Report from Dagstuhl Seminar 17171

Computational Geometry

Edited by

Otfried Cheong[1], Anne Driemel[2], and Jeff Erickson[3]

1 **KAIST – Daejeon, KR,** otfried@kaist.edu
2 **TU Eindhoven, NL,** a.driemel@tue.nl
3 **University of Illinois – Urbana-Champaign, US,** jeffe@illinois.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17171 "Computational Geometry". The seminar was held from 23rd to 28th April 2017 and 47 participants from various countries attended it. Recent advances in computational geometry were presented and new challenges were identified. The report collects the abstracts of talks and open problems presented in the seminar.

Seminar April 23–28, 2017 – http://www.dagstuhl.de/17171
1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity, G.2 Discrete Mathematics, G.4 Mathematical Software
Keywords and phrases algorithms, applications, combinatorics, complexity, geometric computing, high-dimensional computational geometry, implementation, monitoring and shape data
Digital Object Identifier 10.4230/DagRep.7.4.107
Edited in cooperation with Marcel Roeloffzen

1 Executive Summary

Anne Driemel
Otfried Cheong
Jeff Erickson

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, CAD, robotics, computer vision, image processing, spatial databases, GIS, 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 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 new results in the field, the seminar included broad survey talks on monitoring and shape data and on high-dimensional geometric computing, two focus areas that have seen exciting recent progress and that present numerous opportunities for further cross-disciplinary impact.

Computational geometry for monitoring and shape data

The combination of movement and geometry has always been an important topic in computational geometry, initially motivated by robotics and resulting in the study of kinetic data structures. With the advent of widely available location tracking technologies such as GPS sensors, trajectory analysis has become a topic in itself, which has connections to other classical topics in computational geometry such as shape analysis. Still, efficient technologies to perform the most basic operations are lacking. We need data structures supporting similarity queries on trajectory data and geometric clustering algorithms that can handle the infinite-dimensional geometry inherent in the data. A related type of data, namely time series data, has not received much attention in the computational geometry community, despite its universality and its close relation to trajectory data. Shedding light on the interconnections of these topics will promote new results in the field which will address these timely questions.

Computing in high-dimensional and infinite-dimensional spaces

The famous “curse of dimensionality” prevents exact geometric computations in high-dimensional spaces. Most of the data in science and engineering is high-dimensional, rendering classical geometric techniques, such as the sweepline approach, insufficient. One way to address this issue is to use sparsity, but it is not always easy to find a sparse representation of the data. The search of the most efficient representation and how to exploit this representation leads to dimension-reduction techniques, metric embeddings, and approximation algorithms. This line of research has strong ties to machine learning and discrete mathematics as well as computational geometry.

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 third 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. The seminar has now a more balanced attendance in terms of seniority and gender than in the past. This year, 47 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

Open problems

3 Overview of Talks

3.1 Graph Embedding While Preserving Pairwise Distances

Tetsuo Asano (JAIST – Ishikawa, JP)

Main reference T. Cox, M. Cox, "Multidimensional Scaling Second Edition", Chapman & Hall CRC, 2001.

In this talk we consider a distance preserving graph embedding problem: Given a weighted graph, embed vertices into points in d-space so that for each edge the distance between their corresponding endpoints is as close as possible to the weight of the edge. If it is known that there is an exact embedding (without any error) for a full matrix for all pairs of vertices, then the existing algorithm known as Principal Coordinate Analysis (PCO) can find such an exact embedding in polynomial time. Although it is believed that PCO almost always gives a good solution, it is not true. We show a worst example for PCO and compare it with a heuristic algorithm. We also consider a special case where every pair of points has the same distance. We show the problem can be solved in one dimension. We also show some experimental results in two dimensions which reflect the results in one dimension.

3.2 Improved Time-Space Trade-offs for Computing Voronoi Diagrams

Bahareh Banyassady (FU Berlin, DE), Matias Korman (Tohoku University – Sendai, JP), Wolfgang Mulzer (FU Berlin, DE), Marcel Roeloffzen, Paul Seiferth, Yannik Stein, and André van Renssen (National Institute of Informatics – Tokyo, JP)

Joint work of Bahareh Banyassady, Matias Korman, Wolfgang Mulzer, André van Renssen, Marcel Roeloffzen, Paul Seiferth, Yannik Stein
Main reference B. Banyassady, M. Korman, W. Mulzer, A. van Renssen, M. Roeloffzen, P. Seiferth, Y. Stein, "Improved Time-Space Trade-offs for Computing Voronoi Diagrams", in Proc. of the 34th Symposium on Theoretical Aspects of Computer Science (STACS 2017), LIPIcs, Vol. 66, pp. 9:1–9:14, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.
URL http://dx.doi.org/10.4230/LIPIcs.STACS.2017.9

In this talk, we are interested in computing various Voronoi diagrams for a given set of sites P in the memory-constrained model. In this model, we assume that the input is in a read-only array and the algorithm may use an additional workspace memory of size $O(s)$ words, $s \in \{1, \dots, n\}$, for reading and writing intermediate data, and the output is write-only. Clearly, when s increases, the running time of the algorithm decreases. In this model, we provide a time-space trade-off for computing the nearest site Voronoi diagram and the farthest site Voronoi digram with running time $O(n^2/s \log s)$ using $O(s)$ words of workspace. Furthermore, we extend this result to compute the family of all higher-order Voronoi diagrams, up to a given order $K \in \{1, \dots, O(\sqrt{s})\}$, in a pipelined fashion in $O\big(\frac{n^2 K^5}{s}(\log s + K \log K)\big)$ time using $O(s)$ words of workspace. The main idea is to use edges of each diagram to compute the edges of the next diagram. However, this needs to be coordinated carefully, in order to keep the bound on the used memory and to prevent edges from being reported multiple times.

3.3 The Morse theory of Cech and Delaunay Complexes

Ulrich Bauer (TU München, DE)

Joint work of Herbert Edelsbrunner, Ulrich Bauer
Main reference U. Bauer, H. Edelsbrunner, "The Morse theory of Čech and Delaunay complexes", Trans. Amer. Math. Soc. 369(2017), pp. 3741–3762, 2016.
URL http://dx.doi.org/10.1090/tran/6991

Given a finite set of points in R^n and a radius parameter, we study the Čech, Delaunay-Čech, Delaunay (or alpha), and Wrap complexes in the light of generalized discrete Morse theory. Establishing the Čech and Delaunay complexes as sublevel sets of generalized discrete Morse functions, we prove that the four complexes are simple-homotopy equivalent by a sequence of simplicial collapses, which are explicitly described by a single discrete gradient field.

3.4 Conditional Lower Bounds for Similarity Measures on Curves and Strings

Karl Bringmann (MPI für Informatik – Saarbrücken, DE)

Joint work of Marvin Künnemann, Wolfgang Mulzer, Amir Abboud, Karl Bringmann
Main reference K. Bringmann, "Why Walking the Dog Takes Time: Frechet Distance Has No Strongly Subquadratic Algorithms Unless SETH Fails", in Proc. of the 55th Annual Symp. on Foundations of Computer Science (FOCS 2014), pp. 661–670, IEEE, 2014.
URL https://doi.org/10.1109/FOCS.2014.76

This talk is an introduction to the area of Fine-grained Complexity, where we show running time lower bounds conditional on certain conjectures on classic problems, such as the Strong Exponential Time Hypothesis for the Satisfiability problem. Specifically, we survey lower bounds known for similarity measures on curves (e.g. Frechet distance) and strings (e.g. longest common subsequence and edit distance), and we discuss many recent extensions.

3.5 Untangling and Unwinding Curves

Jeff Erickson (University of Illinois – Urbana-Champaign, US)

Joint work of Hsien-Chih Chang, Jeff Erickson
Main reference H.-C. Chang, J. Erickson, "Untangling Planar Curves", in Proc. of the 32nd Int'l Symp. on Computational Geometry (SoCG 2016), LIPIcs, Vol. 51, pp. 29:1–29:16, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
URL https://doi.org/10.4230/LIPIcs.SoCG.2016.29

Any closed curve in the plane can be transformed into a simple closed curve using a finite sequence of local transformations called homotopy moves. We prove that $\Theta(n^{3/2})$ homotopy moves are necessary and sufficient in the worst case, improving the previous best $O(n^2)$ upper bound due to Steinitz in 1916 and the previous best $\Omega(n)$ lower bound, which is trivial. We also prove that $\Omega(n^2)$ moves are necessary in the worst case to simplify a contractible curve in the annulus, and therefore in any surface with non-positive Euler characteristic; a matching $O(n^2)$ upper bound follows from Steinitz's planar results and more recent work by Hass and Scott.

3.6 A Nearly Quadratic Bound for the Decision Tree Complexity of k-SUM

Esther Ezra (Georgia Institute of Technology – Atlanta, US)

Joint work of Esther Ezra, Micha Sharir

We show that the k-SUM problem can be solved by a linear decision tree of depth $O(n^2 \log^2 n)$, improving the recent bound $O(n^3 \log^3 n)$ of Cardinal etal. Our bound depends linearly on k, and allows us to conclude that the number of linear queries required to decide the n-dimensional Knapsack or SubsetSum problems is only $O(n^3 \log n)$, improving the currently best known bounds by a factor of n. Our algorithm extends to the RAM model, showing that the k-SUM problem can be solved in expected polynomial time, for any fixed k, with the above bound on the number of linear queries. Our approach relies on a new point-location mechanism, exploiting "Epsilon-cuttings" that are based on vertical decompositions in hyperplane arrangements in high dimensions. A major side result of our analysis is a sharper bound on the complexity of the vertical decomposition of such an arrangement (in terms of its dependence on the dimension).

3.7 Computing Optimal Flight Patterns with Minimum Turn Cost

Sándor Fekete (TU Braunschweig, DE)

Joint work of Aaron Becker, Mustapha Debboun, Dominik Krupke, An Nguyen, Sándor Fekete
URL https://youtu.be/SFyOMDgdNao

We present results arising from the problem of sweeping a mosquito-infested area with an Unmanned Aerial Vehicle (UAV) equipped with an electrified metal grid. Planning good trajectories is related to a number of classic problems of geometric optimization, in particular the Traveling Salesman Problem, the Lawn Mower Problem and, most closely, Milling with Turn Cost. We describe how planning a good trajectory can be reduced to considering penalty and budget variants of covering a grid graph with minimum turn cost. On the theoretical side, we show the solution of a problem from The Open Problems Project that had been open for more than 15 years, and hint at approximation algorithms. On the practical side, we describe an exact method based on Integer Programming that is able to compute provably optimal instances with over 500 pixels. These solutions are actually used for practical trajectories, as demonstrated in a video.

3.8 Faster Algorithms for the Geometric Transportation Problem

Kyle Jordan Fox (Duke University – Durham, US)

Joint work of Pankaj K. Agarwal, Kyle Fox, Debmalya Panigrahi, Kasturi R. Varadarajan, Allen Xiao
Main reference P. K. Agarwal, K. Fox, D. Panigrahi, K. R. Varadarajan, A. Xiao, "Faster algorithms for the geometric transportation problem", in Proc. of the 33rd International Symposium on Computational Geometry (SoCG 2017), LIPIcs, Vol. 77, pp. 7:1–7:15, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.
URL https://doi.org/10.4230/LIPIcs.SoCG.2017.7

Let $R, B \subset R^d$, for constant d, be two point sets with $|R| + |B| = n$, and let $\lambda : R \cup B \to N$ such that $\sum_{r \in R} \lambda(r) = \sum_{b \in B} \lambda(b)$ be demand functions over R and B. Let $d(\cdot, \cdot)$ be a suitable distance function such as the L_p distance. The transportation problem asks to find a map $\tau : R \times B \to N$ such that $\sum_{b \in B} \tau(r, b) = \lambda(r), \sum_{r \in R} \tau(r, b) = \lambda(b)$, and $\sum_{r \in R, b \in B} \tau(r, b) d(r, b)$ is minimized. We present three new results for the transportation problem when $d(\cdot, \cdot)$ is any L_p metric:

- For any constant $\epsilon > 0$, an $O(n^{1+\epsilon})$ expected time randomized algorithm that returns a transportation map with expected cost $O(\log^2(1/\epsilon))$ times the optimal cost.
- For any $\epsilon > 0$, a $(1 + \epsilon)$-approximation in $O(n^{3/2}\epsilon^d \text{polylog}(U) \text{polylog}(n))$ time, where $U = \max_{p \in R \cup B} \lambda(p)$.
- An exact strongly polynomial $O(n^2 \text{polylog}(n))$ time algorithm, for $d = 2$.

3.9 Don't Collect too Much: Geometric Approaches for Protecting Location and Trajectory Privacy

Jie Gao (Stony Brook University, US)

Joint work of Jiaxin Ding, Chien Chun Ni, Mengyu Zhou, Xiaotian Yin, Wei Han, Dengpan Zhou, David Gu, Jie Gao
Main reference J. Ding, C.-C. Ni, M. Zhou, J. Gao, "MinHash hierarchy for privacy preserving trajectory sensing and query", in Proc. of the 16th ACM/IEEE Int'l Conf. on Information Processing in Sensor Networks, pp. 17–28, ACM, 2017.
URL http://doi.acm.org/10.1145/3055031.3055076
Main reference X. Yin, C.-C. Ni, J. Ding, W. Han, D. Zhou, J. Gao, X. D. Gu, "Decentralized human trajectories tracking using hodge decomposition in sensor networks", in Proc. of the 23rd SIGSPATIAL Int'l Conf. on Advances in Geographic Information Systems (SIGSPATIAL '15), pp. 54:1–54:4, ACM, 2015.
URL http://doi.acm.org/10.1145/2820783.2820844

Large amounts of geometric data are becoming available due to recent technology advances in sensing, communication and computation. GPS traces, location registrations, sensor data in smart environments can be used to infer human behaviors and patterns. While we celebrate the enormous learning opportunities these have enabled, the learned patterns may reveal sensitive or personal identifying information. In this talk I will describe geometric methods to protect location and trajectory information. Our methods focus on reducing location/trajectory data collected by sensors to meet the target privacy requirements using differential forms, Hodge decomposition and MinHash schemes.

3.10 Spatio-Temporal Analysis of Team Sports and Computational Geometry

Joachim Gudmundsson (The University of Sydney, AU)

Joint work of Michael Horton, Joachim Gudmundsson
Main reference J. Gudmundsson, M. Horton, "Spatio-Temporal Analysis of Team Sports", ACM Computing Surveys, Vol. 50(2): 22:1-22:34, ACM, 2017.
URL http://doi.acm.org/10.1145/3054132

Team-based invasion sports such as football, basketball and hockey are similar in the sense that the players are able to move freely around the playing area; and that player and team performance cannot be fully analysed without considering the movements and interactions of all players as a group. State of the art object tracking systems now produce spatio-temporal traces of player trajectories with high definition and high frequency, and this, in turn, has facilitated a variety of research efforts, across many disciplines, to extract insight from the trajectories. In this talk we focus on some of the geometric problems that arise in the area.

3.11 High-dimensional Theta Numbers

Anna Gundert (Universität Köln, DE)

Joint work of Christine Bachoc, Anna Gundert, Alberto Passuello
Main reference C. Bachoc, A. Gundert, A. Passuello, "The Theta Number of Simplicial Complexes", arXiv:1704.01836 [math.CO], 2017.
URL https://arxiv.org/abs/1704.01836

We introduce a generalization of the celebrated Lovász theta number of a graph to simplicial complexes of arbitrary dimension. Our generalization takes advantage of real simplicial cohomology theory, in particular combinatorial Laplacians, and provides a semidefinite programming upper bound of the independence number of a simplicial complex. We consider properties of the graph theta number such as the relationship to Hoffman's ratio bound and to the chromatic number and study how they extend to higher dimensions. Furthermore, we analyze the value of the theta number for dense random simplicial complexes.

3.12 Minkowski Sums of Polyhedra with Holes

Dan Halperin (Tel Aviv University, IL)

Joint work of Alon Baram, Efi Fogel, Dan Halperin, Michael Hemmer, Sebastian Morr

The Minkowski sum of two sets P and Q in Euclidean space is the result of adding every point (position vector) in P to every point in Q. Considering the Minkowski sum of two polyhedra with holes, we show that one can always fill up the holes in one of the summand polyhedra and still get the same Minkowski sum as of the original summands. We present a simple proof of this observation, improving on (our) earlier rather involved proof of a more restricted claim. As we explain, this observation helps in speeding up the computation of Minkwoski sums in practice. We also review additional recent results in computing and using Minkowksi sums.

3.13 Self-Aligning Shapes

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

Joint work of Ashwin Gopinath, David G. Kirkpatrick, Paul W. K. Rothemund and Chris Thachuk
Main reference A. Gopinath, D. G. Kirkpatrick, P. W. K. Rothemund, Chris Thachuk, "Progressive Alignment of Shapes", in Proc. of the 28th Canadian Conference on Computational Geometry (CCCG 2016), pp. 230–236, Simon Fraser University, 2016.
URL http://www.cccg.ca/proceedings/2016/proceedings2016.pdf

A planar shape S is said to be self-aligning if any two overlapping, but otherwise arbitrarily placed, copies of S, can be brought into a unique configuration of maximum overlap by a continuous motion that monotonically increases their overlap. The identification of self-aligning shapes is motivated by applications of self-assembly, driven by molecular forces, in nano-fabrication processes. We are interested in the design and certification of shapes that are self-aligning, but also satisfy certain other constraints that arise due to fabrication issues.

3.14 On the Computational Bottleneck in Sampling-Based Robot Motion Planning

Michal Kleinbort (Tel Aviv University, IL)

Joint work of Michal Kleinbort, Oren Salzman, Dan Halperin
Main reference M. Kleinbort, O. Salzman, D. Halperin, "Collision detection or nearest-neighbor search? On the computational bottleneck in sampling-based motion planning", arXiv:1607.04800v3 [cs.RO], 2016.
URL https://arxiv.org/abs/1607.04800v3

Many sampling-based motion planning algorithms rely heavily on two main components: (i) collision detection and (ii) nearest-neighbor search. The complexity of the latter dominates the asymptotic running time of such algorithms. However, collision detection is often considered to be the computational bottleneck in practice. We describe settings in which the practical computational role of nearest-neighbor search is far from being negligible, i.e., the portion of running time taken up by nearest-neighbor search is comparable to, or sometimes even greater than the portion of time taken up by collision detection. We show that by using efficient, specially-tailored nearest-neighbor data structures, the overall running time of certain motion-planning algorithms in such settings can be significantly reduced.

3.15 Computing Wave Impact in Self-Organised Mussel Beds

Maarten Löffler (Utrecht University, NL)

We model the effects of byssal connections made by mussels within patterned mussel beds on bed stability as a disk graph, and propose a formula for assessing which mussels, if any, would get dislodged from the bed under the impact of a wave. We formulate the computation as a flow problem, giving access to efficient algorithms to evaluate the formula. We then analyse the geometry of the graph, and show that we only need to compute a maximum flow in a restricted part of the graph, giving rise to a near-linear solution in practice.

3.16 The Art Gallery Problem is ETR-complete

Tillmann Miltzow (Free University of Brussels, BE)

Joint work of Mikkel Abrahamsen, Anna Adamaszek, Tillmann Mitlzow
Main reference M. Abrahamsen, A. Adamaszek, T. Miltzow, "The Art Gallery Problem is ∃ℝ-complete", arXiv:1704.06969 [cs.CG], 2017.
URL https://arxiv.org/abs/1704.06969

We prove that the *art gallery problem* is equivalent under polynomial time reductions to deciding whether a system of polynomial equations over the real numbers has a solution. The art gallery problem is a classical problem in computational geometry, introduced in 1973 by Viktor Klee. Given a simple polygon P and an integer k, the goal is to decide if there exists a set G of k *guards* within P such that every point $p \in P$ is seen by at least one guard $g \in G$. Each guard corresponds to a point in the polygon P, and we say that a guard g *sees* a point p if the line segment pg is contained in P.

The art gallery problem has stimulated a myriad of research in geometry and in algorithms. However, despite extensive research, the complexity status of the art gallery problem has not been resolved. It has long been known that the problem is NP-hard, but no one has been able to show that it lies in NP. Recently, the computational geometry community became more aware of the complexity class $\exists\mathbb{R}$. The class $\exists\mathbb{R}$ consists of problems that can be reduced in polynomial time to the problem of deciding whether a system of polynomial equations with integer coefficients and any number of real variables has a solution. It can be easily seen that $\mathrm{NP} \subseteq \exists\mathbb{R}$. We prove that the art gallery problem is $\exists\mathbb{R}$-complete, implying that (1) any system of polynomial equations over the real numbers can be encoded as an instance of the art gallery problem, and (2) the art gallery problem is not in the complexity class NP unless $\mathrm{NP} = \exists\mathbb{R}$. As a corollary of our construction, we prove that for any real algebraic number α there is an instance of the art gallery problem where one of the coordinates of the guards equals α in any guard set of minimum cardinality. That rules out many geometric approaches to the problem.

3.17 TSP With Locational Uncertainty: The Adversarial Model

Joseph S. B. Mitchell (Stony Brook University, US)

Joint work of Gui Citovsky, Tyler Mayer, Joseph S. B. Mitchell

In this paper we study a natural special case of the Traveling Salesman Problem (TSP) with point-locational-uncertainty which we will call the *adversarial TSP* problem (ATSP). Given a metric space (X, d) and a set of subsets $R = \{R_1, R_2, ..., R_n\} : R_i \subseteq X$, the goal is to devise an ordering of the regions, σ_R, that the tour will visit such that when a single point is chosen from each region, the induced tour over those points in the ordering prescribed by σ_R is as short as possible. Unlike the classical locational-uncertainty-TSP problem, which focuses on minimizing the expected length of such a tour when the point within each region is chosen according to some probability distribution, here, we focus on the *adversarial model* in which once the choice of σ_R is announced, an adversary selects a point from each region in order to make the resulting tour as long as possible. In other words, we consider an offline problem in which the goal is to determine an ordering of the regions R that is optimal with respect to the “worst” point possible within each region being chosen by an adversary, who knows the chosen ordering. We give a 3-approximation when R is a set of arbitrary regions/sets of points in a metric space. We show how geometry leads to improved constant factor approximations when regions are parallel line segments of the same lengths, and a polynomial-time approximation scheme (PTAS) for the important special case in which R is a set of disjoint unit disks in the plane.

3.18 Dynamic Planar Voronoi Diagrams for General Distance Functions

Wolfgang Mulzer (FU Berlin, DE), Haim Kaplan, Liam Roditty, Paul Seiferth, and Micha Sharir

Main reference H. Kaplan, W. Mulzer, L. Roditty, P. Seiferth, M. Sharir, “Dynamic Planar Voronoi Diagrams for General Distance Functions and their Algorithmic Applications”, in Proc. of the 28th ACM-SIAM Symp. on Discrete Algorithms (SODA 2017), pp. 2495–2504, ACM, 2017.
URL http://dx.doi.org/10.1137/1.9781611974782.165

We describe a new data structure for dynamic nearest neighbor queries in the plane with respect to a general family of distance functions that includes L_p-norms and additively weighted Euclidean distances, and for general (convex, pairwise disjoint) sites that have constant description complexity (line segments, disks, etc.). Our data structure has a polylogarithmic update and query time, improving an earlier data structure of Agarwal, Efrat and Sharir that required $O(n^\varepsilon)$ time for an update and $O(\log n)$ time for a query.

3.19 Maximizing Volume Subject to Combinatorial Constraints

Aleksandar Nikolov (University of Toronto, CA)

The maximum volume j-simplex problem asks to compute the j-dimensional simplex of maximum volume inside the convex hull of a given set of n points in d-dimensional space. We discuss a deterministic approximation algorithm for this problem which achieves an approximation ratio of $e^{j/2+o(j)}$ in time polynomial in n and d. The problem is known to be NP-hard to approximate within a factor of c^j for some constant $c > 1$. Our algorithm also gives a factor $e^{j+o(j)}$ approximation for the problem of finding the principal $j \times j$ submatrix of a rank d positive semidefinite matrix with the largest determinant. In the final part of the talk we sketch how these results have been generalized to maximizing volume or determinants subject to more complicated combinatorial constraints. Such problems appear naturally in machine learning and experimental design when dealing with models of diversity.

3.20 Colorful simplicial depth, Minkowski sums, and generalized Gale transforms

Zuzana Patáková (IST Austria – Klosterneuburg, AT)

Joint work of Karim Adiprasito, Philip Brinkmann, Arnau Padrol, Pavel Paták, Zuzana Patáková, Raman Sanyal
Main reference K. Adiprasito, P. Brinkmann, A. Padrol, P. Paták, Z. Patáková, R. Sanyal, "Colorful simplicial depth, Minkowski sums, and generalized Gale transforms", arXiv:1607.00347 [math.CO], 2016.
URL https://arxiv.org/abs/1607.00347

Given $d+1$ sets $S_1, S_2, \dots, S_{d+1}$ (called color classes) in $\mathbb{R}^d$, a simplex is called *colorful*, if all its vertices are in different color classes. The number of colorful simplices containing a point $p \in \mathbb{R}^d$ is known as the *colorful simplicial depth of* p. The first result concerning colorful simplicial depth in discrete geometry was the colorful Carathéodory's theorem by Imre Bárány in 1982: "Any point $p \in \mathbb{R}^d$ contained in the convex hull of all color classes has a non-zero simplicial depth provided that each color class has at least $d+1$ points."

In 2006 Deza, Huang, Stephen, and Terlaky asked for the minimal and maximal values of the colorful simplicial depth of the point p in colorful Carathéodory's theorem. We use methods from combinatorial topology to prove a tight upper bound of the form $1 + \prod_{i=1}^{d+1}(|S_i| - 1)$.

The second goal of this talk is to highlight a connection between colorful configurations and faces of Minkowski sums. Considering the Gale transform of the Cayley embedding, we define *colorful Gale transforms* associated to a collection of convex polytopes that capture the facial structure of Minkowski sums in the combinatorics of colorful configurations. This dictionary between Minkowski sums and colorful configurations allows us to resolve a conjecture of Ben Burton from the *normal surface* theory.

3.21 Approximate Range Maximization

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

Joint work of Michael Matheny, Raghvendra Singh, Kaiqiang Wang, Liang Zhang, Jeff M. Phillips
Main reference M. Matheny, R. Singh, K. Wang, L. Zhang, J. M. Phillips, "Scalable Spatial Scan Statistics through Sampling", in Proc. of the 24th ACM SIGSPATIAL Int'l Conf. on Advances in Geographic Information Systems, (GIS 2016), pp. 20:1–20:10, ACM, 2016.
URL http://doi.acm.org/10.1145/2996913.2996939

Consider a geometric range space (X, A) where each data point $x \in X$ has two or more values (say r and b). Also consider a function $\phi(A)$ defined on any subset C in (X, A) on the sum of values in that range e.g., $r_C = \sum_{x \in C} r(x)$ and $b_C = \sum_{x \in C} b(x)$. The maximum range is $A^* = arg\max_{X in(X,A)} \phi(A)$. Our goal is to find some $\hat{A}$ such that $|\phi(\hat{A}) - \phi(A^*)| \leq \varepsilon$. We develop algorithms for this problem for range spaces defined by balls, halfspaces, and axis-aligned rectangles; it has applications in many areas including discrepancy evaluation and spatial scan statistics.

3.22 A Treehouse with Custom Windows: Minimum Distortion Embeddings into Bounded Treewidth Graphs

Benjamin Raichel (University of Texas – Dallas, US)

Joint work of Amir Nayyeri, Benjamin Raichel
Main reference A. Nayyeri, B. Raichel, "A treehouse with custom windows: Minimum distortion embeddings into bounded treewidth graphs", in Proc. of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2017), pp. 724–736, ACM, 2017.
URL https://doi.org/10.1137/1.9781611974782.46

We describe a $(1+\varepsilon)$-approximation algorithm for finding the minimum distortion embedding of an n-point metric space X into the shortest path metric space of a weighted graph G with m vertices. The running time of our algorithm is

$$m^{O(1)} \cdot n^{O(\lambda)} \cdot (\delta_{opt}\Delta)^{\lambda\cdot(1/\varepsilon)^{\lambda+2}\cdot\lambda\cdot(O(\delta_{opt}))^{2\lambda}}$$

parametrized by the values of the minimum distortion, δ_{opt}, the spread, Δ, of the points of X, the treewidth, λ, of G, and the doubling dimension, λ, of G.

In particular, our result implies a PTAS provided an X with polynomial spread, and the doubling dimension of G, the treewidth of G, and δ_{opt}, are all constant. For example, if X has a polynomial spread and δ_{opt} is a constant, we obtain PTAS's for embedding X into the following spaces: the line, a cycle, a tree of bounded doubling dimension, and a k-outer planar graph of bounded doubling dimension (for a constant k).

3.23 Spatial data processing with SAP HANA

Alejandro Salinger (SAP SE – Walldorf, DE)

URL http://help.sap.com/hana/SAP_HANA_Spatial_Reference_en.pdf

The rapid development of geospatial technologies in recent years has led to an increase in the types and amount of enterprise data containing geospatial information. The ability to store and process spatial data can add a valuable dimension to business analytics.

In this talk we give an overview of the spatial data processing capabilities of SAP HANA, SAP's in-memory data platform. We describe the spatial data types and operations supported and give examples of SQL queries using spatial data. These include using simple predicates such as determining if a point is within a given geometry as well as more complex queries that use clustering methods.

We then give some insights about the optimizations used to reduce computational time in window queries, including the use of parallelism, space-filling curves, dictionaries, and indices.

Finally, we show a live demo based on an application for proactive pipeline maintenance used by a major natural gas provider. The application allows visualizing pipelines in a map and finding those that are within a given distance to buildings. We briefly describe the underlying algorithms and data structures used to compute spatial joins that can process millions of geometries in a just a few seconds.

3.24 A Framework for Algorithm Stability

Kevin Verbeek (TU Eindhoven, NL)

We say that an algorithm is stable if small changes in the input result in small changes in the output. Algorithm stability plays an important role when analyzing and visualizing time-varying data. However, so far, there are only few theoretical results on the stability of algorithms, possibly due to a lack of theoretical analysis tools. In this talk we present a framework for analyzing the stability of algorithms. We focus in particular on the tradeoff between the stability of an algorithm and the quality of the solution it computes. Our framework allows for three types of stability analysis with increasing degrees of complexity: event stability, topological stability, and Lipschitz stability. We demonstrate the use of our stability framework by applying it to kinetic Euclidean minimum spanning trees.

3.25 Towards Spectral Sparsification of Simplicial Complexes based on Generalized Effective Resistance

Bei Wang (University of Utah – Salt Lake City, US)

Joint work of Braxton Osting, Sourabh Palande, Bei Wang

As a generalization of the use of graphs to describe pairwise interactions, simplicial complexes can be used to model higher-order interactions between three or more objects in complex systems. There has been a recent surge in activity for the development of data analysis methods applicable to simplicial complexes, including techniques based on computational topology, higher-order random processes, generalized Cheeger inequalities, isoperimetric inequalities, and spectral methods. In particular, spectral learning methods (e.g. label propagation and clustering) that directly operate on simplicial complexes represent a new direction emerging from the confluence of computational topology and machine learning. Similar to the challenges faced by massive graphs, computational methods that operate on simplicial complexes are severely limited by computational costs associated with massive datasets.

To apply spectral methods in learning to massive datasets modeled as simplicial complexes, we work towards the sparsification of simplicial complexes based on preserving the spectrum of the associated Laplacian operators. We show that the theory of Spielman and Srivastava for the sparsification of graphs extends to the generality of simplicial complexes via the up Laplacian. In particular, we introduce a generalized effective resistance for simplexes; provide an algorithm for sparsifying simplicial complexes at a fixed dimension; and gives a specific version of the generalized Cheeger inequalities for weighted simplicial complexes under the sparsified setting. In addition, we demonstrate via experiments the preservation of up Laplacian during sparsification, as well as the utility of sparsification with respect to spectral clustering.

4 Open problems

4.1 Nash equilibria of spanner games

Mohammad Ali Abam (Sharif University of Technology – Tehran, IR)

Starting from a t-sink spanner, that is a directed graph, where from every vertex there is a path to a sink where the length of the path is at most a factor t more than the cost of a direct edge. The vertices (players) take turns and in each turn a player can replace its outgoing edge by a different edge, but it needs to maintain the t-sink spanner property of the graph. Each player pays for its outgoing edge. Does this game have a nash-equilibrium?

A similar problem can be posed for a general t-spanner, where there should be spanning paths between any pair of vertices. In this case during each turn, a vertex can change its set of outgoing edges under the restriction that the graph must remain a spanner. The question is then if this is a process that terminates and if so, with how many edges?

4.2 Negative cycles on surface embedded graphs

Jeff Erickson (University of Illinois – Urbana-Champaign, US)

Given a directed graph on a (possibly non-compatible) surface, where edges have weights that may be negative. The problem is to determine if there is a negative contractible cycle, that is, a closed walk which may reuse edges and vertices. What is the complexity of this decision problem? Is this problem in $\mathbb{P}$ or even in $\mathbb{NP}$ (the cycle might have exponential length), is it $\mathbb{NP}$-hard?

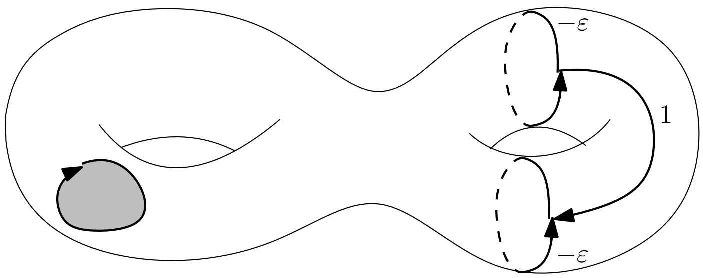

4.3 Maximum-Area Triangle in a Convex Polygon

Maarten Löffler (Utrecht University, NL)

Main reference D. P. Dobkin, L. Snyder, "On a general method for maximizing and minimizing among certain geometric problems", in 20th Annual Symposium on Foundations of Computer Science (FOCS 1979), pp. 9–17, IEEE, 1979.
URL https://doi.org/10.1109/SFCS.1979.28

Given a convex polygon P, find the largest-area inscribed triangle. We revisit the linear-time algorithm proposed by Dobkin and Snyder and ask if there is a counter-example to their claimed proof.

Update: There has been progress on this problem since the problem was posed: Vahideh Keikha, Maarten Löffler, Jérôme Urhausen, Ivor van der Hoog: "Maximum-Area Triangle in a Convex Polygon, Revisited." CoRR abs/1705.11035, 2017.

4.4 Is G area-universal?

Tillmann Miltzow (Free University of Brussels, BE)

Joint work of Linda Kleist, Pawel Rzazewski, Michael Gene Dobbins, Tillmann Miltzow
Main reference L. Kleist, P. Rzazewski, T. Miltzow, "Is Area Universality ∀∃ℝ-complete?", in Proc. of the 33rd European Workshop on Computational Geometry (EuroCG 2017), Malmö University, 2017.
URL http://csconferences.mah.se/eurocg2017/proceedings.pdf

During the Dagstuhl Seminar, I asked whether the graph G (drawn on the blackboard) is area universal. See the EuroCG abstract for details (page 181).

4.5 Finding small polygons with a fixed number of triangulations; simple polygons with a unique triangulation

Joseph S. B. Mitchell (Stony Brook University, US)

Given an integer k, let $N(k)$ denote the minimum number n, so that there is a polygon of n vertices that has exactly k triangulations. This number always exists for any k. The first problem is to find $N(k)$.

A second problem is to detect if a given polygon P has a unique triangulation. This can be done in $O(n)$ time from a starting triangulation. Using Chazelle's linear triangulation algorithm this yields linear time in total; however, since no implementation for linear-time triangulation exists, an interesting question is if there is a different, simple linear-time algorithm.

4.6 Minimum k dimensional cut of an n dimensional polytope

Aleksandar Nikolov (University of Toronto, CA)

Main reference A. Pajor, N. Tomczak-Jaegermann, "Volume ratio and other s-numbers of operators related to local properties of Banach spaces", J. Funct. Anal., 87(2): 273–293, Elsevier, 1989.
URL https://doi.org/10.1016/0022-1236(89)90011-6

Given a convex polytope P in $\mathbb{R}^n$ that is symmetric around the origin (i.e. $P = -P$), specified in H-representation as $P = \{x : b \le Ax \le b\}$ for a matrix A and a vector b with non-negative entries. Note that P can have exponentially many vertices in the size of its representation. More generally, we are given a convex body K symmetric around the origin, specified via a membership or a separation oracle. Can we efficiently find a k-dimensional subspace through the origin that minimizes its cut-volume with P? That is, can we find or approximate $\min\{\mathrm{vol}(W \Cap P)^{1/k} : W$ is a subspace of dimension $k\}$?

The problem is easy for centrally symmetric ellipsoids E, specified by $E = \{x : x^\intercal Bx \le 1\}$, for a positive semidefinite matrix B. Then the minimizing subspace is spanned by eigenvectors associated with the largest k eigenvalues of M. Using this fact, and Milman's M-ellipsoid theorem, we can get a factor $C^{n/k}$-approximation for an arbitrary symmetric convex body K, specified by a membership oracle, and an absolute constant C, independent of n, k, and K. However, this is unsatisfying when k is sublinear in n. On the other hand, for a polytope P in H-representation and $k = 1$, the problem just reduces to computing the closest (in Euclidean distance) facet to the origin. Intermediate values of k seem challenging.

This problem is related to computing volume ratio numbers, which can be used to compute estimates on the Gaussian mean width of a convex body (see Pajor and Tomczak-Jaegermann, JFA 1989).

4.7 VC-dimension of inflated polynomials

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

Main reference J. M. Phillips, Y. Zheng, "Subsampling in Smoothed Range Spaces", in Proc. of the 26th International Conference on Algorithmic Learning Theory (ALT 2015), LNCS, Vol. 9355, pp. 224–238, Springer, 2015.
URL https://doi.org/10.1007/978-3-319-24486-0_15

Is there a bound on the VC-dimension of the range space defined by shapes formed by the Minkowski sum of a ball and a polynomial curve? In more detail, let $\mathcal{S}_p$ be any family of subsets of $\mathbb{R}^d$ which are defined by a shape whose boundary is a polynomially curve of degree p. Simple examples are when $H \in \mathcal{S}_1$ is a halfspace, or when $D \in \mathcal{S}_2$ is a disk; but the problem seems more challenging when we only restrict the degree of the polynomial. Then define the set of subsets $\mathcal{M}_p = \{S \oplus B \mid S \in \mathcal{S}_p\}$ where B is any ball and $\oplus$ is the Minkowski sum. What is the VC-dimension of $(\mathbb{R}^d, \mathcal{M}_p)$?

4.8 Stability of Topological Signatures of Graphs

Bei Wang (University of Utah – Salt Lake City, US)

This open problem arises from a joint work with Mustafa Hajij, Carlos Scheidegger and Paul Rosen [4], where we are interested in quantifying and visualizing structural changes of a time-varying graph using persistent homology. However for simplification purposes here, we describe the open problem by focusing on a single snapshot G of a time-varying graph.

Let $G = (V, E, w)$ be a weighted, undirected graph. To study the topological signatures of G, we first embed G into a metric space (M, d), using the shortest-path distance or other distance metrics based on the graph Laplacian [3] (such as commute-time distance, discrete biharmonic distance [5] and diffusion distance [2]). The topological signatures of such a (metric-space-embedded) graph is then extracted, by computing persistent homology of its corresponding Vietoris-Rips filtration; the topological features of the Vietoris-Rips filtration are captured by its persistence diagrams PD. Such persistence diagrams encoding the topological signatures of G can be compared and structural changes among time-varying snapshots may be detected and visualized.

The question is: *given the above analysis pipeline, can we quantify the stability of topological signatures of graphs*? In other words, if we add random perturbation to a graph, how stable (and therefore potentially trustworthy) are its topological signatures?

The stability of topological signatures of a given graph G can be influenced by at least three factors: the perturbation model, the stability of its metric structure embedding (M, d), and the stability of its corresponding persistence diagram PD. In terms of random perturbation models, we can consider deletion-only and insertion-only perturbation, as well as rewiring (while maintaining degree or joint degree distributions). We would like to know, under a particular perturbation model, and for a certain metric space embedding, whether we can obtain stable topological signatures. As topological methods can be coordinate and deformation invariant, under certain restrictive settings, can we hope for stable topological signatures even if the metric structures are not stable?

The metric structure behind perturbed graphs in the setting of shortest path metric and Erdös-Rényi type perturbation has been studied recently [6]. Given a "noisy" observation G of a true graph G^*, Parthasarathy et al. [6] propose a de-noising procedure to recover (approximately) the "true" shortest path metric of G^* from G. Since the shortest path metric is known to be sensitive to random perturbations, such a de-noising procedure can be used as a way to recover the stable metric structure of a graph. On the other hand, persistent homology stability result [1] implies that if M and N are two finite metric space, then the distance between persistence diagrams constructed by Vietoris-Rips filtration is bounded by twice of their Gromov-Hausdorff distance [7]. Therefore, the shortest-path de-noising procedure of [6] can be combined with persistent homology stability result (e.g. [1, 7]) to potentially address an instance of this open problem.

References

1 Frédéric Chazal, David Cohen-Steiner, Leonidas J Guibas, Facundo Méemoli, and Steve Y. Oudot. Gromov-hausdorff stable signatures for shapes using persistence. Computer Graphics Forum, 28(5):1393–1403, 2009.

2 Ronald R Coifman, Stephane Lafon, Ann B Lee, Mauro Maggioni, Boaz Nadler, Frederick Warner, and Steven W Zucker. Geometric diffusions as a tool for harmonic analysis and

structure definition of data: Diffusion maps. Proceedings National Academy of Sciences of the United States of America, 102(21):7426–7431, 2005.

3 Dragoš M Cvetković, Michael Doob, and Horst Sachs. Spectra of graphs: theory and application, volume 87. Academic Press, 1980.

4 Mustafa Hajij, Bei Wang, Carlos Scheidegger, and Paul Rosen. Persistent homology guided exploration of time-varying graphs. arXiv:1707.06683, 2017.

5 Yaron Lipman, Raif M Rustamov, and Thomas A Funkhouser. Biharmonic distance. ACM Transactions on Graphics, 29(3):27, 2010.

6 Srinivasan Parthasarathy, David Sivakoff, Minghao Tian, and Yusu Wang. A quest to unravel the metric structure behind perturbed networks. Proceedings 33rd International Symposium on Computational Geometry, 2017.

7 Katharine Turner. Generalizations of the rips filtration for quasi-metric spaces with persistent homology stability results. arXiv:1608.00365, 2016.

Participants

- Mohammad Ali Abam
Sharif University of Technology – Tehran, IR
- Tetsuo Asano
JAIST – Ishikawa, JP
- Sang Won Bae
Kyonggi University, KR
- Bahareh Banyassady
FU Berlin, DE
- Ulrich Bauer
TU München, DE
- Karl Bringmann
MPI für Informatik – Saarbrücken, DE
- Otfried Cheong
KAIST – Daejeon, KR
- Man-Kwun Chiu
National Institute of Informatics – Tokyo, JP
- Jinhee Chun
Tohoku University – Sendai, JP
- Kenneth L. Clarkson
IBM Almaden Center – San Jose, US
- Michael Gene Dobbins
Binghamton University, US
- Anne Driemel
TU Eindhoven, NL
- Jeff Erickson
University of Illinois – Urbana-Champaign, US
- Esther Ezra
Georgia Institute of Technology – Atlanta, US
- Sándor Fekete
TU Braunschweig, DE
- Kyle Jordan Fox
Duke University – Durham, US
- Jie Gao
Stony Brook University, US
- Joachim Gudmundsson
The University of Sydney, AU
- Anna Gundert
Universität Köln, DE
- Dan Halperin
Tel Aviv University, IL
- David G. Kirkpatrick
University of British Columbia – Vancouver, CA
- Michal Kleinbort
Tel Aviv University, IL
- Christian Knauer
Universität Bayreuth, DE
- Matias Korman
Tohoku University – Sendai, JP
- Amer Krivosija
TU Dortmund, DE
- Maarten Löffler
Utrecht University, NL
- Tillmann Miltzow
Free University of Brussels, BE
- Joseph S. B. Mitchell
Stony Brook University, US
- Wolfgang Mulzer
FU Berlin, DE
- Aleksandar Nikolov
University of Toronto, CA
- Zuzana Patáková
IST Austria – Klosterneuburg, AT
- Jeff M. Phillips
University of Utah – Salt Lake City, US
- Benjamin Raichel
University of Texas – Dallas, US
- Abhishek Rathod
TU München, DE
- Saurabh Ray
New York University – Abu Dhabi, AE
- Marcel J. M. Roeloffzen
National Institute of Informatics – Tokyo, JP
- Michael Sagraloff
MPI für Informatik – Saarbrücken, DE
- Alejandro Salinger
SAP SE – Walldorf, DE
- Lena Schlipf
FernUniversität in Hagen, DE
- Raimund Seidel
Universität des Saarlandes, DE
- Fabian Stehn
Universität Bayreuth, DE
- Monique Teillaud
INRIA Nancy – Grand Est, FR
- André van Renssen
National Institute of Informatics – Tokyo, JP
- Kevin Verbeek
TU Eindhoven, NL
- Bei Wang
University of Utah – Salt Lake City, US
- David P. Woodruff
IBM Almaden Center – San Jose, US