Improving the stretch factor of a geometric network by edge augmentation

Citation for published version (APA):

Farshi, M., Giannopoulos, P., & Gudmundsson, J. (2008). Improving the stretch factor of a geometric network by edge augmentation. *SIAM Journal on Computing*, 38(1), 226-240. <https://doi.org/10.1137/050635675>

DOI:

[10.1137/050635675](https://doi.org/10.1137/050635675)

Document status and date:

Published: 01/01/2008

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

IMPROVING THE STRETCH FACTOR OF A GEOMETRIC NETWORK BY EDGE AUGMENTATION*

MOHAMMAD FARSHI[†], PANOS GIANNOPOULOS[‡], AND JOACHIM GUDMUNDSSON[§]

Abstract. Given a Euclidean graph G in $\mathbb{R}^d$ with n vertices and m edges, we consider the problem of adding an edge to G such that the stretch factor of the resulting graph is minimized. Currently, the fastest algorithm for computing the stretch factor of a graph with positive edge weights runs in $\mathcal{O}(nm + n^2 \log n)$ time, resulting in a trivial $\mathcal{O}(n^3 m + n^4 \log n)$ -time algorithm for computing the optimal edge. 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 time we consider several approximation algorithms.

Key words. computational geometry, approximation algorithms, geometric networks

AMS subject classifications. 65D18, 68U05, 68Q25

DOI. 10.1137/050635675

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 (u, v) has a positive weight $w(u, v)$. A Euclidean network is a geometric network where the weight of the edge (u, v) is equal to the Euclidean distance $|uv|$ between its two endpoints u and v .

For two vertices u, v in a weighted graph G we use $\delta_G(u, v)$ to denote a shortest path between u and v in G , and the length of the path is denoted by $d_G(u, v)$. Consider a weighted graph $G = (V, E)$ and a graph $G' = (V, E')$ on the same vertex set but with edge set $E' \subseteq E$. We say that G' is a t -spanner of G if for each pair of vertices $u, v \in V$ we have that $d_{G'}(u, v) \leq t \cdot d_G(u, v)$. The minimum t such that G is a t -spanner for V is called the stretch factor, or dilation, of G .

We say that a Euclidean network $G = (V, E)$ is a t -spanner if $G = (V, E)$ is a t -spanner of the complete network on V . In other words, for any two points $p, q \in V$ the graph distance in G is at most t times the Euclidean distance between the two points.

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. Recently spanners found interesting practical applications

*Received by the editors July 11, 2005; accepted for publication (in revised form) November 20, 2007; published electronically March 28, 2008.

<http://www.siam.org/journals/sicomp/38-1/63567.html>

[†]School of Computer Science, Carleton University, Ottawa, ON, K1S 5B6, Canada (mfarshi@cgs.scs.carleton.ca) and Department of Computer Science, Yazd University, P.O. Box 89195-741, Yazd, Iran (mfarshi@yazduni.ac.ir). This author's research was supported by the Ministry of Science, Research and Technology of I.R. Iran.

[‡]Institut für Informatik, Humboldt-Universität zu Berlin, D-10099 Berlin, Germany (panos@informatik.tu-berlin.de). This author's research was partially conducted at IICS, Utrecht University, The Netherlands.

[§]IMAGEN Program, National ICT Australia Ltd., NSW Alexandria 1435, Sydney, Australia (joachim.gudmundsson@nicta.com.au). NICTA is funded through the Australian government's Backing Australia's Ability initiative, in part through the Australian Research Council.

TABLE 1
Complexity bounds for the algorithms presented in the paper.

Input graph	Approximation factor	Time complexity	Space	Section
Weighted graph	1	$\mathcal{O}(n^3m + n^4 \log n)$	$\mathcal{O}(m)$	2.1
Weighted graph	1	$\mathcal{O}(n^4)$	$\mathcal{O}(n^2)$	2.1
Euclidean graph	$1 + \varepsilon$	$\mathcal{O}(n^3/\varepsilon^d)$	$\mathcal{O}(n^2)$	2.2
Weighted graph	3	$\mathcal{O}(nm + n^2 \log n)$	$\mathcal{O}(m)$	3
Euclidean graph	$2 + \varepsilon$	$\mathcal{O}(nm + n^2(\log n + 1/\varepsilon^{3d}))$	$\mathcal{O}(n^2)$	4
Euclidean t -spanner	$1 + \varepsilon$	$\mathcal{O}((t^7/\varepsilon^4)^d \cdot n^2)$	$\mathcal{O}(m + (t^3/\varepsilon^2)^d n \log(tn))$	5

in areas such as metric space searching [29, 30] and broadcasting in communication networks [2, 14, 25].

Several well-known theoretical results also use the construction of t -spanners as a building block; for example, Rao and Smith [32] 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 multiconnectivity problems in geometric graphs. The problem of constructing geometric spanners has received considerable attention from a theoretical perspective; see [1, 3, 4, 5, 8, 9, 10, 17, 20, 21, 23, 24, 33, 36], the surveys [12, 16, 34], and the book by Narasimhan and Smid [28]. Note that considerable research has also been done in the construction of spanners for general graphs; see, for example, the book by Peleg [31] or the recent work by Elkin and Peleg [11] and Thorup and Zwick [35].

All the existing algorithms construct a network from scratch, but in many applications the 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. The problem was first stated by Narasimhan [26] and, surprisingly, it had not been studied earlier, 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.

The results presented in this paper are summarized in Table 1. Note that some of the presented bounds hold for any graph with positive edge weights (*weighted graphs*), while some hold only for Euclidean graphs.

Finally, throughout this paper we will use $G_{\mathcal{P}}$ to denote the optimal solution, while $t_{\mathcal{P}}$ and t denote the stretch factor of $G_{\mathcal{P}}$ and the input graph G , respectively.

2. Three simple algorithms. A naive approach to deciding which edge to add is to test every possible candidate edge. The number of such edges is obviously $(\frac{n(n-1)}{2} - m) = \mathcal{O}(n^2)$. Testing a candidate edge e entails computing the stretch factor of the graph $G' = (V, E \cup \{e\})$, denoted the candidate graph. Therefore we briefly consider the problem of computing the stretch factor of a given graph with

positive edge weights. This problem has recently received considerable attention; see, for example, [13, 22, 27].

2.1. Exact algorithms. 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 trivial upper bound is obtained by computing the length of the shortest paths between every pair of vertices in G' . This can be done by running Dijkstra's algorithm—implemented using Fibonacci heaps— n times, resulting in an $\mathcal{O}(mn + n^2 \log n)$ -time algorithm using linear space. This approach 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, 22, 27]. Applying the stated bound to the problem of computing the exact stretch factor of G' gives that $G_{\mathcal{P}}$ can be computed in time $\mathcal{O}(n^3(m + n \log n))$ using linear space.

A small 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 whether there is a shorter path using the edge (u, v) . That is, we only have to compute $d_G(x, u) + w(u, v) + d_G(v, y)$, $d_G(x, v) + w(v, u) + d_G(u, y)$, and $d_G(x, y)$, which can be done in constant time since the length of a shortest path between every pair of vertices in G has already been computed (provided that we store this information). Hence, by first computing all-pair-shortest paths of G we obtain the following lemma.

LEMMA 1. *Given a graph G with positive edge weights, 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 all the distances are stored in an $n \times n$ matrix. The $\mathcal{O}(n^2)$ edges are tested for insertion: for each candidate edge compute the length of the shortest path between every pair of points in G , each of which can be done in constant time as described above. $\square$

2.2. A $(1 + \varepsilon)$ -approximation for Euclidean graphs. In the previous section we showed that an optimal solution can be obtained by testing a quadratic number of candidate edges. Testing each candidate edge entails $\mathcal{O}(n^2)$ distance queries, where a distance query asks for the length of a shortest path in the graph between two query points. 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 $t_G \leq t' \leq \beta \cdot t_G$, where t_G is the stretch factor of G . The problem of computing an approximate stretch factor of a geometric graph was considered by Narasimhan and Smid in [27]. They showed the following fact.

Fact 1 (Narasimhan and Smid [27]). *Given a Euclidean graph G and a real value $\tau > 0$, a $(1 + \tau)^2$ -approximate stretch factor of G can be computed by performing $\mathcal{O}(n/\tau^d)$ many $(1 + \gamma)$ -approximate distance queries, where γ is a positive constant smaller than τ .*

The algorithm is almost as stated in the previous section with the exception that when the stretch factor of the candidate graph is computed we approximate it by only performing $\mathcal{O}(n/\tau^d)$ shortest path queries as stated in Fact 1. As a result the time to compute the stretch factor decreases from $\mathcal{O}(n^2)$ to $\mathcal{O}(n/\tau^d)$; thus the total running time decreases from $\mathcal{O}(n^4)$ to $\mathcal{O}(n^3/\tau^d)$.

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

Proof. The time bound follows from the above discussion setting $\tau = \sqrt{1+\varepsilon} - 1$, where τ is as stated in Fact 1. It remains to prove that G' has stretch factor $((1+\varepsilon) \cdot t_{\mathcal{P}})$.

For each candidate graph G'_i , let t'_i be its approximate stretch factor as computed by the algorithm, and let t_i be its exact stretch factor. From Fact 1 it follows that for each candidate graph G'_i , $t'_i \leq (1+\tau)^2 \cdot t_i$. Assume that $t_{\mathcal{P}} = t_j$ and that $t' = t'_k = \min_i t'_i$, for some indices j and k . As a result we have

$$t' = t'_k \leq t'_j \leq (1+\tau)^2 \cdot t_j = (1+\tau)^2 \cdot t_{\mathcal{P}} = (1+\varepsilon) \cdot t_{\mathcal{P}} \quad \text{and} \quad t_{\mathcal{P}} \leq t_k \leq t'_k = t'.$$

Thus, $t_{\mathcal{P}} \leq t' \leq (1+\varepsilon) \cdot t_{\mathcal{P}}$. $\square$

3. Adding a bottleneck edge. Consider a graph $G = (V, E)$ with positive edge weights and stretch factor t . In this section we analyze the following simple algorithm: Add an edge between a pair of vertices in G with stretch factor t ; this edge is called a *bottleneck edge* of G .

Let $G_{\mathcal{B}}$ be a graph obtained from G by adding a 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 can be decided, i.e., in $\mathcal{O}(mn + n^2 \log n)$ time for graphs with positive edge weights.

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

LEMMA 3. *Given a graph G with positive edge weights, it holds that $t_{\mathcal{B}} < 3t_{\mathcal{P}}$.*

Proof. Recall that t denotes the stretch factor of G and that $G_{\mathcal{P}}$ denotes the optimal graph. Let (x, y) be the edge added to G to obtain $G_{\mathcal{P}}$, and let (u, v) be the edge added to G to obtain $G_{\mathcal{B}}$; i.e., (u, v) is a bottleneck edge of G , as illustrated in Figure 1.

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 are endpoints of a bottleneck edge of $G_{\mathcal{B}}$. Fix a path $\delta_{G_{\mathcal{P}}}(a, b)$. If this path does not include the edge (x, y) , then $d_{G_{\mathcal{P}}}(a, b) = d_G(a, b) \geq d_{G_{\mathcal{B}}}(a, b)$ and we are done. Therefore, we may assume that the path $\delta_{G_{\mathcal{P}}}(a, b)$ includes (x, y) . 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 ; otherwise the labels a and b may be switched. Note that $\delta_{G_{\mathcal{P}}}(u, v)$ must pass through (x, y) ; otherwise we have $t_{\mathcal{P}} \geq d_{G_{\mathcal{P}}}(u, v)/|uv| = d_G(u, v)/|uv| = t$, which means that $t = t_{\mathcal{P}}$, which contradicts the assumption that $t_{\mathcal{P}} \leq t/3$. Furthermore, we assume that a shortest path in $G_{\mathcal{P}}$ from u to v goes from u to x and then to v via y ; otherwise the labels u and v may be switched.

As a first step we bound the distance between the endpoints of the bottleneck edge u and v . This is done by bounding the length of the path in G between x and y

as follows (see Figure 1):

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

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

$$(1) \quad |uv| < 3/2 \cdot |xy|.$$

Also,

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

which implies that

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

and

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

which gives that

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

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

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

Now we are ready to put together the results:

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

This completes the proof of the lemma since $t_B < 3t_P$. $\square$

We conclude by stating the main result of this section followed by a lower bound for the bottleneck approach.

THEOREM 4. *Given a graph $G = (V, E)$ with positive edge weights, a t_B -spanner $G' = (V, E \cup \{e\})$ with $t_B < 3t_P$ can be computed in $\mathcal{O}(mn + n^2 \log n)$ time using $\mathcal{O}(m)$ space.*

OBSERVATION 1. *There exists a Euclidean graph G such that $(2 - \varepsilon) \cdot t_P \leq t_B$ for any $0 < \varepsilon < 1$.*

Proof. Consider the graph G , as in Figure 2(a). More specifically, G is a graph with ten vertices $p_i = ((i - 1) \bmod 5, \lfloor (i - 1)/5 \rfloor \cdot \delta)$, $1 \leq i \leq 10$, 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 $\delta \leq 1$, (p_1, p_6) is a bottleneck edge in G and $t_B = \frac{4+\delta}{\delta}$; see Figure 2(b).

In the case where edge (p_2, p_7) is added to G , as shown in Figure 2(c), the resulting graph has stretch factor $(2 + \delta)/\delta$. Combining the upper and lower bounds gives $\frac{t_B}{t_P} \geq \frac{4+\delta}{2+\delta} = (2 - \varepsilon)$, where the last equality follows if we set $\delta = \min\{1, \frac{2\varepsilon}{1-\varepsilon}\}$. $\square$

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

FIG. 2. (a) The input graph G , (b) the graph G_B , and (c) the graph G_P .

4. A $(2 + \varepsilon)$ -approximation for Euclidean graphs. In the remainder of the paper we will develop approximation algorithms for Euclidean graphs. In this section we present a fast approximation algorithm which guarantees an approximation factor of $(2 + \varepsilon)$. The algorithm is similar to the algorithms presented in section 2 in the sense that it tests candidate edges. Testing a candidate edge entails computing the stretch factor of the input graph augmented with the candidate edge. The main difference is that we will show, in section 4.2, that only a linear number of candidate edges need to be tested to obtain a solution that gives a $(2 + \varepsilon)$ -approximation, instead of a quadratic number of edges.

Moreover, in section 4.3 we show 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 5 (see [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 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 5.

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

DEFINITION 7 (see [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 nonempty subsets of S , $(A_1, B_1), \dots, (A_m, B_m)$, such that*

1. $A_i \cap B_i = \emptyset$ for all $i = 1, \dots, m$,
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 with respect to 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. Here we will show that only a linear number of candidate edges need to be tested to obtain a solution that gives a $(2 + \varepsilon)$ -approximation.

The approach is straightforward. 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 WSPD is computed. Note that, in step 5, the candidate edges will be chosen using the WSPD. In step 6, the function STRETCHFACTOR returns the stretch factor of the graph on V with edge set $E \cup (a_i, b_i)$; i.e., in steps 5–8, a candidate edge is tested by computing the stretch factor of G with the candidate edge (a_i, b_i) added to G .

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. $t_i \leftarrow$ STRETCHFACTOR(a_i, b_i, M).
 7. **if** $t_i < t'$
 8. **then** $t' \leftarrow t_i$ and $e \leftarrow (a_i, b_i)$
 9. **return** $G' = (V, E \cup \{e\})$.
-

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

LEMMA 8. *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 6, is straightforward to calculate. Recall that step 1 requires $\mathcal{O}(mn + n^2 \log n)$ time and quadratic space, and step 2 requires $\mathcal{O}(n/\varepsilon^{2d} + n \log n)$ time according to section 4.1. Thus, it remains to consider step 6 of the algorithm. Note that the number of times step 6 is executed is $\mathcal{O}(n/\varepsilon^{2d})$.

Let $G_i = (V, E \cup \{(a_i, b_i)\})$. Since we computed the all-pair-shortest distances of G and stored the results in a matrix M , it holds that shortest path 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, a quadratic number of 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 WSPD depends only 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 . In Lemma 9, we show the existence of a “short” edge e' that is a good approximation of the optimal edge and then, in Lemma 10, we show that EXPANDGRAPH computes a good approximation of e' .

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, $\Delta(p, q)$ is the set of point pairs for which a shortest path between them in $G \cup \{(p, q)\}$ passes through (p, q) .

LEMMA 9. *For any given constant $0 < \lambda \leq 1$, there exists a point pair $p, q \in V$ such that*

- (I) $|uv| \geq \frac{\lambda}{2}|pq|$ for every pair $(u, v) \in \Delta(p, q)$, and
- (II) 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 (I). Then we prove that this pair will also fulfill (II), i.e., that the stretch factor of $G \cup \{(p_j, q_j)\}$ is bounded by $(2 + \lambda) \cdot t_{\mathcal{P}}$.

Consider an optimal solution $G_1 = G \cup \{(p_1, q_1)\}$. If (p_1, q_1) fulfills (I), then we are done; i.e., we have found the point pair $(p = p_1, q = q_1)$ we are searching for. Otherwise, let $e_2 = (p_2, q_2)$ denote the closest pair in $\Delta(p_1, q_1)$. Since there exists a pair $(u, v) \in \Delta(p_1, q_1)$ such that $|uv| < \frac{\lambda}{2} \cdot |p_1 q_1|$ and since (p_2, q_2) is the closest pair in $\Delta(p_1, q_1)$ we have $|p_2 q_2| < \frac{\lambda}{2} \cdot |p_1 q_1|$, as illustrated in Figure 3(a).

If (p_2, q_2) fulfills (I), then $(p = p_2, q = q_2)$ and we are done. Otherwise, let $e_3 = (p_3, q_3)$ denote the closest pair in $\Delta(p_2, q_2)$. We continue this procedure until we find a point pair (p_j, q_j) that satisfies (I). Since, for each $i > 0$, $|p_{i+1} q_{i+1}| < \frac{\lambda}{2} \cdot |p_i q_i|$, the process must terminate.

Now for each $1 \leq i \leq j$, let $G_i = G \cup \{(p_i, q_i)\}$ where (p_i, q_i) are the point pairs constructed above. We claim that G_j has stretch factor at most $(2 + \lambda) \cdot t_{\mathcal{P}}$. Before we continue we need to prove

$$(6) \quad d_{G_i}(p_{i+1}, q_{i+1}) \leq t_{\mathcal{P}} \cdot |p_{i+1} q_{i+1}|.$$

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 $d_G(p_{i+1}, q_{i+1}) = d_{G_1}(p_{i+1}, q_{i+1}) \leq t_{\mathcal{P}} \cdot |p_{i+1}, q_{i+1}|$. Since G is a subgraph of G_i , 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 $t_{\mathcal{P}} \cdot |p_{i+1} q_{i+1}|$. Thus, inequality (6) holds.

We continue with the second part of the proof. 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

FIG. 3. (a) Illustrating the proof of Lemma 9. (b) Illustrating the proof of Lemma 10.

following holds (see Figure 3(a) for an illustration):

$$\begin{aligned}
 d_{G_j}(u, v) &\leq d_{G_1}(u, v) - |p_1 q_1| + (d_{G_1}(p_2, q_2) - |p_1 q_1|) + \dots \\
 &\quad + (d_{G_{j-1}}(p_j, q_j) - |p_{j-1} q_{j-1}|) + |p_j q_j| \\
 &< t_{\mathcal{P}} \cdot |uv| - |p_1 q_1| + (t_{\mathcal{P}} \cdot |p_2 q_2| - |p_1 q_1|) + \dots \\
 &\quad + (t_{\mathcal{P}} \cdot |p_j q_j| - |p_{j-1} q_{j-1}|) + |p_j q_j| \quad (\text{cf. (6)}) \\
 &= t_{\mathcal{P}} \cdot |uv| - 2|p_1 q_1| + ((t_{\mathcal{P}} - 1) \cdot |p_2 q_2|) + \dots + ((t_{\mathcal{P}} - 1) \cdot |p_j q_j|) + 2|p_j q_j| \\
 &< t_{\mathcal{P}} \cdot |uv| + (t_{\mathcal{P}} - 1)(|p_2 q_2| + \dots + |p_j q_j|) \quad (\text{since } |p_j q_j| < |p_1 q_1|) \\
 &< t_{\mathcal{P}} \cdot |uv| + t_{\mathcal{P}} \cdot \sum_{i=2}^j \left(\frac{\lambda}{2}\right)^{i-2} |p_2 q_2| \quad (\text{since } |p_{i+1} q_{i+1}| \leq (\lambda/2) \cdot |p_i q_i|) \\
 &\leq t_{\mathcal{P}} \cdot |uv| + t_{\mathcal{P}} \cdot |uv| \cdot \sum_{i=0}^{j-2} \left(\frac{\lambda}{2}\right)^i \quad (\text{since } |p_2 q_2| \leq |uv|) \\
 &= 2t_{\mathcal{P}} \cdot |uv| + t_{\mathcal{P}} \cdot |uv| \cdot \sum_{i=1}^{j-2} \left(\frac{\lambda}{2}\right)^i \\
 &\leq 2t_{\mathcal{P}} \cdot |uv| + t_{\mathcal{P}} \cdot |uv| \cdot \lambda \cdot \sum_{i=1}^{j-2} \frac{1}{2^i} \quad (\text{since } \lambda \leq 1) \\
 &< (2 + \lambda) \cdot t_{\mathcal{P}} \cdot |uv|.
 \end{aligned}$$

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

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 10. *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 9, 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 7 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' the algorithm performs a shortest path distance 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 have to consider only the pairs x, y for which it holds that $(x, y) \in \Delta(p, q)$; see Figure 3(b). Now the claim is

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

Lemma 9 states that if $(x', y') \in \Delta(p, q)$, then $|x'y'| \geq \frac{\varepsilon}{8}|pq|$. But by Lemma 6 the distances $|ap|$ and $|bq|$ are less than $\frac{2}{s}|pq| = \frac{\varepsilon^2}{128}|pq|$, which is less than $\frac{\varepsilon}{8}|pq|$ since $\varepsilon \leq 1$. As a consequence, $(a, p) \notin \Delta(p, q)$ and $(b, q) \notin \Delta(p, q)$; thus $(p, q) \notin \delta_H(a, p)$ 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 Figure 3(b). Recall that x and y are two arbitrary points of V for which it holds that $(x, y) \in \Delta(p, q)$. Without loss of generality we have

$$\begin{aligned} d_{G'}(x, y) &\leq d_G(x, p) + d_G(p, a) + |ab| + d_G(b, q) + d_G(q, y) \\ &\leq d_G(x, p) + d_H(p, a) + |ab| + d_H(b, q) + d_G(q, y) && \text{(cf. (7))} \\ &\leq d_G(x, p) + |ab| + d_G(q, y) + t_H \cdot (|pa| + |bq|) \\ &\leq d_G(x, p) + (1 + 4/s) \cdot |pq| + d_G(q, y) + \frac{4t_H}{s} \cdot |pq| && \text{(Lemma 6)} \\ &\leq d_H(x, y) + \frac{8t_H}{s} \cdot |pq| \\ &\leq d_H(x, y) + \frac{64t_H}{\varepsilon s} \cdot |xy| && \text{(Lemma 9)} \\ &= d_H(x, y) + \frac{\varepsilon}{4} \cdot t_H \cdot |xy|. \end{aligned}$$

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 9 and the fact that $\lambda = \varepsilon/4$, it holds that $t_H \leq (2 + \varepsilon/4) \cdot t_P$. 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 11. *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 $t' \leq (2 + \varepsilon) \cdot t_P$.*

4.3. Speeding up Algorithm EXPANDGRAPH. 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. As in section 2.2 we will use Fact 1 by Narasimhan and Smid [27].

Their idea is to compute a WSPD of size $\mathcal{O}(s^d n)$ with respect to $s = 4(1 + \tau)/\tau$, 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 $(1 + \tau)^2$ -approximate stretch factor needs to be computed.

We will use their idea to speed up step 6 of EXPANDGRAPH from $\mathcal{O}(n^2)$ to $\mathcal{O}(n/\varepsilon^d)$; i.e., we check a linear number of pairs in order to compute an approximate stretch factor using Fact 1. However, we will not use the fact that only approximate distance queries are needed; instead the exact shortest distance will be computed, and thus $\gamma = 0$ where γ is as stated in Fact 1. There will be two main changes in the EXPANDGRAPH algorithm; two WSPDs will be computed, and the computation of the stretch factor will be different. Instead of computing the exact stretch factor of G with the candidate edge (a_i, b_i) added to G , we compute the approximate stretch factor. This is done by a call to APPROXIMATESTRETCHFACTOR, or ASF for short, with parameters (a_i, b_i) , M , and $\mathcal{S}$. The ASF algorithm is stated in more detail below. Note that the number of point pairs in $\mathcal{S}$ is bounded by $\mathcal{O}(n/\varepsilon^d)$.

ALGORITHM EXPANDGRAPH2(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 $s = 256/\varepsilon^2$.
 3. $\{(C_j, D_j)\}_{j=1}^\ell \leftarrow$ WSPD of the set V with respect to $s' = 4(1 + \varepsilon)/\varepsilon$.
 4. **for** $j \leftarrow 1$ **to** ℓ
 5. Select an arbitrary point c_j of C_j and an arbitrary point d_j of D_j .
 6. $\mathcal{S} = \{(c_1, d_1), \dots, (c_\ell, d_\ell)\}$.
 7. $t' \leftarrow \infty$.
 8. **for** $i \leftarrow 1$ **to** k
 9. Select an arbitrary point a_i of A_i and an arbitrary point b_i of B_i .
 10. $t_i \leftarrow \text{ASF}((a_i, b_i), M, \mathcal{S})$.
 11. **if** $t_i < t'$
 12. **then** $t' \leftarrow t_i$ and $e \leftarrow (a_i, b_i)$
 13. **return** $G' = (V, E \cup \{e\})$.
-

For completeness we also state the ASF algorithm.

ALGORITHM ASF($(a, b), M, \mathcal{S}$)

Input: Vertex pair $(a, b) \in V^2$, 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. $\text{dist} \leftarrow \min\{M[c_j, d_j], M[c_j, a] + |ab| + M[b, d_j], M[c_j, b] + |ba| + M[a, d_j]\}$
 4. $\mathcal{D} \leftarrow \max\{\mathcal{D}, \text{dist}/|c_j d_j|\}$
 5. **return** $\mathcal{D}$.
-

THEOREM 12. *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 $t' \leq (2 + \epsilon) \cdot t_P$.*

Proof. The complexity of all steps of the algorithm, except step 10, is as in Lemma 8. Steps 1–7 require $\mathcal{O}(mn + n^2 \log n + n/\varepsilon^{2d})$ time. It remains to consider

step 10 of the algorithm. Note that the number of times step 10 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 10 is $\mathcal{O}(\frac{n}{\varepsilon^{2d}} \cdot \frac{n}{\varepsilon^d})$. Summing up the running times gives the stated time complexity.

In Lemma 10 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 10 can be applied to prove that the algorithm EXPANDGRAPH2 returns a graph with stretch factor at most $(1 + \varepsilon)^2 \cdot (2 + \varepsilon) \cdot t\rho$. Setting $\varepsilon = \min\{\varepsilon/10, 1\}$ concludes the proof of the theorem. $\square$

5. A special case: G has constant stretch factor. 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 both the time complexity and the space complexity of the algorithms and improve the approximation factor.

Fact 2 (see [18]). Let V be a set of n points in $\mathbb{R}^d$, let $t > 1$ and $0 < \varepsilon \leq 1$ 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 2 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 are 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 . Also note that this algorithm only needs one WSPD.

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. $M' \leftarrow \text{QUERYSTRUCTURE}(G, t, \varepsilon)$ using Fact 2.
 2. $\{(A_i, B_i)\}_{i=1}^k \leftarrow \text{WSPD of } V \text{ with respect to the separation constant } s = 8(t + 1)/\varepsilon$.
 3. **for** $j \leftarrow 1$ **to** k
 4. Select an arbitrary point a_j of A_j and an arbitrary point b_j of B_j .
 5. $\mathcal{S} = \{(a_1, b_1), \dots, (a_k, b_k)\}$.
 6. $t_C \leftarrow \infty$.
 7. **for** $i \leftarrow 1$ **to** k
 8. $t_i \leftarrow \text{ASF}'((a_i, b_i), M', \mathcal{S})$.
 9. **if** $t_i < t_C$
 10. **then** $t_C \leftarrow t_i$ and $e_C \leftarrow (a_i, b_i)$
 11. **return** $G' = (V, E \cup \{e_C\})$.
-

LEMMA 13. *Algorithm 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 1–3 is dominated by step 1; 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 2. Summing up the time bounds gives the time bound stated in the algorithm.

The space bound follows since the approximate distance oracle stated in Fact 2 uses only $\mathcal{O}((t^3/\varepsilon^2)^d n \log tn)$ 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 on both ε and t , which is the main difference compared to the previous algorithms. This allows us to improve the approximation factor.

LEMMA 14. *Let $G = (V, E)$ be a Euclidean graph with constant stretch factor t and a positive real constant ε , and let $\{(A_i, B_i)\}_{i=1}^k$ be a WSPD of V with respect to $s = \frac{8(t+1)}{\varepsilon}$. For every pair (A_i, B_i) and any 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 $G_2 = (V, E \cup \{(a_2, b_2)\})$, and let t_1 and t_2 denote the stretch factor of G_1 and G_2 , respectively. It holds that $t_1 \leq (1+\varepsilon)t_2$.*

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 $(1+\varepsilon) \cdot d_{G_2}(u, v)$. Without loss of generality we may assume that the shortest path between u and v in G_2 goes from u to a_2 and to v via b_2 . We have

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

In the second inequality we used Lemma 6, in the fifth inequality we used the fact that $s = 8(t+1)/\varepsilon$, and in the final step we used that $d_{G_2}(u, v) \geq |a_2 b_2|$ since $(u, v) \in \Delta(a_2, b_2)$. The lemma follows. $\square$

LEMMA 15. *Algorithm EXPANDGRAPH3 returns a graph with stretch factor at most $(1+\varepsilon)^3 \cdot t_{\mathcal{P}}$.*

Proof. Assume that $t_{\mathcal{P}}$ is the stretch factor of an optimal solution $G \cup \{(p, q)\}$, and let G' with stretch factor $t_{\mathcal{C}}$ be the output of the above algorithm.

We will use the same notation as in the algorithm. For each i let t_i^* be the stretch factor of $G_i = G \cup \{(a_i, b_i)\}$. According to Fact 1 we have $t_i^* \leq t_i \leq (1+\varepsilon)^2 \cdot t_i^*$ for each i .

Let (A_j, B_j) be the pair in the WSPD such that $p \in A_j$ and $q \in B_j$, or $p \in B_j$ and $q \in A_j$. From Lemma 14 it follows that $t_j^* \leq (1+\varepsilon) \cdot t_{\mathcal{P}}$. As a result it follows that $t_{\mathcal{C}} \leq t_j \leq (1+\varepsilon)^2 \cdot t_j^* \leq (1+\varepsilon)^3 \cdot t_{\mathcal{P}}$. Therefore $t_{\mathcal{P}} \leq t_{\mathcal{C}} \leq (1+\varepsilon)^3 \cdot t_{\mathcal{P}}$, which completes the lemma. $\square$

The following theorem follows by setting $\varepsilon = \min\{\varphi/15, 1\}$ and combining Lemmas 13 and 15.

THEOREM 16. *Let V be a set of n points in $\mathbb{R}^d$, let $t > 1$ and $\varphi > 0$ be real numbers, and let $G = (V, E)$ be a t -spanner of V . One can in $\mathcal{O}((t^7/\varphi^4)^d \cdot n^2)$ time, using $\mathcal{O}((t^3/\varphi)^d n \log tn)$ space, compute a t' -spanner $G' = (V, E \cup \{e\})$ such that $t' \leq (1+\varphi) \cdot t_{\mathcal{P}}$.*

6. Concluding remarks. We considered the problem of adding an edge 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 12?
3. A natural extension is to allow more than one edge to be added. Can we generalize our results to this case?

Acknowledgments. The authors would like to thank René van Oostrum for fruitful discussions during the early stages of this work, Mohammad Ali Abam for discussions about section 2.2, and Sergio Cabello for simplifying the algorithm in section 5. Finally, we thank the anonymous referees for many insightful comments and suggestions on how to improve the paper.

REFERENCES

- [1] I. ALTHÖFER, G. DAS, D. P. DOBKIN, D. JOSEPH, AND J. SOARES, *On sparse spanners of weighted graphs*, Discrete Comput. Geom., 9 (1993), pp. 81–100.
- [2] K. ALZOUBI, X.-Y. LI, Y. WANG, P.-J. WAN, AND O. FRIEDER, *Geometric spanners for wireless ad hoc networks*, IEEE Trans. Parallel Distrib. Syst., 14 (2003), pp. 408–421.
- [3] S. ARYA, G. DAS, D. M. MOUNT, J. S. SALOWE, AND M. SMID, *Euclidean spanners: Short, thin, and lanky*, in Proceedings of the 27th Annual ACM Symposium on Theory of Computing, 1995, pp. 489–498.
- [4] S. ARYA, D. M. MOUNT, AND M. SMID, *Randomized and deterministic algorithms for geometric spanners of small diameter*, in Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science, 1994, pp. 703–712.
- [5] P. BOSE, J. GUDMUNDSSON, AND P. MORIN, *Ordered theta graphs*, Comput. Geom., 28 (2004), pp. 11–18.
- [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*, J. ACM, 42 (1995), pp. 67–90.
- [7] A. CZUMAJ AND A. LINGAS, *Fast approximation schemes for Euclidean multi-connectivity problems*, in Proceedings of the 27th International Colloquium on Automata, Languages and Programming, Lecture Notes in Comput. Sci. 1853, Springer-Verlag, Berlin, 2000, pp. 856–868.
- [8] G. DAS, P. HEFFERNAN, AND G. NARASIMHAN, *Optimally sparse spanners in 3-dimensional Euclidean space*, in Proceedings of the Ninth Annual ACM Symposium on Computational Geometry, 1993, pp. 53–62.
- [9] G. DAS AND G. NARASIMHAN, *A fast algorithm for constructing sparse Euclidean spanners*, Internat. J. Comput. Geom. Appl., 7 (1997), pp. 297–315.
- [10] G. DAS, G. NARASIMHAN, AND J. SALOWE, *A new way to weigh malnourished Euclidean graphs*, in Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, 1995, pp. 215–222.
- [11] M. ELKIN AND D. PELEG, *$(1+\epsilon, \beta)$ -spanner constructions for general graphs*, SIAM J. Comput., 33 (2004), pp. 608–631.
- [12] D. EPPSTEIN, *Spanning trees and spanners*, in Handbook of Computational Geometry, J.-R. Sack and J. Urrutia, eds., North-Holland, Amsterdam, 2000, pp. 425–461.
- [13] D. EPPSTEIN AND K. WORTMAN, *Minimum dilation stars*, Comput. Geom., 37 (2007), pp. 27–37.
- [14] A. M. FARLEY, A. PROSKUROWSKI, D. ZAPPALA, AND K. J. WINDISCH, *Spanners and message distribution in networks*, Discrete Appl. Math., 137 (2004), pp. 159–171.
- [15] A. GRÜNE, *Tightness of Upper Bound on t_B/t_P* , manuscript, 2005.
- [16] J. GUDMUNDSSON AND C. KNAUER, *Dilation and detour in geometric networks*, in Handbook on Approximation Algorithms and Metaheuristics, T. Gonzales, ed., Chapman & Hall/CRC, Boca Raton, FL, 2007, pp. 52-1–52-17.

- [17] J. GUDMUNDSSON, C. LEVCOPOULOS, AND G. NARASIMHAN, *Fast greedy algorithms for constructing sparse geometric spanners*, SIAM J. Comput., 31 (2002), pp. 1479–1500.
- [18] J. GUDMUNDSSON, C. LEVCOPOULOS, G. NARASIMHAN, AND M. SMID, *Approximate distance oracles revisited*, in Proceedings of the 13th International Symposium on Algorithms and Computation, Lecture Notes in Comput. Sci. 2518, Springer-Verlag, Berlin, 2002, pp. 357–368.
- [19] J. GUDMUNDSSON, G. NARASIMHAN, AND M. SMID, *Fast pruning of geometric spanners*, in Proceedings of the 22nd International Symposium on Theoretical Aspects of Computer Science, Lecture Notes in Comput. Sci. 3404, Springer-Verlag, Berlin, 2005, pp. 508–520.
- [20] J. M. KEIL, *Approximating the complete Euclidean graph*, in Proceedings of the 1st Scandinavian Workshop on Algorithmic Theory, Lecture Notes in Comput. Sci. 318, Springer-Verlag, Berlin, 1988, pp. 208–213.
- [21] J. M. KEIL AND C. A. GUTWIN, *Classes of graphs which approximate the complete Euclidean graph*, Discrete Comput. Geom., 7 (1992), pp. 13–28.
- [22] S. LANGERMAN, P. MORIN, AND M. A. SOSS, *Computing the maximum detour and spanning ratio of planar paths, trees, and cycles*, in Proceedings of the 19th International Symposium on Theoretical Aspects of Computer Science, Lecture Notes in Comput. Sci. 2285, Springer-Verlag, Berlin, 2002, pp. 250–261.
- [23] 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 (1992), pp. 251–256.
- [24] C. LEVCOPOULOS, G. NARASIMHAN, AND M. SMID, *Improved algorithms for constructing fault-tolerant spanners*, Algorithmica, 32 (2002), pp. 144–156.
- [25] X.-Y. LI, *Applications of computational geometry in wireless ad hoc networks*, in Ad Hoc Wireless Networking, X.-Z. Cheng, X. Huang, and D.-Z. Du, eds., Kluwer Academic Publishers, Dordrecht, The Netherlands, 2003, pp. 197–264.
- [26] G. NARASIMHAN, *Geometric Spanner Networks: Open Problems*, Invited talk at the 1st Utrecht-Carleton Workshop on Computational Geometry, 2002.
- [27] G. NARASIMHAN AND M. SMID, *Approximating the stretch factor of Euclidean graphs*, SIAM J. Comput., 30 (2000), pp. 978–989.
- [28] G. NARASIMHAN AND M. SMID, *Geometric Spanner Networks*, Cambridge University Press, Cambridge, UK, 2007.
- [29] G. NAVARRO AND R. PAREDES, *Practical construction of metric t -spanners*, in Proceedings of the Fifth Workshop on Algorithm Engineering and Experiments, SIAM, Philadelphia, 2003, pp. 69–81.
- [30] G. NAVARRO, R. PAREDES, AND E. CHÀVEZ, *t -spanners as a data structure for metric space searching*, in Proceeding of the 9th International Symposium on String Processing and Information Retrieval, Lecture Notes in Comput. Sci. 2476, Springer-Verlag, Berlin, 2002, pp. 298–309.
- [31] D. PELEG, *Distributed Computing: A Locality-Sensitive Approach*, SIAM, Philadelphia, 2000.
- [32] S. RAO AND W. D. SMITH, *Approximating geometrical graphs via “spanners” and “banyans,”* in Proceedings of the 30th Annual ACM Symposium on Theory of Computing, 1998, pp. 540–550.
- [33] J. S. SALOWE, *Constructing multidimensional spanner graphs*, Internat. J. Comput. Geom. Appl., 1 (1991), pp. 99–107.
- [34] M. SMID, *Closest-point problems in computational geometry*, in Handbook of Computational Geometry, J.-R. Sack and J. Urrutia, eds., North-Holland, Amsterdam, 2000, pp. 877–935.
- [35] M. THORUP AND U. ZWICK, *Approximate distance oracles*, in Proceedings of the 33rd Annual ACM Symposium on Theory of Computing, 2001, pp. 183–192.
- [36] P. M. VAIDYA, *A sparse graph almost as good as the complete graph on points in K dimensions*, Discrete Comput. Geom., 6 (1991), pp. 369–381.