Report from Dagstuhl Seminar 11111

Computational Geometry

Edited by

Pankaj Kumar Agarwal[1], Kurt Mehlhorn[2], and Monique Teillaud[3]

1 **Duke University, US,** pankaj@cs.duke.edu
2 **MPI für Informatik - Saarbrücken, DE,** mehlhorn@mpi-inf.mpg.de
3 **INRIA Sophia Antipolis - Méditerranée, FR,** Monique.Teillaud@inria.fr

Abstract

This report documents the outcomes of Dagstuhl Seminar 11111 "Computational Geometry". The Seminar gathered fifty-three senior and younger researchers from various countries in the unique atmosphere offered by Schloss Dagstuhl. Abstracts of talks are collected in this report as well as a list of open problems.

Seminar 13.–18. March, 2011 – www.dagstuhl.de/11111
1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity (Geometrical problems and computations), G.4 Mathematical Software (Algorithm design and Analysis)
Keywords and phrases Algorithms, geometry, combinatorics, topology, theory, applications, implementation
Digital Object Identifier 10.4230/DagRep.1.3.19
Edited in cooperation with Saurabh Ray

1 Executive Summary

Pankaj Kumar Agarwal
Kurt Mehlhorn
Monique Teillaud

Computational Geometry and its Evolution

The field of computational geometry is concerned with the design, analysis, and implementation of algorithms for geometric problems, which arise in a wide range of areas, including computer graphics, CAD, robotics computer vision, image processing, spatial databases, GIS, molecular biology, and sensor networks. Since the mid 1980s, computational geometry has arisen as an independent field, with its own international conferences and journals.

In the early years mostly theoretical foundations of geometric algorithms were laid and fundamental research remains an important issue in the field. Meanwhile, as the field matured, researchers have started paying close attention to applications and implementations of geometric algorithms. Several software libraries for geometric computation (e.g. LEDA, CGAL, CORE) have been developed. Remarkably, this emphasis on applications and implementations has emerged from the originally theoretically oriented computational geometry community itself, so many researchers are concerned now with theoretical foundations as well as implementations.

Computational Geometry, *Dagstuhl Reports*, Vol. 1, Issue 3, pp. 19–41
Editors: Pankaj Kumar Agarwal, Kurt Mehlhorn, and Monique Teillaud
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Seminar Topics

The emphasis of the seminar was on presenting the recent developments in the field as well as identifying new challenges. We have identified a few broad topics, listed below, that cover both theoretical and practical issues in computational geometry and that we believe are some of the most interesting subareas in the field.

- *Theoretical foundations* of computational geometry lie in combinatorial geometry and its algorithmic aspects. They are of an enduring relevance for the field, particularly the design and the analysis of efficient algorithms require deep theoretical insights.
- Various *applications* such as robotics, GIS, or CAD lead to interesting variants of the *classical topics* originally investigated, including convex hulls, Voronoi diagrams and Delaunay triangulations, and geometric data structures. For example, Voronoi diagrams and nearest-neighbor data structures under various metrics have turned out to be useful for many applications and are being investigated intensively.
- Because of applications in molecular biology, computer vision, geometric databases, *shape analysis* has become an important topic. Not only it It raises many interesting geometric questions ranging from modeling and reconstruction of surfaces to shape similarity and classification, but it has also led to the emergence of the so-called field *computational topology*.
- In many applications the data lies in very high dimensional space and typical geometric algorithms suffer from the curse of dimensionality. This has led to extensive work on dimension-reduction and embedding techniques.
- Massive geometric data sets are being generated by networks of sensors at unprecedented spatial and temporal scale. How to store, analyze, query, and visualize them has raised several algorithmic challenges. New computational models have been proposed to meet these challenges, e.g., streaming model, communication-efficient algorithms, and maintaining geometric summaries.
- *Implementation issues* have become an integral part of the research in computational geometry. Besides general software design questions especially *robustness* of geometric algorithms is important. Several methods have been suggested and investigated to make geometric algorithms numerically robust while keeping them efficient, which lead to interaction with the field of computer algebra, numerical analysis, and topology.

Participants

53 researchers from various countries and continents attended the meeting. This high number shows the strong interest of the community for this event. The feedback from participants was very positive.

Dagstuhl seminars on computational geometry have been organized since 1990, lately in a two year rhythm. They have been extremely successful both in disseminating the knowledge and identifying new research thrusts. Many major results in computational geometry were first presented in Dagstuhl seminars, and interactions among the participants at these seminars have led to numerous new results in the field. These seminars have also played an important role in bringing researchers together and fostering collaboration. They have arguably been the most influential meetings in the field of computational geometry.

A session of this Seminar was dedicated to our dear friend Hazel Everett, deceased on July 20th, 2010.

The place itself is a great strength of the Seminar. Dagstuhl allows people to really meet and socialize, providing them with a wonderful atmosphere of a unique closed and pleasant environment, which is highly beneficial to interactions.

Therefore, we warmly thank the scientific, administrative and technical staff at Schloss Dagstuhl!

2 Table of Contents

3 Overview of Talks

3.1 Computing the depth of an arrangement of axes-parallel rectangles in parallel

Helmut Alt (FU Berlin, DE)

Joint work of Alt, Helmut; Hagerup, Torben; Scharf, Ludmila

We consider the problem of determining the depth of an arrangement of axes parallel rectangles in the plane, i.e., the highest number of rectangles intersecting in one point. Sequentially, standard procedures have quadratic runtime since the size of the arrangement is quadratic in the number of rectangles. However, $O(n \log n)$ algorithms are known for this problem.

We design *parallel* algorithms for this problem. We consider a structure related to interval trees which is traversed top-down level by level propagating and adjusting the information belonging to the interval associated to a node from parent to child. We obtain a parallel runtime of $O(log^2 n)$ with a total of O(n log n) operations or a parallel runtime of $O(logn)$ with $O(n^(1 + \varepsilon))$ operations for any constant $\varepsilon > 0$.

3.2 Memory-constrained algorithms

Tetsuo Asano (JAIST – Nomi, JP)

Main reference T. Asano, W. Mulzer, and Y. Wang, "Constant-Work-Space Algorithm for a Shortest Path in a Simple Polygon," Invited talk, Proc. 4th International Workshop on Algorithms and Computation, WALCOM, Dhaka, Bangladesh, 2010, LNCS 5942, pp.9–20.s
URL http://dx.doi.org/10.1007/978-3-642-11440-3_2

In this talk I will introduce algorithms with limited work space, which are desired for applications to highly functional hardware such as scanners, digital cameras, and Android cellular phones. One extreme set of memory- constrained algorithms have been studied under the name of log-space algorithms which use only $O(\log n)$ bits for their work space. This talk starts with a simple example of a memory-constrained algorithms based on a general paradigm for designing such algorithms. Then, it is extended to a problem on image processing, where work space of square root of an input size is more realistic. A practically efficient algorithm with work space of $O(\sqrt{n} \log n)$ bits is given for a problem of extracting an interesting region in a given color image.

3.3 Rips complexes for Shape Reconstruction

Dominique Attali (GRIPSA Lab – Saint Martin d'Hères, FR)

Joint work of Attali, Dominique; Lieutier, André; Salinas, David
Main reference D. Attali, A. Lieutier and D. Salinas, "Vietoris-Rips complexes also provide topologically correct reconstructions of sampled shapes," Proc. 27th Ann. Sympos. Comput. Geom., Paris, France, June 13–15 2011
URL http://hal.archives-ouvertes.fr/hal-00579864/en/

We associate with each compact set X of $\mathbb{R}^n$ two real-valued functions c_X and h_X defined on $\mathbb{R}^+$ which provide two measures of how much1 the set X fails to be convex at a given scale. First, we show that, when P is a finite point set, an upper bound on $c_P(t)$ entails that the Rips complex of P at scale r collapses to the Čech complex of P at scale r for some suitable values of the parameters t and r. Second, we prove that, when P samples a compact set X, an upper bound on h_X over some interval guarantees a topologically correct reconstruction of the shape X either with a Čech complex of P or with a Rips complex of P. Regarding the reconstruction with Čech complexes, our work compares well with previous approaches when X is a smooth set and surprisingly enough, even improve constants when X has a positive μ-reach. Most importantly, our work shows that Rips complexes can also be used to provide topologically correct reconstruction of shapes. This may be of some computational interest in high dimension.

3.4 The Height of a Homotopy

Erin Moriarty Wolf Chambers (St. Louis University, US)

In our paper in CCCG 2009, we examined the problem of computing the height of a homotopy. This can be phrased as a very combinatorial problem, and indeed has been studied in at least one very different context, submodular percolation. We proved several properties of homotopies that obtain the minimum height value, but were unable to completely characterize or compute the minimum height homotopy. More recently, new work has been completed to compute an $O(\log n)$ approximation, but again no exact algorithms or hardness results are known. We will survey known results and techniques for this problem.

3.5 A Generalization of Kakeya's Problem

Otfried Cheong (KAIST – Daejeon, KR)

Given a not necessarily finite family F of line segments, we show that the smallest-area convex figure P such that every segment in F can be translated to lie in P can always be chosen to be a triangle.

This generalizes the result by Pal from 1921 for the case where F contains a unit-length segment of every possible orientation.

Our result can be rephrased as follows: Given a convex figure P, there is always a triangle T of area at most the area of P such that for every possible direction, the width of T in that direction is not less than the width of P in that direction.

We also given an algorithm that computes the smallest-area triangle as above when the input F is a set of n line segments in time $O(n \log n)$.

3.6 Star Trek Replicators via Staged Assembly

Erik D. Demaine (MIT – Cambridge, US)

Tile self-assembly is an intriguing approach to manufacturing desired shapes with nano-scale feature size. A recent direction in this theory allows the use of multiple stages—operations performed by the experimenter, such as mixing two self-assembling systems together. This flexibility transforms the experimenter from a passive entity into a parallel algorithm, and vastly reduces the number of distinct parts required to construct a desired shape, possibly making the systems practical to build. The staged-assembly perspective also enables the possibility of additional operations, such as adding an enzyme that destroys all tiles with a special label. By enabling destruction in addition to the usual construction, we can perform tasks impossible in a traditional self-assembly system, such as replicating many copies of a given object's shape, without knowing anything about that shape, and building an efficient nano computer.

3.7 The Effect of Noise on the Number of Extreme Points

Olivier Devillers (INRIA Sophia Antipolis, FR)

Joint work of Attali, Dominique; Devillers, Olivier; Goaoc, Xavier
Main reference D. Attali, O. Devillers, X. Goaoc, "The Effect of Noise on the Number of Extreme Points", Research Report 7134, INRIA, 2009.
URL http://hal.inria.fr/inria-00438409/

Assume that Y is a noisy version of a point set X in convex position. How many vertices does the convex hull of Y have, that is, what is the number of extreme points of Y?

We consider the case where X is an (ϵ, κ)-sample of a sphere in $\mathbb{R}^d$ and the noise is random and uniform: Y is obtained by replacing each point $x \in X$ by a point chosen uniformly at random in some region $S(x)$ of size δ around x. We give upper and lower bounds on the expected number of extreme points in Y when $S(x)$ is a ball (in arbitrary dimension) or an axis-parallel square (in the plane). Our bounds depend on the size n of X and δ, and are tight up to a polylogarithmic factor. These results naturally extend in various directions (more general point sets, other regions $S(x)\ldots$).

We also present experimental results, showing that our bounds for random noise provide good estimators of the behavior of snap-rounding, that is when Y is obtained by rounding each point of X to the nearest point on a grid of step δ.

3.8 Improved Bound for the Union of Fat Triangles

Esther Ezra (New York University, US)

Joint work of Ezra, Esther; Aronov, Boris; Sharir, Micha
Main reference E. Ezra, B. Aronov, M. Sharir, "Improved Bound for the Union of Fat Triangles," Proc. 22nd Annual ACM-SIAM Symp. on Discrete Algorithms (SODA'11), to appear.
URL http://www.cims.nyu.edu/ esther/Publications/fatri.pdf

We show that, for any fixed $\delta > 0$, the combinatorial complexity of the union of n triangles in the plane, each of whose angles is at least δ, is $O(n2^{\alpha(n)} \log^* n)$, with the constant of proportionality depending on δ. This considerably improves the twenty-year-old bound $O(n \log \log n)$, due to Matousek etal.

3.9 Exact Solutions and Bounds for General Art Gallery Problems

Sandor Fekete (TU Braunschweig, DE)

Joint work of Fekete, Sándor; Kröller, Alexander; Schmidt, Christiane; Kamphans, Tom; Baumgartner, Tobias
Main reference Proceedings of the SIAM-ACM Workshop on Algorithm Engineering and Experiments (ALENEX'10), pp. 11–22.
URL http://www.siam.org/proceedings/alenex/2010/alx10_002_baumgartnert.pdf

The classical Art Gallery Problem asks for the minimum number of guards that achieve visibility coverage of a given polygon. This problem is known to be NP-hard, even for very restricted and discrete special cases. For the general problem (in which both the set of possible guard positions and the point set to be guarded are uncountable), neither constant-factor approximation algorithms nor exact solution methods are known. We present a primal-dual algorithm based on linear programming that provides lower bounds on the necessary number of guards in every step and (in case of convergence and integrality) ends with an optimal solution. We describe our implementation and give results for an assortment of polygons, including non-orthogonal polygons with holes.

3.10 Range Queries in Distributed Networks

Jie Gao (SUNY – Stony Brook, US)

Joint work of Gao, Jie; Sarkar, Rik

Consider mobile targets moving in a plane and their movements being monitored by a network such as a field of sensors. We develop distributed algorithms for in-network tracking and range queries for aggregated data (for example returning the number of targets within any user given region). Our scheme stores the target detection information locally in the network, and answers a query by examining the perimeter of the given range. The cost of updating data about mobile targets is proportional to the target displacement.

The key insight is to maintain in the sensor network a function with respect to the target detection data on the graph edges that is a differential one-form such that the integral of this one-form along any closed curve C gives the integral within the region bounded by C.

3.11 Intersection patterns of convex sets

Xavier Goaoc (INRIA Lorraine, FR)

Joint work of Colin de Verdière, Éric; Ginot, Grégory; Goaoc, Xavier
Main reference Eric Colin de Verdiere, Gregory Ginot, Xavier Goaoc, "Helly numbers of acyclic families," arXiv:1101.6006v2
URL http://arxiv.org/abs/1101.6006

The Helly number of a family of sets (with empty intersection) is the size of its largest inclusion-wise minimal sub-family with empty intersection. We show how techniques from homology theory lead to fairly general conditions under which the Helly number of a family can be bounded.

Our typical result is along the following lines: "if F is a family of sets in $\mathbb{R}^d$ such that the intersection of any subfamily has at most r connected components, each of which is a homology cell then the Helly number of F is at most $r(d+1)$". This result can be generalized so as to imply, in a unified way, bounds on Helly numbers in geometric transversal theory that had previously been studied by ad hoc techniques.

3.12 On Path Quality in Sampling-Based Motion Planning

Dan Halperin (Tel Aviv University, IL)

Joint work of Angela Enosh, Angela; Halperin, Dan; Nechsushtan, Oren; Raveh, Barak
Main reference (1) Raveh-Enosh-Halperin, A Little More, a Lot Better: Improving Path Quality by a Path-Merging Algorithm, IEEE Trans. on Robotics, 27/2, 2011, pp 365–371.
(2) Nechushtan-Raveh-Halperin, Sampling-Diagram Automata: A Tool for Analyzing Path Quality in Tree Planners, WAFR 2010.
URL http://acg.cs.tau.ac.il/projects

Sampling-based motion planners are a central tool for solving motion-planning problems in a variety of domains, but the theoretical understanding of their behavior remains limited, in particular with respect to the quality of the paths they generate (in terms of path length, clearance, etc.). We prove, for a simple family of obstacle settings, that the popular dual-tree planner Bi-RRT may produce low-quality paths that are arbitrarily worse than optimal with modest but significant probability, and overlook higher-quality paths even when such paths are easy to produce. At the core of our analysis are probabilistic automata designed to reach an accepting state when a path of significantly low quality has been generated.

Complementary experiments suggest that our theoretical bounds are conservative and could be further improved. We also present a method to improve path quality by merging an arbitrary number of input motion paths into a hybrid output path of superior quality, for a broad and general formulation of path quality. Our approach is based on the observation that the quality of certain sub-paths within each solution may be higher than the quality of the entire path.

3.13 Algorithms for Persistent Homology

Michael Kerber (IST Austria – Klosterneuburg, AT)

Joint work of Kerber, Michael; Chen, Chao

Persistent homology is a quickly-growing area of research in the analysis of topological spaces. One reason for its success is the existence of an efficient algorithm to solve the problem using Gaussian elimination. I will present two recent results in this context. First, I show a simple optimization technique of the default algorithm that avoids column operations on roughly half of the columns.

This yields both significant practical improvements, and provides new insights on the complexity of persistence for certain special cases.

Second, I will present a divide-and-conquer approach to compute persistence based on rank computations of submatrices instead of Gaussian elimination.

The algorithm only outputs homology classes with persistence larger than a given threshold, and permits an output-sensitive complexity analysis. In particular, using a Monte-Carlo algorithm for rank computation and assuming that the number of returned classes is logarithmic in the input size, an quadratic algorithm for persistence computation is achieved, modulo logarithmic factors.

3.14 Polygonal paths of bounded curvature

David G. Kirkpatrick (University of British Columbia – Vancouver, CA)

Joint work of Kirkpatrick, David G.; Polishchuk, Valentin

In one of the seminal papers [Dubins57] in non-holonomic motion planning, L.E. Dubins developed a strong characterization of minimum length curves of bounded curvature joining fixed initial and final configurations (specified by position and direction): in the absence of obstacles, such paths consist of two (possibly degenerate) circular arcs joined by either a straight line segment or a third circular arc. Dubins' original proof uses advanced calculus; subsequently the same result was reproved using control-theoretic techniques.

We revisit bounded-curvature in the context of polygonal paths (paths consisting of a sequence of straight line segments) and formulate a natural notion of bounded-curvature for such paths. While polygonal paths clearly violate curvature bounds in sufficiently small neighbourhoods, our notion still manages to capture the less restrictive constraint that they "do not turn too sharply". We present an elementary and purely geometric proof of a characterization result, analogous to that of Dubins, for minimum length polygonal paths satisfying our bounded curvature property. This not only provides a discrete analogue of continuous motion of bounded curvature, but it also gives a fundamentally new proof of Dubins' original result as a limiting case of our constructions.

3.15 VC Dimension of Visibility

Rolf Klein (Universität Bonn, DE)

Main reference A. Gilbers and R. Klein, "A New Upper Bound For the VC-Dimension of Visibility Regions," Proc. 27th ACM Symp. on Computational Geometry (SoCG'11), pp.380–386.
URL http://doi.acm.org/10.1145/1998196.1998259

We lower the upper bound for the VC dimension of visibility in simple polygons from 23 to 14.

3.16 The complexity of Ham-Sandwich cuts in high dimension

Christian Knauer (Universität Bayreuth, DE)

We study a canonical decision problem arising from the Ham-Sandwich cut theorem. We show it to be W[1]-hard (and NP-hard) if the dimension is part of the input. This is done by an fpt-reduction (which is actually a ptime-reduction) from the d-SUM problem. Our reduction also implies that the problem cannot be solved in time $n^{o(d)}$ unless the Exponential-Time Hypothesis (ETH) is false.

3.17 On the topology of plane algebraic curves

Sylvain Lazard (INRIA - Nancy, FR)

Joint work of Lazard, S.; Bouzidi, Y.; Cheng, J.; Penaranda, L.; Pouget, M.; Rouillier, F.; Tsigaridas, E.

I will present some recent results on the computation of the topology of planar algebraic curves. Our approach is based on a rectangular decomposition of the plane, instead of the standard CAD, which has the advantage of being oblivious to degenerate configurations. It is also based on a new algorithm for computing the critical points of the input curve. We first decompose the corresponding system into subsystems according to the number of roots (counted with multiplicities) in vertical lines, as presented by Gonzalez-Vega and Necula in 2002. We then show how these systems can be efficiently solved by computing their lexicographic Grobner bases and Rational Univariate Representations. We also show how this approach can be performed using modular arithmetic, while remaining deterministic. We finally demonstrate that our approach yields a substantial gain of a factor between 1 to 200 on curves of degree up 40 compared to state-of-the-art implementation.

3.18 The Phsyarum Computer

Kurt Mehlhorn (MPI für Informatik - Saarbrücken, DE)

Physarum is a slime mold. It was observed over the past 10 years that the mold is able to solve shortest path problems and to construct good Steiner networks (Nakagaki-Yamada-Toth, Tero-Takagi-et al). In a nutshell, the shortest path experiment is as follows: A maze is built and the mold is made to cover the entire maze. Food is then provided at two positions s and t and the evolution of the slime is observed. Over time, the slime retracts to the shortest s-t-path.

A mathematical model of the slime's dynamic behaviour was proposed by Tero-Kobayashi-Nakagaki.

Extensive computer simulations confirm the experimental findings; the slime retracts to the shortest path. We (joint work with Vincenzo Bonifaci, Girish Varma) have recently proved convergence.

I will start with a video showing the mold in action. Then I review the mathematical model, explain the computer simulation, and run some computer experiments. Next, I will discuss how we formulated the right conjecture based on computer experiments. Finally, I will briefly discuss the convergence proof.

3.19 Recent Progress on Some Covering Tour Problems

Joseph S. Mitchell (SUNY – Stony Brook, US)

Joint work of Mitchell, Joseph S.; Arkin, E.; Polishchuk, V.; Yang, S.

We give some recent results on some covering tour problems:

(1) We answer the question initially posed by Arik Tamir at the Fourth NYU Computational Geometry Day (March, 1987):

"Given a collection of compact sets, can one decide in polynomial time whether there exists a convex body whose boundary intersects every set in the collection?"

We prove that when the sets are segments in the plane, deciding existence of the convex stabber is NP-hard. We also show that in 3D the stabbing problem is hard when the sets are balls. On the positive side, we give a polynomial-time algorithm to find a convex polygonal transversal (on a given discrete set of vertices) of a maximum number of segments in 2D if the segments are pairwise-disjoint. Our algorithm also finds a convex stabber of the maximum number of a set of polygonal pseudodisks in the plane.

The stabbing problem is related to "convexity" of point sets and polygons, measured as the minimum distance by which the points/polygons must be shifted in order that there is a convex stabber.

(2) Watchman routes in polygons with holes: We prove a lower bound of $\Omega(\log n)$ on approximation, from SET-COVER. We also give an $O(\log n)$-approximation algorithm for the capae that a certain "bounded perimeter assumption" holds.

3.20 Natural Neighbor Interpolation Based Grid DEM Construction Using a GPU

Thomas Molhave (Duke University, US)

Joint work of Beutel, Alex; Mølhave, Thomas; Agarwal, Pankaj K.
Main reference Natural Neighbor Interpolation Based Grid DEM Construction Using a GPU. Alex Beutel, Thomas Malhave, Pankaj K. Agarwal. GIS '10: Proceedings of the 18th ACM SIGSPATIAL International Symposium on Advances in Geographic Information Systems, 2010.
URL http://dx.doi.org/10.1145/1869790.1869817

With modern LiDAR technology the amount of topographic data, in the form of massive point clouds, has increased dramatically.

One of the most fundamental GIS tasks is to construct a grid digital elevation model (DEM) from these 3D point clouds. In this paper we present a simple yet very fast algorithm for constructing a grid DEM from massive point clouds using natural neighbor interpolation (NNI). We use a graphics processing unit (GPU) to significantly speed up the computation. To handle the large data sets and to deal with graphics hardware limitations clever blocking schemes are used to partition the point cloud. For example, using standard desktop computers and graphics hardware, we construct a high-resolution grid with 150 million cells from two billion points in less than thirty-seven minutes. This is about one-tenth of the time required for the same computer to perform a standard linear interpolation, which produces a much less smooth surface

3.21 Tight bounds for epsilon-nets

Janos Pach (EPFL – Lausanne, CH)

According to a well known theorem of Haussler and Welzl (1987), any range space of bounded VC-dimension admits an ϵ-net of size $O\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$. Using probabilistic techniques, Pach and Woeginger (1990) showed that there exist range spaces of VC-dimension 2, for which the above bound is sharp. The only known range spaces of small VC-dimension, in which the ranges are geometric objects in some Euclidean space and the size of the smallest ϵ-nets is superlinear in $\frac{1}{\epsilon}$, were found by Alon (2010). In his examples, the size of the smallest ϵ-nets is $\Omega\left(\frac{1}{\epsilon}g(\frac{1}{\epsilon})\right)$, where g is an extremely slowly growing function, closely related to the inverse Ackermann function.

We show that there exist geometrically defined range spaces, already of VC-dimension 2, in which the size of the smallest ϵ-nets is $\Omega\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$. We also construct range spaces induced by axis-parallel rectangles in the plane, in which the size of the smallest ϵ-nets is $\Omega\left(\frac{1}{\epsilon}\log\log\frac{1}{\epsilon}\right)$. By a theorem of Aronov, Ezra, and Sharir (2010), this bound is tight.

3.22 Double Permutation Sequences and Arrangements of Planar Families of Convex Sets

Richard Pollack (New York University, US)

We (re)introduce Double Permutation Sequences, which provide a combinatorial encoding of arrangements of convex sets in the plane. We also recall the notion of a topological affine plane and several (some new) of its properties.

In particular, that there is a universal topological affine plane P (i.e. any finite arrangement of pseudolines is isomorphic to some arrangement of finitely many lines of P).

All of this work is joint with Jacob E. Goodman and some involves numerous other people, among whom are Raghavan Dhandapani, Andreas Holmsen, Shakhar Smorodinsky, Rephael Wenger, and Tudor Zamfirescu.

3.23 Shortest Paths in Time-Dependent FIFO Networks

Joerg-Ruediger Sack (Carleton University – Ottawa, CA)

Main reference F.Dehne, M.T.Omran, and J.-R.Sack, "Shortest paths in time-dependent FIFO networks," Algorithmica, to appear

In this talk, we study the time-dependent shortest paths problem for two types of time-dependent FIFO networks. First, we consider networks where the availability of links, given by a set of disjoint time intervals for each link, changes over time. Here, each interval is assigned a non-negative real value which represents the travel time on the link during the corresponding interval.

The resulting shortest path problem is the time-dependent shortest path problem for availability intervals, which asks to compute all shortest paths to any (or all) destination node(s) d for all possible start times at a given source node s. Second, we study time-dependent networks where the cost of using a link is given by a non-decreasing piece-wise linear function of a real-valued argument.

Here, each piece-wise linear function represents the travel time on the link based on the time when the link is used. The resulting shortest paths problem is the time-dependent shortest path problem for piece-wise linear functions which asks to compute, for a given source node s and destination d, the shortest paths from s to d, for all possible starting times.

We present an algorithm for both problems which improve significantly on the previously known algorithms.

3.24 Can Nearest Neighbor Search be Simple and always Fast?

Raimund Seidel (Universität des Saarlandes, DE)

We show that any method for nearest neighbor search that is simple, in the sense that the only operations involving query points are distance to site comparisons, cannot be fast on all inputs. In particular we show that any such methods can be forced to make $n-1$ such comparisons for some input, where n ist the number of sites.

3.25 From joints to distinct distances and beyond: The dawn of an algebraic era in combinatorial geometry

Micha Sharir (Tel Aviv University, IL)

In November 2010 the earth has shaken, when Larry Guth and Nets Hawk Katz posted a nearly complete solution to the distinct distances problem of Erdős, open since 1946. The excitement was twofold:

(a) The problem was one of the most famous problems, as well as one of the hardest nuts in the area, resisting solution in spite of many attempts (which only produced partial improvements).

(b) The proof techniques were algebraic in nature, drastically different from anything tried before.

The distinct distances problem is to show that any set of n points in the plane determine $Omega(n/\sqrt{\log n})$ distinct distances.

(Erdős showed that the grid attains this bound.) Guth and Katz obtained the lower bound $Omega(n/\log n)$.

Algebraic techniques of this nature were introduced into combinatorial geometry in 2008, by the same pair Guth and Katz. At that time they gave a complete solution to another (less major) problem, the so-called joints problem, posed by myself and others back in 1992.

Since then these techniques have led to several other developments, including an attempt, by Elekes and myself, to reduce the distinct distances problem to an incidence problem between points and lines in 3-space. Guth and Katz used this reduction and gave a complete solution to the reduced problem.

One of the old-new tools that Guth and Katz bring to bear is the Polynomial Ham Sandwich Cut, due to Stone and Tukey (1942). I will discuss this tool, including a "1-line" proof thereof, and its potential applications in geometry. One such application, just noted by Matoušek, is an algebraic proof of the classical Szemerédi-Trotter incidence bound for points and lines in the plane.

In the talk I will review all these developments, as time will permit.

Only very elementary background in algebra and geometry will be assumed.

3.26 Fixed Points of the Restricted Delaunay Triangulation Operator

Jonathan Shewchuk (Univ. of California – Berkeley, US)

Consider the problem of reconstructing an unknown smooth 2-manifold Σ embedded in three dimensions from a set of points S sampled from the surface. It is well known that if the point sample is sufficiently dense, then the restriction of the Voronoi diagram of S to the surface Σ dualizes to a triangulation that is homeomorphic to Σ. This triangulation is called the restricted Delaunay triangulation of S with respect to Σ. What if we take the restricted Delaunay triangulation of S with respect to this triangulation? And then repeat with the new triangulation? I show that the iteration always "converges" to a "fixed point," that is, a fixed set of triangles; and that for sufficiently dense samples, this set of triangles is likely to be a particularly nice reconstruction of Σ. Best of all, the idea works equally well for 2-manifolds embedded in much higher-dimensional spaces.

3.27 The potential to improve the choice: List coloring for geometric hypergraphs

Shakhar Smorodinsky (Ben Gurion University – Beer Sheva, IL)

Joint work of Smordinsky, Shakhar; Cheilaris, Panagiotis; Sulovsky, Marek

Given a hypergraph $H = (V, \mathcal{E})$, a coloring of its vertices is said to be conflict-free if for every hyperedge $S \in \mathcal{E}$ there is at least one vertex whose color is distinct from the colors of all other vertices in S.

This notion has been studied for several geometric hypergraphs and in various generalizations.

We study the list version of this notion:

In this version we are interested in the minimum number $k = k(H)$ such that if each vertex v is associated with a set of colors L_v of size k then one can pick a color for each vertex v from its “list" L_v such that the resulting coloring is conflict-free. Denote this number by $ch(H)$ (the conflict free “choice" number of H) It is easy to see that the minimum number of colors needed for a conflict-free coloring of H is bounded by $ch(H)$.

Let C be some absolute constant. We prove that for hypergraphs H with n vertices which are hereditarily C colorable (in the non-monochromatic sense) we have $ch(H) = O(\log n)$ and this bound is asymptotically tight. More over, we show that one can color the vertices of such a hypergraph from lists of size $O(\log n)$ such that the maximum color in any hyperedge is unique. The proof is constructive and uses a suitable potential function for constructing such coloring and analyzing the size of lists needed.

3.28 Shapely Measures, and Measures on Shapes: The story of the Kernel distance

Suresh Venkatasubramanian (University of Utah, US)

Joint work of Joshi, Sarang; Kommaraju, Raj Varma; Phillips, Jeff M.; Venkatasubramanian, Suresh
Main reference S.C. Joshi, R.V. Kommaraju, J.M. Phillips, S. Venkatasubramanian, "Comparing Distributions and Shapes using the Kernel Distance," Proc. 27th ACM Symp. on Computational Geometry, 2011, pp. 47–56.
URL http://doi.acm.org/10.1145/1998196.1998204

There are many ways to compute a distance between probability measures. However, if the underlying domain of the measures itself carries a metric, then the standard way to compare probability measures in a way that respects this is via the earthmover distance (or the transportation distance). This is a popular distance, especially in computer vision, but is expensive to compute.

There are also many ways to compare shapes, using methods from computational geometry, topology, and differential geometry. But these distances are all expensive to compute, are limited in important ways, and are difficult to use for applications that involve analyzing collections of shapes, rather than just pairs.

It turns out that there is a single mechanism that makes comparing distributions over metrics easy, and makes comparing (noisy) shapes (even surfaces) easy. It uses a kernel-based method to embed distributions (or shapes) in an (infinite-dimensional) Hilbert space, so that distance computation is a matter of computing the induced Hilbertian metric.

In this talk, I will describe algorithms for (i) estimating this "kernel distance" approximately and efficiently, (ii) computing sparse representations of shape s that are close to the original under the kernel distance, (iii) comparing shapes under rigid transformations using the kernel distance. I'll also briefly mention applications in the realm of clustering and metaclustering.

3.29 Kasteleyn and the Number of Crossing-Free Matchings and Cycles on a Planar Point Set

Emo Welzl (ETH Zürich, CH)

In 1967 Piet Kastelyn showed how to count the number of perfect matchings in a planar graph, simply by looking at the determinant of an appropriate skew symmetric variant of the adjacency matrix. Raimund Seidel observed an extremal combinatorics implication, namely that a planar graph on n vertices has at most $\sqrt[4]{6}^n$ perfect matchings (via the so-called Hadamard bound). We show how, for a planar set P of n points, that can be used to relate the number, $sc(P)$, of crossing-free straight-line spanning cycles (simple polygonizations) of P to the number of triangulations, $tr(P)$, of P: $sc(P) < 8\sqrt[4]{12}^n tr(P)$. We conjecture, though, that this can be significantly improved, actually to $sc(P) = O(c^n)tr(P)$ for some constant $c < 1$.

3.30 Computing the Frechet Distance Between Folded Polygons

Carola Wenk (Univ. of Texas at San Antonio, US)

Joint work of Cook, Atlas F. IV; Driemel, Anne; Har-Peled, Sariel; Sherette, Jessica; Wenk, Carola;

We present the first results showing that the Frechet distance between non-flat surfaces can be approximated within a constant factor in polynomial time. Computing the Frechet distance for surfaces is a surprisingly hard problem. It is not known whether it is computable, it has been shown to be NP-hard, and the only known algorithm computes the Frechet distance for flat surfaces (Buchin et al.). We adapt this algorithm to create one for computing the Frechet distance for a class of surfaces which we call folded polygons. Unfortunately, if extended directly the original algorithm no longer guarantees that a homeomorphism exists between the surfaces. We present three different methods to address this problem. The first of which is a fixed-parameter tractable algorithm.

The second is a polynomial-time approximation algorithm which approximates the optimum mapping. Finally, we present a restricted class of folded polygons for which we can compute the Frechet distance in polynomial time.

3.31 Triangular Meshes: Registration and GPU-based Distance Computation

Nicola Wolpert (Hochschule für Technik – Stuttgart, DE)

We present two projects with industrial partners. The first one is on the registration of triangular meshes. The motivation comes from sheet metal forming where after the stamping process the formed metal has to be compared with the original CAD-data. The main step here, based on the ICP-algorithm, is to compute the translation and rotation which best possibly aligns the two meshes.

The second project is called RASAND: Robust Algorithms for Distance Computation of large moving Triangular Meshes. The motivation comes from the digital mock up process where early in the construction phase of a mechanical part constructions and motions have to be evaluated. Mathematically we are given two large triangular meshes, one moving in discrete steps over time.

The question is to find all time steps and all pairs of triangles that become closer than a given epsilon. The challenge is the mass data and we are looking for data structures and algorithms that can answer this question using massively the computing-capability of modern graphics cards.

3.32 Cxyz: Isotopic Subdivision Algorithm for Non-Singular Surfaces

Chee K. Yap (New York University, US)

Joint work of Yap, Chee K.; Long Lin
URL http://cs.nyu.edu/exact/papers/

We consider exact algorithms for isotopic approximating of implicit surfaces given as the zero set of a C1 real function $f(x, y, z)$.

We focus on (domain) subdivision methods based on numerical predicates because they tend to have adaptive complexity, locality, and are easy to implement.

Numerical predicates, unlike algebraic or geometric predicates, admits a trade-off between efficacy (how tight is the bound) and efficiency (how quickly it can be computed); this can be exploited. The challenge of numerical predicates is how to ensure global correctness.

We describe an algorithm called Cxyz Algorithm (the name derives from the key predicate Cxyz in the algorithm).

It is a generalization our earlier Cxy algorithm for plane curves.

The algorithm exploits two properties of two previous approaches: non-local isotopy [Plantinga and Vegter, 2004] and parameterizability [Snyder, 1994].

Our preliminary implementations in Core Library suggest that our algorithm is very efficient compared to previous methods.

The proof of correctness is quite non-trivial. Termination is guaranteed for non-singular surface. Following Plantinga-Vegter, we first prove the correctness of a simpler algorithm called Regular Cxyz, and extend it to the Cxyz Algorithm which involves balancing of boxes.

As in the 2-D case, we need to deal with ambiguity.

A new phenomenon is that the arc connection rules require some global consistency properties.

We also briefly extend the method to discuss boundary processing, admitting a general Region-of-Interest (ROI), and anisotropic subdivision (boxes can be half-split, quarter-split or full-split producing arbitrary aspect ratio boxes).

3.33 Median trajectories

Marc van Kreveld (Utrecht University, NL)

Given a collection of more or less similar trajectories, what would constitute a "middle" trajectory? We address this problem by constructing a median trajectory from pieces of the input trajectories, while always staying in the middle. We consider two possible methods. One is simple, often fails when the trajectories have self-intersections. The other one works well and uses the concept of homotopy. We motivate our choices, give algorithms, and show some experimental results.

4 Open Problems

▶ Problem 1 (Jack Snoeyink). *Solvent Inaccessible Volume, or Nature Abhors a Vacuum:* Given a set $S = \{B_1, \ldots, B_n\}$ of n balls in 3D, give a very fast method to estimate the volume of the portion of 3-space that is inaccessible to a ball W that is allowed to be placed anywhere that it is not intersecting a ball of S. (i.e., give a very fast method to estimate the volume of an α-hull determined by n balls) The motivating application (solvent-accessibility for molecules) has balls B_i of radii in some range (1–2Å), W of radius 1.4Å, and density constraints that no pair of balls is within .8Å. One could also fix 2/3 of the atom positions, and choose from a fixed set of positions for the rest (using what is known as a rotamer library). (Danny Halperin mentions the related work he and Mark Overmars had in SoCG'94, "Spheres, Molecules, and Hidden Surface Removal".)

▶ Problem 2 (Shakhar Smorodinsky). *Coloring vertices of a simplicial complex:* Given a set P of n points in 3D and a set S of triangles spanned by P, with no two triangles having their interiors intersecting. How many colors suffice to color P so that no triangle of S is monochromatic? (This is a natural generalization of the problem of coloring the vertices of a planar (straight-line) graph, for which the answer is 4 colors always suffice.) Shakhar knows an upper bound of $O(\sqrt{n})$. The conjecture is that a constant or polylog number of colors suffice. Günter Rote asked if the special case of a simplicial complex of triangles given by a tetrahedralization of the convex hull of a set of points may be of interest (even if the points P are in convex position).

▶ Problem 3 (Janos Pach). We define the *obstacle number*, $f(G)$, of a graph $G = (V, E)$ as the minimum number of obstacles needed in a "representation" of G as a visibility graph of a set of $|V|$ points in the plane. Let $F(n)$ be the maximum obstacle number of graphs having n vertices. Clearly, $F(n) \leq \binom{n}{2}$. Janos knows that $F(n) \geq \Omega(n/\log n)$. He asks if there exists an n-vertex graph G having $f(G) = \Omega(n)$. He conjectures that "most" graphs (e.g., random graphs, expanders) have nearly quadratic obstacle number.

▶ Problem 4 (Janos Pach, and Gabor Tardos). Is it true that for $S = \{p_1, \ldots, p_n\} \subset \Re^2$ one can always find $n/2$ points not necessarily in S such that every axis-parallel rectangle that avoids T contains at most 1000 points of S? (i.e., are $n/2$ points enough to stab every rectangle containing at least 1000 elements of S?)

▶ Problem 5 (Suresh Venkatasubramanian). (Shout out to cstheory.stackexchange.com) Given a set S of n points in 2D. Consider the complete graph whose edge weights are the squared Euclidean distances between pairs of points of S. Is it always possible to find a cut whose weight is at least 2/3 of the total edge weight? (Note that it is not possible in 3D.)

▶ Problem 6 (Ferran Hurtado (6.1)). Let $S = \{p_1, \ldots, p_n\}$ be a set of n "blue" (not white!) points in 2D. Let W be a set of "red" points in 2D; these are "witness" points. Define the "witness Delaunay graph", $WDG(S, W)$ of S with respect to W that joins blue points $p_i \in S$ and $p_j \in S$ with an edge if and only if there exists a circle through p_i and p_j whose interior has no red points of W. Note that $WDG(S, \emptyset) = K_n$ and $WDG(S, S) = Del(S)$, the usual Delaunay graph of S. Let $f(n)$ be the number of witness points W that always suffice, and are sometimes necessary, to make $WDG(S, W) = \emptyset$, for an n-element set S (equivalently, no pair of blue points have adjacent regions in $Vor(S \cup W)$, as they have been "surrounded" by the red witnesses). Note that it always suffice to place a red point interior to each Delaunay edge of $Del(S)$ in order to "kill" all edges and make $WDG(S, W)$ empty; thus, $f(n) \leq 3n$. Better bounds are known: $n \leq f(n) \leq (3/2)n$, where the lower bound is from (6.1) and the upper bound from (6.2). Close the gap! The authors of (6.2) conjecture that $f(n) = n$. Also, what can be said algorithmically about computing a smallest set W to "kill" all Delaunay edges of points S?

(6.1) Witness (Delaunay) Graphs. B. Aronov, M. Dulieu, F. Hurtado. Computational Geometry: Theory and Applications, Volume 44, Issues 6-7, August 2011, Pages 329–344.
(6.2) O. Aichholzer, R. Fabila-Monroy, T. Hackl, M. van Kreveld, A. Pilz, P. Ramos, B. Vogtenhuber, Blocking Delaunay triangulations, in: Proc. 22nd Annual Canadian Conference on Computational Geometry CCCG 2010, Winnipeg, Manitoba, Canada, 2010, pp. 21–24.

▶ Problem 7 (Boris Aronov). (Posed previously, but still open.) Let K be a convex body in 3D. Consider n translates of K, $K_1, \ldots, K_n$. What is the maximum possible number, $c(n)$, of connected components of the set, $(K_1 \cup \cdots \cup K_n)^c$, the complement of the union of the translates? He knows that $\Omega(n^2) \leq c(n) \leq \binom{n}{3}$ and conjectures that $c(n) = \Theta(n^2)$. (Some special classes of convex bodies have a linear number of components in the complement of the union.)

▶ Problem 8 (Boris Aronov). Let S be a set of n blue points in 2D, and let W be a set of n red points in 2D. Let D_{pq} be the diametrical disk determined by $p, q \in S$. We want to compute all m pairs (p, q) for which $D_{pq} \cap W \neq \emptyset$. How efficiently can it be done? ($O(n^2 \log n)$, and possibly $O(n^2)$, is easy) Give an output-sensitive algorithm, e.g., taking time $O(m + n \log n)$.

▶ Problem 9 (Joe Mitchell). The following problem arises in the "Last Byte" column of Peter Winkler in CACM in 2010. Given n points in 2D, can they always be covered by n disjoint unit disks ("pennies")? The answer is "yes" for $n = 10$ (by a simple probabilistic method argument), for $n = 11$, and for $n = 12$. It is not known for $n = 13$. Janos mentions that there is a paper (by physicists) showing that for $n \geq 50$ there are configurations of points for which coverage by n disjoint disks is not possible. There remains a gap between 12 and 50.

▶ Problem 10 (Pankaj Agarwal). Let P be a set of n points in $\Re^2$, and let D be a set of m disks in $\Re^2$. We are given a positive integer k. Our goal is to pick a subset $R \subseteq P$ of k points in order to maximize the number of disks that are "double-stabbed" (i.e., that contain at least 2 points of R). THe problem is known to be NP-hard. No nontrivial approximation is known. (The related problem of minimizing the size of R in order that *every* disk is double-stabbed has an $O(\log OPT)$-approximation, by Chekuri, Clarkson, and Har Peled. The abstract set version is known not to have any nontrivial approximation algorithm.)

Participants

- Pankaj Kumar Agarwal
Duke University, US
- Oswin Aichholzer
TU Graz, AT
- Helmut Alt
FU Berlin, DE
- Victor Alvarez
Universität des Saarlandes, DE
- Lars Arge
Aarhus University, DK
- Boris Aronov
Polytech. Univ. – New York, US
- Tetsuo Asano
JAIST – Nomi, JP
- Dominique Attali
GRIPSA Lab – Saint Martin d'Héres, FR
- Erin Moriarty Wolf Chambers
St. Louis University, US
- Bernard Chazelle
Princeton University, US
- Otfried Cheong
KAIST – Daejeon, KR
- Kenneth L. Clarkson
IBM Almaden Center – San José, US
- Erik D. Demaine
MIT – Cambridge, US
- Olivier Devillers
INRIA Sophia Antipolis, FR
- Herbert Edelsbrunner
IST Austria – Klosterneuburg, AT
- Alon Efrat
Univ. of Arizona – Tucson, US
- Esther Ezra
New York University, US
- Sándor Fekete
TU Braunschweig, DE
- Jie Gao
SUNY – Stony Brook, US
- Joachim Giesen
Universität Jena, DE
- Xavier Goaoc
INRIA Lorraine, FR
- Dan Halperin
Tel Aviv University, IL
- Ferran Hurtado
UPC – Barcelona Tech, ES
- Michael Kerber
IST Austria – Klosterneuburg, AT
- David G. Kirkpatrick
University of British Columbia – Vancouver, CA
- Rolf Klein
Universität Bonn, DE
- Christian Knauer
Universität Bayreuth, DE
- Sylvain Lazard
INRIA – Nancy, FR
- Kurt Mehlhorn
MPI für Informatik – Saarbrücken, DE
- Joseph S. Mitchell
SUNY - Stony Brook, US
- Thomas Molhave
Duke University, US
- Nabil Hassan Mustafa
LUMS – Lahore, PK
- Janos Pach
EPFL – Lausanne, CH
- Richard Pollack
New York University, US
- Saurabh Ray
EPFL – Lausanne, CH
- Günter Rote
FU Berlin, DE
- Jörg-Rüdiger Sack
Carleton Univ. – Ottawa, CA
- Michael Sagraloff
MPI für Informatik - Saarbrücken, DE
- Raimund Seidel
Universität des Saarlandes, DE
- Micha Sharir
Tel Aviv University, IL
- Jonathan Shewchuk
University of California – Berkeley, US
- Shakhar Smorodinsky
Ben Gurion University – Beer Sheva, IL
- Jack Snoeyink
University of North Carolina – Chapel Hill, US
- Bettina Speckmann
TU Eindhoven, NL
- Ileana Streinu
Smith Coll. – Northampton, US
- Monique Teillaud
INRIA Sophia Antipolis, FR
- Marc van Kreveld
Utrecht University, NL
- Suresh Venkatasubramanian
University of Utah – Salt Lake City, US
- Emo Welzl
ETH Zürich, CH
- Carola Wenk
University of Texas at San Antonio, US
- Nicola Wolpert
Hochschule für Technik – Stuttgart, DE
- Chee K. Yap
New York University, US
- Ke Yi
The Hong Kong University of Science & Technology, HK