Report from Dagstuhl Seminar 21181

Computational Geometry

Edited by

Siu-Wing Cheng[1], Anne Driemel[2], and Jeff M. Phillips[3]

1 **HKUST – Kowloon, HK,** `scheng@cse.ust.hk`
2 **Universität Bonn, DE,** `driemel@cs.uni-bonn.de`
3 **University of Utah – Salt Lake City, US,** `jeffp@cs.utah.edu`

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 21181 "Computational Geometry". The seminar was held from May 2 to May 7, 2021. Because of COVID, the seminar was held online in a virtual manner, and 36 participants from various countries attended it. New advances and directions in computational geometry were presented and discussed. The report collects the abstracts of talks and open problems presented in the seminar.

Seminar May 2–7, 2021 – http://www.dagstuhl.de/21181
2012 ACM Subject Classification Theory of computation → Computational geometry
Keywords and phrases algorithms, computational geometry, Computational topology, data structures, Discrete geometry
Digital Object Identifier 10.4230/DagRep.11.4.1
Edited in cooperation with Conradi, Jacobus

1 Executive Summary

Siu-Wing Cheng
Anne Driemel
Jeff M. Phillips

Computational Geometry

The field of 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.

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

Seminar Topics

The emphasis of this 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 Computational Topology on Surfaces and Graphs as well as Combinatorial Complexity of Geometric Structures.

Computational Topology on Surfaces and Graphs

Computational topology has seen exciting advances in a number of topics. Indeed, best paper awards in several recent SoCGs went to papers on these topics. In 2019, Cohen-Addad et al. give a lower bound to a cutting problem in embedded graphs, essentially matching the running time of the fastest algorithm known and settling a 17-year old question. In 2018, Goaoc et al. proved that it is NP-complete to decide if a d-dimensional simplicial complex is shellable for $d \geq 2$, resolving a question of Danaraj and Klee in 1978. In 2017, Despré and Lazarus presented simple quasi-linear algorithms for questions regarding geometric intersection number of a curve on a surface. Progress in these and related topics have had influences in problems on graphs embedded on surfaces, maximum flows and multiple-source shortest paths in planar graphs, collapsibility of simplicial complexes, metric learning, etc. The seminar highlighted these topics with two overview talks. The first by Hsien-Chih Chang was on Tightening Curves on Surfaces, and provided a overview of recent advancements in this area, and exciting directions for future work on flipping triangulations and morphing planar multicurves using electrical moves. The second by Uli Wagner discussed Embeddability of Simplicial Compexes, and also the flurry of recent research in this area, and pinpointed the several remaining questions and where the community has not yet been able to resolve the embeddability and why the challenges remain. These talks, and other on recent advances, helped summarize the state of this area, and generate new avenues towards moving the field further forward.

Combinatorial Complexity of Geometric Structures

The understanding of the combinatorial properties of geometric structures is at the core of computational geometry. A lot of these structures such as union of shapes, cuttings, arrangements, Delaunay triangulation, Voronoi diagram have found numerous applications in algorithm design. For example, the analysis of the complexity of the union of translates of a convex body allows us to understand the complexity of the free space in planning the motion of that convex body under translation. Their studies have also triggered the development of new theoretical tools such as the polynomial method that has been gaining a lot of attention lately. There are also new applications that require the modeling of uncertain data and hence call for a study of many geometric structures under a stochastic setting. The seminar promoted these topics via two overview talks. The first overview talk was by Mikkel Abrahamsen on Minimum Fence Enclosure and Separation Problems; this line of work generalizes the notion of convex hull by identifies other minimally enclosing structures called fences, and the interesting combinatorial properties that arise. The second overview talk by Evanthia Papadopoulou was on Problems in Voronoi and Voronoi-like diagrams. This talk discussed the advancement in generalizations of the classic geometric object of Voronoi diagrams to be defined among geometric objects beyond points, and to higher-order complexes. In addition to providing snapshots of these exciting subareas, they provided future directions for research within these topics and in how they can interact across the broader computational geometry landscape.

Participants and Participation

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 and vice versa. They have arguably been the most influential meetings in the field of computational geometry. The organizers held a lottery for the fifth 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, 36 researchers from various countries and continents attended the seminar, despite the virtual nature due to COVID-19, showing the strong interest of the community for this event.

Due to the COVID-19 pandemic, the seminar was held entirely virtually. Talks were held over four days. Each day had 2 two-hour blocks of talks, separated by a 2-hour meal break. They were held in the late-afternoon and evening in Europe, which allowed for participants from North America to join in during their morning hours. Unfortunately, this timing was late for those in Asia. The talks were held on Zoom, a Slack server was set up for a more persistent text-based discussion, and a Wonder.me instance was arranged for dynamic forming of group discussions before and after each session. All of these settings were used to communicate research, form collaborations, and attack open problems. Although not as wonderful as actually being at Schloss Dagstuhl, these online mechanisms provided for a workable replacement for what a normal Dagstuhl seminar provides in this abnormal time.

The feedback from participants was very positive. The participants viewed the composition of the group positively, remarking how it was well-balanced in terms of seniority and gender. They also praised the quality of the talks as of very high quality – making the virtual-only participation worthwhile.

We warmly thank the scientific, administrative and technical staff at Schloss Dagstuhl! Dagstuhl made virtual hosting possible and easy in a time filled with complications. Despite not providing a physical space to meet, socialize, and collaborate, their help in organizing the event made it a success despite the less than ideal circumstances.

2 Table of Contents

3 Overview of Talks

3.1 Minimum Fence Enclosure and Separation Problems

Mikkel Abrahamsen (University of Copenhagen, DK)

The classical problem of computing the convex hull of a given set of points in the plane can be formulated in a natural way as a fence enclosure problem: Find the shortest fence that encloses the points. In this talk, we survey our recent work on related problems that appear when the formulation is changed slightly. We will touch upon the following problems:

1. Given a set of points in the plane, find the two fences of minimum total length that together enclose all the points. We will outline an algorithm with $O(n \log^2 n)$ running time.
2. Given a set of points and a number k, find a system of at most k fences of minimum total length that enclose the points. We will explain how this problem can be attacked via dynamic programming, which leads to a polynomial-time, although very slow, algorithm.
3. Given a set of unit disks, find a system of fences of minimum total length that enclose all the disks (with no restriction on the number of fences). We report on a near-linear time algorithm for this and related problems.
4. Finally, we consider the problem where the input consists of pairwise interior-disjoint polygons in the plane, and each polygon has a color. We want to compute the fence of minimum total length that separates all pairs of polygons of different colors. This problem can be solved in polynomial time when there are just two colors, but it becomes NP-hard already for three colors. We report on an approximation algorithm. During the talk, we will suggest directions for future research.

3.2 On the Union of Cubes in 3D

Pankaj Kumar Agarwal (Duke University – Durham, US)

Joint work of Pankaj K. Agarwal, Micha Sharir, Alex Steiger
Main reference Pankaj K. Agarwal, Micha Sharir, Alex Steiger: "Decomposing the Complement of the Union of Cubes in Three Dimensions", in Proc. of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021, pp. 1425–1444, SIAM, 2021.
URL http://dx.doi.org/10.1137/1.9781611976465.86

Let C be a set of n axis-aligned cubes of arbitrary sizes in 3D, let $K = \mathbb{R}^3 \setminus U(C)$ be the complement of their union (i.e. free space). The complexity of K, denoted by k, can vary between $O(1)$ and $O(n^2)$. This talk presents two main results: (i) An output-sensitive algorithm to compute K in time $O(n\text{polylog}(n) + k)$ time; and (ii) an output-sensitive algorithm to partition K into $O((n+k)\text{polylog}(n))$ boxes in the same time bound. These results can be slightly improved if the cubes in C have roughly the same size or if they have bounded depth (i.e. any point in $\mathbb{R}^3$ lies in $O(1)$ cubes).

3.3 Fine-grained Complexity of Nearest Neighbors for Fréchet Distance

Karl Bringmann (Universität des Saarlandes – Saarbrücken, DE)

Joint work of Karl Bringmann, Anne Driemel, André Nusser, Ioannis Psarros

Fine-grained complexity theory is the area of theoretical computer science that proves conditional lower bounds based on the 3-SUM Hypothesis, the Strong Exponential Time Hypothesis, and similar conjectures. This talk is an introduction to recent fine-grained lower bounds in computational geometry, with a focus on lower bounds for polynomial-time problems based on the Orthogonal Vectors Hypothesis. Specifically, we discuss conditional lower bounds for nearest neighbor search under the Euclidean distance and Fréchet distance. We see that lower bounds for the Bichromatic Closest Pair problem follow from the Orthogonal Vectors Hypothesis by simple embeddings. This implies a near-linear lower bound for the query time of nearest neighbor data structures. Then we see Unbalanced Orthogonal Vectors, a simple trick to even rule out any polynomial preprocessing time and near-linear query time for nearest neighbors. Finally, we discuss recent, unpublished work on approximate nearest neighbor data structures for the Fréchet distance.

3.4 Around k-fold filtrations

Mickaël Buchet (TU Graz, AT)

Joint work of Mickaël Buchet, Bianca B. Dornelas, Michael Kerber

Given a point set P, a number k and a radius r, the k-fold cover is defined as the union of all intersections of k balls of radius r around points of P. This cover has two parameters (k and r) and defines a very natural bi-filtration of the space. This bi-filtration represents one natural occurrence for multi-parameter persistent homology. Unfortunately, the usual combinatorial objects used to represent the bi-filtration are not bi-filtration themselves and are of larger size. I will talk about several ways to tackle this issue through various approximations, constructions and sparsifications, mostly adapted from techniques used in the one-parameter case but where the nature of the k-fold cover raises new challenges and interesting open questions.

3.5 Computing the inverse geodesic length in graphs of bounded treewidth

Sergio Cabello (University of Ljubljana, SI)

Main reference Sergio Cabello: "Computing the inverse geodesic length in planar graphs and graphs of bounded treewidth", CoRR, Vol. abs/1908.01317, 2019.
URL https://arxiv.org/abs/1908.01317

The inverse geodesic length of a graph G is the sum of the inverse of the distances between all pairs of distinct vertices of G. In some domains it is known as the Harary index or the global efficiency of the graph. We show that, if G has n vertices and constant treewidth, then the inverse geodesic length of G can be computed in near-linear time. To achieve this we use techniques developed for computing the sum of the distances, which does not have "inverse" component, together with batched evaluations of rational functions.

3.6 Tightening Curves on Surfaces

Hsien-Chih Chang (Dartmouth College – Hanover, US)

In this talk we survey the recent advancements in tightening curves on surfaces under different categories of curves and deformations in the past few decades. We then present two possible directions for future research, one on computing geodesics by flipping triangulations and the other on the complexity of morphing planar multicurves using electrical moves. Open questions are provided for the curious to ponder.

3.7 Multicuts in planar and surface-embedded graphs

Éric Colin de Verdière (CNRS, LIGM, Marne-la-Vallée, FR)

Joint work of Vincent Cohen-Addad, Vincent, Éric Colin de Verdière, Dániel Marx, Arnaud de Mesmay

The Multicut problem is defined as follows. Given an edge-weighted graph G and pairs of vertices (s_1, t_1), ..., (s_k, t_k), compute a minimum-weight subset of edges whose removal disconnects each pair (s_i, t_i).

This problem is NP-hard, APX-hard, and W[1]-hard in the number of pairs of terminals, even in very simple cases, such as planar graphs.

We will survey some recent results on this problem, on planar graphs and more generally on graphs embedded on a fixed surface: An exact algorithm, whose running time is a polynomial in the genus and the number of terminals [1]; a matching lower bound assuming ETH [2]; and an approximation scheme with running time $O(n \log n)$ if the approximation factor, the genus, and the number of terminals are fixed [3].

All these results rely on topological methods: The subgraph of the dual of G, made of the edges dual to a multicut, has nice properties, which can be exploited using classical tools from algebraic topology such as homotopy, homology, and universal covers of surfaces.

References

1 Éric Colin de Verdière. Multicuts in planar and bounded-genus graphs with bounded number of terminals. *Algorithmica*, 78:1206–1224, 2017.
2 Vincent Cohen-Addad, Éric Colin de Verdière, Dániel Marx, and Arnaud de Mesmay. Almost tight lower bounds for hard cutting problems in embedded graphs. In *Proc. Int. Symp. on Computational Geometry*, pages 27:1–27:16, 2019. Full version to appear in *J. ACM.*
3 Vincent Cohen-Addad, Éric Colin de Verdière, and Arnaud de Mesmay. A near-linear approximation scheme for multicuts of embedded graphs with a fixed number of terminals. *SIAM J. Comput.*, 50(1):1–33, 2021.

3.8 Fine-grained Complexity of the k-Shortcut Fréchet distance

Jacobus Conradi (Universität Bonn, DE)

Joint work of Jacobus Conradi, Anne Driemel

The Fréchet distance is a popular measure of dissimilarity for polygonal curves. It is defined as a min-max formulation that considers all direction-preserving continuous bijections of the two curves. Because of its susceptibility to noise, Driemel and Har-Peled introduced the shortcut Fréchet distance in 2012, where one is allowed to take shortcuts along one of the curves, similar to the edit distance for sequences. We analyse the parametrized version of this problem, where the number of shortcuts is bounded by a parameter k. The corresponding decision problem can be stated as follows: Given two polygonal curves T and B of at most n vertices, a parameter k and a distance threshold δ, is it possible to introduce k shortcuts along B such that the Fréchet distance of the resulting curve and the curve T is at most δ? We study this problem for polygonal curves in the plane. We provide a complexity analysis for this problem with the following results: (i) assuming the exponential-time-hypothesis (ETH), there exists no algorithm with running time bounded by $n^{o(k)}$; (ii) there exists a decision algorithm with running time in $O(kn^{2k+2} \log n)$. In contrast, we also show that efficient approximate decider algorithms are possible, even when k is large. We present a $(3 + \varepsilon)$-approximate decider algorithm with running time in $O(kn^2 \log^2 n)$ for fixed ε. In addition, we can show that, if k is a constant and the two curves are c-packed for some constant c, then the approximate decider algorithm runs in near-linear time.

3.9 Contractibility on 3-manifold boundaries and compressed problems on surfaces

Arnaud de Mesmay (University Paris-Est – Marne-la-Vallée, FR)

Joint work of Erin W. Chambers, Arnaud de Mesmay, Francis Lazarus, Salman Parsa
Main reference Erin Wolf Chambers, Francis Lazarus, Arnaud de Mesmay, Salman Parsa: "Algorithms for Contractibility of Compressed Curves on 3-Manifold Boundaries", CoRR, Vol. abs/2012.02352, 2020.
URL https://arxiv.org/abs/2012.02352

We show that the problem of deciding whether a closed curve on the boundary of a 3-manifold is contractible is in NP, and furthermore we provide an algorithm that is FPT in the complexity of the manifold. This relies on techniques to solve various topological problems for curves on surfaces with compressed inputs. The talk assumes no topological background and focuses on explaining why issues with compression appear naturally in this line of work.

3.10 Practical volume approximation of H, V, and Z-polytopes

Ioannis Emiris (University of Athens & Athena Research Center, GR)

Joint work of Apostolos Chalkis, Ioannis Z. Emiris, Vissarion Fisikopoulos
Main reference Apostolos Chalkis, Ioannis Z. Emiris, Vissarion Fisikopoulos: "Practical Volume Estimation by a New Annealing Schedule for Cooling Convex Bodies", CoRR, Vol. abs/1905.05494, 2019.
URL http://arxiv.org/abs/1905.05494

We tackle the problem of efficiently approximating the volume of convex polytopes, when these are given in 3 different representations: H-polytopes, which have been studied extensively, V-polytopes, and zonotopes (Z-polytopes). We design a novel practical Multiphase Monte Carlo (MMC) algorithm that leverages geometric random walks. Our algorithmic contributions include: (i) a uniform sampler employing billiard walk for the first time in volume computation, showing it mixes much faster than Hit-and-Run variants, (ii) a new simulated annealing schedule, generalizing existing MMC, by introducing adaptive convex bodies which, moreover, (iii) probabilistically restricts volume ratios to a target interval, thus drastically reducing the number of bodies in MMC. Extensive experiments indicate that our method requires about $O(d^2)$ oracle calls compared to the best theoretical bound of $O^*(d^3)$, where d is the dimension. For zonotopes, we appropriately use centrally symmetric polytopes which yield an MMC with constant number of phases, when the ratio of generators over dimension is small. We present a detailed experimental evaluation of our algorithm using Birkhoff polytopes and polytopes of all 3 classes. Our open-source C++ software offers the first method that scales up to thousands of dimensions for H-polytopes and in the hundreds for V- and Z-polytopes on moderate hardware. We illustrate it on SDP optimization, by means of sampling spectrahedra, and on sampling structured polytopes obtained from modeling financial portfolios.

3.11 Consistent Digital Line Segments

Matias Korman (Siemens EDA – Wilsonville, US)

Joint work of Man-Kwun Chiu, Matias Korman, Martin Suderland, Takeshi Tokuyama
Main reference Man-Kwun Chiu, Matias Korman, Martin Suderland, Takeshi Tokuyama: "Distance Bounds for High Dimensional Consistent Digital Rays and 2-D Partially-Consistent Digital Rays", in Proc. of the 28th Annual European Symposium on Algorithms, ESA 2020, September 7-9, 2020, Pisa, Italy (Virtual Conference), LIPIcs, Vol. 173, pp. 34:1–34:22, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
URL http://dx.doi.org/10.4230/LIPIcs.ESA.2020.34

In this talk I will introduce the concept of "consistent digital segments": In short, we look for an axiomatic construction of segments in discrete spaces, akin to the construction that we have in Euclidean segments. After discussing movivation, we will focus on known results in two and higher dimensions. Each result will be followed with discussion on what are the big open problems that remain and what are possible lines of research that could be followed.

3.12 Approximating Maximum Independent Set in the Plane

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

Main reference Joseph S. B. Mitchell: "Approximating Maximum Independent Set for Rectangles in the Plane", CoRR, Vol. abs/2101.00326, 2021.
URL https://arxiv.org/abs/2101.00326

We give a polynomial-time constant-factor approximation algorithm for maximum independent set for (axis-aligned) rectangles in the plane. Using a polynomial-time algorithm, the best approximation factor previously known is $O(\log \log n)$. The results are based on a new form of recursive partitioning in the plane, in which faces that are constant-complexity and orthogonally convex are recursively partitioned in a constant number of such faces.

3.13 Efficient Near-Neighbor Search via Average Distortion Embeddings

Aleksandar Nikolov (University of Toronto, CA)

Joint work of Deepanshu Kush, Aleksandar Nikolov, Haohua Teng
Main reference Deepanshu Kush, Aleksandar Nikolov, Haohua Tang: "Near Neighbor Search via Efficient Average Distortion Embeddings", CoRR, Vol. abs/2105.04712, 2021.
URL https://arxiv.org/abs/2105.04712

A recent series of papers by Andoni, Naor, Nikolov, Razenshteyn, and Waingarten (STOC 2018, FOCS 2018) has given approximate near neighbour search (ANN) data structures for a wide class of distance metrics, including all norms. In particular, these data structures achieve approximation on the order of p for ℓ_p norms with space complexity nearly linear in the dataset size n and polynomial in the dimension d, and query time sub-linear in n and polynomial in d. The main shortcoming is the exponential in d pre-processing time required for their construction. In this talk, we describe a more direct framework for constructing ANN data structures for general norms. More specifically, we show via an algorithmic reduction that an efficient ANN data structure for a given metric is implied by an efficient average distortion embedding of the metric into the Manhattan norm or into Euclidean space. In particular, the resulting data structures require only polynomial pre-processing time, as long as the embedding can be computed in polynomial time. As a concrete instantiation of this framework, we give an ANN data structure for ℓ_p with efficient pre-processing that matches the approximation factor, space and query complexity of the aforementioned data structure of Andoni et al.

3.14 Problems in Voronoi and Voronoi-like diagrams

Evanthia Papadopoulou (University of Lugano, CH)

Joint work of Kolja Junginger, Evanthia Papadopoulou
Main reference Kolja Junginger, Evanthia Papadopoulou: "Deletion in Abstract Voronoi Diagrams in Expected Linear Time", in Proc. of the 34th International Symposium on Computational Geometry, SoCG 2018, June 11-14, 2018, Budapest, Hungary, LIPIcs, Vol. 99, pp. 50:1–50:14, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
URL http://dx.doi.org/10.4230/LIPIcs.SoCG.2018.50
Main reference Kolja Junginger, Evanthia Papadopoulou: "Deletion in abstract Voronoi diagrams in expected linear time", CoRR, Vol. abs/1803.05372, 2018.
URL http://arxiv.org/abs/1803.05372

Differences between classical Voronoi diagrams of points, versus segments, circles, or polygons are often forgotten or underestimated. Abstract Voronoi diagrams (AVDs) offer a unifying framework for many such Voronoi diagrams in the plane; however, diagrams of points are not a representative concrete structure for AVDs. In this talk, I will first survey fundamental differences between higher order Voronoi diagrams of points and their counterparts of segments or AVDs. I will then address the problem of site-deletion in abstract Voronoi diagrams in expected linear time. Although linear-time algorithms for site-deletion in planar point Voronoi diagrams had been well-known to exist since the late 80's, the corresponding problems for non-point Voronoi diagrams remained open, until recently. As a byproduct, I will introduce *abstract Voronoi-like diagrams*, a relaxed Voronoi structure of independent interest, which leads to a very simple randomized incremental technique to perform site-deletion in abstract Voronoi diagrams. The technique extends to computing various tree-like Voronoi diagrams such as constructing the farthest abstract Voronoi diagram, after the order of its regions at infinity is known, constructing the order-$(k+1)$ subdivision within an order-k Voronoi region, and others. The time analysis introduces a simple alternative to backwards analysis applicable to order-dependent structures.

3.15 Stronger Bounds for Weak Epsilon-Nets in Higher Dimensions

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

Main reference Natan Rubin: "Stronger Bounds for Weak Epsilon-Nets in Higher Dimensions", CoRR, Vol. abs/2104.12654, 2021.
URL https://arxiv.org/abs/2104.12654

Given a finite point set P in $\mathbb{R}^d$, and $\epsilon > 0$ we say that $N \subseteq \mathbb{R}^d$ is a weak ϵ-net if it pierces every convex set K with $|K \cap P| \geq \epsilon |P|$.

Let $d \geq 3$. We show that for any finite point set in $\mathbb{R}^d$, and any $\epsilon > 0$, there exist a weak ϵ-net of cardinality $O\left(\frac{1}{\epsilon^{d-1/2+\gamma}}\right)$, where $\gamma > 0$ is an arbitrary small constant.

This is the first improvement of the bound of $O^*\left(\frac{1}{\epsilon^d}\right)$ that was obtained in 1993 by Chazelle, Edelsbrunner, Grigni, Guibas, Sharir, and Welzl for general point sets in dimension $d \geq 3$.[1]

[1] $O^*(\cdot)$-notation hides multiplicative factors that are polylogarithmic in $\log 1/\epsilon$.

3.16 Terrain prickliness: theoretical grounds for low complexity viewsheds

Maria Saumell (The Czech Academy of Sciences – Prague & Czech Technical University in Prague, CZ)

Joint work of Ankush Acharyya, Ramesh K. Jallu, Maarten Löffler, Gert G. T. Meijer, Maria Saumell, Rodrigo I. Silveira, Frank Staals, Hans Raj Tiwary
Main reference Ankush Acharyya, Ramesh K. Jallu, Maarten Löffler, Gert G. T. Meijer, Maria Saumell, Rodrigo I. Silveira, Frank Staals, Hans Raj Tiwary: "Terrain prickliness: theoretical grounds for low complexity viewsheds", CoRR, Vol. abs/2103.06696, 2021.
URL https://arxiv.org/abs/2103.06696

An important task when working with terrain models is computing viewsheds: the parts of the terrain visible from a given viewpoint. When the terrain is modeled as a polyhedral terrain, the viewshed is composed of the union of all the triangle parts that are visible from the viewpoint. The complexity of a viewshed can vary significantly, from constant to quadratic in the number of terrain vertices, depending on the terrain topography and the viewpoint position.

In this work we study a new topographic attribute, the *prickliness*, that measures the number of local maxima in a terrain from all possible perspectives. We show that the prickliness effectively captures the potential of 2.5D terrains to have high complexity viewsheds, and we present near-optimal algorithms to compute the prickliness of 1.5D and 2.5D terrains. We also report on some experiments relating the prickliness of real word 2.5D terrains to the size of the terrains and to their viewshed complexity.

3.17 Guarding Problems

Christiane Schmidt (Linköping University, SE)

Joint work of Sarah Cannon, Ovidiu Daescu, Thomas Fai, Stephan Friedrichs, Justin Iwerks, Undine Leopold, Hemant Malik, Bengt J. Nilsson, Valentin Polishchuk, Christiane Schmidt

The classical Art Gallery Problem (AGP) asks for the minimum number of guards that are necessary to visually cover a polygon P, where visibility between two points is defined as the line segment between these points being fully contained in P. In this talk, we highlight some recent works and open problems for different variants of the AGP.

First, we consider k-transmitters, for which the definition of visibility is altered: two points, p, q, can see each other if the line segment $\overline{pq}$ intersects P's boundary at most k times. We review results on stationary point and edge $k-/2$-transmitters – Art Gallery theorems and complexity results; we present several properties of $k-/2$-transmitters and recent complexity results on 2-transmitter watchman routes. Finally, we highlight an open problem on Art Gallery theorems for 2-transmitters in simple polygon: we have a lower bound of $\lfloor n/5 \rfloor$ guards, but the best known upper bound essentially stems from "normal" guards/0-transmitters: $\lfloor (n-1)/3 \rfloor$.

We then consider guarding problems in special polygon classes (altering the environment to be guarded rather than the capabilities of the guards). We show that we can find an optimal guard set for uni-monotone polygons in linear time and that the size of a minimum cardinality guard set equals the size of a maximum cardinality witness set for this class –

uni-monotone polygons are perfect. We survey for which guarding problems discretizations have been obtained. Finally, we review under which types of visibility definition and for which polygon classes similar results on perfectness have been obtained and give the basic idea of these results. This leads to an open question of AGP in monotone polygons: can we show perfectness for, e.g., staircase or O-visibility. Can we discretize in this special polygon class?

3.18 Sketching Persistence Diagrams

Donald Sheehy (North Carolina State University – Raleigh, US)

Joint work of Donald R. Sheehy, Siddharth Sheth
Main reference Donald R. Sheehy, Siddharth Sheth: "Sketching Persistence Diagrams", in Proc. of the 37th International Symposium on Computational Geometry, SoCG 2021, June 7-11, 2021, Buffalo, NY, USA (Virtual Conference), LIPIcs, Vol. 189, pp. 57:1–57:15, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
URL http://dx.doi.org/10.4230/LIPIcs.SoCG.2021.57

Given a persistence diagram with n points, we give an algorithm that produces a sequence of n persistence diagrams converging in bottleneck distance to the input diagram, the ith of which has i distinct (weighted) points and is a 2-approximation to the closest persistence diagram with that many distinct points. For each approximation, we precompute the optimal matching between the ith and the $(i + 1)$st. Perhaps surprisingly, the entire sequence of diagrams as well as the sequence of matchings can be represented in $O(n)$ space. The main approach is to use a variation of the greedy permutation of the persistence diagram to give good Hausdorff approximations and assign weights to these subsets. We give a new algorithm to efficiently compute this permutation, despite the high implicit dimension of points in a persistence diagram due to the effect of the diagonal. The sketches are also structured to permit fast (linear time) approximations to the Hausdorff distance between diagrams – a lower bound on the bottleneck distance. For approximating the bottleneck distance, sketches can also be used to compute a linear-size neighborhood graph directly, obviating the need for geometric data structures used in state-of-the-art methods for bottleneck computation.

3.19 Optimal bounds for the colorful fractional Helly theorem

Martin Tancer (Charles University – Prague, CZ)

Joint work of Denys Bulavka, Afshin Goodarzi, Martin Tancer
Main reference Denys Bulavka, Afshin Goodarzi, Martin Tancer: "Optimal Bounds for the Colorful Fractional Helly Theorem", in Proc. of the 37th International Symposium on Computational Geometry, SoCG 2021, June 7-11, 2021, Buffalo, NY, USA (Virtual Conference), LIPIcs, Vol. 189, pp. 19:1–19:14, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
URL http://dx.doi.org/10.4230/LIPIcs.SoCG.2021.19

The well known fractional Helly theorem and colorful Helly theorem can be merged into the so called colorful fractional Helly theorem. It states: for every $\alpha \in (0, 1]$ and every non-negative integer d, there is $\beta = \beta(\alpha, d) \in (0, 1]$ with the following property. Let $\mathcal{F}_1, \ldots, \mathcal{F}_{d+1}$ be finite nonempty families of convex sets in $\mathbb{R}^d$ of sizes $n_1, \ldots, n_{d+1}$, respectively. If at least $\alpha n_1 n_2 \cdots n_{d+1}$ of the colorful $(d + 1)$-tuples have a nonempty intersection, then there is

$i \in [d+1]$ such that $\mathcal{F}_i$ contains a subfamily of size at least βn_i with a nonempty intersection. (A colorful $(d+1)$-tuple is a $(d+1)$-tuple $(F_1, \ldots, F_{d+1})$ such that F_i belongs to $\mathcal{F}_i$ for every i.)

The colorful fractional Helly theorem was first stated and proved by Bárány, Fodor, Montejano, Oliveros, and Pór in 2014 with $\beta = \alpha/(d+1)$. In 2017 Kim proved the theorem with better function β, which in particular tends to 1 when α tends to 1. Kim also conjectured what is the optimal bound for $\beta(\alpha, d)$ and provided the upper bound example for the optimal bound. The conjectured bound coincides with the optimal bounds for the (non-colorful) fractional Helly theorem proved independently by Eckhoff and Kalai around 1984.

We verify Kim's conjecture by extending Kalai's approach to the colorful scenario. Moreover, we obtain optimal bounds also in a more general setting when we allow several sets of the same color.

3.20 Light Euclidean Spanners

Csaba Tóth (California State University – Los Angeles, US)

Main reference Sujoy Bhore, Csaba D. Tóth: "On Euclidean Steiner (1+ϵ)-Spanners", in Proc. of the 38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021, March 16-19, 2021, Saarbrücken, Germany (Virtual Conference), LIPIcs, Vol. 187, pp. 13:1–13:16, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
URL http://dx.doi.org/10.4230/LIPIcs.STACS.2021.13

Lightness is a fundamental parameter for Euclidean spanners; it is the ratio of the spanner weight to the weight of the minimum spanning tree of a finite set of points in $\mathbb{R}^d$. In a recent breakthrough, Le and Solomon (2019) established the precise dependencies on $\varepsilon > 0$ and $d \in \mathbb{N}$ of the minimum lightness of a $(1+\varepsilon)$-spanner, and observed that additional Steiner points can substantially improve the lightness. Le and Solomon (2020) constructed Steiner $(1+\varepsilon)$-spanners of lightness $O(\varepsilon^{-1} \log n)$ for n points in the plane. They also constructed spanners of lightness $\tilde{O}(\varepsilon^{-(d+1)/2})$ in dimensions $d \geq 3$.

We established a lower bound of $\Omega(\varepsilon^{-d/2})$ for the lightness of Steiner $(1+\varepsilon)$-spanners in $\mathbb{R}^d$, for all $d \geq 2$. We also prove that this bound is the best possible for $d = 2$, that is, for every finite set of points in the plane and every $\varepsilon > 0$, there exists a Euclidean Steiner $(1+\varepsilon)$-spanner of lightness $O(\varepsilon^{-1})$. We generalize the notion of shallow light trees, which may be of independent interest, and use directional spanners and a modified window partitioning scheme to achieve a tight weight analysis. (Joint work with Sujoy Bhore.)

3.21 Embeddability of Simplicial Complexes

Uli Wagner (IST Austria – Klosterneuburg, AT)

Consider the following decision problem in computational topology, which we refer to as the *embeddability problem*: Given a finite k-dimensional simplicial complex K, does K admit a (piecewise-linear) embedding into $\mathbb{R}^d$? More generally, the *extendability of embeddings problem* asks: Given K, a subcomplex $A \subseteq K$, and an embedding $f\colon A \to \mathbb{R}^d$, can f be extended to an embedding $F\colon K \to \mathbb{R}^d$? (The embeddability problem is the special case $A = \emptyset$.)

We survey what is known about the decidability and computational complexity of these problems in higher dimensions (for fixed positive integers $k \leq d$). Some of the main results and open questions are:

- For $d = 3$, the embeddability problem is known to be algorithmically decidable [7] as well as NP-hard [2]; the exact complexity of the problem (including whether it lies in NP) remains unknown.
- In the so-called *metastable range* $d \geq \frac{3(k+1)}{2}$, both embeddability and extendability of embeddings can be decided in polynomial time (for fixed k and d). Indeed, in this dimension range, by classical work of Haefliger and Weber [4, 8], both problems reduce to questions about the existence of equivariant maps from the *deleted product* $(K \times K) \setminus \{(x, x) \colon x \in K\}$ to S^{d-1}, and the latter can be decided in polynomial time by a series of works on computational homotopy theory culminating in [1].
- Outside the metastable range, it is known that the embeddability problem is NP-hard if $d \geq 4$, and algorithmically undecidable if $k + 1 \geq d \geq 5$ [6]. Moreover, extendability of embeddings is algorithmically undecidable for almost all dimensions outside the metastable range, namely for $8 \leq d < \lfloor \frac{3(k+1)}{2} \rfloor$ [3]. In [3], it is claimed that this also implies undecidability of the embeddability problem in the same range of dimensions, but the proof of this implication contains a gap [5]. Fixing this gap would require constructing suitable so-called *linking gadgets.* E.g., in the special case $k = 5, d = 8$, this would require constructing a 5-dimensional complex L containing copies of S^5 and S^2 as vertex-disjoint subcomplexes, such that in any embedding of G, the images of S^5 and S^2 are linked with linking number ± 1. (Currently, it is only known that there are examples of complexes L that force an odd linking number.)

References

1. M. Čadek, M. Krčál, and L. Vokřínek. Algorithmic solvability of the lifting-extension problem. *Discrete Comput. Geom.*, 57(4):915–965, 2017.
2. A. de Mesmay, Y. Rieck, E. Sedgwick, and M. Tancer. Embeddability in $\mathbb{R}^3$ is NP-hard. In *Proc. 20th Ann. ACM-SIAM Symp. Discr. Algorithms (SODA)*, pp. 1316–1329, 2018.
3. M. Filakovský, S. Zhechev, and U. Wagner. Embeddability of Simplicial Complexes is Undecidable. In *Proc. 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 767–785, 2020.
4. A. Haefliger. Plongements différentiables dans le domaine stable. *Comment. Math. Helv.*, 37:155–176, 1962/1963.
5. R. Karasev and A. Skopenkov. Some 'converses' to intrinsic linking theorems. Preprint, https://arxiv.org/abs/2008.02523.
6. J. Matoušek, M. Tancer, and U. Wagner. Hardness of embedding simplicial complexes in $mathbbR^d$. *J. Eur. Math. Soc.*, 13(2):259–295, 2011.
7. J. Matoušek, E. Sedgwick, M. Tancer, and U. Wagner. Embeddability in the 3-sphere is decidable. *J. ACM*, 65(1):Art. 5, 49, 2018.
8. C. Weber. Plongements de polyedres dans le domaine metastable. *Comment. Math. Helv.*, 42:1–27, 1967.

3.22 Comparing Embedded and Immersed Graphs

Carola Wenk (Tulane University – New Orleans, US)

Data in the form of one-dimensional structures, embedded or immersed in an ambient space, arise in a variety of applications, including GIS analysis, trajectory clustering, protein alignment, plant morphology, commodity networks such as electrical grids, and geographic networks of roads or rivers. Often one is interested in comparing such structures. But since data collection introduces noise and errors, distance measure need to be robust to different issues in the data.In this talk we will focus on graphs and provide a survey of distance measures suitable for comparing embedded or immersed graphs. Oftentimes these graphs are not isomorphic, nor is one interested in true subgraph isomorphism. However, it is desirable to have a mapping of one graph to the other, which should fulfill certain properties such as continuity. And the distances should ideally measure differences in geometry and topology. We will examine several distances from mathematical, algorithmic, and applied point of views, and pose open problems for comparing embedded or immersed graphs.

4 Open problems

4.1 Twin-width of String Graphs

Édouard Bonnet (ENS – Lyon, FR)

A tri-graph is a graph consisting of vertices, edges and red edges. Contractions on a tri-graph of two vertices u and v recolour edges according to the following rules depending on the sets of neighbours $N(u)$ and $N(v)$ of u and v: 1) edges from $N(u) \triangle N(v)$ to the new vertex are always red 2) edges from $N(u) \cap N(v)$ to the new vertex are not red, only if both original edges were not red. The twin-width of any graph G is then definied as the smallest integer d, that allows a contraction sequence on G, where the maximum red degree during the sequence is bound by d. Do $K_{t,t}$-free string graphs have bounded twin-width? Similarly, do $H_{t,t}$-free string graphs have bounded twin-width?

4.2 Expected Volume of Stochastic Bounding Box

Sergio Cabello (University of Ljubljana, SI)

Given n points in $\mathbb{R}^d$ together with probabilities for each point, we want to compute the expected volume of the bounding box of these points. Is this problem fixed parameter tractible with respect to d? Is it $\#W[1]$-hard w.r.t. d? Is the dependency on d in the degree needed?

4.3 Average Distortion Embeddings

Aleksandar Nikolov (University of Toronto, CA)

Joint work of Deepanshu Kush, Aleksandar Nikolov, Haohua Tang
Main reference Deepanshu Kush, Aleksandar Nikolov, Haohua Tang: "Near Neighbor Search via Efficient Average Distortion Embeddings", CoRR, Vol. abs/2105.04712, 2021.
URL https://arxiv.org/abs/2105.04712

Suppose (M, dist) is a metric space, and P is a finite set of points in M. A function $f : M \to \ell_2$ is called an embedding with average distortion D w.r.t. P, if $\|f(x)-f(y)\|_2 \leq \mathrm{dist}(x,y)$ for all $x, y \in M$, and $\sum_{x\in P}\sum_{y\in P} ||f(x) - f(y)||_2^2 \geq D^{-2} \sum_{x\in P}\sum_{y\in P} \mathrm{dist}(x,y)^2$. Naor [1] showed, that for any d-dimensional normed space $(X, \|\cdot\|)$, defining the metric $\mathrm{dist}(x,y) = \sqrt{||x-y||_X}$, then any $P \subset X$ embeds into ℓ_2 with average distortion $O(\sqrt{\log d})$. Naor's proof proceeds via duality, and does not give an explicit embedding f. Can we find an explicit f given M and P, and, in particular, can we find an f that is efficiently computable from P?

References

1 Naor, Assaf. An average John theorem. arXiv, 1905.01280, 2020

Remote Participants

- Mikkel Abrahamsen
University of Copenhagen, DK
- Pankaj Kumar Agarwal
Duke University – Durham, US
- Helmut Alt
FU Berlin, DE
- Elena Arseneva
St. Petersburg State University, RU
- Édouard Bonnet
ENS – Lyon, FR
- Karl Bringmann
Universität des Saarlandes – Saarbrücken, DE
- Mickaël Buchet
TU Graz, AT
- Maike Buchin
Ruhr-Universität Bochum, DE
- Sergio Cabello
University of Ljubljana, SI
- Hsien-Chih Chang
Dartmouth College – Hanover, US
- Siu-Wing Cheng
HKUST – Kowloon, HK
- Man-Kwun Chiu
FU Berlin, DE
- Jinhee Chun
Tohoku University – Sendai, JP
- Éric Colin de Verdière
University Gustav Eiffel – Marne-la-Vallée, FR
- Jacobus Conradi
Universität Bonn, DE
- Arnaud de Mesmay
University Paris-Est – Marne-la-Vallée, FR
- Anne Driemel
Universität Bonn, DE
- Ioannis Emiris
University of Athens & Athena Research Center, GR
- Matias Korman
Siemens EDA – Wilsonville, US
- Joseph S. B. Mitchell
Stony Brook University, US
- Aleksandar Nikolov
University of Toronto, CA
- Eunin Oh
POSTECH – Pohang, KR
- Steve Y. Oudot
INRIA Saclay – Île-de-France, FR
- Evanthia Papadopoulou
University of Lugano, CH
- Zuzana Patáková
Charles University – Prague, CZ
- Jeff M. Phillips
University of Utah – Salt Lake City, US
- Benjamin Raichel
University of Texas – Dallas, US
- Natan Rubin
Ben Gurion University – Beer Sheva, IL
- Maria Saumell
The Czech Academy of Sciences – Prague & Czech Technical University in Prague, CZ
- Lena Schlipf
Universität Tübingen, DE
- Christiane Schmidt
Linköping University, SE
- Donald Sheehy
North Carolina State University – Raleigh, US
- Kolay Sudeshna
Indian Institute of Technology – Kharagpur, IN
- Martin Tancer
Charles University – Prague, CZ
- Csaba Tóth
California State University – Los Angeles, US
- Uli Wagner
IST Austria – Klosterneuburg, AT
- Carola Wenk
Tulane University – New Orleans, US
- Sue Whitesides
University of Victoria, CA