Fast approximation schemes for Euclidean multi-connectivity problems

Czumaj, Artur; Lingas, Andrzej

Published in:
Automata, languages and programming / Lecture notes in computer science

2000

Fast Approximation Schemes for Euclidean Multi-connectivity Problems *

(Extended Abstract)

Artur Czumaj[1**] and Andrzej Lingas[2]

[1] Department of Computer and Information Science, New Jersey Institute of Technology, Newark, NJ 07102-1982, USA. czumaj@cis.njit.edu

[2] Department of Computer Science, Lund University, Box 118, S-22100 Lund, Sweden. Andrzej.Lingas@cs.lth.se

Abstract. We present new polynomial-time approximation schemes (PTAS) for several basic minimum-cost multi-connectivity problems in geometrical graphs. We focus on low connectivity requirements. Each of our schemes either significantly improves the previously known upper time-bound or is the first PTAS for the considered problem.

We provide a randomized approximation scheme for finding a biconnected graph spanning a set of points in a multi-dimensional Euclidean space and having the expected total cost within $(1+\varepsilon)$ of the optimum. For any constant dimension and ε, our scheme runs in time $\mathcal{O}(n \log n)$. It can be turned into Las Vegas one without affecting its asymptotic time complexity, and also efficiently derandomized. The only previously known truly polynomial-time approximation (randomized) scheme for this problem runs in expected time $n \cdot (\log n)^{\mathcal{O}((\log\log n)^9)}$ in the simplest planar case. The efficiency of our scheme relies on transformations of nearly optimal low cost special spanners into sub-multigraphs having good decomposition and approximation properties and a simple subgraph connectivity characterization. By using merely the spanner transformations, we obtain a very fast polynomial-time approximation scheme for finding a minimum-cost k-edge connected multigraph spanning a set of points in a multi-dimensional Euclidean space. For any constant dimension, ε, and k, this PTAS runs in time $\mathcal{O}(n \log n)$.

Furthermore, by showing a low-cost transformation of a k-edge connected graph maintaining the k-edge connectivity and developing novel decomposition properties, we derive a PTAS for Euclidean minimum-cost k-edge connectivity. It is substantially faster than that previously known.

Finally, by extending our techniques, we obtain the first PTAS for the problem of Euclidean minimum-cost Steiner biconnectivity. This scheme runs in time $\mathcal{O}(n \log n)$ for any constant dimension and ε. As a byproduct, we get the first known non-trivial upper bound on the number of Steiner points in an optimal solution to an instance of Euclidean minimum-cost Steiner biconnectivity.

* Research supported in part by the DFG Sonderforschungsbereich 376, by DFG grant Me872/7-1, and by TFR grant 96-278. Full version of the paper is available from the authors.

** Work partly done while the author was with the Heinz Nixdorf Institute and Dept. of Mathematics & Computer Science at the University of Paderborn, D-33095 Paderborn, Germany.

U. Montanari et al. (Eds.): ICALP 2000, LNCS 1853, pp. 856–868, 2000.

1 Introduction

Multi-connectivity graph problems are central in algorithmic graph theory and have numerous applications in computer science and operation research [2,10,22]. They are also very important in the design of networks that arise in practical situations [2,10]. Typical application areas include telecommunication, computer and road networks. Low degree connectivity problems for geometrical graphs in the plane can often closely *approximate* such practical connectivity problems (see, e.g., the discussion in [10,22]).

In this paper, we provide a thorough *theoretical study* of these problems in Euclidean space (i.e., for geometrical graphs). We consider several basic connectivity problems of the following form: for a given set S of n points in the Euclidean space $\mathbb{R}^d$, find a minimum-cost subgraph of a complete graph on S that satisfies a priori given connectivity requirements. The cost of such a subgraph is equal to the sum of the Euclidean distances between adjacent vertices.

The most classical problem we investigate is the *(Euclidean) minimum-cost k-vertex connected spanning subgraph problem*. We are given a set S of n points in the Euclidean space $\mathbb{R}^d$ and the aim is to find a minimum-cost k-vertex connected graph spanning points in S (i.e., a subgraph of the complete graph on S). By substituting the requirement of k-edge connectivity for that of k-vertex connectivity, we obtain the corresponding *(Euclidean) minimum-cost k-edge connected spanning subgraph problem*. We term the generalization of the latter problem which allows for parallel edges in the output graph spanning S as the *(Euclidean) minimum-cost k-edge connected spanning sub-multigraph problem*.

The concept of minimum-cost k-connectivity naturally extends to include that of *Euclidean Steiner k-connectivity* by allowing the use of additional vertices, called *Steiner points*. The problem of *(Euclidean) minimum-cost Steiner k-vertex- (or, k-edge-) connectivity* is to find a minimum-cost graph on a *superset* of the input point set S in $\mathbb{R}^d$ which is k-vertex- (or, k-edge-) connected with respect to S. For $k = 1$, it is simply the famous *Steiner minimal tree* (SMT) problem, which has been very extensively studied in the literature (see, e.g., [11,16]).

Since all the aforementioned problems are known to be $\mathcal{NP}$-hard when restricted to even two-dimensions for $k \geq 2$ [8,18], we focus on efficient constructions of good approximations. We aim at developing a *polynomial-time approximation scheme*, a *PTAS*. This is a family of algorithms $\{\mathcal{A}_\varepsilon\}$ such that, for each fixed $\varepsilon > 0$, $\mathcal{A}_\varepsilon$ runs in time polynomial in the size of the input and produces a $(1 + \varepsilon)$-approximation [15].

Previous work. Despite the practical relevance of the multi-connectivity problems for geometrical graphs and the vast amount of practical heuristic results reported (see, e.g., [9,10,22,23]) very little theoretical research has been done towards developing efficient approximation algorithms for these problems. This contrasts with the very rich and successful theoretical investigations of the corresponding problems in general metric spaces and for general weighted graphs (see, e.g., [10,12,15,17]). Even for the simplest (and most fundamental) problem considered in our paper, that of finding a minimum-cost biconnected graph spanning a given set of points in the Euclidean plane, for a long time obtaining approximations achieving better than a $\frac{3}{2}$ ratio had been elusive and only very recently has a PTAS been developed [6]. For any fixed $\varepsilon > 0$, this algorithm outputs a $(1+\varepsilon)$-approximation in expected time $n\,(\log n)^{\mathcal{O}((\log\log n)^9)}$. The approximation scheme developed in [6] can be extended to arbitrary k and d, but in the general case the dependence on k and particularly on d makes the algorithm impractical. For an $\varepsilon > 0$, the algorithm runs in expected time $n \cdot 2^{(\log\log n)^{2\,d+\binom{2\,d}{2}} \cdot ((\mathcal{O}(d\,k^2/\varepsilon))^d \; \log(\varepsilon^{-1}))!}$. Note that d plays a large role in the running time of these schemes. In fact, the result from [6] implies that for every $d = \Omega(\log n)$, even for $k = 2$ no PTAS exists unless $\mathcal{P} = \mathcal{NP}$. Thus, the problem of finding a minimum-cost biconnected spanning subgraph does not have a PTAS (unless $\mathcal{P} = \mathcal{NP}$) even in the metric case. Hence, our restriction to Euclidean graphs in low dimensions plays an essential role in these schemes.

A related, but significantly weaker result has been also presented in [5]. Here an optimal solution to the problem without allowing Steiner points is approximated to an arbitrarily close degree via the inclusion of Steiner points.

When Steiner points are allowed in the minimum-cost Steiner k-vertex- (or k-edge-) connectivity problem, the only non-trivial results are known for $k = 1$, i.e., for the minimum Steiner tree problem (SMT). In the breakthrough paper [3], Arora designed a PTAS for SMT for all constants d. Mitchell independently obtained a similar result for $d = 2$ [19]. Soon after Rao and Smith [20] offered a significantly faster PTAS for SMT running in time $\mathcal{O}(n\,\log n)$ for a constant d. For $k \geq 2$, the only result we are aware of is a $\sqrt{3}$-approximation in polynomial-time for $k = 2$ [14].

New results. In this paper we present new polynomial-time approximation schemes for several of the aforementioned connectivity problems in geometric graphs. We focus on low connectivity requirements. Each of our approximation schemes either significantly improves the previously known upper time-bound or is the first PTAS for the considered problem.

Our *main new result* is a fast polynomial-time (randomized) approximation scheme for finding a biconnected graph spanning a set of points in a d-dimensional Euclidean space and having expected cost within $(1+\varepsilon)$ of optimum. For any constant d and ε, our algorithm runs in expected time $\mathcal{O}(n\,\log n)$. Our scheme is a PTAS for the problem in $\mathbb{R}^d$ for all d such that $2^{d^{c\,d^2}} = \text{poly}(n)$, for some absolute constant c. We can turn our randomized scheme into a Las Vegas one without affecting its asymptotic time complexity. With a very slight increase of the running time (a constant factor provided that d and ε are constant) we can also obtain a deterministic $(1+\varepsilon)$-approximation. Our scheme is significantly, i.e., by a factor of at least $(\log n)^{\mathcal{O}((\log\log n)^9)}$, faster than that from [6].

Since a minimum-cost biconnected graph spanning a set of points in a metric space is also a minimum-cost two-edge connected graph spanning this set, our PTAS yields also the corresponding PTAS for the Euclidean minimum-cost two-edge connectivity.

We extend the techniques developed for the biconnectivity algorithms and present a fast randomized PTAS for finding a minimum cost k-edge connected *multigraph* spanning a set of points in a d-dimensional Euclidean space. The running time of our Las Vegas scheme is $\mathcal{O}(n\,\log n) + n\,2^{k^{\mathcal{O}(k)}}$ for any constant d and ε.

We are also able to improve upon the k-edge connectivity results from [6] significantly. By showing a low-cost transformation of a k-edge connected graph maintaining the k-edge connectivity and developing novel decomposition properties, we derive a

PTAS for Euclidean minimum-cost k-edge connectivity which for any constant d, ε, and k, runs in expected time $n\,(\log n)^{O(1)}$. The corresponding scheme in [6] requires $n\,(\log n)^{\mathcal{O}((\log\log n)^9)}$ time when $d = 2$ and for any constant $\varepsilon,\ k$.

Furthermore, we present a series of new structural results about minimum-cost biconnected Euclidean Steiner graphs, e.g., a decomposition of a minimum-cost biconnected Steiner graph into minimal Steiner trees. We use these results to derive the *first* PTAS for the minimum-cost Steiner biconnectivity and Steiner two-edge connectivity problems. For any constant d and ε, our scheme runs in expected time $\mathcal{O}(n\ \log n)$. As a byproduct of the aforementioned decomposition, we also obtain the first known non-trivial upper bound on the minimum number of Steiner points in an optimal solution to an n-point instance of Euclidean minimum-cost Steiner biconnectivity, which is $3n - 2$.

Techniques. The only two known PTAS approaches to Euclidean minimum-cost k-vertex- (or, k-edge-) connectivity (see [5,6]) are based on decompositions of k-connected Euclidean graphs combined with the general framework proposed recently by Arora [3] for designing PTAS for Euclidean versions of TSP, Minimum Steiner Tree, Min-Cost Perfect Matching, k-TSP, etc. (For another related framework for geometric PTAS see [19].) In contrast to all previous applications of Arora's framework using Steiner points in the so-called patching procedures [3,5], a patching method free of Steiner points is given [6]. (Steiner points of degree at least three are difficult to remove for k-connectivity when $k \geq 2$. This should be compared to the problems considered by Arora [3] and Rao and Smith [20] where the output graphs have very simple connectivity structure.) This disallowance in [6] makes it hard to prove strong global structural properties of close approximations with respect to a given geometric partition.

Structural theorems in Arora's framework typically assert the existence of a recursive partition of a box containing the n input points (perturbed to nearest grid points) into cubes such that the optimal solution can be closely approximated by a so called (m, r)-light solution in which, for every cube, there are very few edges crossing its boundaries. The structural theorem in [6] yields only weaker structural properties of approximate solutions ((m, r)-grayness and (m, r)-blueness) which bound solely the number of crossings between the cube boundaries and the edges having exactly one endpoint within the cube. That bound is constant for "short edges" (i.e., edges having length within a constant factor of the side-length of the cube) and it is $\mathcal{O}(\log\log n)$ for "long edges" (assuming that k, d, and ε are constant). Furthermore, most of the crossings are located in one of $2\,d$ prespecified points. The weaker structural properties (especially the fact that there might be as many as $\Theta(\log\log n)$ edges having exactly one endpoint in a cube in the partition) lead to the high time complexity of the main dynamic programming procedure in the PTAS presented in [6].

We take a novel approach in order to guarantee stronger structural properties of approximate solutions disallowing Steiner points. Our approach is partly inspired by the recent use of spanners to speed-up PTAS for Euclidean versions of TSP by Rao and Smith [20]. In effect, for $k = 2$, we are able to prove a substantially stronger structural property (r-local-lightness, see Theorem 3.1) than that in [6]. It yields a constant upper bound on the number of long edges with exactly one endpoint in a cube in the partition provided that d and ε are constant.

Our proof relies on a series of transformations of a $(1+\delta)$-spanner for the input point set, having low cost and the so called *isolation property* [1], into an r-locally-light k-edge connected multigraph spanning the input set and having nearly optimal cost. Without any increase in the cost, in case $k = 2$, the aforementioned multigraph is efficiently transformed into a biconnected graph spanning the input point set. Furthermore, for the purpose of dynamic programming, we succeed to use a more efficient subgraph connectivity characterization in case $k = 2$ than that used in [5,6].

By using merely the aforementioned spanner transformations, we also obtain the fast randomized PTAS for finding a minimum-cost k-edge connected *multigraph* spanning a set of points in a multi-dimensional Euclidean space.

It seems unlikely that a cost-efficient transformation of a k-edge connected multigraph into a k-edge connected graph on the same point set exists. For this reason, in case of k-edge connectivity, we consider an arbitrary k-edge connected graph on the input point set instead of a spanner, and derive a series of cost-efficient transformations of the former into an r-locally-light k-edge connected graph on the input set. The transformations yield the fastest known randomized PTAS for Euclidean minimum-cost k-edge connectivity.

Our investigations of spanners with the isolation property, the *explicit* use of multigraphs instead of graphs, and the proof that nearly optimal, low cost spanners possessing the isolation property induce r-locally-light sub-multigraphs having good approximation properties are the main sources of the efficiency of the approximation schemes for Euclidean minimum-cost connectivity problems (without Steiner points) presented in this paper.

By extending the aforementioned techniques to include Steiner points, deriving the decomposition of a minimum-cost biconnected Steiner graph into minimal Steiner trees, and using the generalization of $(1+\varepsilon)$-spanner to include Steiner points called $(1+\varepsilon)$-*banyans* in [20,21], we obtain the first PTAS for minimum-cost Steiner biconnectivity and Steiner two-edge connectivity.

Organization of the paper. Section 2 provides basic terminology used in our approximation schemes. In Section 3 we outline our new PTAS for Euclidean minimum-cost biconnectivity. Section 4 sketches the PTAS for Euclidean minimum-cost k-edge connectivity in multigraphs. In Section 5 we derive our PTAS for Euclidean minimum-cost k-edge connectivity in graphs. Section 6 presents the PTAS for minimum-cost Steiner biconnectivity and Steiner two-edge connectivity. Due to space limitations most of our technical claims and their proofs are postponed to the full version of the paper.

2 Definitions

We consider geometrical graphs. A *geometrical (multi-)graph* on a set of points S in $\mathbb{R}^d$ is a weighted (multi-)graph whose set of vertices is exactly S and for which the *cost* of an edge is the Euclidean distance between its endpoints. The (total) cost of the (multi-)graph is the sum of the costs of its edges. A (multi-)graph G on S *spans* S if it is connected (i.e., there is a path in G connecting any two points in S). As in [3,5,6,20], we allow edges to deviate from straight-line segments and specify them as straight-lines paths (i.e., paths consisting of straight-line segments) connecting the endpoints. This relaxation enables the edges to pass through some prespecified points (called *portals*) where they may be "*bent*." When all edges are straight-line segments, G is called a *straight-line graph*. For a multigraph G, the *graph induced by* G is the graph obtained by reducing the multiplicity of each edge of G to one.

We shall denote the cost of the minimum spanning tree on a point set X by $\ell(\mathrm{MST}(X))$. A t-*spanner* of a set of points S in $\mathbb{R}^d$ is a subgraph of the complete straight-line graph on S such that for any two points $x, y \in S$ the length of the shortest path from x to y in the spanner is at most t times the Euclidean distance between x and y [1].

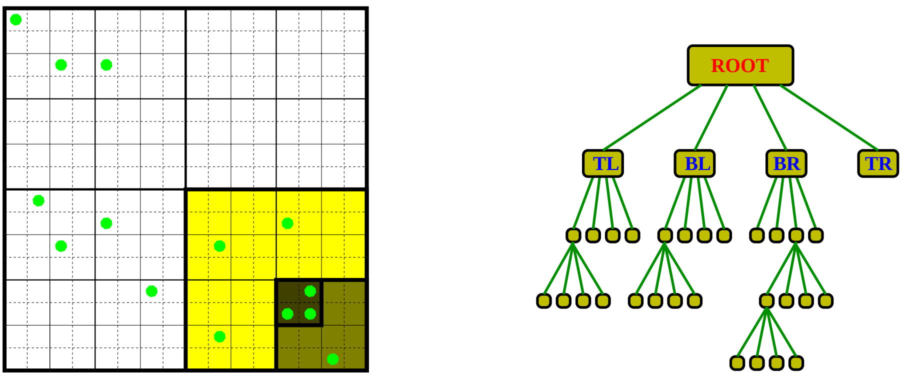

Fig. 1. Dissection of a bounding cube in $\mathbb{R}^2$ (left) and the corresponding 2^2-ary tree (right). In the tree, the children of each node are ordered from left to right: **T**op/**L**eft square, **B**ottom/**L**eft square, **B**ottom/**R**ight square, and **T**op/**R**ight square.

We hierarchically partition the space as in [3]. A *bounding box* of a set S of points in $\mathbb{R}^d$ is a smallest d-dimensional axis-parallel cube containing the points in S. A *(*2^d*-ary) dissection* [3] (see Figure 1) of a set of points in a cube L^d in $\mathbb{R}^d$ is the recursive partitioning of the cube into smaller sub-cubes, called *regions*. Each *region* U^d of volume > 1 is recursively partitioned into 2^d regions $(U/2)^d$. A 2^d*-ary tree* (for a given 2^d-ary dissection) is a tree whose root corresponds to L^d, and whose other non-leaf nodes correspond to the regions containing at least two points from the input set (see Figure 1). For a non-leaf node v of the tree, the nodes corresponding to the 2^d regions partitioning the region corresponding to v, are the children of v in the tree.

For any d-vector $\mathbf{a} = (a_1, \ldots, a_d)$, where all a_i are integers $0 \le a_i \le L$, the $\mathbf{a}$-*shifted dissection* [3,6] of a set X of points in the cube L^d in $\mathbb{R}^d$ is the dissection of the set X^* in the cube $(2\,L)^d$ in $\mathbb{R}^d$ obtained from X by transforming each point $\mathbf{x} \in X$ to $\mathbf{x} + \mathbf{a}$. A *random shifted dissection* of a set of points X in a cube L^d in $\mathbb{R}^d$ is an $\mathbf{a}$-shifted dissection of X with $\mathbf{a} = (a_1, \ldots, a_d)$ and the elements $a_1, \ldots, a_d$ chosen independently and uniformly at random from $\{0, 1, \ldots, L\}$.

A crossing of an edge with a region facet of side-length W in a dissection is called *relevant* if it has exactly one endpoint in the region and its length is at most $2\sqrt{d}\,W$. A graph is r-*gray* with respect to a shifted dissection if each facet of each region in the dissection has at most r relevant crossings. A graph is *r-locally-light* with respect to a shifted dissection if for each region in the dissection there are at most r edges having *exactly one endpoint in the region*. A graph is *r-light* [3,20,6] with respect to a shifted dissection if for each region in the dissection there are at most r edges *crossing any of its facets*. (It is important to understand the difference between these two latter notions.)

An *m-regular* set of *portals* in a $(d-1)$-dimensional region facet N^{d-1} is an orthogonal lattice of m points in the facet where the spacing between the portals is $(N+1) \cdot m^{-1/(d-1)}$ (cf. [3]). If a graph is r-locally-light and for each facet in any region every edge crosses through one of the m portals in the facet then the graph is called *(m, r)-locally-light*.

3 Algorithm for Euclidean Biconnectivity

In this section we *sketch* a randomized algorithm that finds a biconnected graph spanning a set S of n points in $\mathbb{R}^d$ whose cost is at most $(1+\varepsilon)$ times of the minimum. We specify here only a key lemma and the structural theorem and defer the detailed description of the algorithm and its analysis to the full version of the paper.

Our algorithm starts by finding a smallest bounding box for the input point set S, rescaling the input coordinates so the bounding box is of the form $[0, \mathcal{O}(n\sqrt{d}\,(1 + 1/\varepsilon))]^d$, and moving the points in S to the closest unit grid points. We shall term the perturbed point set as *well-rounded* (see also [3,6]). Next, the algorithm finds an appropriate $(1+\Theta(\varepsilon))$-spanner of the well-rounded point set which has the so-called *(κ, c)-isolation property* for appropriate parameters κ and c [1].

Definition 3.1. *Let c, $0 < c < 1$, be a constant and let $\kappa \geq 1$. A geometrical graph G on a set of points in $\mathbb{R}^d$ satisfies the (κ, c)-isolation property, if, for each edge e of G of length l, there is a cylinder $\mathcal{C}$ of length and radius $c\,l$ and with its axis included in e such that $\mathcal{C}$ intersects at most κ edges of G.*

In the next step, the algorithm chooses a random shifted dissection and builds the corresponding shifted 2^d-ary tree for the perturbed S.

The following key lemma yields a low upper bound on the number of crossings of the boundaries of a region in the dissection by long edges of the spanner having exactly one endpoint within the region.

Lemma 3.1. *Let G be a geometrical graph spanning a set S of points in $\mathbb{R}^d$ and satisfying the (κ, c)-isolation property, where $0 < c < 1$ is a constant. There exists a constant c' such that for any region in a shifted dissection of S the number of edges of G of length at least $c'\sqrt{d}$ times the side-length of the region that have precisely one endpoint within the region is $\kappa \cdot (d/\varepsilon)^{\mathcal{O}(d)}$.*

Combining this lemma, with several lemmas that describe graph transformations reducing the number of crossings of the boundaries of a region in the dissection by short edges, we obtain the following structural theorem.

Theorem 3.1. *Let ε, λ be any positive reals and let k be any positive integer. Next, let $\mathfrak{S}$ be a $(1+\varepsilon)$-spanner for a well-rounded set S of points in $\mathbb{R}^d$ that satisfies the (κ, c)-isolation property with constant c, $0 < c < 1$, $\kappa = (d/\varepsilon)^{\mathcal{O}(d)}$, and has $n \cdot (d/\varepsilon)^{\mathcal{O}(d)}$*

edges whose total cost equals $L_{\mathfrak{G}}$. Choose a shifted dissection uniformly at random. Then one can modify $\mathfrak{G}$ to a graph $\mathcal{G}$ spanning S such that

- *$\mathcal{G}$ is r-locally-light with respect to the shifted dissection chosen, where $r = 2^{d+1} \cdot d \cdot k^2 + d \cdot (\mathcal{O}(\lambda \cdot d^{3/2}))^d + (d/\varepsilon)^{\mathcal{O}(d)}$, and*
- *there exists a k-edge connected multigraph H which is a spanning subgraph of $\mathcal{G}$ with possible parallel edges (of multiplicity at most k), whose expected (over the choice of the shifted dissection) cost is at most $(1 + \varepsilon + \frac{k \cdot L_{\mathfrak{G}}}{\lambda \cdot \ell(\mathrm{MST})})$ times cost of the minimum-cost of k-edge connected multigraph spanning S.*

Furthermore, this modification can be performed in time $\mathcal{O}(d \cdot L \cdot \log L) + n \cdot 2^{d^{\mathcal{O}(d)}} + n \cdot (d/\varepsilon)^{\mathcal{O}(d)}$, where L is the side-length of the smallest bounding box containing S.

Further, our algorithm modifies the spanner according to Theorem 3.1 producing an r-locally-light graph G where r is constant for constant ϵ and d. In the consecutive step, the algorithm runs a dynamic programming subroutine for finding a minimum-cost two-edge connected multigraph for which the induced graph is a subgraph of G (note that by Theorem 3.1 the multigraph has expected cost very close to that of the minimum-cost of a k-edge connected multigraph spanning the perturbed S). The efficiency of the subroutine relies on a new, forest-like characterization of the so called connectivity type of a multigraph within a region of the dissection. It is substantially more concise than the corresponding one used in [5,6]. Next, the algorithm transforms the multigraph to a biconnected (straight-line) graph without any increase in cost. Finally, it modifies the biconnected graph to a biconnected graph on the input set by re-perturbing its vertices.

Theorem 3.2. *The algorithm finds a biconnected graph spanning the input set of n points in $\mathbb{R}^d$ and having expected cost within $(1 + \varepsilon)$ from the optimum. The running time of the algorithm is $\mathcal{O}(n \cdot d^{3/2} \cdot \varepsilon^{-1} \cdot \log(n\, d/\varepsilon)) + n \cdot 2^{(d/\varepsilon)^{\mathcal{O}(d^2)}}$. In particular, when d and ε are constant, then the running time is $\mathcal{O}(n \log n)$. For a constant d and arbitrary $s = \frac{1}{\varepsilon} \geq 1$ the running time is $\mathcal{O}(n\, s \log(n\, s) + n\, 2^{s^{\mathcal{O}(1)}})$. The algorithm can be turned into a Las Vegas one without affecting the stated asymptotic time bounds.*

Although we have used many ideas from [6] in the design of our algorithm, we have chosen the method of picking a random shifted dissection given in [3,20,21]. Therefore we can apply almost the same arguments as those used by Rao and Smith [21, Sections 2.2 and 2.3] to derandomize our algorithm at small increase of the cost

Theorem 3.3. *For every positive ε there exists a deterministic algorithm running in time $n\,(d/\varepsilon)^{\mathcal{O}(1)} \log n + n\, 2^{(d/\varepsilon)^{\mathcal{O}(d^2)}}$ that for every set of n points in $\mathbb{R}^d$ produces a biconnected graph spanning the points and having the cost within $(1+\varepsilon)$ of the minimum. In particular, when d and ε are constant, the running time is $\mathcal{O}(n \log n)$. For a constant d and arbitrary $s = \frac{1}{\varepsilon} \geq 1$ the running time is $\mathcal{O}(n\, s^{\mathcal{O}(1)} \log n + n\, 2^{s^{\mathcal{O}(1)}})$.*

4 Euclidean k-Edge Connectivity in Multigraphs

We can extend the techniques developed in the previous sections to the problem of finding a low-cost k-edge connected *multigraph* spanning a set of points in $\mathbb{R}^d$ for $k \geq 2$. To begin with, we follow the PTAS from Section 3 up and inclusive the spanner-modification step, only changing some parameters. The resulting graph G is r-locally-light graph for $r = k^d \cdot (d/\varepsilon)^{\mathcal{O}(d^2)}$. By Theorem 3.1, there exists a k-edge connected multigraph H such that the induced graph is a subgraph of G and the expected cost of H is at most $(1 + \frac{\varepsilon}{4})$ times larger than the minimum-cost k-edge connected multigraph spanning S. As was the case for the PTAS from Section 3, we can apply dynamic programming to find a minimum-cost k-edge connected multigraph H^* for which the induced graph is a subgraph of G. This time we use a more general (but less efficient) connectivity characterization from [5] yielding $2^{\mathcal{O}(r!)}$ different connectivity types. In effect, the dynamic programming is more expensive, and the total running time is $n \cdot 2^{\mathcal{O}(r!)}$. Now, it is sufficient to re-perturb the vertices of the multigraph H^* (correspondingly to the last step of the PTAS from Section 3) in order to obtain the following theorem.

Theorem 4.1. *Let k be an arbitrary positive integer. There exists an algorithm that finds a k-edge connected multigraph spanning the input set of n points in $\mathbb{R}^d$ and having expected cost within $(1+\varepsilon)$ from the optimum. The running time of the algorithm is $\mathcal{O}(n \cdot d^{3/2} \cdot \varepsilon^{-1} \cdot \log(n\, d/\varepsilon)) + n \cdot 2^{\mathcal{O}((k^d \cdot (d/\varepsilon)^{\mathcal{O}(d^2)})!)}$. In particular, when d and ε are constant, then the running time is $\mathcal{O}(n \log n) + n\, 2^{k^{\mathcal{O}(k)}}$. For a constant d and an arbitrary $s = \frac{1}{\varepsilon} \geq 1$ the running time is $\mathcal{O}(n\, s \log(n\, s)) + n\, 2^{(k\, s)^{\mathcal{O}(k\, s)}}$. The algorithm can be turned into a Las Vegas one without affecting the stated asymptotic time bounds.*

Recall the use of Steiner points in the first attempt of deriving PTAS for the Euclidean minimum-cost k-connectivity in [5] by allowing them solely on the approximation side. As a byproduct, we can substantially subsume the results on minimum-cost k-connectivity from [5] in the complexity aspect by using Theorem 4.1.

5 Euclidean *k*-Edge Connectivity in Graphs

Our spanner approach to biconnectivity relies on an efficient transformation of a two-edge connected multigraph into a biconnected graph on the same point set without any cost increase. Unfortunately, it seems that for $k > 2$ there is no any similar cost-efficient transformation between k-edge connected *multigraphs* and k-vertex- or k-edge connected *graphs*. We show in this section that an arbitrary (in particular, minimum-cost) k-edge connected graph spanning a well-rounded point set admits a series of transformations resulting in an r-locally-light k-edge connected graph on this set with a small increase in cost. By some further small cost increase, we can make the latter graph (m, r)-locally-light in order to facilitate efficient dynamic programming.

The following lemma plays a crucial role in our analysis. Intuitively, it aims at providing a similar transformation of the input graph as that presented in Lemma 3.1. The main difficulty here is in the fact that the input graph may be arbitrary (while in Lemma 3.1 we have analyzed graphs having the isolation property).

Lemma 5.1. *Let G be a spanning graph of a well-rounded point set S in $\mathbb{R}^d$. For any shifted dissection of S, G can be transformed into a graph G' satisfying the following three conditions:*

– *if G' is different from G then the total cost of G' is smaller than that of G,*

– *for any region of size W in the dissection there are positive reals x_i, $i = 1, \ldots, 2^{\mathcal{O}(d)}$, greater than $4\sqrt{d}\,W$ such that the number of edges having theirs lengths outside any of the intervals $[x_i, 2\,x_i)$, and each having precisely one endpoint within the region is $2^{\mathcal{O}(d)}$,*
– *if G is k-edge connected then so is G'.*

Lemma 5.1 guarantees that in the graph resulting from the transformation provided in the lemma no region in the dissection is crossed by too many long edges having their length outside finite number of intervals of the form $[x, 2x)$, $x > 4\sqrt{d}\,W$ and one point inside the region. The following lemma reduces the number of the remaining edges.

Lemma 5.2. *Let G be an r-gray spanning graph of a well-rounded set S of points in $\mathbb{R}^d$. Let Q be a region of size W in the dissection of S, and let x_i, $i = 1, \ldots, 2^{\mathcal{O}(d)}$, be positive reals greater than $4\sqrt{d}\,W$. If there are $\mathfrak{L}$ edges having their lengths outside the intervals $[0, 2\sqrt{d}\,W]$, $[x_i, 2\,x_i)$, $i = 1, \ldots, 2^{\mathcal{O}(d)}$, and such that each has precisely one endpoint in Q, then there are at most $\mathfrak{L} + r\,2^{\mathcal{O}(d)}$ edges crossing the facets of Q and having one endpoint in Q.*

By these two lemmas, we obtain our structure theorem for k-edge connectivity.

Theorem 5.1. *Let $\varepsilon > 0$, and let S be a well-rounded set of n points in $\mathbb{R}^d$. A minimum-cost k-edge connected graph spanning S can be transformed to a k-edge connected graph H spanning S such that*
– *H is (m, r)-locally-light with respect to the shifted dissection, where $m = (\mathcal{O}(\varepsilon^{-1} \cdot \sqrt{d} \cdot \log n))^{d-1}$ and $r = (\mathcal{O}(k^2 \cdot d^{3/2} / \varepsilon))^d$, and*
– *the expected (over the choice of the shifted dissection) cost of H is at most $(1 + \varepsilon)$ times larger than that of the minimum-cost graph.*

The concept of (m, r)-local-lightness is a very simple case of that of (m, r)-blueness used in [6]. Therefore, we can use a simplified version of the dynamic programming method of [6] involving the k-connectivity characterization from [5] in order to find an (m, r)-locally-light k-edge connected graph on a well-rounded point set satisfying the requirements of Theorem 5.1. This yields a substantially faster PTAS for the Euclidean minimum-cost k-edge connectivity than that presented in [6].

Theorem 5.2. *Let k be an arbitrary positive integer and let $\varepsilon > 0$. There exists a randomized algorithm that finds a k-edge connected graph spanning the input set of n points in $\mathbb{R}^d$ and having expected cost within $(1 + \varepsilon)$ from the optimum. The expected running time of the algorithm is $n \cdot (\log n)^{(\mathcal{O}(k^2 \cdot d^{3/2} / \varepsilon))^d} \cdot 2^{\mathcal{O}(((k^2 \cdot d^{3/2} / \varepsilon)^d)!)}$. In particular, when k, d, and ε are constant, then the running time is $n\,(\log n)^{\mathcal{O}(1)}$.*

6 Euclidean Steiner Biconnectivity

In this section we provide the first PTAS for Euclidean minimum-cost Steiner biconnectivity and Euclidean minimum-cost two-edge connectivity. For any constant dimension and ε, our scheme runs in time $\mathcal{O}(n \log n)$. Our proof relies on a decomposition of a minimum-cost biconnected Steiner graph into minimal Steiner trees and the use of the so called $(1+\varepsilon)$-*banyans* introduced by Rao and Smith [20,21]. As a byproduct of the decomposition, we derive the first known non-trivial upper bound on the minimum number of Steiner points in an optimal solution to an n-point instance of Euclidean minimum-cost Steiner biconnectivity which is $3n-2$.

Since for any set of points X in $\mathbb{R}^d$ the minimum-cost of a biconnected graph spanning X is the same as the minimum-cost of a two-edge connected graph spanning X, in the remaining part of this section we shall focus only on the Euclidean minimum-cost Steiner biconnectivity problem. By a series of technical lemmas, we obtain the following characterization of any optimal graph solution to the Euclidean minimum-cost Steiner biconnectivity problem.

Theorem 6.1. *Let G be a minimum-cost Euclidean Steiner biconnected graph spanning a set S of $n \geq 4$ points in $\mathbb{R}^d$. Then G satisfies the following conditions:*
(i) Each vertex of G (inclusive Steiner points) is of degree either two or three.
(ii) By splitting each vertex v of G corresponding to an input point into $\deg(v)$ independent endpoints of the edges, graph G can be decomposed into a number of minimal Steiner trees.
(iii) G has at most $3n-2$ Steiner points.

6.1 PTAS

Our spanner-based method for Euclidean minimum-cost biconnectivity cannot be extended directly to include Euclidean minimum-cost Steiner biconnectivity since spanners do not include Steiner points. Nevertheless, the decomposition of an optimal Steiner solution into minimum Steiner trees given in Theorem 6.1 opens the possibility of using the aforementioned banyans to allow Steiner points for the purpose of approximating the Euclidean minimum Steiner tree problem in [20].

Definition 6.1. [20] *A $(1+\varepsilon)$-banyan of a set S of points in $\mathbb{R}^d$ is a geometrical graph on a superset of S (i.e., Steiner points are allowed) such that for any subset U of S, the cost of the shortest connected subgraph of the banyan which includes U, is at most $(1+\varepsilon)$ times larger than the minimum Steiner tree of U.*

Rao and Smith have proved the following useful result on banyans in [20].

Lemma 6.1. *Let $0 < \varepsilon < 1$. One can construct a $(1+\varepsilon)$-banyan of an n-point set in $\mathbb{R}^d$ which uses only $d^{\mathcal{O}(d^2)}\,(d/\varepsilon)^{\mathcal{O}(d)}\,n$ Steiner points and has cost within a factor of $d^{\mathcal{O}(d^2)}\,\varepsilon^{-\mathcal{O}(d)}$ of the minimum Steiner tree of the set. The running time of the construction is $d^{\mathcal{O}(d^2)}\,(d/\varepsilon)^{\mathcal{O}(d)}\,n + \mathcal{O}(d\,n\,\log n)$.*

By combining the definition of $(1+\varepsilon)$-banyan with Theorem 6.1 (2) and Lemma 6.1, we get the following lemma.

Lemma 6.2. *For a finite point set S in $\mathbb{R}^d$ let $\mathfrak{B}$ be a $(1+\varepsilon/4)$-banyan constructed according to Lemma 6.1. Let $\mathfrak{B}_m$ be the multigraph obtained from $\mathfrak{B}$ by doubling its edges. There is a two-edge connected sub-multigraph of $\mathfrak{B}_m$ which includes S and whose cost is within $(1+\varepsilon/4)$ of the minimum cost of two-edge connected multigraph on any superset of S.*

Let $\mathfrak{B}_m^*$ be the multigraph resulting from scaling and perturbing all the vertices (i.e., also the Steiner points) of the multigraph $\mathfrak{B}_m$ specified in Lemma 6.2 according to the first step in Section 3. The vertices of $\mathfrak{B}_m^*$ are on a unit grid $[0, L]^d$ and, by arguing analogously as in Section 3, a minimum-cost two-edge connected sub-multigraph of $\mathfrak{B}_m^*$ that includes S is within $(1 + \varepsilon/4)$ of a minimum-cost two-edge connected sub-multigraph of $\mathfrak{B}_m$ that includes S.

The patching method of [5] applied to $\mathfrak{B}_m^*$ yields the following structure theorem.

Theorem 6.2. *Choose a shifted dissection of the set of vertices of the banyan $\mathfrak{B}$ at random. Then $\mathfrak{B}_m^*$ can be modified to a multigraph $\mathfrak{B}'_m$ such that:*

- *$\mathfrak{B}'_m$ is r-light with respect to the shifted dissection, where $r = (\mathcal{O}(\sqrt{d}/\varepsilon))^{d-1}$,*
- *the set of vertices of $\mathfrak{B}'_m$ includes that of $\mathfrak{B}_m^*$ and some additional vertices placed at the crossings between the edges of $\mathfrak{B}_m^*$ and the boundaries of the regions in the shifted dissection,*
- *there exists a two-edge connected sub-multigraph of $\mathfrak{B}'_m$ including S whose expected cost is within $(1+\varepsilon/4)$ of the minimum-cost of two-edge connected sub-multigraph of $\mathfrak{B}_m^*$ that includes S.*

To find such a subgraph of $\mathfrak{B}'_m$ efficiently we apply a simplification of the dynamic programming method used by the PTAS from Section 3 to the set of vertices of $\mathfrak{B}_m^*$ (it would be even simpler to use a modification of the dynamic programming approach from [5]). In effect we can find $\mathfrak{B}'_m$ in expected time $\mathcal{O}(n \log n)$ for constant ε and d. By combining this with Lemma 6.2, Theorem 6.2, and the efficient transformation of two-edge connected multigraphs into biconnected graphs, we obtain the main result in this section.

Theorem 6.3. *There exists an approximation algorithm for the minimum-cost Steiner biconnectivity (and two-edge connectivity) which for any $\varepsilon > 0$ returns a Euclidean Steiner biconnected (or two-edge connected) graph spanning the input set of n points in $\mathbb{R}^d$ and having expected cost within $(1 + \varepsilon)$ from the optimum. The running time of the algorithm is $\mathcal{O}(n\, d^{3/2}\, \varepsilon^{-1} \log(n\, d/\varepsilon)) + d^{\mathcal{O}(d^2)}\, (d/\varepsilon)^{\mathcal{O}(d)}\, n + n\, 2^{(\sqrt{d}/\varepsilon)^{d\,(\mathcal{O}(\sqrt{d}/\varepsilon))^{d-1}}}$. In particular, when d and ε are constant, then the running time is $\mathcal{O}(n \log n)$. The algorithm can be turned into a Las Vegas one without affecting asymptotic time bounds.*

Acknowledgments

We thank unknown referees for their valuable comments and suggestions (because of space considerations not all of the suggestions could be implemented).

References

1. S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. Smid. Euclidean spanners: Short, thin, and lanky. *Proc. 27th ACM STOC*, pp. 489–498, 1995.
2. R. K. Ahuja, T. L. Magnanti, J. B. Orlin, and M. R. Reddy. Applications of network optimization. In Nemhauser, editors, *Handbooks in Operations Research and Management Science*, volume 7: *Network Models*, ch. 1, pp. 1–83. North-Holland, Amsterdam, 1995.

3. S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *J. Assoc. Comput. Mach.*, 45(5):753–782, 1998.
4. D. Bienstock, E. F. Brickell, and C. L. Monma. On the structure of minimum-weight k-connected spanning networks. *SIAM J. Disc. Math.*, 3(3):320–329, 1990.
5. A. Czumaj and A. Lingas. A polynomial time approximation scheme for Euclidean minimum cost k-connectivity. *Proc. 25th ICALP*, pp. 682–694, 1998.
6. A. Czumaj and A. Lingas. On approximability of the minimum-cost k-connected spanning subgraph problem. *Proc. 10th ACM-SIAM SODA*, pp. 281–290, 1999.
7. G. N. Frederickson and J. JáJá. On the relationship between the biconnectivity augmentation and Traveling Salesman Problem. *Theoret. Comput. Sci.*, 19(2):189–201, 1982.
8. M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of $\mathcal{NP}$-completeness*. Freeman, New York, NY, 1979.
9. M. Grötschel, C. L. Monma, and M. Stoer. Computational results with a cutting plane algorithm for designing communication networks with low-connectivity constraints. *Operations Research*, 40(2):309–330, 1992.
10. M. Grötschel, C. L. Monma, and M. Stoer. Design of survivable networks. *Handbooks in Operations Research and Management Science*, volume 7: *Network Models*, ch. 10, pp. 617–672. North-Holland, Amsterdam, 1995.
11. E. N. Gilbert and H. O. Pollak. Steiner minimal trees. *SIAM J. Appl. Math.*, 16(1):1–29, 1968.
12. M. X. Goemans and D. P. Williamson. The primal-dual method for approximation algorithms and its application to network design problems. Chapter 4 in [15], pp. 144–191.
13. J. Gudmundsson, C. Levcopoulos, and G. Narasimhan. Improved greedy algorithms for constructing sparse geometric spanners. To appear in *Proc. 7th SWAT*, 2000.
14. D. F. Hsu and X-D. Hu. On short two-connected Steiner networks with Euclidean distance. *Networks*, 32(2):133–140, 1998.
15. D. S. Hochbaum, editor. *Approximation Algorithms for $\mathcal{NP}$-Hard Problems*. PWS Publishing Company, Boston, MA, 1996.
16. F. K. Hwang, D. S. Richards, and P. Winter. *The Steiner Tree Problem*. North-Holland, Amsterdam, 1992.
17. S. Khuller. Approximation algorithms for finding highly connected subgraphs. Chapter 6 in [15], pp. 236–265.
18. D. Krznaric, C. Levcopoulos, G. Narasimhan, B. J. Nilsson, and M. Smid. The minimum 2-connected subgraph of a set of points. Manuscript, September 15, 1997.
19. J. S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. *SIAM J. Comput.*, 28(4):1298–1309, 1999.
20. S. B. Rao and W. D. Smith. Approximating geometrical graphs via "spanners" and "banyans." *Proc. 30th ACM STOC*, pp. 540–550, 1998. A full version appeared in [21].
21. S. B. Rao and W. D. Smith. Improved approximation schemes for geometrical graphs via "spanners" and "banyans." May 1998. An extended abstract appeared in [20].
22. M. Stoer. *Design of Survivable Networks*, volume 1531 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1992.
23. K. Steiglitz, P. Weiner, and D. J. Kleitman. The design of minimum cost survivable networks. *IEEE Trans. Circuit Theor.*, 16:455–460, 1969.