The Piano Movers @ 40: Geometry for Robotics

Dan Halperin
School of CS and AI
Tel Aviv University

SharirFest75, נחשולים , June 2025

Riddle

Given two Roomba's, each has to move from given start to goal positions, in a room without obstacles. What are the joint shortest paths (minimum total length)?

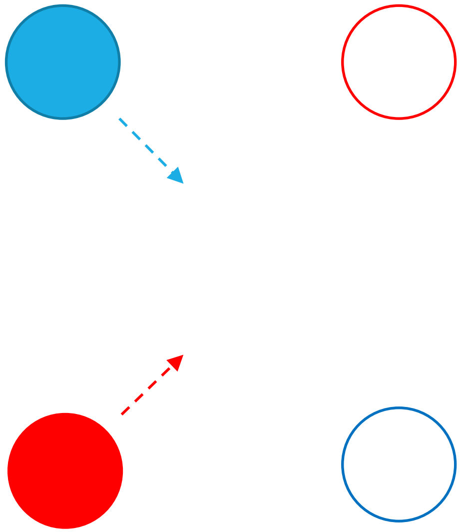

Talk overview

- The Piano Movers
- My favorite Micha paper
- CGAL arrangements and Micha
- CG and Robotics: a tale of two cultures
- Sampling based planners: a parallel universe
- Optimizing fleet motion
- Final notes

The Piano Movers by Schwartz and Sharir

1983-1986

Motion planning: the basic problem

Let B be a system (the robot/s) with k degrees of freedom moving in a known environment cluttered with obstacles. Given free start and goal placements for B decide whether there is a collision free motion for B from start to goal and if so plan such a motion.

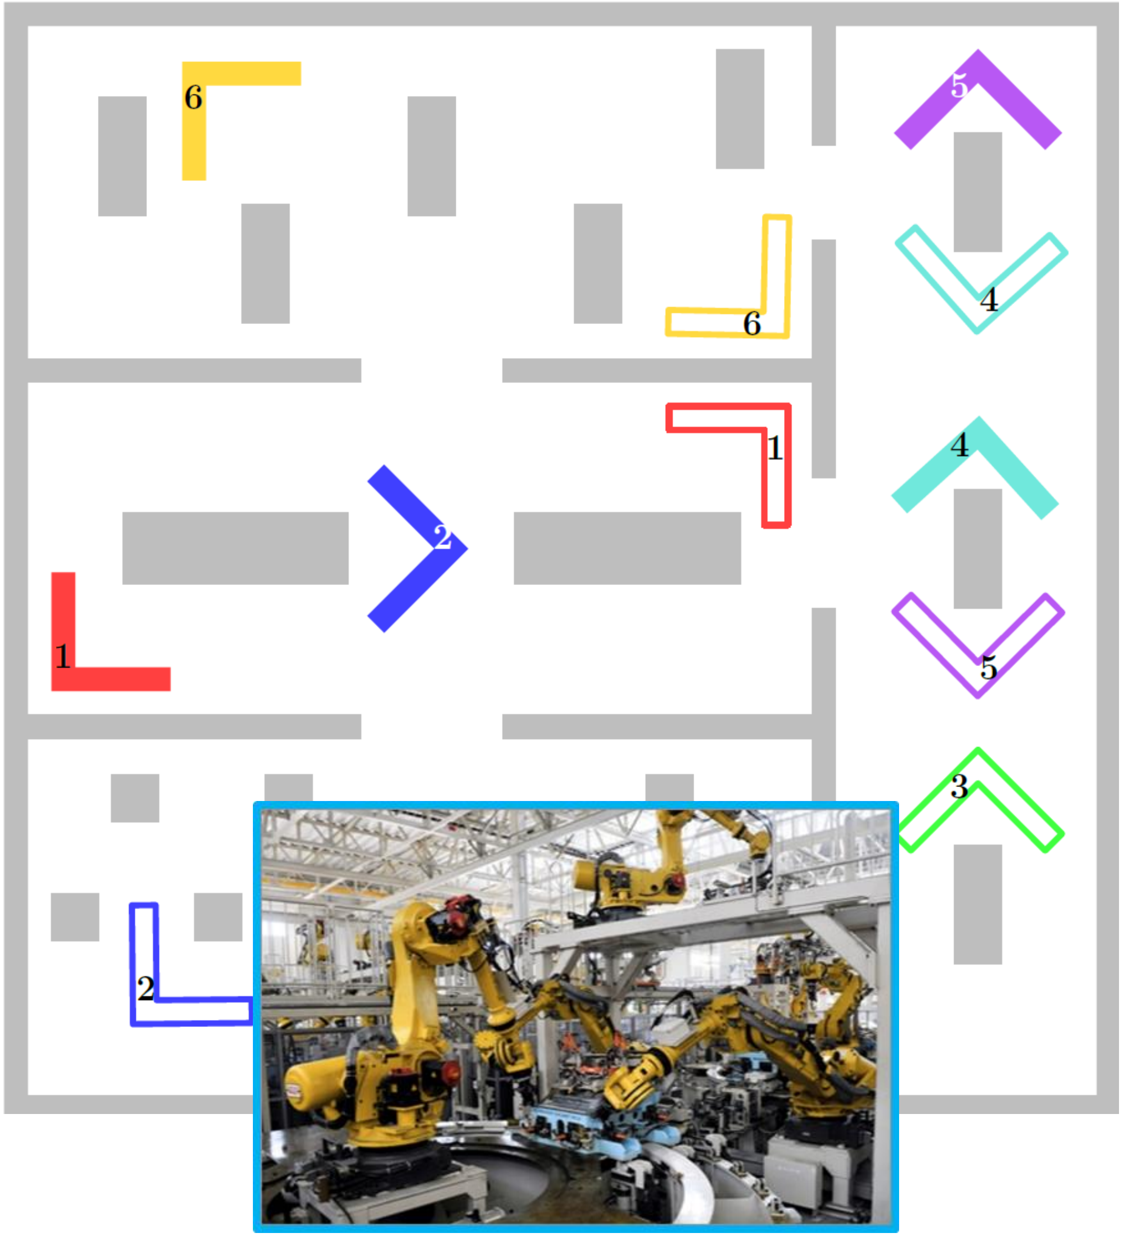

Two key terms:
(i) degrees of freedom (dof), and
(ii) configuration space

Configuration space

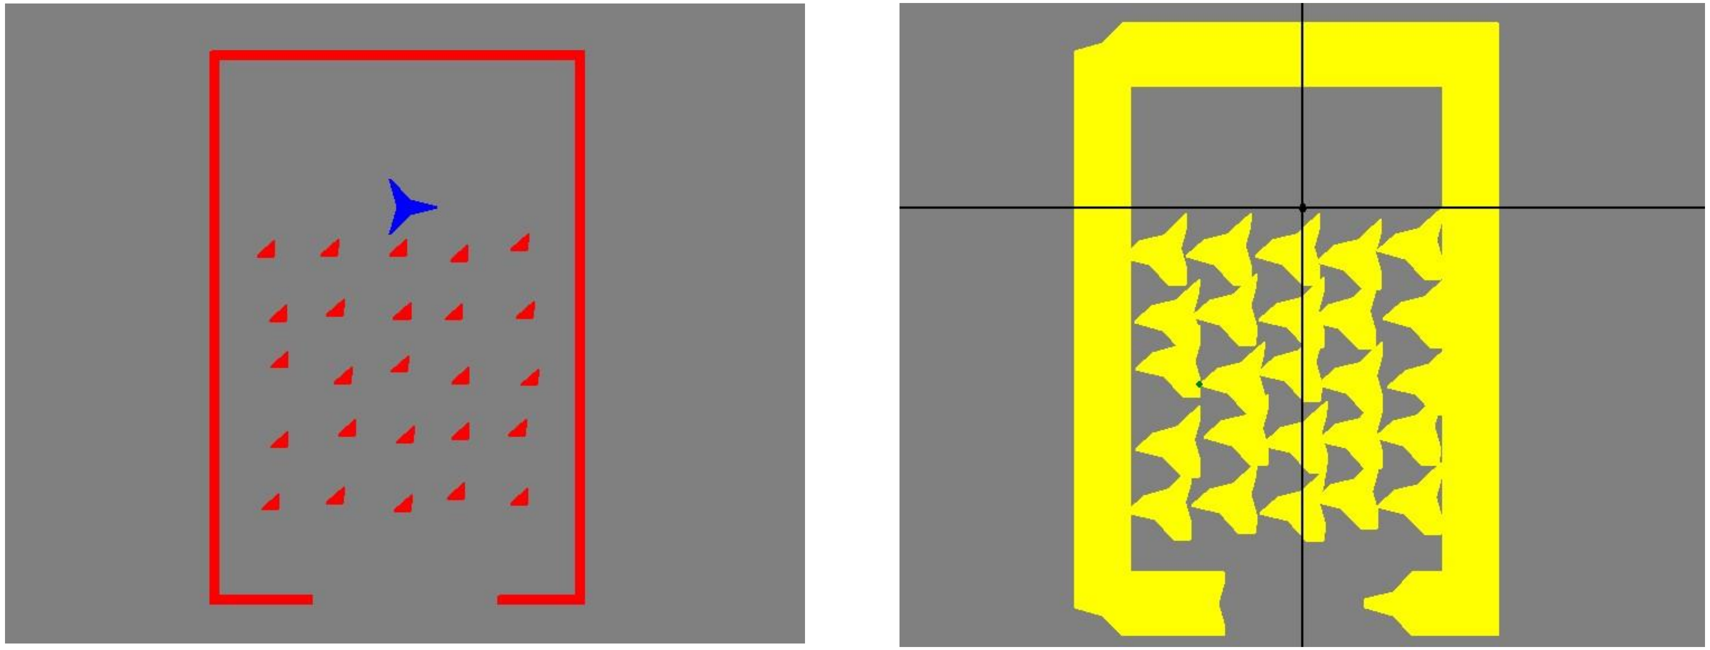

[Lozano-Perez, late 70s]

Complete solutions I

- the problem is hard when the number of degrees of freedom (# dofs) is part of the input [Reif 79], [Hopcroft-Schwartz-Sharir 84], …
- cell decomposition the Piano movers series [Schwartz-Sharir 83]: a doubly-exponential solution

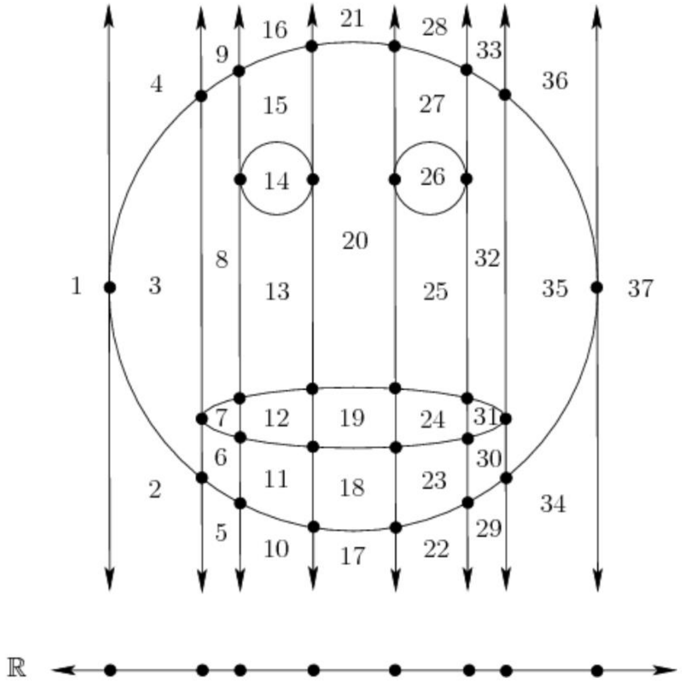

[LaValle]

To be continued soon

Piano Movers I:
A rod translating and rotating amidst polygonal obstacles in the plane

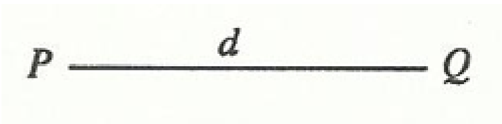

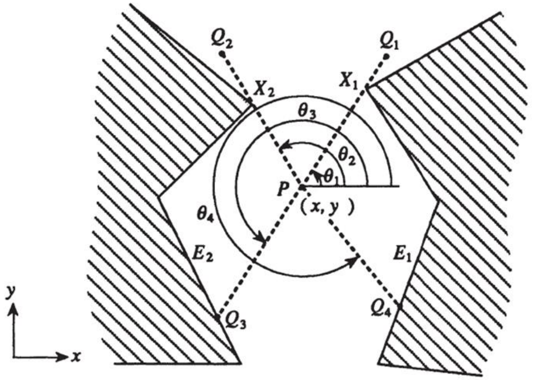

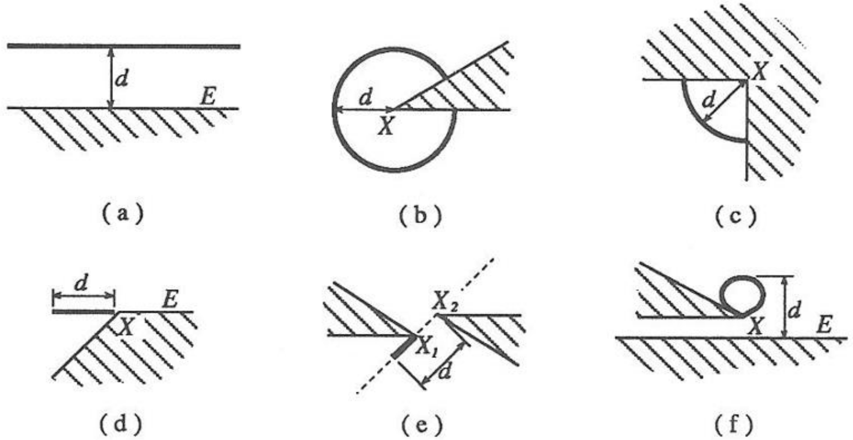

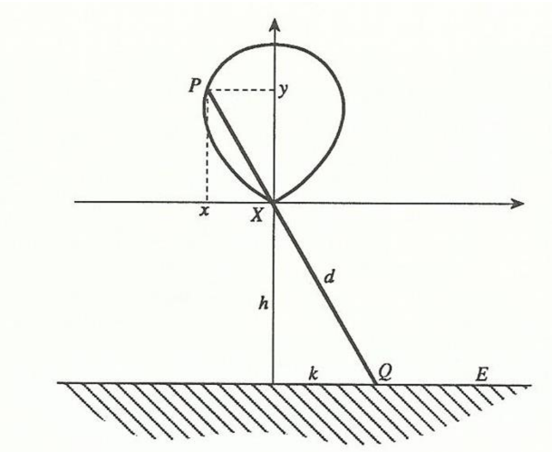

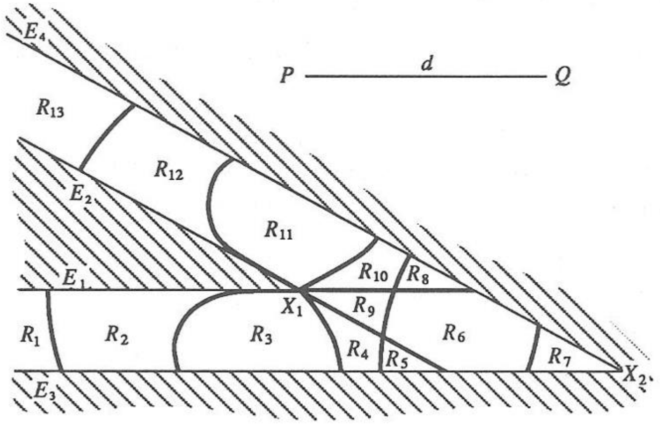

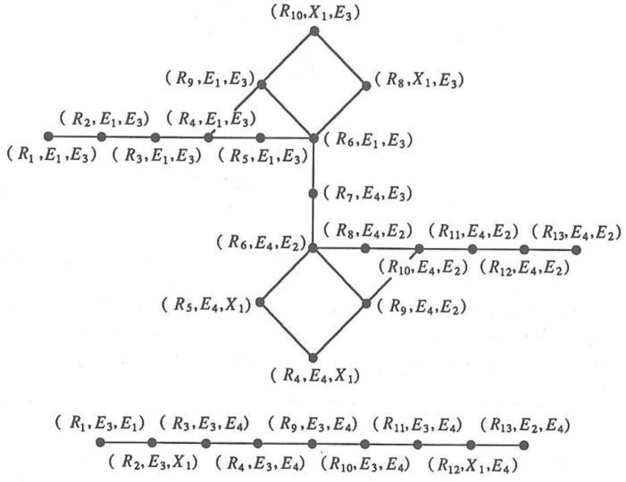

[Drawings: Latombe's Robot Motion book]

Piano Movers II, IV, V

- General framework for *solving almost any motion planning problem* with arbitrary number of degrees of freedom
- Tarski '51: Theoretical guarantee that problems defined by real polynomial constraints (semi-algebraic sets) are decidable

  <the first-order theory of real closed fields is decidable>
- Collins '75: Cylindrical Algebraic Decomposition (CAD), a constructive realization of Tarski's theory
- Put together:
  - Tarski's result implies that motion planning problems are decidable — since configuration space obstacles can be described by semi-algebraic sets
  - They use Collins' CAD as a practical method to construct decompositions of the configuration space, which allows for path connectivity and decision procedures in configuration spaces

Impact of the Piano Movers papers

- Robotics: Mathematical and algorithmic foundations for robot motion planning and more generally for algorithmic robotics
- Computational Geometry: Opens up to the real, curved world (previously almost exclusively linear), arrangements of curves and surfaces

Complete solutions II

- the problem is hard when the number of degrees of freedom (# dofs) is part of the input [Reif 79], [Hopcroft-Schwartz-Sharir 84], …
- cell decomposition the Piano movers series [Schwartz-Sharir 83]: a doubly-exponential solution
- roadmap [Canny 87], [Basu-Pollack-Roy]: a singly-exponential solution
- few dofs, the Sharir school: very efficient, near-optimal, solutions (mid 80s – mid 90s)

Few degrees of freedom: The piano movers school, mid 1980s – mid 1990s

- The robot has fixed descriptive complexity
- The obstacles have complexity n (e.g., # of vertices of the polygons)

2 degrees of freedom

- The configuration space is a 2D arrangement of well-behaved curves
- Hence the complexity of the full C-space is $O(n^2)$

2 dofs, optimal combinatorial results

- If the robot is convex (translation), the complexity of the entire C-space is $O(n)$

  Klara Kedem, Ron Livne, János Pach, Micha Sharir: On the Union of Jordan Regions and Collision-Free Translational Motion Amidst Polygonal Obstacles. Discret. Comput. Geom. 1: 59-70 (1986)

- Otherwise (general), the complexity of a single free cell in C-space is near-linear <Davenport-Schinzel related functions>, e.g. $O(n\alpha(n))$

  Leonidas J. Guibas, Micha Sharir, Shmuel Sifrony: On the General Motion-Planning Problem with Two Degrees of Freedom. Discret. Comput. Geom. 4: 491-521 (1989)

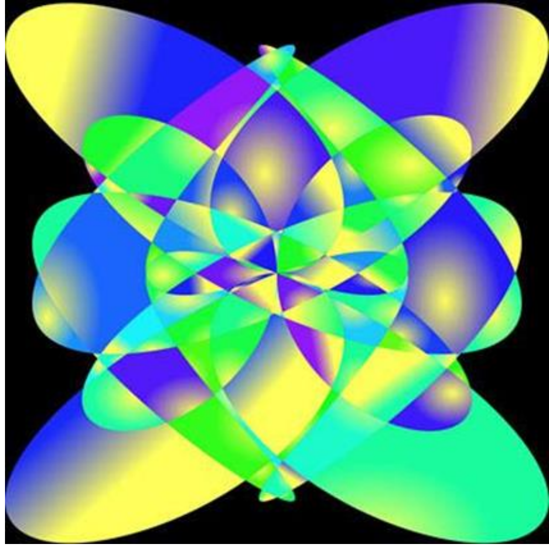

My favorite Micha paper

3 dofs, combinatorics

- The configuration space is a 3D arrangement of well-behaved surfaces
- Hence the complexity of the full C-space is $O(n^3)$

Expectations

- If the robot is convex, the complexity of the entire C-space would be near-quadratic
- Otherwise, the complexity of a single free cell in C-space would be near-quadratic

Discrete Comput Geom 2:255–270 (1987)

Discrete & Computational
Geometry

On the Number of Critical Free Contacts of a Convex Polygonal Object Moving in Two-Dimensional Polygonal Space*

Daniel Leven[1] and Micha Sharir[1,2]*

[1] School of Mathematical Sciences, Tel-Aviv University, Tel-Aviv, Israel

[2] Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA

Abstract. We show that the number of critical positions of a convex polygonal object B moving amidst polygonal barriers in two-dimensional space, at which it makes three simultaneous contacts with the obstacles but does not penetrate into any obstacle is $O(kn\lambda_s(kn))$ for some $s \leq 6$, where k is the number of boundary segments of B, n is the number of wall segments, and $\lambda_s(q)$ is an almost linear function of q yielding the maximal number of "breakpoints" along the lower envelope (i.e., pointwise minimum) of a set of q continuous functions each pair of which intersect in at most s points (here a breakpoint is a point at which two of the functions simultaneously attain the minimum). We also present an example where the number of such critical contacts is $\Omega(k^2n^2)$, showing that in the worst case our upper bound is almost optimal.

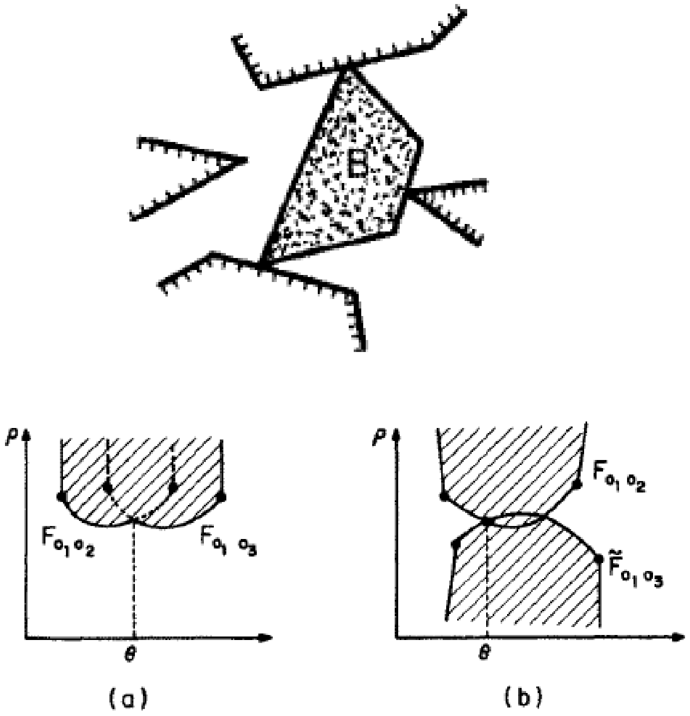

Type (iv) Critical Contacts

How about 3D work space?

- What is the combinatorial complexity of the free C-space for a convex polyhedron translation among polyhedral in 3-space?
- Or just a box?

- Box: $O(n^2\alpha(n))$

  Dan Halperin, Chee-Keng Yap: Combinatorial Complexity of Translating a Box in Polyhedral 3-Space. SoCG 1993: 29-37 applying Leven-Sharir '87 as above

- Convex polyhedron among k convex polydedra (n is tricky): $O(nk \log k)$

  Boris Aronov, Micha Sharir: On Translational Motion Planning of a Convex Polyhedron in 3-Space. SIAM J. Comput. 26(6): 1785-1803 (1997)

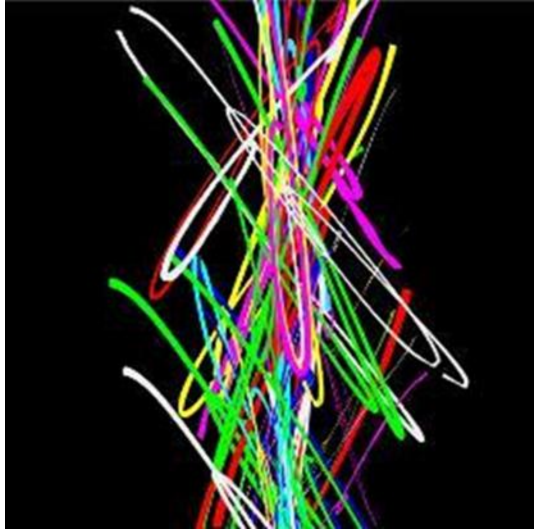

CGAL arrangements and Micha

CGAL

- Computational Geometry Algorithms Library
- A collection of software packages written in C++
- Adheres the generic programming paradigm
- Development started in 1995
- Funded by a succession of EU projects
- An open source library

The arrangement package and its relatives

- Arrgs of curves on 2D surfaces
- Boolean operations on curved objects
- Minkowski sums
- Lower envelopes in 3D

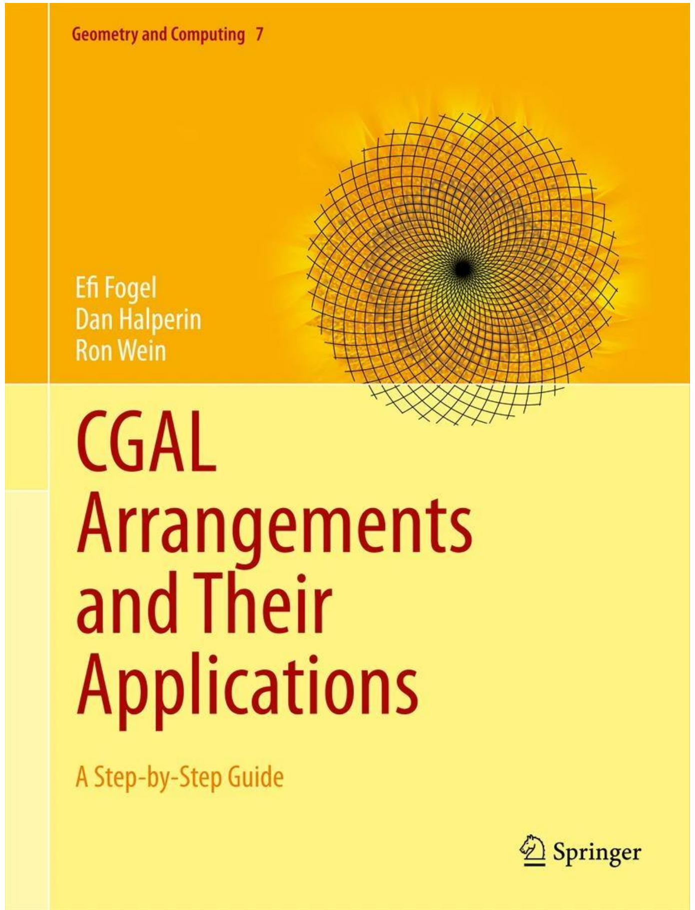

Lessons from Micha

- Arrangements *can solve almost everything*
- Separate the topology (combinatorics) from the algebra
- Vertical decomposition

In practice:
general CG algorithms and data structures for well-behaved curves
together with
Computational algebra toolbox for curves, Exacus (Mehlhorn et al) then CGAL

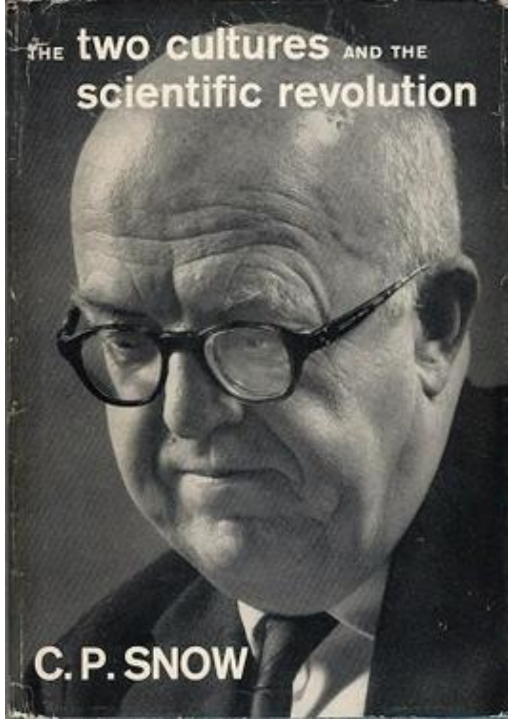

CG and robotics:
A tale of two cultures

and the rise of SB planners

A meeting at ADEPT in 1993

- Silicon valley
- Jean-Claude Latombe, Ken Goldberg, Brian Carlisle, …
- The topic: The future of robotics
- Crisis, robotics winter
- Computational geometry summer (gliding toward autumn)
- 12 problems to work on in robotics

[CASE 2015]

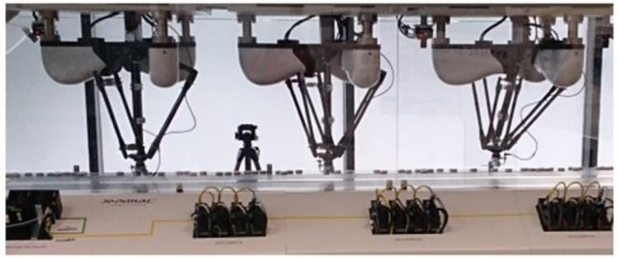

[Omron offices 2018]

Two events in 1994

- Bernard Chazelle's speech of admonishment (SoCG)

Does Computational Geometry Have A Future?

BERNARD CHAZELLE
Department of Computer Science
Princeton University
Princeton, NJ 08544, USA
chazelle@cs.princeton.edu

I believe that we, CG'ers, need to initiate a serious debate regarding the future of our field, in general, and the ACM CG conference, in particular.

The last few years have witnessed impressive advances in CG on the theoretical front. There are smarter people working in CG today than at any time before. The community can take collective credit for its sustained focus and its success in providing CG with solid foundations. In the area of algorithms and data struc-

impressed that CG has something to teach them. There's still no comprehensive geometric software library. Evety programming effort must start from scratch. Most geometric codes are still intended as existence proofs, with little or no follow-up on the user end. There are, as you well know, several first-rate projects going on now to turn this around, but these are still by and large individual, scattered efforts.

- The first WAFR

WAFR: Algorithmic Foundations of Robotics, 1994

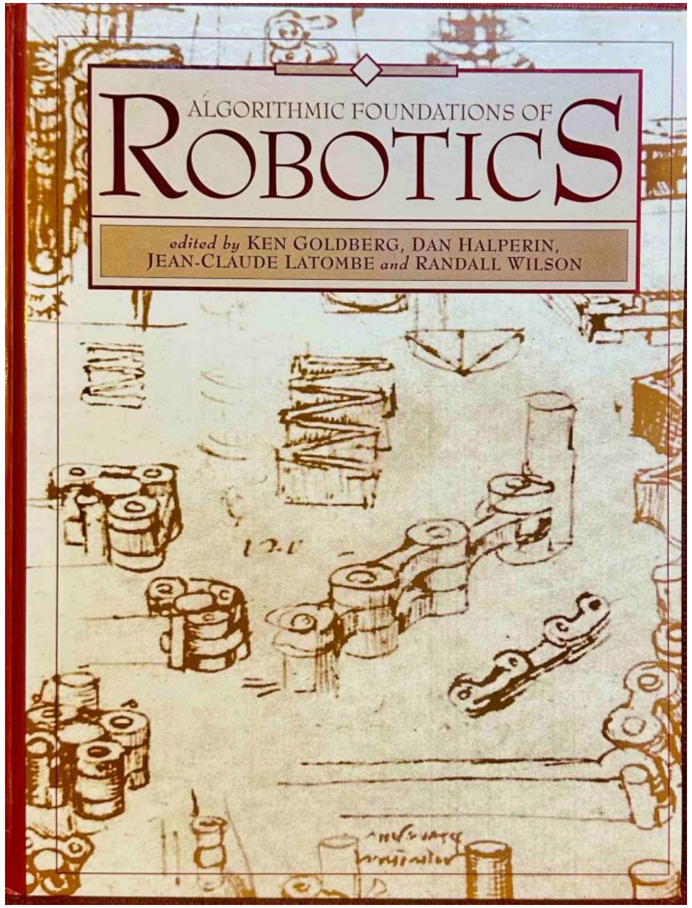

- Micha Sharir
- Mark Overmars
- Richard Pollack
- Leo Guibas
- Jean-Daniel Boissonnat
- …

- Jean-Claude Latombe
- Ken Goldberg
- John Canny
- Matt Mason
- Lydia Kavraki
- …

computational geometry

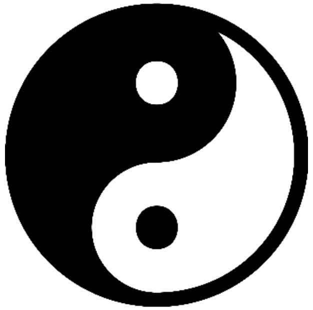

robotics

Participants

Rachid Alami	LAAS–CNRS
Ruzena Bajcsy	University of Pennsylvania
Saugata Basu	New York University
Pierre Bessière	LIFIA
Jean-Daniel Boissonnat	INRIA
Randy C. Brost	Sandia National Laboratories
John Canny	University of California, Berkeley
Raja Chatila	LAAS–CNRS
Alan Christiansen	Tulane University
Thomas Dean	Brown University
Bruce Donald	Cornell University
Michael Erdmann	Carnegie Mellon University
Ayman O. Farahat	Texas A&M University
Ken Goldberg	University of Southern California
Leonidas J. Guibas	Stanford University
Dan Halperin	Stanford University
Seth Hutchinson	University of Illinois, Urbana-Champaign
Daniel P. Huttenlocher	Cornell University
Leo Joskowicz	IBM T.J. Watson Research Center
Takeo Kanade	Carnegie Mellon University
Lydia Kavraki	Stanford University
Oussama Khatib	Stanford University
Pradeep K. Khosla	Carnegie Mellon University
Daniel E. Koditschek	University of Michigan
Yotto Koga	Stanford University
Krasimir Kolarov	Interval Research
David J. Kriegman	Yale University
Jean-Claude Latombe	Stanford University
Christian Laugier	LIFIA
Jean-Paul Laumond	LAAS–CNRS
Tsai-Yen Li	Stanford University
Kevin M. Lynch	Carnegie Mellon University
Matthew T. Mason	Carnegie Mellon University
Brian Mirtich	University of California, Berkeley
Rajeev Motwani	Stanford University
Balas K. Natarajan	Hewlett Packard Labs
Mark H. Overmars	Utrecht University
Christos Papadimitriou	University of California, San Diego
Christiaan J. J. Paredis	Carnegie Mellon University
Richard Pollack	New York University
Jean Ponce	University of Illinois, Urbana-Champaign
Anil Rao	Utrecht University
John H. Reif	Duke University
Ari Requicha	University of Southern California
Daniela Rus	Cornell University
Micha Sharir	Tel Aviv University
Camillo J. Taylor	Yale University
Marek Teichmann	New York University
Bob Tilove	General Motors
Jeffrey C. Trinkle	Texas A&M University
Hongyan Wang	Duke University
Randall H. Wilson	Sandia National Laboratories

From the mid 1990s, a dramatic change

- Sampling-base planners appear/catch on (PRM, RRT, …)

Roughly:

- They *can solve almost everything*
- Easy to implement
- Work well for a variety of real-life problems

- Roboticists did not need complex CG algorithms anymore
- CGers did not find interest in the new `simplistic' techniques

A tale of two cultures

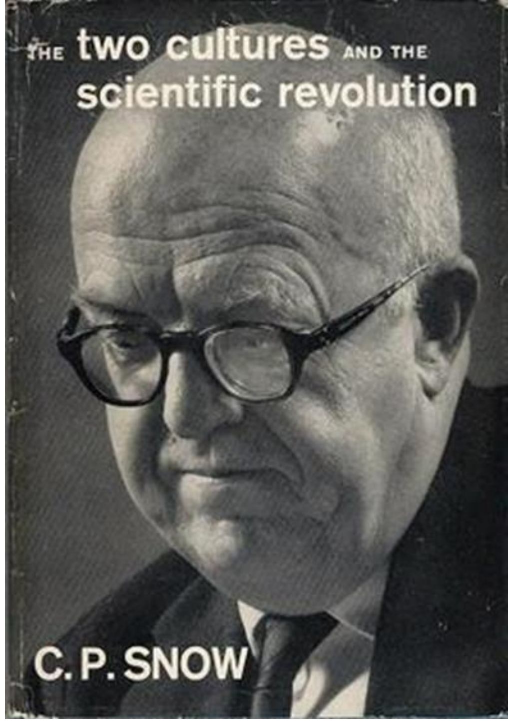

- The thesis of Snow’s lecture and book,1959 (ruthlessly snipped to suit the current setting): […] had become split into "two cultures" and that this division was a major handicap to both in solving the […] problems

Very brief intro to SB planners

and their theoretical guarantees

Sampling-based planning

- Probabilistic Roadmaps illustrated

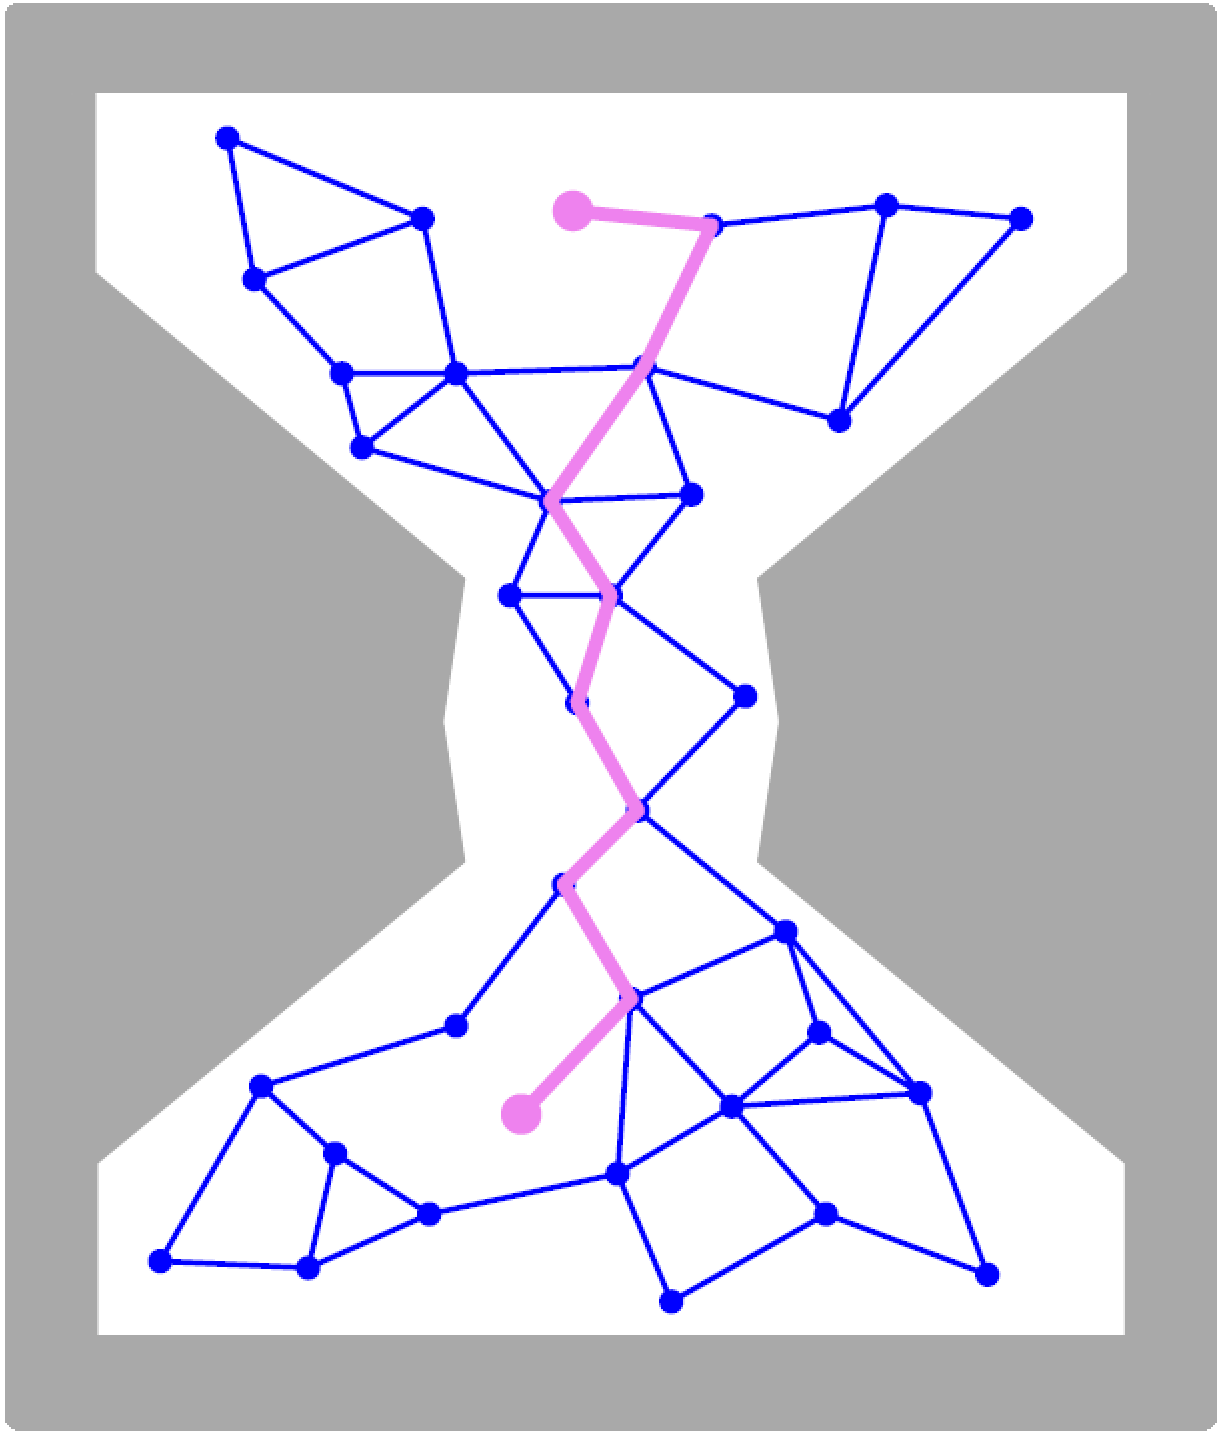

Milestones

- Probabilistic roadmaps for path planning in high-dimensional configuration spaces [Kavraki-Svestka-Latombe-Overmars. IEEE Trans. Robotics Autom 1996] PRM
- RRT-Connect: An efficient approach to single-query path planning [Kuffner-LaValle. ICRA 2000] RRT
- Sampling-based algorithms for optimal motion planning [Karaman-Frazzoli IJRR 2011] RRT*

Type of guarantees

- A motion planner is said to be complete if the planner in finite time either produces a solution or correctly reports that there is none
- Probabilistic completeness is the property that as more "work" is performed, the probability that the planner fails to find a path, if one exists, asymptotically approaches zero
- Asymptotic optimality is the property of almost-sure convergence to optimal solutions with increasing number of samples

Type of guarantees

- A motion planner is said to be complete if the planner in finite time either produces a solution or correctly reports that there is none
- Probabilistic completeness is the property that as more "work" is performed, the probability that the planner fails to find a path, if one exists, asymptotically approaches zero
  - PC of RRT [Kleinbort-Solovey-Littlefield-Bekris-H, IEEE RA-L 2019]
  - Non-optimality of bi-RRT [Nechushtan-Raveh-H, WAFR 2010]
- Asymptotic optimality is the property of almost-sure convergence to optimal solutions with increasing number of samples
  - The critical radius [Solovey-Kleinbort, RSS 2018, IJRR 2020]
  - Asymptotically near-optimal (LBT) RRT [Salzman-H, ICRA 2015, IEEE TOR 2016]

Near-optimality for MRMP with finite sampling

- Optimality criteria: the total length travelled by the robots from start to goal, or max length travelled by a robot
- Given δ and ε, if a multi-robot motion planning has a δ-clear solution then we have a finite-sampling scheme that guarantees to find a solution that is at most (1+ ε) longer than the optimal solution
- Guarantee to find a solution within bounded time using A*, dRRT, dRRT*, Conflict Based Search, M* or any other tensor-product based algorithm

[Dayan-Solovey-Pavone-H, ICRA 2021,

IEEE TOR 2023]

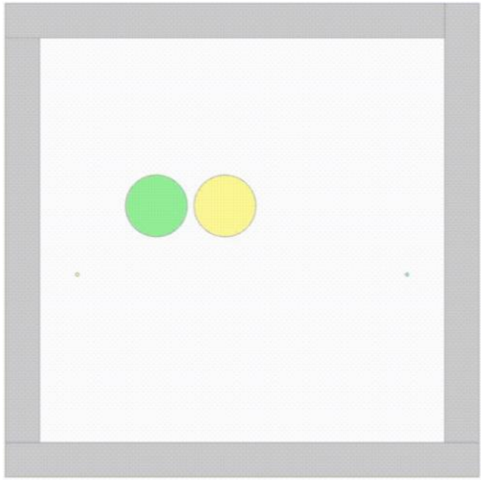

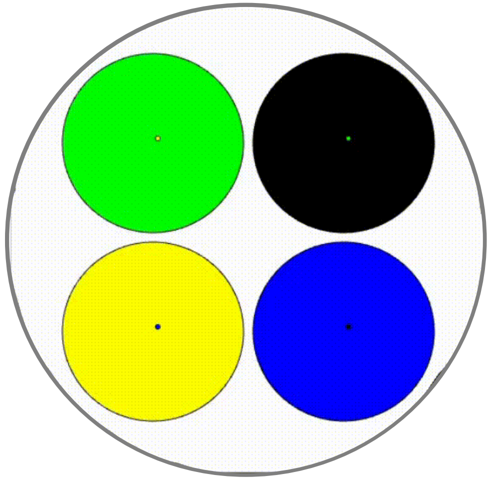

In the d-dimensional unit cube C-space

Theorem 4 (Sufficient conditions for MRMP $(\varepsilon, \vec{\delta})$-completeness)**.** Let $\varepsilon > 0$ be a stretch factor, let $\vec{\delta}$ be a clearance vector $(\delta_1, \ldots, \delta_R)$, and denote $\omega = \frac{\varepsilon}{2(\varepsilon+2)}$. Define the sampling distributions $\vec{\mathcal{X}} = (\mathcal{X}_1, \ldots, \mathcal{X}_R)$ and radii vector $\vec{r} = (r_1, \ldots, r_R)$, as

$$\mathcal{X}_i = \mathcal{X}_{\omega\delta_i, \delta_i}, \quad r_i = \delta_i(\varepsilon + 1)/(\varepsilon + 2),$$

for every robot $1 \leq i \leq R$. Then $\left(\vec{\mathcal{X}}, \vec{r}\right)$ is $\left(\varepsilon, \vec{\delta}\right)$-complete.

Underlying sampling per robot: ε-cover with staggered grids

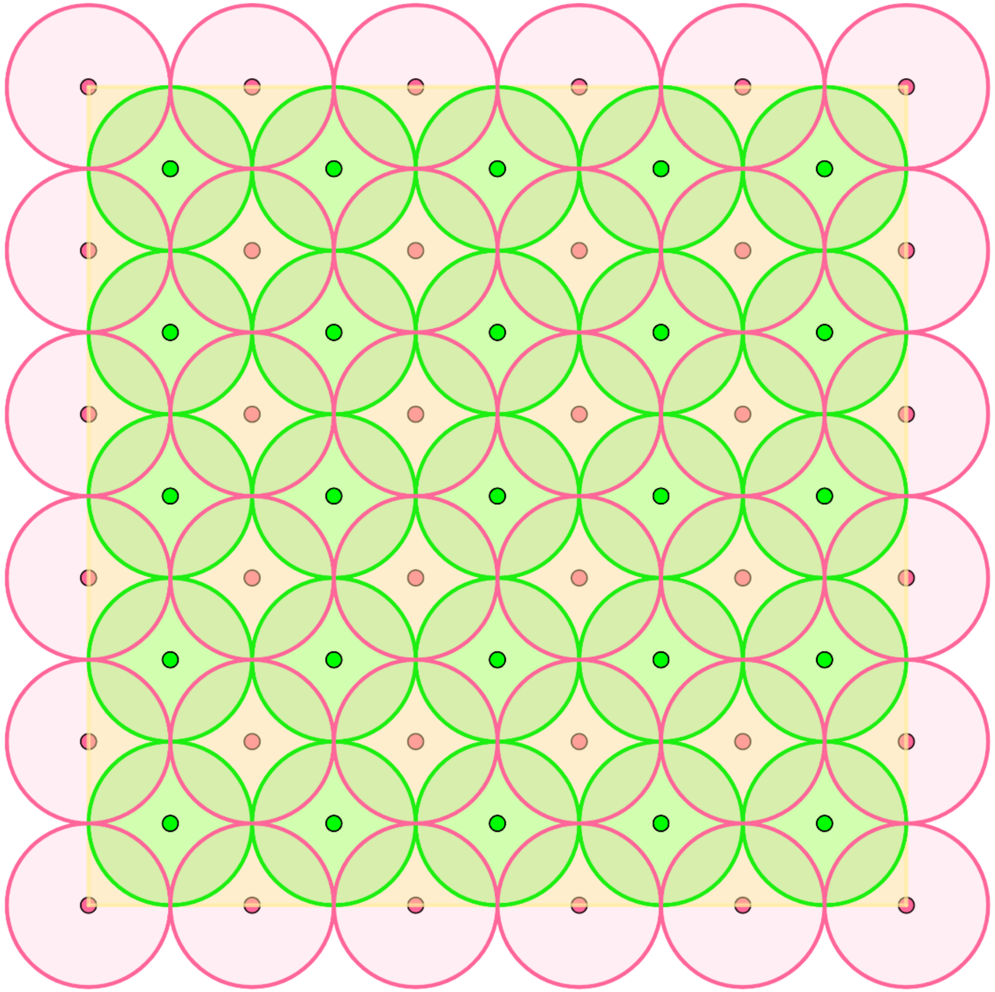

Surprise and news

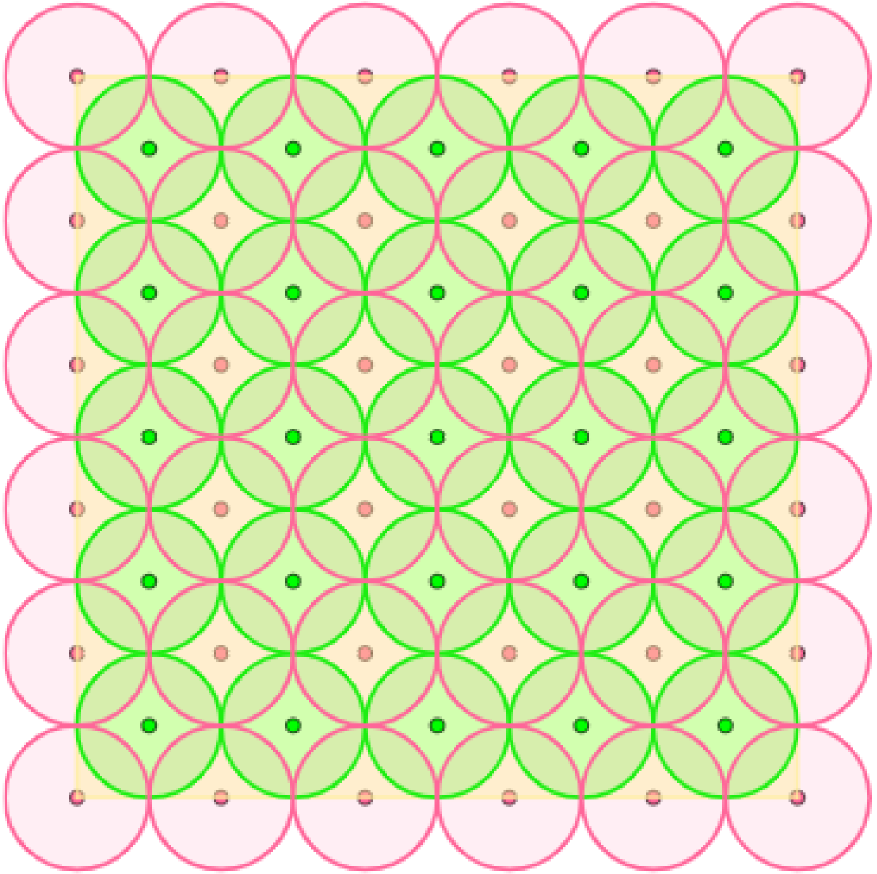

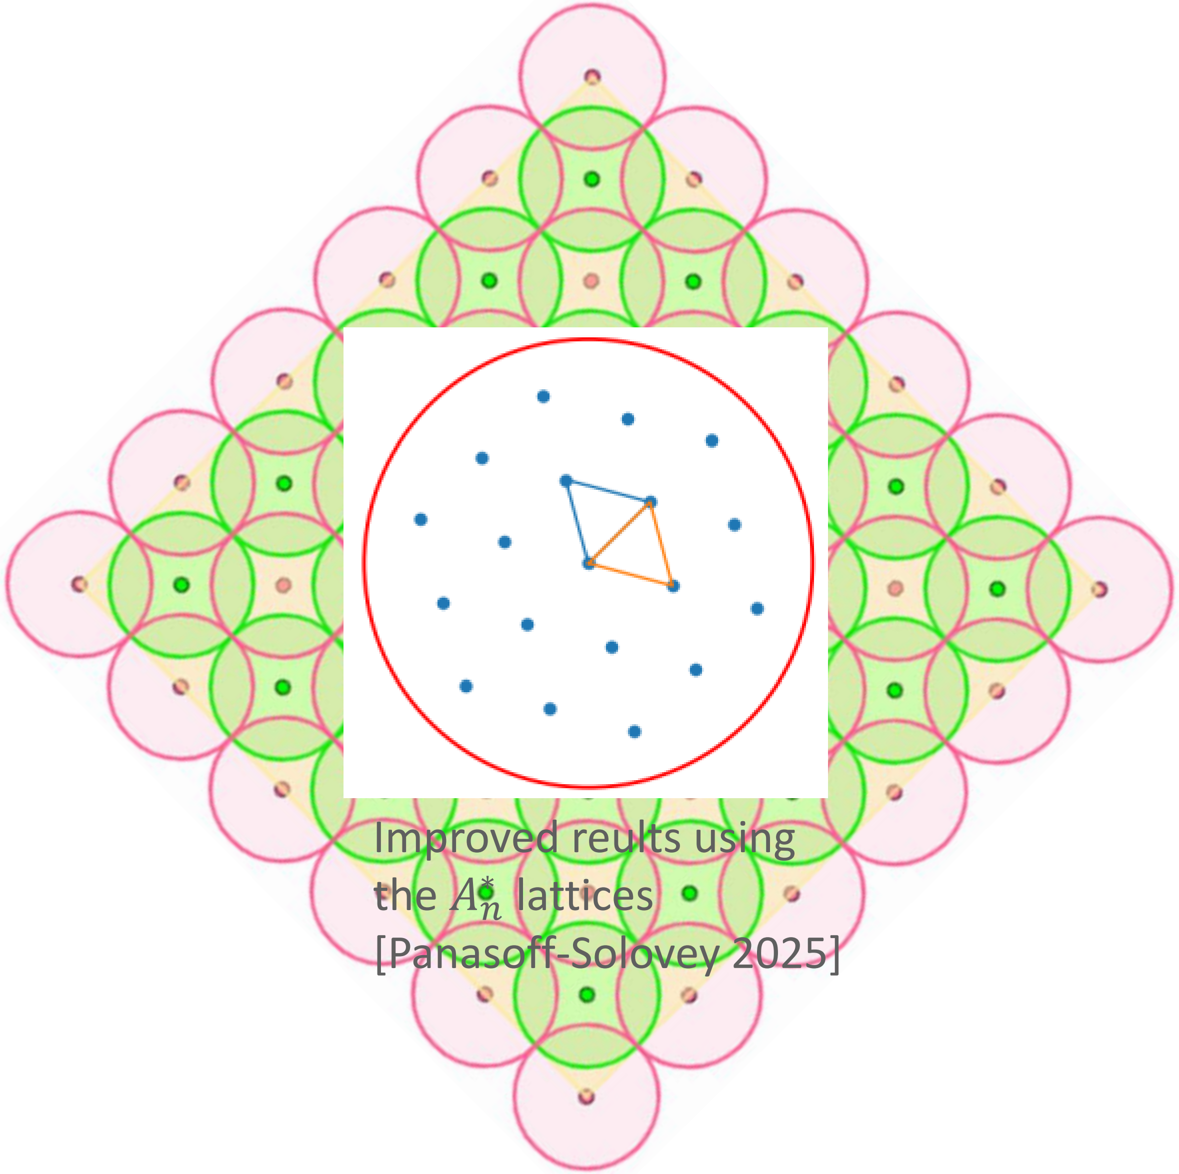

Improved reults using the A_n^* lattices [Panasoff-Solovey 2025]

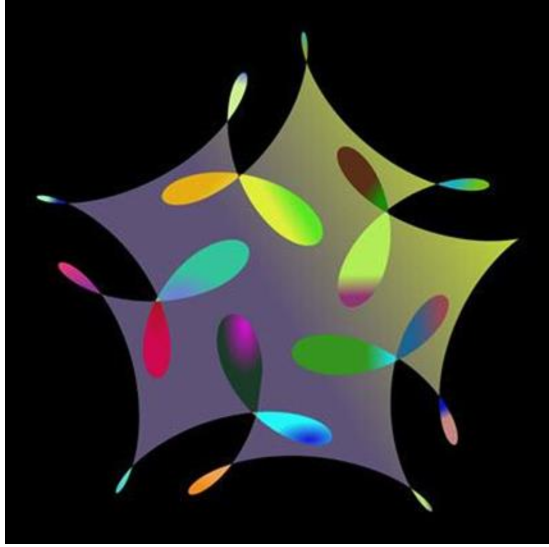

Optimizing the coordination of fleets of robots

and a recent paper with Micha

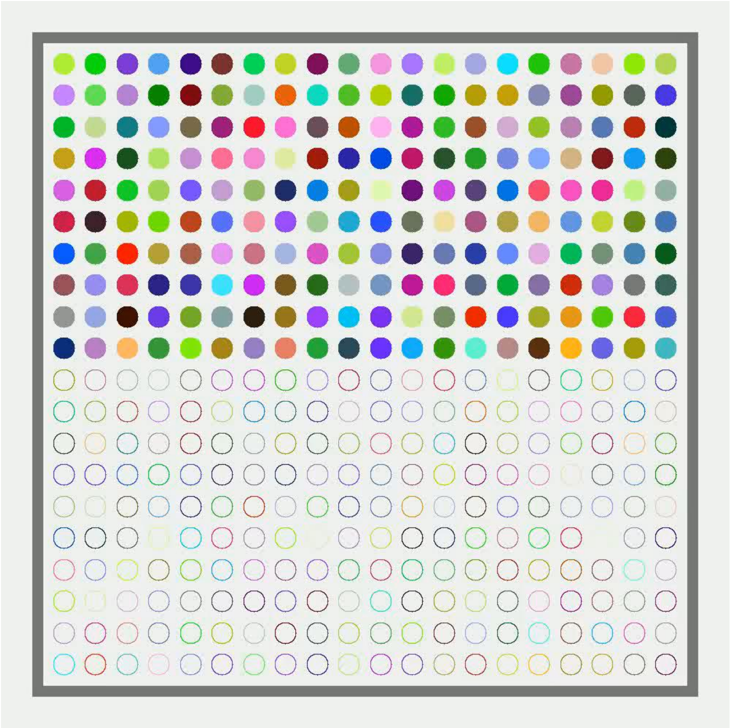

- How about 2 robots?

Riddle: MinSum

Given two Roomba’s, each has to move from given start to goal positions, in a room without obstacles. What are the joint shortest paths (minimum total length)?

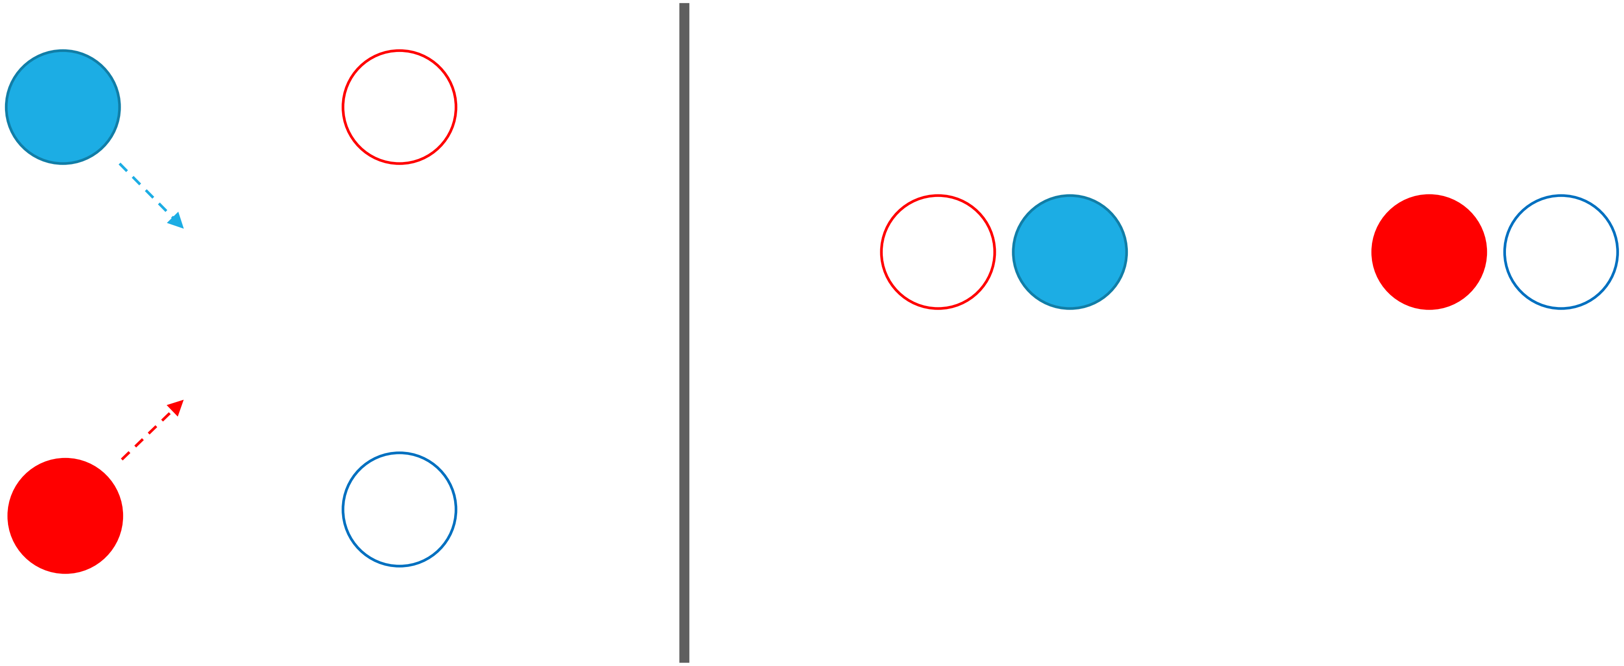

Riddle: MinSum

Given two Roomba's, each has to move from given start to goal positions, in a room without obstacles. What are the joint shortest paths (minimum total length)?

Answer:

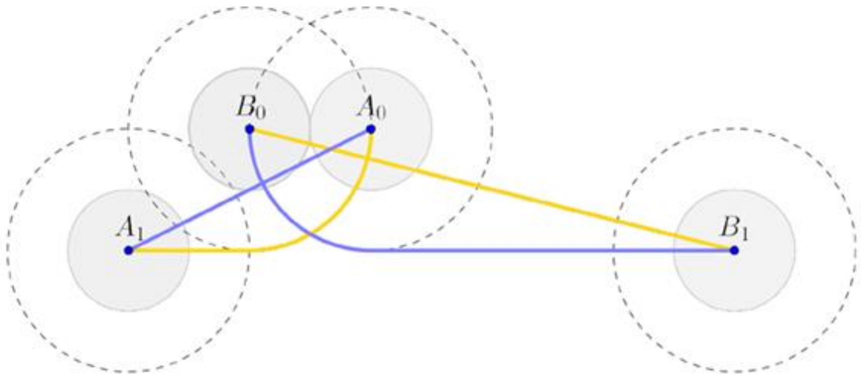

[Kirkpatrick-Liu, CCCG 2016]

2 squares: [Esteban-H-Silveira, Autonomous Robots 2025]

Open problem: MinTime

Given two Roomba's, each has to move from given start to goal positions, in a room without obstacles moving in unit speed. What is the shortest time plan (makespan)?

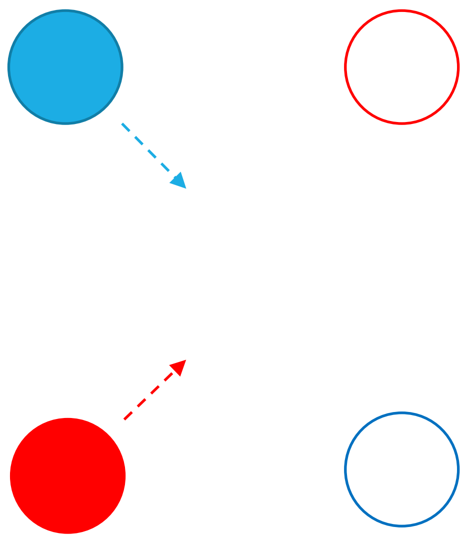

Open problem: MinSum with obstacles

If we add obstacles we not know how to efficiently compute the joint length (NP-hard?)

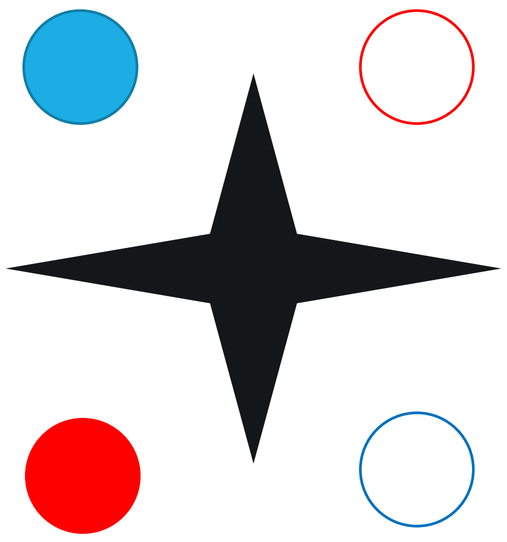

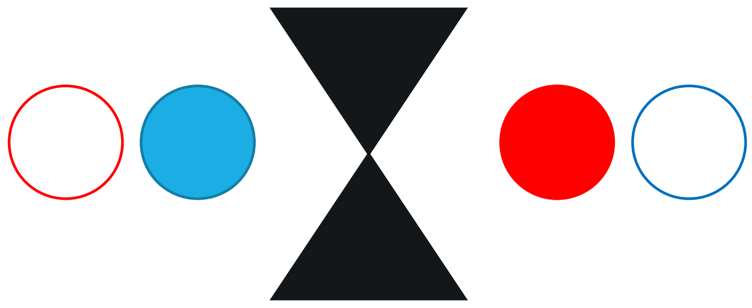

Motion planning for (2) discs without optimization

- Jacob T. Schwartz and Micha Sharir: On the Piano Movers' Problem: III. Coordinating the Motion of Several Independent Bodies: The Special Case of Circular Bodies Moving Amidst Polygonal Barriers. The International Journal of Robotics Research, 1983
- Micha Sharir, Shmuel Sifrony: Coordinated Motion Planning for Two Independent Robots. Ann. Math. Artif. Intell. 3(1): 107-130 (1991)

Algorithm for optimal motion of 2 squares among obstacles

Pankaj K. Agarwal, Dan Halperin, Micha Sharir, Alex Steiger: Near-Optimal Min-Sum Motion Planning for Two Square Robots in a Polygonal Environment. SODA 2024: 4942-4962

- An $n^2 e^{-O(1)} \log n$-time $(1 + \varepsilon)$-approximation algorithm for this problem
- The first polynomial-time $(1 + \varepsilon)$-approximation algorithm for an optimal motion-planning problem involving two robots moving in a polygonal environment

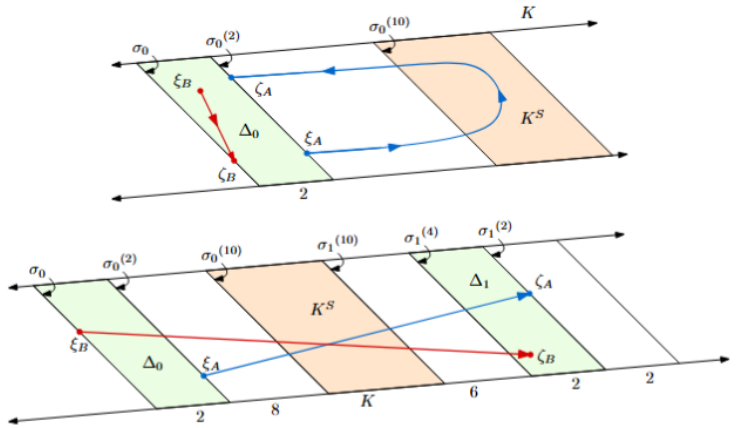

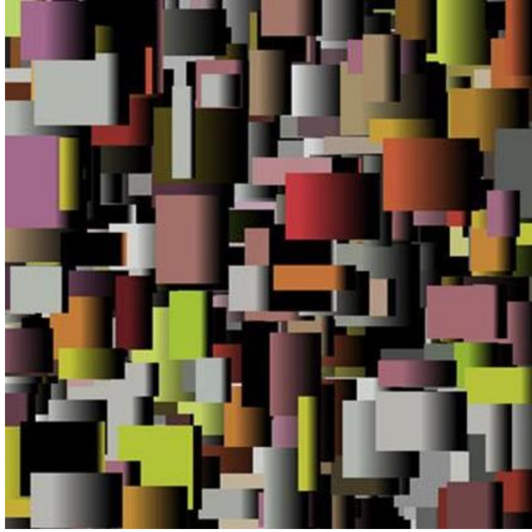

Final notes

A little history from the 1950s

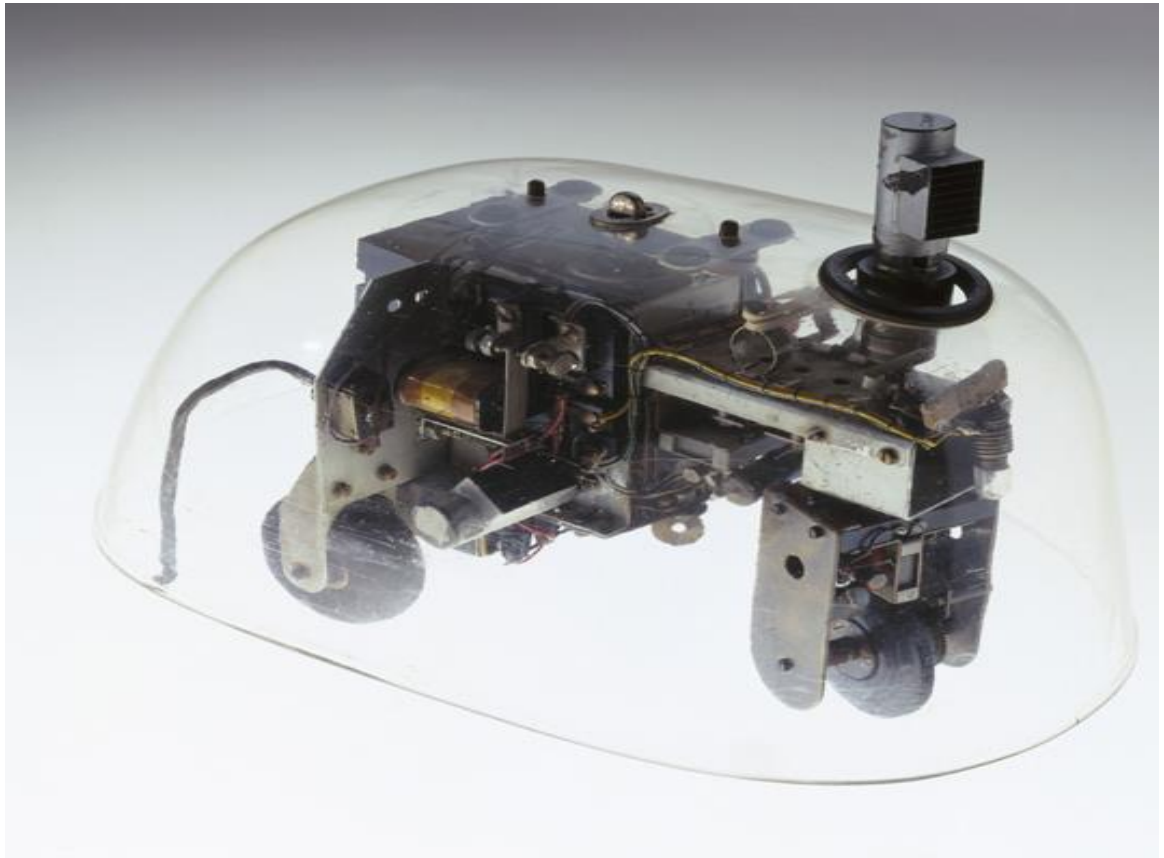

- Grey Walter's tortoises
- “Codebreaker – Alan Turing's life and legacy” at the Science Museum 2012
- Turing’s visit to the Science Museum 1951

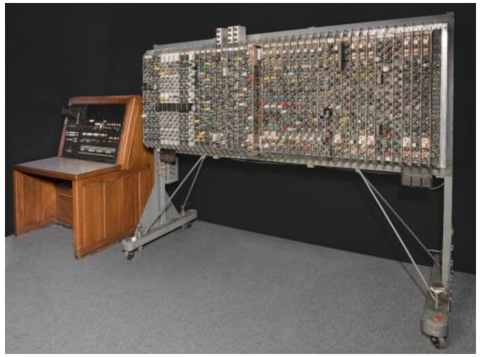

1994 (WAFR established) → 2024, 2026

- Winds of change

- WAFR 2026: Oulu Finland, June 15-17, 2026
- Deadline: January 15, 2026

More hard problems

Ten Problems in Geobotics*

Mikkel Abrahamsen Dan Halperin

August 2024

Abstract

Robots sense, move and act in the physical world. It is therefore natural that algorithmic problems in robotics and automation have a geometric component, often central to the problem. Below we review ten challenging problems at the intersection of robotics and computational geometry—let's call this intersection *Geobotics*. What is common to most of these problems is that the prevalent algorithmic techniques used in robotics do not seem suitable for solving them, or at least do not suggest quality guarantees for the solution. Solving some of them, even partially, can shed light on less well-understood aspects of computation in robotics.

0 Introduction

Robotics has persistently raised interesting problems for computational geometry. A pioneering and exemplary case in point is the 1980s series of papers "On the Piano Movers

From a website by Nina Amenta

me = Nina Amenta:

“Here I am at a workshop with a lot of Computational Geometers who are taller than I am. You can just see my eye peeking out behind Micha Sharir.”

Thank you Micha!

[Art by Al. Jeb Gaither using CGAL arrangements]