

Computational Geometry

Siu-Wing Cheng^{*1}, Maarten Löffler^{*2}, Jeff M. Phillips^{*3}, and Aleksandr Popov^{†4}

1 HKUST – Hong Kong, CN. scheng@cse.ust.hk

2 Utrecht University, NL & Tulane University – New Orleans, LA, US.
m.loffler@uu.nl

3 University of Utah – Salt Lake City, UT, US. jeffp@cs.utah.edu

4 TU Eindhoven, NL. a.popov@tue.nl

Abstract

This report documents the program and the outcomes of the Dagstuhl Seminar 23221 “Computational Geometry”. The seminar was held from May 29th to June 2nd, 2023, and 39 participants from various countries attended it, including two remote participants. Recent advances in computational geometry were presented and discussed, and new challenges were identified. This report collects the abstracts of the talks and the open problems presented at the seminar.

Seminar May 29 – June 2, 2023 – <https://www.dagstuhl.de/23221>

2012 ACM Subject Classification Theory of computation → Computational geometry; Theory of computation → Design and analysis of algorithms; Mathematics of computing → Discrete mathematics

Keywords and phrases Algorithms, Combinatorics, Geometric Computing, Reconfiguration, Uncertainty

Digital Object Identifier 10.4230/DagRep.13.5.165

1 Executive Summary

Siu-Wing Cheng (HKUST – Hong Kong, CN, scheng@cse.ust.hk)

Maarten Löffler (Utrecht University, NL & Tulane University – New Orleans, US, m.loffler@uu.nl)

Jeff M. Phillips (University of Utah – Salt Lake City, US, jeffp@cs.utah.edu)

License  Creative Commons BY 4.0 International license
© Siu-Wing Cheng, Maarten Löffler, and Jeff M. Phillips

This Dagstuhl Seminar constituted a biennial gathering of computational geometers at the Dagstuhl venue to share recent results, and further research on some of the most important problems of the time in that field. This year, the seminar focused on two of the most exciting sub-areas within computational geometry: (1) reconfiguration, and (2) processing and applications of uncertain and probabilistic geometric data. Within the reconfiguration topic, two overview talks focused on triangulation and graph reconfiguration, and on how reconfiguration plays a role in puzzle complexity. A highlight of the seminar were the set of three-dimensional reconfiguration puzzles brought by Ryuhei Uehara, which occupied attendees endlessly, and brought the challenge of modelling such puzzles to life. In the second theme on uncertainty, one overview talk covered uncertainty issues in spatial data, and another focused on how uncertainty connects to differential privacy in geometric settings.

* Editor / Organizer

† Editorial Assistant / Collector



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

Computational Geometry, *Dagstuhl Reports*, Vol. 13, Issue 5, pp. 165–181

Editors: Siu-Wing Cheng, Maarten Löffler, Jeff M. Phillips, and Aleksandr Popov



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Other results were shared by participants connecting these topics to diverse motivations ranging from robotics to data analysis to graph drawing. Exciting open problems were proposed, and were used to focus the discussion for the span of the seminar.

2 Table of Contents

Executive Summary

Siu-Wing Cheng, Maarten Löffler, and Jeff M. Phillips 165

Overview of Talks

Uncertain Points and Trajectories	
<i>Kevin Buchin</i>	169
Recent Progress in Geometric Random Walks and Sampling	
<i>Ioannis Emiris</i>	169
Geometric Reconfiguration: Triangulations, Spanning Trees, Graphs	
<i>Anna Lubiw</i>	170
Oblivious Sketching for Sparse Linear Regression	
<i>Alexander Munteanu</i>	170
Parameterized Algorithm for the Planar Disjoint Path Problem	
<i>Eunjin Oh</i>	171
Homology of Reeb Spaces and the Borsuk–Ulam Theorem	
<i>Salman Parsa</i>	171
Random Projections for Curves in High Dimensions	
<i>Ioannis Psarros</i>	172
Using SAT Solvers in Combinatorics, Combinatorial Geometry, and Graph Drawing	
<i>Manfred Scheucher</i>	172
Modular Robot Reconfiguration: Sliding Squares	
<i>Willem Sonke</i>	173
Deep Neural Network Training Acceleration with Geometric Data Structures	
<i>Konstantinos Tsakalidis</i>	173
Computational Complexity of Puzzles (In the Context of Reconfiguration)	
<i>Ryuhei Uehara</i>	174
Combinatorial Reconfiguration via Triangle Flips	
<i>Birgit Vogtenhuber</i>	174
Differential Privacy and Computational Geometry	
<i>Ke Yi</i>	175

Open Problems

Largest Precise Subset (Making Points Precise)	
<i>Peyman Afshani</i>	176
Partial Vertex Cover in Planar Bipartite Graphs	
<i>Peyman Afshani, Kevin Buchin, Fabian Klute, Willem Sonke, and Ryuhei Uehara</i>	177
Flipping Spanning Trees	
<i>Anna Lubiw</i>	177
Shortest Path Retaining k Obstacles	
<i>Subhash Suri</i>	178


Why Are Polytime and NP-Hard Switched for Fréchet and Hausdorff?	
<i>Carola Wenk</i>	179
Participants	180
Remote Participants	181

3 Overview of Talks

This section contains short abstracts of the various talks that were given throughout the seminar. Four speakers (K. Buchin, A. Lubiw, R. Uehara, and K. Yi) were invited to give an overview of the state of the art and to highlight challenges in the two seminar themes, and helped set the stage for the rest of the seminar. In addition, nine participants gave shorter talks throughout the week, which helped exchange ideas and focus the discussions.

3.1 Uncertain Points and Trajectories

Kevin Buchin (TU Dortmund, DE, kevin.buchin@tu-dortmund.de)


License  Creative Commons BY 4.0 International license
© Kevin Buchin

Location data often comes with some error. For instance, the measurements might be imprecise, or the exact location might be obfuscated to preserve privacy. Nonetheless, most algorithms treat these data as if the exact locations are known. The computational geometry of uncertain points is about designing algorithms that deal with the uncertainty explicitly.

In my talk, I first present how uncertain points are commonly modelled and give an overview of algorithmic results on uncertain points for a variety of geometric problems. Then I focus on uncertain trajectories and discuss in particular the problem of computing the Fréchet distance between uncertain curves in detail. I finish my talk with a discussion of open challenges.

3.2 Recent Progress in Geometric Random Walks and Sampling

Ioannis Z. Emiris (Athena Research Center – Marousi, GR & National and Kapodistrian University of Athens – Zografou, GR, emiris@athenarc.gr)


License  Creative Commons BY 4.0 International license
© Ioannis Emiris

Joint work of Ioannis Emiris, Apostolos Chalkis, Vissarion Fisikopoulos

We overview recent advances in geometric random walks for producing a sample of points in convex polyhedra, and applications of such random samples. The main motivation comes from polynomial-time randomized approximation schemes for computing the volume of H-polytopes, i.e. presented as intersection of halfspaces; our methods tackle inputs whose dimension is in the thousands. We extend these algorithms to V-polytopes and zonotopes and present algorithmic results and implementations that show that in practice such polytopes can also be handled efficiently in dimension around 100. Here we have used the existing random walks but optimized the sequence of bodies employed in the standard multiphase Monte Carlo approach. The second part of the talk concentrates on non-linear bodies such as spectrahedra, and an application in systems biology. We conclude with open questions on employing these methods to optimization, and to devising efficient approximate polytope oracles.

3.3 Geometric Reconfiguration: Triangulations, Spanning Trees, Graphs

Anna Lubiw (University of Waterloo, CA, alubiw@uwaterloo.ca)

License  Creative Commons BY 4.0 International license
© Anna Lubiw

I discuss recent results and open questions on three topics in geometric reconfiguration.


1. Reconfiguring triangulations of a set of n points in the plane. A flip, or reconfiguration step, replaces one edge by another to give a new triangulation. It is known that if some edges are *fixed* (i.e. required to be present in all the triangulations) then reconfiguration is still possible via the constrained Delaunay triangulation. I show some results about *forbidding* some edges. Reconfiguration is no longer always possible, i.e. the flip graph may become disconnected, even by forbidding a single edge. However, for points in convex position, we must forbid $n - 3$ edges in order to disconnect the flip graph.
2. Reconfiguring non-crossing spanning trees on a set of n points in the plane. A flip replaces one edge by a new edge to give a new non-crossing spanning tree. Reconfiguring one non-crossing spanning tree to another takes at most $2n$ flips [1] and (in the worst case) at least $1.5n$ flips, even for points in convex position. I show some improvements in the upper bound for points in convex position, and for some cases where one tree is a path. There are still gaps, and the complexity of finding the minimum flip distance (in P or NP-complete) is open.
3. Reconfiguring planar graph drawings by moving the points (aka morphing). One straight-line planar drawing of an n -vertex graph can be morphed to another (with the same combinatorial embedding) via $\mathcal{O}(n)$ *unidirectional morphs* that move vertices along parallel lines at uniform speeds. If the original drawings lie on a small grid, can the $\mathcal{O}(n)$ intermediate drawings be constrained to a small grid?

References

- 1 David Avis and Komei Fukuda. Reverse Search for Enumeration. *Discrete Applied Mathematics*, 65(1–3):21–46, 1996. doi:10.1016/0166-218X(95)00026-N.

3.4 Oblivious Sketching for Sparse Linear Regression

Alexander Munteanu (TU Dortmund, DE, alexander.munteanu@tu-dortmund.de)

License  Creative Commons BY 4.0 International license
© Alexander Munteanu

Joint work of Alexander Munteanu, Tung Mai, Cameron Musco, Anup B. Rao, Chris Schwiegelshohn, David P. Woodruff

Main reference Tung Mai, Alexander Munteanu, Cameron Musco, Anup Rao, Chris Schwiegelshohn, David Woodruff: “Optimal Sketching Bounds for Sparse Linear Regression”, in Proc. of The 26th International Conference on Artificial Intelligence and Statistics, Proceedings of Machine Learning Research, Vol. 206, pp. 11288–11316, PMLR, 2023.

URL <https://proceedings.mlr.press/v206/mai23a.html>

Oblivious sketching enables efficient algorithms for analysing data streams and distributed data by reducing the number of data points while preserving a $(1 + \varepsilon)$ -approximation for various regression loss functions such as l_p -norms, or logistic loss. For dense models, a sketching complexity of $\Omega(d)$ is immediate. For very high-dimensional data, model sparsity is a common assumption, where we do not regress on all, but only on at most $k \ll d$ dimensions. While the computational complexity for solving this problem becomes hard, the assumption allows us to sketch to a smaller $o(d)$ size. We show that reducing to $\frac{k \log(d/k)}{\varepsilon^2}$ is

essentially optimal for any sketching algorithm under several regression losses, combining high-dimensional probability and geometry, information-theoretic arguments, and metric embedding theory.

3.5 Parameterized Algorithm for the Planar Disjoint Path Problem

Eunjin Oh (POSTECH – Pohang, KR, eunjin.oh@postech.ac.kr)

License © Creative Commons BY 4.0 International license
© Eunjin Oh

Joint work of Eunjin Oh, Kyeongjin Cho, Seunghyeok Oh

Main reference Kyungjin Cho, Eunjin Oh, Seunghyeok Oh: “Parameterized Algorithm for the Disjoint Path Problem on Planar Graphs: Exponential in k^2 and Linear in n ”, in Proc. of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023, pp. 3734–3758, SIAM, 2023.

URL <https://doi.org/10.1137/1.9781611977554.ch144>

In this talk, I present a parameterized algorithm for the planar disjoint paths problem. Given a planar graph $G = (V, E)$ and a set $T = \{(s_1, t_1), \dots, (s_k, t_k)\}$ of terminal pairs, the goal is to compute a set of vertex-disjoint paths, each connecting s_i and t_i . This problem is NP-hard if we measure the running time as a function of the input size, but if we measure the running time as a function of n and k , we can achieve a non-trivial bound. In this talk, I give a sketch of this algorithm with running time of $2^{\mathcal{O}(k^2)}n$. This improves the two previously best known algorithms running in $2^{\mathcal{O}(k^2)}n^6$ and $2^{2^{\mathcal{O}(k)}}n$ time, respectively.

3.6 Homology of Reeb Spaces and the Borsuk–Ulam Theorem

Salman Parsa (DePaul University – Chicago, IL, US, s.parsa@depaul.edu)

License © Creative Commons BY 4.0 International license
© Salman Parsa

Joint work of Salman Parsa, Sarah Percival

In this talk, I prove an extension of the Borsuk–Ulam theorem for maps from 2-sphere into \mathbb{R} . The extension says that there are always two antipodal points $S^2 \ni x, -x$ such that $f(x) = f(-x)$ and the two points are connected in the preimage.

The proof uses the concept of the Reeb graph. We also consider the relationship between extra homology of the Reeb space of $f : S^n \rightarrow \mathbb{R}^{n-1}$ and the existence of the analogous extensions of the Borsuk–Ulam theorem.

3.7 Random Projections for Curves in High Dimensions

Ioannis Psarros (Athena Research Center – Marousi, GR, ipsarros@athenarc.gr)

License © Creative Commons BY 4.0 International license
© Ioannis Psarros

Joint work of Ioannis Psarros, Dennis Rohde

Main reference Ioannis Psarros, Dennis Rohde: “Random Projections for Curves in High Dimensions”, in Proc. of the 39th International Symposium on Computational Geometry, SoCG 2023, June 12-15, 2023, Dallas, Texas, USA, LIPIcs, Vol. 258, pp. 53:1–53:15, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023.

URL <https://doi.org/10.4230/LIPIcs.SocG.2023.53>

Modern time series analysis requires the ability to handle datasets that are inherently high-dimensional; examples include applications in climatology, where measurements from numerous sensors must be taken into account, or inventory tracking of large shops, where the dimension is defined by the number of tracked items. The standard way to mitigate computational issues arising from the high dimensionality of the data is by applying some dimension reduction technique that preserves the structural properties of the ambient space. The dissimilarity between two time series is often measured by “discrete” notions of distance, e.g. the dynamic time warping or the discrete Fréchet distance. Since all these distance functions are computed directly on the points of a time series, they are sensitive to different sampling rates or gaps. The continuous Fréchet distance offers a popular alternative which aims to alleviate this by taking into account all points on the polygonal curve obtained by linearly interpolating between any two consecutive points in a sequence.

We study the ability of random projections à la Johnson and Lindenstrauss to preserve the continuous Fréchet distance of polygonal curves by effectively reducing the dimension.

3.8 Using SAT Solvers in Combinatorics, Combinatorial Geometry, and Graph Drawing

Manfred Scheucher (TU Berlin, DE, scheucher@math.tu-berlin.de)

License © Creative Commons BY 4.0 International license
© Manfred Scheucher

In this talk, we discuss how modern SAT solvers can be used to tackle mathematical problems. We discuss various problems to give the audience a better understanding, which might be tackled in this fashion, and which might not. Besides the naïve SAT formulation further ideas are sometimes required to tackle problems. Additional constraints such as statements which hold “without loss of generality” might need to be added so that solvers terminate in reasonable time.

3.9 Modular Robot Reconfiguration: Sliding Squares

Willem Sonke (TU Eindhoven, NL, w.m.sonke@tue.nl)

License © Creative Commons BY 4.0 International license
© Willem Sonke

Joint work of Willem Sonke Hugo A. Akitaya, Erik D. Demaine, Matias Korman, Irina Kostitsyna, Irene Parada, Bettina Speckmann, Ryuhei Uehara, Jules Wulms

Main reference Hugo A. Akitaya, Erik D. Demaine, Matias Korman, Irina Kostitsyna, Irene Parada, Willem Sonke, Bettina Speckmann, Ryuhei Uehara, Jules Wulms: “Compacting Squares: Input-Sensitive In-Place Reconfiguration of Sliding Squares”, in Proc. of the 18th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2022, June 27-29, 2022, Tórshavn, Faroe Islands, LIPIcs, Vol. 227, pp. 4:1–4:19, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.

URL <https://doi.org/10.4230/LIPIcs.SWAT.2022.4>

Modular robots consist of a large number (say n) of small identical modules that can move along each other to change the overall shape of the robot. A well-established model for modular robots is called sliding squares. Here, each module is represented by a square in the 2D grid, which form an edge-connected configuration. Modules can perform two types of moves: slides and convex transitions.

The main question now is: given a source and target configuration, can we find a short move sequence to transform the source into the target? In this talk I show an algorithm named Gather&Compact that finds such a move sequence of length $\mathcal{O}(Pn)$ where P is the circumference of the configurations’ bounding box. Furthermore, I show that it is NP-hard to find a move sequence of minimum length.

3.10 Deep Neural Network Training Acceleration with Geometric Data Structures

Konstantinos Tsakalidis (University of Liverpool, GB, tsakalid@liverpool.ac.uk)

License © Creative Commons BY 4.0 International license
© Konstantinos Tsakalidis

The efficiency of deep learning applications deteriorates significantly as the sizes of the training data and of the neural networks grow larger. In this talk we identify beyond-state-of-the-art open problems in the intersection of deep learning with computational geometry. Motivated by the recent application of dynamic data structures for geometric halfspace range searching in the acceleration of deep neural networks’ training and preprocessing complexity, we revisit efficient algorithms for constructing geometric multi-dimensional data structures and maintaining them dynamically.

3.11 Computational Complexity of Puzzles (In the Context of Reconfiguration)

Ryuhei Uehara (JAIST – Nomi, Ishikawa, JP, uehara@jaist.ac.jp)

License © Creative Commons BY 4.0 International license

© Ryuhei Uehara

Main reference Ryuhei Uehara: “Computational Complexity of Puzzles and Related Topics”, Interdisciplinary Information Sciences, Vol. 29(2), pp. 119–140, 2023.

URL <https://doi.org/10.4036/iis.2022.R.06>

I first give a short history of computational complexity of puzzles and games, including combinatorial reconfiguration ones. In this context, there are three open problems.

1. The Rubik’s cube has a similar property to one of the $n^2 - 1$ puzzles. Then what is the counterpart of the sliding block puzzle?
 - Is there some PSPACE-complete problem in general?
 - What about rectangular faces?
2. Reconfiguration of triangulations of a simple polygon also has a similar property. Then, again, can we have a counterpart of the sliding block puzzle that leads us to PSPACE-complete variant in general?
 - How about non-simple polygon with holes?
 - Can the diameter of the configuration space be super-poly? (Otherwise, it is in NP since we have a poly-length witness.)
3. Computational complexity of *slide-and-pack* puzzles. By some observations, it seems to be PSPACE-complete in general. Do we have some tractable restrictions?

3.12 Combinatorial Reconfiguration via Triangle Flips

Birgit Vogtenhuber (TU Graz, AT, bvogt@ist.tugraz.at)

License © Creative Commons BY 4.0 International license

© Birgit Vogtenhuber

Joint work of Birgit Vogtenhuber, Oswin Aichholzer, Man-Kwun Chiu, Stefan Felsner, Hung P. Hoang, Michael Hoffmann, Yannic Maus, Johannes Obenaus, Sandro Roch, Manfred Scheucher, Alexandra Weinberger

In this talk we discuss the reconfiguration of arrangements of simple curves in the plane or on the sphere via triangle flips (a.k.a. Reidemeister moves of Type 3). This operation refers to the act of moving one edge of a triangular cell formed by three pairwise crossing curves over the opposite crossing of the cell, via a local transformation. We study two types of arrangements, namely, arrangements of pseudocircles in the plane and simple drawings of graphs on the sphere.

An arrangement of pseudocircles is a finite collection of simple closed curves in the plane such that every pair of curves is either disjoint or intersects in two crossing points. We show that triangle flips induce a connected flip graph on intersecting arrangements, i.e. on arrangements where every pair of pseudocircles intersects. To obtain this result, we first show that every intersecting arrangement can be flipped into some cylindrical arrangement, i.e. an arrangement where a single point stabs the interior of every pseudocircle. Then we show that the cylindrical arrangement can be flipped to a canonical arrangement, which also shows the connectivity of cylindrical intersecting arrangements of pseudocircles under triangle flips. With a careful analysis we obtain that the diameter of both flip graphs is cubic in

the number of pseudocircles. The construction of the two flipping sequences makes essential use of variants of the sweeping lemma for pseudocircle arrangements due to Snoeyink and Hershberger [1].


A simple drawing of a labelled graph is a drawing in which the vertices are distinct points and the edges are simple curves connecting their endpoints. Moreover, any two edges share at most one point, which is either a common endpoint or a crossing. The extended rotation system of such a drawing is the collection of the rotations of all vertices and crossings of the drawing. The rotation of a vertex or crossing is the cyclic order in which the edges emanate from it. Gioan's Theorem states that for any two simple drawings of the complete graph K_n with the same crossing edge pairs, one drawing can be transformed into the other by a sequence of triangle flips. We investigate to what extent Gioan-type theorems can be obtained for wider classes of graphs. A necessary (but in general not sufficient) condition for two drawings of a graph to be transformable into each other by a sequence of triangle flips is that they have the same ERS. We show that for the large class of complete multipartite graphs, this necessary condition is in fact also sufficient and give bounds on the diameter of the resulting flip graph. We further show that this result is essentially tight in the sense that there exist drawings of multipartite graphs plus one edge or minus two edges which cannot be transformed into each other via triangle flips.

References

- 1 Jack Scott Snoeyink and John E. Hershberger. Sweeping Arrangements of Curves. In *Proceedings of the 5th Annual Symposium on Computational Geometry (SCG 1989)*, pages 354–363, New York, NY, US, 1989. Association for Computing Machinery. doi:10.1145/73833.73872.

3.13 Differential Privacy and Computational Geometry

Ke Yi (HKUST – Hong Kong, CN, yike@cse.ust.hk)

License  Creative Commons BY 4.0 International license
© Ke Yi


Differential privacy has become the de facto standard for personal information privacy, and has been recently widely adopted in both governments and the industry. Roughly speaking, a differentially private algorithm should have indistinguishable output distributions on neighbouring instances, which differ by one individual's data. Such a definition also applies to geometric data, where one input point set has one more point than the other. In this talk, I give an overview of existing differentially private algorithms for geometric data, including range counting, mean estimation, and convex hull. I also discuss geometric privacy, a variant of differential privacy that is more suitable for certain geometric problems.

4 Open Problems

This section contains short summaries of specific relevant open questions that were proposed at the start of the seminar and discussed in more depth during the week. Partial progress was made towards resolving these questions, and we expect there will be further collaboration on these topics between seminar participants beyond the duration of the seminar.

4.1 Largest Precise Subset (Making Points Precise)

Peyman Afshani (Aarhus University, DK, peyman@cs.au.dk)

License  Creative Commons BY 4.0 International license
© Peyman Afshani

Imprecision through Precision

In geometric computation, it is useful to be able to answer geometric predicates, e.g. a sidedness test with respect to points in the plane. To model imprecision, it is common to replace points with disks in 2D.

First we fix our geometric predicate. If we replace every point with an imprecise point modelled as a circle, then the test can still be resolved on the points if there is no line that intersects three circles; let S be the set of input circles. We call S a *precise point set* (w.r.t. the sidedness test). The motivation is that any algorithm that can use only sidedness test can be run without any modifications on the set S .

Some Open Problems

Selecting a large precise set. One open question could be to select a large precise subset of an imprecise points. Let S be a set of circles such that any line intersects at most ℓ circles. What is the largest precise subset of $H \subset S$ one can select, so that no three circles in H should be stabbed by any line?

- Using standard techniques (random sample and refine), it is easy to show that we can take $H = \Omega(\sqrt{n}/\ell)$.
- There are also previous results on similar questions but when input is a set of points rather than circles. This line of work is studied under the name of *general position subset selection problem*. For example, it is known that given a set of n points such that no ℓ of them are on a line, one can select a subset of size $\Omega(\sqrt{n/\log \ell})$ that is in general position [1]. Observe that this bound is much better than the our trivial bound. However, to prove it, the authors use incidence bounds that we do not have for a set of circles.

Open question 1. Improve the bound for the circles.

Open question 2. Motivated by the techniques in the above mentioned paper, we can also ask the following question. Let S be a set of circles such that no ℓ of them are on a line. What is the maximum number of triples of circles (c_i, c_j, c_k) such that all three can be stabbed by a line? The trivial upper bound is $\mathcal{O}(n^2\ell^3)$, because we have at most $\mathcal{O}(n^2)$ combinatorially different lines and each line can stab ℓ , so we can select $\mathcal{O}(\ell^3)$ triples out of them. The trivial lower bound is $\Omega(\min(n^2, n\ell^2))$ obtained by either using a $(3 \times n/3)$ -grid or placing n/ℓ groups of ℓ circles on a line.

Open question 3. For the general position subset selection problem, there exists an $o(n)$ upper bound through Hales–Jewett theorem (referenced in the paper [1]). Can we improve the upper bound for the circles? Note that for this upper bound problem, having the possibility of replacing a point by a circle should make the upper bound problem easier.

References

- 1 Michael S. Payne and David R. Wood. On the General Position Subset Selection Problem. *SIAM Journal on Discrete Mathematics*, 27(4):1727–1733, 2013. doi:10.1137/120897493.

4.2 Partial Vertex Cover in Planar Bipartite Graphs

Peyman Afshani (Aarhus University, DK, peyman@cs.au.dk)

Kevin Buchin (TU Dortmund, DE, kevin.buchin@tu-dortmund.de)

Fabian Klute (UPC Barcelona Tech, ES, fmklute@gmail.com)

Willem Sonke (TU Eindhoven, NL, w.m.sonke@tue.nl)

Ryuhei Uehara (JAIST – Nomi, Ishikawa, JP, uehara@jaist.ac.jp)

License © Creative Commons BY 4.0 International license

© Peyman Afshani, Kevin Buchin, Fabian Klute, Willem Sonke, and Ryuhei Uehara

Partial Vertex Cover asks whether we can cover at least t edges by selecting k vertices. In contrast to vertex cover, this problem is NP-hard in bipartite graphs [1]. How about planar bipartite graphs?

It seems that the paper by Caskurlu et al. [1] asks this as an open problem, although “bipartite” is missing in the problem statement.

This problem is inspired by the obstacle-deletion problem by Subhash Suri. If this problem is NP-hard, we could use this fact in the following way for the obstacle-deletion open problem. Since the graph is bipartite, we can find a path in the plane (not in the graph) that crosses every edge exactly once. Now we interpret every vertex as obstacles, and for every edge, we make the corresponding obstacles interlock where the path would like to pass. Now partial vertex cover corresponds to allowing to delete k obstacles while trying to get rid of as many interlockings as possible.

For the obstacle-deletion problem we can give weights to vertices, we can also give weights to edges. Therefore, even weighted versions of partial vertex cover in planar bipartite graphs would be of interest.

References

- 1 Bugra Caskurlu, Vahan Mkrtchyan, Ojas Parekh, and K. Subramani. Partial Vertex Cover and Budgeted Maximum Coverage in Bipartite Graphs. *SIAM Journal on Discrete Mathematics*, 31(3):2172–2184, 2017. doi:10.1137/15M1054328.

4.3 Flipping Spanning Trees

Anna Lubiw (University of Waterloo, CA, alubiw@uwaterloo.ca)

License © Creative Commons BY 4.0 International license

© Anna Lubiw

During the seminar, the following question from the talk given by Anna Lubiw was investigated. Given a set P of n points in the plane and two non-crossing spanning trees T_r , T_b on P , what is their flip distance $\text{dist}(T_r, T_b)$ in terms of n in the worst case?

The following bounds are known:

1. lower bound of $\frac{3}{2} \cdot n - 5$ [3];
2. upper bound of $2n - 4$ [1];
3. upper bound of $2n - \Omega(\sqrt{n})$ if P is in convex position [2].


Can these bounds be improved? Are there interesting special cases with tighter bounds?

References

- 1 David Avis and Komei Fukuda. Reverse Search for Enumeration. *Discrete Applied Mathematics*, 65(1–3):21–46, 1996. doi:10.1016/0166-218X(95)00026-N.
- 2 Nicolas Bousquet, Valentin Gledel, Jonathan Narboni, and Théo Pierron. A Note on the Flip Distance between Non-crossing Spanning Trees. 2023. arXiv:2303.07710v1 [cs.CG].
- 3 M. Carmen Hernando, Ferrán Hurtado, Alberto Márquez, Mercè Mora, and Marc Noy. Geometric Tree Graphs of Points in Convex Position. *Discrete Applied Mathematics*, 93(1):51–66, 1999. doi:10.1016/S0166-218X(99)00006-2.

4.4 Shortest Path Retaining k Obstacles

Subhash Suri (University of California – Santa Barbara, US, suri@cs.ucsb.edu)

License  Creative Commons BY 4.0 International license
© Subhash Suri

Suppose we are in polygonal domain in the plane with n vertices in total. The obstacles (holes) do not have to be convex. Given two points s and t , what is the Euclidean length of the shortest path between s and t that may ignore k obstacles? In other words, if we are allowed to remove k obstacles, what is the length of the shortest obstacle-free path between s and t ? Formulated differently, let \mathcal{B} be the set of m disjoint simple polygons that are the obstacles. Define $d_{\mathcal{B}'}(s, t)$ as the shortest path distance from s to t with only $\mathcal{B}' \subseteq \mathcal{B}$ present. For given k and s and t , we wish to find

$$d_k(s, t) = \min_{|\mathcal{B} \setminus \mathcal{B}'|=k} d_{\mathcal{B}'}(s, t).$$


It is known how to compute the solution in polynomial time for convex obstacles [2]. A similar problem with overlapping obstacles is known to be intractable even for very simple obstacle shapes [2]. The weighted problem, where each obstacle has a cost and we are given a budget, is NP-hard even for vertical line segments as obstacles [1]. We would like to find out if the problem presented here is hard.

References

- 1 Pankaj K. Agarwal, Neeraj Kumar, Stavros Sintos, and Subhash Suri. Computing Shortest Paths in the Plane with Removable Obstacles. In *Proceedings of the 16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018)*, number 101 in Leibniz International Proceedings in Informatics (LIPIcs), article 5, Dagstuhl, DE, 2018. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SWAT.2018.5.
- 2 John E. Hersberger, Neeraj Kumar, and Subhash Suri. Shortest Paths in the Plane with Obstacle Violations. *Algorithmica*, 82:1813–1832, 2020. doi:10.1007/s00453-020-00673-y.

4.5 Why Are Polytime and NP-Hard Switched for Fréchet and Hausdorff?

Carola Wenk (Tulane University – New Orleans, US, cwenk@tulane.edu)

License  Creative Commons BY 4.0 International license
© Carola Wenk

Consider a common definition of an uncertain trajectory as a sequence of uncertain points, where each uncertain point is some compact connected region, most commonly a disk in \mathbb{R}^2 or an interval in \mathbb{R} . A true location must be inside the region, but we do not know where. In this setting, it is natural to generalise the polyline metrics to ask for the maximum and minimum possible values. In particular, we say a polyline realises an uncertain trajectory if it is formed by a sequence of precise points, taken in the correct order from the uncertainty regions. Then we can ask for the lower bound and the upper bound distance, that is, given two uncertain curves, we want to find the minimum and the maximum distance between them over all realisations. Within this framework, we can use different uncertainty models. We can also use different distance metrics, most commonly, the Hausdorff or the Fréchet distance (or their variants).

Some of these combinations have previously been studied [1, 2, 3]; all either have a relatively simple polynomial-time solution, or are NP-hard. We can ask the following questions.

1. For the directed Fréchet distance, the lower bound problem is in P, and the upper bound problem is NP-hard. For the directed Hausdorff distance, the situation is essentially reversed. Similar dichotomies exist for the weak (discrete) Fréchet distance, where different settings turn out to be NP-hard compared to the Fréchet distance. Is there something in common among the NP-hard variants? Why are the lower and upper bound problems reversed for NP-hardness?
2. Can we fill in the gaps by studying the remaining variations, both for the Hausdorff and the Fréchet distance, and conclude for each whether it is NP-hard or solvable in polynomial time?

References

- 1 Kevin Buchin, Chenglin Fan, Maarten Löffler, Aleksandr Popov, Benjamin Raichel, and Marcel Roeloffzen. Fréchet Distance for Uncertain Curves. *ACM Transactions on Algorithms*, 19(3), article 29, 2023. doi:10.1145/3597640.
- 2 Kevin Buchin, Maarten Löffler, Tim Ophelders, Aleksandr Popov, Jérôme Urhausen, and Kevin Verbeek. Computing the Fréchet Distance between Uncertain Curves in One Dimension. *Computational Geometry: Theory and Applications*, 109, article 101923, 2023. doi:10.1016/j.comgeo.2022.101923.
- 3 Christian Knauer, Maarten Löffler, Marc Scherfenberg, and Thomas Wolle. The Directed Hausdorff Distance between Imprecise Point Sets. *Theoretical Computer Science*, 412(32):4173–4186, 2011. doi:10.1016/j.tcs.2011.01.039.

Participants

- Peyman Afshani
Aarhus University, DK
- Hee-Kap Ahn
POSTECH – Pohang, KR
- Håvard Bakke Bjerkevik
TU Graz, AT
- Kevin Buchin
TU Dortmund, DE
- Maike Buchin
Ruhr-Universität Bochum, DE
- Siu-Wing Cheng
HKUST – Hong Kong, CN
- Guilherme D. da Fonseca
Aix-Marseille University, FR
- Vincent Despré
LORIA – Nancy, FR
- Ioannis Emiris
Athena Research Center –
Maroussi, GR
- Linda Kleist
TU Braunschweig, DE
- Fabian Klute
UPC Barcelona Tech, ES
- Maarten Löffler
Utrecht University, NL & Tulane
University – New Orleans, US
- Zuzana Masárová
IST Austria –
Klosterneuburg, AT
- Bojan Mohar
Simon Fraser University –
Burnaby, CA
- Alexander Munteanu
TU Dortmund, DE
- Martin Nöllenburg
TU Wien, AT
- Eunjin Oh
POSTECH – Pohang, KR
- Irene Parada
UPC Barcelona Tech, ES
- Salman Parsa
DePaul University – Chicago, US
- Zuzana Patáková
Charles University – Prague, CZ
- Jeff M. Phillips
University of Utah –
Salt Lake City, US
- Aleksandr Popov
TU Eindhoven, NL
- Ioannis Psarros
Athena Research Center –
Maroussi, GR
- Manfred Scheucher
TU Berlin, DE
- Raimund Seidel
Universität des Saarlandes –
Saarbrücken, DE
- Willem Sonke
TU Eindhoven, NL
- Subhash Suri
University of California –
Santa Barbara, US
- Monique Teillaud
INRIA Nancy Grand Est –
Villers-lès-Nancy, FR
- Konstantinos Tsakalidis
University of Liverpool, GB
- Torsten Ueckerdt
Karlsruher Institut für
Technologie, DE
- Ryuhei Uehara
JAIST – Nomi, Ishikawa, JP
- Birgit Vogtenhuber
TU Graz, AT
- Hubert Wagner
University of Florida –
Gainesville, US
- Emo Welzl
ETH Zürich, CH
- Carola Wenk
Tulane University –
New Orleans, US
- Sang Duk Yoon
Sungshin Women's University –
Seoul, KR
- Yelena Yuditsky
Université Libre de Bruxelles, BE



Remote Participants

■ Anna Lubiw
University of Waterloo, CA

■ Ke Yi
HKUST – Hong Kong, CN