Geometric Spanners

Joachim Gudmundsson, National ICT Australia Ltd, www.gudmundsson.biz/
Giri Narasimhan, Florida International University, www.cs.fiu.edu/˜giri
Michiel Smid, Carleton University, www.scs.carleton.ca/˜michiel/
entry editor: Andrzej Lingas

INDEX TERMS: Computational geometry, geometric networks, spanners, dilation.

1 PROBLEM DEFINITION

Consider a set S of n points in d-dimensional Euclidean space. A network on S can be modeled as an undirected graph G with vertex set S of size n and an edge set E where every edge (u, v) has a weight. A geometric (Euclidean) network is a network where the weight of the edge (u, v) is the Euclidean distance $|uv|$ between its endpoints. Given a real number $t > 1$ we say that G is a *t-spanner* for S, if for each pair of points $u, v \in S$, there exists a path in G of weight at most t times the Euclidean distance between u and v. The minimum t such that G is a t-spanner for S is called the stretch factor, or dilation, of G. For a more detailed description of the construction of t-spanners see the book by Narasimhan and Smid [17]. The problem considered is the construction of t-spanners given a set S of n points in $\mathcal{R}^d$ and a positive real value $t > 1$, where d is a constant. The aim is to compute a good t-spanner for S with respect to the following quality measures:

size: the number of edges in the graph.
degree: the maximum number of edges incident on a vertex.
weight: the sum of the edge weights.
spanner diameter: the smallest integer k such that for any pair of vertices u and v in S, there is a path in the graph of length at most $t \cdot |uv|$ between u and v containing at most k edges.
fault-tolerance: the resilience of the graph to edge, vertex or region failures.

Thus, good t-spanners require large fault-tolerance and small size, degree, weight and spanner diameter.

2 KEY RESULTS

In this section we describe the three most common approaches for constructing a t-spanner of a set of points in Euclidean space. We also describe the construction of fault-tolerant spanners, spanners among polygonal obstacles and, finally, a short note on dynamic and kinetic spanners.

2.1 Spanners of points in Euclidean space

The most well-known classes of t-spanner networks for points in Euclidean space include: Θ-graphs, WSPD-graphs and Greedy-spanners. In the following sections we give the main idea of each of these classes, together with the known bounds on the quality measures.

2.1.1 The Θ-graph

The Θ-graph was discovered independently by Clarkson and Keil in the late 80's. The general idea is to process each point $p \in S$ independently as follows: partition $\mathcal{R}^d$ into k simplicial cones of angular diameter at most θ and apex at p, where $k = O(1/\theta^{d-1})$. For each non-empty cone C, an edge is added between p and the point in C whose orthogonal projection onto some fixed ray in C emanating from p is closest to p, see Fig. 1a. The resulting graph is called the Θ-graph on S

Theorem 1. *The Θ-graph is a t-spanner of S for $t = \frac{1}{\cos\theta - \sin\theta}$ with $O(\frac{n}{\theta^{d-1}})$ edges and can be computed in $O(\frac{n}{\theta^{d-1}} \log^{d-1} n)$ time using $O(\frac{n}{\theta^{d-1}} + n\log^{d-2} n)$ space.*

The following variants of the Θ-graph also give bounds on the degree, diameter and weight.

Skip-list spanners: The idea is to generalize skip-lists and apply them to the construction of spanners. Construct a sequence of h subsets, $S_1, \ldots, S_h$, where $S_1 = S$ and S_i is constructed from S_{i-1} as follows (reminiscent of the levels in a skip list). For each point in S_{i-1}, flip a fair coin. The set S_i is the set of all points of S_{i-1} whose coin flip produced heads. The construction stops if $S_i = \emptyset$. For each subset a Θ-graph is constructed. The union of the graphs is the skip-list spanner of S with dilation t, having $O(\frac{n}{\theta^{d-1}})$ edges and $O(\log n)$ spanner diameter with high probability [3].

Gap-greedy: A set of directed edges is said to satisfy the *gap* property if the sources of any two distinct edges in the set are separated by a distance that is at least proportional to the length of the shorter of the two edges. Arya and Smid [5] proposed an algorithm that uses the gap property to decide whether or not an edge should be added to the t-spanner graph. Using the gap property the constructed spanner can be shown to have degree $O(1/\theta^{d-1})$ and weight $O(\log n \cdot wt(MST(S)))$, where $wt(MST(S))$ is the weight of the minimum spanning tree of S.

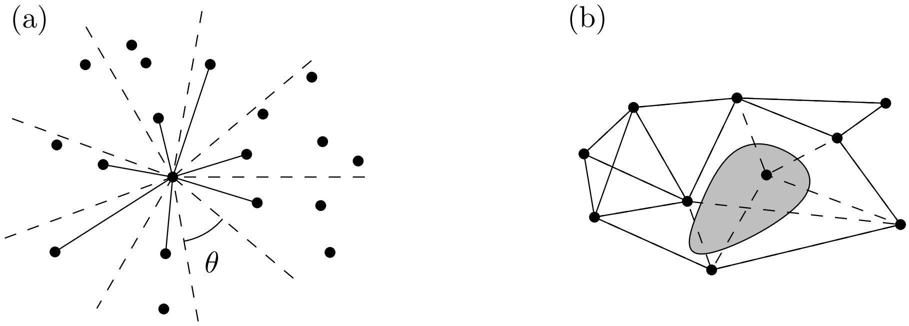

Figure 1: (a) Illustrating the Θ-graph, and (b) a graph with a region-fault.

2.1.2 The WSPD-graph

The well-separated pair decomposition (WSPD) was developed by Callahan and Kosaraju [6]. The construction of a t-spanner using the well-separated pair decomposition is done by first constructing a WSPD of S with respect to a separation constant $s = \frac{4(t+1)}{(t-1)}$. Initially set the spanner graph $G = (S, \emptyset)$ and add edges iteratively as follows. For each well-separated pair $\{A, B\}$ in the decomposition, an edge (a, b) is added to the graph, where a and b are arbitrary points in A and B, respectively. The resulting graph is called the WSPD-graph on S.

Theorem 2. *The WSPD-graph is a t-spanner for S with $O(s^d \cdot n)$ edges and can be constructed in time $O(s^d n + n\log n)$, where $s = 4(t+1)/(t-1)$.*

There are modifications that can be made to obtain bounded diameter or bounded degree.

Bounded diameter: Arya, Mount and Smid [3] showed how to modify the construction algorithm such that the diameter of the graph is bounded by $2\log n$. Instead of selecting an arbitrary point in each well-separated set, their algorithm carefully choose a representative point for each set.

Bounded degree: A single point v can be part of many well-separated pairs and each of these pairs may generate an edge with an endpoint at v. Arya et al. [2] suggested an algorithm that retains only the shortest edge for each cone direction, thus combining the Θ-graph approach with the WSPD-graph. By adding a post-processing step that handles all high-degree vertices, a t-spanner of degree $O(\frac{1}{(t-1)^{2d-1}})$ is obtained.

2.1.3 The Greedy-spanner

The greedy algorithm was first presented in 1989 by Bern and since then the greedy algorithm has been subject to considerable research. The graph constructed using the greedy algorithm is called a Greedy-spanner and the general idea is that the algorithm iteratively builds a graph G. The edges in the complete graph are processed in order of increasing edge length. Testing an edge (u, v) entails a shortest path query in the partial spanner graph G. If the shortest path in G between u and v is at most $t \cdot |uv|$ then the edge (u, v) is discarded, otherwise it is added to the partial spanner graph G.

Das, Narasimhan and Salowe [11] proved that the greedy-spanner fulfills the so-called *leapfrog property*. A set of undirected edges E is said to satisfy the t-leapfrog property, if for every $k \geq 2$, and for every possible sequence $\{(p_1, q_1), \ldots, (p_k, q_k)\}$ of pairwise distinct edges of E,

$$t \cdot |p_1 q_1| < \sum_{i=2}^{k} |p_i q_i| + t \cdot \Big(\sum_{i=1}^{k-1} |q_i p_{i+1}| + |p_k q_1| \Big).$$

Using the leapfrog property it is possible to bound weight of the graph. Das and Narasimhan [10] observed that the Greedy-spanner can be approximated while maintaining the leapfrog property. This observation allowed for faster construction algorithms.

Theorem 3. *[14] The greedy-spanner is a t-spanner of S with $O(\frac{n}{(t-1)^d} \log(\frac{1}{t-1}))$ edges, maximum degree $O(\frac{1}{(t-1)^d} \log(\frac{1}{t-1}))$, weight $O(\frac{1}{(t-1)^{2d}} \cdot wt(MST(S))$, and can be computed in time $O(\frac{n}{(t-1)^{2d}} \log n)$.*

2.2 Fault-tolerant spanners

The concept of fault-tolerant spanners was first introduced by Levcopoulos et al. [15] in 1998, i.e., after one or more vertices or edges fail, the spanner should retain its good properties. In particular, there should still be a short path between any two vertices in what remains of the spanner after the fault. Czumaj and Zhao [8] showed that a greedy approach produces a k-vertex (or k-edge) fault tolerant geometric t-spanner with degree $O(k)$ and total weight $O(k^2 \cdot wt(MST(S)))$; these bounds are asymptotically optimal.

For geometric spanners it is natural to consider *region faults*, i.e., faults that destroy all vertices and edges intersecting some geometric fault region. For a fault region F let $G \ominus F$ be the part of G that remains after the points from S inside F and all edges that intersect F have been removed from the graph, see Fig. 1b. Abam et al. [1] showed how to construct region-fault tolerant t-spanners of size $O(n \log n)$ that are fault-tolerant to any convex region-fault. If one is allowed to use Steiner points then a linear size t-spanner can be achieved.

2.3 Spanners among obstacles

The visibility graph of a set of pairwise non-intersecting polygons is a graph of intervisible locations. Each polygonal vertex is a vertex in the graph and each edge represents a visible connection between them, that is, if two vertices can see each other, an edge is drawn between them. This graph is useful since it contains the shortest obstacle avoiding path between any pair of vertices.

Das [9] showed that a t-spanner of the visibility graph of a point set in the Euclidean plane can be constructed by using the Θ-graph approach followed by a pruning step. The obtained graph has linear size and constant degree.

2.4 Dynamic and kinetic spanners

Not much is known in the areas of dynamic or kinetic spanners. Arya et al. [4] showed a data structure of size $O(n \log^d n)$ that maintains the skip-list spanner, described in Section 2.1.1, in $O(\log^d n \log\log n)$ expected amortized time per insertion and deletion in the model of random updates.

Gao et al. [13] showed how to maintain a t-spanner of size $O(\frac{n}{(t-1)^d})$ and maximum degree $O(\frac{1}{(t-2)^d}\log\alpha)$ in time $O(\frac{\log\alpha}{(t-1)^d})$ per insertion and deletion, where α denotes the aspect ratio of S, i.e., the ratio of the maximum pairwise distance to the minimum pairwise distance. The idea is to use an hierarchical structure T with $O(\log\alpha)$ levels, where each level contains a set of centers (subset of S). Each vertex v on level i in T is connected by an edge to all other vertices on level i within distance $O(\frac{2^i}{t-1})$ of v. The resulting graph is a t-spanner of S and it can be maintained as stated above. The approach can be generalized to the kinetic case so that the total number of events in maintaining the spanner is $O(n^2 \log n)$ under pseudo-algebraic motion. Each event can be updated in $O(\frac{\log\alpha}{(t-1)^d})$ time.

3 APPLICATIONS

The construction of sparse spanners has been shown to have numerous applications areas such as metric space searching [18], which includes query by content in multimedia objects, text retrieval, pattern recognition and function approximation. Another example is broadcasting in communication networks [16]. Several well-known theoretical results also use the construction of t-spanners as a building block, for example, Rao and Smith [19] made a breakthrough by showing an optimal $O(n \log n)$-time approximation scheme for the well-known Euclidean *traveling salesperson problem*, using t-spanners (or banyans). Similarly, Czumaj and Lingas [7] showed approximation schemes for minimum-cost multi-connectivity problems in geometric networks.

4 OPEN PROBLEMS

There are many open problems in this area. We only mention a few:

1. Design a dynamic t-spanner that can be updated in $O(\log^c n)$ time, for some constant c.

2. Determine if there exists a fault-tolerant t-spanner of linear size for convex region faults.

3. The k-vertex fault tolerant spanner by Czumaj and Zhao [8] produces a k-vertex fault tolerant t-spanner of degree $O(k)$ and weight $O(k^2 \cdot wt(MST(S)))$. However, it is not known how to implement it efficiently. Can such a spanner be computed in $O(n \log n + kn)$ time?

4. Bound the weight of skip-list spanners.

5 EXPERIMENTAL RESULTS

The problem of constructing spanners has received considerable attention from a theoretical perspective but not much attention from a practical, or experimental perspective. Navarro and Paredes [18] presented four heuristics for point sets in high-dimensional space ($d = 20$) and showed by empirical methods that the running time was $O(n^{2.24})$ and the number of edges in the produced graphs was $O(n^{1.13})$. Recently Farshi and Gudmundsson [12] performed a thorough comparison of the construction algorithms discussed in Section 2.1.

6 CROSS REFERENCES

Plane geometric spanners, Well-separated pair decomposition and Applications of geometric spanners.

References

[1] M. A. Abam, M. de Berg, M. Farshi, and J. Gudmundsson, *Region-fault tolerant geometric spanners*, in Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms, 2007.

[2] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. Smid, *Euclidean spanners: short, thin, and lanky*, in Proceedings of the 27th ACM Symposium on Theory of Computing, 1995, pp. 489–498.

[3] S. Arya, D. M. Mount, and M. Smid, *Randomized and deterministic algorithms for geometric spanners of small diameter*, in Proceedings of the 35th IEEE Symposium on Foundations of Computer Science, 1994, pp. 703–712.

[4] S. Arya, D. M. Mount, and M. Smid, *Dynamic algorithms for geometric spanners of small diameter: Randomized solutions*, Computational Geometry - Theory and Applications, 13 (1999), pp. 91–107.

[5] S. Arya and M. Smid, *Efficient construction of a bounded-degree spanner with low weight*, Algorithmica, 17 (1997), pp. 33–54.

[6] P. B. Callahan and S. R. Kosaraju, *A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields*, Journal of the ACM, 42 (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, vol. 1853 of Lecture Notes in Computer Science, Springer-Verlag, 2000, pp. 856–868.

[8] A. Czumaj and H. Zhao, *Fault-tolerant geometric spanners*, Discrete & Computational Geometry, 32 (2004), pp. 207–230.

[9] G. Das, *The visibility graph contains a bounded-degree spanner*, in Proceedings of the 9th Canadian Conference on Computational Geometry, 1997.

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

[11] G. Das, G. Narasimhan, and J. Salowe, *A new way to weigh malnourished Euclidean graphs*, in Proceedings of the 6th ACM-SIAM Symposium on Discrete Algorithms, 1995, pp. 215–222.

[12] M. Farshi and J. Gudmundsson, *Experimental study of geometric t-spanners*, in Proceedings of the 13th Annual European Symposium on Algorithms, vol. 3669 of Lecture Notes in Computer Science, Springer-Verlag, 2005, pp. 556–567.

[13] J. Gao, L. J. Guibas, and A. Nguyen, *Deformable spanners and applications*, in Proceedings of the 20th ACM Symposium on Computational Geometry, 2004, pp. 190–199.

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

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

[16] X.-Y. Li, *Applications of computational geomety in wireless ad hoc networks*, in Ad Hoc Wireless Networking, X.-Z. Cheng, X. Huang, and D.-Z. Du, eds., Kluwer, 2003, pp. 197–264.

[17] G. Narasimhan and M. Smid, *Geometric spanner networks*, Cambridge University Press, 2006.

[18] G. Navarro and R. Paredes, *Practical construction of metric t-spanners*, in Proceedings of the 5th Workshop on Algorithm Engineering and Experiments, SIAM Press, 2003, pp. 69–81.

[19] S. Rao and W. D. Smith, *Approximating geometrical graphs via spanners and banyans*, in Proceedings of the 30th ACM Symposium on Theory of Computing, 1998, pp. 540–550.