06481 Abstracts Collection
Geometric Networks and Metric Space Embeddings
— Dagstuhl Seminar —

Joachim Gudmundsson[1], Rolf Klein[2], Giri Narasimhan[3], Michiel Smid[4] and Alexander Wolff[5]

[1] National ICT Australia, AU
joachim.gudmundsson@nicta.com.au
[2] Univ. Bonn, DE
rolf.klein@uni-bonn.de
[3] Florida Int. Univ. Miami, US
giri@cs.fiu.edu
[4] Carleton Univ. Ottawa, CA
michiel@scs.carleton.ca
[5] TU Eindhoven, NL
awolff@win.tue.nl

Abstract. The Dagstuhl Seminar 06481 "Geometric Networks and Metric Space Embeddings" was held from November 26 to December 1, 2006 in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. In this paper we describe the seminar topics, we have compiled a list of open questions that were posed during the seminar, there is a list of all talks and there are abstracts of the presentations given during the seminar. Links to extended abstracts or full papers are provided where available.

Keywords. Geometric networks, metric space embeddings, phylogenetic networks, spanners, dilation, distortion

Dagstuhl Seminar Proceedings 06481
Geometric Networks and Metric Space Embeddings
http://drops.dagstuhl.de/opus/volltexte/2007/1029

1 Summary of the seminar

This seminar has, for the first time, brought together scientists from three different communities who are actively working on distance problems.

Geometric networks are at the core of any model for the flows of goods, traffic or information. They also play an important role in telecommunication, VLSI design, motion planning (robotics), pattern matching, data compression, bioinformatics (gene analysis), and sensor networks. One is interested in spanners with other useful properties like a linear number of edges, small total edge length, small node degree, few crossings, or small link diameter. Apart from these applications, geometric spanners have had great impact on the construction of approximation algorithms, e.g., for the traveling salesman problem. Such problems have been investigated by researchers in computational geometry and combinatorial optimization. For storage, visualization, and retrieval of high-dimensional data the question of reducing the dimension plays a crucial rôle. This has led to a theory of *metric space embeddings* that was most actively developed by scientists in discrete geometry and mathematics. Finally, mathematicians and biologists are interested in metric properties and the visualization of *phylogenetic networks*.

For each of the three fields, a survey talk was given at the seminar (Michiel Smid: *Geometric Spanner Networks*, Anupam Gupta: *Metric Embeddings*, Daniel Huson: *Application of Phylogenetic Networks in Evolutionary Studies*). In addition, twenty-two regular talks were given by the seminar participants. They encouraged a fruitful exchange of ideas and stimulated interesting discussions and co-operations.

The range of interesting topics that came up during this seminar is well documented by the following list of problems that were discussed in two open-problem sessions and attacked by small groups of participants in the corresponding problem-solving sessions.

2 List of open problems

(compiled by Alexander Wolff)

Problem 1 (Günter Rote) Given a set S of n points in the plane, find a weighted graph $G(S', E)$ with $S' \supseteq S$ and few edges such that for each pair $\{p, q\}$ of points in S it holds that the shortest-path distance $d_G(p, q)$ in G equals the Manhattan distance $d_1(p, q)$ of p and q. There is a simple divide-and-conquer algorithm that produces graphs with $O(n \log n)$ edges.

Question: Is it possible to construct such graphs with a linear number of edges?

Answer: For the case that $S' \subset \mathbb{R}^2$ and edge lengths are Euclidean distances, Gudmundsson et al. [GKKS07] have recently shown that $\Omega(n \log n)$ is in fact a lower bound.

Problem 2 (Pat Morin) An ordered θ-graph G is a $(1+\varepsilon)$-spanner for random vertex insertion orders. There is an order (e.g., random) such that G has hop diameter $O(\log n)$. It is also known that there is an order such that G has maximum degree $O(\log n)$, although the expected maximum degree is linear.

Questions:

- Is there an insertion order that yields a $(1 + \varepsilon)$-spanner with logorithmic hop diameter *and* maximum degree?
- Is there an order that yields maximum degree in $o(\log n)$?

Problem 3 (Victor Chepoi) Given a set S of points in the plane and a set F of pairs of points in S, the *F-restricted minimum Manhattan network problem* consists in computing a set N of axis-parallel line segments in the plane such that each pair $\{p, q\}$ in F is connected in $\bigcup N$ by a Manhattan path, i.e., a path of length $d_1(p, q)$. It can be shown that the problem has an intgrality gap of 1.5.

Question: Is there a constant-factor approximation for the F-restricted minimum Manhattan network problem?

Problem 4 (Victor Chepoi) Given a finite metric space (X, d), find a minimum-weight network $N(V, E)$ such that $X \subseteq V$ and for all $u, v \in X$ it holds that $d_N(u, v) = d(u, v)$. It is known that this problem is NP-hard even if d is integral [Alt88].

Question: Is there a constant-factor approximation for the *metric minimum-weight network problem*?

Problem 5 (Victor Chepoi) Let $c_p(n)$ be the smallest d such that any n-point metric space (X, δ) has an isometric (zero-distortion) embedding into

$(\mathbb{R}^d, L_p)$. The following is known about $c_p(n)$:

p	n	$c_p(n)$	reference
1	2	6	Bandelt & Chepoi
∞	2	6	— " —
∞	3	∞	Jeff Edmonds
2	n	$n+3$	Menger
1	3	≥ 10	

Question: Is there a finite upper bound on $c_1(3)$?

Problem 6 (Victor Chepoi & Marc Benkert) A set S of n points in the plane induces a grid $\Gamma(S)$, which consists of all line segments that result from intersecting each horizontal and each vertical line through a point in S with the bounding box of S. Given S, we want to select a minimum-weight subset N of grid edges in $\Gamma(S)$ and to orient each edge in the subset such that each pair in $S \times S$ is connected by a directed Manhattan path in the network formed by the edges in N. Since the problem does not have a solution if S contains points on the same horizontal or vertical line, we assume that this is not the case. We say that a rectangle $R(p,q)$ spanned by a pair $\{p,q\}$ of points in S is empty if $R \cap S = \{p,q\}$. It is not hard to see that the network N_{empty} constructed as follows is a bidirected Manhattan network: go through all empty rectangles spanned by points in S and orient the edges on the boundary of each empty rectangle $R(p,q)$ in (counter)clockwise order if the line segment pq has positive (negative) slope.

Question: Is N_{empty} always a *minimum-weight* bidirected Manhattan network?

Problem 7 (Anupam Gupta) An *equilateral set* in a metric space is a set in which all pairwise distances are the same. Let $e(l_p^n)$ denote the maximum cardinality of an equilateral set in the metric space $l_p^n = (\mathbb{R}^n, l_p)$. Then trivially $e(l_2^n) = n+1$. Using the vectors of the standard basis shows that $e(l_1^n) \geq 2n$. Alon and Pudlak [AP03] have shown that $e(l_1^n) = O(n \log n)$. Kusner [Guy83] conjectured $e(l_1^n) = 2n$ and $e(l_p^n) = n+1$ for every $1 < p < \infty$. Swanepoel [Swa04] showed that in fact $e(l_4^n) = n+1$ for every $n \geq 1$, but that $e(l_p^n) \geq (1+\varepsilon_p)n$ if $1 < p < 2$, where $\varepsilon_p > 0$ depends on p. He also showed that $e(l_p^n) \leq (2\lceil p/4 \rceil - 1)n + 1$ if p is an even integer.

Task: Find better bounds for $e(l_p^n)$!

Problem 8 (Rolf Klein) Given a regular pentagon, find the minimum-dilation triangulation (allowing Steiner points). Taking just the center as Steiner point yields dilation $\delta \approx 1.05$. It is known that $1 < \delta < 1.0204$. The upper bound uses five Steiner points.

Task: Improve the bounds!

Problem 9 (Joachim Gudmundsson) Given a geometric spanner $G(S, E)$ with an set $S \subset \mathbb{R}^2$ of n vertices and with m edges, find the edge whose insertion decreases the stretch-factor of G most. This can be done in $O(n^3 m + n^4 \log n)$ time by repeatedly applying Dijkstra's algorithm [FGG05]. There is a $(2+\varepsilon)$-approximation that runs in $O(nm + n^2(\log n + 1/\varepsilon^{3d}))$ time [FGG05].

Question: What about the k best improvements? (The problem is hard if $k \in \Omega(n)$ [CHL07].)

Problem 10 (Anastasios Sidiropoulos) Let G be a graph that is planar w.r.t. a surface S_g of genus g.

Question: Does there always exist a subgraph H of G such that $S_g \setminus H$ is homothetic to the disk, and H has dilation depending only on g? (For example, if $g = 1$, a cycle yields dilation 1.)

Problem 11 (Anupam Gupta) Consider the problem of embedding weighted graphs with small distortion into l_1^d, where the dimension d can be arbitrary. For example, $K_{2,3}$ does not embed isometrically into l_1^d for any dimension d. In fact the distortion is $4/3$. In l_1^d the pentagon inequality holds. It says that given a set of two and a set of three points, the sum of the (six) distances between points of different sets must be at least the sum of the (four) distances between points of same set. This shows that $K_{2,3}$ cannot be realized with unit distances in l_1^d.

Similar techniques can be used to show that $K_{2,n}$ does not embed into ℓ_1 with distortion less than $3/2$ [ADG+03]. On the other hand, it is known that all *planar* graphs embed into ℓ_1 with distortion $O(\sqrt{\log n})$ [Rao99], and that series-parallel graphs embed into ℓ_1 with distortion $O(1)$ [GNRS99]. In recent as-yet-unpublished work, James Lee and colleagues have shown that the Imase-Waxman "diamond" graphs (see, e.g., [GNRS99]) require a distortion of $2 - o(1)$ to embed into ℓ_1.

Question: What are the best lower bounds for embedding planar graphs into ℓ_1?

Problem 12 (Giri Narasimham) Given a set S of n points in the plane, find the minimum-dilation path whose vertex set is S.

Questions:

– Is the problem NP-hard?
– What about constant-factor approximations?

Problem 13 (Victor Chepoi) Given a cycle C_n where each edge e has a length and a cost associated with it, and a list of demands (u_i, v_i, d_i) $(i = 1, \ldots, m)$ that require that $d_{C_n}(u_i, v_i) \leq d_i$, contract edges such that all demands are satisfied and the sum of the costs of the contracted edges is minimum.

There is a 2-approximation for the problem, and the problem is NP-hard if double edges are allowed. It is also hard for cycles of cycles.

Questions:

- What about hardness of the simple problem?
- What about unit cost?

Problem 14 (Matthew Katz) Given a set S of n points in the plane, the *minimum-area spanning tree* (MAST) is a spanning tree of S that minimizes the area of the union of the $n-1$ disks whose diameters are the edges of the tree. It is known that the Euclidean minimum spanning tree of S is a constant-factor approximation for MAST [CK07].

Question: Is the MAST problem NP-hard?

Problem 15 (Piotr Indyk) Consider the problem of finding a minimum-distortion embedding of a given n-point subset of S^2 into $\mathbb{R}^2$.

Question: Is this problem NP-hard?

Problem 16 (Piotr Indyk) Consider the problem of embedding metric spaces into the real line.

Questions:

- Is there a polynomial time algorithm which, given an input metric that is embeddable with distortion c into the line, produces a $c^{O(1)}$-distortion embedding into the line? In other words, is there a $c^{O(1)}$-approximation algorithm for this problem?
- Is there a constant-factor approximation algorithm for embedding unweighted graph metrics into the line?

Problem 17 (Marc van Kreveld) It is known that embedding a star on the line has distortion $O(c^3)$.

Task: Find a better bound for the distortion.

Problem 18 (Piotr Indyk) It is known that there are $(1+\varepsilon)$-spanners with $O(n/\varepsilon^d)$ edges for any set of n points in $\mathbb{R}^d$. It is also known that any n-point metric has a c-spanner of size $n^{1+O(1)/c}$.

Question: Is there a, say 100-spanner, with $O(n \operatorname{polylog} n)$ edges for any n-point set in L_2 (with or without Steiner points)?

Problem 19 (Piotr Indyk) Consider the distortion of embedding n-point tree metrics into the Euclidean plane. Stars give a lower bound of $\Omega(\sqrt{n})$ via a packing argument. There is an upper bound of $O(\sqrt{n})$ for unweighted trees, for unweighted outerplanar graphs, for weighted stars (and, more generally, for ultrametrics). For weighted trees only the trivial upper bound n is known.

Question: Is $O(\sqrt{n})$ the upper bound for weighted trees?

Problem 20 (Kirk Pruhs) Let (X, d) be a finite metric space, G a greedy path $O(1)$-spanner, S a spanning tree of G, and T a maximum spanning tree of X.

Question: Does $\text{weight}(S) = O(\text{weight}(T))$ hold?

Problem 21 (Anupam Gupta) It is known to be hard to embed a weighted tree optimally into the real line, and hard to approximate. There is a $\tilde{O}(n^{1/3})$-approximation algorithm for minimum-distortion embedding of unweighted trees into the line [BDG⁺05].

Question: Is it also NP-hard to embed *unweighted* trees optimally?

Problem 22 (Marc Benkert) Given a set S of points in the plane, find the shortest *strongly connected* Manhattan network $M(S)$ of S. An *empty* rectangle of S is an axis-parallel rectangle that is spanned by two points in S and does not contain any further points in S.

Question: Is $M(S)$ the union of all empty rectangles of S?

Problem 23 (Christian Knauer) One can find a $(1+O(1/\sqrt{\log n}))$-approximation of the minimum-dilation triangulation of the regular n-gon in $O(n\sqrt{\log n})$ time [KM05].

Question: Can the minimum-dilation triangulation of a convex n-gon be computed in polynomial time? Is the problem NP-hard?

Problem 24 (Giri Narasimhan) Let $t^\star$ be the smallest dilation that allows plane $t^\star$-spanners (without Steiner points) for any point set. The vertices of the unit square yield $t^\star \geq \sqrt{2}$. A slightly larger lower bound has been proven by Mulzer [Mul04]. On the other hand Chew has shown that $t^\star \leq 2$, and the conjecture is that the Delaunay triangulation has dilation $\pi/2$. (Allowing Steiner points, Ebbers-Baumann et al. [EBGK⁺05] have shown that $t^\star_{\text{Steiner}} \leq 1.57$.)

Task: Improve the bounds for $t^\star$!

3 List of Talks

(in chronological order)

presenter	*title*
Michiel Smid	Geometric spanner networks: A survey
Otfried Cheong	Computing a minimum-dilation spanning tree is NP-hard
Christian Knauer	Optimal edge deletion in polygonal cycles
Mohammad Farshi	Region-fault tolerant geometric spanners
Sergio Cabello	Multiple-source shortest paths in a genus-g graph
Anupam Gupta	Metric embeddings: A brief survey
Yuri Rabinovich	Hard metrics and Abelian Cayley graphs
Victor Chepoi	A constant factor approximation algorithm for fitting Robinson structures to distance matrices
Hubert Chan	Small hop-diameter sparse spanners for doubling metrics
Günter Rote	The geometric dilation for three points
Christian Sohler	A Fast PTAS for k-Means Clustering
Hans Burkhardt	Invariants for discrete structures – an extension of Haar integrals over transformation groups to Dirac delta-functions
Vincent Moulton	Embeddings, tight-spans and phylogenetic networks
Sándor Fekete	Geometric distance estimation for sensor networks and unit disk graphs
Daniel Huson	Survey talk: Application of Phylogenetic Networks in Evolutionary Studies
Hans-Jürgen Bandelt	Translating DNA data tables into quasi-median networks
Alexander Wolff	Computing 1-spanners – in Manhattan
Sergey Bereg	Rigid graphs and pseudo-triangulations
Matthew Katz	On two variants of the power assignment problem in radio networks
Jan Vahrenhold	Pruning dense spanners in the presence of hierarchical memory
Martin Zachariasen	Minimum-length two-connected networks
Meera Sitharam	Partial metric spaces, rigidity and geometric constraint decomposition
Mattias Andersson	Approximate distance oracles for graphs with dense clusters
Anastasios Sidiropoulos	Probabilistic embeddings of bounded genus graphs into planar graphs
Piotr Indyk	Approximation algorithms for minimum-distortion embeddings into low-dimensional spaces

4 List of abstracts

(in alphabetical order)

Approximate Distance Oracles for Graphs with Dense Clusters

Mattias Andersson (Lund University, S)

Let $H_1 = (V, E_1)$ be a collection of N pairwise vertex disjoint $O(1)$-spanners where the weight of an edge is equal to the Euclidean distance between its endpoints. Let $H_2 = (V, E_2)$ be the graph on V with M edges of non-negative weight. The union of the two graphs is denoted $G = (V, E_1 \cup E_2)$. We present a data structure of size $O(M^2 + n \log n)$ that answers $(1+\varepsilon)$-approximate shortest path queries in G in constant time, where $\varepsilon > 0$ is constant.

Keywords: Computational geometry, shortest paths, t-spanners, disjoint

Joint work of: Andersson, Mattias; Gudmundsson, Joachim; Levcopoulos, Christos

Full Paper: doi:10.1016/j.comgeo.2004.12.011

See also: *Computational Geometry: Theory & Applications*, 37(3):142–154, 2007.

Translating DNA data tables into quasi-median networks

Hans-Jürgen Bandelt (Universität Hamburg, D)

The entire k-dimensional Hamming space (network) of sequences over an alphabet with 5 letters is the host space of any data set of aligned DNA sequences of length k. To investigate combinatorial properties of such data, solving e.g. instances of the Steiner problem alias parsimony problem, only a small portion of the Hamming space needs to be considered. In fact, every DNA data table can be turned into its Ploščica dual, the quasi-median network, that faithfully represents the data. For (condensed) data tables the associated quasi-median network harbours all solutions of the Steiner problem relative to a given tree topology, i.e. all most parsimonious reconstructions for any tree imposed to connect the sampled sequences.

Structural features of this network can be computed directly from the data table. The key principle repeatedly used is that the quasi-median network is uniquely determined by the sub-tables for pairs of characters. The translation of a table into a network enhances the understanding of the properties of the data in regard to homoplasy (recurrent mutations) and potential artifacts. A pertinent example drawn from human mtDNA illustrates these points.

Keywords: Maximum parsimony; Strong compatibility; Quasi-median network; Filter analysis; Error detection

Joint work of: Bandelt, Hans-Jürgen; Dür, Arne

Full Paper: doi:10.1016/j.ympev.2006.07.013

See also: Molecular Phylogenetics and Evolution 42 (2007) 256-271

Rigid Graphs and Pseudo-triangulations

Sergey Bereg (University of Texas at Dallas, USA)

Graph embeddings into Euclidean spaces are related to the theory of rigid graphs. We discuss classical results from the theory of rigid graphs including body-and-bar and body-and-hinge structures, rigidity matroid, characterizations of minimally rigid graphs in the plane and Henneberg construction. We also talk about pseudo-triangulations as embeddings of rigid graphs in the plane.

Keywords: Graph embedding, rigid graph, pseudo-triangulation

Invariants for Discrete Structures – An Extension of Haar Integrals over Transformation Groups to Dirac Delta Functions

Hans Burkhardt (Universität Freiburg, D)

Due to the increasing interest in 3D models in various applications there is a growing need to support e.g. the automatic search or the classification in such databases. As the description of 3D objects is not canonical it is attractive to use invariants for their representation. We recently published a methodology to calculate invariants for continuous 3D objects defined in the real domain by integrating over the group of Euclidean motion with monomials of a local neighbourhood of voxels as kernel functions and we applied it successfully for the classification of scanned pollen in 3D. In this paper we are going to extend this idea to derive invariants from discrete structures, like polygons or 3D-meshes by summing over monomials of discrete features of local support. This novel result for a space-invariant description of discrete structures can be derived by extending Haar integrals over the Euclidean transformation group to Dirac delta functions. Examples will be given for the search within 2D-polygonal objects and for searching in protein databases. We see possible extensions of the methodology to general weighted graphs.

Keywords: Invariants, discrete structures, retrieval

See also: http://lmb.informatik.uni-freiburg.de/papers/

Multiple Source Shortest Paths in a Genus-g Graph

Sergio Cabello (University of Ljubljana, SLO)

We give an $O(g^2 n \log n)$ algorithm to represent the shortest path tree from all the vertices on a single specified face f in a genus g graph. From this representation, any query distance from a vertex in f can be obtained in $O(\log n)$ time. The algorithm uses a kinetic data structure, where the source of the tree iteratively moves across edges in f. In addition, we give applications using these shortest path trees to compute the shortest non-contractible cycle and the shortest non-separating cycle in a combinatorial surface in $O(g^3 n \log n)$ time.

Keywords: Shortestest path, topological graph, combinatorial surface, non-contractible, non-separating

Joint work of: Cabello, Sergio; Chambers, Erin W.

Small Hop-diameter Sparse Spanners for Doubling Metrics

Hubert Chan (CMU, Pittsburgh, USA)

Given a metric $M = (V, d)$, a graph $G = (V, E)$ is a t-spanner for M if every pair of nodes in V has a "short" path (i.e., of length at most t times their actual distance) between them in the spanner. Furthermore, this spanner has a *hop diameter* bounded by D if every such short path also uses at most D edges. We consider the problem of constructing sparse $(1+\varepsilon)$-spanners with small hop diameter for metrics of low doubling dimension.

We show that given any metric with constant doubling dimension k, and any $0 < \varepsilon < 1$, one can find a $(1+\varepsilon)$-spanner for the metric with nearly linear number of edges (i.e., only $O(n \log^* n + n\varepsilon^{-O(k)})$ edges) and a *constant* hop diameter, and also a $(1+\varepsilon)$-spanner with linear number of edges (i.e., only $n\varepsilon^{-O(k)}$ edges) which achieves a hop diameter that grows like the functional inverse of the Ackermann's function. Moreover, we prove that such tradeoffs between the number of edges and the hop diameter are asymptotically optimal.

Keywords: Spanners, doubling dimension, hop diameter

Joint work of: Chan, T-H. Hubert; Anupam, Gupta

Full Paper: http://www.cs.cmu.edu/~hubert/soda06.pdf

See also: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm. Pages: 70 - 78. 2006

Computing a Minimum-Dilation Spanning Tree is NP-hard

Otfried Cheong (KAIST, Daejeon, ROK)

Given a set S of n points in the plane, a minimum-dilation spanning tree of S is a tree with vertex set S of smallest possible dilation. We show that given a set S of n points and a dilation $\delta > 1$, it is NP-hard to determine whether a spanning tree of S with dilation at most δ exists.

Keywords: Minimum dilation, spanning tree, NP-hardness

Joint work of: Cheong, Otfried; Haverkort, Herman; Lee, Mira

See also: In Computing: The Australasian Theory Symposium (CATS), CRPIT Vol. 65, 2007.

A constant factor approximation algorithm for fitting Robinson structures to distance matrices

Victor Chepoi (Université de la Méditerrannée, Marseille, F)

The classical seriation problem consists in finding a permutation of the rows and the columns of the dissimilarity matrix d on a finite set X with the objective of revealing an underlying one-dimensional structure (d is a dissimilarity if $d(x, y) = d(y, x) \geq 0$ and $d(x, y) = 0$ if $x = y$). The basic idea is that small values should be concentrated around the main diagonal as closely as possible, whereas large values should fall as far from it as possible. This goal is best achieved by considering the so-called Robinson property: a dissimilarity matrix d_R on X is Robinsonian if its matrix can be symmetrically permuted so that its elements do not decrease when moving away from the main diagonal along any row or column.

If the dissimilarity matrix between the objects fails to satisfy the Robinson property, then we are lead to the problem of finding a matrix reordering which is as close as possible to a Robinsonian matrix. In the talk, we will describe a factor 16 approximation algorithm for the following NP-hard fitting problem: given a finite set X and a dissimilarity d on X, we wish to find a Robinsonian dissimilarity d_R on X minimizing the l_∞-error $||d - d_R||_\infty = \max_{x,y \in X} |d(x, y) - d_R(x, y)|$.

Keywords: Fitting, seriation problem, Robinson dissimilarity, approximation algorithm

Joint work of: Chepoi, Victor; Seston, Morgan

Full Paper: http://www.lif-sud.univ-mrs.fr/˜chepoi/robinson16approx.pdf

Region-Fault Tolerant Geometric Spanners

Mohammad Farshi (TU Eindhoven, NL)

We introduce the concept of region-fault tolerant spanners for planar point sets, and prove the existence of region-fault tolerant spanners of small size.

For a geometric graph G on a point set P and a region F, we define $G \ominus F$ to be what remains of G after the vertices and edges of G intersecting F have been removed. A $\mathcal{C}$*-fault tolerant* t*-spanner* is a geometric graph G on P such that for any convex region F, the graph $G \ominus F$ is a t-spanner for $G_c(P) \ominus F$, where $G_c(P)$ is the complete geometric graph on P. We prove that any set P of n points admits a $\mathcal{C}$-fault tolerant $(1+\varepsilon)$-spanner of size $O(n \log n)$, for any constant $\varepsilon > 0$; if adding Steiner points is allowed then the size of the spanner reduces to $O(n)$, and for several special cases we show how to obtain region-fault tolerant spanners of $O(n)$ size without using Steiner points.

We also consider *fault-tolerant geodesic* t*-spanners*: this is a variant where, for any disk D, the distance in $G \ominus D$ between any two points $u, v \in P \setminus D$ is at most t times the geodesic distance between u and v in $R^2 \setminus D$. We prove that for any P we can add $O(n)$ Steiner points to obtain a fault-tolerant geodesic $(1+\varepsilon)$-spanner of size $O(n)$.

Keywords: Fault-Tolerant, Geometric network, spanner

Joint work of: Abam, Mohammad Ali; de Berg, Mark; Farshi, Mohammad; Gudmundsson, Joachim

Full Paper: http://www.win.tue.nl/~mfarshi/SODA07.pdf

See also: The Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2007.

Geometric Distance Estimation for Sensor Networks and Unit Disk Graphs

Sándor Fekete (TU Braunschweig, D)

We present an approach to estimating distances in sensor networks. It works by counting common neighbors, high values indicating closeness.

Such distance estimates are needed in many self-localization algorithms. Other than many other approaches, ours does not rely on special equipment in the devices.

Keywords: Sensor networks, distance estimation, unit disk graphs.

Joint work of: Fekete, Sándor; Kröller, Alexander; Buschmann, Carsten; Fischer, Stefan

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2007/1028

Full Paper: http://maven.smith.edu/~streinu/FwCG/e-proc/kroller.pdf

See also: Proc. 16th Fall Workshop on Computational Geometry, 2006

Metric Embeddings: A Brief Survey

Anupam Gupta (CMU, Pittsburgh, USA)

We give a survey of some salient results, techniques and open problems in the algorithmic study of metric finite metric spaces. Some recent research directions are also discussed.

Keywords: Metric spaces, metric embeddings, approximation algorithms,

Application of Phylogenetic Networks in Evolutionary Studies

Daniel H. Huson (Universität Tübingen, D)

Phylogenetic trees are usually used to describe evolutionary relationships between different species. Phylogenetic networks are used when reticulate events such as hybridization, horizontal gene transfer or recombination is believed to play an important role. Phylogenetic networks come in two flavors, as implicit networks that are used to show different, incompatible evolutionary signals simutaneously, and explicit networks that explicitly describe evolution in terms of reticulate events. This survey talk presents split networks as an important example of implicit networks and reticulate networks as an important example of explicit networks. Some algorithmic details are discussed.

Keywords: Phylogenetic networks, split networks, evolution, graphs

Approximation algorithms for minimum distortion embeddings into low-dimensional spaces

Piotr Indyk (MIT, Cambridge, USA)

In this talk I will give an overview of known approximation algorithms for minimum distortion embeddings of metrics into low-dimensional spaces, such as the line and the plane.

On Two Variants of the Power Assignment Problem in Radio Networks

Matthew Katz (Ben Gurion University, IL)

We study the power assignment problem in radio networks, where each radio station can transmit in one of two possible power levels, corresponding to two ranges – short and long. We show that this problem is NP-hard, and present an $O(n^2)$-time assignment algorithm such that the number of transmitters that are assigned long range by the algorithm is at most (11/6) times the number of transmitters that are assigned long range by an optimal algorithm. We also present an (9/5)-approximation algorithm for this problem whose running time is considerably higher.

Next, we formulate and study the Minimum-Area Spanning Tree (MAST) problem: Given a set P of n points in the plane, find a spanning tree of P of minimum "area," where the area of a spanning tree T is the area of the union of the $n-1$ disks whose diameters are the edges in T. We prove that the Euclidean minimum spanning tree of P is a constant-factor approximation for MAST. We then apply this result to obtain constant-factor approximations for the Minimum-Area Range Assignment (MARA) problem, for the Minimum-Area Connected Disk Graph (MACDG) problem, and for the Minimum-Area Tour (MAT) problem. The first problem is a variant of the power assignment problem in radio networks, the second problem is a related natural problem, and the third problem is a variant of the traveling salesman problem.

Keywords: Radio networks, power assignment, approximation algorithms, minimum spanning tree, disk graph

Joint work of: Carmi, Paz; Katz, Matthew; Mitchell, Joseph

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2007/1027

Optimal edge deletion in polygonal cycles

Christian Knauer (FU Berlin, D)

We give an $O(n \log^3 n)$ randomized algorithm for detecting an edge e of a closed planar polygonal cycle P on n vertices whose removal results in a polygonal curve $P \setminus \{e\}$ of smallest possible dilation.

Keywords: Dilation, Polygonal cycle, Edge deletion

Embeddings, tight-spans and phylogenetic networks

Vincent Moulton (Univ. of East Anglia, Norwich, GB)

One of the main problems in phylogenetics is to find good approximations of metrics by weighted trees. As an aid to solving this problem, it could be tempting to consider optimal realizations of metrics – the guiding principle being that, the (necessarily unique) optimal realization of a tree metric is the weighted tree that realizes this metric. And, although optimal realizations of arbitrary metrics are not trees in general, but rather weighted networks, one could still hope to obtain a phylogenetically informative representation of a given metric, maybe even more informative than the best approximating tree.

However, optimal realizations are not only difficult to compute, but they may also be non-unique. In this talk we discuss a possible way out of this dilemma: hereditarily optimal, or h-optimal, realizations. These are essentially unique, and can be described in a rather explicit way. We shall recall what a h-optimal realization of a metric is, how it is related to the tight-span of that metric, and describe under what conditions it coincides with the 1-skeleton of the tight-span.

As a consequence, we will show that h-optimal realizations can be computed in a straight-forward way for a large class of phylogentically relevant metrics. In addition, we give a family of metrics that provides some counter-examples to previously published conjectures concerning the relationship between optimal realizations, h-optimal realizations and the tight-span.

Keywords: Phylogenetic tree, phylogenetric network, optimal realization, tight-span

Hard metrics and Abelian Cayley Graphs

Yuri Rabinovich (Haifa University, IL)

Hard metrics are the class of extremal metrics with respect to embedding into Euclidean Spaces: their distortion is as bad as it possibly gets, which is $\Omega(\log n)$. Besides being very interesting objects akin to expanders and good codes, with rich structure of independent interest, such metrics are important for obtaining lower bounds in Combinatorial Optimization, e.g., on the value of Min-Cut/MaxFlow ratio for multicommodity flows.

For more than a decade, a single family of hard metrics was known. Recently, a different such family was found, causing a considerable excitement among the researchers in the area.

In this talk we present another general and clear construction of hard metrics.

Keywords: Abelian Groups, Cayley Graph, Metric Distortion

Joint work of: Rabinovich, Yuri; Newman, Ilan

The geometric dilation for 3 points

Günter Rote (Freie Universität Berlin, D)

For 3 given points A, B, C in the plane, we want to find a connecting network N for which the *geometric dilation*

$$\min_{x,y \in N, x \neq y} \frac{d_N(x,y)}{\|x-y\|}$$

is minimized, where $d_N(x,y)$ is the distance between x and y in the network.

The optimal network has two possible structures:

(a) a tree with an additional vertex, the so-called Fermat-Torricelli point, where three straight edges meet at an angle of 120°.

This network has geometric dilation $\sqrt{4/3} \approx 1.15$.

(b) a path formed from the straight segment AB, another straight segment BP of the same length, forming the legs of an isosceles triangle with base angle

$\tau \leq 30°$, and an arc PC of a logarithmic spiral around A forming the constant angle τ with the radial lines.

The geometric dilation is $1/\cos\tau$.

Keywords: Small network, logarithmic spiral

Joint work of: Rote, Günter; Klein, Rolf; Ebbers-Baumann, Annette; Knauer, Christian

Probabilistic Embeddings of Bounded Genus Graphs Into Planar Graphs

Anastasios Sidiropoulos (MIT, Cambridge, USA)

We show that every graph of bounded genus can be $O(1)$-probabilistically approximated by a distribution over planar graphs. For a graph of genus g, if an embedding of the graph in a surface of genus g is given, then the distribution can be computed efficiently.

Using this probabilistic approximation we can reduce a large class of graph optimization problems over graphs of bounded genus, to corresponding problems over planar graphs.

Moreover, we obtain a strengthening of a conjecture of [GNRS] about embedding planar graphs into ℓ_1.

Keywords: Embeddings, bounded genus graphs, planar graphs

Joint work of: Indyk, Piotr; Sidiropoulos, Anastasios

Partial Metric Spaces, Rigidity and Geometric Constraint Decomposition

Meera Sitharam (University of Florida, USA)

Geometric Constraint Systems (GCSs) arise in numerous applications from mechanical computer aided design to sensor network location to molecular modeling.

Algebraically solving such systems corresponds to finding an embedding of a partially specified metric space in (typically, but not always) Euclidean space of fixed dimension (typically 2 or 3).

Distortion within specified tolerances is often allowed. "Good" decompositions of GCSs are crucial not just for efficient solving and updates, but also for describing and sampling the entire set of solutions or embeddings etc.

Combinatorial rigidity theory studies graph properties of partial metric spaces that generically result in finitely many (or unique) isometric embeddings into (typically, but not always) Euclidean space of fixed dimension.

This talk will show the close connection between obtaining a "good" decomposition of a GCS and the combinatorial characterization of generic rigidity and related notions such as d-realizability - the source of long sought-after open problems.

Specifically, our decomposition approach results (a) in 2 tractable, approximate characterizations of generic rigidity, for 3D distance GCSs and 2D angle GCSs; (b) in a close connection between d-realizability and the descriptive complexity of the set of embeddings.

Geometric Spanner Networks: A Survey

Michiel Smid (Carleton Univ., Ottawa, CDN)

Given a set S of n points in d-dimensional space and a real number $t > 1$, a t-spanner for S is a graph with vertex set S, such that any two points p and q of S are connected by a path whose length is at most t times the Euclidean distance between p and q. I will present an overview of the three main constructions for computing t-spanners:

1. The Θ-graph and its variants.
2. The WSPD-spanner, which is based on the well-separated pair decomposition.
3. The path-greedy spanner.

Techniques used in the path-greedy spanner can be used to design a data structure of size $O(n \log n)$, such that an approximate shortest path between any two query points can be computed in $O(1)$ time.

Keywords: Computational geometry, spanners, shortest paths

A Fast PTAS for k-Means Clustering

Christian Sohler (Universität Paderborn, D)

Given a point set $P \subseteq \mathbb{R}^d$ the k-means clustering problem is to find a set $C = \{c_1, \ldots, c_k\}$ of k points and a partition of P into k clusters $C_1, \ldots, C_k$ such that the sum of squared errors $\sum_{i=1}^{k} \sum_{p \in C_i} \|p - c_i\|_2^2$ is minimized. For given centers this costfunction is minimized by assigning points to the nearest center.

The k-means cost function is probably the most widely used cost function in the area of clustering.

In this paper we show that every unweighted point set P has a weak (ϵ, k)-coreset of size $\text{poly}(k, 1/\epsilon)$ for the k-means clustering problem, i.e. its size is *independent* of the cardinality $|P|$ of the point set and the dimension d of the Euclidean space $\mathbb{R}^d$. A weak coreset is a weighted set $C \subseteq P$ together with a set X such that X contains a $(1 + \epsilon)$-approximation for the optimal cluster centers from P and for every set of k centers from X the cost of the centers for C is a $(1 \pm \epsilon)$-approximation of the cost for P.

We apply our weak coreset to obtain a PTAS for the k-means clustering problem with running time $O(nkd + d \cdot \text{poly}(k/\epsilon) + 2^{\widetilde{O}(k/\epsilon)})$.

Keywords: Clustering, approximation algorithms, randomized algorithms, computational geometry, coresets

Joint work of: Feldman, Dan; Monemizadeh, Morteza; Sohler, Christian

Pruning Dense Spanners in the Presence of Hierarchical Memory

Jan Vahrenhold (Universität Münster, D)

Given a geometric graph $G = (S, E)$ in R^d with constant dilation t, and a positive constant ϵ, we show how to construct a $(1 + \varepsilon)$-spanner of G with $\mathcal{O}(|S|)$ edges using $\mathcal{O}(\text{sort}(|E|))$ I/O operations.

We also dicsuss how the resulting algorithm can be made cache-oblivious in the model of Frigo *et al.* (1999).

As part of this process, we observe that the algorithm of Govindarajan *et al.* (2000) can be modifed to produce a linear-size well-separated pair decomposition using $\mathcal{O}(\text{sort}(|E|))$ memory transfers.

This modification allows us to obtain an $\mathcal{O}(\text{sort}(|E|))$ bound for the overall algorithm also in the cache-oblivious setting.

Keywords: I/O-efficient, cache-oblivious, well-separated pair decomposition

Joint work of: Gieseke, Fabian; Gudmundsson, Joachim; Vahrenhold, Jan

Constructing 1-Spanners - in Manhattan

Alexander Wolff (Universität Karlsruhe, D)

Given a set of points in the plane and a constant $t \geq 1$, a Euclidean *t-spanner* is a network in which, for any pair of points, the ratio of the network distance and the Euclidean distance of the two points is at most t. Such networks have applications in transportation or communication network design and have been studied extensively.

In this paper we study 1-spanners under the Manhattan (or L_1-) metric. Such networks are called *Manhattan networks*. A Manhattan network for a set of points is a set of axis-parallel line segments whose union contains an x- and y-monotone path for each pair of points. It is not known whether it is NP-hard to compute minimum Manhattan networks (MMN), i.e. Manhattan networks of minimum total length. In this paper we present an approximation algorithm for this problem. Given a set P of n points, our algorithm computes in $O(n \log n)$

time and linear space a Manhattan network for P whose length is at most 3 times the length of an MMN of P.

We also establish a mixed-integer programming formulation for the MMN problem. With its help we extensively investigate the performance of our factor-3 approximation algorithm on random point sets.

Keywords: Manhattan metric, spanner, geometric network

Joint work of: Benkert, Marc; Wolff, Alexander; Widmann, Florian; Shirabe, Takeshi

Full Paper: doi:10.1016/j.comgeo.2005.09.004

See also: Marc Benkert, Alexander Wolff, Florian Widmann, and Takeshi Shirabe: The Minimum Manhattan Network Problem: Approximations and Exact Solution, Computational Geometry: Theory and Applications, p. 188-208, 35 (3) 2006

Minimum-Length Two-Connected Networks

Martin Zachariasen (University of Copenhagen, DK)

We consider the problem of constructing a shortest Euclidean 2-connected Steiner network (SMN) for a set of terminals. This problem has natural applications in the design of survivable communication networks. A SMN decomposes into components that are full Steiner trees. We use the notion of reduced block-bridge trees of Luebke to show that no component in a SMN spans more than approximately one-third of the terminals. Algorithmic implications of this result are discussed.

Keywords: Survivable networks, Euclidean 2-connected Steiner networks, full Steiner trees

Joint work of: Hvam, Kenneth; Reinhardt, Line; Winter, Pawel; Zachariasen, Martin

5 Conclusion

We are grateful to the Dagstuhl research center for hosting this successful seminar and for providing such a stimulating and creative environment.

References

[ADG+03] Alexandr Andoni, Michel Marie Deza, Anupam Gupta, Piotr Indyk, and Sofya Raskhodnikova. Lower bounds for embedding edit distance into normed spaces. In *Proc. 14th ACM-SIAM Sympos. Discrete Algorithms (SODA'03)*, pages 523–526, 2003. 5

[Alt88] Ingo Althöfer. On optimal realizations of finite metric spaces by graphs. *Discrete Comput. Geom.*, 3(2):103–122, 1988. 3

[AP03] Noga Alon and Pavel Pudlák. Equilateral sets in l_p^n. *Geometric and Functional Analysis*, 13(3):467–482, 2003. 4

[BDG+05] Mihai Bădoiu, Kedar Dhamdhere, Anupam Gupta, Yuri Rabinovich, Harald Räcke, R. Ravi, and Anastasios Sidiropoulos. Approximation algorithms for embeddings into low-dimensional spaces. In *Proc. 16th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA'05)*, pages 119–128, 2005. 7

[CHL07] Otfried Cheong, Herman Haverkort, and Mira Lee. Computing a minimum-dilation spanning tree is NP-hard. In *Proc. Computing: the Australasian Theory Symposium (CATS'07)*, volume 65 of *CRPIT*, pages 15–24. Australian Comput. Soc., 2007. 5

[CK07] Paz Carmi and Matthew J. Katz. Power assignment in radio networks with two power levels. *Algorithmica*, 47(2):183–201, 2007. 6

[EBGK+05] Annette Ebbers-Baumann, Ansgar Grüne, Marek Karpinski, Rolf Klein, Christian Knauer, and Andrzej Lingas. Embedding point sets into plane graphs of small dilation. In *Proc. 16th Annu. Intern. Sympos. Algorithms Comput. (ISAAC'05)*, volume 3827 of *Lecture Notes Comput. Sci.*, pages 5–16. Springer-Verlag, 2005. 7

[FGG05] Mohammad Farshi, Panos Giannopoulos, and Joachim Gudmundsson. Finding the best shortcut in a geometric network. In *Proc. 21st Annu. ACM Symp. Comput. Geom. (SoCG'05)*, pages 327–335, 2005. 5

[GKKS07] Joachim Gudmundsson, Oliver Klein, Christian Knauer, and Michiel Smid. Small manhattan networks and algorithms for the earth mover's distance. In *Proc. 23rd European Workshop Comput. Geom. (EWCG'07)*, pages 174–177, Graz, 2007. 3

[GNRS99] Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Alistair Sinclair. Cuts, trees and ℓ_1-embeddings of graphs. In *Proc. 40th Annu. IEEE Sympos. Found. Comput. Sci. (FOCS'99)*, pages 399–408, 1999. 5

[Guy83] Richard K. Guy. Unsolved problems: An olla-podrida of open problems, often oddly posed. *Amer. Math. Monthly*, 90:196–200, 1983. 4

[KM05] Christian Knauer and Wolfgang Mulzer. Minimum dilation triangulations. Technical Report B-05-06, Institut für Informatik, Freie Universität Berlin, April 2005. 7

[Mul04] Wolfgang Mulzer. Umwegprobleme in Graphen und Polygonen. Master's thesis, Institut für Informatik, Freie Universität Berlin, 2004. 7

[Rao99] Satish B. Rao. Small distortion and volume preserving embeddings for planar and Euclidean metrics. In *Proc. 15th Annu. ACM Sympos. Comput. Geom. (SoCG'99)*, pages 300–306, 1999. 5

[Swa04] Konrad J. Swanepoel. A problem of Kusner on equilateral sets. *Archiv der Mathematik*, 83(2), 2004. 4