Algorithmic Robotics and Motion Planning

Multi robot motion planning: Extended review

Dan Halperin
School of Computer Science
Tel Aviv University

Fall 2019-2020

Alternative settings/approaches

- distributed, swarm
- the discrete version: MAPF= multi agent path finding
- machine learning

we will review central-control algorithms in continuous domains

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.

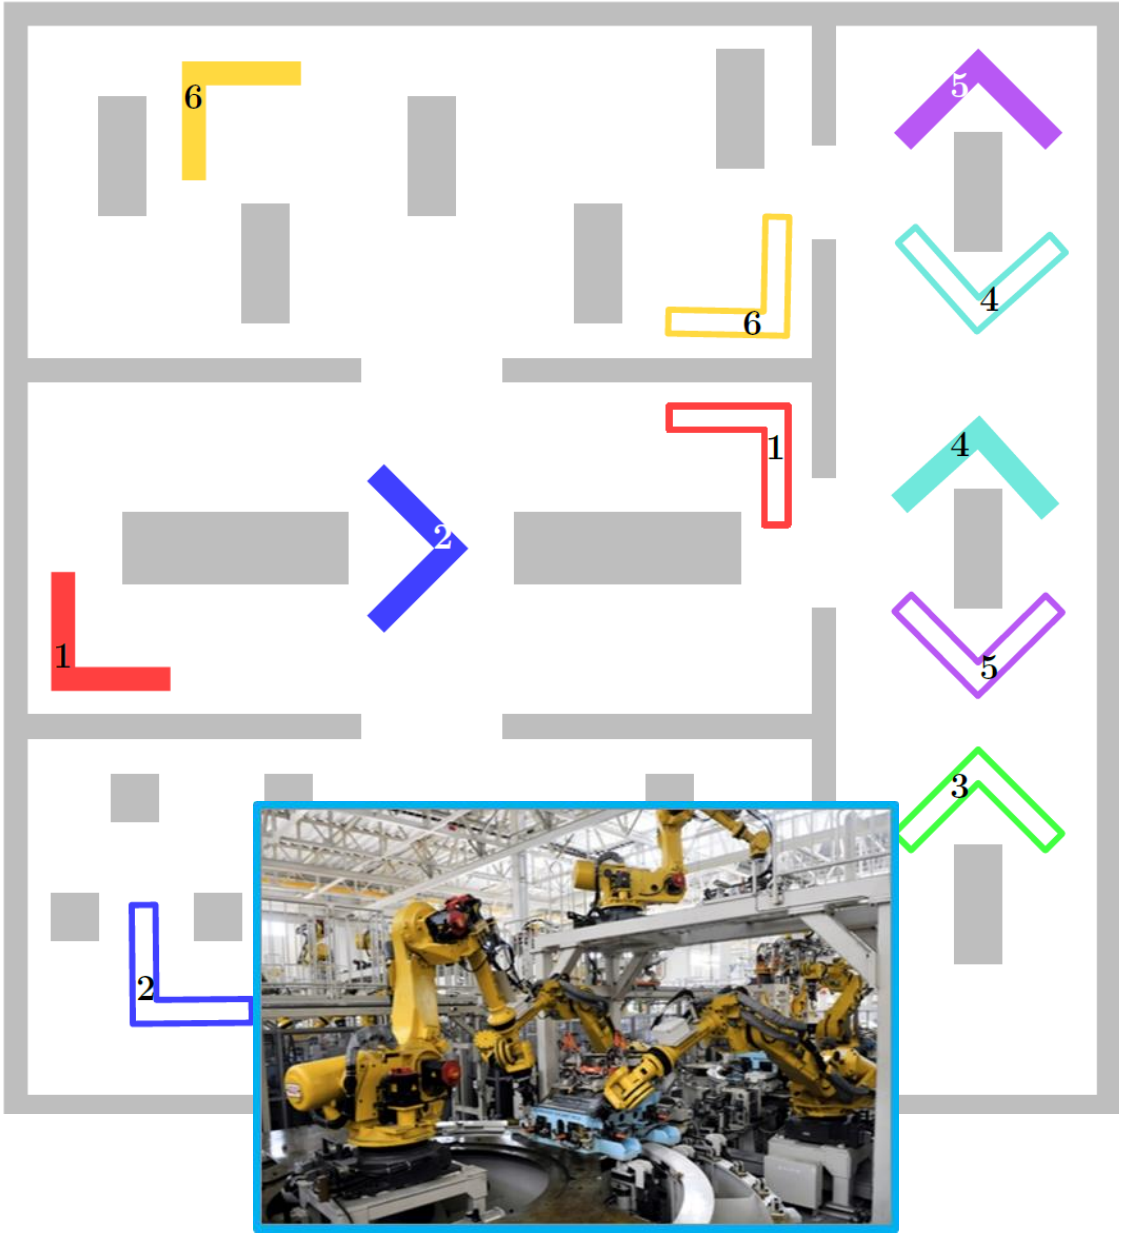

Two key terms:

(i) degrees of freedom (dof), and

(ii) configuration space

Review overview

- motion planning, an ultra brief history, hard-vs-easy perspective
- Hard vs. easy:

  unlabeled motion planning for many discs
- multi-robot planning in tight settings
- summary and outlook

Motion planning,
an ultra brief history

Complete solutions

- the problem is hard when the number of degrees of freedom (# dof) is part of the input [Reif 79], [Hopcroft et al. 84], …

- cell decomposition the Piano movers series [Schwartz-Sharir 83]: a doubly-exponential solution

- roadmap [Canny 87], [Basu-Pollack-Roy]: a singly-exponential solution

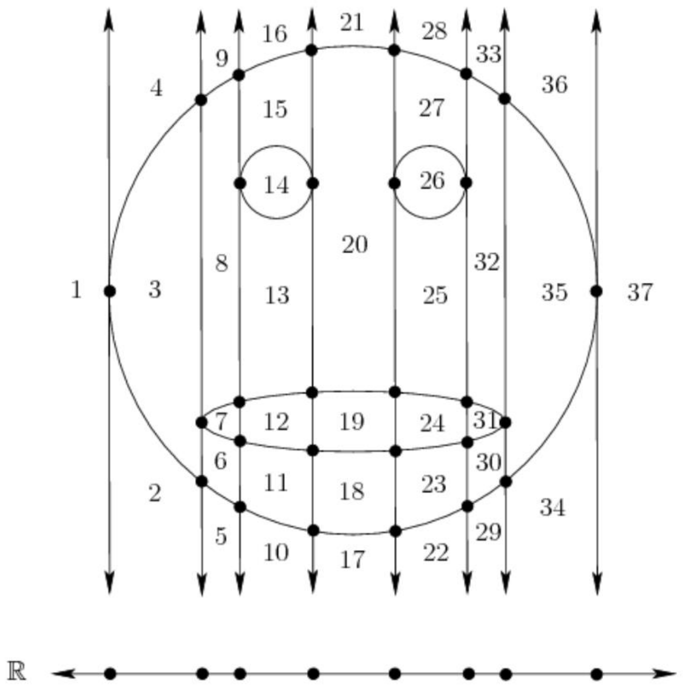

[LaValle]

- few dof: very efficient, near-optimal, solutions (mid 80s – mid 90s)

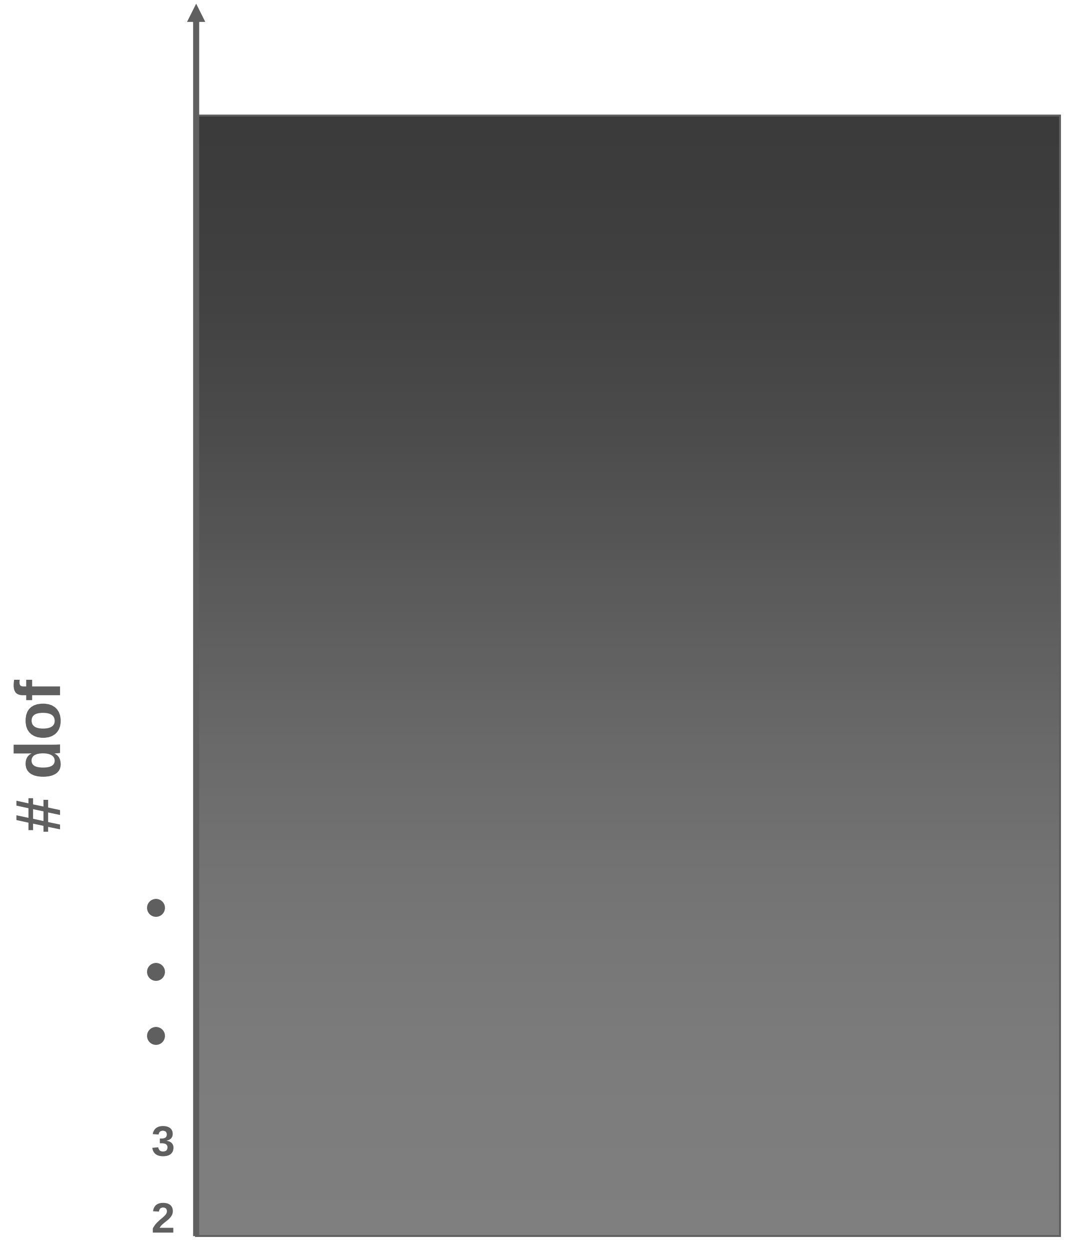

dof
3
2

Meanwhile in robotics

- potential field methods [Khatib 86]
  attractive potential (goal), repulsive potential (obstacles)
- random path planner (RPP)
  [Barraquand-Latombe 90]

- and then, around 1995
  PRM (Probabilistic RoadMaps)
  [Kavraki, Svestka, Latombe,Overmars]
- RRT (Rapidly Exploring Random Trees) [LaValle-Kuffner 99]
- many variants followed
- numerous uses, also for many dof

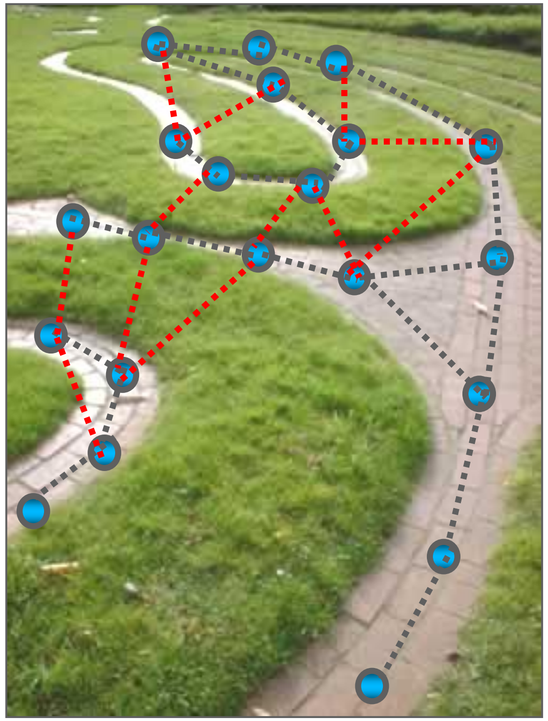

Hard or easy?

- when is motion planning hard or easy?
- (modern) folklore: it's hard when there are narrow passages in the C-space on the way to the goal

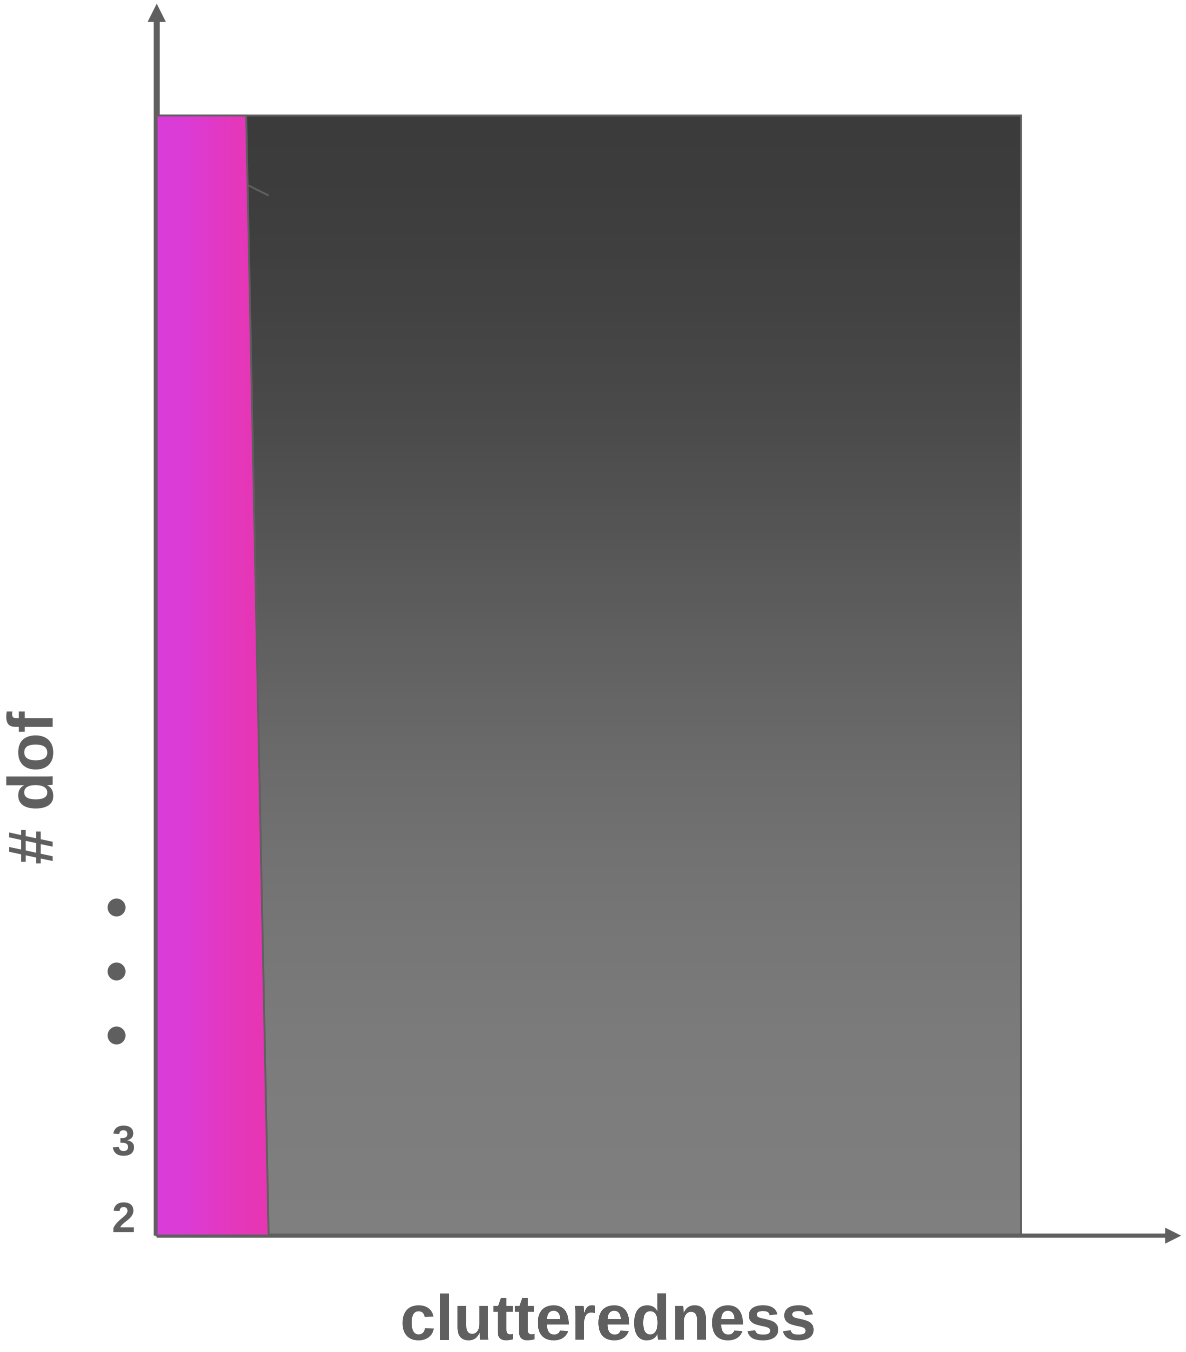

dof
3
2
clutteredness

The role of clearance

- probabilistic completeness proofs require an empty sleeve around the solution path
  - the needed number of samples is inversely proportional to the width of this empty sleeve
  - it seems equally hard to compute this width a priori

Hard vs. easy:
Unlabeled motion planning for many discs

k-Color multi robot motion planning

- m robots arranged in k groups
- The extreme cases:
  - k=m, the standard, fully colored problem
  - k=1, the unlabeled case
    - [Kloder and Hutchinson T-RO 2006]
    - [Turpin-Mohta-Michael-Kumar AR 2014 (ICRA 2013)]

[Solovey-H, WAFR 2012, IJRR 2014]

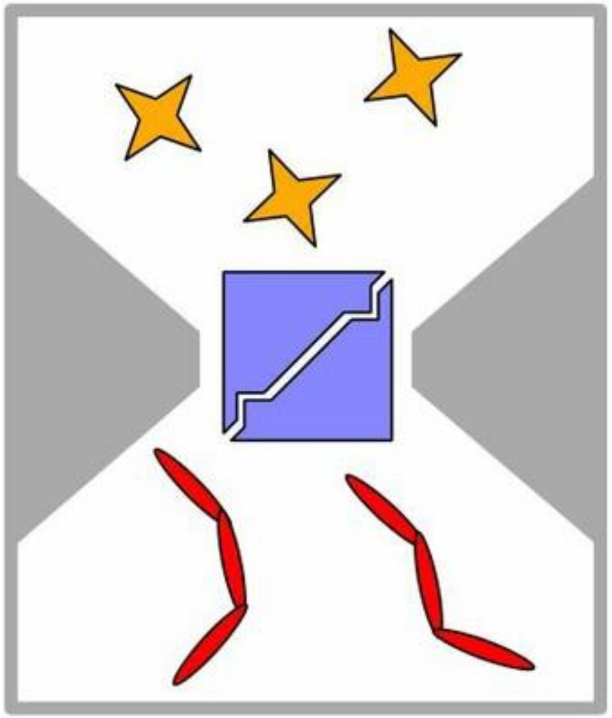

m=7, k=3

Unlabeled motion planning

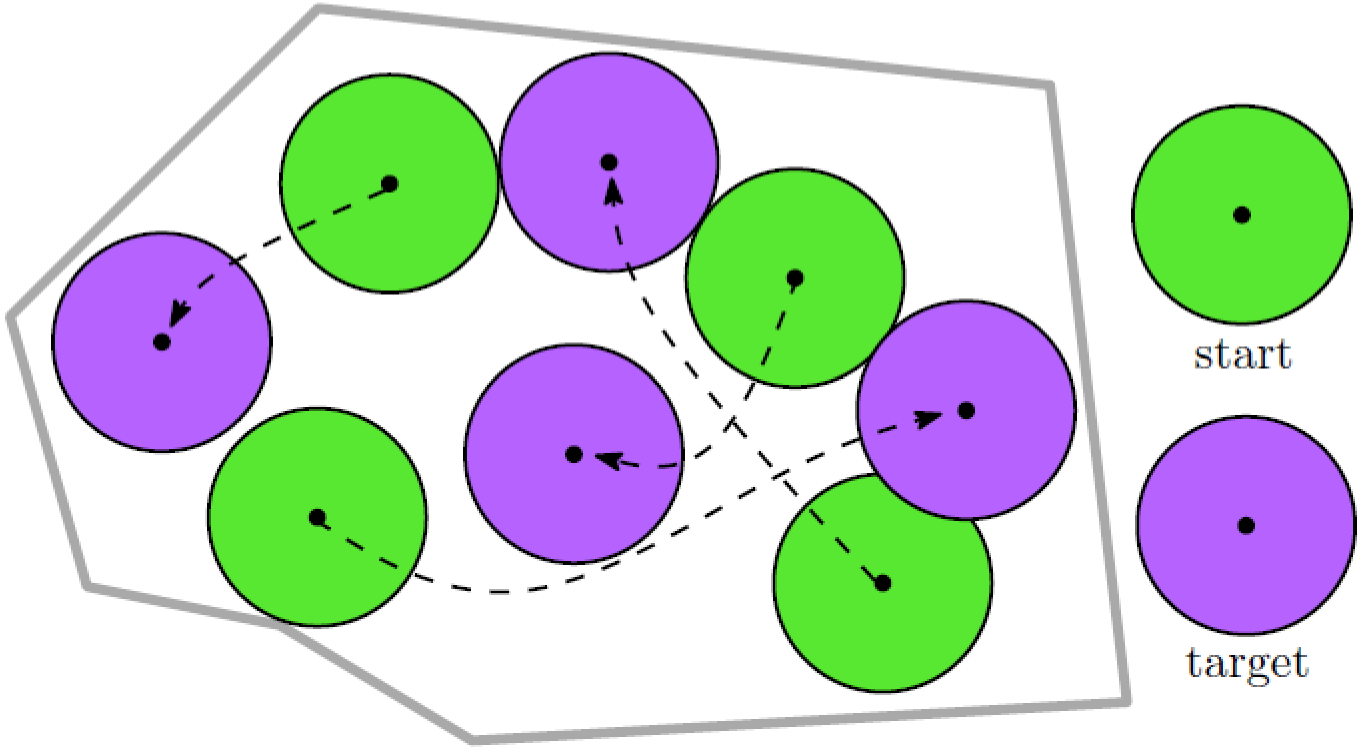

Unlabeled discs in the plane: the problem

Plan the motion from start to goal:

- m interchangeable unit disc robots
- moving inside a simple polygon with n sides
- each of the m goal positions needs to be occupied by some robot at the end of the motion
- the robots at the start and goal positions are pair-wise 2 units apart, or 4 unit apart from center to center

Unlabeled discs in the plane: the problem

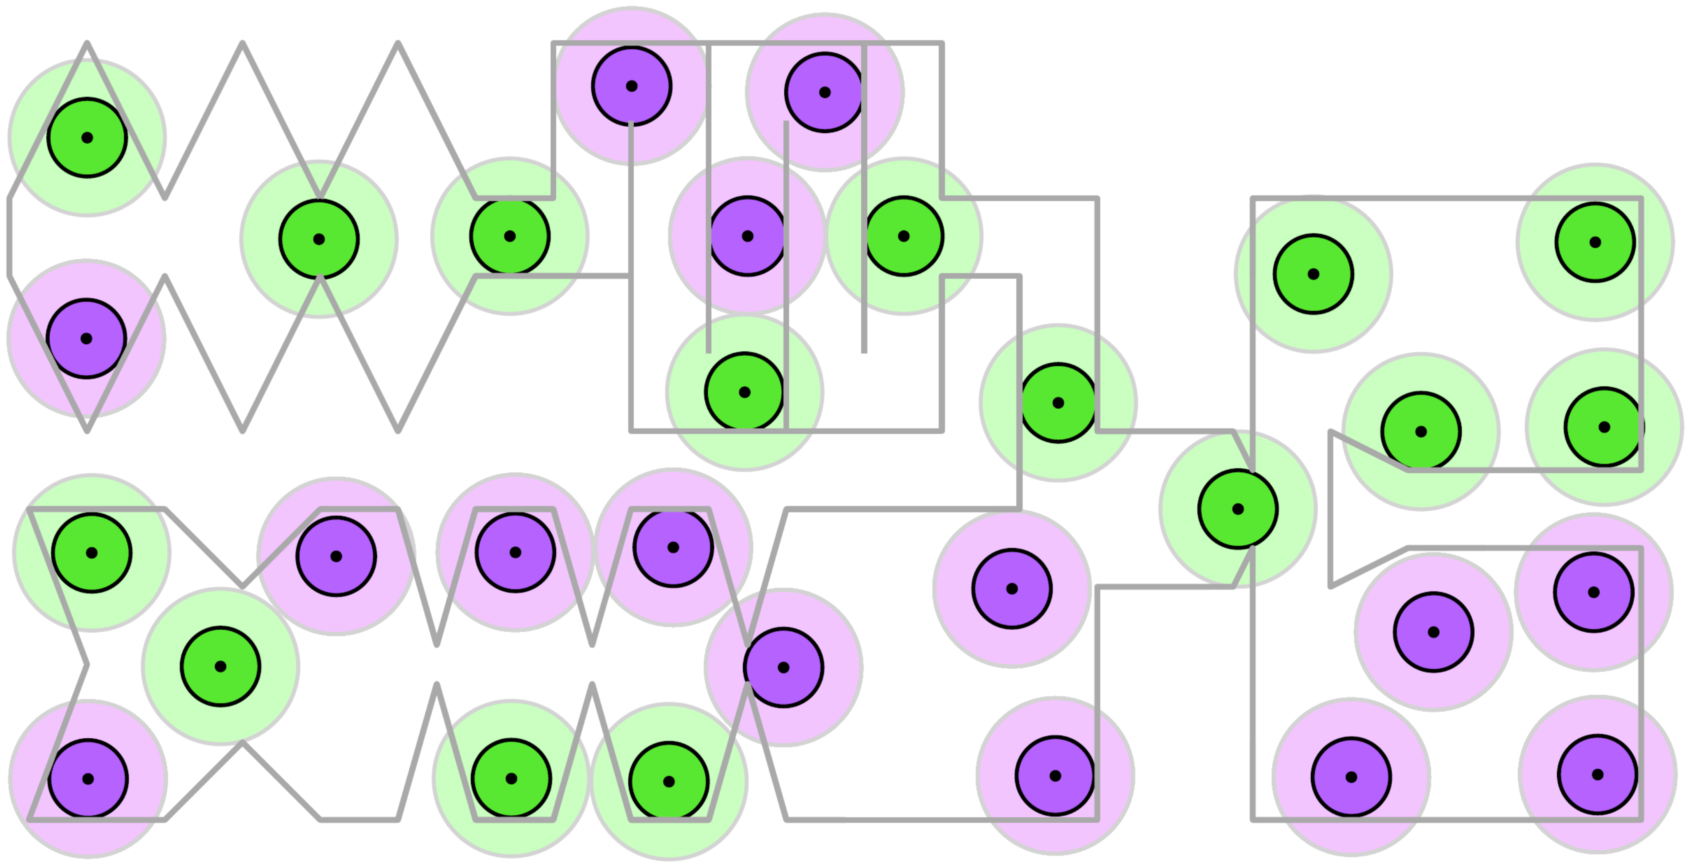

Unlabeled discs in the plane: the solution

A complete combinatorial algorithm running in $O(n \log n + mn + m^2)$ time, m is the number of robots and n is the complexity of the polygon

[Adler-de Berg-H-Solovey, WAFR 2014, IEEE T-ASE 2015]

Unlabeled discs in the plane: the solution

A complete combinatorial algorithm running in $O(n \log n + mn + m^2)$ time, m is the number of robots and n is the complexity of the polygon

F is the free space of a single robot, F = $\cup_i F_i$

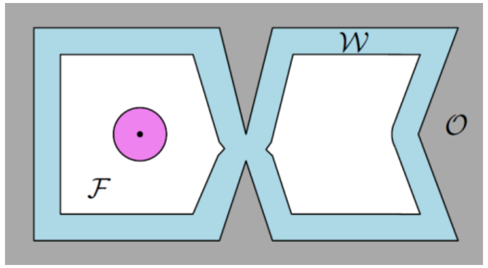

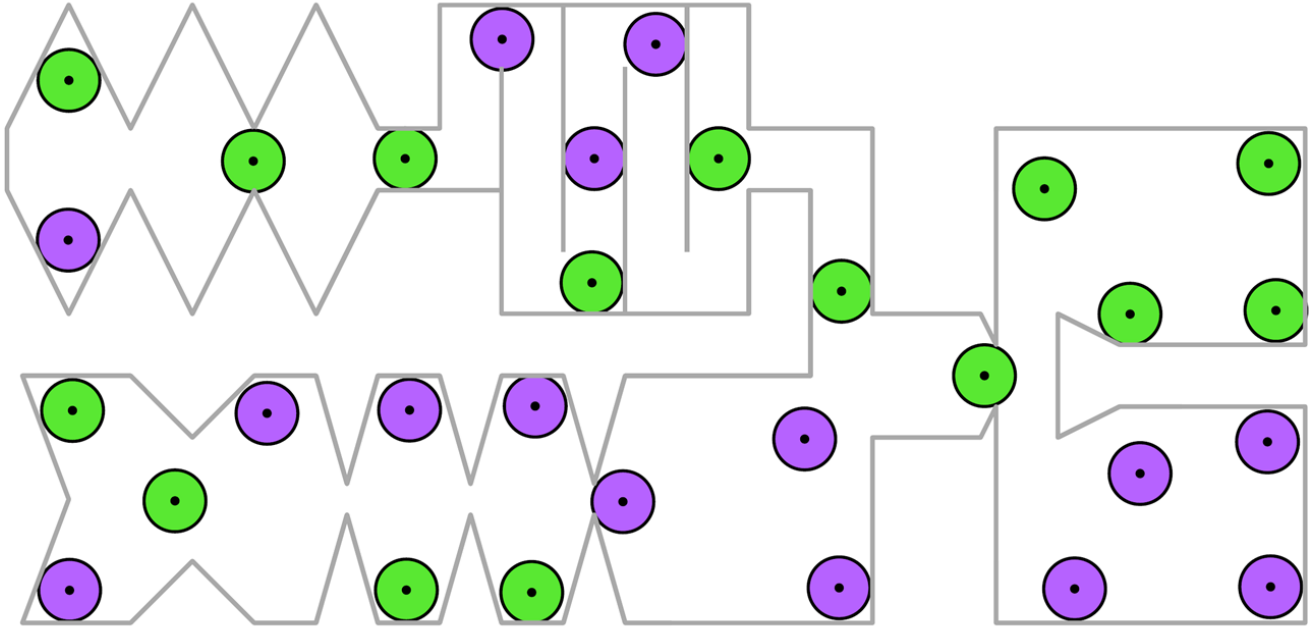

[Adler-de Berg-H-Solovey, WAFR 2014, IEEE T-ASE 2015]

Unlabeled discs in the plane: behind the scenes

- nice behavior in a single connected component of F

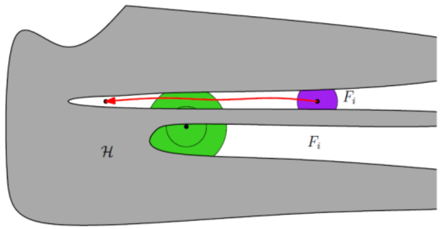

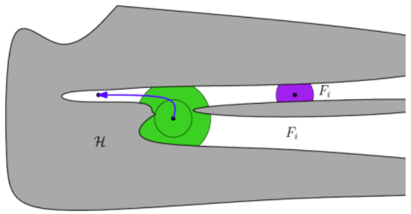

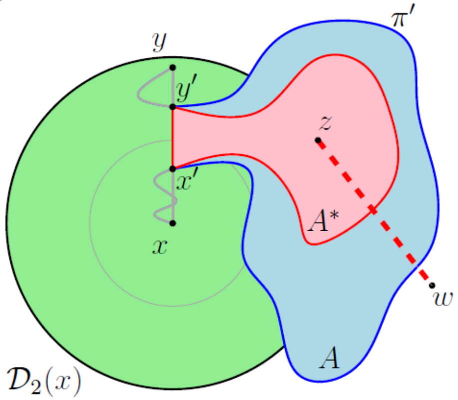

- impossibility of cycle of effects between connected components >> topological order of handling components

Unlabeled discs in the plane: why is it (so) easy?

- ☐ because the workspace is homeomorphic to a disc?
- ☐ because it is the unlabeled variant?
- ☐ because the robots are so simple?
- ☐ because of the separation assumption?

☐ Because the workspace is homeomorphic to a disc?

NO

Motion planning for discs
in a simple polygon is
NP-hard [Spirakis-Yap 1984]

Reduction from the strong NP-C 3-partition

Labeled, different radii

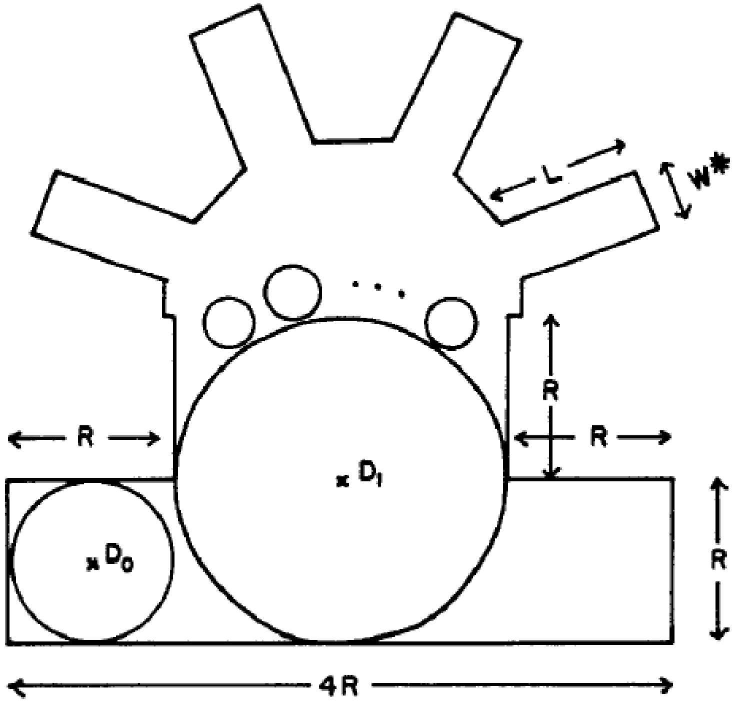

☐ Because it is the unlabeled variant?

NO

Motion planning for unlabeled unit squares in the plane is PSPACE-hard

[Solovey-H RSS 2015 best student paper award, IJRR 2016]

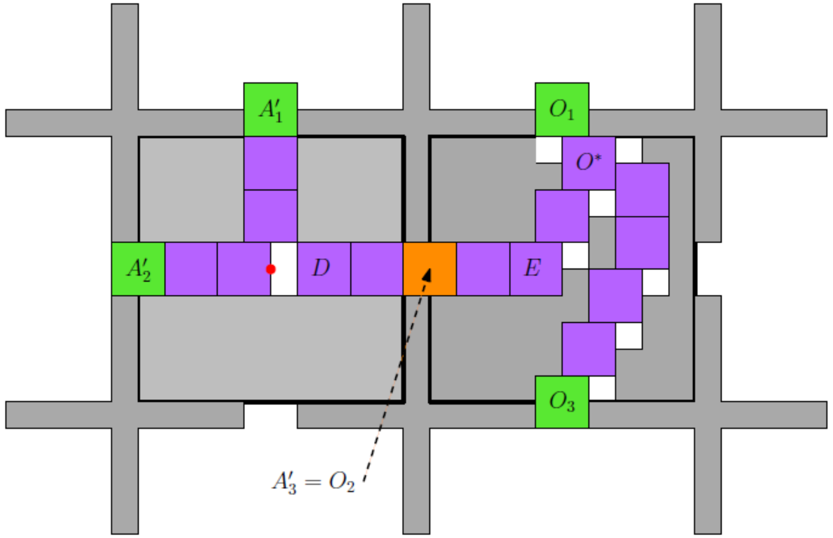

PSPACE-hardness, cont'd

- the first hardness result for unlabeled motion planning
- applies as well to labeled motion planning: the first multi-robot hardness result that uses only one type of robot geometry
- four variants, including "move any robot to a single target"

[Solovey-H RSS 2015 best student paper, IJRR 2016]

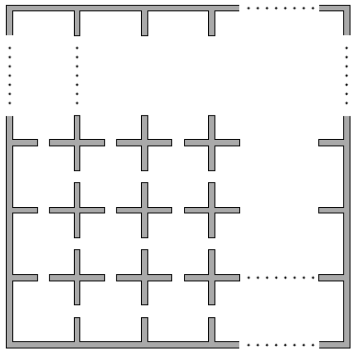

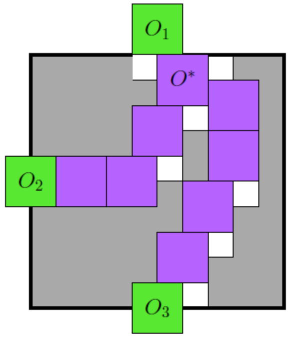

side note

a powerful gem:

PSPACE-Completeness of Sliding-Block Puzzles and other Problems through the Nondeterministic Constraint Logic Model of Computation

[Hearn and Demaine 2005]

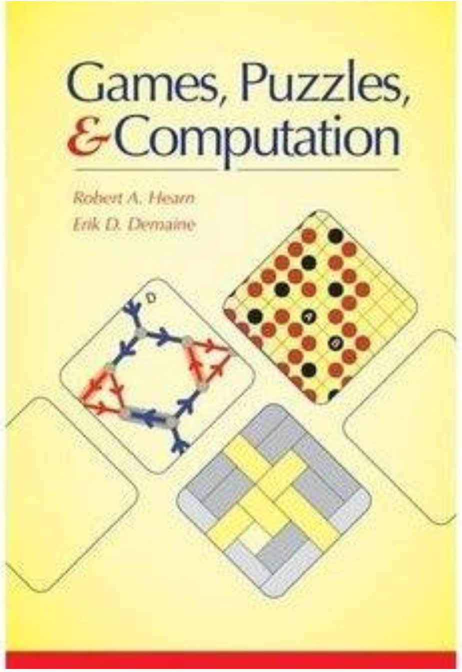

☐ Because the robots are so simple?

NO

Motion planning for unlabeled unit squares in the plane is PSPACE-hard

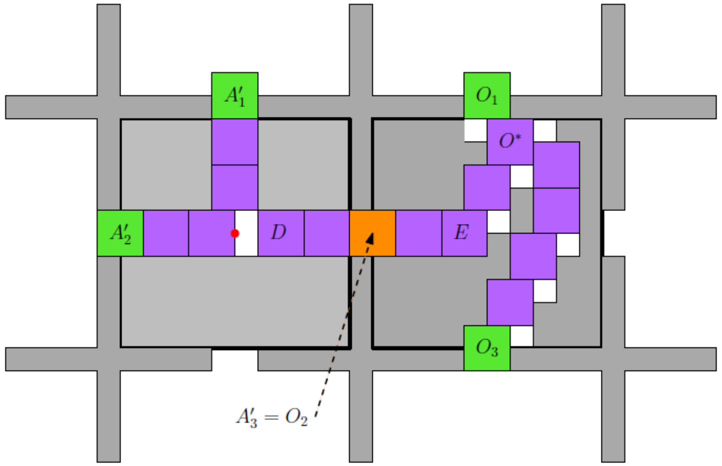

☐ Because of the separation assumption?

YES

- Recall that
  - the separation relates to two static configurations and not to a full path
  - no clearance from the obstacles is required

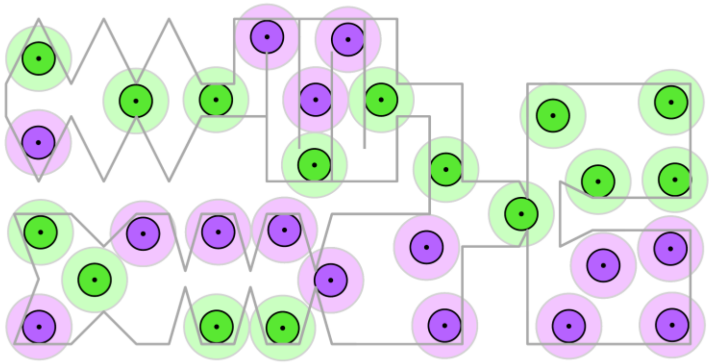

An exercise in separation

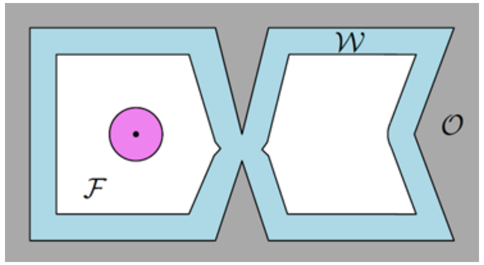

- a side effect of the analysis [Adler et al] is a simple decision procedure: there is a solution iff in each F_i (connected component of the free space) there is an equal number of start and goal positions

- Q: what is the minimum separation distance λ that guarantees a solution?
- A: $4\sqrt{2}-2$ (≈3.646) ≤ λ ≤ 4

  [Adler-de Berg-H-Solovey, T-ASE 2015]
- new A: λ = 4

  [Bringmann, 2018]

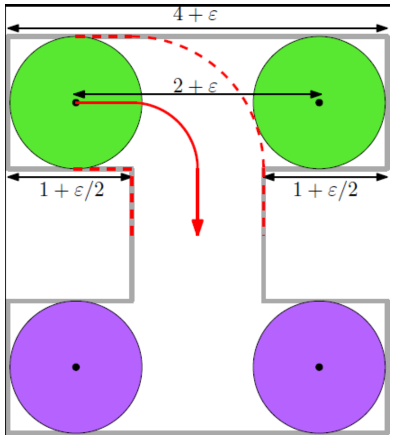

Challenges

- Q I: Does the unlabeled hardness proof still hold for unit discs (instead of unit squares)?
- Q II: Is it possible to solve the problem with separation 2+epsilon in time polynomial in m,n, and 1/epsilon?

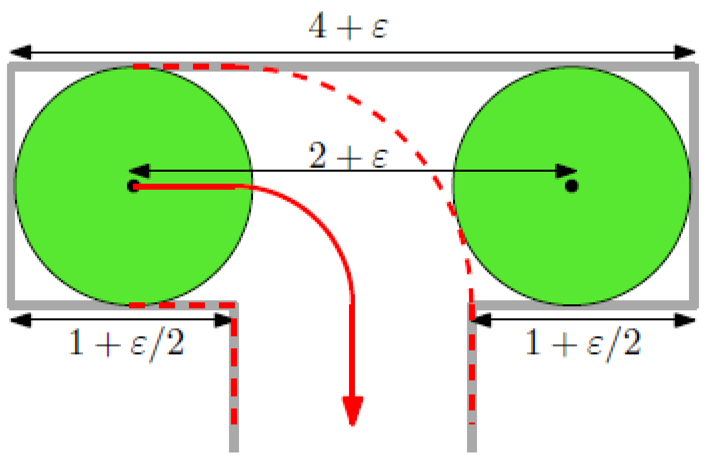

Multi-robot planning in tight settings

Compactifying a multi-robot packaging station

- Before: disjoint workspaces

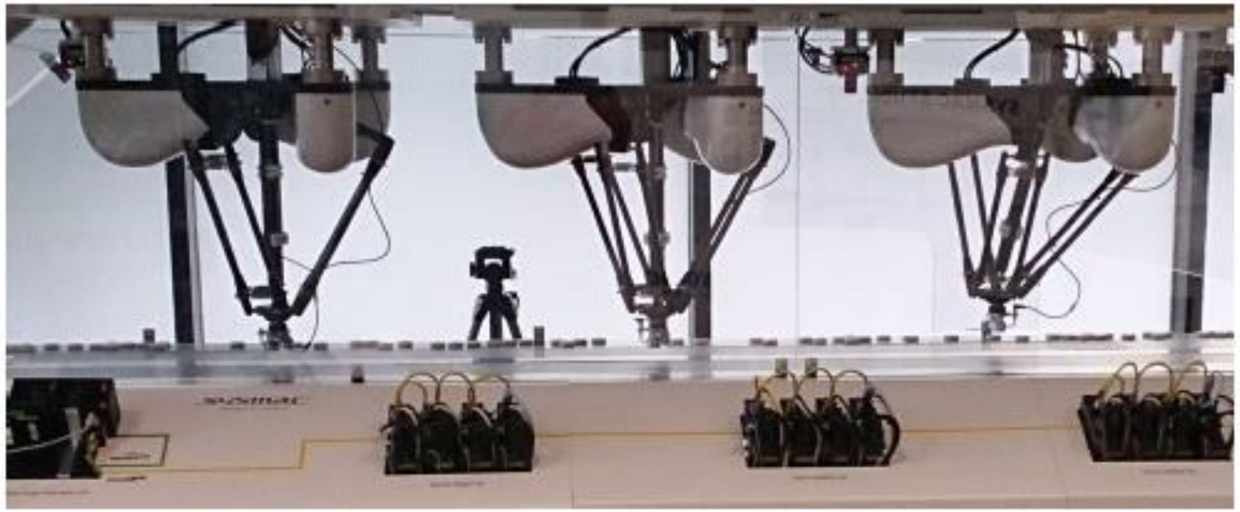

- After: overlapping workspaces
- Real-time collision detection [van Zon et al CASE 2015]

Multi robot, complex settings

- Common belief: view as a compound robot with many dofs and use single-robot sampling-based planning to solve coordinated motion problems

modest roadmap with 1K nodes per robot means tensor product for 6 robots with quintillion nodes dRRT, slides by Kiril Solovey ,5-13

Complex multi-robot settings

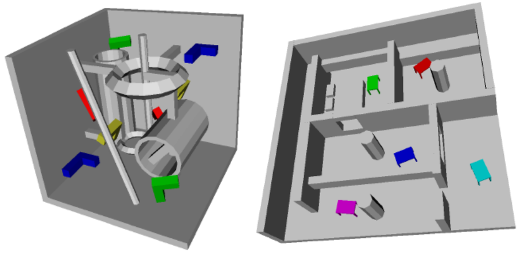

- Discrete RRT (dRRT)

[Solovey-Salzman-H WAFR 2014, IJRR 2016]

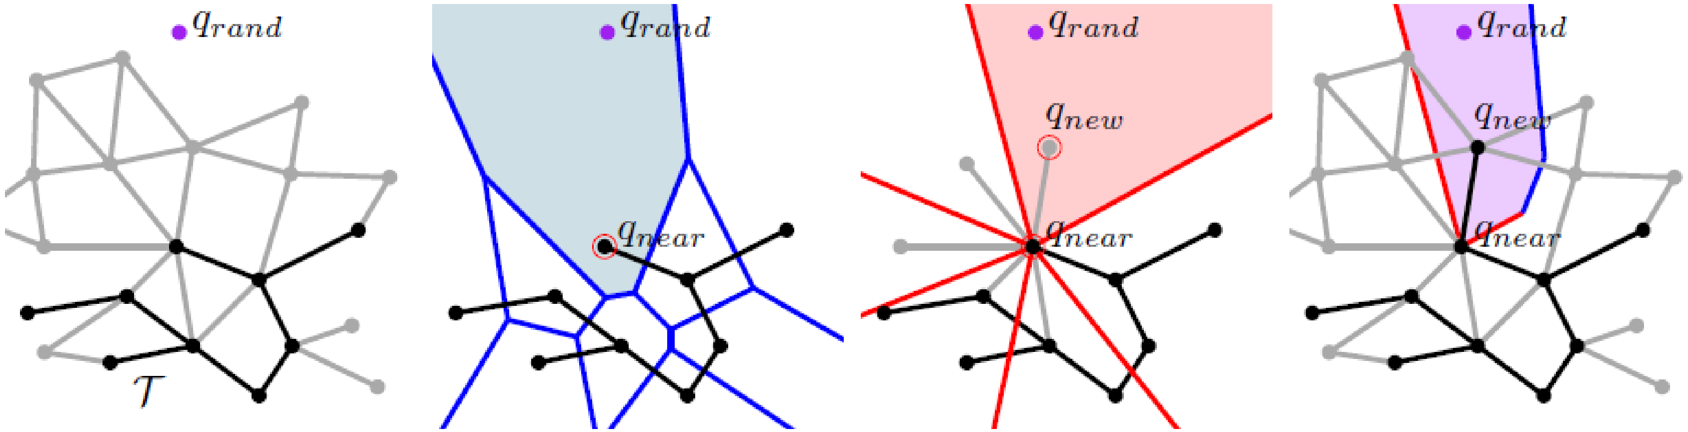

[probabilistic completeness]

- M*

[Wagner-Choset IROS 2010, AI 2015]

Complex multi-robot settings, cont'd

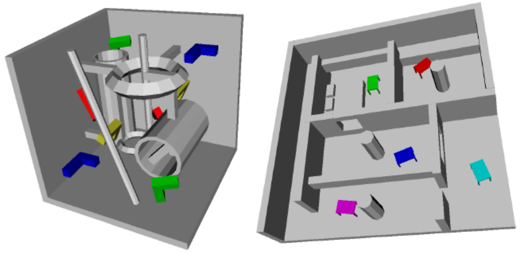

dRRT*

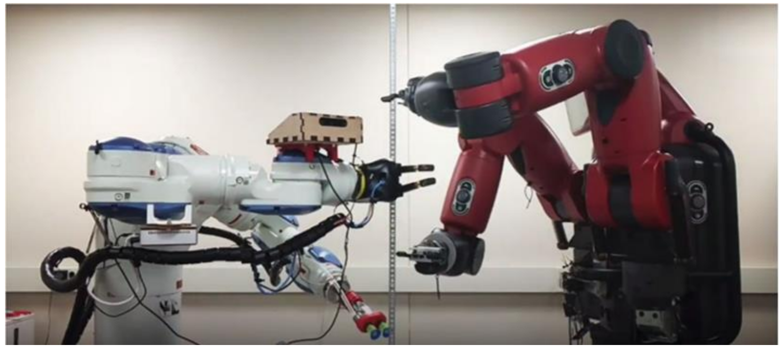

- Asymptotically optimal [KF11] version of dRRT

[Dobson et al, MRS 2017, best paper award]

- Applied for dual-arm object re-arrangement

[Shome et al, 2018]

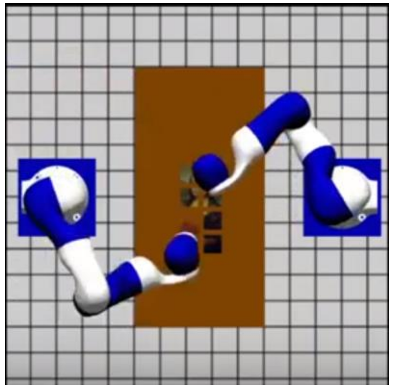

clip72 > sec 37

Side note
Effective metrics for multi-robot motion-planning

- When are two multi-robot configurations close by?
- Metric is key to guaranteeing probabilistic completeness and asymptotic optimality

- Novel metrics tailored to multi-robot planning
- Tools to assess the efficacy of metrics

[Atias-Solovey-H RSS 2017, IJRR 2018]

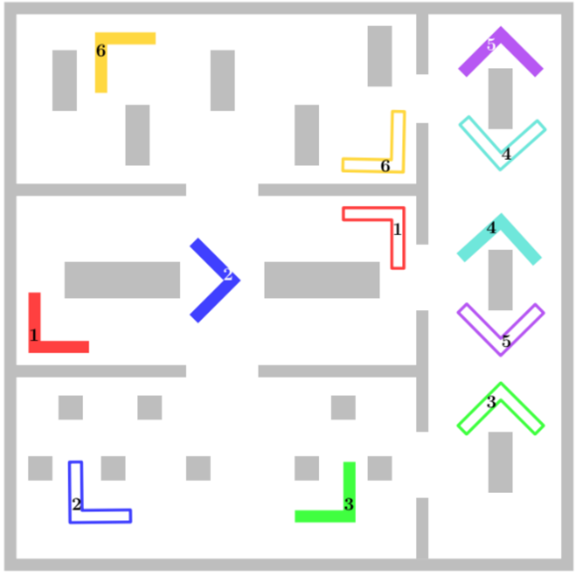

Multiple unit balls in R^d

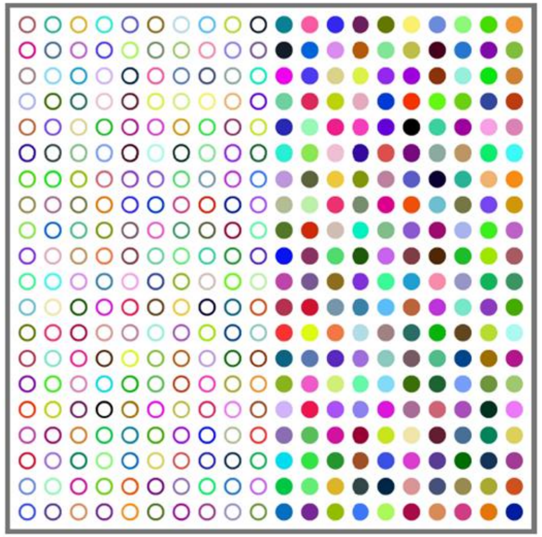

- Fully colored, decoupled (prioritized)
- Revolving areas with non-trivial separation

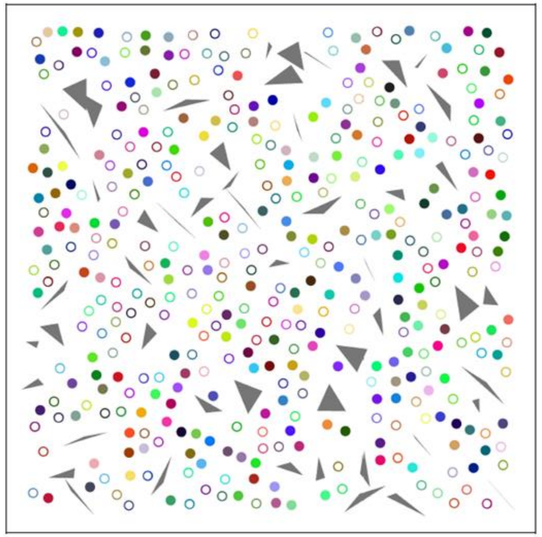

- Handling hundreds of discs in seconds, CGAL
- Finding the optimal order of execution in decoupled algorithms that locally solve interferences is NP-hard

[Solomon-H WAFR 2018]

clip18

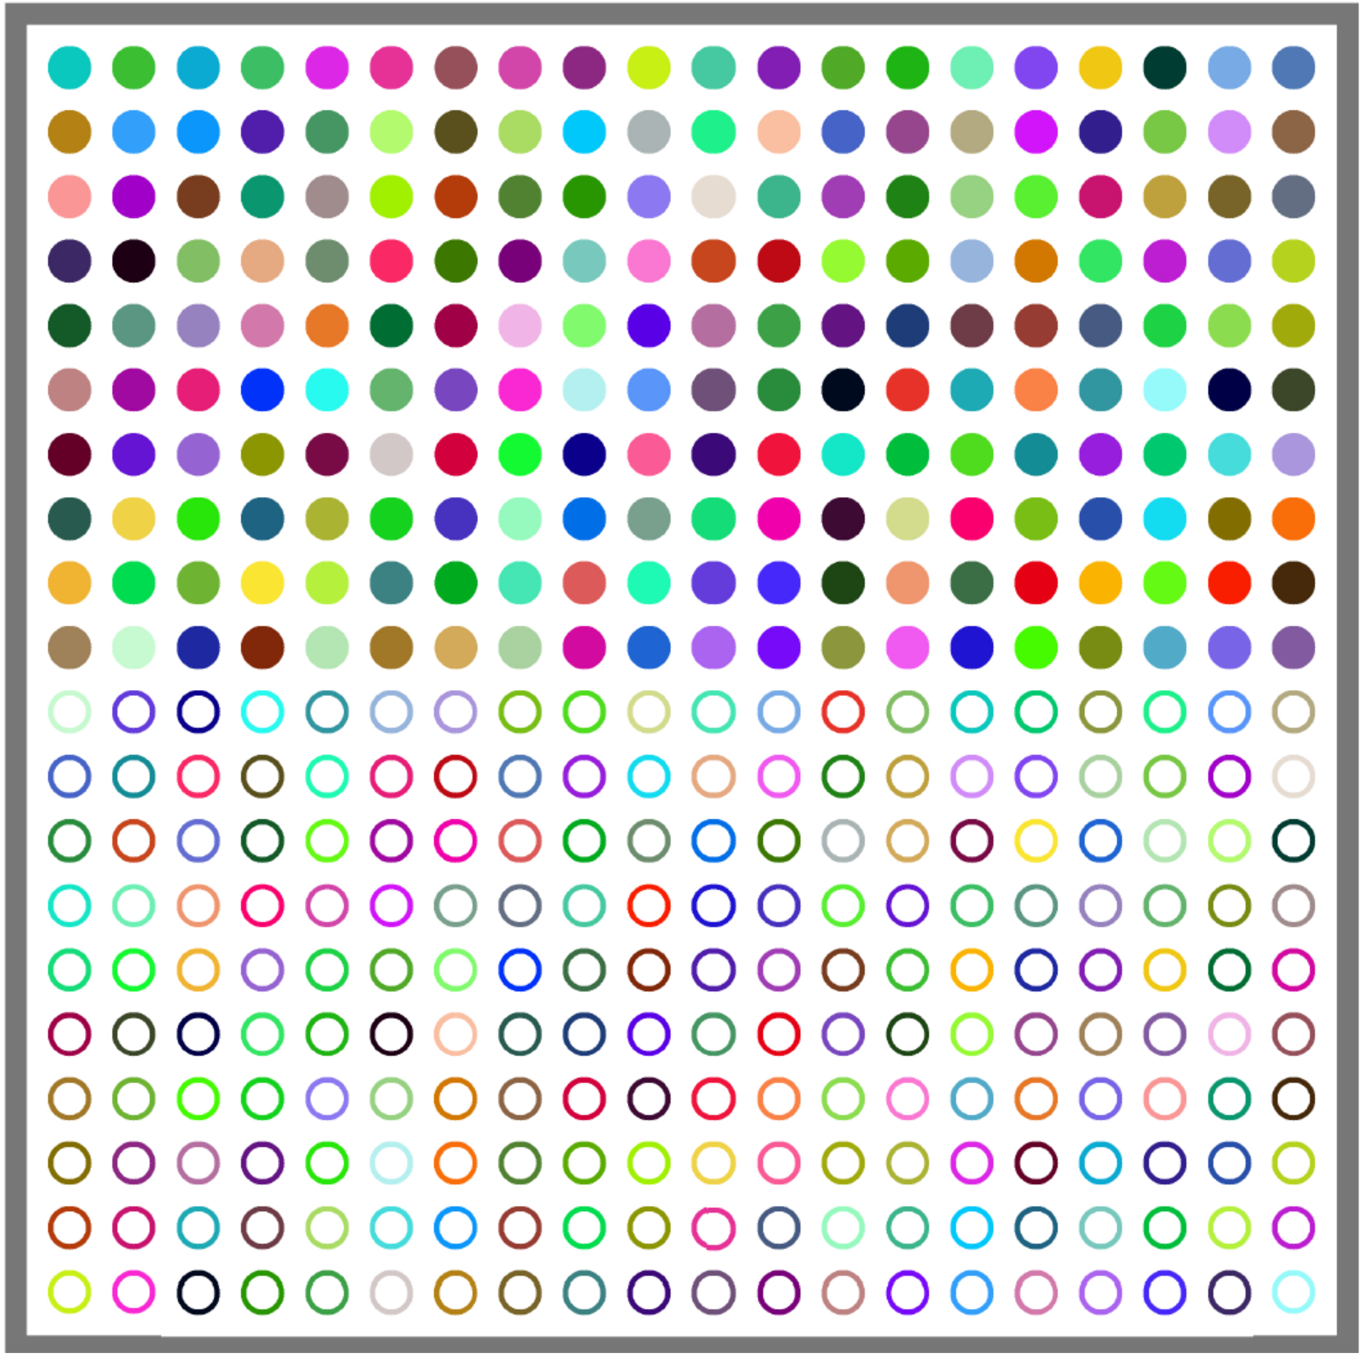

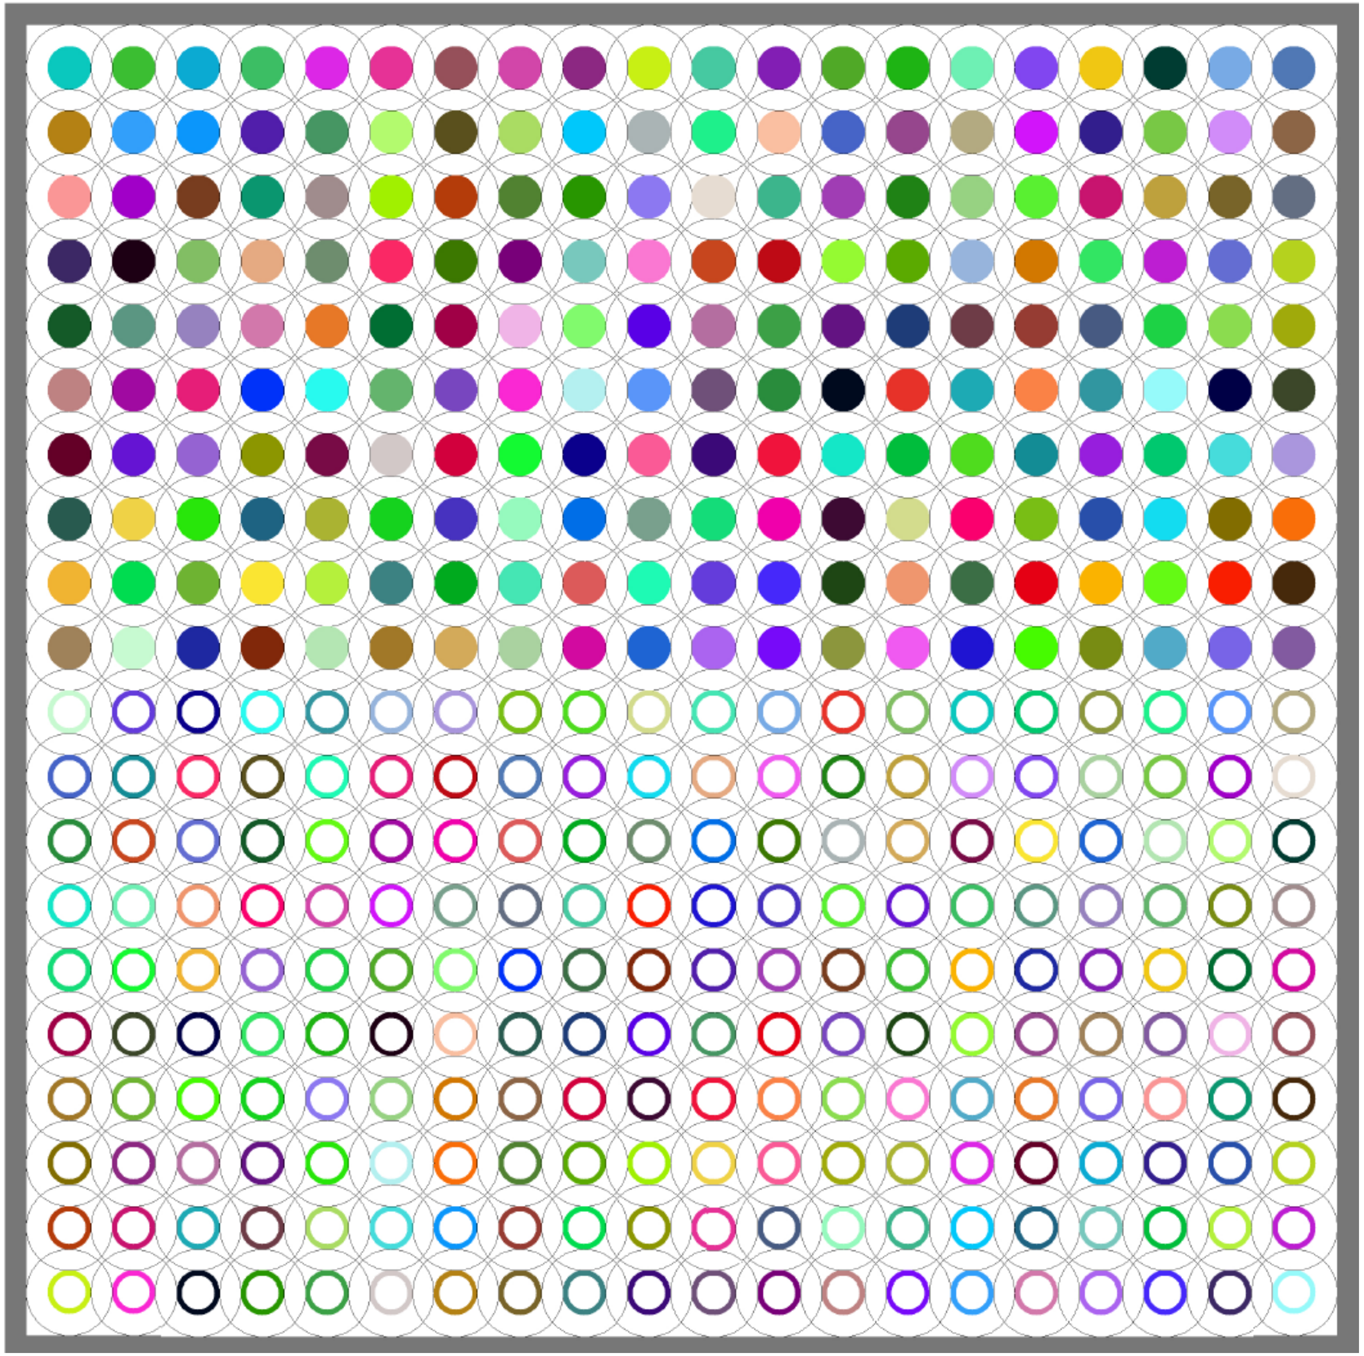

Optimality guarantees in unlabeled multi-robot planning

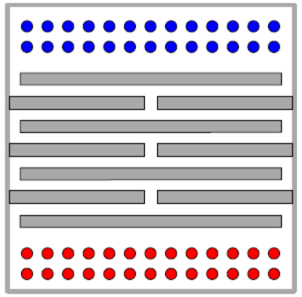

- Each result requires some extra separation and other conditions
- [Turpin-Mohta-Michael-Kumar AR 2014]:

optimizing min-max

- [Solovey-Yu-Zamir-H RSS 2015]:

optimizing total travel, approx.

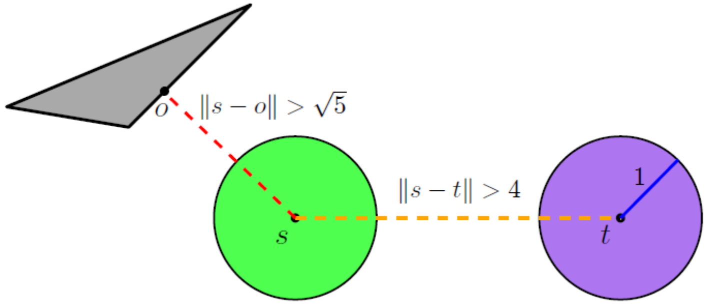

assuming 4 separation as before

and minimum distance of start/goal to obstacles

- discrete version pebble problems on graphs [Yu and LaValle]

Optimizing total travel in unlabeled multi-robot planning, cont'd

- full fledged exact implementation using CGAL for free space computation: arrangements, Minkowski sums, point location, etc.

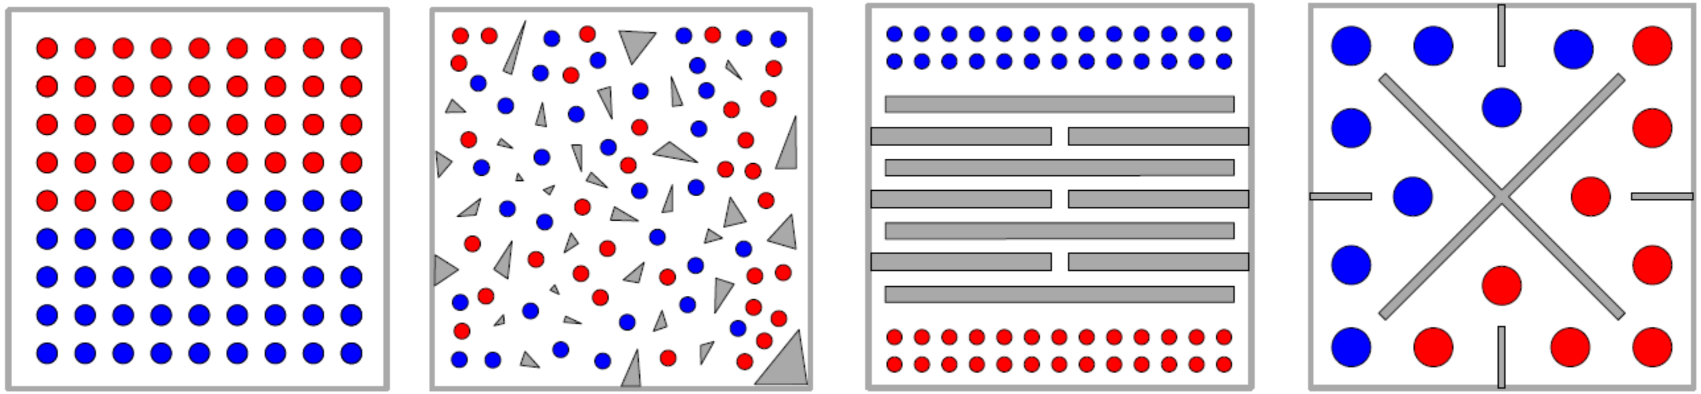

[Solovey-Yu-Zamir-H RSS 2015]

Multi-robot?
How about two robots?

Coordinating the motion of two discs in the plane

- Problem: Given two (unit) discs moving in the plane among polygonal obstacles, plan a joint free motion from start to goal of minimum total path length

- Efficient algorithm?
- Hardness?

Coordinating the motion of two discs in the plane, cont'd

- Characterization of optimal paths in the absence of obstacles (Reeds-Shepp style) [Kirkpatrick-Liu 2016]: at most six [straight,circular arc] segments
- Adaptation to translating squares [H-Ruiz-Sacristan-Silveira 2019]

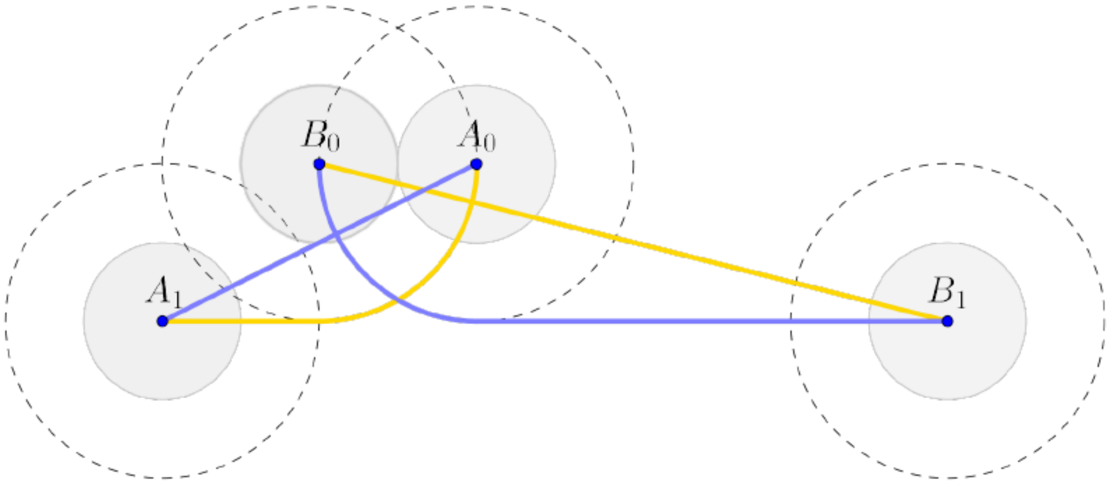

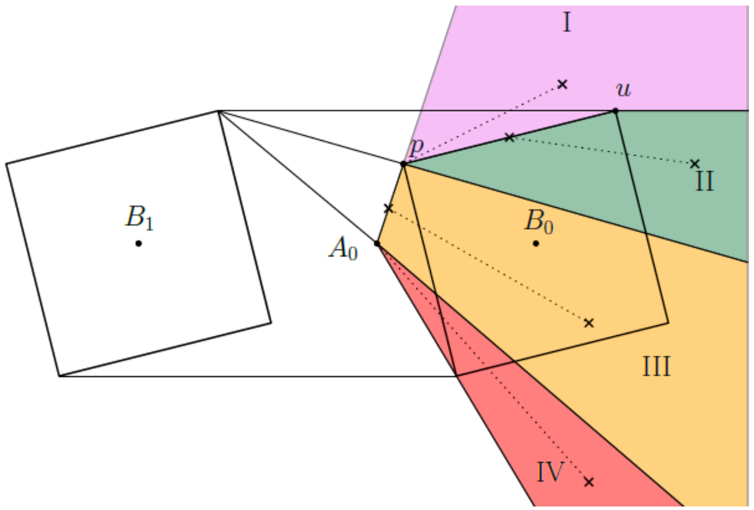

Rigid motion of two polygons: The limits of sampling-based planning

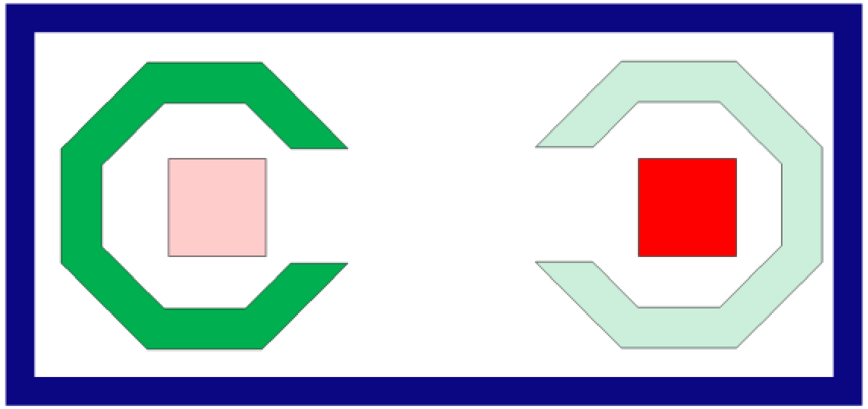

[Salzman-Hemmer-H TASE 2015]

- Each robot translates and rotates: system w/ 6 dofs
- Start position in bright colors, goal in pale colors
- Pacman needs to swallow the square before rotating to target

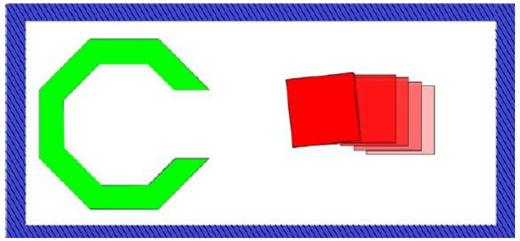

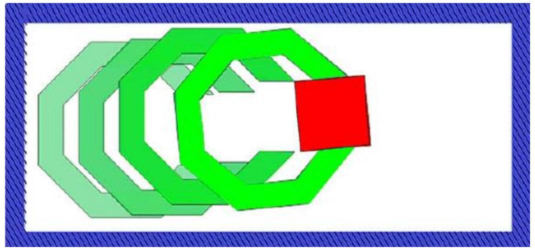

Rigid motion of two polygons, cont'd

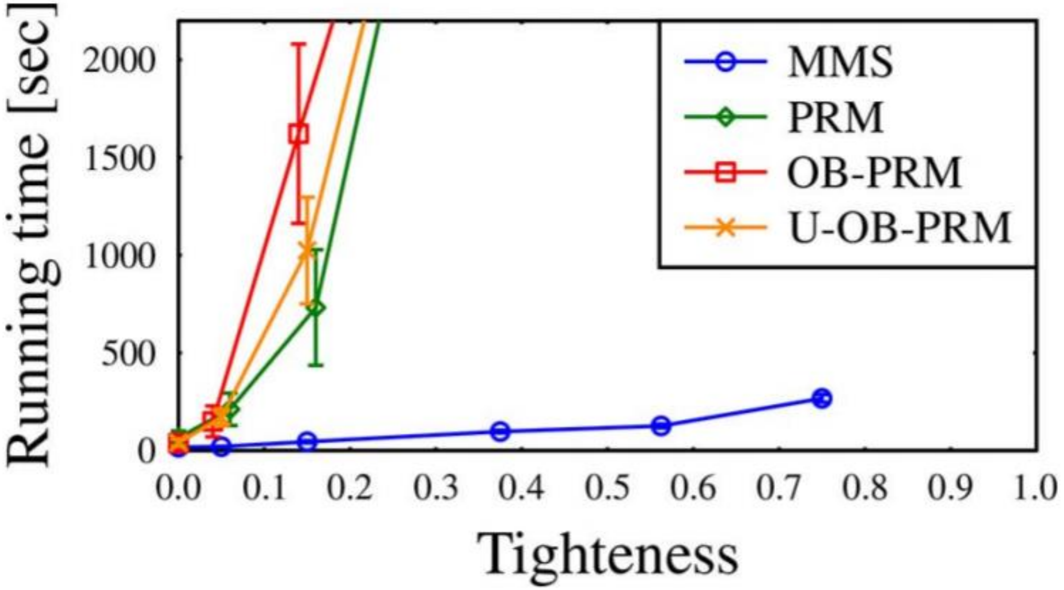

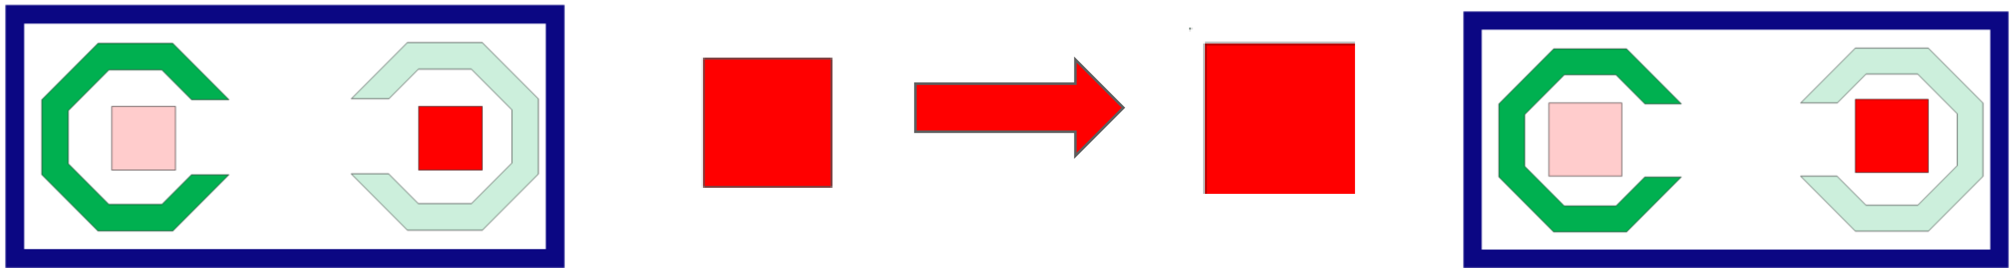

MMS: Motion planning via manifold samples

[Salzman-Hemmer-Raveh-H Algorithmica 2013]

Example: polygon translating and rotating among polygons

- sampling the 3D configuration space by strong geometric primitives, including exact arrangements of curves CGAL

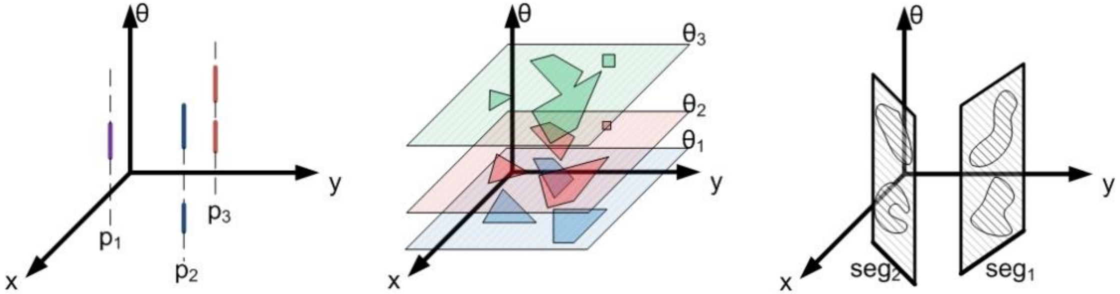

- combinatorial analysis of primitives yields *free space cells*
- path planning by intersecting free space cells

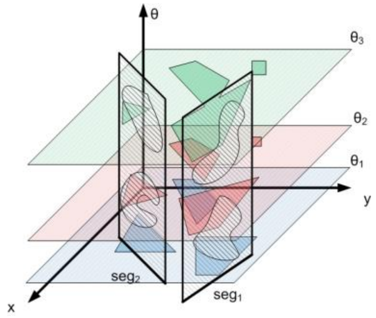

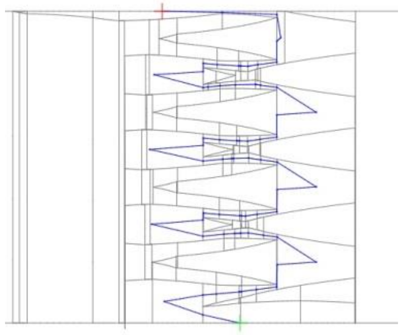

side note
k-handed assembly planning and multi-robot

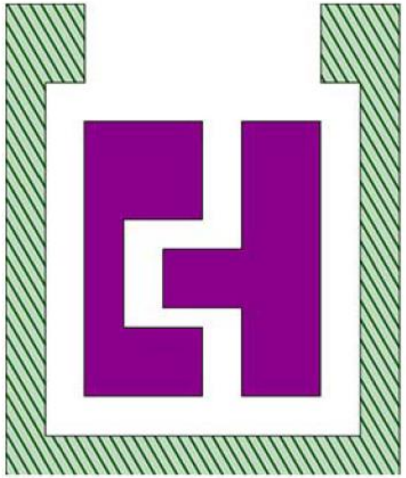

[Salzman-Hemmer-H]

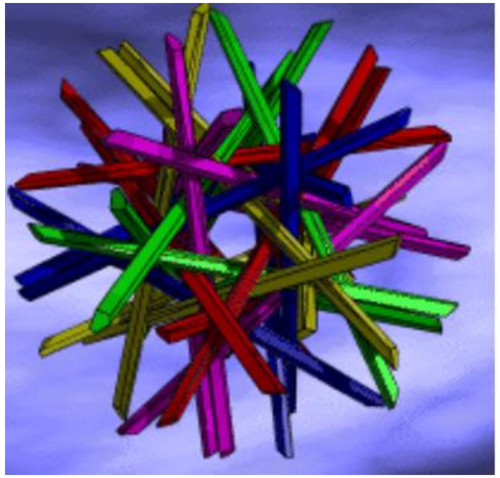

[Snoeyink-Stolfi]

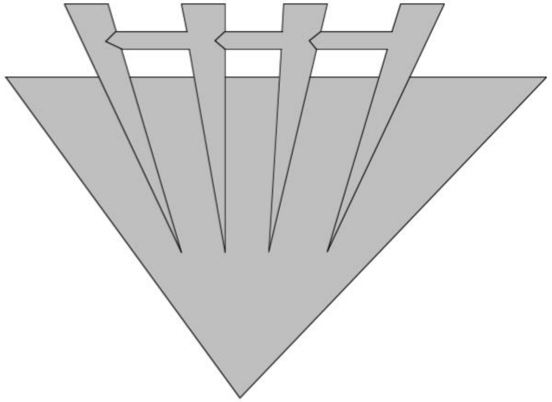

[Natarajan/Wilson]

Summary and outlook

Tools for MRMP

- Multi two-dimensional robots, with separation: complete deterministic algorithms, CGAL
- Complex robot, complex environment: sampling based planners, probabilistic completeness, asymptotic optimality, OMPL
- Multi complex robots: sampling based planners, probabilistic completeness, asymptotic optimality

Challenges

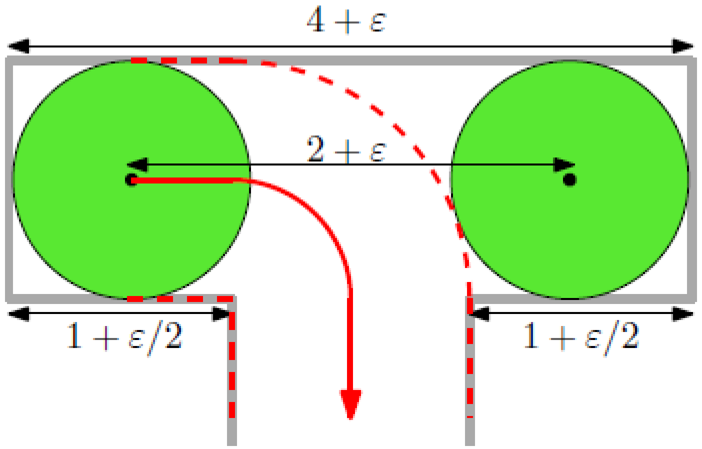

- Predictive analysis for finite time, which will interpolate between easy and hard
- Identifying the inherent difficulties in multi-robot motion planning
- Optimality!
- Assembly planning, k-handed

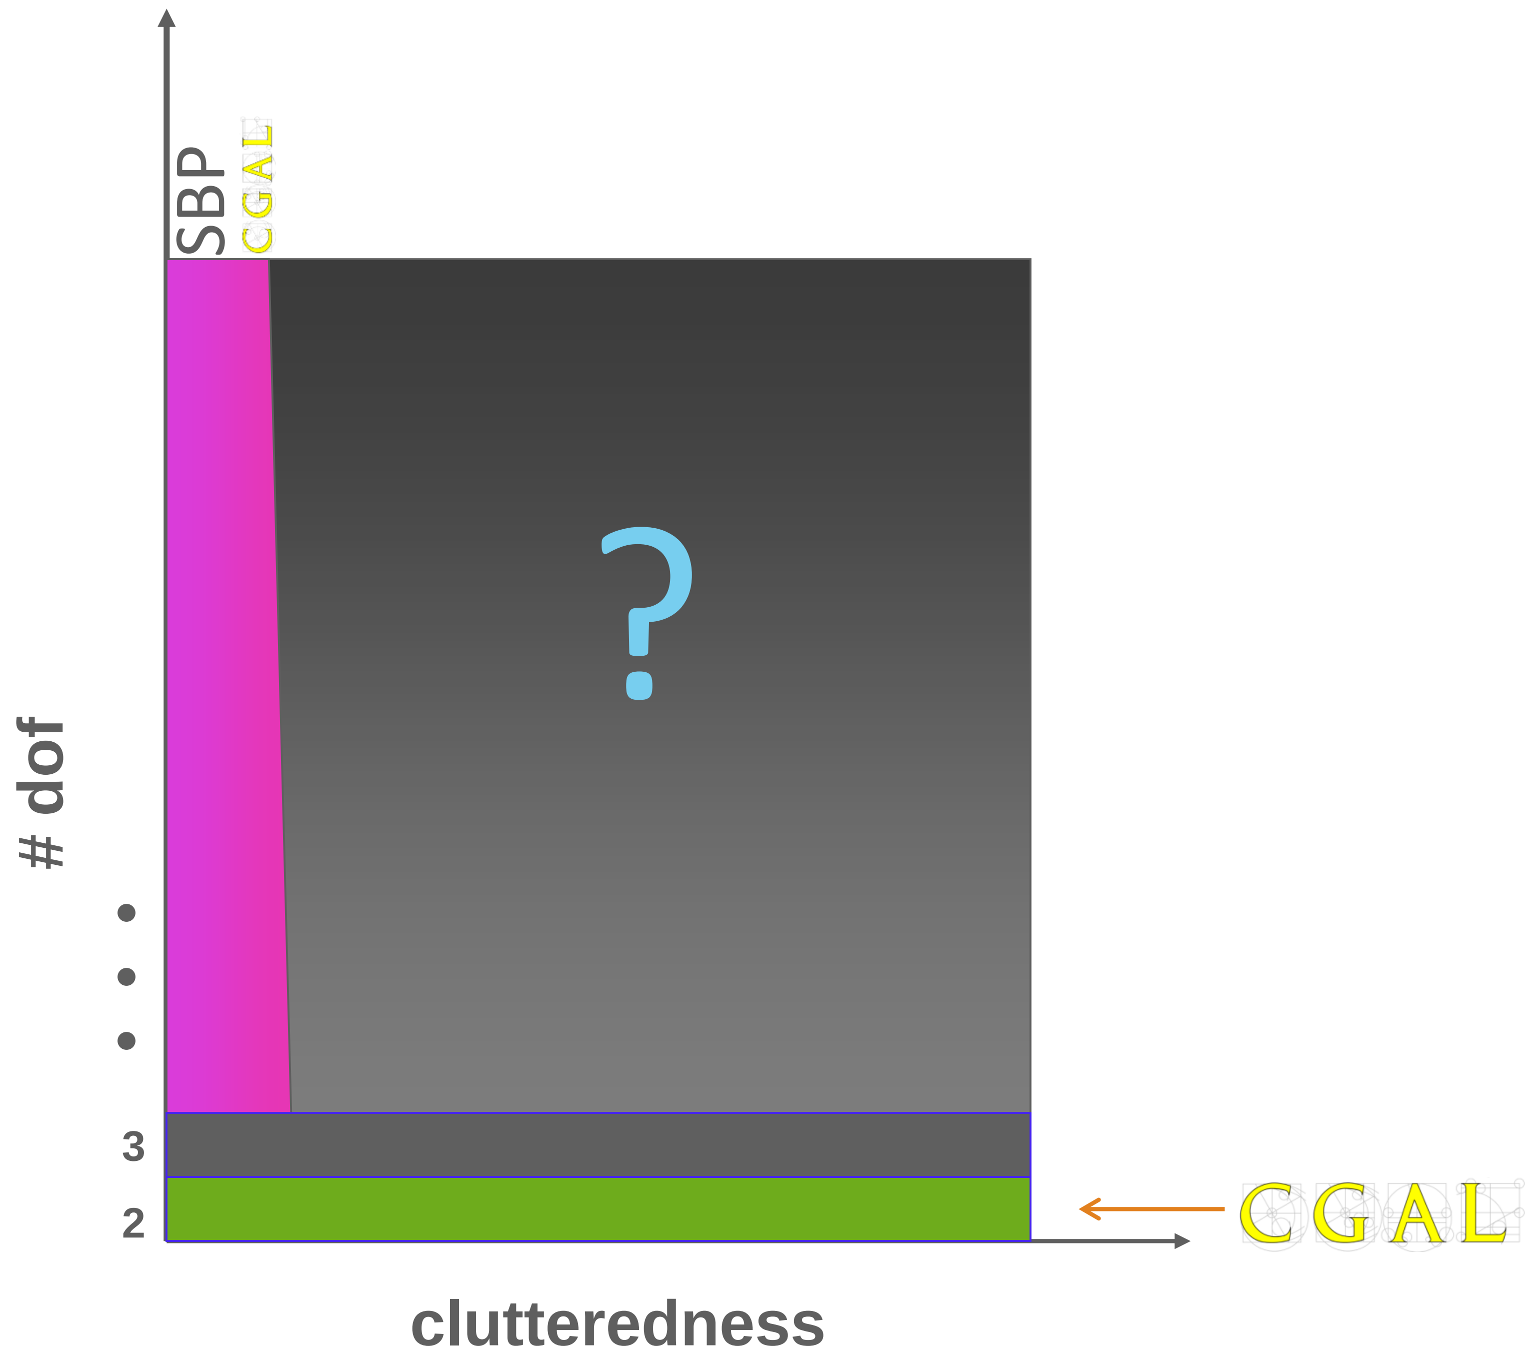
SBP
CGAL
?
dof
3
2
CGAL
clutteredness

References: SB planners for multi robot

- Petr Svestka, Mark H. Overmars: Coordinated path planning for multiple robots. Robotics and Autonomous Systems 23(3): 125-152 (1998)
- (**M***) Glenn Wagner, Howie Choset: Subdimensional expansion for multirobot path planning. Artif. Intell. 219: 1-24 (2015)
- (**dRRT**) Kiril Solovey, Oren Salzman, Dan Halperin: Finding a needle in an exponential haystack: Discrete RRT for exploration of implicit roadmaps in multi-robot motion planning. I. J. Robotics Res. 35(5): 501-513 (2016)
- Rahul Shome, Kiril Solovey, Andrew Dobson, Dan Halperin, Kostas E. Bekris: **dRRT***: Scalable and Informed Asymptotically-Optimal Multi-Robot Motion Planning. CoRR abs/1903.00994 (2019).Also in Autonomous Robots.

References, cont’d

- Aviel Atias, Kiril Solovey, Oren Salzman, Dan Halperin: Effective **metrics** for multi-robot motion-planning. I. J. Robotics Res. 37(13-14) (2018)

THE END