Approximating the Geometric Minimum-Diameter Spanning Tree

Michael J. Spriggs* J. Mark Keil† Sergei Bespamyatnikh‡ Michael Segal§ Jack Snoeyink¶

Abstract

Given a set P of points in the plane, a geometric minimum-diameter spanning tree (GMDST) of P is a spanning tree of P such that the longest path through the tree is minimized. In this paper, we present an approximation algorithm that generates a tree whose diameter is no more than $(1+\epsilon)$ times that of a GMDST, for any $\epsilon > 0$. Our algorithm reduces the problem to several grid-aligned versions of the problem and runs within time $O(\epsilon^{-3}+n)$ and space $O(n)$ improving the result by Gudmundsson et al. [4].

1 Introduction

Given a set P of points in the plane, define the weight of an edge between two points of P as the Euclidean (L_2) distance between the points. Compute a spanning tree of P such that the longest path through the tree between any two points is minimized. This tree is known as a *geometric minimum diameter spanning tree* (GMDST) of the point set.

In the most efficient known algorithm Timothy Chan [2] gave $O(n^{\frac{17}{6}})$ runtime solution improving Ho et al. [5] result that uses $O(n^3)$ time to generate a GMDST of n points in the plane. The (almost) cubic time bound on GMDST generation may be too large to be practical for some applications. We are interested in finding an efficent approximating algorithm that generates a spanning tree whose diameter is no more than $(1+\epsilon)$ times the length of a GMDST, for any $\epsilon > 0$. The efficiency of algorithm is measured in terms of n and ϵ. One of our goals is to reduce the dependency of ϵ as much as possible. Using such a heavy machinery as well-separated pair decomposition (WSPD) introduced first by Callahan and Kosaraju [1] Gudmundsson et al. [4] were able to obtain $O(n/\epsilon^3 + n\log n/\epsilon)$ time algorithm with $O(n/\epsilon^2 + n\log n)$ space. In this paper, we describe an algorithm that performs the required task in $O(n+\frac{1}{\epsilon^3})$ time and $O(n)$ space.

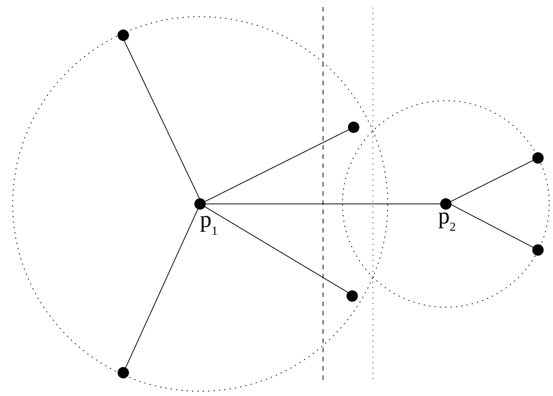

Figure 1: A 2-star GMDST for which connecting each point to the closer internal vertex would produce a spanning tree of greater diameter. This tree satisfies the stability condition, since neither circle contains all points.

2 Preliminaries

A *geometric graph* is a graph that is composed of a set of *points*, and a set of *edges*. Each edge can be represented by a line segment between two points. The *weight* of an edge is the Euclidean, straight-line distance between its endpoints. The *diameter* of a graph is the sum of the edge-weights of the longest path through the graph.

Given a point set P, a k-star, $1 \le k \le n$ for P is a spanning tree of the complete Euclidean graph on P which has k internal nodes. Ho et al. [5] show that every point set admits a GMDST that is either a 1-star or a 2-star, by a nice geometric application of the triangle inequality.

For a given spanning tree of a point set P, $\mathcal{T}(P)$, let $|\mathcal{T}(P)|$ denote the diameter of $\mathcal{T}(P)$. For a path $p_1p_2\dots p_k$ through a tree, let $|p_1p_2\dots p_k|$ denote the length of the path.

Figure 1 illustrates a condition of Jones[6] for how the

*School of Computer Science, University of Waterloo, 200 University Avenue West, Waterloo Ontario N2L 3G1 Canada, mjsprigg@math.uwaterloo.ca.

†Department of Computer Science, University of Saskatchewan, Saskatoon, Saskatchewan, Canada, S7N 5A9, keil@cs.usask.ca, 306-966-4894, FAX 306-966-4884 Supported by NSERC.

‡Dept. of Computer Science, University of Texas at Dallas, Box 830688, Richardson TX 75083 U.S.A., besp@utdallas.edu

§Communication Systems Engineering Dept., Ben Gurion University, Beer-Sheva, 84105 Israel, segal@cse.bgu.ac.il.

¶Dept. of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-3175 U.S.A., snoeyink@cs.unc.edu

remaining points in P can be connected to the given interior vertices of a 2-star.

Lemma 2.1 *[6]*

Suppose that points $p_1, p_2 \in P$ form a horizontal line $\overline{p_1p_2}$ from left to right. There exists a vertical line L such that connecting points left of L to p_1 and right of L to p_2 produces a minimum diameter spanning tree among all spanning trees with interior vertices p_1 and p_2.

Ho et al. [5] also establish a *stability condition* for 2-stars, which can also be illustrated in Figure 1. Let $P_i \subset P \setminus \{p_1, p_2\}$ denote the points joined to interior vertex p_i.

Lemma 2.2 *[5] The diameter of a 2-star is determined by a three-edge path if, for $i \in \{1, 2\}$, some point not joined to p_i is farther than all points joined to p_i. That is,*

$$\max_{q \in P_i} |p_i q| < \max_{q' \in P_{3-i}} |p_i q'|.$$

A *uniform grid* is a grid composed of an infinite number of horizontal and vertical lines in the plane, such that adjacent lines are placed at uniform intervals. The grid breaks up the plane into square regions that we call *grid-squares*. The *center* of a grid-square is the point in the middle of the square that lies equidistant from all four corners of the grid-square. Define a *grid-aligned point set* as a set of points in the plane, such that the points lie only at the centers of grid-squares, with respect to some underlying grid.

3 A Simpler Problem

Instead of a general point set, we first consider a grid-aligned point set P contained in an m row and m column bounding box. We also restrict the spanning tree to be either a 1-star, or a 2-star such that both interior vertices lie in a single row of the grid. We call this special version of the GMDST a restricted geometric minimum diameter spanning tree (RGMDST).

Suppose further that we have been given, for each row j of the m rows, two candidate points $p_{j1}, p_{j2} \in P$ such that (a) if the RGMDST is a 1-star, then the interior vertex is p_{ja} for some $1 \le j \le m$ and $a \in \{1, 2\}$, and (b) if the RGMDST of P is a 2-star then the interior vertices are p_{j1} and p_{j2} for some $1 \le j \le m$. We now show how we can efficiently exactly solve the RGMDST problem when we are given the candidate interior vertices. To do this we first find the optimal 1-star, then find the optimal 2-star and take the minimum of the two.

Lemma 3.1 *Let P be a grid-aligned set of n points contained in an $m \times m$ bounding box. Given two candidate interior vertices $p_{j1}p_{j2} \in P$ in each row j, $1 \le j \le m$, such that the optimal interior vertex is among the candidates, the minimum diameter 1-star of P can be found in $O(m \log m + n)$ time and $O(m + n)$ space.*

Lemma 3.2 *Let P be a grid-aligned set of n points contained in an m row and m column bounding box. Given candidate interior vertices $p_{j1}p_{j2} \in P$ in each row j, $1 \le j \le m$, such that the optimal horizontally restricted 2-star has interior vertices among the candidate pairs, then the optimal 2-star RGMDST of P can be found in $O(m^2 + n)$ time and $O(m + n)$ space.*

4 Solving the RGMDST problem

Imagine a 1-star of a planar point set. Without changing the structure of the graph, move the interior vertex to the left and to the right along a horizontal line.

Lemma 4.1 *Given a point set P, and three collinear points p_1, p_2, p_3, such that p_2 lies between p_1 and p_3, the diameter of the 1-star of P in which p_2 is the interior vertex is smaller than the maximum of the diameter of the 1-star with interior vertex p_1, and the diameter of the 1-star with interior vertex p_3.*

For a given line l and point set P, a *Steiner monopole* of l and P is a point $s \in l$ such that the diameter of the 1-star of $P \cup \{s\}$ with interior vertex s is minimum, among such trees with interior Steiner vertices on l. When we speak of a Steiner monopole of a particular grid-row, we refer to the Steiner monopole that lies on the line that passes through the centers of every grid-square in the grid-row.

Lemma 4.2 *Let P be a grid-aligned point set and let r be a grid-row containing points of P. Let $\Delta_r(P)$ be the minimum diameter spanning tree of P which is restricted to be either a 1-star with interior vertex in row r, or a 2-star with both interior vertices in row r. If $\Delta_r(P)$ is a 2-star then no point of P lies on the horizontal line segment between the two interior vertices, p_1 and p_2, and the Steiner monopole s of the grid-row r must lie between p_1 and p_2.*

4.1 Computing candidate interior vertices

We now show how to compute, for each grid-row in the bounding box of P, the Steiner monopole of that row. From this information, we can compute a set of candidate interior vertices in each row. These are the two points of each row, that are nearest the Steiner monopole, such that one lies left of the Steiner monopole, and the other lies to the right, by Lemma 4.2.

The second order furthest point Voronoi Diagram of P can be computed in $O(m \log m)$ time. Once this diagram is computed, the furthest two points from any point p are those two points corresponding with the Voronoi cell that contains p. For each row of the grid, we can compute the furthest two points from every grid-aligned point in that row in an additional $O(m)$ time by traversing Voronoi cells, from a cell to an adjacent cell, along the line that passes through the points in the row. Once we know the furthest two points from a point p, in constant time we can compute the diameter of the 1-star with interior vertex p. There are m rows, and so it will take a total of $O(m^2)$ time to process all rows in this manner. This gives us the best possible grid-aligned Steiner monopole and its cost in $O(m^2)$ time. The candidate interior vertices of each row can be computed in an additional $O(m)$ time per row.

Theorem 4.1 *Given a set P of n points aligned with a grid $\mathcal{G}$ and contained within a bounding box of $m \times m$ grid-squares, there is an $O(m^2)$-time algorithm that generates a RGMDST of P.*

5 The approximation algorithm

In order to approximate an optimum GMDST of a general point set, we transform the problem to several instances of the RGMDST problem. Let P be an arbitrary set of points in the plane with a GMDST $\Delta(P)$. Suppose that $\Delta(P)$ is a 2-star. We seek to overlay the plane with a grid $\mathcal{G}$, with grid-square edges of length ϕ, so that that the two interior vertices of $\Delta(P)$, p_1 and p_2, lie in a single grid-row. Let $\mathcal{D}$ denote the distance between the furthest two points of P. Therefore, $|p_1p_2| \leq \mathcal{D}$. Let θ denote the smaller angle between the line through p_1 and p_2, and the horizontal grid-lines of $\mathcal{G}$.

Using trigonometry, we find that if $\sin\theta < \frac{\phi}{\mathcal{D}}$ then p_1 and p_2 can reside in a single grid-row of a grid oriented like $\mathcal{G}$.

By insisting that $\sin\theta < \frac{\phi/2}{\mathcal{D}} = \frac{\phi}{2\mathcal{D}}$ we need consider only two such grids, offset from one another by a vertical distance of $\frac{\phi}{2}$.

Figure 2 shows that if θ is small enough the segment joining the interior vertices will lie in one of the two grids whose horizontal grid lines are offset by $\frac{\phi}{2}$.

In order to account for all possible orientations of the line through p_1 and p_2, several orientations of grids are used.

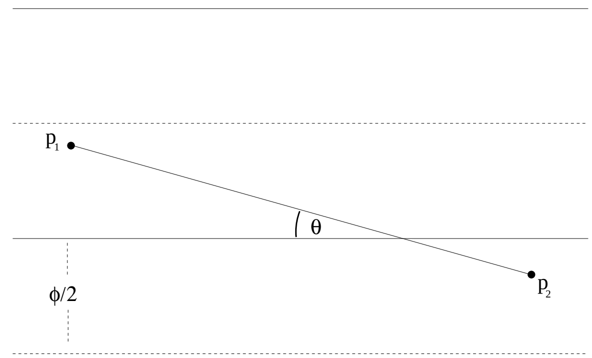

Figure 2: The line through p_1 and p_2, closely aligned with the horizontal grid-lines of two offset grids.

Lemma 5.1 *The number of orientations of grids can be bounded by $\frac{\pi}{\arcsin(\phi/2\mathcal{D})} = O\left(\frac{\mathcal{D}}{\phi}\right)$.*

5.1 Grid transformation

One of the $O(\frac{\mathcal{D}}{\phi})$ oriented grids will contain the two interior vertices of an optimal 2-star of P in a single grid-row. For each such grid, we need to generate a grid-aligned point set P' from P and analyze how the gridding changes the GMDST. If any point $p \in P$ lies on the boundary of two or more grid-squares, the point is moved a negligible distance in some direction until it no longer lies on a boundary. We generate the grid-aligned point set P' as follows. For each grid-square in $\mathcal{G}$, if the grid-square contains a single point of P, then add a single point in the center of the grid-square to P'. If two or more points of P reside in the grid-square, then add exactly two points to P', such that both points lie at the center of the grid-square. Notice that set P' is aligned with grid $\mathcal{G}$. Let $\Delta(P')$ be a GMDST of P'. Generate a spanning tree of P, $\mathcal{T}(P)$, as follows. If $\Delta(P')$ is a 1-star or a 2-star with both interior vertices in the same grid-square, then make $\mathcal{T}(P)$ a 1-star such that the interior vertex of $\mathcal{T}(P)$ is any point $p \in P$ that lies in the same grid-square as does the interior vertex of $\Delta(P')$. Otherwise $\Delta(P')$ is a 2-star with interior vertices p'_1 and p'_2 in different grid-squares.

Choose two interior vertices $p_1, p_2 \in P$ such that p_1 lies in the same grid-square as p'_1 and p_2 lies in the same grid-square as p'_2. In $\mathcal{T}(P)$, generate an edge between p_1 and p_2. For every point $p \in P$, where $p \neq p_i$ for $i \in \{1, 2\}$, such that p resides in the same grid-square as p_i, add edge pp_i to $\mathcal{T}(P)$. For any grid-square that contains at least one point of P', if one or both of these points are linked by an edge to vertex p'_1 in $\Delta(P')$, then for each point $p \in P$ that resides in the grid-square add edge p_1p to $\mathcal{T}(P)$. Otherwise, for each point $p \in P$ that resides in the grid-square, add edge p_2p to $\mathcal{T}(P)$.

We call the above procedure of converting P to P', and then using a GMDST of P' to generate a spanning tree of P, the *grid transformation*.

Lemma 5.2 *Given a set P of points in the plane with GMDST $\Delta(P)$, and some value $\phi > 0$, the grid transformation generates a spanning tree of P, $\mathcal{T}(P)$, such that $|\mathcal{T}(P)| \leq |\Delta(P)| + 6\sqrt{2}\phi$.*

5.2 Putting it all together

Lemma 5.3 *Given a set P of points in the plane such that $\mathcal{D}$ is the largest distance between any two points of P, any GMDST of P must be of size $\mathcal{D}$ or larger.*

Theorem 5.1 *Given a set P of n points in the plane, there exists an algorithm such that, for any $0 < \epsilon$, the algorithm generates a $(1+\epsilon)$-approximate GMDST of P within time $O(\epsilon^{-3} + n)$ and space $O(n)$.*

References

[1] P. Callahan and R. Kosaraju, "A decomposition of multidimensional point sets with applications to k-nearest neighbors and n-body potential fields", *Journal of ACM*, 42(1), pp. 67–90, 1995.

[2] T. Chan, Semi-online maintenance of geometric optima and measures, Proc. 13th ACM-SIAM Symposium on Discrete Algorithms (SODA), 474-483 (2002).

[3] J. Cong, L. He, C. Koh, and P. Madden, Performance optimization of VLSI interconnection layout, Integration, the VLSI Journal, 21, 1–94 (1996).

[4] J. Gudmundsson, H. Haverkort, S.-M. Park, C.-S. Shin, A. Wolff, "Facility location and the geometric minimum-diameter spanning tree", in *Approximation Algorithms for Combinatorial Optimization (APPROX'02)*, pp. 146-160, 2002.

[5] J. Ho and D. Lee and C. Chang and C. Wong, Minimum diameter spanning trees and related problems, SIAM J. on Comput., 20, 987–997 (1991).

[6] W. D. Jones, Euclidean Communication Spanning Trees, M.Sc. Thesis, Univ. of Saskatchewan, 1994.