepsilon^{3d}))$ using $\mathcal{O}(n^2)$ space.

Categories and Subject Descriptors

F.2 [**Theory of Computation**]: Analysis of algorithms and problem complexity; G.2.2 [**Mathematics of Computing**]: Graph theory

General Terms

Algorithms, Theory

Keywords

Computational Geometry, Approximation algorithms, Geometric networks

1. INTRODUCTION

Consider a set V of n points in $\mathbb{R}^d$. A network on V can be modeled as an undirected graph G with vertex set V of size

n and an edge set E of size m where every edge $e = (u,v)$ has a weight $wt(e)$. A Euclidean network is a geometric network where the weight of the edge $e = (u,v)$ is equal to the Euclidean distance $d(u,v)$ between its two endpoints u and v. Let $t > 1$ be a real number. We say that G is a *t-spanner* for V, if for each pair of points $u,v \in V$, there exists a path in G of weight at most t times the Euclidean distance between u and v. The minimum t such that G is a t-spanner for V is called the stretch factor, or dilation, of G.

Complete graphs represent ideal communication networks, but they are expensive to build; sparse spanners represent low-cost alternatives. The weight of the spanner network is a measure of its sparseness; other sparseness measures include the number of edges, the maximum degree, and the number of Steiner points. Spanners for complete Euclidean graphs as well as for arbitrary weighted graphs find applications in robotics, network topology design, distributed systems, design of parallel machines, and many other areas and have been a subject of considerable research. Recently spanners found interesting practical applications in areas such as metric space searching [26, 27] and broadcasting in communication networks [2, 14, 24].

Several well-known theoretical results also use the construction of t-spanners as a building block, for example, Rao and Smith [28] made a breakthrough by showing an optimal $\mathcal{O}(n \log n)$-time approximation scheme for the well-known Euclidean *traveling salesperson problem*, using t-spanners (or banyans). Similarly, Czumaj and Lingas [7] showed approximation schemes for minimum-cost multi-connectivity problems in geometric graphs. The problem of constructing spanners has received considerable attention from a theoretical perspective, see [1, 3, 4, 5, 8, 9, 10, 16, 19, 20, 22, 23, 29, 31], and the surveys [12, 30].

All the above algorithms construct a network from scratch but in many applications the geometric network is already given, and the problem at hand is to extend the network with an additional edge, or edges, while minimizing the stretch factor of the resulting graph. Surprisingly this problem has not been studied previously, to the best of the authors' knowledge. In this paper we study the following problem:

Problem. Given a graph G construct a graph G' by adding an edge to G such that the stretch factor of G' is minimized.

We present one exact algorithm and several approximation algorithms. The results presented in this paper are summarized in Table 1.

We will denote by $|uv|$ the Euclidean distance between

Input graph	Apx. factor	Time complexity	Space	Section
Metric	1	$\mathcal{O}(n^3 m + n^4 \log n)$	$\mathcal{O}(n)$	2
Metric	1	$\mathcal{O}(n^4)$	$\mathcal{O}(n^2)$	2
Metric	3	$\mathcal{O}(nm + n^2 \log n)$	$\mathcal{O}(n)$	3
Euclidean	$2 + \varepsilon$	$\mathcal{O}(nm + n^2(\log n + 1/\varepsilon^{3d}))$	$\mathcal{O}(n^2)$	4
t-spanner	$1 + \varepsilon$	$\mathcal{O}((t^7/\varepsilon^4)^d \cdot n^2)$	$\mathcal{O}((t^3/\varepsilon^2)^d \, n \log(tn))$	5

Table 1: Complexity bounds for the algorithms presented in the paper.

u and v, and $\delta_G(u, v)$ denotes the shortest path between u and v in G with length $d_G(u, v)$. Finally, $G_\mathcal{P}$ will denote the optimal solution, while $t_\mathcal{P}$ and t denotes the stretch factor of $G_\mathcal{P}$ and G respectively.

2. FINDING AN OPTIMAL SOLUTION

We consider the problem of computing an optimal solution $G_\mathcal{P}$. That is, we are given a t-spanner $G = (V, E)$, and the aim is to compute a $t_\mathcal{P}$-spanner $G_\mathcal{P} = (V, E \cup \{e\})$.

A naïve approach to decide which edge to add is to test every possible candidate edge. The number of such edges is obviously $\left(\frac{n(n-1)}{2} - m\right) = \mathcal{O}(n^2)$. Testing a candidate edge e entails computing the stretch factor of the graph $G' = (V, E \cup \{e\})$, therefore we briefly consider the problem of computing the stretch factor of a given Euclidean graph. This problem has recently received considerable attention, see for example [11, 13, 21, 25].

A trivial upper bound is obtained by computing the All-Pairs-Shortest-Path for the given graph G. Running Dijkstra's algorithm – implemented using Fibonacci heaps – gives the stretch factor of G in time $\mathcal{O}(mn + n^2 \log n)$ using linear space. This algorithm is quite slow and we would like to be able to compute the stretch factor more efficiently, but no faster algorithm is known for any graphs except planar graphs, paths, cycles, stars and trees [13, 21, 25].

Applying the above bounds for computing the exact stretch factor of a Euclidean graph gives us that $G_\mathcal{P}$ can be computed in time $\mathcal{O}(n^3(m + n \log n)))$ using linear space.

An improvement can be obtained by observing that when an edge (u, v) is about to be tested we do not have to check all possible shortest paths between two vertices $x, y \in V$ again, it suffices to check if there is a shorter path using the edge (u, v). That is, we only have to check the length of the paths $\delta_G(x, u) + |uv| + \delta_G(v, y)$, $\delta_G(x, v) + |vu| + \delta_G(u, y)$ and $\delta_G(x, y)$, which can be done in constant time since shortest path distances between every pair of vertices in G already have been computed (provided that we store this information). Hence by first computing all-pair-shortest paths of G we obtain:

LEMMA 1. *Given a Euclidean graph G, an optimal solution $G_\mathcal{P}$ can be computed in time $\mathcal{O}(n^4)$ using $\mathcal{O}(n^2)$ space.*

PROOF. Computing the all-pair-shortest path requires cubic time and $\mathcal{O}(n^2)$ space. The $\mathcal{O}(n^2)$ edges are tested for insertion, for each candidate edge one performs $\mathcal{O}(n^2)$ shortest path queries, which each can be answered in constant time as described above. □

The above lemma can be generalized to hold for any graph for which the weight of the edges obey the triangle inequality.

3. ADDING THE BOTTLENECK EDGE

In this section we study the approach of adding an edge between a pair of vertices in G that decides the stretch factor of G.

Consider an optimal solution $G_\mathcal{P}$ and denote by x and y the two endpoints of an edge added to G to obtain $G_\mathcal{P}$. Assume that a pair of vertices deciding the stretch factor of G is (u, v), i.e., the length of the path between u and v in G is exactly $t \cdot |uv|$. We call this edge a bottleneck edge of G. Let $G_\mathcal{B}$ be a graph obtained from G by adding one bottleneck edge, and let $t_\mathcal{B}$ be the stretch factor of $G_\mathcal{B}$. Note that $G_\mathcal{B}$ can be computed in the same time as the stretch factor of G, i.e., in $\mathcal{O}(mn + n^2 \log n)$ time for Euclidean graphs.

LEMMA 2. *Given a Euclidean graph G in $\mathbb{R}^d$ it holds that $t_\mathcal{B} < 3t_\mathcal{P}$.*

PROOF. Recall that t denotes the stretch factor of G. First note that if $t_\mathcal{P} > t/3$ then the lemma holds and we are done. Thus we may assume that $t_\mathcal{P} \leq t/3$. The proof of the lemma is done by considering a pair of vertices, denoted (a, b), that decides the stretch factor of $G_\mathcal{B}$. Note that the path $\delta_{G_\mathcal{P}}(a, b)$ must include the added edge (x, y), otherwise $d_{G_\mathcal{P}}(a, b) = d_G(a, b) \geq d_{G_\mathcal{B}}(a, b)$ and we are done. Also, we will assume without loss of generality that a shortest path in $G_\mathcal{P}$ from a to b goes from a to x and then to b via y. The same is assumed about a shortest path in $G_\mathcal{P}$ between u and v, i.e., from u to x and then to v via y.

Our first step is to bound the distance between the bottleneck vertices u and v. This is done by bounding the length of the path in G between x and y as follows, see Fig. 1.

$$
\begin{aligned}
d_G(u, v) & \leq d_{G_\mathcal{P}}(u, v) - |xy| + d_G(x, y) \\
& \leq t_\mathcal{P} \cdot |uv| - |xy| + t \cdot |xy| \\
& < \frac{t}{3} \, |uv| - |xy| + t|xy| \\
& < \frac{t}{3} \, |uv| + t|xy|.
\end{aligned}
$$

Since $d_G(u, v) = t \cdot |uv|$ it follows that

$$|uv| < 3/2 \cdot |xy|. \tag{1}$$

Figure 1: (x, y) is the optimal edge added to G and (u, v) is a bottleneck edge.

Also,

$$\begin{aligned}
t \cdot |uv| &= d_G(u, v) \\
&\leq d_G(u, a) + d_G(a, b) + d_G(b, v) \\
&\leq d_G(u, a) + t|ab| + d_G(b, v)
\end{aligned}$$

which implies that

$$t \cdot (|uv| - |ab|) \leq d_G(u, a) + d_G(b, v), \qquad (2)$$

and

$$\begin{aligned}
d_G(a, u) + 2|xy| + d_G(v, b) &\leq \\
\leq d_G(a, x) + d_G(x, u) + 2|xy| &+ d_G(v, y) + d_G(y, b) \\
= d_{G_\mathcal{P}}(a, b) + d_{G_\mathcal{P}}(u, v) & \\
\leq t_\mathcal{P}(|ab| + |uv|), & \qquad (3)
\end{aligned}$$

which gives that

$$d_G(a, u) + d_G(v, b) \leq t_\mathcal{P}(|ab| + |uv|) - 2|xy|. \qquad (4)$$

By putting together (2) and (4) we have

$$\begin{aligned}
t(|uv| - |ab|) &\leq d_G(a, u) + d_G(v, b) \\
&\leq t_\mathcal{P}(|ab| + |uv|) - 2|xy| \\
&< t_\mathcal{P}(|ab| + |uv|),
\end{aligned}$$

which implies that

$$|ab|(t_\mathcal{P} + t) > |uv|(t - t_\mathcal{P})$$

and

$$|ab| > \frac{t - t_\mathcal{P}}{t_\mathcal{P} + t}|uv| > \frac{t - \frac{t}{3}}{\frac{t}{3} + t}|uv| = \frac{1}{2}|uv|. \qquad (5)$$

Now we are ready to put together the results:

$$\begin{aligned}
t_\mathcal{B} \cdot |ab| &= d_{G_\mathcal{B}}(a, b) \\
&\leq d_G(a, u) + |uv| + d_G(v, b) \\
(1) \quad &< d_G(a, u) + \frac{3}{2}|xy| + d_G(v, b) \\
&< d_G(a, u) + 2|xy| + d_G(v, b) \\
(3) \quad &\leq t_\mathcal{P}(|ab| + |uv|) \\
(5) \quad &< 3t_\mathcal{P}|ab|
\end{aligned}$$

This completes the lemma since $t_\mathcal{B} < 3t_\mathcal{P}$. □

The above lemma can be generalized to hold for any graph for which the edge weights obey the triangle inequality. We conclude by stating the main result of this section.

THEOREM 1. *Given a Euclidean graph $G = (V, E)$ one can in $\mathcal{O}(mn + n^2 \log n)$ time, using $\mathcal{O}(n)$ space, compute a $t_\mathcal{B}$-spanner $G' = (V, E \cup \{e\})$ where $t_\mathcal{B} < 3t_\mathcal{P}$.*

We end this section by giving a lower bound for the bottleneck approach.

OBSERVATION 1. *There exists a graph G such that $(2-\varepsilon) \cdot t_\mathcal{P} \leq t_\mathcal{B}$, for any $0 < \varepsilon < 1$.*

PROOF. Consider the graph G, as in Fig. 2a. More specifically, G is a graph with ten vertices $p_i = ((i-1) \mod 5, \lfloor i/5 \rfloor \cdot \delta)$, $1 \leq i \leq 10$, and nine edges (p_5, p_{10}) and (p_j, p_{j+1}), for $1 \leq j \leq 4$ and $6 \leq j \leq 9$. For any value δ: (p_1, p_6) is the bottleneck in G and $t_\mathcal{B} = \frac{4+\delta}{\delta}$, see Fig. 2b.

In the case when edge (p_2, p_7) is added to G, as shown in Fig. 2c, the resulting graph has stretch factor $(2 + \delta)/\delta$. Combining the upper and lower bounds gives $\frac{t_\mathcal{B}}{t_\mathcal{P}} = \frac{4+\delta}{2+\delta} = (2-\varepsilon)$, where the last equality follows if we set $\delta = \frac{2\varepsilon}{1-\varepsilon}$. □

Recently, Grüne [15] improved the lower bound in Observation 1 to $(3 - \varepsilon)$, so the upper bound stated in Lemma 2 is tight.

4. A $(2 + \varepsilon)$-APPROXIMATION

In this section we will present a fast approximation algorithm which guarantees an approximation factor of $(2 + \varepsilon)$. The algorithm is similar to the algorithm presented in Section 2 in the sense that it tests candidate edges. Testing a candidate edge entails computing the stretch factor of the graph. The main difference is that we will show, in Section 4.2, that only a linear number of candidate edges needs to be tested to obtain a solution that gives a $(2 + \varepsilon)$-approximation, instead of a quadratic number of edges.

Moreover, Section 4.3 shows that the same approximation bound can be achieved by performing only a linear number of shortest path queries for each candidate edge. The candidate edges are selected by using the well-separated pair decomposition, which we briefly define below.

4.1 Well-separated pair decomposition

Our algorithm uses the well-separated pair decomposition defined by Callahan and Kosaraju [6]. We briefly review this decomposition before we state the algorithms.

DEFINITION 1 ([6]). *Let $s > 0$ be a real number, and let A and B be two finite sets of points in $\mathbb{R}^d$. We say that A and B are* well-separated *with respect to s, if there are two*

Figure 2: (a) The input graph G. (b) The bottleneck solution compared to (c) the optimal solution.

disjoint d-dimensional balls C_A and C_B, having the same radius, such that (i) C_A contains the bounding box $R(A)$ of A, (ii) C_B contains the bounding box $R(B)$ of B, and (iii) the minimum distance between C_A and C_B is at least s times the radius of C_A.

The parameter s will be referred to as the *separation constant*. The next lemma follows easily from Definition 1.

LEMMA 3 ([6]). *Let A and B be two finite sets of points that are well-separated w.r.t. s, let x and p be points of A, and let y and q be points of B. Then (i) $|xy| \leq (1+2/s) \cdot |xq|$, (ii) $|xy| \leq (1+4/s) \cdot |pq|$, and (iii) $|px| \leq (2/s) \cdot |pq|$.*

DEFINITION 2 ([6]). *Let S be a set of n points in $\mathbb{R}^d$, and let $s > 0$ be a real number. A well-separated pair decomposition (WSPD) for S with respect to s is a sequence of pairs of non-empty subsets of S, $\{A_1, B_1\}, \ldots, \{A_m, B_m\}$, such that*

1. *$A_i \cap B_i = \emptyset$, for all $i = 1, \ldots, m$,*

2. *for any two distinct points p and q of S, there is exactly one pair $\{A_i, B_i\}$ in the sequence, such that (i) $p \in A_i$ and $q \in B_i$, or (ii) $q \in A_i$ and $p \in B_i$,*

3. *A_i and B_i are well-separated w.r.t. s, for $1 \leq i \leq m$.*

The integer m is called the size *of the WSPD.*

Callahan and Kosaraju showed that a WSPD of size $m = \mathcal{O}(s^d n)$ can be computed in $\mathcal{O}(s^d n + n \log n)$ time.

4.2 Linear number of candidate edges

In this section we show how to obtain a $(2+\varepsilon)$-approximation in cubic time. As mentioned above, the algorithm is similar to the algorithm presented in Section 2 in the sense that it tests candidate edges. It will below be shown that only a linear number of candidate edges is needed to be tested to obtain a solution that gives a $(2 + \varepsilon)$-approximation.

The approach is straight-forward. First the algorithm computes the length of the shortest path in G between every pair of points in V. The distances are saved in a matrix M. Next, the well-separated pair decomposition is computed. Note that, in Step 5, the candidate edges will be chosen using the well-separated pair decomposition. Finally, steps 4–9, each candidate edge is tested by computing the stretch factor of the candidate graph.

Algorithm EXPANDGRAPH(G, ε)
Input: Euclidean graph $G = (V, E)$ and a real constant $\varepsilon > 0$.
Output: Euclidean graph $G' = (V, E \cup \{e\})$.
1. $M \leftarrow$ All-Pairs-Shortest-Path distance matrix of G.
2. $\{(A_i, B_i)\}_{i=1}^k \leftarrow$ WSPD of the set V with respect to separation constant $s = \frac{256}{\varepsilon^2}$.
3. $t' \leftarrow \infty$.
4. **for** $i \leftarrow 1$ to k
5.   Select arbitrary points $a_i \in A_i$ and $b_i \in B_i$.
6. $G_i \leftarrow G = (V, E \cup (a_i, b_i))$.
7. $t_i \leftarrow$ STRETCHFACTOR(G_i, M).
8.   **if** $t_i < t'$
9.   **then** $t' \leftarrow t_i$ and $e \leftarrow (a_i, b_i)$
10. **return** $G' = (V, E \cup \{e\})$.

Next we bound the running time of the approximation algorithm and then prove the approximation bound.

LEMMA 4. *Algorithm* EXPANDGRAPH *requires $\mathcal{O}(n^3/\varepsilon^{2d})$ time and $\mathcal{O}(n^2)$ space.*

PROOF. The complexity of all steps of the algorithm, except step 7, is straight-forward to calculate. Recall that step 1 requires $\mathcal{O}(mn + n^2 \log n)$ time and quadratic space, and step 2 requires $\mathcal{O}(n/\varepsilon^{2d} + n \log n)$ according to Section 4.1. Thus, it remains to consider step 7 of the algorithm. Note that the number of times step 7 is executed is $\mathcal{O}(n/\varepsilon^{2d})$.

Since we computed the all-pair shortest distances of G, and stored the results in a matrix M it holds that shortest-distance queries in G_i can be computed in constant time. That is, for a query (p, q) return $\min\{M[p, q], M[p, a_i] + |a_i b_i| + M[b_i, q], M[p, b_i] + |b_i a_i| + M[a_i, q]\}$. For each candidate edge, $\mathcal{O}(n^2)$ queries are performed, thus summing up we get $\mathcal{O}(\frac{n}{\varepsilon^{2d}} \cdot n^2)$, as stated in the lemma. $\square$

It remains to analyze the quality of the solution obtained from algorithm EXPANDGRAPH. We need to compare the graph resulting from adding an optimal edge to G and the graph G' resulting from EXPANDGRAPH. Let $e = (a, b)$ be an optimal edge and let (A_i, B_i) be the well separated pair such that $a \in A_i$ and $b \in B_i$. At first sight, it seems that the edge (a_i, b_i) tested by the algorithm should be a good candidate. However the separation constant of our well separated pair decomposition only depends on ε which implies that the shortest path between a and a_i, and between b and b_i could be very long compared to the distance between a and b. Therefore we will, in Lemma 5, prove the existence of a "short" edge e' that is a good approximation of the optimal edge and then, in Lemma 6, we show that Ex-PANDGRAPH computes a good approximation of e'. Note

Figure 3: (a) Illustrating the proof of Lemma 5. (b) Illustrating the proof of Lemma 6.

also that in Section 5 it will be shown that if the separation constant also depends on the stretch factor of G then a $(1+\varepsilon)$-approximation can be proven.

Let $\Delta(p,q)$ denote the set of point pairs in V such that the point pair (u,v) belongs to $\Delta(p,q)$ if and only if $(p,q) \in \delta_{G\cup\{(p,q)\}}(u,v)$. That is, the set of point pairs for which the shortest path between them in $G \cup \{(p,q)\}$ passes through (p,q).

LEMMA 5. *For any given constant $0 < \lambda \leq 1$, there exists a point pair $p,q \in V$ such that for every pair $(u,v) \in \Delta(p,q)$ it holds that $|uv| \geq \frac{\lambda}{2}|pq|$, and the stretch factor of $G \cup \{(p,q)\}$ is bounded by $(2+\lambda) \cdot t_{\mathcal{P}}$.*

PROOF. The proof is done in two steps. First a point pair $p_j, q_j \in V$ is selected that fulfills the first requirement of the lemma. Then, we prove the second requirement, i.e., the stretch factor of $G \cup \{(p_j, q_j)\}$ is bounded by $(2+\lambda) \cdot t_{\mathcal{P}}$.

Consider an optimal solution $G_1 = G \cup \{(p_1, q_1)\}$, with stretch factor $t_1 = t_{\mathcal{P}}$. If for each point pair $(u,v) \in \Delta(p_1, q_1)$ it holds that $|uv| \geq \lambda/2 \cdot |p_1 q_1|$ then we have found the point pair $(p_j = p_1, q_j = q_1)$ that we are searching for. Otherwise, let $e_2 = (p_2, q_2)$ denote the closest pair in $\Delta(p_1, q_1)$, and continue the search for (p_j, q_j). See Fig. 3a for an illustration.

We define e_3 in a similar way, that is, if for each point pair $(u,v) \in \Delta(p_2, q_2)$ it holds that $|uv| \geq \lambda/2 \cdot |p_2 q_2|$ then we have found the point pair $(p_j = p_2, q_j = q_2)$ that we are searching for. Otherwise, let $e_3 = (p_3, q_3)$ denote the closest pair in $\Delta(p_2, q_2)$.

We continue to define the edges $e_4, \ldots, e_j$ and the corresponding graphs $G_4, \ldots, G_j$ until we find a point pair (p_j, q_j) for which there is no edge e_{j+1} such that $(p_{j+1}, q_{j+1}) \in \Delta(p_j, q_j)$ and $|p_{j+1}q_{j+1}| < \lambda/2 \cdot |p_j q_j|$. We claim that G_j has stretch factor at most $(2+\lambda) \cdot t_{\mathcal{P}}$.

Before we continue we need to prove that:

$$d_{G_i}(p_{i+1}, q_{i+1}) \leq t_1 \cdot |p_{i+1}q_{i+1}|. \qquad (6)$$

The inequality is obviously true for $i = 1$. For $i > 1$ it holds that $|p_{i+1}q_{i+1}| < |p_2 q_2|$ which implies that $(p_{i+1}, q_{i+1}) \notin \Delta(p_1, q_1)$ since (p_2, q_2) is the closest pair in $\Delta(p_1, q_1)$. This in turn implies that $d_G(p_{i+1}, q_{i+1}) = d_{G_1}(p_{i+1}, q_{i+1}) \leq t_1 \cdot$

$|p_{i+1}, q_{i+1}|$. Finally, since G is a subgraph of G_i it holds that the length of the shortest path in G_i between p_{i+1} and q_{i+1} must be bounded by the length of the shortest path in G between p_{i+1} and q_{i+1}, which is bounded by $t_1 \cdot |p_{i+1}q_{i+1}|$. Thus, inequality (6) holds.

We continue with the second part of the proof. It will be shown that for every pair $u,v \in V$ it holds that $d_{G_j}(u,v) \leq (2+\lambda) \cdot d_{G_1}(u,v)$. If $(u,v) \notin \Delta(p_1, q_1)$ then we are done since $d_{G_j}(u,v) \leq d_G(u,v) = d_{G_1}(u,v)$. Otherwise, if $(u,v) \in \Delta(p_1, q_1)$, the following holds (see also Fig. 3a for an illustration):

$$
\begin{aligned}
d_{G_j}(u,v) \leq\ & d_{G_1}(u,v) - |p_1 q_1| + (d_{G_1}(p_2, q_2) - |p_1 q_1|) + \\
& + \ldots + (d_{G_{j-1}}(p_j, q_j) - |p_{j-1}q_{j-1}|) + |p_j q_j| \\
<\ & t_1 \cdot |uv| - |p_1 q_1| + ((t_1 - 1) \cdot |p_2 q_2|) + \\
& + \ldots + ((t_1 - 1) \cdot |p_j q_j|) + |p_j q_j| \quad \text{(from (6))} \\
\leq\ & 2t_1 \cdot |uv| - |p_1 q_1| + \\
& + (t_1 - 1)(|p_3 q_3| + \ldots + |p_k q_k|) + |p_j q_j| \\
<\ & 2t_1 \cdot |uv| + (t_1 - 1)(|p_3 q_3| + \ldots + |p_j q_j|) \\
\leq\ & 2t_1 \cdot |uv| + t_1 \cdot \sum_{i=3}^{j} \left(\frac{\lambda}{2}\right)^{i-2} |p_2 q_2| \\
\leq\ & 2t_1 \cdot |uv| + t_1 \cdot |uv| \cdot \sum_{i=1}^{j-2} \left(\frac{\lambda}{2}\right)^{i} \\
<\ & (2+\lambda) \cdot t_1 \cdot |uv|
\end{aligned}
$$

Thus, $t_j < (2+\lambda) \cdot t_1$ which concludes the lemma, since $t_1 = t_{\mathcal{P}}$. □

In the previous lemma we showed the existence of a "short" candidate edge (p,q) for which the resulting graph has small stretch factor. Note that algorithm EXPANDGRAPH might not test (p,q). However, in the following lemma it will be shown that algorithm EXPANDGRAPH will test an edge (a,b) that is almost as good as (p,q).

LEMMA 6. *For any given constant $0 < \varepsilon \leq 1$ it holds that the graph G' returned by algorithm EXPANDGRAPH has stretch factor at most $(2+\varepsilon) \cdot t_{\mathcal{P}}$.*

PROOF. According to Lemma 5 there exists an edge (p,q) such that for every pair $(u,v) \in \Delta(p,q)$ it holds that $|uv| \geq \frac{\lambda}{2}|pq|$, and the stretch factor t_H of $H = G \cup \{(p,q)\}$ is bounded by $(2+\lambda) \cdot t_{\mathcal{P}}$. Let $\{A_i, B_i\}$ be the well-separated pair computed in step 2 of the algorithm such that $p \in A_i$ and $q \in B_i$. According to Definition 2 such a well-separated pair must exist. Next consider the candidate edge (a_i, b_i) tested by the algorithm, such that $a_i, p \in A_i$ and $b_i, q \in B_i$. For simplicity of writing we will use a and b to denote a_i and b_i respectively.

Our claim is that the stretch factor t' of $G' = G \cup \{(a,b)\}$ is bounded by $(1+\varepsilon/4) \cdot t_H$. Thus setting $\lambda = \varepsilon/4$ would then prove the lemma since $(2+\varepsilon/4)(1+\varepsilon/4) < (2+\varepsilon)$, for $\varepsilon \leq 1$.

Now we are ready to prove the claim. To compute the stretch factor of G' one performs a shortest path query between each pair of points in V. If it holds that $(x,y) \notin \Delta(p,q)$ for every pair of points $x,y \in V$ then the claim is obviously true, thus we only have to consider the pairs x,y for which it holds that $(x,y) \in \Delta(p,q)$, see Fig 3b. It holds that:

$$d_G(a,p) = d_H(a,p) \quad \text{and} \quad d_G(b,q) = d_H(b,q). \tag{7}$$

This follows from the fact that the closest pair x', y' for which it holds that $(x',y') \in \Delta(p,q)$ has inter point distance at least $|x'y'| \geq \frac{\varepsilon}{8}|pq|$, according to Lemma 5. It holds that $|ap|$ and $|bq|$ are bounded by $\frac{2}{s}|pq| = \frac{\varepsilon^2}{128}|pq|$ which is less than $\frac{\varepsilon}{8}|pq|$ since $\varepsilon < 1$. As a consequence $(p,q) \notin \delta_H(a,p)$ and $(p,q) \notin \delta_H(b,q)$. Hence, claim (7) holds, which we will need below.

Next, we consider the length of the path in G' between x and y as illustrated in Fig. 3b. Recall that x and y are two arbitrary points of V for which it holds that $(x,y) \in \Delta(p,q)$.

$$
\begin{aligned}
d_{G'}(x,y) \leq & \\
\leq & \; d_G(x,p) + d_G(p,a) + |ab| + d_G(b,q) + d_G(q,y) \\
\leq & \; d_G(x,p) + d_H(p,a) + |ab| + d_H(b,q) + d_G(q,y) \\
\leq & \; d_G(x,p) + |ab| + d_G(q,y) + t_H \cdot (|pa| + |bq|) \\
< & \; d_G(x,p) + (1+4/s) \cdot |pq| + d_G(q,y) + \frac{4t_H}{s} \cdot |pq| \\
\leq & \; d_H(x,y) + \frac{8t_H}{s} \cdot |pq| \\
\leq & \; d_H(x,y) + \frac{64t_H}{\varepsilon s} \cdot |xy| \\
= & \; d_H(x,y) + \frac{\varepsilon}{4} \cdot t_H \cdot |xy|
\end{aligned}
$$

The se3cond inequality follows from (7), the fourth inequality follows from Lemma 3 and the sixth inequality follows from the fact that $(x,y) \in \Delta(p,q)$.

The stretch factor of the path in G' between x and y is:

$$\frac{d_{G'}(x,y)}{|xy|} \leq \frac{d_H(x,y)}{|xy|} + \frac{\frac{\varepsilon}{4}t_H|xy|}{|xy|} \leq \left(1 + \frac{\varepsilon}{4}\right) \cdot t_H.$$

Finally, according to Lemma 5 and the fact that $\lambda = \varepsilon/4$ it holds that $t_H \leq (2+\varepsilon/4) \cdot t_{\mathcal{P}}$. This completes the lemma since $(2+\varepsilon/4)(1+\varepsilon/4) < (2+\varepsilon)$. $\square$

We may now conclude this section with the following theorem.

THEOREM 2. Given a Euclidean graph $G = (V,E)$ in $\mathbb{R}^d$ one can in time $\mathcal{O}(n^3/\varepsilon^{2d})$, using $\mathcal{O}(n^2)$ space, compute a t'-spanner $G' = (V, E \cup \{e\})$, where $t' \leq (2+\varepsilon) \cdot t_{\mathcal{P}}$.

4.3 Speed-up the algorithm

In the previous section we showed that a $(2+\varepsilon)$-approximate solution can be obtained by testing a linear number of candidate edges. Testing each candidate edge entails $\mathcal{O}(n^2)$ shortest path queries. One way to speed up the computation is to compute an approximate stretch factor. t' is said to be a β-approximate stretch factor of G if $t_G \leq t' \leq \beta \cdot t_G$. The problem of computing the approximate stretch factor was considered by Narasimhan and Smid in [25]. They showed the following fact:

FACT 1. (Narasimhan and Smid [25]) Given a Euclidean graph G and a real value $\varepsilon > 0$, a $(1+\varepsilon)^2$-approximative stretch factor of G can be computed by performing $\mathcal{O}(n/\varepsilon^d)$ many $(1+\gamma)$-approximate distance queries, where γ is a positive constant smaller than ε.

Their idea is to compute a well-separated pair decomposition of size $\mathcal{O}(s^d n)$ with respect to $s = 4(1+\varepsilon)/\varepsilon$, and then for each well-separated pair $\{A_i, B_i\}$ select an arbitrary pair $a_i \in A_i$ and $b_i \in B_i$. They prove that these are the only pairs for which the stretch factor needs to be computed.

We will use their idea to speed up step 7 of the algorithm from $\mathcal{O}(n^2)$ to $\mathcal{O}(n/\varepsilon^d)$. On the other hand we will not use the fact that only approximate distance queries are needed, instead the exact shortest distance will be computed, thus $\gamma = 0$. There will be two changes in the EXPANDGRAPH algorithm. First, between steps 2 and 3, the following four lines are inserted:

- $\{(C_j, D_j)\}_{j=1}^{\ell} \leftarrow$ WSPD of the set V w.r.t. $s' = 4(1+\varepsilon)/\varepsilon$.
- **for** $j \leftarrow 1$ **to** ℓ
    Select arbitrary points $c_j \in C_j$ and $d_j \in D_j$.
- $\mathcal{S} = \{(c_1, d_1), \ldots, (c_\ell, d_\ell)\}$

Then, in step 7 of EXPANDGRAPH we will instead of computing the exact stretch factor of G_i make a call to APPROXIMATESTRETCHFACTOR, or ASF for short, with parameters G_i, (a_i, b_i), M, and $\mathcal{S}$. Note that the number of point pairs in $\mathcal{S}$ is bounded by $\mathcal{O}(n/\varepsilon^d)$.

Algorithm ASF$(e, M, \mathcal{S})$
Input: Edge $e = (a,b) \in E$, distance matrix M and a set of point pairs $\mathcal{S}$.
Output: A real value $\mathcal{D}$.
1. $\mathcal{D} \leftarrow 1$
2.   **for** each point pair (c_j, d_j) in $\mathcal{S}$
3.   dist $\leftarrow \min\{M[c_j, d_j]$,
            $M[c_j, a] + |ab| + M[b, d_j]$,
            $M[c_j, b] + |ba| + M[a, d_j]\}$
4. $\mathcal{D} \leftarrow \max\{\mathcal{D}, \text{dist}/|c_j d_j|\}$
5.   **return** $\mathcal{D}$.

We denote the modified algorithm EXPANDGRAPH2.

THEOREM 3. Given a Euclidean graph $G = (V,E)$ and a real constant $\epsilon > 0$ one can in $\mathcal{O}(nm + n^2(\log n + 1/\epsilon^{3d}))$ time, using $\mathcal{O}(n^2)$ space, compute a t'-spanner $G' = (V, E \cup \{e\})$ such that $t' \leq (2+\epsilon) \cdot t_{\mathcal{P}}$.

PROOF. The complexity of all steps of the algorithm, except step 7, is as in Lemma 4. Steps 1–6 requires $\mathcal{O}(mn + n^2 \log n + n/\varepsilon^{2d})$ time. It remains to consider step 7 of the algorithm. Note that the number of times step 7 is executed

is $\mathcal{O}(n/\varepsilon^{2d})$. Procedure ASF performs $\mathcal{O}(n/\varepsilon^d)$ shortest-path queries, instead of $\mathcal{O}(n^2)$, thus the total time needed by step 7 is $\mathcal{O}(\frac{n}{\varepsilon^{2d}} \cdot \frac{n}{\varepsilon^d})$. Summing up the running times gives the stated time complexity.

In Lemma 6 it was proven that the solution returned by algorithm EXPANDGRAPH had a stretch factor that was at most a factor $(2 + \varepsilon)$ worse than the stretch factor of an optimal solution. Since the modified algorithm does not compute the exact stretch factor of a candidate graph, but instead computes a $(1+\varepsilon)^2$-approximate stretch factor it is not hard to verify that the same arguments as in Lemma 6 can be applied to prove that the algorithm EXPANDGRAPH2 returns a graph with stretch factor at most $(1+\varepsilon)^2 \cdot (2 + \varepsilon) \cdot t_{\mathcal{P}}$. Setting $\varepsilon = \min\{\epsilon/10, 1\}$ concludes the proof of the theorem. $\square$

5. A SPECIAL CASE

In the special case when the stretch factor of a graph G is known to be constant there are well-known tools that we can use to decrease the complexity of the algorithms and improve the approximation factor.

FACT 2. *(Corollary 1 in [18]) Let V be a set of n points in $\mathbb{R}^d$, let $t > 1$ and $\varepsilon > 0$ be real numbers, and let $G = (V, E)$ be a t-spanner for V. A $(1+\varepsilon)$-spanner for G having $O((t/\varepsilon)^d n)$ edges can be computed in $O((t/\varepsilon)^d n + m + n \log n)$ time.*

Also, in [17] the following fact was shown.

FACT 3. *([17]) Let V be a set of n points in $\mathbb{R}^d$, let $t > 1$ and $0 < \varepsilon \le 1$ be real numbers, and let $G = (V, E)$ be a t-spanner for V. In $\mathcal{O}(m + \frac{nt^{5d}}{\varepsilon^{2d}}(\log n + (t/\varepsilon)^d))$ time, we can preprocess G into a data structure of size $\mathcal{O}(\frac{t^{3d}}{\varepsilon^{2d}} n \log(tn))$ such that for any two distinct points p and q in V, a $(1+\varepsilon)$-approximation to the shortest-path distance between p and q in G can be computed in time $\mathcal{O}((t^5/\varepsilon^2)^d)$.*

The query structure in Fact 3 is denoted M' and is constructed by algorithm QUERYSTRUCTURE. We have to use a modified version of ASF, denoted ASF$'$, that takes the query structure M' as input instead of the matrix M. The shortest path distance queries using M in ASF is replaced in ASF$'$ by performing approximate shortest path distance queries using M'.

Next we state the main algorithm. Recall that the parameter t is a constant and an upper bound on the stretch factor of the input graph G.

Algorithm EXPANDGRAPH3(G, t, ε)
Input: Euclidean t-spanner $G = (V, E)$ and two real constants $t > 1$ and $\varepsilon > 0$.
Output: Euclidean graph $G' = (V, E \cup \{e\})$.
1. $G' = (V, E') \leftarrow$ PRUNEGRAPH(G, ε) using Fact 2.
2. $M' \leftarrow$ QUERYSTRUCTURE(G, t, ε) using Fact 3.
3. $\mathcal{S} = \{(A_i, B_i)\}_{i=1}^k \leftarrow$ WSPD of V with respect to the separation constant $s = 8(t+1)/\varepsilon$.
4. $t_{\mathcal{C}} \leftarrow \infty$
5.   **for** $i \leftarrow 1$ **to** k
6.   Select arbitrary points $a_i \in A_i$ and $b_i \in B_i$.
7. $G_i' \leftarrow G' \cup (a_i, b_i)$.
8. $t_i \leftarrow$ ASF$'((a_i, b_i), M', \mathcal{S})$.
9.   **if** $t_i < t_{\mathcal{C}}$

10.   **then** $t_{\mathcal{C}} \leftarrow t_i$ and $e_{\mathcal{C}} \leftarrow (a_i, b_i)$
11.   **return** $G' = (V, E \cup \{e_{\mathcal{C}}\})$.

LEMMA 7. EXPANDGRAPH3 *runs in $\mathcal{O}((t^7/\varepsilon^4)^d \cdot n^2)$ time and uses $\mathcal{O}((t^3/\varepsilon^2)^d n \log(tn))$ space.*

PROOF. The time complexity of steps 1–4 is dominated by step 2, thus $\mathcal{O}(m + n(t^5/\varepsilon^2)^d(\log n + (t/\varepsilon)^d))$ time. Step 8 is executed $\mathcal{O}((t/\varepsilon)^d n)$ times, and each iteration requires $\mathcal{O}((t/\varepsilon)^d n \cdot (t^{5d}/\varepsilon^{2d}))$ time according to Facts 1 and 3. Summing up the time bounds gives the time bound stated in the algorithm.

The space bound follows sice the approximate distance oracle stated in Fact 3 only uses $\mathcal{O}((t^3/\varepsilon^2)^d n \log(tn))$ space, instead of the quadratic space needed earlier. $\square$

Now, we show that this algorithm computes a $(1 + \varepsilon)$-approximation of the optimal solution. Note that in EXPANDGRAPH3 the separation constant depends both on ε and t which is the main different compared to the previous algorithms. This allows us to improve the approximation factor.

LEMMA 8. *Given a Euclidean graph $G = (V, E)$ with constant stretch factor t and a positive real constant ε. Let $\{(A_i, B_i)\}_{i=1}^k$ be a well-separated pair decomposition of V with respect to $s = \frac{8(t+1)}{\varepsilon}$. For every pair (A_i, B_i) and all elements $a_1, a_2 \in A_i$ and $b_1, b_2 \in B_i$, let $G_1 = (V, E \cup (a_1, b_1))$ and $G_2 = (V, E \cup (a_2, b_2))$. Let t_1 and t_2 denote the stretch factor of G_1 and G_2, respectively. It holds that $t_1 \le (1+\varepsilon)t_2$.*

PROOF. It suffices to prove that for every pair of points $(u, v) \in \Delta(a_2, b_2)$ there exists a path in G_1 of length at most $(1 + \varepsilon) \cdot d_{G_2}(u, v)$. We have,

$$d_{G_1}(u, v) \le$$
$$\le d_G(u, a_2) + d_G(a_2, a_1) + |a_1 b_1| + d_G(b_1, b_2) + d_G(b_2, v)$$
$$\le d_G(u, a_2) + t|a_1 a_2| + |a_1 b_1| + t|b_1 b_2| + d_G(b_2, v)$$
$$\le d_G(u, a_2) + \frac{4t}{s}|a_2 b_2| + (1 + 4/s) \cdot |a_2 b_2| + d_G(b_2, v)$$
$$< d_G(u, a_2) + |a_2 b_2| + d_G(b_2, v) + \frac{8t}{s}|a_2 b_2|$$
$$= d_{G_2}(u, v) + \frac{t\varepsilon}{t+1}|a_2 b_2|$$
$$< (1 + \varepsilon) \cdot d_{G_2}(u, v)$$

In the second inequality we used Lemma 3, in the fifth inequality we used the fact that $s = 8(t+1)/\varepsilon$ and in the final step we used that $d_{G_2}(u, v) \ge |a_2 b_2|$ since $(u, v) \in \Delta(a_2, b_2)$. The lemma follows. $\square$

Figure 4: Illustrating Lemma 8.

LEMMA 9. *Algorithm EXPANDGRAPH3 returns a graph with stretch factor at most $(1 + \varepsilon)^4 \cdot t_{\mathcal{P}}$.*

PROOF. The first step of the algorithm returns a subgraph G' of G such that G' only has $\mathcal{O}(n)$ edges and G' is a $(1+\varepsilon)$-spanner of G. Assume that $t_{\mathcal{P}}$ is the stretch factor of an optimal solution. Let $t'_{\mathcal{P}}$ denote the stretch factor of the optimal solution given G' as input, and let $t_{\mathcal{C}}$ be the output of the above algorithm. Since G' is a $(1+\varepsilon)$-spanner of G it holds that $t_{\mathcal{P}} \leq t'_{\mathcal{P}} \leq (1+\varepsilon) \cdot t_{\mathcal{P}}$.

We will use the same notations as in the algorithm. Let i be the index such that $t_{\mathcal{C}} = t_i$. If t_i^* is the actual stretch factor of $G'_i = G' \cup (a_i, b_i)$, then by Fact 1, it holds that $t_i^* \leq t_i \leq (1+\varepsilon)^2 \cdot t_i^*$. Hence $t'_{\mathcal{P}} \leq t_i^* \leq t_i = t_{\mathcal{C}}$.

Next assume that $G' \cup (p, q)$ has stretch factor t'_P. Let (A_i, B_i) be the pair in the well-separated pair decomposition such that $p \in A_i$ and $q \in B_i$, or $p \in B_i$ and $q \in A_i$. Without loss of generality we may assume that the first holds. From Fact 1 and Lemma 8 it follows that $t_{\mathcal{C}} \leq t_i \leq (1+\varepsilon)^2 \cdot t_i^* \leq (1+\varepsilon)^3 \cdot t'_{\mathcal{P}}$. Therefore $t'_{\mathcal{P}} \leq t_{\mathcal{C}} \leq (1+\varepsilon)^3 \cdot t'_{\mathcal{P}}$. Since $t_{\mathcal{P}} \leq t'_{\mathcal{P}} \leq (1+\varepsilon) \cdot t_{\mathcal{P}}$ it follows that $t_{\mathcal{P}} \leq t'_{\mathcal{P}} \leq t_{\mathcal{C}} \leq (1+\varepsilon)^3 \cdot t'_{\mathcal{P}} \leq (1+\varepsilon)^4 \cdot t_{\mathcal{P}}$. $\square$

The following theorem follows by setting $\varepsilon = \min\{\epsilon/15, 1\}$ and combining Lemmas 7 and 9.

THEOREM 4. *Given a Euclidean t-spanner $G = (V, E)$ and two real constants $t > 1$ and $\epsilon > 0$ it holds that a graph G' can be computed in time $\mathcal{O}((t^7/\epsilon^4)^d \cdot n^2)$ with stretch factor $(1+\epsilon) \cdot t_{\mathcal{P}}$ using $\mathcal{O}((t^3/\epsilon^2)^d n \log(tn))$ space.*

6. CONCLUDING REMARKS

We considered the problem of adding a shortcut to a Euclidean graph such that the stretch factor of the resulting graph is minimized, and gave several algorithms. Our main result is a $(2+\varepsilon)$-approximation algorithm with running time $\mathcal{O}(nm + n^2(\log n + 1/\varepsilon^{3d}))$ using $\mathcal{O}(n^2)$ space. Several problems remain open.

1. Is there an exact algorithm with running time $o(n^4)$ using linear space?

2. Can we achieve a $(1+\varepsilon)$-approximation within the same time bound as in Theorem 3?

3. A natural extension is to allow more than one edge to be added. Can we generalize our results to this case?

7. REFERENCES

[1] I. Althöfer, G. Das, D. P. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9(1):81–100, 1993.

[2] 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.

[3] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. Smid. Euclidean spanners: short, thin, and lanky. In *Proc. 27th ACM Symposium on Theory of Computing*, pages 489–498, 1995.

[4] S. Arya, D. M. Mount, and M. Smid. Randomized and deterministic algorithms for geometric spanners of

[5] P. Bose, J. Gudmundsson, and P. Morin. Ordered theta graphs. *Computational Geometry: Theory and Applications*, 28:11–18, 2004.

[6] P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields. *Journal of the ACM*, 42:67–90, 1995.

[7] A. Czumaj and A. Lingas. Fast approximation schemes for Euclidean multi-connectivity problems. In *Proc. 27th International Colloquium on Automata, Languages and Programming*, volume 1853 of *Lecture Notes in Computer Science*, pages 856–868. Springer-Verlag, 2000.

[8] G. Das, P. Heffernan, and G. Narasimhan. Optimally sparse spanners in 3-dimensional Euclidean space. In *Proc. 9th Annual ACM Symposium on Computational Geometry*, pages 53–62, 1993.

[9] G. Das and G. Narasimhan. A fast algorithm for constructing sparse Euclidean spanners. *International Journal of Computational Geometry and Applications*, 7:297–315, 1997.

[10] G. Das, G. Narasimhan, and J. Salowe. A new way to weigh malnourished Euclidean graphs. In *Proc. 6th ACM-SIAM Symposium on Discrete Algorithms*, pages 215–222, 1995.

[11] A. Ebbers-Baumann, R. Klein, E. Langetepe, and A. Lingas. A fast algorithm for approximating the detour of a polygonal chain. In *Proc. 9th European Symposium on Algorithms*, volume 2161 of *Lecture Notes in Computer Science*, pages 321–332, Springer-Verlag, 2001.

[12] D. Eppstein. Spanning trees and spanners. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 425–461. Elsevier Science Publishers, Amsterdam, 2000.

[13] D. Eppstein and K. Wortman. Minimum dilation stars. In *Proc. 21th ACM Symposium on Computational Geometry*, 2005.

[14] A. M. Farley, A. Proskurowski, D. Zappala, and K. J. Windisch. Spanners and message distribution in networks. *Discrete Applied Mathematics*, 137(2):159–171, 2004.

[15] A. Grüne. Tightness of upper bound on t_B/t_P. Manuscript, March 2005.

[16] J. Gudmundsson, C. Levcopoulos, and G. Narasimhan. Improved greedy algorithms for constructing sparse geometric spanners. *SIAM Journal of Computing*, 31(5):1479–1500, 2002.

[17] J. Gudmundsson, C. Levcopoulos, G. Narasimhan, and M. Smid. Approximate distance oracles revisited. In *Proc. 13th International Symposium on Algorithms and Computation*, volume 2518 of *Lecture Notes in Computer Science*, pages 357–368, Springer-Verlag, 2002.

[18] J. Gudmundsson, G. Narasimhan, and M. Smid. Fast pruning of geometric spanners. In *Proc. 22nd International Symposium on Theoretical Aspects of Computer Science*, volume 3404 of *Lecture Notes in Comp. Science*, pages 508–520, Springer-Verlag, 2005.

small diameter. In *Proc. 35th IEEE Symposium on Foundations of Computer Science*, pages 703–712, 1994.

[19] J. M. Keil. Approximating the complete Euclidean graph. In *Proc. 1st Scandinavian Workshop on Algorithmic Theory*, volume 318 of *Lecture Notes in Computer Science*, pages 208–213, Springer-Verlag, 1988.

[20] J. M. Keil and C. A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete and Computational Geometry*, 7:13–28, 1992.

[21] S. Langerman, P. Morin, and M. A. Soss. Computing the maximum detour and spanning ratio of planar paths trees, and cycles. In *Proc. 19th International Symposium on Theoretical Aspects of Computer Science*, volume 2285 of *Lecture Notes in Computer Science*, pages 250–261, Springer-Verlag, 2002.

[22] C. Levcopoulos and A. Lingas. There are planar graphs almost as good as the complete graphs and almost as cheap as minimum spanning trees. *Algorithmica*, 8:251–256, 1992.

[23] C. Levcopoulos, G. Narasimhan, and M. Smid. Improved algorithms for constructing fault-tolerant spanners. *Algorithmica*, 32:144–156, 2002.

[24] X.-Y. Li. Applications of computational geomety in wireless ad hoc networks. In X.-Z. Cheng, X. Huang, and D.-Z. Du, editors, *Ad Hoc Wireless Networking*. Kluwer, 2003.

[25] G. Narasimhan and M. Smid. Approximating the stretch factor of Euclidean graphs. *SIAM Journal of Computing*, 30(3):978–989, 2000.

[26] G. Navarro and R. Paredes. Practical construction of metric *t*-spanners. In *Proc. 5th SIAM Workshop on Algorithm Engineering and Experiments*, pages 69–81, 2003.

[27] G. Navarro, R. Paredes, and E. Chvez. *t*-spanners as a data structure for metric space searching. In *Proc. 9th International Symposium on String Processing and Information Retrieval*, volume 2476 of *Lecture Notes in Comp. Science*, pages 298–309. Springer-Verlag, 2002.

[28] S. Rao and W. D. Smith. Approximating geometrical graphs via spanners and banyans. In *Proc. 30th ACM Symposium on the Theory of Computing*, pages 540–550, 1998.

[29] J. S. Salowe. Constructing multidimensional spanner graphs. *International Journal of Computational Geometry & Applications*, 1:99–107, 1991.

[30] M. Smid. Closest point problems in computational geometry. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 877–935. Elsevier Science Publishers, Amsterdam, 2000.

[31] P. M. Vaidya. A sparse graph almost as good as the complete graph on points in K dimensions. *Discrete Computational Geometry*, 6:369–381, 1991.