ISSN 2186-7437

NII Shonan Meeting Report

No. 234

Path Planning and Routing in Geometric Environments

Sang Won Bae
Yoshio Okamoto
Haitao Wang

June 30–July 4, 2025

National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-Ku, Tokyo, Japan

Path Planning and Routing in Geometric Environments

Organizers:
Sang Won Bae (Kyonggi University, Republic of Korea)
Yoshio Okamoto (The University of Electro-Communications, Japan)
Haitao Wang (University of Utah, U.S.A.)

June 30–July 4, 2025

Background and Introduction

Path planning and routing in computational geometry are essential for optimizing a wide array of real-world applications, significantly impacting industries such as logistics, robotics, telecommunications, and urban planning. These problems often involve devising efficient routes or strategies to achieve specific goals, such as minimizing travel time, reducing costs, or ensuring safety. In logistics, advanced routing algorithms enhance the efficiency of delivery services by optimizing routes, thereby cutting operational expenses and improving customer satisfaction. In robotics, precise path planning enables autonomous systems to navigate complex environments safely and efficiently. Telecommunications benefit from optimized routing by improving data transmission across networks, leading to faster and more reliable communication. Urban planners use these computational techniques to design better traffic flow systems, reducing congestion and improving public transportation. Overall, advancements in path planning and routing drive technological innovation and contribute to societal progress by solving complex real-world challenges.

Emerging applications of path planning and routing in computational geometry are revolutionizing various cutting-edge fields. In autonomous driving, sophisticated path planning algorithms enable self-driving cars to navigate dynamically changing environments with precision and safety. In drone technology, efficient routing ensures optimal flight paths for tasks ranging from delivery services to agricultural monitoring and disaster response. The healthcareindustry is leveraging these advancements for robotic surgery, where precise path planning allows minimally invasive procedures with increased accuracy. In environmental monitoring, autonomous underwater vehicles (AUVs) use advanced routing to conduct extensive marine surveys, aiding in oceanographic research and resource management. Additionally, smart city initiatives utilize computational geometry for optimizing public transportation routes, reducing traffic congestion, and enhancing urban mobility. These applications demonstrate the expanding influence of path planning and routing in addressing complex, real-world challenges across diverse domains.

Recently, significant theoretical advancements have been made in the field

of computational geometry to address certain path planning and routing problems. For instance, Wang developed an algorithm to compute Euclidean shortest paths in triangulated polygonal domains with optimal running time [7]. The longstanding two-point shortest path query data structures in polygonal domains, established more than two decades ago by Chiang and Mitchell [4], have been improved by S. de Berg, Miltzow, and Staals [6]. Moreover, the traveling salesman problem in d-dimensional Euclidean space can now be solved exactly in time optimal up to the exponential-time hypothesis, thanks to the work of M. de Berg, Bodlaender, Kisfaludi-Bak, and Kolay [5]. These breakthroughs highlight the ongoing progress and potential of computational geometry in addressing complex path routing and planning problems.

This seminar aims to unite emerging and established experts in the field to share recent developments and identify future research directions. On the first day, overview lectures will cover significant recent developments, algorithmic methodologies, and emerging applications. The second and third days will focus on group discussions to pinpoint future research directions. On the fourth day, participants will be regrouped to exchange insights from the previous days and concentrate on more specific research questions. The final day will be dedicated to a wrapup discussion, consolidating ideas to foster future research collaborations.

Overview of the Meeting

The meeting began on Monday with an introduction of all participants. This was followed by three tutorial talks.

- Evanthia Papadopoulou (Universita della Svizzera italiana) presented on "Abstract Voronoi-Like Diagrams with Applications."
- Eunjin Oh (POSTECH) spoke about "Path Planning and Routing in Geometric Intersection Graphs."
- Csaba Tóth (California State University Northridge) gave a talk on "Towards Instance-Optimal Euclidean Spanners."

The day concluded with an open problem session, where participants proposed and discussed new research directions related to the meeting's topics.

Tuesday's session started with a research talk by Sujoy Bhore on "Metric Sketching: Towards Ideal Trades Among Parameters." Afterward, participants voluntarily formed several working groups to discuss the problems. The group discussions continued for the rest of the day.

On Wednesday, the morning began with a research talk by Maarten Löffler entitled "A Curve with Rotation Number One That Is Not Universal for Beacon Routing." The group discussions on open problems continued throughout the morning. In the afternoon, the group went on an excursion to Kamakura.

Thursday was dedicated to further group discussions. At the end of the day, all groups reconvened to share and disseminate the ideas that had been generated.

The meeting concluded on Friday with a short final round of group discussions, followed by a wrap-up session to summarize the week's outcomes.

Overview of Talks

Abstract Voronoi-Like Diagrams with Applications

Evanthia Papadopoulou, Universita della Svizzera italiana (USI)

I will give an overview of abstract Voronoi diagrams (AVDs) and define a more relaxed Voronoi-like version. AVDs offer a unifying framework to various classical concrete Voronoi diagrams involving point sites in various metrics, but also generalized sites including segments, disks, polygons, or clusters of points. Surprisingly, some basic problems regarding the latter diagrams have only recently been addressed, and some still remain unresolved. Voronoi-like graphs can help ease these differences. For example, we can extend Delaunay's theorem from points in the Euclidean plane to any Voronoi diagram under the abstract framework; we can obtain simple expected linear-time algorithms to compute Voronoi trees and forests such as those occurring when updating a Voronoi diagram after site deletion, and others, which until recently were known only for points and restricted AVDs. I will further review a recent variant of backwards analysis applicable to order-dependent randomized incremental constructions that was motivated by these problems.

Path Planning and Routing in Geometric Intersection Graphs

Eunjin Oh, POSTECH

Geometric intersection graphs provide a natural framework for modeling wireless communication and motion planning problems. In this talk, I will present recent algorithmic advances in computing shortest paths and distance oracles in geometric intersection graphs, particularly focusing on unit disk graphs and general disk graphs. For the single-source shortest path problem in weighted unit disk graphs and disk graphs, I will introduce a framework combining geometric data structures with graph algorithms. I will also discuss distance oracles for unweighted unit disk graphs.

Towards Instance-Optimal Euclidean Spanners

Csaba Tóth, California State University Northridge

Spanners were introduced for an edge-weighted graph G to approximate pairwise distances between all pairs of vertices up to a multiplicative distortion (stretch) of at most $t \geq 1$. In general, subquadratic t-spanners are possible for $t \geq 3$. Metric spanners approximate pairwise distances in a finite metric space: The graph G is the complete graph on the points of the metric space and the edge weights are the metric distances between the endpoints. For many important metrics, such as constant-dimensional Euclidean spaces or planar graph metrics, linear-size spanners can achieve stretch $t = 1 + \varepsilon$ for any $\varepsilon > 0$.

The two most basic "compactness" measures of a metric spanner $H = (V, E)$ are the size (number of edges) $|E|$ and the weight (sum of edge weights) $\|E\|$. The state-of-the-art constructions of Euclidean $(1 + \varepsilon)$-spanners in $\mathbb{R}^d$ have $O_d\left(n \cdot \varepsilon^{-d+1}\right)$ edges (or *sparsity* $O_d(\varepsilon^{-d+1})$) and weight $O_d\left(\varepsilon^{-d} \log \varepsilon^{-1}\right) \cdot$

$w(E_{\mathrm{MST}})$ (or *lightness* $O_d(\varepsilon^{-d}\log\varepsilon^{-1})$); here O_d suppresses a factor of $d^{O(d)}$ and $w(E_{\mathrm{MST}})$ denotes the weight of a minimum spanning tree of the input point set. These two upper bounds are (near-)optimal (up to the $d^{O(d)}$ factor and disregarding the factor of $\log(\varepsilon^{-1})$ in the lightness bound) for some *extremal instances*, and therefore they are (near-)optimal in the worst case.

Recent research focused on beating the worst-case bounds, that is, optimization problems to find instance-optimal spanners: given n metric points and a parameter $\varepsilon > 0$, find a $(1+\varepsilon)$-spanner of minimum size or weight. This talk will survey some recent results and open problems.

Metric Sketching: Towards Ideal Trades Among Parameters

Sujoy Bhore, IIT Bombay

In many applications — including network design, machine learning, robotics, and distributed systems — the input data can be modeled as points in a finite metric space and is often massive in scale. Efficient processing of such data requires compact representations that preserve essential structural properties. Sketching compresses large datasets into smaller summaries that approximately retain key metric properties. Among the most fundamental sketching structures are spanners and tree covers. This talk will present recent advances in metric sketching, focusing on optimal trade-offs among parameters such as size, weight, degree, and diameter.

A Curve with Rotation Number One That Is Not Universal for Beacon Routing

Maarten Löffler, Utrecht University

Beacon routing is the study of how to indirectly route objects through various domains. Consider two points a —which we think of as the attractor— and b —which we think of as the ball— lying on a smooth closed curve in the plane. In curve-restricted beacon routing, b moves along the curve as long as its Euclidean straight-line distance to a decreases, until this distance is locally minimal. Assuming b moves infinitely faster than a, the goal is to move a along the curve in such a way that a ends up meeting b. We say that a curve is universal if there always exists a strategy to catch the ball from every initial configuration of the attractor and the ball. Recent work of Abrahamsen et al. has shown that every simple curve is universal. The authors also conjectured that all curves with rotation number one are universal. In this note, we disprove their conjecture and present a curve with rotation number one that is not universal.

List of Participants

- Hee-Kap Ahn (POSTECH)
- Taehoon Ahn (National Institute for Mathematical Sciences)
- Sang Won Bae (Kyonggi University)
- Sujoy Bhore (Indian Institute of Technology (IIT) Bombay)
- Jie Gao (Rutgers University)
- Michael Hoffmann (ETH Zurich)
- Tonan Kamata (Japan Advanced Institute of Science and Technology)
- Sándor Kisfaludi-Bak (Aalto University)
- Chih-Hung Liu (National Taiwan University)
- Maarten Löffler (Utrecht University)
- Anurag Naredla (University of Bonn)
- Yakov Nekrich (Michigan Technological University)
- Eunjin Oh (POSTECH)
- Yoshio Okamoto (The University of Electro-Communications)
- Evanthia Papadopoulou (Universita della Svizzera italiana (USI))
- Thomas Schibler (University of California Santa Barbara)
- Frank Staals (Utrecht University)
- Frederick Stock (University of Massachusetts Lowell)
- Csaba Tóth (California State University Northridge)
- André van Renssen (University of Sydney)
- Antoine Vigneron (Ulsan National Institute of Science and Technology)
- Haitao Wang (University of Utah)
- Jie Xue (New York University Shanghai)
- Sang Duk Yoon (Sungshin Women's University)

Meeting Schedule

Check-in Day: June 29 (Sun)

- 19:00–21:00 Welcome Reception

Day 1: June 30 (Mon)

- 9:00–9:40 Introduction
- 9:40–10:30 Tutorial 1: Evanthia Papadopoulou
- 10:30–11:00 Coffee Break
- 11:00–11:50 Tutorial 2: Eunjin Oh
- 12:00–13:30 Lunch
- 14:00–14:50 Tutorial 3: Csaba Tóth
- 14:50–15:20 Coffee Break
- 15:20–17:00 Open Problem Session
- 18:00– Dinner

Day 2: July 1 (Tue)

- 9:00–9:30 Talk 1: Sujoy Bhore
- 9:30–10:30 Problem Solving
- 10:30–11:00 Coffee Break
- 11:00–11:50 Problem Solving
- 12:00–13:30 Lunch
- 14:00–15:00 Problem Solving
- 15:00–15:30 Coffee Break
- 15:30–17:00 Problem Solving
- 18:00– Dinner

Day 3: July 2 (Wed)

- 9:00–9:30 Talk 2: Maarten Löffler
- 9:30–10:30 Problem Solving
- 10:30–11:00 Coffee Break
- 11:00–11:50 Problem Solving
- 12:00–13:15 Lunch
- 13:15–13:30 Group Photo

- 13:30–21:00 Excursion (Kamakura) and Dinner (Minemoto Asahina)

Day 4: July 3 (Thu)

- 9:00–10:30 Problem Solving
- 10:30–11:00 Coffee Break
- 11:00–11:50 Problem Solving
- 12:00–13:30 Lunch
- 14:00–15:00 Problem Solving
- 15:00–15:30 Coffee Break
- 15:30–16:00 Dissemintation of Ideas
- 16:00–17:00 Problem Solving
- 18:00– Dinner

Day 5: July 4 (Fri)

- 9:00–10:30 Problem Solving
- 10:30–11:00 Coffee Break
- 11:00–11:50 Wrap-Up
- 12:00–13:30 Lunch

Summary of Discussions

Path TSP of Line Neighborhoods

We studied the following problem: given a set of lines in the Euclidean plane, we want to find the shortest open curve intersecting all the lines. This is a variant of the Euclidean traveling salesperson (Euclidean TSP) problem, where the salesperson is tasked with visiting lines instead of points.

The problem is also equivalent to the so-called watchman route problem. These problems have polynomial-time algorithms in case the optimum is a closed curve (see [3]), however the open curve variant have eluded all attempts so far.

During the workshop, we first recalled how the closed curve variant is solved, and the combinatorial and geometric properties that surprisingly allow for a polynomial-time solution (while most variants of TSP are NP-hard). Then we studied the structural properties of the optimum solution: we observed that its vertices are in convex position, but unlike the closed curve variant, the optimum tour can switch between two convex chains several times. Despite this possibility, we struggled to construct an instance that requires more than a constant number of such zig-zags, and conjectured that the optimum solution never does more than constantly many swicthes between the two convex chains. Indeed, if this conjecture is true, then using other techniques we would immediately obtain a polynomial-time algorithm.

The rest of our discussions was spent on trying to model the hard cases more accurately, to get a better grasp on optima.

The TSP with neighborhoods is a fundamental problem, and generally understood to be even harder than TSP with points; its variants are all NP-hard with very few exceptions. Good approximations are only known for special objects. However, this curious case of line neighborhoods likely remains polynomial-time solvable; we hope to make further progress towards proving this result.

The Point Enclosure Problem

The point enclosure problem is a geometric separation question where the input is a set of points P and a set of geometric objects (obstacles) O in the plane. The question is then to find a subset of obstacles $o \subseteq O$ so for any point p of P, every path from p to a point at infinity intersects at least one object of o. This problem has a number of variants, delineated by what class of geometric object the obstacles are drawn from. In this vein, we studied this problem with many classes of objects.

There is an $O(1)$-approximation algorithm when the obstacles are unit disks. However it is unknown if this variant is indeed NP-hard or if a polynomial time algorithm is possible. We were unable to provide conclusive evidence either way for this question. But the discussion lead to two related results. We proved the 3D case is NP-hard. That is, when points are from $\mathcal{R}^3$ and the obstacles are unit balls point enclosure is NP-hard. We also considered an alternative 2D problem, where P consists red and blue points, and the goal is to select a minimal set of obstacles that separates all of the red points from the blue. Note, the original instance of point enclosure is a special case of this, where there is only one red or blue point. In the broader two color version, we proved that

finding this set of obstacles is hard for unit disks.

Additionally, this problem is known to be NP-hard when the obstacles are line segments. But the case when obstacles are lines in the plane is unstudied. On the final day of the meeting we started investigating this case, establishing a few combinatorial bounds on the size of a solution to this instance. Giving hope of finding a polynomial time algorithm for this variant.

Visibility Polygon Queries in a Simple Polygon

Let P be a simple polygon with n vertices. A point $p \in P$ is visible to point $q \in P$ if and only if the line segment $\overline{pq}$ is fully contained in P. The visibility polygon $V_P(q)$ of q (with respect to P) is the set of all points in P visible to q. It is well known that $V_P(q)$ is a star-shaped polygon whose vertices are either (i) vertices of P, or (ii) intersection points of an edge of P with a ray from q through a reflex vertex v of P. The *combinatorial representation* of $V_P(q)$ is a list of these vertices in CCW order around q. Our goal is to store P so that given a query point $q \in P$ we can efficiently report the (combinatorial representation of the) visibility polygon of q.

Bose et al. provide an $O(n^3)$ space data structure that can answer queries in optimal $O(\log n + k)$ time, where $k = |V_P(q)|$ is the output complexity [2]. Aronov et al. provide an $O(n^2)$ size data structure that answers queries in $O(\log^2 n + k)$ time [1].

We achieve the following results:

- There is an $O(n \operatorname{polylog} n)$ space data structure that can answer visibility polygon queries in $O(\sqrt{n} \log n + k)$ time.
- Any data structure providing $O(\log n + k)$ query time uses at least $\Omega(n^2)$ space, by reduction from 2D fractional cascading.
- An $O(n^{2+\varepsilon})$ space data structure that achieves optimal $O(\log n + k)$. Here, $\varepsilon > 0$ is an arbitrarily small constant.
- An $O(n^2)$ space structure that achieves $O(\log n \log \log n)$ query time.
- An $O(n \log n)$ size data structure that can answer queries in $O(\log^2 n + k)$ time, provided that the query point lies on the boundary ∂P of P.
- For the boundary case, we can also obtain $O(n^{1+\varepsilon})$ space and $O(\log n + k)$ query time.

Direction Sensitivity

Problem Background

Let (T, G) be an infinite *periodic* graph, where T is a *tile* that can periodically tile the plane, and G is a partial embedded graph on T where edges on opposite sides match up.

For a given (finite or infinite) graph G and two points p,q on G, the *detour* $\delta(p, q)$ is the length of the shortest path from p to q in G divided by their Euclidean distance.

Now, for a given infinite graph G, and a given direction $\vec{d}$, we put the detour of every pair of points with direction $\vec{d}$ in a point cloud, with the Euclidean distance on one axis and the detour on the other axis.

Observation 1. *As the Euclidean distance tends to infinity, the detour approaches a single value that is representative for that direction.*

Let this value be $\ell(G, \vec{d})$, the *limit detour* for a given graph and fixed direction.

We can observe the following behaviour in practice:

Observation 2. *When we plot the limit detours as a function of their direction, we get a continuous function from S^1 to $\mathbb{R}$.*

The flatter this function is, the less sensitive to the direction is the detour of G.

One way to capture the direction sensitivity of a graph as a single number is to simply take the difference (or, alternatively, ratio) between its highest and lowest directional detours.

Problem Questions

We consider the following questions.

Question 1. *Can we prove Observations 1 and 2?*

Question 2. *Is there an algorithm to compute $\sigma(G)$ for a fixed periodic graph G?*

Question 3. *What values can $\sigma(G)$ have, over all periodic graphs G?*

Question 4. *What values can $\sigma(G)$ have, over all periodic graphs G with n vertices?*

Progress

This problem was explored from several directions during the week.

The first direction that was explored is the design of arbitrarily complex tiles with low direction sensitivity. Here, a first observation is that it is not too difficult to design tiles which have a detour arbitrarily close to zero in all directions. For interestingly, for any positive value x, it is possible to design a tile that has a detour arbitrarily close to x in all directions. Even more surprisingly, this is even possible with a construction that only has convex faces. This settles the question of how low the direction sensitivity can theoretically be, even for restricted geometry, when the complexity of G is allowed to be arbitrary large.

The second direction that was explored is to study tilings with bounded complexity. For periodic graphs with very small complexity, we have started to understand the smallest and largest possible detours over all directions. For the "plus" tile, the minimum detour is 1, and the maximum detour is $\sqrt{2}$. For the "diamond", we have the same values, but with reversed directions. For the "mixed" tile, the minimum detour is larger: it is $\frac{1+\sqrt{2}}{2}$. Based on these early observations, the group formulated a conjecture, that the minimum detour of the mixed tile is a weighted average of the minimum detour of the two simple tiles, where the weights relate to the size of the diamond.

The third direction that was explored is to consider the computational question: given a graph, can we compute its direction sensitivity? Here, the first result is that the limit detour in the horizontal direction can be computed in $O(n^4)$ time for a graph G of complexity n. This restricted case is already somewhat non-trivial. Next, this was extended to all direction with slope a/b for constant integers a,b. For irrational slopes, there is hope that the algorithm may be extensible by alternating between rational slopes, but it is still unclear how big the ratio needs to be to give any guarantees on the runtime, and the algebraic complexity involved may prohibit exact answers.

References

[1] B. Aronov, L. J. Guibas, M. Teichmann, and L. Zhang, *Visibility queries and maintenance in simple polygons*, Discret. Comput. Geom., 27 (2002), pp. 461–483.

[2] P. Bose, A. Lubiw, and J. I. Munro, *Efficient visibility queries in simple polygons*, Comput. Geom., 23 (2002), pp. 313–335.

[3] S. Carlsson, H. Jonsson, and B. J. Nilsson, *Finding the shortest watchman route in a simple polygon*, Discret. Comput. Geom., 22 (1999), pp. 377–402.

[4] Y. Chiang and J. S. B. Mitchell, *Two-point euclidean shortest path queries in the plane*, in Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms, 17-19 January 1999, Baltimore, Maryland, USA, R. E. Tarjan and T. J. Warnow, eds., ACM/SIAM, 1999, pp. 215–224.

[5] M. de Berg, H. L. Bodlaender, S. Kisfaludi-Bak, and S. Kolay, *An eth-tight exact algorithm for euclidean TSP*, SIAM J. Comput., 52 (2023), pp. 740–760.

[6] S. de Berg, T. Miltzow, and F. Staals, *Towards space efficient two-point shortest path queries in a polygonal domain*, in 40th International Symposium on Computational Geometry, SoCG 2024, June 11-14, 2024, Athens, Greece, W. Mulzer and J. M. Phillips, eds., vol. 293 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024, pp. 17:1–17:16.

[7] H. Wang, *A new algorithm for Euclidean shortest paths in the plane*, J. ACM, 70 (2023), pp. 11:1–11:62.