Report from Dagstuhl Seminar 15111

Computational Geometry

Edited by

Otfried Cheong[1], Jeff Erickson[2], and Monique Teillaud[3]

1 **KAIST – Daejeon, KR,** otfried@kaist.edu
2 **University of Illinois – Urbana, US,** jeffe@cs.uiuc.edu
3 **INRIA Nancy – Grand Est, FR,** Monique.Teillaud@inria.fr

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 15111 "Computational Geometry". The seminar was held from 8th to 13th March 2015 and 41 senior and young researchers from various countries and continents attended it. Recent developments in the field were presented and new challenges in computational geometry were identified.

This report collects abstracts of the talks and a list of open problems.

Seminar March 8–13, 2015 – http://www.dagstuhl.de/15111

1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity, G.2 Discrete Mathematics, G.4 Mathematical Software

Keywords and phrases Algorithms, geometry, theory, approximation, implementation, combinatorics, topology

Digital Object Identifier 10.4230/DagRep.5.3.41

Edited in cooperation with Maria Saumell

1 Executive Summary

Otfried Cheong
Jeff Erickson
Monique Teillaud

Computational Geometry

Computational geometry is concerned with the design, analysis, and implementation of algorithms for geometric and topological problems, which arise naturally in a wide range of areas, including computer graphics, robotics, geographic information systems, molecular biology, sensor networks, machine learning, data mining, scientific computing, theoretical computer science, and pure mathematics. Computational geometry is a vibrant and mature field of research, with several dedicated international conferences and journals, significant real-world impact, and strong intellectual connections with other computing and mathematics disciplines.

Seminar Topics

The emphasis of the seminar was on presenting recent developments in computational geometry, as well as identifying new challenges, opportunities, and connections to other fields of computing. In addition to the usual broad coverage of emerging results in the field, the seminar included invited survey talks on two broad and overlapping focus areas that cover a wide range of both theoretical and practical issues in geometric computing. Both focus areas have seen exciting recent progress and offer numerous opportunities for further cross-disciplinary impact.

Computational topology and topological data analysis. Over the last decade, computational topology has grown from an important subfield of computational geometry into a mature research area in its own right. Results in this field combine classical mathematical techniques from combinatorial, geometric, and algebraic topology with algorithmic tools from computational geometry and optimization. Key developments in this area include algorithms for modeling and reconstructing surfaces from point-cloud data, algorithms for shape matching and classification, topological graph algorithms, new generalizations of persistent homology, practical techniques for experimental low-dimensional topology, and new fundamental results on the computability and complexity of embedding problems. These results have found a wide range of practical applications in computer graphics, computer vision, robotics, sensor networks, molecular biology, data analysis, and experimental mathematics.

Geometric data analysis. Geometric data sets are being generated at an unprecedented scale from many different sources, including digital video cameras, satellites, sensor networks, and physical simulations. The need to manage, analyze, and visualize dynamic, large-scale, high-dimensional, noisy data has raised significant theoretical and practical challenges not addressed by classical geometric algorithms. Key developments in this area include new computational models for massive, dynamic, and distributed geometric data; new techniques for effective dimensionality reduction; approximation algorithms based on coresets and other sampling techniques; algorithms for noisy and uncertain geometric data; and geometric algorithms for information spaces. Results in this area draw on mathematical tools from statistics, linear algebra, functional analysis, metric geometry, geometric and differential topology, and optimization, and they have found practical applications in spatial databases, clustering, shape matching and analysis, machine learning, computer vision, and scientific visualization.

Participants. Dagstuhl seminars on computational geometry have been organized in a two year rhythm since a start in 1990. 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, fostering collaboration, and exposing young talent to the seniors of the field. They have arguably been the most influential meetings in the field of computational geometry.

The organizers held a *lottery* for the second time this year; the lottery allows to create space to invite younger researchers, rejuvenating the seminar, while keeping a large group of senior and well-known scholars involved. Researchers on the initial list who were not selected by the lottery were notified by us separately per email, so that they knew that they were not forgotten, and to reassure them that—with better luck—they will have another chance in future seminars. The seminar has now a more balanced attendance in terms of seniority and gender than in the past.

This year, 41 researchers from various countries and continents attended the seminar, showing the strong interest of the community for this event. The feedback from participants was very positive.

No other meeting in our field allows young researchers to meet with, get to know, and work with well-known and senior scholars to the extent possible at the Dagstuhl Seminar.

We warmly thank the scientific, administrative and technical staff at Schloss Dagstuhl! 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, Schloss Dagstuhl itself is a great strength of the seminar.

2 Table of Contents

3 Overview of Talks

3.1 Untraditional Geometric Queries

Peyman Afshani (Aarhus University, DK)

We consider some geometric queries that do not fit in the traditional semigroup searching/reporting model. After a brief review of the classical roots of range searching and existing classical results, we will examine a few recent results (both published and unpublished).

First, we look at recent results on "Concurrent Queries" where each point is associated with nomial data fields (e.g., "color") and the query includes both a geometric region and a list of colors. The output should be the points inside the geometric region with the specified colors. These results were presented in SODA'14 and they are joint works with Bryan Wilkinson, Yufei Tao, Cheng Sheng.

Next, we look at two different geometric queries: range summary queries and range sampling queries. After reviewing their definitions, we will briefly mention some yet unpublished results obtained in a joint work with Zhewei Wei.

We will finish with a list of open problems.

3.2 Surface Patches from Unorganized Space Curves

Annamaria Amenta (University of California – Davis, US)

Recent 3D sketch tools produce networks of three-space curves that suggest the contours of shapes. The shapes may be non-manifold, closed three-dimensional, open two-dimensional, or mixed. We describe a system that automatically generates intuitively appealing piecewise-smooth surfaces from such a curve network, and an intelligent user interface for modifying the automatically chosen surface patches. Both the automatic and the semi-automatic parts of the system use a linear algebra representation of the set of surface patches to track the topology. On complicated inputs from ILoveSketch [1], our system allows the user to build the desired surface with just a few mouse-clicks.

References

1 Seok-Hyung Bae, Ravin Balakrishnan, Karan Singh. *ILoveSketch: as-natural-as-possible sketching system for creating 3d curve models*. Proc. UIST'08, pp. 151–160.

3.3 Voronoi Diagrams of Parallel Halflines in 3D

Franz Aurenhammer (TU Graz, AT)

The Voronoi diagram for n lines and/or line segments in 3D is a complicated structure. Bisectors are complex geometric objects, and the combinatorial size is still unclear. Things get somewhat easier when the line segments are confined to have only a constant number of orientations. We consider the special case of n parallel (vertical) halflines in 3D. In this case, the intersection of the 3D diagram with any horizontal plane can be shown to be a power diagram of n weighted point sites. This enables us to study the structural properties of the Voronoi diagram of parallel halflines, and to design a relatively simple and output-sensitive algorithm for constructing it.

3.4 Faster DBSCAN and HDBSCAN in Low-Dimensional Euclidean Spaces

Mark de Berg (TU Eindhoven, NL)

DBSCAN is one of the most widely used density-based clustering methods. The clustering it produces depends on two parameters, MinPoints and ε, where MinPoints is typically fixed at a small constant, and ε essentially determines the scale at which we perform the clustering.

We present a new algorithm for DBSCAN in Euclidean spaces, whose running time is much less sensitive to the value of the parameter ε than previous approaches. As a result, our algorithm computes a DBSCAN-clustering in subquadratic time in the worst case when MinPoints is a constant, irrespective of the choice of ε. The worst-case running time of our algorithm in $\mathbb{R}^d$ is $O(n \log n)$ for $d = 2$ and $O(n^{2-\frac{2}{\lceil d/2 \rceil+1}+\gamma})$ for $d \geq 3$, where $\gamma > 0$ is an arbitrarily small constant. Our experiments show that the new algorithm is not only faster in theory, but also in many practical settings.

We also present a novel algorithm for HDBSCAN, a hierarchical version of DBSCAN introduced recently. In $\mathbb{R}^2$ our algorithm computes the HDBSCAN hierarchy in $O(n \log n)$ time in the worst case when MinPoints is a constant.

Finally, we introduce δ-approximate DBSCAN* and δ-approximate HDBSCAN, and we show how to compute these approximate versions of DBSCAN and HDBSCAN in near-linear time in any fixed dimension, for any given approximation error $\delta > 0$.

3.5 Segmentation and Classification of Trajectories

Maike Buchin (Ruhr-Universität Bochum, DE)

We consider segmentation and classification of trajectories, that is splitting and grouping trajectories such that they have similar movement characteristics. Our approach is based on a movement model parameterized by a single parameter, like the Brownian bridge movement model. We define an optimal segmentation (resp. classification) to be one that minimizes an information criterion balancing the likelihood of the model and its size. We give an efficient algorithm to compute the optimal classification for a discrete set of parameter values. For continuous parameters the problem becomes NP-hard. But we also present an algorithm that solves the problem in polynomial time under mild assumptions on the input.

3.6 Shortest Paths on Polyhedral Surfaces and Terrains

Siu-Wing Cheng (HKUST – Kowloon, HK)

Joint work of Cheng, Siu-Wing; Jin, Jiongxin; Vigneron, Antoine
Main reference S.-W. Cheng, J. Jin, "Shortest paths on polyhedral surfaces and terrains," in Proc. of the 46th Annual ACM Symp. on Theory of Computing (STOC'14), pp. 373–382, ACM, 2014.
URL http://doi.acm.org/10.1145/2591796.2591821

We present an algorithm for computing shortest paths on polyhedral surfaces under convex distance functions. Let n be the total number of vertices, edges and faces of the surface. Our algorithm can be used to compute L_1 and L_∞ shortest paths on a polyhedral surface in $O(n^2 \log^4 n)$ time. Given an $\epsilon \in (0, 1)$, our algorithm can find $(1 + \epsilon)$-approximate shortest paths on a terrain with gradient constraints and under cost functions that are linear combinations of path length and total ascent. The running time is $O(\frac{1}{\sqrt{\varepsilon}} n^2 \log n + n^2 \log^2 n \log^2(n/\epsilon)$. This is the first efficient PTAS for such a general setting of terrain navigation.

3.7 Walking in Random Delaunay Triangulations

Olivier Devillers (INRIA Nancy – Grand Est, FR)

Joint work of Broutin, Nicolas; Devillers, Olivier; Hemsley, Ross
Main reference N. Broutin, O. Devillers, R. Hemsley, "Efficiently navigating a random Delaunay triangulation," in Proc. of the 25th Int'l Conf. on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms (AofA'14), DMTCS-HAL Proceedings Series, pp. 49–60, 2014.
URL https://hal.inria.fr/hal-01077251

Walking in triangulation is a widely used strategy for point location in triangulation. There are several strategies to walk between neighboring vertices or neighboring faces of a triangulation, but the analysis of such strategies under random distribution hypotheses for the point set is very difficult. This is due to the fact that the probability for an edge to be part of the walk depends on the whole set of points, thus you get dependence between these probabilities that are difficult to deal with. All these kind of walks are conjectured to have length $O(\sqrt{n})$. We propose the analysis of two walking strategies.

The cone walk is a walk amongst vertices where the dependence is reduced. The visibility is the most commonly used strategy to walk amongst faces and we analyze it using percolation theory.

3.8 Toward Parameter-Free (Friendly?) Topology Inference

Tamal K. Dey (Ohio State University – Columbus, US)

Joint work of Dey, Tamal K.; Dong, Zhe; Wang, Yusu

In topological inference from point data, a simplicial complex such as Vietoris-Rips is built on top of the data to carry out the topological analysis. This requires a user-supplied global parameter, which in some cases may be impossible to determine for the purpose of correct topology inference. We show that when the underlying space is a smooth manifold of known dimension embedded in an Euclidean space, a parameter-free sparsification of the data leads to a correct homology inference. This follows from the fact that we can compute a function called lean-set feature size over the data points with which it can be made locally uniform. The construction of the Vietoris-Rips complex on such data can be done adaptively without requiring any user-supplied parameter from which homology of the hidden manifold can be inferred. Preliminary experiments suggest that the strategy achieves correct topological (homology) inference with effective sparsification in practice.

3.9 Realization Spaces of Arrangements of Convex Bodies

Michael Gene Dobbins (Postech – Pohang, KR)

Joint work of Dobbins, Michael Gene; Holmsen, Andreas; Hubard, Alfredo
Main reference M. G. Dobbins, A. Holmsen, A. Hubard, "Realization spaces of arrangements of convex bodies," arXiv:1412.0371v2 [math.MG], 2015.
URL http://arxiv.org/abs/1412.0371v2

In this talk I introduce combinatorial types of arrangements of convex bodies, extending order types of point sets to arrangements of convex bodies, and present some results on their realization spaces. Our main results witness a trade-off between the combinatorial complexity of the bodies and the topological complexity of their realization space. First, we show that every combinatorial type is realizable and its realization space is contractible under mild assumptions. Second, we prove a universality theorem that says the restriction of the realization space to arrangements polygons with a bounded number of vertices can have the homotopy type of any primary semialgebraic set. This is joint work with Andreas Holmsen and Alfredo Hubard.

3.10 Clustering Time Series under the Frechet Distance

Anne Driemel (TU Eindhoven, NL)

The Frechet distance is a popular distance measure for curves. We study the problem of clustering time series under the Frechet distance. In particular, we give $(1+\varepsilon)$-approximation algorithms for variations of the following problem with parameters k and l. Given n univariate time series P, each of complexity at most m, we find k time series, not necessarily from P, which we call cluster centers and which each have complexity at most l, such that (a) the maximum distance of an element of P to its nearest cluster center or (b) the sum of these distances is minimized. Our algorithms have running time near-linear in the input size. To the best of our knowledge, our algorithms are the first clustering algorithms for the Frechet distance which achieve an approximation factor of $(1+\varepsilon)$ or better.

3.11 Low-quality Dimension Reduction and High-dimensional Approximate Nearest Neighbor

Ioannis Z. Emiris (University of Athens, GR)

Joint work of Anagnostopoulos, Evangelos; Emiris, Ioannis Z.; Psarros, Ioannis
Main reference E. Anagnostopoulos, I. Z. Emiris, I. Psarros, "Low-quality dimension reduction and high-dimensional Approximate Nearest Neighbor," in Proc. of the 31st Int'l Symp. on Computational Geometry (SoCG'14), LIPIcs, Vol. 34, pp. 436–450, Schoss Dagstuhl – Leibniz-Zentrum für Informatik, 2015; pre-print available at arXiv:1412.1683v1 [cs.CG].
URL http://dx.doi.org/10.4230/LIPIcs.SOCG.2015.436
URL http://arxiv.org/abs/1412.1683v1

The approximate nearest neighbor problem (ϵ-ANN) in a Euclidean space is a fundamental question, which has been addressed by two main approaches: Data-dependent space partitioning techniques, typically tree-based such as kd-trees or BBD-trees, perform well when the dimension is bounded, but are affected by the curse of dimensionality. On the other hand, Locality Sensitive Hashing (LSH) has polynomial dependence in the dimension, sublinear query time with an exponent inversely proportional to $(1+\epsilon)^2$, and subquadratic space requirement.

In this paper, we generalize the celebrated Johnson-Lindenstrauss Lemma to define "low-quality" mappings to a Euclidean space of significantly lower dimension than previously considered, such that they satisfy a requirement weaker than approximately preserving all distances or even preserving the nearest neighbor. This mapping guarantees, with high probability, that an ANN lies among the k ANN's in the projected space: the latter can be efficiently retrieved by a tree-based data structure, such as BBD-trees. Our algorithm, given n points in dimension d, achieves optimal space usage in $O(dn)$, preprocessing time in $O(dn \log n)$, and query time in $O(dn^{\rho} \log n)$, where ρ is proportional to $1 - 1/\ln\ln n$, for fixed $\epsilon \in (0,1)$. Moreover, our method is quite simple and easy to implement. The dimension reduction is larger if one assumes that pointsets possess some structure, namely bounded expansion rate.

We implemented our method using projection matrices whose entries are i.i.d. Gaussian variables and solve the k-ANN problem in the projected space by using software library ANN. We present experimental results in up to 500 dimensions and 10^6 points, which show that the practical performance is better than that predicted by the theoretical analysis. In particular, k seems to grow like $\sqrt{n}$ rather than n^{ρ}. In addition, we compare our approach to E2LSH: our method requires less space but is somewhat slower than E2LSH on the examined datasets.

3.12 The Offset Filtration of Convex Objects

Michael Kerber (MPI für Informatik – Saarbrücken, DE)

Joint work of Kerber, Michael; Halperin, Dan; Shaharabani, Doron
Main reference D. Halperin, M. Kerber, D. Shaharabani, "The Offset Filtration of Convex Objects," arXiv:1407.6132v2 [cs.CG], 2015.
URL http://arxiv.org/abs/1407.6132v2

We consider offsets of a union of convex objects. We aim for a filtration, a sequence of nested simplicial complexes, that captures the topological evolution of the offsets for increasing radii. We describe methods to compute a filtration based on the Voronoi partition with respect to the given convex objects. The size of the filtration and the time complexity for computing it are proportional to the size of the Voronoi diagram and its time complexity, respectively. Our approach is inspired by alpha-complexes for point sets, but requires more involved machinery and analysis primarily since Voronoi regions of general convex objects do not form a good cover. We show by experiments that our approach results in a similarly fast and topologically more stable method for computing a filtration compared to approximating the input by a point sample.

3.13 Minimizing Co-location Potential for Moving Points

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

Joint work of Kirkpatrick, David G.; Evans, Will; Löffler, Maarten; Staals, Frank; Busto, Daniel

Imagine a collection of entities that move in d-dimensional space each with some bound on their speed. If we know the location of an individual entity at a particular time then its location lies in a region of uncertainty at all subsequent times. We consider the problem of minimizing the ply of the uncertainty regions (defined as the maximum, over all points p in the space, of the number of uncertainty regions that contain p) by means of queries to individual entities that are restricted to one query per unit of time. This notion of co-location potential is studied in two settings, one where ply is measured at some fixed time in the future, and the other where ply is measured continuously (i.e. at all times). Competitive query strategies are described in terms of a notion of intrinsic ply (the minimum ply achievable by any query strategy, even one that knows the trajectories of all entities).

Based on joint work with Will Evans, Maarten Löffler, Frank Staals, and Daniel Busto.

3.14 Fire

Rolf Klein (Universität Bonn, DE)

Joint work of Klein, Rolf; Langetepe, Elmar; Levcopoulos, Christos

Suppose that a circular fire spreads in the plane at unit speed. A fire fighter can build a barrier at speed $v > 1$. How large must v be to ensure that the fire can be contained, and how should the fire fighter proceed? We provide two results. First, we analyze the natural strategy where the fighter keeps building a barrier along the frontier of the expanding fire. We prove that this approach contains the fire if $v > v_c = 2.6144\ldots$ holds. Second, we show that any "spiralling" strategy must have speed $v > 1.618$, the golden ratio, in order to succeed.

3.15 Approximating the Colorful Caratheodory Theorem

Wolfgang Mulzer (FU Berlin, DE)

Joint work of Mulzer, Wolfgang; Stein, Yannik
Main reference W. Mulzer, Y. Stein, "Computational Aspects of the Colorful Caratheodory Theorem," arXiv:1412.3347v1 [cs.CG], 2014.
URL http://arxiv.org/abs/1412.3347v1

Given $d+1$ point sets $P_1, \ldots, P_{d+1}$ in $\mathbb{R}^d$ (the color classes) such that each set P_i contains the origin in its convex hull, the colorful Caratheodory theorem states that there is a colorful choice C which also contains the origin in its convex hull. Here, a colorful choice means a set containing at most one point from each color class. So far, the computational complexity of computing such a colorful choice is unknown.

We consider a new notion of approximation: a set C' is called a c-colorful choice if it contains at most c points from each color class. We show that for all $\varepsilon > 0$, an $\varepsilon(d+1)$-colorful choice containing the origin in its convex hull can be found in polynomial time.

3.16 The Cosheaf-Less Reeb Graph Interleaving Distance

Elizabeth Munch (University of Albany, US)

Joint work of de Silva, Vin; Munch, Elizabeth; Patel, Amit
Main reference V. de Silva, E. Munch, A. Patel, "Categorified Reeb Graphs," arXiv:1501.04147v1 [cs.CG], 2015.
URL http://arxiv.org/abs/1501.04147v1

The interleaving distance was recently defined in order to give a method for comparison of Reeb graphs. The definition draws inspiration from the interleaving distance for persistence modules via category theory and cosheaves. Here, we present this distance using the equivalent yet concrete definition which looks for function preserving maps on graphs and checks for commutativaty of a particular diagram. The distance definition also yields as a substep a new definition for the smoothed Reeb graph. This later construction can be performed in polynomial time, while the general computation of the distance is graph isomorphism hard. This is joint work with Vin de Silva and Amit Patel.

3.17 On a Line-symmetric Puzzle

Yota Otachi (JAIST – Ishikawa, JP)

Joint work of Korman, Matias; Otachi, Yota; Roeloffzen, Marcel; Uehara, Ryuhei; van Renssen, André

Given k simple polygons, the goal of the line-symmetric puzzle is to find a polygon that can be exactly covered by the k polygons without overlap. We study the computational complexity of this puzzle and show a hardness result.

3.18 Geometric Data Analysis: Matrix Sketching to Kernels

Jeff M. Phillips (University of Utah – Salt Lake City, US)

I overview some recent developments in geometric data analysis. The initial focus will be in describing how geometric analysis has been essential and central to many core problems in data mining and machine learning. Then I overview recent developments in the area of matrix sketching which has broad applications within these core data mining and machine learning problems. I highlight the geometric connections, recent developments, and broad future directions. Finally, I talk about the uses of kernels and kernel density estimates for geometric data analysis. These enforce certain analyses to be robust, and in some cases have computational advantages. In this area I identify a number of open computational geometry problems which while easy to state may have important implications in data analysis.

3.19 Richter-Thomassen Conjecture about Pairwise Intersecting Curves (and Beyond)

Natan Rubin (Ben Gurion University – Beer Sheva, IL)

A long standing conjecture of Richter and Thomassen states that the total number of intersection points between any n simple closed (i.e., Jordan) curves in the plane which are in general position and any pair of them intersect, is at least $(2 - o(1))n$.

Very recently, we established an even stronger form of the above conjecture, which states that the overall number of proper intersection points must exceed, in asymptotic terms, the number of the touching pairs of curves.

If time permits, we discuss this result in connection with other fundamental questions concerning string graphs and arrangements of curves in the plane.

This is joint work in progress with Janos Pach and Gabor Tardos.

3.20 Controlling Modular Robotic Systems: Some Ideas from Computational Geometry

Vera Sacristan (UPC – Barcelona, ES)

A self-reconfiguring modular robot consists of a large number of independent units that can rearrange themselves into a structure best suited for a given environment or task. For example, it may reconfigure itself into a thin, linear shape to facilitate passage through a narrow tunnel, transform into an emergency structure such as a bridge, or surround and manipulate objects in outer space. Since modular robots comprise groups of identical units, they can also repair themselves by replacing damaged units with functional ones. Such robots are especially well-suited for working in unknown and remote environments.

In this talk I will introduce various types of units for modular robots that have been designed and prototyped by the robotics community, present the current challenges in the field, discuss how computational geometry can help in solving some of them, and present some current results and strategies, as well as open problems.

3.21 A Dynamic Programming Algorithm to Find Subsets of Points in Convex Position Optimizing some Parameter

Maria Saumell (University of West Bohemia – Pilsen, CZ)

Given a set S of n points in the plane, we may consider the problem of finding a subset of S of maximum cardinality such that they are the vertices of a convex polygon and their convex hull is empty of other points of S. This problem can be solved in cubic time by a dynamic programming algorithm [1]. We show that this algorithm can be adapted to solve a variety of other optimization problems related to convex polygons, in particular, the problem of computing largest monochromatic islands in a bicolored point set [2], or the problem of finding cliques of maximum size in the visiblity graph of a simple polygon [3, 4].

References

1 David Avis, David Rappaport. *Computing the largest empty convex subset of a set of points*. Proc. SoCG'85, pp. 161–167.
2 Crevel Bautista-Santiago, José Miguel Díaz-Báñez, Dolores Lara, Pablo Pérez-Lantero, Jorge Urrutia, Inmaculada Ventura. *Computing optimal islands*. Oper. Res. Lett. 39(4):246–251 (2011).
3 Sergio Cabello, Maria Saumell. *A randomized algorithm for finding a maximum clique in the visibility graph of a simple polygon*. Discrete Math. Theor. Comput. Sci. 17(1):1–12 (2015).
4 Sergio Cabello, Josef Cibulka, Jan Kynčl, Maria Saumell, Pavel Valtr. *Peeling potatoes near-optimally in near-linear time*. Proc. SoCG'14, pp. 224–231.

3.22 A Middle Curve Based on Discrete Fréchet Distance

Ludmila Scharf (FU Berlin, DE)

Joint work of Ahn, Hee-Kap; Alt, Helmut; Buchin, Maike; Scharf, Ludmila; Wenk, Carola
Main reference H.-K. Ahn, H. Alt, M. Buchin, L. Scharf, C. Wenk, "A Middle Curve Based on Discrete Fréchet Distance," to appear in Proc. of the 2015 European Workshop on Computational Geometry (EuroCG'15).

Given a set of polygonal curves we seek to find a "middle curve" that represents the set of curves. We ask that the middle curve consists of points of the input curves and that it minimizes the discrete Fréchet distance to the input curves. We develop algorithms for three different variants of this problem.

3.23 On Perturbations of the Expansion Cone

André Schulz (Universität Münster, DE)

Joint work of Schulz, André; Igamberdiev, Alexander

An expansive motion is an assignment of infinitesimal velocities to points in the plane such that all pairwise distances are (infinitesimal) nondecreasing. The space of the infinitesimal velocities forms a polyhedral cone. After a perturbation we obtain a polyhedron, whose corners represent geometric graphs induced by the tight inequalities. One set of perturbation parameters gives the polytope of pointed pseudo-triangulations. We reprove this result and show how a different set of parameters can be used to define a polyhedron whose corners represent a different class of planar Laman graphs. These graphs have no nonempty convex polygon (a necessary but not a sufficient condition). As a consequence we obtain a new description of the associahedron.

3.24 Topological Data Analysis

Donald Sheehy (University of Connecticut – Storrs, US)

URL http://donsheehy.net

I will present a top-down survey of some topics in topological data analysis (TDA). Consider the following model of data analysis.

$$U \longrightarrow (X \to R) \longrightarrow S$$

U is the universe, a population, or some "underlying" thing to be studied. The "data" comes in the form of real-valued functions on some (possibly unknown) space X. S is for signatures or summaries. A major goal of TDA is to define and compute signatures that are "topologically invariant" in the sense that

$$Sig(f(X)) = Sig(f(h(X)))$$

whenever h is a homeomorphism. I will show how many of the known results and many open research directions in TDA can be understood by systematically adding noise, error, discretization, or new hypotheses into this model.

3.25 Restricted Constrained Delaunay Triangulations

Jonathan Shewchuk (University of California – Berkeley, US)

Joint work of Shewchuk, Jonathan; Lévy, Bruno; Khoury, Marc; van Kreveld, Marc

The constrained Delaunay triangulation is a geometric structure that adapts the Delaunay triangulation to enforce the presence of specified edges. The restricted Delaunay triangulation is a geometric structure drawn on a smooth surface embedded in three-dimensional space, having properties similar to those of the Delaunay triangulation in the plane. We combine these two structures to address a question of Bruno Levy: can we define mathematically well-behaved constrained Delaunay triangulations on smooth surfaces?

We define the restricted constrained Delaunay triangulation to be the dual of a restricted extended Voronoi diagram, which is a generalization of the extended Voronoi diagram introduced by Raimund Seidel as a dual of the constrained Delaunay triangulation. The topological space on which we define the restricted extended Voronoi diagram is a 2-manifold created by cutting slits in the input surface (one slit for each specified edge constraint) and gluing two extrusions onto each slit. We define a metric on this 2-manifold that is similar to the three-dimensional Euclidean metric but is modified so that vertices on one "side" of an edge constraint cannot influence the portion of the Voronoi diagram on the other "side". The Voronoi diagram on the 2-manifold under this metric dualizes to a triangulation of the original surface if certain sampling conditions are met.

3.26 Beyond the Euler Characteristic: Approximating the Genus of General Graphs

Anastasios Sidiropoulos (Ohio State University – Columbus, US)

Main reference K.-i. Kawarabayashi, A. Sidiropoulos, "Beyond the Euler Characteristic: Approximating the Genus of General Graphs," in Proc. of the 47th Annual ACM Symp. on Theory of Computing (STOC'15), pp. 675–682, ACM, 2015; pre-print available at arXiv:1412.1792v1 [cs.DS].
URL http://dx.doi.org/10.1145/2746539.2746583
URL http://arxiv.org/abs/1412.1792v1

Computing the Euler genus of a graph is a fundamental problem in graph theory and topology. It has been shown to be NP-hard by [Thomassen 1989] and a linear-time fixed-parameter algorithm has been obtained by [Mohar 1999]. Despite extensive study, the approximability of the Euler genus remains wide open. While the existence of an $O(1)$-approximation is not ruled out, the currently best-known upper bound is a trivial $O(n/g)$-approximation that follows from bounds on the Euler characteristic.

In this paper, we give the first non-trivial approximation algorithm for this problem. Specifically, we present a polynomial-time algorithm which given a graph G of Euler genus g outputs an embedding of G into a surface of Euler genus $g^{O(1)}$. Combined with the above $O(n/g)$-approximation, our result also implies a $O(n^{1-\alpha})$-approximation, for some universal constant $\alpha > 0$.

Our approximation algorithm also has implications for the design of algorithms on graphs of small genus. Several of these algorithms require that an embedding of the graph into a surface of small genus is given as part of the input. Our result implies that many of these algorithms can be implemented even when the embedding of the input graph is unknown.

3.27 The Cosheaf Reeb-graph Interleaving Distance

Vin de Silva (Pomona College – Claremont, US)

Topological data analysis is typically carried out in a persistent framework [1]: a data set is converted to a filtered family of topological spaces, and the homological invariants of this system (rather than of any individual space in the family) are provably stable [2,3]. The family is typically parametrized by a real variable, which represents the scale at which the discrete data set is blurred to make it into a space.

Taking a more general view of persistence [4] as the study of functors on small sites and certain 'interleaving' relationships between them, we see that merge trees and Reeb graphs are susceptible to the same treatment. A merge tree can be viewed as a set-valued functor on the real line, and a Reeb greeph can be viewed as a set-valued cosheaf on the category of real intervals. In both cases there is defined an interleaving metric [5,6] that is provably stable with respect to perturbations of the initial data.

References

1 Herbert Edelsbrunner, David Letscher, Afra Zomorodian. *Topological persistence and simplification.* Discrete & Computational Geometry, 28:511–533 (2002).
2 David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. *Stability of persistence diagrams.* Discrete & Computational Geometry, 37(1):103–120 (2007).
3 Frederic Chazal, David Cohen-Steiner, Marc Glisse, Leonidas Guibas, Steve Oudot. *Proximity of persistence modules and their diagrams.* Proceedings of the 25th annual Symposium on Computational Geometry: 237–246 (2009).
4 Peter Bubenik and Jonathan A. Scott. *Categorification of persistent homology.* Discrete & Computational Geometry, 51(3):600–627 (2014).
5 Kenes Beketayev, Damir Yeliussizov, Dmitriy Morozov, Gunther Weber, Bernd Hamann. *Measuring the distance between merge trees.* Topological Methods in Data Analysis and Visualization III: Theory, Algorithms, and Applications, Mathematics and Visualization, 151–166 (2014).
6 Vin de Silva, Elizabeth Munch, Amit Patel. *Categorified Reeb Graphs.* Unpublished manuscript, arXiv:1501.04147 (2015).

3.28 Augmenting Embedded Paths and Trees to Optimize their Diameter

Fabian Stehn (Universität Bayreuth, DE)

We consider the problem of augmenting a graph with n vertices embedded in a metric space, by inserting one additional edge in order to minimize the diameter of the resulting graph. We present algorithms for the cases when the input graph is a path (running in $O(n \log^3 n)$ time) or a tree (running in $O(n^2 \log n)$ time). For the case when the input graph is a path in $\mathbb{R}^d$, where d is a constant, we present an algorithm that computes a $(1+\varepsilon)$-approximation in $O(n + 1/\varepsilon^3)$ time.

3.29 Flip Distances in Triangulations and Rectangulation

Csaba Toth (California State University – Northridge, US)

Joint work of Ackerman, Eyal; Allen, Michelle; Barequet, Gill; Cardinal, Jean; Hoffmann, Michael; Kusters, Vincent; Löffler, Maarten; Mermelstein, Joshua; Souvaine, Diane; Toth, Csaba; Wettstein, Manuel
Main reference J. Cardinal, M. Hoffmann, V. Kusters, C. D. Tóth, M. Wettstein, "Arc Diagrams, Flip Distances, and Hamiltonian Triangulations," in Proc. of the 32nd Symp. on Theoretical Aspects of Computer Science (STACS'15), LIPIcs, Vol. 30, pp. 197–210, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015.
URL http://dx.doi.org/10.4230/LIPIcs.STACS.2015.197

It is shown that every triangulation (maximal planar graph) on $n \geq 6$ vertices can be flipped into a Hamiltonian triangulation using a sequence of less than $n/2$ combinatorial edge flips. The previously best upper bound uses 4-connectivity as a means to establish Hamiltonicity. But in general about $3n/5$ flips are necessary to reach a 4-connected triangulation. Our result improves the upper bound on the diameter of the flip graph of combinatorial triangulations on n vertices from $5.2n - 33.6$ to $5n - 23$. We also show that for every triangulation on n vertices there is a simultaneous flip of less than $2n/3$ edges to a 4-connected triangulation. The bound on the number of edges is tight, up to an additive constant.

For n noncorectilinear points in a unit square $[0, 1]^2$, a rectangulation is a subdivision of $[0, 1]^2$ into $n + 1$ rectangles by n axis-aligned line segments, one passing through each point. It is shown that a sequence of $O(n \log n)$ elementary *flip* and *rotate* operations can transform any rectangulation to any other rectangulation on the same set of n points. This bound is the best possible for some point sets, while $\Theta(n)$ operations are sufficient and necessary for others.

3.30 Road Map Construction and Comparison

Carola Wenk (Tulane University, US)

Map construction is a new type of geometric reconstruction problem in which the task is to extract the underlying geometric graph structure described by a set of movement-constrained trajectories, or in other words reconstruct a geometric domain that has been sampled with continuous curves that are subject to noise. Due to the ubiquitous availability of geo-referenced trajectory data, the map construction task has widespread applications ranging from a variety of location-based services on street maps to the analysis of tracking data for hiking trail map generation or for studying social behavior in animals.

Several map construction algorithms have recently been proposed in the literature, however it remains a challenge to measure the quality of the reconstructed maps. We present an incremental map construction algorithm based on the Frechet distance. And we present different distance measures for comparing two road maps which amounts to comparing two uncertain embedded geometric graphs. One approach is based on comparing the set of paths in the graphs, and the other uses persistent homology of the offset filtration to compare the local topology of the graphs. We also introduce local signatures based on these distance measures, which allow us to identify regions where the maps differ the most.

3.31 Completely Randomized RRT-Connect: A Case Study on 3D Rigid Body Motion Planning

Nicola Wolpert (University of Applied Sciences – Stuttgart, DE)

Joint work of Schneider, Daniel; Schoemer, Elmar; Wolpert, Nicola
Main reference D. Schneider, E. Schömer, N. Wolpert, "Completely Randomized RRT-Connect: A Case Study on 3D Rigid Body Motion Planning," in Proc. of the 2015 IEEE Int'l Conf. on Robotics and Automation (ICRA'15), pp. 2944–2950, IEEE, 2015.
URL http://dx.doi.org/10.1109/ICRA.2015.7139602

Nowadays sampling-based motion planners use the power of randomization to compute multidimensional motions at high performance. Nevertheless the performance is based on problem-dependent parameters like the weighting of translation versus rotation and the planning range of the algorithm. Former work uses constant user-adjusted values for these parameters which are defined a priori. Our new approach extends the power of randomization by varying the parameters randomly during runtime. This avoids a preprocessing step to adjust parameters and moreover improves the performance in comparison to existing methods in the majority of the benchmarks. Our method is simple to understand and implement. In order to compare our approach we present a comprehensive experimental analysis about the parameters and the resulting performance. The algorithms and data structures were implemented in our own library RASAND, but we also compare the results of our work with OMPL and the commercial software KineoTM Kite Lab.

4 Open Problems

On Monday evening (19:15–20:30), March 9, 2015, we held an open problem discussion. The session scribe was Joe Mitchell and the session chair was Jeff Erickson. The problems span a range of topics, including fundamental algorithms, discrete geometry, combinatorics, and optimization.

▶ Problem 1 (Don Sheehy). *A metric problem:* Given n points P in $\mathbb{R}^d$. For a curve γ, define $len(\gamma) = \int_\gamma N(x)dx$, where $N(x)$ is the Euclidean distance from the nearest point of P to the point x. Let $d_N(p,q) = \inf_\gamma len(\gamma)$, where the infimum is over all paths starting at p and ending at q. Define $W_{ab} = (1/4)||a-b||^2$, for $a, b \in P$. Define $d_S(p,q) = \inf_{p=v_0,\ldots,v_k=q} \sum_i W_{v_{i-1}v_i}$, where the points v_i are all points of P. Conjecture: $d_N = d_S$.

Note that it is true for 2 points; this is the source of the "1/4" in the definition. From this it follows that the piecewise linear path γ that determines d_S has $len(\gamma) = d_S(p,q)$. So, for all $p, q \in P$, $d_S(p,q) \geq d_N(p,q)$. Moreover, it's easy to check that any edge traversed in the piecewise linear path determining d_S must be Gabriel, i.e. it must have a diametral ball empty of other points of P. This follows from the Pythagorean theorem, as any point inside the diametral ball would create a shortcut and, thus, a shorter path.

The problem is motivated by density-based distances. The metric d_N is a natural density-based distance arising from the nearest neighbor density estimator. We originally believed d_S would be a good approximation, but never found an example where they differ.

▶ Problem 2 (Jeff Erickson). *A question in elementary topology:* Any generic closed curve in the plane can be continuously deformed into a simple closed curve through a series of elementary local transformations resembling Reidemeister moves:

- Remove an empty loop: $\propto \Rightarrow ($
- Remove an empty bigon: $\between \Rightarrow)($
- Flip an empty triangle: $\forall \Rightarrow \rotatebox[origin=c]{180}{$\forall$}$

How many moves are required in the worst case, as a function of the number of self-intersection points? A proof of Steinitz's theorem[1] by Grünbaum[2] yields an $O(n^2)$ upper bound. A more recent algorithm of Feo and Provan[3] yields an upper bound of $O(nD)$ moves, where D is the diameter of the graph. On the other hand, the $\sqrt{n} \times (\sqrt{n}+1)$ "torus knot" curve provably requires at least $\binom{\sqrt{n}}{3} = \Omega(n^{3/2})$ moves. I conjecture that the lower bound is tight.

▶ Problem 3 (Tamal Dey). *Deciding triviality of cycles:* Let K be a finite simplicial complex linearly emedded in $\mathbb{R}^3$. Let C be any given 1-cycle in K. We are interested in detecting if C is *trivial* in the first homology group, that is, if there is a set of triangles in K whose boundaries when summed over $\mathbb{Z}_2$ give C. This problem can be solved in $O(M(n))$ time by first reducing the boundary matrix of K (triangle-edge matrix) to Echelon form and then reducing a column corresponding to C to see if it becomes empty column or not. Here $M(n)$ is the matrix multiplication time whose current best bound is $O(n^{2.37..})$.

▶ Conjecture 1. Let K be a finite simplicial complex linearly embedded in $\mathbb{R}^3$ with a total of n simplices. Given a 1-cycle C in K, one can detect if C is trivial in the first homology group (with $\mathbb{Z}_2$ coefficient) in $O(n^2)$ time.

If K is a 2-manifold, the detection can be performed in $O(n)$ time by a simple depth-first walk in K. If K is a 3-manifold with connected boundary, the algorithm in "An efficient computation of handle and tunnel loops via Reeb graphs [D.-Fan-Wang] *ACM Trans. Graphics (SIGGRAPH 2013)*, Vol. 32(4), 2013" can be modified to accomplish the task in $O(n^2)$ time. The question remains open for general simplicial complexes. Although, the conjecture is posed here for K embedded in $\mathbb{R}^3$ and for a 1-cycle C, it can be posed for a finite simplicial complex embedded linearly in $\mathbb{R}^d$ and a given p-cycle C in it.

▶ Problem 4 (Nina Amenta). *A problem of unique polyhedron determination:* Let P be a simplicial (triangulated) three-dimensional polyhedron, not necessarily convex. Given the combinatorial structure of P, that is, the graph of its 1-skeleton, and the dihedral angle at every edge. Assume the dihedrals are all bounded away from 0, although they could be positive (convex) or negative (concave). Does this uniquely determine the vertex positions (up to rotation, translation, scale)? (Mazzeo and Montcouquiol, 2011, Journal of Differential Geometry, proved that uniqueness holds for convex polyhedra; highly nontrivial proof.)

▶ Problem 5 (Michael Gene Dobbins). *Realizing order types by k-gons:* We say an arrangement of convex bodies is orientable when the bodies do not pair-wise cross (each pair of bodies has exactly 2 common supporting tangents) and among every three bodies, each body appears exactly once on the boundary of their convex hull. We define the order type of an orientable arrangement as the orientation of each triple of bodies: $(+)$ if the bodies appear in counter-clockwise order around the boundary of their convex hull, and $(-)$ if they appear in clockwise order.

For a fixed integer k, how complicated can the set of arrangements of k-gons of a fixed order type be?

---

[1] Every 3-connected planar graph is the 1-skeleton of a 3-polytope, and vice versa.
[2] Branko Grünbaum. *Convex Polytopes.* John Wiley & Sons, 1967.
[3] Thomas A. Feo and J. Scott Provan. Delta-wye transformations and the efficient reduction of two-terminal planar graphs. *Operations Research* 41(3):572–582, 1993.

With Andreas Holmsen and Alfredo Hubard, we were able to show that the k-gon realization space of an arrangement can have the homotopy type of any primary semialgebraic set, but the arrangement used for this construction was not orientable. Orientable arrangements are a natural class of arrangements to consider, since the orientations on triples in such an arrangement satisfy the chirotope axioms, and as such are more closely related to configurations of points. We conjecture that universality also holds for orientable arrangements. That is, we conjecture that the set of arrangements of k-gons of a fixed order type modulo projectivities can have the homotopy type of any primary semialgebraic set.

▶ Problem 6 (Joe Mitchell). *Two problems:* (a) Given n points in $\mathbb{R}^3$ in general position, is it always the case that there exists a triangulation (tetrahedralization) of S whose dual graph is Hamiltonian? (The dual graph has a node for each tetrahedron, and an edge between facet-sharing tetrahedra. We look for a Hamiltonian path.) In $\mathbb{R}^2$ it is always the case that a Hamiltonian triangulation exists. In $\mathbb{R}^3$ it suffices to consider points in convex position (after which, if a Hamiltonian triangulation is found, the interior points can be inserted, one by one, and the corresponding tetrahedra repartitioned to maintain Hamiltonicity).

(b) Given a unit-radius ball ("planet") in $\mathbb{R}^3$, find a minimum-length set X (path, cycle, or tree), outside the ball, such that X does not penetrate the interior of the ball and all of the surface of the ball is illuminated by X. The shortest known path (see SoCG video paper Timothy M. Chan, Alexander Golynski, Alejandro López-Ortiz, Claude-Guy Quimper, "The asteroid surveying problem and other puzzles". SoCG 2003:372-373) consists of a union of two segments and a connecting spiral curve; the shortest known cycle is the "baseball curve" consisting of 4 semicircles on the surface of the bounding cube; is the shortest tree any different from the shortest path?

▶ Problem 7 (Michael Kerber). *A problem of well centeredness:* A d-simplex σ in $\mathbb{R}^d$ is *well-centered* if the circumsphere of σ is inside $CH(\sigma)$. Is there a point set P of n points in $\mathbb{R}^d$ such that the Delaunay diagram of P has at least $c \cdot n^{\lceil d/2 \rceil}$ well-centered d-simplices? What if $d = 3$?

(Related to Pitteway triangulations.)

▶ Problem 8 (Joe Mitchell). *The guarding game:* In 2014 I posed the "guarding game": For a given set S of n points in the plane, player 1 (the "guarder") is to pick a subset, G, of S, of size $k = |G|$, at which he places guards; separately, without seeing what play 1 does, play 2 (the "polygonalizer") is to give a simple polygonalization, P, of S (the set S is the vertex set of P). The guarder wins if G guards P; otherwise, the polygonalizer wins. What is a reasonable value for k (as a function of n, or possibly of the number, i, of points of S interior to $CH(S)$) to make the game close to "fair"? What is the best strategy for each player?

Participants

- Peyman Afshani
Aarhus University, DK
- Annamaria Amenta
Univ. of California – Davis, US
- Franz Aurenhammer
TU Graz, AT
- Maike Buchin
Ruhr-Universität Bochum, DE
- Sergio Cabello
University of Ljubljana, SI
- Siu-Wing Cheng
HKUST – Kowloon, HK
- Otfried Cheong
KAIST – Daejeon, KR
- Jinhee Chun
Tohoku University – Sendai, JP
- Mark de Berg
TU Eindhoven, NL
- Vin de Silva
Pomona College – Claremont, US
- Olivier Devillers
INRIA Nancy – Grand Est, FR
- Tamal K. Dey
Ohio State University – Columbus, US
- Michael Gene Dobbins
Postech – Pohang, KR
- Anne Driemel
TU Eindhoven, NL
- Ioannis Z. Emiris
University of Athens, GR
- Jeff Erickson
Univ. of Illinois – Urbana, US
- Jie Gao
SUNY – Stony Brook, US
- Marc Glisse
INRIA Saclay – Île-de-France, FR
- Leonidas J. Guibas
Stanford University, US
- Michael Kerber
MPI für Informatik – Saarbrücken, DE
- David G. Kirkpatrick
University of British Columbia – Vancouver, CA
- Rolf Klein
Universität Bonn, DE
- Joseph S. B. Mitchell
SUNY – Stony Brook, US
- Wolfgang Mulzer
FU Berlin, DE
- Elizabeth Munch
University of Albany, US
- Yota Otachi
JAIST – Ishikawa, JP
- Jeff M. Phillips
University of Utah – Salt Lake City, US
- Natan Rubin
Ben Gurion University – Beer Sheva, IL
- Vera Sacristan
UPC – Barcelona, ES
- Maria Saumell
University of West Bohemia – Pilsen, CZ
- Ludmila Scharf
FU Berlin, DE
- André Schulz
Universität Münster, DE
- Raimund Seidel
Universität des Saarlandes, DE
- Donald Sheehy
University of Connecticut – Storrs, US
- Jonathan Shewchuk
University of California – Berkeley, US
- Anastasios Sidiropoulos
Ohio State University – Columbus, US
- Fabian Stehn
Universität Bayreuth, DE
- Monique Teillaud
INRIA Nancy – Grand Est, FR
- Csaba Toth
California State University – Northridge, US
- Carola Wenk
Tulane University, US
- Nicola Wolpert
University of Applied Sciences – Stuttgart, DE