Report from Dagstuhl Seminar 25372

Precision in Geometric Algorithms

Mikkel Abrahamsen*[1], Sándor Kisfaludi-Bak*[2], Linda Kleist*[3], and Till Miltzow*[4]

1 **University of Copenhagen, DK.** miab@di.ku.dk
2 **Aalto University, FI.** sandor.kisfaludi-bak@aalto.fi
3 **Universität Hamburg, DE.** linda.kleist@uni-hamburg.de
4 **Utrecht University, NL.** t.miltzow@googlemail.com

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 25372 "Precision in Geometric Algorithms". This seminar was an opportunity for a get together of researchers interested in geometric problems that require high precision of the coordinates to find a correct solution.

Seminar September 7–12, 2025 – https://www.dagstuhl.de/25372
2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms
Keywords and phrases Computational Geometry, Real Complexity Theory
Digital Object Identifier 10.4230/DagRep.15.9.21

1 Executive Summary

Till Miltzow (Utrecht University, NL)

The Dagstuhl Seminar "Precision in Geometric Algorithms" (25372) brought together researchers working on computational geometry, real-complexity theory, and geometric computation models that require high-precision reasoning. The seminar aimed to understand how geometric problems behave when precision, real-number computation, and continuous models become central, and to explore the algorithmic, structural, and complexity-theoretic consequences.

The invited talks covered a broad spectrum: geometric graph theory in hyperbolic spaces; optimal convex-hull reconstruction from imprecise data; ER-complete recognition problems for geometric intersection graphs; oracle separations in the real polynomial hierarchy; new approximation schemes for geometric multimatching; the complexity of the boundary–boundary art gallery problem; and dynamic Steiner spanners in curved spaces. Together, these contributions showcased how precision constraints shape both the geometry and the complexity of algorithmic problems.

Several working groups produced substantial progress. One group extended Fréchet-distance techniques to more than two curves and proved meaningful lower bounds via reductions from 3OV. Another initiated the study of "Devil's Games," a class of infinite-move combinatorial games linked to the first-order theory of the reals. Others explored realization spaces of geometric graph representations, online packing of convex objects, sparse geometric emulators, and the flip distance needed to eliminate crossings in geometric matchings.

* Editor / Organizer

Precision in Geometric Algorithms, *Dagstuhl Reports*, Vol. 15, Issue 9, pp. 21–37
Editors: Mikkel Abrahamsen, Sándor Kisfaludi-Bak, Linda Kleist, and Till Miltzow
DAGSTUHL REPORTS Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The open-problem session highlighted future challenges: recognizing strongly hyperbolic disk graphs, understanding polygonal knot realization spaces and their potential universality, establishing $\exists\mathbb{R}$-completeness of continuous distance problems, efficiently computing weak circle representations of planar graphs, and bounding fixed points of compositions of monotone polynomials.

Beyond the technical program, the week was marked by a warm and collaborative atmosphere. Discussions continued naturally over meals, hikes, and informal gatherings, helping participants strengthen existing collaborations and spark new ones. Many attendees commented that the social setting – relaxed yet intellectually charged – played a major role in enabling deep, productive exchanges.

Overall, the seminar strengthened connections between computational geometry, real-number computation, and complexity theory, identifying multiple promising directions where precision requirements fundamentally reshape classical algorithmic questions.

2 Table of Contents

3 Overview of Talks

3.1 Graphs in Hyperbolic Geometry

Thomas Bläsius (KIT – Karlsruher Institut für Technologie, DE)

Hyperbolic geometry is a non-Euclidean geometry where the parallel axiom is negated. While the hyperbolic plane behaves locally like the Euclidean plane, it behaves very different beyond that. One crucial difference is that the hyperbolic plane expands exponentially. This has various interesting effects for studying graphs in the hyperbolic plane. When embedding graphs into the hyperbolic plane, the exponential expansion can be used to, for example, achieve successful greedy routing. Moreover, when defining graphs from geometric objects, like intersection graphs of equally sized disks, the properties of the hyperbolic plane translate to interesting graphs properties.

3.2 Instance-Optimal Imprecise Convex Hull

Sarita de Berg (Utrecht University, NL)

Joint work of Sarita de Berg, Ivor van der Hoog, Eva Rotenberg, Daniel Rutschmann, Sampson Wong
Main reference Sarita de Berg, Ivor van der Hoog, Eva Rotenberg, Daniel Rutschmann, Sampson Wong: "Instance-Optimal Imprecise Convex Hull", in Proc. of the 33rd Annual European Symposium on Algorithms (ESA 2025), Leibniz International Proceedings in Informatics (LIPIcs), Vol. 351, pp. 25:1–25:15, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025.
URL https://doi.org//10.4230/LIPIcs.ESA.2025.25

Imprecise measurements of a point set $P = (p_1, \dots, p_n)$ can be modelled by a family of regions $F = (R_1, \dots, R_n)$, where each imprecise region $R_i \in F$ contains a unique point $p_i \in P$. A *retrieval* models an accurate measurement by replacing an imprecise region R_i with its corresponding point p_i.

We construct the convex hull of an imprecise point set in the plane, by determining the cyclic ordering of the convex hull vertices of P as efficiently as possible. Efficiency is interpreted in two ways: (i) minimising the number of retrievals, and (ii) the computation time to determine the set of regions that must be retrieved.

Previous works focused on only one of these two aspects: either minimising retrievals or optimising algorithmic runtime. Our contribution is the first to simultaneously achieve both. Let $r(F, P)$ denote the minimal number of retrievals required by any algorithm to determine the convex hull of P for a given instance (F, P). For a family F of n constant-complexity polygons, our main result is a reconstruction algorithm that performs $\Theta(r(F, P))$ retrievals in $O(r(F, P) \log^3 n)$ time.

Compared to previous approaches that achieve optimal retrieval counts, we improve the runtime per retrieval from polynomial to polylogarithmic. We extend the generality of previous results to simple k-gons, to pairwise disjoint disks with radii in $[1, k]$, and to unit disks where at most k disks overlap in a single point. Our runtime scales linearly with k.

3.3 Calculating with Pennies and Marbles

Anna Lubiw (University of Waterloo, CA) and Marcus Schaefer (DePaul University – Chicago, US)

Joint work of Anna Lubiw, Marcus Schaefer
Main reference Anna Lubiw, Marcus Schaefer: "Recognizing Penny and Marble Graphs is Hard for Existential Theory of the Reals", CoRR, Vol. abs/2508.10136, 2025.
URL https://doi.org//10.48550/ARXIV.2508.10136

Penny graphs are contact graphs of unit disks in the plane. We show that recognizing penny graphs is ER-complete, that is as hard as deciding truth in the existential theory of the reals. The problem remains ER-complete even if a combinatorial embedding of the penny graph is given. Penny graphs which are trees can be realized with the centers of the pennies belonging to a grid of double-exponential size. We can also show that recognizing marble graphs, contact graphs of unit balls in three-dimensional space, is ER-complete.

3.4 Separations for RPH

Lucas Meijer (Utrecht University, NL)

Joint work of Thekla Hamm, Lucas Meijer, Till Miltzow, Subhasree Patro
Main reference Thekla Hamm, Lucas Meijer, Tillmann Miltzow, Subhasree Patro: "Oracle Separations for RPH", CoRR, Vol. abs/2502.09279, 2025.
URL https://doi.org//10.48550/ARXIV.2502.09279

While theoretical computer science primarily works with discrete models of computation, like the Turing machine and the wordRAM, there are many scenarios in which introducing real computation models is more adequate. For example, when working with continuous probability distributions for say smoothed analysis, in continuous optimization, computational geometry or machine learning. We want to compare real models of computation with discrete models of computation. We do this by means of oracle separation results.

We define the notion of a *real Turing machine* as an extension of the (binary) Turing machine by adding a real tape. Using those machines, we define and study the real polynomial hierarchy $\mathbb{R}PH$. We are interested in $\mathbb{R}$PH as the first level of the hierarchy corresponds to the well-known complexity class $\exists\mathbb{R}$. It is known that $NP \subseteq \exists\mathbb{R} \subseteq PSPACE$ and furthermore $PH \subseteq \mathbb{R}\text{PH} \subseteq PSPACE$. We are interested to know if any of those inclusions are tight. In the absence of unconditional separations of complexity classes, we turn to oracle separation. We develop a technique that allows us to transform oracle separation results from the binary world to the real world. As applications, we show there are oracles such that:

- $\mathbb{R}PH^O$ proper subset of $PSPACE^O$,
- BQP^O not contained in $\mathbb{R}PH^O$.

Our results hint that $\exists\mathbb{R}$ is strictly contained in PSPACE and that there is a separation between the different levels of the real polynomial hierarchy. We also bound the power of real computations by showing that NP-hard problems are unlikely to be solvable using polynomial time on a realRAM. Furthermore, our oracle separations hint that polynomial-time quantum computing cannot be simulated on an efficient real Turing machine.

3.5 Approximation Algorithm for the Geometric Multimatching Problem

Eunjin Oh (POSTECH – Pohang, KR)

Joint work of Eunjin Oh, Shinwoo An, Jie Xue

Abstract: Let S and T be two sets of points in a metric space with a total of n points. Each point in S and T has an associated value that specifies an upper limit on how many points it can be matched with from the other set. A multimatching between S and T is a way of pairing points such that each point in S is matched with at least as many points in T as its assigned value, and vice versa for each point in T. The cost of a multimatching is defined as the sum of the distances between all matched pairs of points. The geometric multimatching problem seeks to find a multimatching that minimizes this cost. A special case where each point is matched to at most one other point is known as the geometric many-to-many matching problem.

We present two results for these problems when the underlying metric space has a bounded doubling dimension. Specifically, we provide the first near-linear-time approximation scheme for the geometric multimatching problem in terms of the output size. Additionally, we improve upon the best-known approximation algorithm for the geometric many-to-many matching problem, previously introduced by Bandyapadhyay and Xue (SoCG 2024), which won the best paper award at SoCG 2024.

3.6 The Boundary-Boundary Art-Gallery Problem is in NP

Jack Stade (University of Copenhagen, DK)

Main reference Jack Stade: "NP-membership for the boundary-boundary art-gallery problem", CoRR, Vol. abs/2511.01562, 2025.
URL https://doi.org//10.48550/ARXIV.2511.01562

The X-Y art-gallery problem asks to find a minimum set of guards that guard a polygon P, where the guards are restricted to lie in X and must see all of Y. For X,Y∈ {Point, Boundary, Vertex}, this gives 9 different variants. Recent work has determined the complexity of each of these variants; all but the Boundary-Boundary variant were known to be either NP-complete or $\exists\mathbb{R}$-complete.

We complete this classification, showing that the Boundary-Boundary variant is NP-complete. This is somewhat surprising, since the coordinates of guards in an optimal solution might need to be irrational. We show that each coordinate is at worst $p + q\sqrt{r}$, where p, q and r are rational numbers with polynomially many bits. These coordinates give a certificate that can be verified in polynomial time.

3.7 Near-Optimal Dynamic Steiner Spanners for Constant-Curvature Spaces

Geert van Wordragen (Aalto University, FI)

Joint work of Sándor Kisfaludi-Bak, Geert van Wordragen
Main reference Sándor Kisfaludi-Bak, Geert van Wordragen: "Near-Optimal Dynamic Steiner Spanners for Constant-Curvature Spaces", CoRR, Vol. abs/2509.01443, 2025.
URL https://doi.org//10.48550/ARXIV.2509.01443

We consider Steiner spanners in Euclidean and non-Euclidean geometries. In the Euclidean setting, a recent line of work initiated by Le and Solomon and further improved by Chang et al. obtained Steiner $(1+\varepsilon)$-spanners of size $O_d(\varepsilon^{(1-d)/2}\log(1/\varepsilon)n)$, nearly matching the lower bounds of Bhore and Tóth.

We obtain Steiner $(1+\varepsilon)$-spanners of size $O_d(\varepsilon^{(1-d)/2}\log(1/\varepsilon)n)$ not only in d-dimensional Euclidean space, but also in d-dimensional spherical and hyperbolic space. For any fixed dimension d, the obtained edge count is optimal up to an $O(\log(1/\varepsilon))$ factor in each of these spaces. Unlike earlier constructions, our Steiner spanners are based on simple quadtrees, and they can be dynamically maintained, leading to efficient data structures for dynamic approximate nearest neighbours and bichromatic closest pair.

In the hyperbolic setting, we also show that 2-spanners in the hyperbolic plane must have $\Omega(n\log n)$ edges, and we obtain a 2-spanner of size $O_d(n\log n)$ in d-dimensional hyperbolic space, matching our lower bound for any constant d. Finally, we give a Steiner spanner with *additive* error ε in hyperbolic space with $O_d(\varepsilon^{(1-d)/2}\log(\alpha(n)/\varepsilon)n)$ edges, where $\alpha(n)$ is the inverse Ackermann function.

Our techniques generalize to closed orientable surfaces of constant curvature as well as to some quotient spaces.

4 Working groups

4.1 The Fréchet distance of several curves

Peyman Afshani (Aarhus University, DK), Karl Bringmann (Universität des Saarlandes – Saarbrücken, DE), Mark de Berg (TU Eindhoven, NL), Omrit Filtser (The Open University of Israel – Ra'anana, IL), Dan Halperin (Tel Aviv University, IL), André Nusser (INRIA – Sophia Antipolis, FR), and Günter Rote (FU Berlin, DE)

Given k polygonal curves $c_1,\ldots,c_k$ in the plane or in some higher-dimensional space with endpoints s_i, t_i and k point robots $r_1,\ldots,r_k$, each robot r_i has to move from s_i to t_i with its center constrained to c_i, without backtracking. The objective is to compute a coordinated motion that minimizes the maximum pairwise distance between robots along their trajectories.

For $k=2$, this is the classic Fréchet distance. The problem formulation can be generalized in various ways, some of which are motivated by applications in robotics and transportation. These include minimum clearance conditions or collisions avoidance (the easiest case being circular robots), and alternative objective functions, several of which give rise to especially challenging variants of the problem.

Background. Computing the Fréchet distance between two curves has been studied extensively in computational geometry since its introduction by Alt and Godau. The vast majority of subsequent work in this area concerns the case of two curves. For this setting, efficient implementations are publicly available. Dumitrescu and Rote presented a 2-approximation algorithm for the case of k curves. They claimed without justification an exact algorithm of running time $O(n^k)$ if each curve has at most n edges. Along different lines, a general solution for the case of k curves has been devised and implemented, adapting sampling-based planning – the standard workhorse of robot algorithms This approach comes with provable guarantees on the quality of the approximation. However, it is currently not competitive in practice: In experiments, already for $k = 2$, its running time was several orders of magnitude slower than the efficient implementations for two curves reported by Bringmann.

Results obtained during the seminar. The classic approach to the computation of the Fréchet distance solves the decision version of the problem (with a given threshold on the maximum distance between the robots) by looking for a monotone path in the free-space diagram F. We discussed the extension of this approach to $k > 2$ curves. For k curves with $n_1, \ldots, n_k$ edges, respectively, the free-space diagram lives in a k-dimensional box consisting of $n_1 \cdot n_2 \cdots n_k$ subboxes (cells). Inside each cell, the free space is the intersection of $\binom{k}{2}$ cylindrical prisms.

For $k = 3$ we managed so solve the decision problem in $O(n_1 n_2 n_3) = O(n^3)$ time. We could show that the reachable set on each rectangular face of a subbox has a restricted structure: It is the intersection of a staircase polygon with an ellipse. Although the staircase polygon may have up to n steps, it can be computed in amortized constant time from the "incoming" faces on each box.

Lower bounds. We showed that a substantially better algorithm with a truly subcubic runtime is unlikely to exist, even if only an approximation of the Fréchet distance with an approximation factor of about 1.1 is desired. For this purpose, we reduced the 3OV problem to the (approximate) Fréchet distance for three curves. In the 3OV (three-orthogonal vectors) problem, we are given three sets A, B, C of n vectors in $\{0, 1\}^d$, and the task is to decide if there are three vectors $x \in A$, $y \in B$, $z \in C$ such that $x_i y_i z_i = 0$ for $i = 1, \ldots, d$.

Variations. For a larger number of $k > 3$ curves, we explored ideas that might lead to an algorithm of running time $O(n^{k^2})$. We also considered an alternate objective function: the radius of the smallest circle enclosing the moving points. This makes the free-space more complicated. For curves in the plane, the free-space on each rectangular face of the free-space diagram is a convex region that is bounded by pieces of line segments, ellipses, and a degree-6 curve.

4.2 Devilish Games and QR

Arnaud de Mesmay (Gustave Eiffel University – Marne-la-Vallée, FR), Lucas Meijer (Utrecht University, NL), Till Miltzow (Utrecht University, NL), Marcus Schaefer (DePaul University – Chicago, US), and Jack Stade (University of Copenhagen, DK)

We worked on a new complexity class denoted by $Q\mathbb{R}$. This complexity class can be defined as all problems that are equivalent to deciding the First Order Theory of the Reals. We describe a framework to show $Q\mathbb{R}$-completeness of Devil's games. Devil's games have two key properties.

- Players alternate in taking turns and
- each turn gives an infinite continuum of possible turns.

There is this very cute classical puzzle, which goes as follows: *After a career of elegant proofs occasionally sabotaged by overlooked edge cases, you find yourself in hell's quiet reading room, where the devil greets repentant theoreticians with a friendly smile. He gestures to a round table and proposes a simple challenge: you and he will take turns placing identical coins on the tabletop, and coins may not overlap. Whoever cannot place a new coin on the table loses. The devil insists you move first, confident that impatience will cloud your reasoning just as it did in life. Win, and he'll grant you a brief return to correct that final paper; lose, and you will spend eternity proofreading the edge cases of others.* Interestingly if you place your first coin precisely at the center and then mirror every move of the devil across that center, you can always respond and never be the one to run out of space first. This idea uses symmetry and is a standard technique in combinatorial game theory. But note, if you had not placed the first coin in the middle. It seems impossible to analyze how to win this game. The reasons being the two defining properties of Devil's games. While there are plenty of results that show that combinatorial games are PSPACE-complete, this seems not to capture the second aspect of the Devil's game. And intuitively the second property makes Devil game's quite distinct from most other known combinatorial games. This intuition motivates us to study Devil's games more broadly. We use the term Devil's games, for two reasons. One is a reference to the old puzzle from above and the second is that they are devilishly difficult to analyze.

4.3 Geometric realization spaces of paths, trees and cycles

Arnaud de Mesmay (Gustave Eiffel University – Marne-la-Vallée, FR), Gargi Lather (Indian Institute of Techology Madras, IN), Anna Lubiw (University of Waterloo, CA), Marcus Schaefer (DePaul University – Chicago, US), and Alexandra Wesolek (TU Berlin, DE)

The topic of this working group was to study the realization spaces of simple graphs (paths, trees or cycles) when representing them geometrically in two and three dimensions. A main motivation for this was to complement the recent results of Lubiw and Schaefer that imply universality for the realization spaces of penny and marble graphs in general. More precisely, we have investigated the following problems.

Realizing trees as penny graphs. It follows from celebrated results on the Carpenter's rule problem that the realization space of a path as a penny graph, i.e., a contact graph of unit disks in the plane, is connected. In contrast, this realization space can become disconnected for trees. Known hardness proofs imply the existence of such disconnected examples, but we have obtained a simpler and arguably cleaner construction.

Realizing trees and paths as marble graphs. Moving up to three dimensions, it is easy to see using simple knots that the realization space of trees as marble graphs, i.e., contact graphs of unit balls in $\mathbb{R}^3$, can be disconnected. For paths, we have worked with some precise realizations (both physical and virtual) of overhand and fisherman's knots with chains of marbles and preliminary evidence suggests that they also yield disconnected realization spaces. This is ongoing work.

Spaces of polygonal knots. A third topic of investigation was the space of realizations of a cycle as a polygonal chain with a fixed number of segments of variable length in three dimensions. This topic naturally involves a knot-theoretical aspect, since different knot types necessarily lead to disconnected components in the realization space. Attempts to prove universality and $\exists\mathbb{R}$-hardness for polygonal realizations of fixed knot types raised new questions about intersection and linking graphs of triangles in three dimensions, which we are still looking at.

4.4 Online Packing of Convex Objects

Arindam Khan (Indian Institute of Science – Bangalore, IN), Anders Aamand (University of Copenhagen, DK), Mikkel Abrahamsen (University of Copenhagen, DK), Linda Kleist (Universität Hamburg, DE), Eunjin Oh (POSTECH – Pohang, KR), and Csaba Tóth (California State University – Northridge, US)

We study the packing of convex polygons in the online setting. Here, convex polygons (a total of n in number) arrive one by one, and need to be packed (immediately and irrevocably, using translation and without overlapping) in a horizontal unbounded strip of unit height. Our goal is to minimize the width of the strip to pack all polygons. We have some promising initial results that might lead to a $(\log n)^{O(1)}$-competitive algorithm. Our approach exploits connections with the online sorting problem, where n elements are revealed one by one and have to be placed in an immediate and irrevocable manner into empty cells of an array. The objective is to minimize the sum of absolute differences between elements in the consecutive cells. We also hope to extend the approach to obtain a $(\log n)^{O(d)}$-competitive algorithm for packing d-dimensional convex polytopes into a d-dimensional strip (with $(d-1)$-dimensional unit cube base and unbounded length in the d-th dimension).

4.5 Sparse $(1+\varepsilon)$-emulators for Euclidean point sets

Sándor Kisfaludi-Bak (Aalto University, FI), Sujoy Bhore (Indian Institute of Technology Bombay – Mumbai, IN), Karl Bringmann (Universität des Saarlandes – Saarbrücken, DE), Hung Le (University of Massachusetts Amherst, US), André Nusser (INRIA – Sophia Antipolis, FR), and Karol Wegrzycki (MPI für Informatik – Saarbrücken, DE)

Let P be a set of points in the Euclidean plane (or more generally, in $\mathbb{R}^d$. Consider an edge weighted graph $G=(P\cup S,E), w:E\to\mathbb{R}_{\geq 0}$ and let $\mathrm{dist}_{G,w}$ be the induced shortest-path distance. If for some >0 the distance $\mathrm{dist}_{G,w}$ satisfies for all $u,v\in P$ that

$$\|u-v\|_2 \leq \mathrm{dist}_{G,w}(u,v) \leq (1+\varepsilon)\|u-v\|_2,$$

then (G,w) is called a $(1+\varepsilon)$-*emulator* and the points of S are called *Steiner vertices*. Our goal in this problem has been to minimize the number of edges in the emulator.

The variant of the problem where $S \subset \mathbb{R}^d$ and the weight function is forced to be the Euclidean weight function is well-understood: the resulting graphs are called *Steiner spanners*, and the sparsest possible Steiner spanners in $\mathbb{R}^d$ are known to have $\Theta(n/\varepsilon^{\frac{d-1}{2}})$ edges up to logarithmic factors of $1/\varepsilon$.

During the seminar, we started working on the specific point set P that is used in the lower bound for Euclidean Steiner spanners: this corresponds to a bipartite setting. The bipartite construction can be used as a basis for general constructions, and the best Euclidean construction is simply the complete bipartite graph, so it should be possible to improve to find a sparser graph using emulators where we have more freedom to choose the edge weights. We have identified several structural properties of a potential edge-minimal solution based on the geometric constraints, and have found some examples of graphs that had fewer edges than the complete bipartite graph. Deciding if these constructions can be assigned the desired weights is a work in progress.

4.6 Flip Distance to a Crossing-Free Matching

Lucas Meijer (Utrecht University, NL), Thomas Bläsius (KIT – Karlsruher Institut für Technologie, DE), Sarita de Berg (Utrecht University, NL), Aye Chan May (Thammasat University – Pathum Thani, TH), Arturo Merino (O'Higgins University – Rancagua, CL), and Jack Stade (University of Copenhagen, DK)

At Dagstuhl, we worked on the "Flip Distance to a Geometric Crossing-Free Perfect Matching" problem, which we informally called "Uncrossing Line Segments". In this problem, you are given a set of 2n points in the plane and a perfect matching of them. You create a line segment between any two matched points, which gives a configuration of n line segments. Some of these line segments may intersect. We want to transform the instance into one without any intersections using the "flip" operator. Practically, we may "flip" two crossing line segments, and match the four points involved in these line segments in the other two possible ways, neither of which will cause these two lines to intersect. Prior to the seminar, the best known results were that there exists a configuration that always takes $\Omega(n)$ flips to uncross, all configurations have a sequence of $O(n^2)$ flips to uncross them, and an adversary can take at most $O(n^3)$ flips on any configuration. Additionally, unbeknownst to us, it was also known that if the points are in convex position, there is always a sequence of $O(n)$ flips to uncross it.

During the seminar, we did not manage to make much progress on the problem. Our main novel result is that if there are k points not on the convex hull of the pointset, then there exists a sequence of $O(nk)$ flips to uncross the configuration. Otherwise, we looked into many potential-based strategies to keep track of the 'progress' we make upon performing a specific flip; we looked into search trees; into the dual; into divide and conquer strategies; and many more subideas.

5 Open problems

5.1 Recognition of Strongly Hyperbolic Uniform Disk Graphs

Thomas Bläsius (KIT – Karlsruher Institut für Technologie, DE)

A *hyperbolic uniform disk graph (HUDG)* is the intersection graph of disks of equal radius r in the hyperbolic plane. Recognizing HUDGs is $\exists\mathbb{R}$ complete. However, the hardness is essentially based on the fact that Euclidean UDGs are a subclass of HUDGs. One way of looking at a subset of HUDGs that are very hyperbolic are *strongly hyperbolic UDGs*: here all vertices have to be placed into a disk of radius at most $2r$.

Question: Is recognizing SHUDGs in NP (or even in P)? It might be interesting to restrict SHUDGs further to only sparse graphs. An alternative way of restricting HUDGs to graphs that are rather hyperbolic is to require $r \in \Omega(\log n)$.

5.2 Polygonal representations of knots

Arnaud de Mesmay (Gustave Eiffel University – Marne-la-Vallée, FR)

A *polygonal knot*[1] is a closed polygonal curve in $\mathbb{R}^3$. Two polygonal knots made of k segments are *isotopic* if there is a continuous (non-necessarily polygonal) deformation from one to the other without crossings, and are *polygonally isotopic* if there is such a deformation that respects the polygonal structure, i.e., such that all the intermediate curves are also polygonal knots with k segments. The lengths of the segments do *not* have to be respected.

Question 1. Does there exist a k and two polygonal knots made of k segments which are isotopic to the unknot but not polygonally isotopic to each other?

This question was mentioned by Calvo [calvo] as an open problem and there seems to have been no progress since then. Note that if the length of the segments is fixed, such examples of *stuck unknots* are known [toussaint]. Also, in the same paper, Calvo showed that there are isotopic trefoil knots with 6 segments which are *not* polygonally isotopic.

One way to reformulate this question is to wonder whether the realization space of the unknot within the space of polygonal knots with k segments is connected. From that perspective, this reminds of the classical Ringel isotopy conjecture, which asked whether the realization spaces of arrangements of pseudolines are connected. This conjecture was shattered by the M'nev universality theorem showing that such realization spaces can be very pathological (formally, can be stably equivalent to any semi-algebraic set). This leads to:

Question 2. Are there universality phenomena in the polygonal realization spaces of knots?

Universality phenomena often have algorithmic/complexity implications in that the corresponding problems are often $\exists\mathbb{R}$-hard. The *stick number* of a knot K is the minimum number of segments to make a polygonal knot isotopic to K.

[1] This terminology is nonstandard, don't google it.

Question 3. Is the stick number ∃ℝ-hard to compute?

This is also open for the equilateral stick number (where the lengths of the segments are required to be equal). Actually, it is open whether the stick number equals the equilateral stick number in general [cantarella].

Knot theory is hard and scary but polygonal knots are so different from the standard ones that I think that these problems do not require prior knowledge in knot theory. For all these problems, a natural angle of attack could be to look first at polygonal *links* (disjoint unions of knots) and even graphs embedded polygonally in $\mathbb{R}^3$.

5.3 Precision of continuous distance problems

Sándor Kisfaludi-Bak (Aalto University, FI)

Consider the following distance problems:

Pr1 Given a set of open convex pairwise disjoint obstacles in $\mathbb{R}^3$ and $a, b \in \mathbb{R}^3$, find the shortest path connecting a and b that avoids the obstacles.

Pr2 Given a polygonal surface (e.g., boundary of a 3-dimensional convex polyhedron, but could be just an abstract surface), find its diameter.

Pr3 Given a set of (integer) points in $\mathbb{R}^3$, find the shortest tree connecting them (Euclidean Steiner tree).

Pr1 and Pr3 are known to be NP-hard, and I suspect that Problem 2 might be NP-hard for high-genus surfaces. Pr2 has an exact solution for the convex polyhedron case, but it requires an oracle for solving low-degree equations (note that the diameter can be realized by a pair of points that are in the interior of their respective faces). I am not aware of exact solutions to 1 and 3, but they have approximation schemes with $2^{\text{poly}(n)} \cdot \text{polylog}(1/\varepsilon)$ time via brute force structure guess + second-order cone programming (SOCP). Moreover, Pr1 and Pr2 have an FPTAS [Har-Peled99] and Pr3 has an EPTAS.

Question 1. Is any of these problems ∃ℝ-complete? If not, can we reduce some of them to each other or reduce to them from a common third problem?

There are some problems such as geometric median (given a set of points in the plane, find the point q that minimizes the sum of distances to the given points) that can be solved in $\text{poly}(n, \log(1/\varepsilon))$ time [CohenLMPS16]. They do not give off an NP-hard vibe, and they should probably fall within some class of "polynomial-reals", some polynomial-time analogue of ∃ℝ.

Question 2. Is there any relationship between approximability and ∃ℝ-completeness? Can we at least provide justification that geometric median is not ∃ℝ-complete?

Question 3. Is there any reasonable notion of an exact algorithm that can solve Pr1 and Pr3? What kind of oracle access should be granted?

Question 4. Can we prove (conditional) lower bounds to rule out an FPTAS for Pr3 and a $\text{poly}(n, \log(1/\varepsilon))$ algorithm for Pr1 and Pr2?

5.4 Weak circle representations of planar graphs

Günter Rote (FU Berlin, DE)

Given a plane graph G, how fast can we find a set of disks whose intersection graph is G? By the celebrated Koebe–Andreyev–Thurston Theorem, every planar graph G has a *disk-packing* representation as a touching graph of nonoverlapping disks, see [Felsner and Rote 2019] for a relatively easy proof through a converging infinite algorithm. In many applications of this theorem, it is not necessary to require adjacent disks to touch; overlapping disks are also fine, see for example [Felsner 2016]. I call this a *weak circle representation.*

Background. If G is triangulated, the disk-packing representation is essentially unique, i. e., unique up to Möbius transformations (circle-preserving transformations). It is not hard to see that for some instances, the ratio between the largest and the smallest disk is necessarily exponential, and likewise, the algebraic degree of the solution coordinates and radii can be exponential in the size of G.

Mohar [1997, 2000] has given polynomial-time *approximation* algorithms for disk-packing representation. For a graph embedded on a surface of constant negative curvature, like a Klein bottle, the algorithm computes ε-approximations of the centers and radii of a true circle-packing representation [Mohar 2000, Theorem 5.5]. In the plane, there is an algorithm that computes a vector $r = (r_1, \dots, r_n)$ of radii for the n disks such that the resulting maximum *angle defect* $\mu(r)$ at the centers is smaller than some given bound ε [Mohar 2000, Algorithm A]. In 2019, Dong, Lee and Quanrud gave an improved algorithm to compute an ε-approximation of the centers and radii of a true circle packing in the plane in $O(n \log \frac{U}{\varepsilon})$ time where $U = 2^{O(n)}$ is the ratio between the largest and the smallest disk, and the true circle packing is normalized so that the largest radius is 1.

It remains to work out bounds on ε that guarantee the desired drawing can be obtained by slightly blowing up the radii. Perhaps there is also a more direct and faster approach.

As a stronger variation, we may require that no three disks intersect, or that no disk center is covered by another disk (to remain closer in spirit to a packing representation).

5.5 Fixed points of compositions of monotone polynomials

Jack Stade (University of Copenhagen, DK)

Let f and g be polynomials of degree d each with positive derivative on the whole real line. How many fixed points can the composition $f \circ g$ have?

Background

Miltzow and Schmiermann [MiltzowSchmiermann2022] have asked about the complexity of continuous constraint satisfaction problems when each constraint involves at most 2 variables. In the discrete setting, many-valued 2SAT is polynomial time solvable when the constraints are *monotone* (see [BeckertHanleManya2000]). For $x, y \in \mathbb{Z}$, a constraint $C(x, y)$ is one that can be written as a conjunction of constraints of form $\pm x \geq x_0 \vee \pm y \geq y_0$.

In the continuous setting, it seems natural to study the monotone case since it isn't obviously NP-hard. However, we can compose polynomials, and the degree of the composition grows exponentially with the number of polynomials. What is less clear is whether the combinatorial complexity grows exponentially: if we have two compositions $f_1 \circ \cdots \circ f_k$ and $g_1 \circ \cdots \circ g_j$ of monotone polynomials, can their graphs intersect exponentially many times? If so, this could provide a route towards showing that monotone continuous 2SAT is NP-hard.

The problem I pose above is essentially the simplest non-trivial example of this problem. Going by the degree of the polynomials, it seems that $f \circ g$ could have as many as d^2 fixed points. But there are only $2(d + 1)$ coefficients total, so there aren't enough degrees of freedom to easily construct examples with more than $\mathcal{O}(d)$ fixed points. When $d = 3$, I think I can prove that $f \circ g$ has at most 5 fixed points.

Participants

- Anders Aamand
University of Copenhagen, DK
- Mikkel Abrahamsen
University of Copenhagen, DK
- Peyman Afshani
Aarhus University, DK
- Sujoy Bhore
Indian Institute of Technology Bombay – Mumbai, IN
- Thomas Bläsius
KIT – Karlsruher Institut für Technologie, DE
- Karl Bringmann
Universität des Saarlandes – Saarbrücken, DE
- Mark de Berg
TU Eindhoven, NL
- Sarita de Berg
Utrecht University, NL
- Arnaud de Mesmay
Gustave Eiffel University – Marne-la-Vallée, FR
- Omrit Filtser
The Open University of Israel – Ra'anana, IL
- Dan Halperin
Tel Aviv University, IL
- Arindam Khan
Indian Institute of Science – Bangalore, IN
- Sándor Kisfaludi-Bak
Aalto University, FI
- Linda Kleist
Universität Hamburg, DE
- Gargi Lather
Indian Institute of Techology Madras, IN
- Hung Le
University of Massachusetts Amherst, US
- Anna Lubiw
University of Waterloo, CA
- Aye Chan May
Thammasat University – Pathum Thani, TH
- Lucas Meijer
Utrecht University, NL
- Arturo Merino
O'Higgins University – Rancagua, CL
- Till Miltzow
Utrecht University, NL
- André Nusser
INRIA – Sophia Antipolis, FR
- Eunjin Oh
POSTECH – Pohang, KR
- Günter Rote
FU Berlin, DE
- Marcus Schaefer
DePaul University – Chicago, US
- Jack Stade
University of Copenhagen, DK
- Csaba Tóth
California State University – Northridge, US
- Geert van Wordragen
Aalto University, FI
- Karol Wegrzycki
MPI für Informatik – Saarbrücken, DE
- Alexandra Wesolek
TU Berlin, DE