Title:	Applications of Geometric Spanner Networks
Name:	Joachim Gudmundsson[1], Giri Narasimhan[2], Michiel Smid[3]
Affil./Addr. 1:	School of Information Technologies, The University of Sydney, NSW, Australia
Affil./Addr. 2:	School of Computing & Information Sciences Florida International University, FL, USA
Affil./Addr. 3:	School of Computer Science Carleton University, ON, Canada
Keywords:	distance oracles; spanners; dilation; shortest paths; approximation algorithms; cluster graphs
SumOriWork:	2002; Gudmundsson, Levcopoulos, Narasimhan, Smid 2005; Gudmundsson, Narasimhan, Smid 2008; Gudmundsson, Levcopoulos, Narasimhan, Smid

Applications of Geometric Spanner Networks

Joachim Gudmundsson[1], Giri Narasimhan[2], Michiel Smid[3]

[1] School of Information Technologies, The University of Sydney, NSW, Australia
[2] School of Computing & Information Sciences Florida International University, FL, USA
[3] School of Computer Science Carleton University, ON, Canada

Years aud Authors of Summarized Original Work

2002; Gudmundsson, Levcopoulos, Narasimhan, Smid
2005; Gudmundsson, Narasimhan, Smid
2008; Gudmundsson, Levcopoulos, Narasimhan, Smid

Keywords

distance oracles; spanners; dilation; shortest paths; approximation algorithms; cluster graphs

Problem Definition

Given a geometric graph in d-dimensional space, it is useful to preprocess it so that distance queries, exact or approximate, can be answered efficiently. Algorithms that can report distance queries in constant time are also referred to as "distance oracles". With unlimited preprocessing time and space, it is clear that exact distance oracles can be easily designed. This entry sheds light on the design of approximate distance oracles with limited preprocessing time and space for the family of geometric graphs with constant dilation.

Notation and Definitions If p and q are points in $\mathcal{R}^d$, then the notation $|pq|$ is used to denote the Euclidean distance between p and q; the notation $\delta_G(p,q)$ is used to denote the Euclidean length of a shortest path between p and q in a geometric network G. Given a constant $t > 1$, a graph G with vertex set S is a t-spanner for S if $\delta_G(p,q) \leq t|pq|$ for any two points p and q of S. A t-spanner network is said to have *dilation* (or *stretch factor*) t. A $(1+\varepsilon)$-approximate shortest path between p and q is defined to be any path in G between p and q having length Δ, where $\delta_G(p,q) \leq \Delta \leq (1+\varepsilon)\delta_G(p,q)$. For a comprehensive overview of geometric spanners, see the book by Narasimhan and Smid [14].

All networks considered in this entry are simple and undirected. The model of computation used is the traditional algebraic computation tree model with the added power of indirect addressing. In particular, the algorithms presented here do not use the non-algebraic floor function as a unit-time operation. The problem is formalized below.

Problem 1 (Distance Oracle). Given an arbitrary real constant $\epsilon > 0$, and a geometric graph G in d-dimensional Euclidean space with constant dilation t, design a data structure that answers $(1+\epsilon)$-approximate shortest path length queries in constant time.

The data structure can also be applied to solve several other problems. These include (a) the problem of reporting approximate distance queries between vertices in a planar polygonal domain with "rounded" obstacles, (b) query versions of *closest pair* problems, and (c) the efficient computation of the approximate dilations of geometric graphs.

Survey of Related Research The design of efficient data structures for answering distance queries for general (non-geometric) networks was considered by Thorup and Zwick [17] (unweighted general graphs), Baswanna and Sen [3] (weighted general graphs, i.e., arbitrary metrics), and Arikati et al. [2] and Thorup [16] (weighted planar graphs).

For the geometric case, variants of the problem have been considered in a number of papers (for a recent paper see, for example, Chen et al. [5]). Work on the approximate version of these variants can also be found in many articles (for a recent paper see, for example, Agarwal et al. [1]). The focus of this entry are the results reported in the work of Gudmundsson et al. [10; 11; 12; 13]. Similar results on distance oracles were proved subsequently for *unit disk graphs* [7]. Practical implementations of distance oracles in geometric networks have also been investigated [15].

Key Results

The main result of this entry is the existence of approximate distance oracle data structures for geometric networks with constant dilation (see Theorem 4 below). As preprocessing, the network is "pruned" so that it only has a linear number of edges. The data structure consists of a series of "cluster graphs" of increasing coarseness each of which helps answer approximate queries for pairs of points with interpoint distances of different scales. In order to pinpoint the appropriate cluster graph to search in for a given query, the data structure uses the bucketing tool described below. The idea of using cluster graphs to speed up geometric algorithms was first introduced by Das and Narasimhan [6] and later used by Gudmundsson et al. [9] to design an efficient algorithm to compute $(1+\varepsilon)$-spanners. Similar ideas were explored by Gao et al. [8] for applications to the design of mobile networks.

Pruning If the input geometric network has a superlinear number of edges, then the preprocessing step for the distance oracle data structure involves efficiently "pruning" the network so that it has only a linear number of edges. The pruning may result in a small increase of the dilation of the spanner. The following theorem was proved by Gudmundsson et al. [12].

Theorem 1. *Let $t > 1$ and $\varepsilon' > 0$ be real constants. Let S be a set of n points in $\mathcal{R}^d$, and let $G = (S, E)$ be a t-spanner for S with m edges. There exists an algorithm to compute in $O(m + n \log n)$ time, a $(1+\varepsilon')$-spanner of G having $O(n)$ edges and whose weight is $O(wt(MST(S)))$.*

The pruning step requires the following technical theorem proved by Gudmundsson et al. [12].

Theorem 2. *Let S be a set of n points in $\mathcal{R}^d$, and let $c \geq 7$ be an integer constant. In $O(n \log n)$ time, it is possible to compute a data structure $D(S)$ consisting of:*

1. *a sequence $L_1, L_2, \ldots, L_\ell$ of real numbers, where $\ell = O(n)$, and*
2. *a sequence $S_1, S_2, \ldots, S_\ell$ of subsets of S such that $\sum_{i=1}^{\ell} |S_i| = O(n)$,*

such that the following holds. For any two distinct points p and q of S, it is possible to compute in $O(1)$ time an index i with $1 \leq i \leq \ell$ and two points x and y in S_i such that (a) $L_i/n^{c+1} \leq |xy| < L_i$, and (b) both $|px|$ and $|qy|$ are less than $|xy|/n^{c-2}$.

Despite its technical nature, the above theorem is of fundamental importance to this work. In particular, it helps to deal with networks where the interpoint distances are not confined to a polynomial range, i.e., there are pairs of points that are very close to each other and very far from each other.

Bucketing Since the model of computation assumed here does not allow the use of floor functions, an important component of the algorithm is a "bucketing tool" that allows (after appropriate preprocessing) constant-time computation of a quantity referred to as BINDEX, which is defined to be the floor of the logarithm of the interpoint distance between any pair of input points.

Theorem 3. *Let S be a set of n points in $\mathcal{R}^d$ that are contained in the hypercube $(0, n^k)^d$, for some positive integer constant k, and let ε be a positive real constant. The set S can be preprocessed in $O(n \log n)$ time into a data structure of size $O(n)$, such that for any two points p and q of S, with $|pq| \geq 1$, it is possible to compute in constant time the quantity* $\text{BINDEX}_\varepsilon(p, q) = \lfloor \log_{1+\varepsilon} |pq| \rfloor$.

The constant-time computation mentioned in Theorem 3 is achieved by reducing the problem to one of answering least common ancestor queries for pairs of nodes in a tree, a problem for which constant-time solutions were devised most recently by Bender and Farach-Colton [4].

Main Results Using the bucketing and the pruning tools, and using the algorithms described by Gudmundsson et al. [13], the following theorem can be proved.

Theorem 4. *Let $t > 1$ and $\varepsilon > 0$ be real constants. Let S be a set of n points in $\mathcal{R}^d$, and let $G = (S, E)$ be a t-spanner for S with m edges. The graph G can be preprocessed into a data structure of size $O(n \log n)$ in time $O(mn \log n)$, such that for any pair of query points $p, q \in S$, it is possible to compute a $(1+\varepsilon)$-approximation of the shortest-path distance in G between p and q in $O(1)$ time. Note that all the big-Oh notations hide constants that depend on d, t and ε.*

Additionally, if the traditional algebraic model of computation (without indirect addressing) is assumed, the following weaker result can be proved.

Theorem 5. *Let S be a set of n points in $\mathcal{R}^d$, and let $G = (S, E)$ be a t-spanner for S, for some real constant $t > 1$, having m edges. Assuming the algebraic model of computation, in $O(m \log\log n + n \log^2 n)$ time, it is possible to preprocess G into a data structure of size $O(n \log n)$, such that for any two points p and q in S, a $(1+\varepsilon)$-approximation of the shortest-path distance in G between p and q can be computed in $O(\log\log n)$ time.*

Applications

As mentioned earlier, the data structure described above can be applied to several other problems. The first application deals with reporting distance queries for a planar domain with polygonal obstacles. The domain is further constrained to be t-rounded, which means that the length of the shortest obstacle-avoiding path between any two points in the input point set is at most t times the Euclidean distance between them. In other words, the visibility graph is required to be a t-spanner for the input point set.

Theorem 6. *Let $\mathcal{F}$ be a t-rounded collection of polygonal obstacles in the plane of total complexity n, where t is a positive constant. One can preprocess $\mathcal{F}$ in $O(n \log n)$ time into a data structure of size $O(n \log n)$ that can answer obstacle-avoiding $(1+\varepsilon)$-approximate shortest path length queries in time $O(\log n)$. If the query points are vertices of $\mathcal{F}$, then the queries can be answered in $O(1)$ time.*

The next application of the distance oracle data structure includes query versions of *closest pair* problems, where the queries are confined to specified subset(s) of the input set.

Theorem 7. *Let $G = (S, E)$ be a geometric graph on n points and m edges, such that G is a t-spanner for S, for some constant $t > 1$. One can preprocess G in time $O(m + n \log n)$ into a data structure of size $O(n \log n)$ such that given a query subset S' of S, a $(1+\varepsilon)$-approximate closest pair in S' (where distances are measured in G) can be computed in time $O(|S'| \log |S'|)$.*

Theorem 8. *Let $G = (S, E)$ be a geometric graph on n points and m edges, such that G is a t-spanner for S, for some constant $t > 1$. One can preprocess G in time $O(m + n \log n)$ into a data structure of size $O(n \log n)$ such that given two disjoint query subsets X and Y of S, a $(1+\varepsilon)$-approximate bichromatic closest pair (where distances are measured in G) can be computed in time $O((|X| + |Y|) \log(|X| + |Y|))$.*

The last application of the distance oracle data structure includes the efficient computation of the approximate dilations of geometric graphs.

Theorem 9. *Given a geometric graph on n vertices with m edges, and given a constant C that is an upper bound on the dilation t of G, it is possible to compute a $(1+\varepsilon)$-approximation to t in time $O(m + n \log n)$.*

Open Problems

Two open problems remain unanswered.

1. Improve the space utilization of the distance oracle data structure from $O(n \log n)$ to $O(n)$.
2. Extend the approximate distance oracle data structure to report not only the approximate distance, but also the approximate shortest path between the given query points.

Cross-References

Geometric Spanners
All-pairs Shortest Paths
Sparse Graph Spanners
Synchronizers, Spanners

Recommended Reading

1. Agarwal PK, Har-Peled S, Karia M (2000) Computing approximate shortest paths on convex polytopes. In: Proceedings of the 16th ACM Symposium on Computational Geometry, pp 270–279
2. Arikati S, Chen DZ, Chew LP, Das G, Smid M, Zaroliagis CD (1996) Planar spanners and approximate shortest path queries among obstacles in the plane. In: Proceedings of the 4th Annual European Symposium on Algorithms, Springer-Verlag, Berlin, Lecture Notes in Computer Science, vol 1136, pp 514–528
3. Baswana S, Sen S (2004) Approximate distance oracles for unweighted graphs in $\tilde{O}(n^2)$ time. In: Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms, pp 271–280
4. Bender MA, Farach-Colton M (2000) The LCA problem revisited. In: Proceedings of the 4th Latin American Symposium on Theoretical Informatics, Springer-Verlag, Berlin, Lecture Notes in Computer Science, vol 1776, pp 88–94
5. Chen DZ, Daescu O, Klenk KS (2001) On geometric path query problems. International Journal of Computational Geometry & Applications 11:617–645
6. Das G, Narasimhan G (1997) A fast algorithm for constructing sparse Euclidean spanners. International Journal of Computational Geometry & Applications 7:297–315
7. Gao J, Zhang L (2005) Well-separated pair decomposition for the unit-disk graph metric and its applications. SIAM J Comput 35(1):151–169
8. Gao J, Guibas LJ, Hershberger J, Zhang L, Zhu A (2003) Discrete mobile centers. Discrete & Computational Geometry 30:45–63
9. Gudmundsson J, Levcopoulos C, Narasimhan G (2002) Fast greedy algorithms for constructing sparse geometric spanners. SIAM Journal on Computing 31:1479–1500
10. Gudmundsson J, Levcopoulos C, Narasimhan G, Smid M (2002) Approximate distance oracles for geometric graphs. In: Proceedings of the 13th ACM-SIAM Symposium on Discrete Algorithms, pp 828–837
11. Gudmundsson J, Levcopoulos C, Narasimhan G, Smid M (2002) Approximate distance oracles revisited. In: Proceedings of the 13th International Symposium on Algorithms and Computation, Springer-Verlag, Berlin, Lecture Notes in Computer Science, vol 2518, pp 357–368
12. Gudmundsson J, Narasimhan G, Smid M (2005) Fast pruning of geometric spanners. In: Proceedings of the 22nd Symposium on Theoretical Aspects of Computer Science, Springer-Verlag, Berlin, Lecture Notes in Computer Science, vol 3404, pp 508–520
13. Gudmundsson J, Levcopoulos C, Narasimhan G, Smid M (2008) Approximate distance oracles for geometric spanners. ACM Transactions on Algorithms 4(1.10)
14. Narasimhan G, Smid M (2007) Geometric Spanner Networks. Cambridge University Press, Cambridge, UK
15. Sankaranarayanan J, Samet H (2010) Query processing using distance oracles for spatial networks. IEEE Trans Knowl Data Eng 22(8):1158–1175
16. Thorup M (2004) Compact oracles for reachability and approximate distances in planar digraphs. Journal of the ACM 51:993–1024
17. Thorup M, Zwick U (2001) Approximate distance oracles. In: Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing, pp 183–192