Automated Planning

Path Planning and Motion Planning: An Overview

Jonas Kvarnström
Department of Computer and Information Science
Linköping University

jonas.kvarnstrom@liu.se – 2019

Path/Motion Planning (1)

- Perhaps the *easiest* form of **path planning** / **motion planning**:
  - A robot should move in **two dimensions** between start and goal
  - Avoiding known **obstacles** – or it would be *too* easy…

Path/Motion Planning (2)

- Problem: Generating an **optimal continuous path** is hard!
  - First step (often): **Discretize**
    - Choose a finite number of **potential waypoints** in the map
    - Create a **graph**: Waypoints are **nodes**, short obstacle-free paths are **edges**
    - Use **discrete search algorithms** to decide which waypoints to use

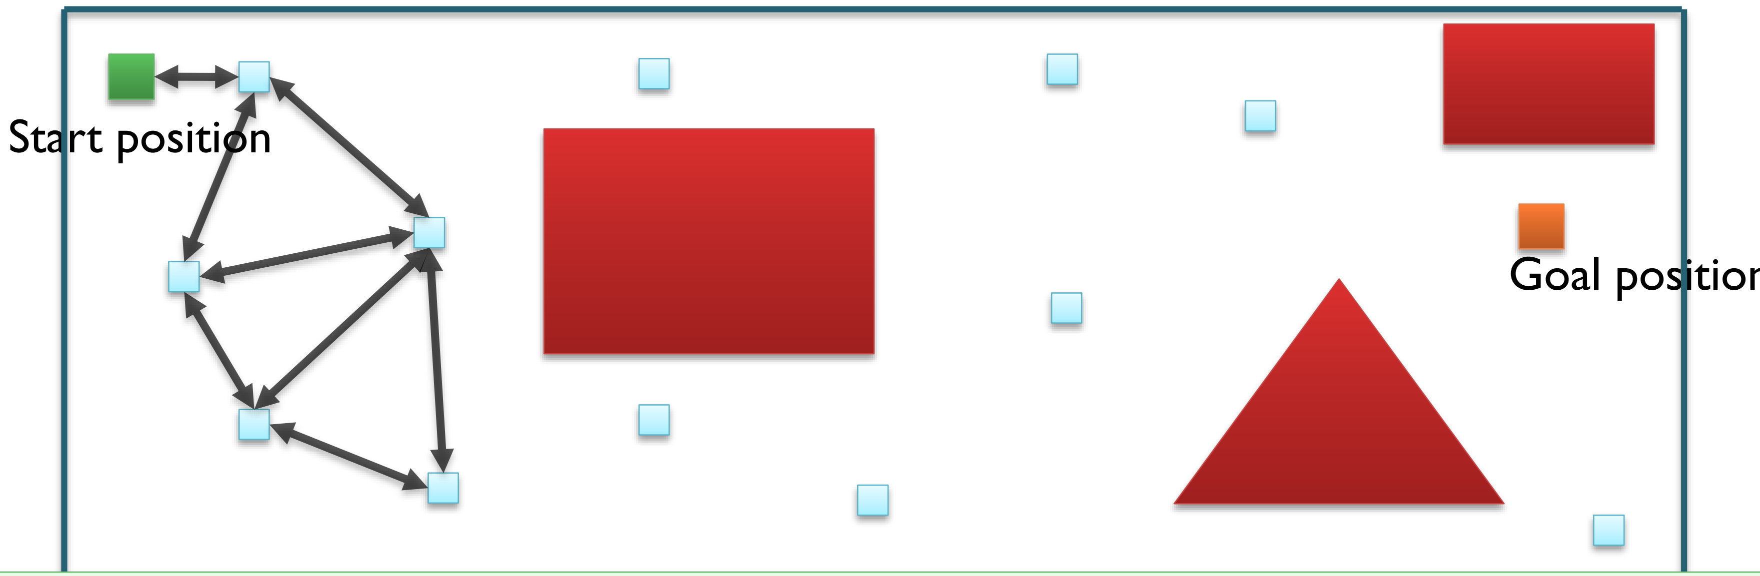

To do: create **nodes / potential waypoints**,
generate appropriate **edges**,
find a path in the resulting graph

Choosing Potential Waypoints: Grid-Based Methods

Regular 2D Grid

- The simplest type of discretization: A **regular grid**
  - Robots only **move** *north, east, south* or *west*
  - **Assumption**: Can deal with details (geometry / terrain) later…

Grid = rutnät, the whole thing
Cell = ruta, a single rectangle

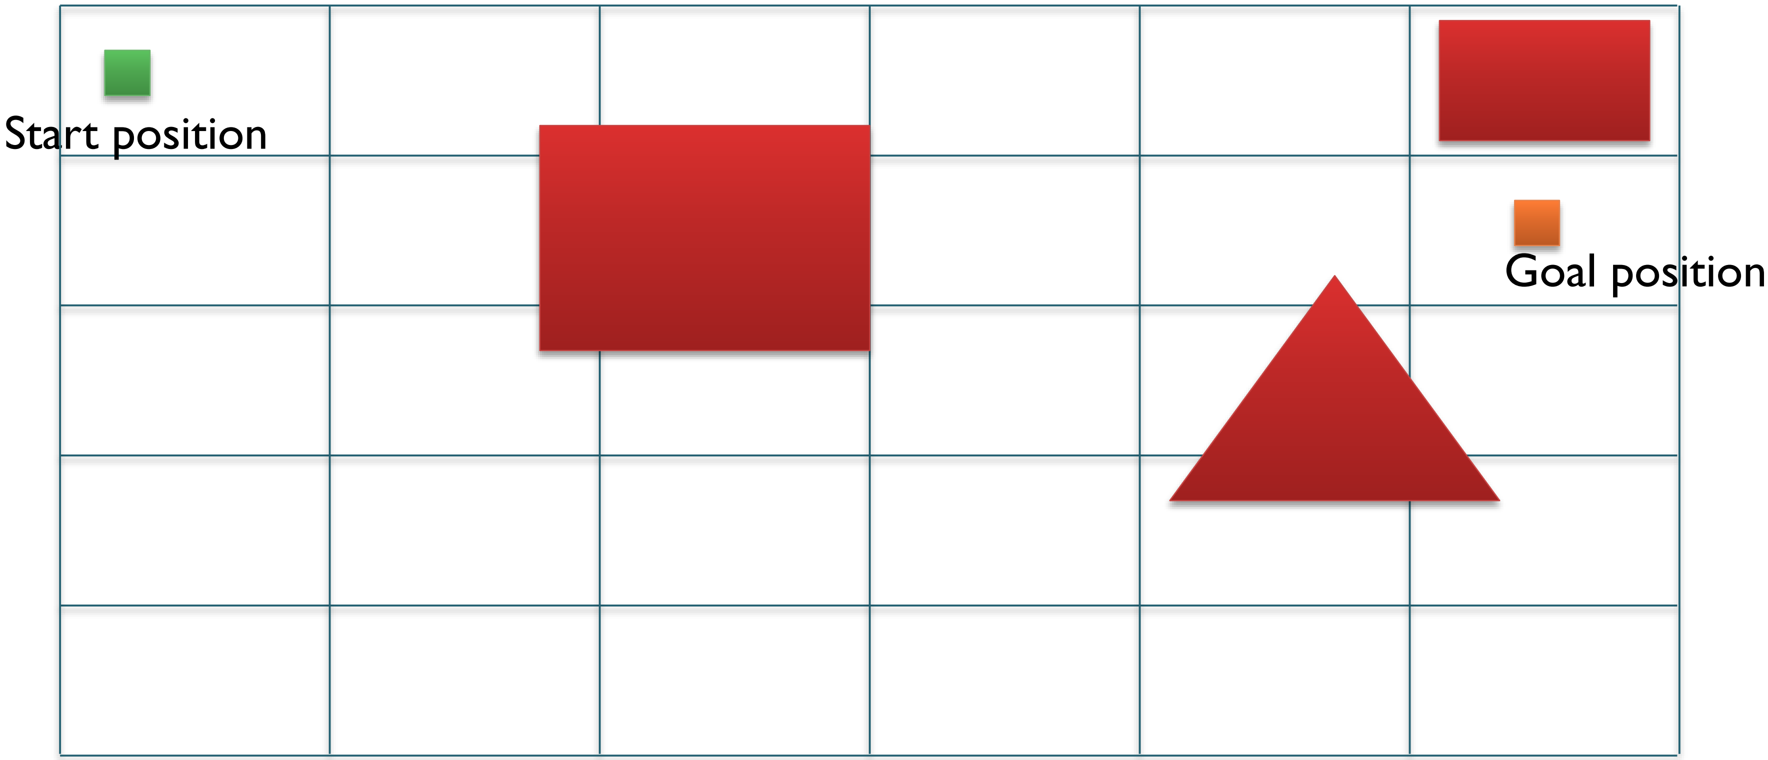

Regular 2D Grid: Real Obstacles

- Real obstacles do not correspond to square / rectangular cells…
  - But we can *cover* them with cells

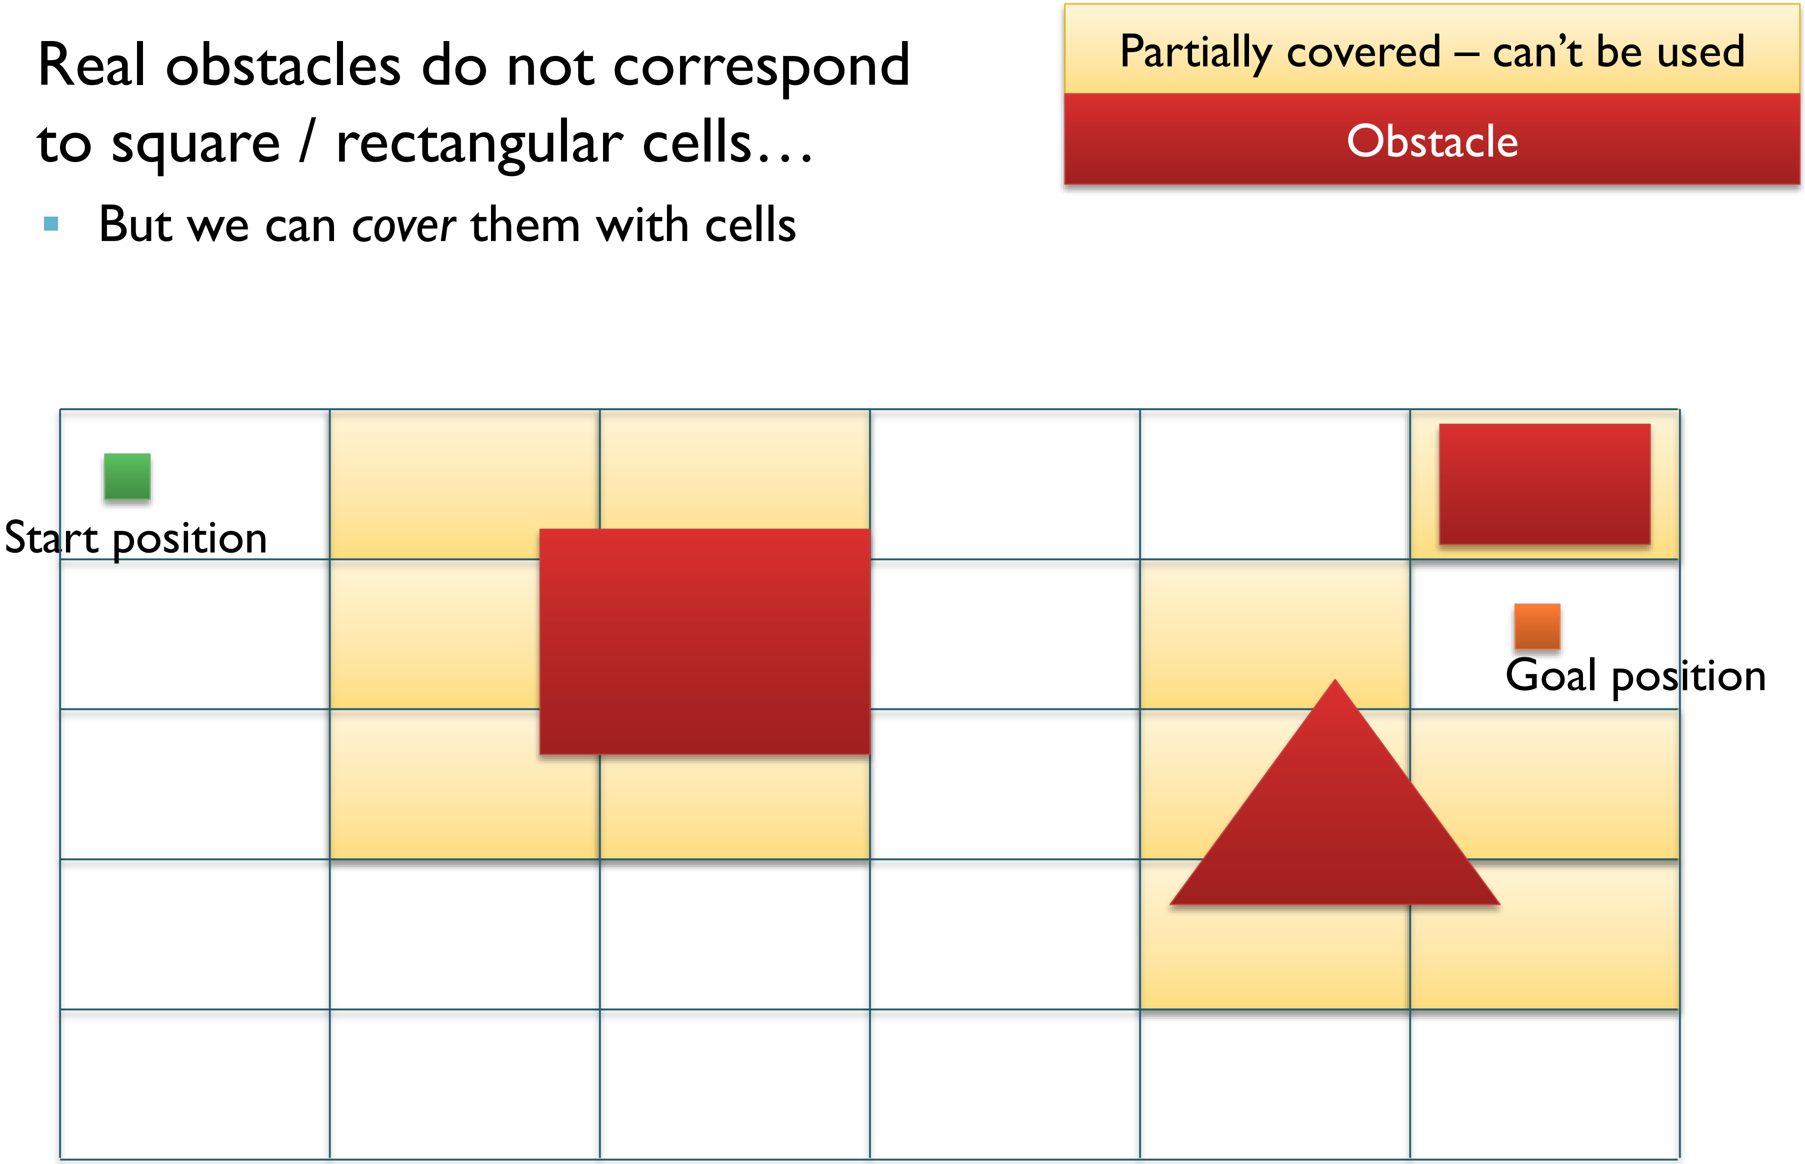

Regular 2D Grid: Nodes

- Each **cell** is associated with a single **node**
  - Corresponding to **2D point**
  - Could be the **center** of the 2D cell

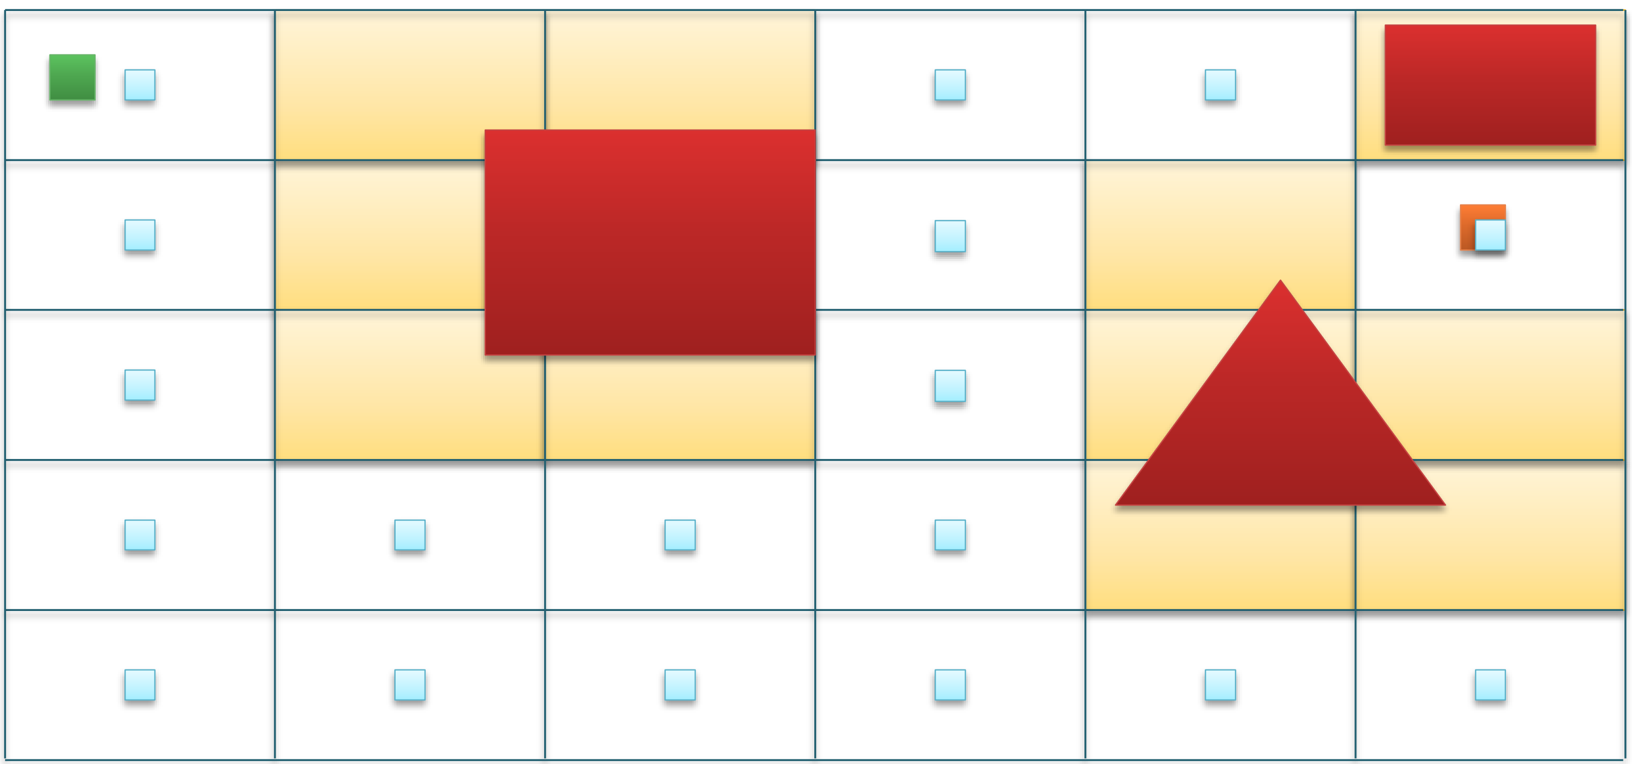

Regular 2D Grid: Edges

jonkv@ida

- Which nodes are **connected** in the discrete graph?
  - Let's **simplify** in the beginning
    - **Straight lines in 2D space**
    - Through **free** cells (completely grid-based, no complex geometry!)
    - **4-connectivity** (north, south, west, east)…

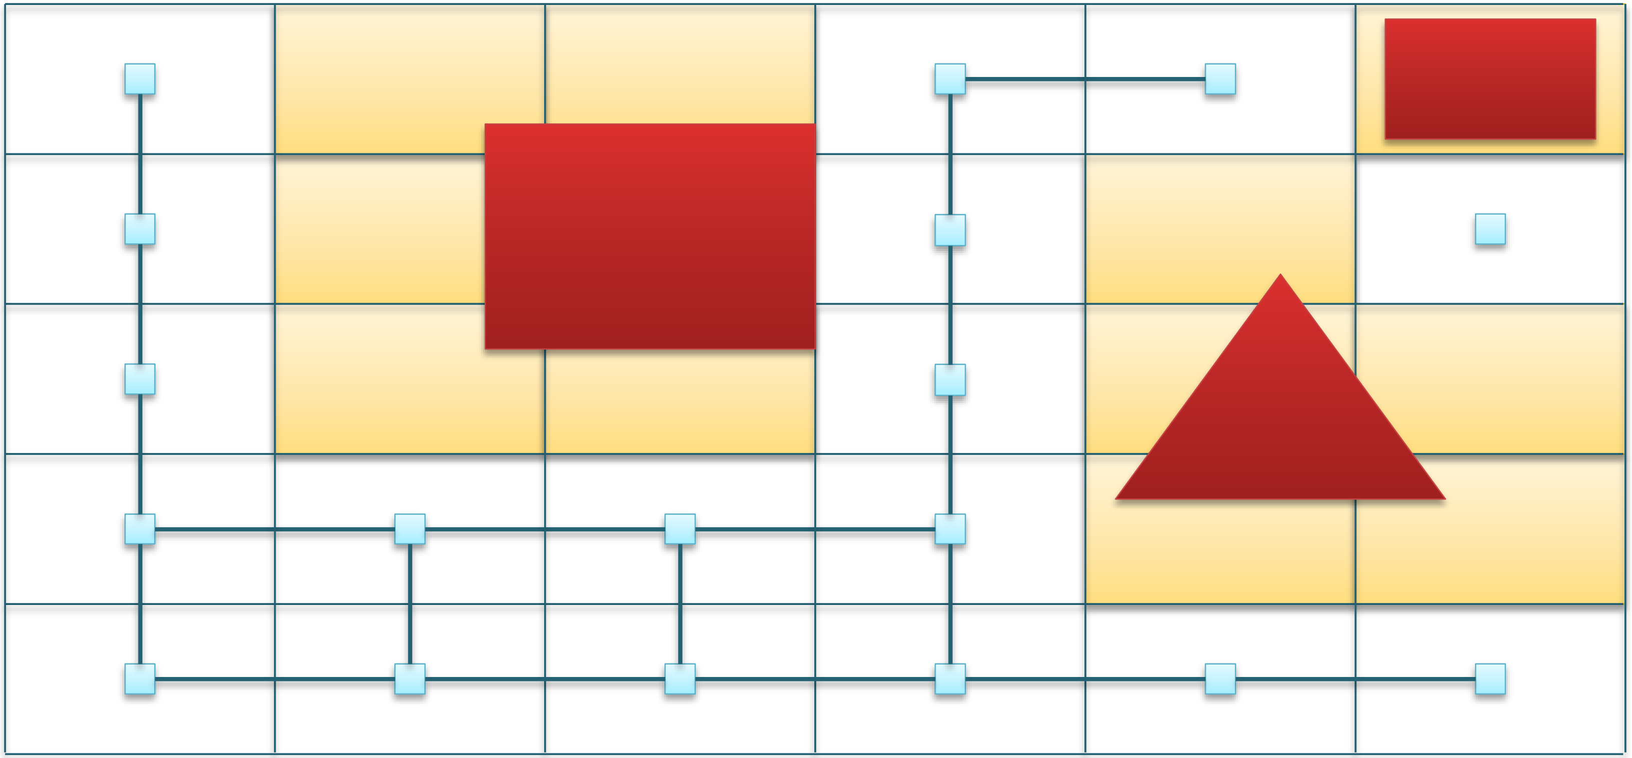

Regular 2D Grid: More Edges

- …or **8-connectivity**

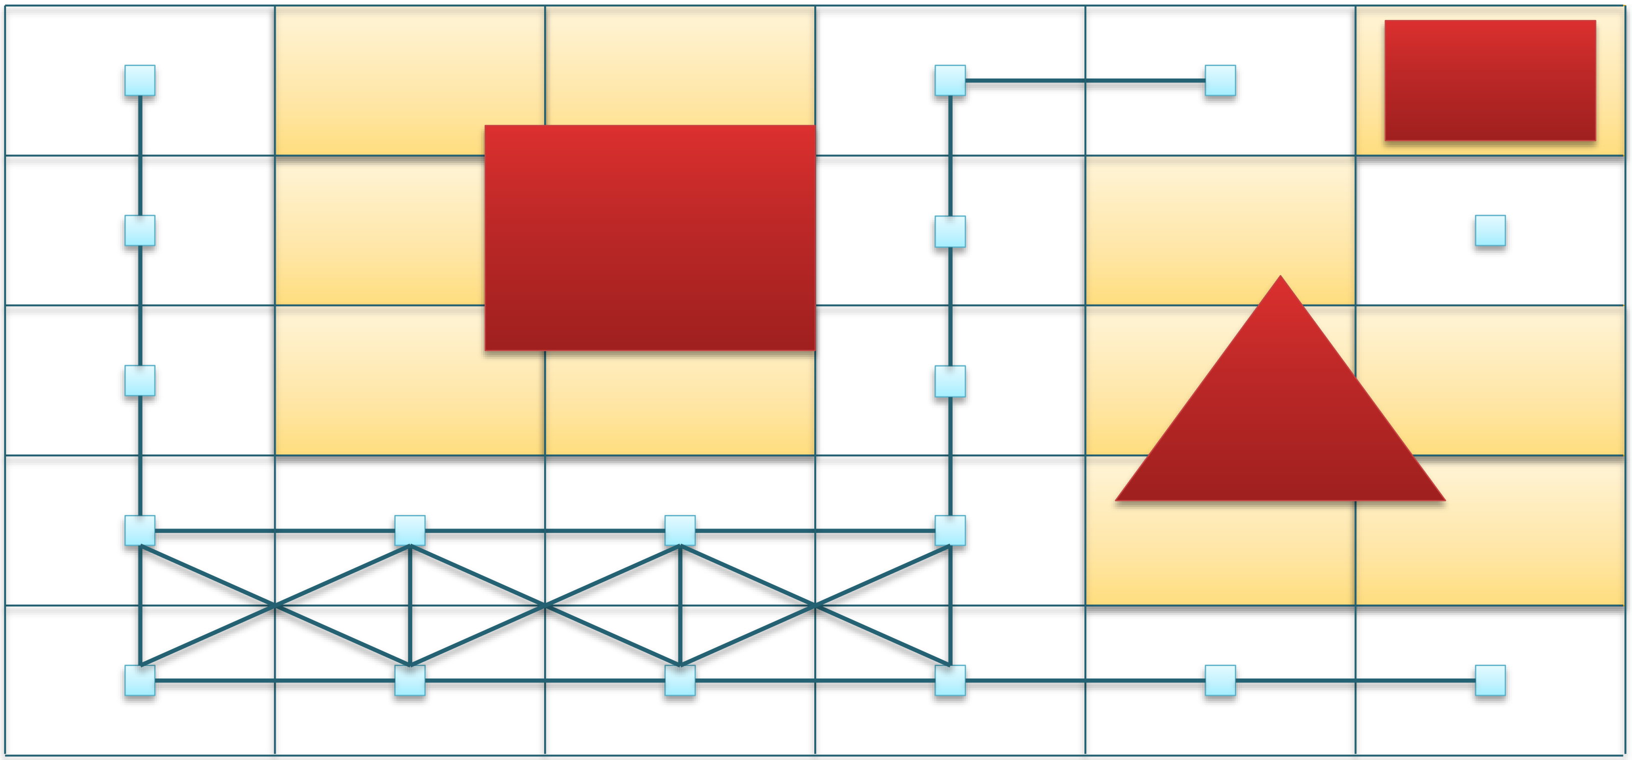

Finding a Solution

Discrete Graph Search

jonkv@ida

- Connect **start/goal** configurations to the nodes in their cells
  - Results in a **discrete graph search** problem

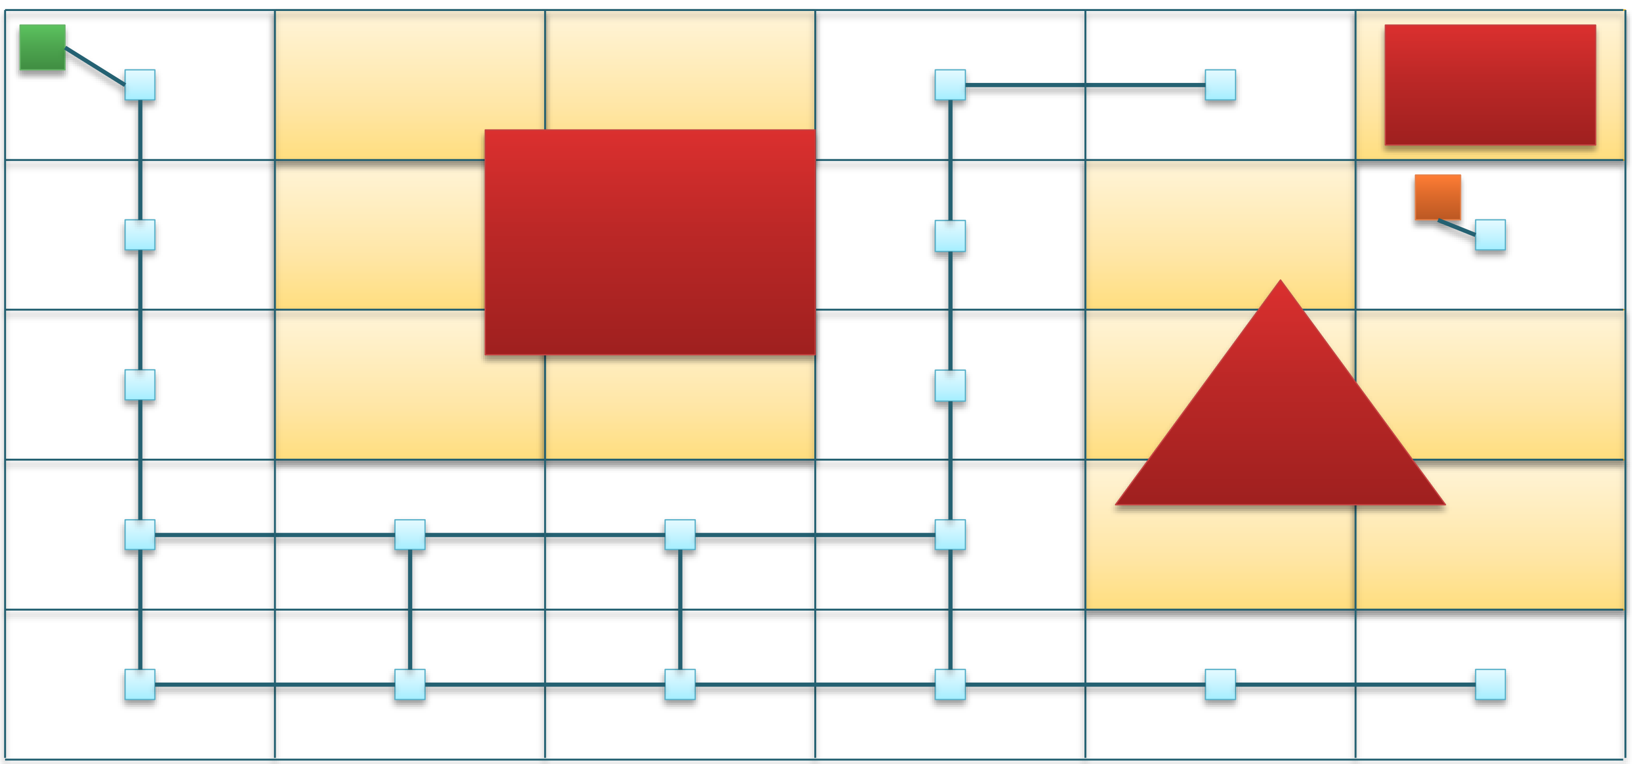

A*, Heuristics

- Finding a path: Any graph search algorithm
  - For example: A*
  - Heuristics in **simple geometric paths**: Manhattan distance (4 directions), Chebyshev distance (moving in 8 directions), Euclidean distance (in general), …

But there is no solution for this discretization!

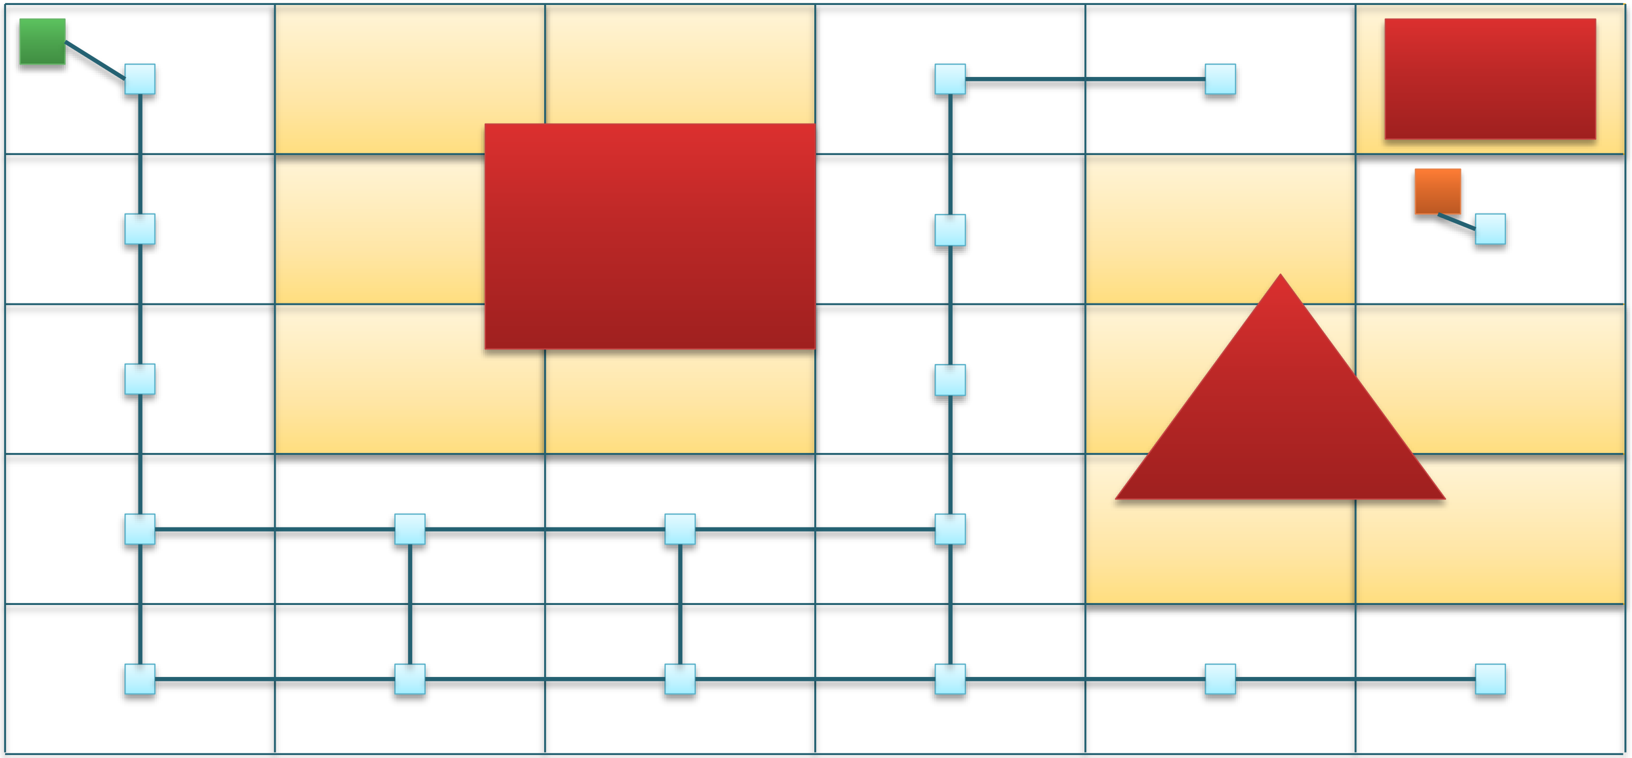

Grid Density

- **Grid density** matters!
  - Here: 4 times as many **cells**
  - Better approximation of the true obstacles, but many more nodes to search

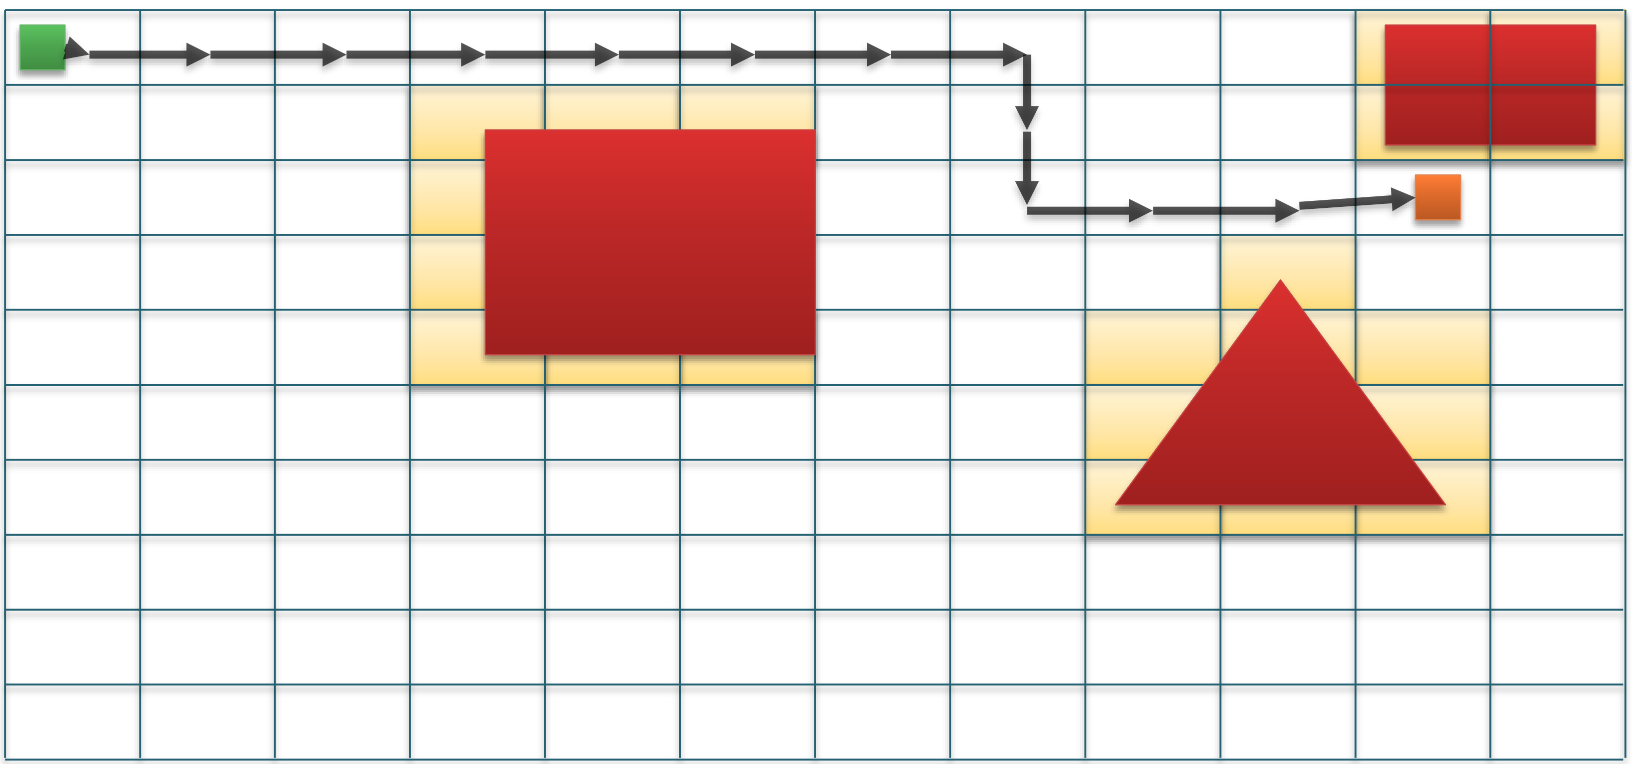

Solutions

- Solutions are correct under certain assumptions
  - The robot **<u>can</u>** turn 90° in place,
    or all free grid cells provide terrain
    where we can actually follow ***curves***

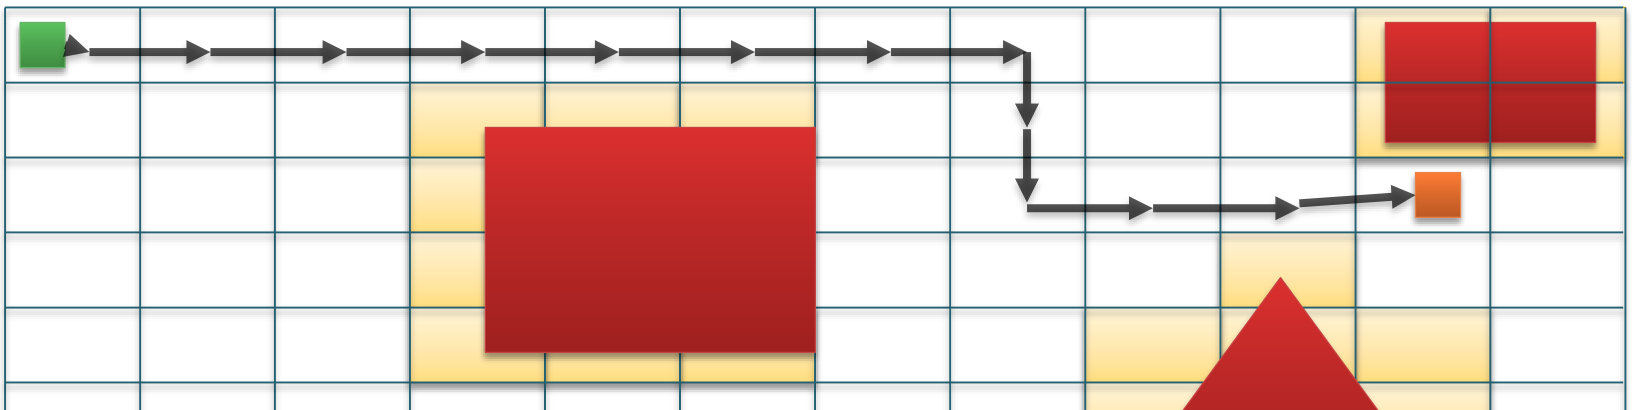

Irregular Node Placement

Non-Regular Grids

- Alternative to high regular density: **<u>Non-regular</u>** grids
  - For example, denser cells around obstacles

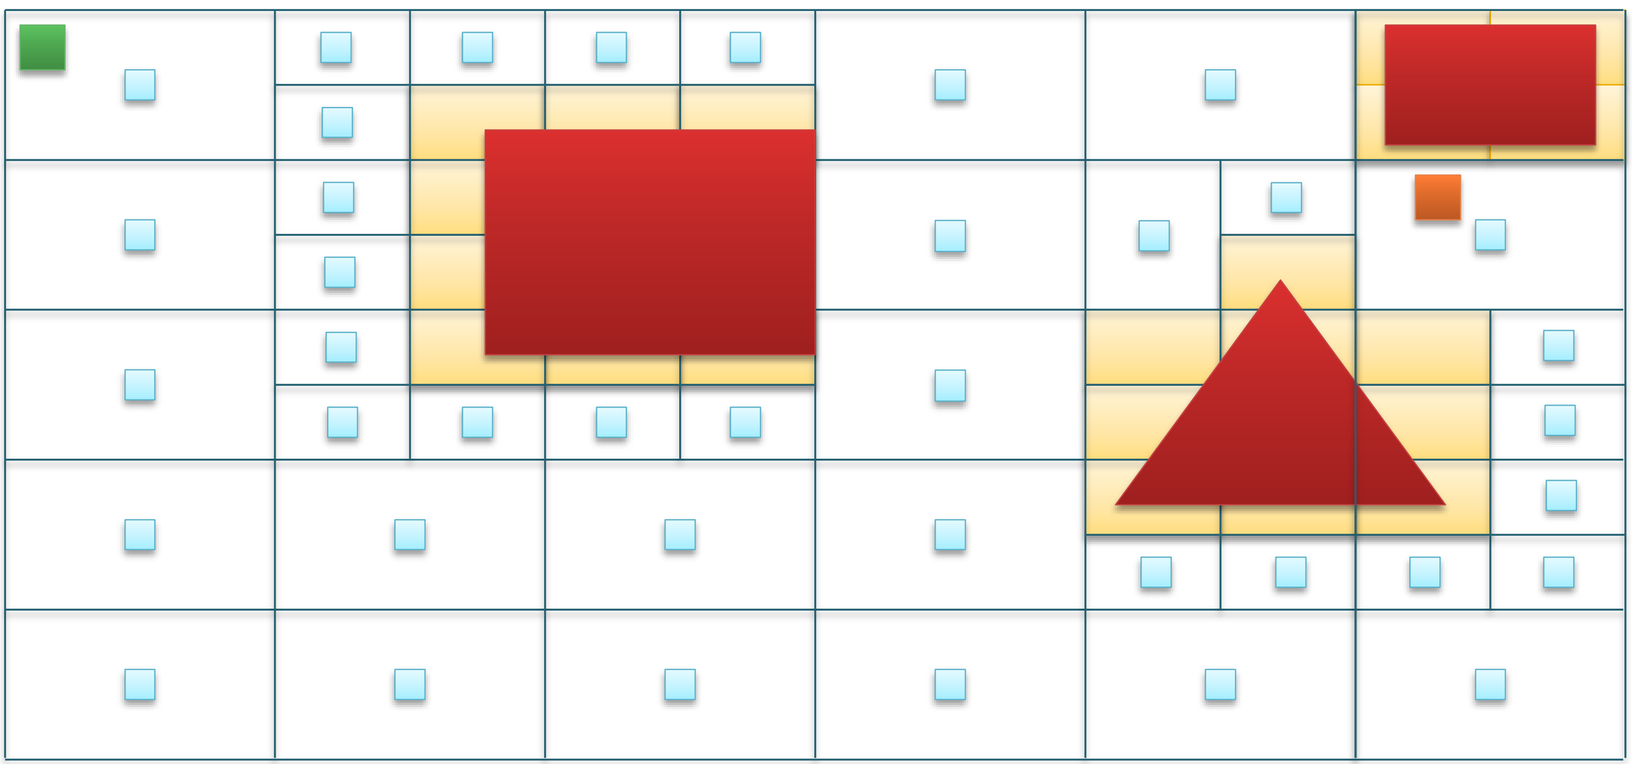

Grid Representations

- Space-efficient data structure: **quadtree**
  - Each node keeps track of:
    - Whether it is *completely* covered, *partially* covered or *non-covered*
  - Each non-leaf node has exactly four children

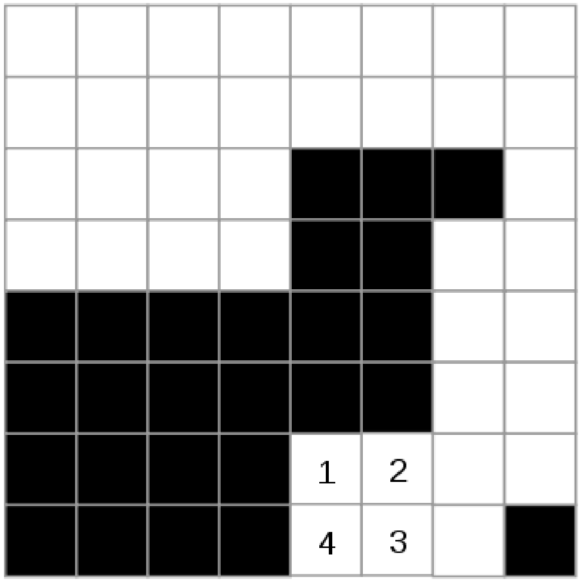

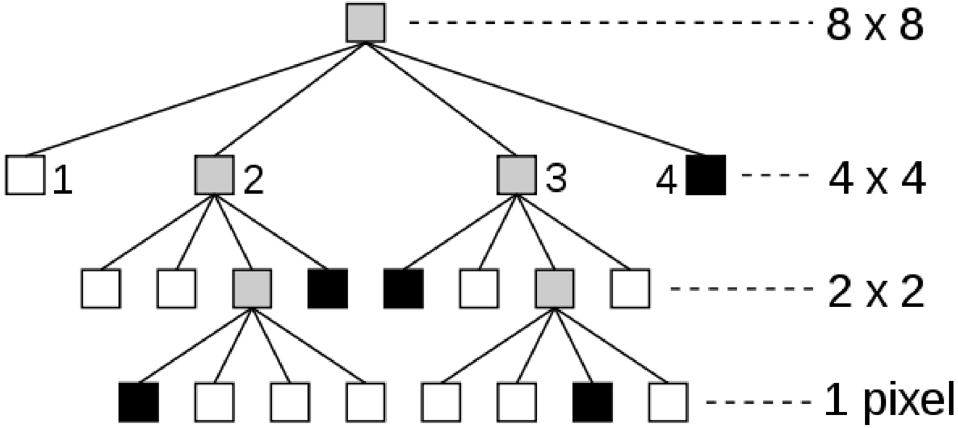

Grid Representations

- Can be generalized to 3D (octree), …

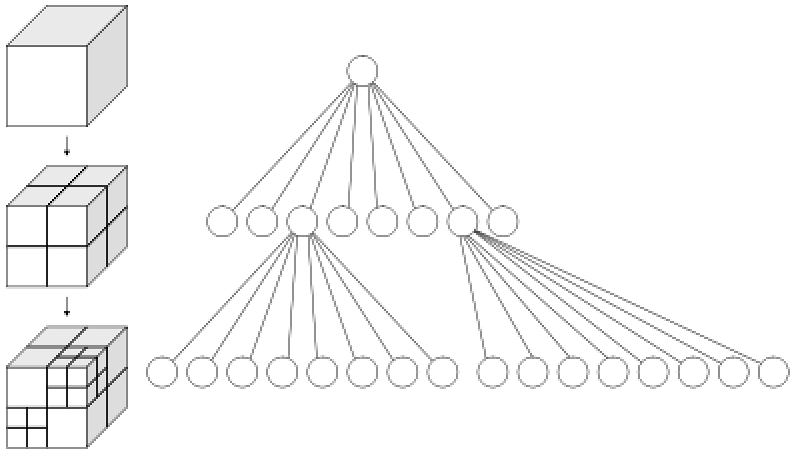

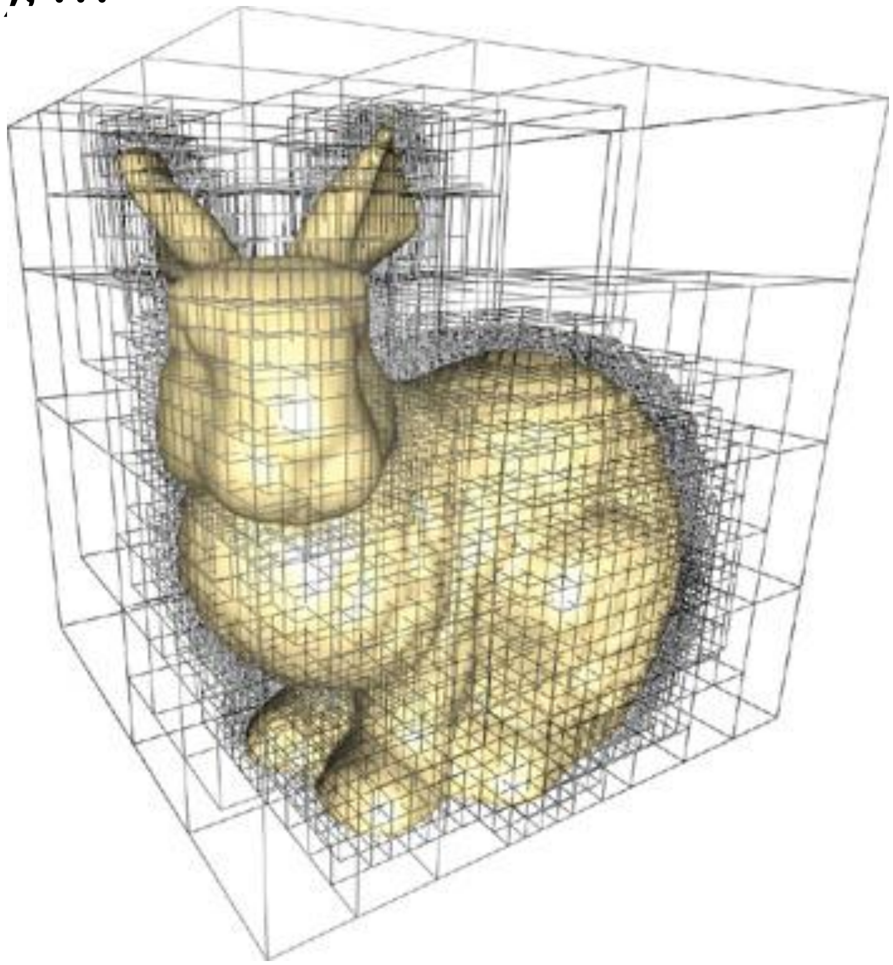

Choosing Potential Waypoints: Geometry-Based Methods

Non-Grid Placement

- Grid-based methods can result in many nodes
  - Even with efficient representation, *searching* the graph takes time
  - Alternative idea: **Place** nodes **depending** on obstacles

- Simple case: Known road map
  - Model all non-road areas as obstacles, then add a dense **grid**?

  - Or **place** a node in each **intersection**?

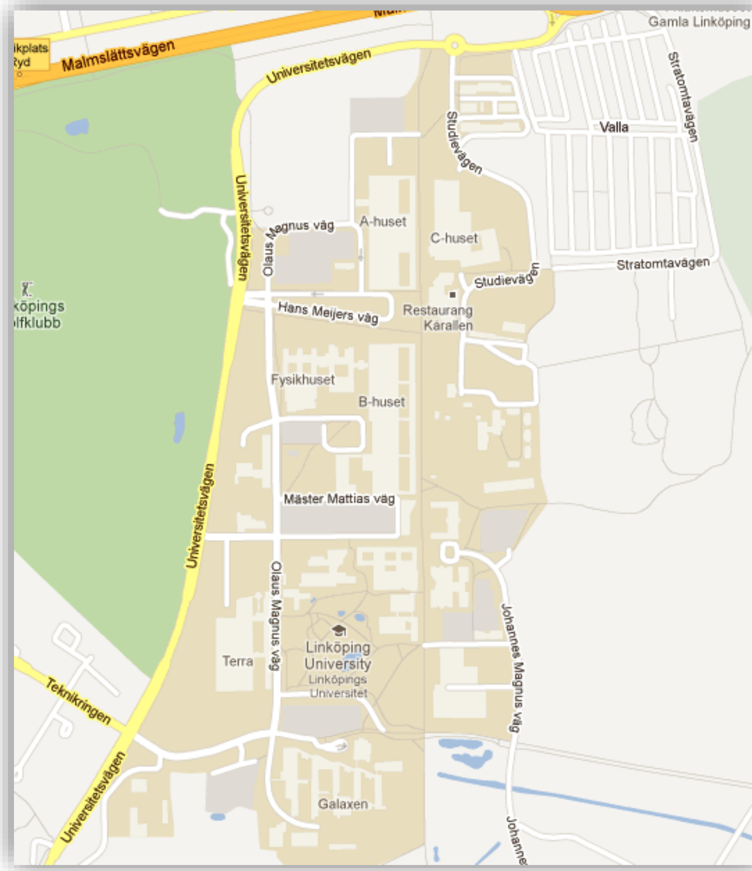

If we only know the **obstacles** (no roads), where to place the **nodes**?

Visibility Graphs

- **Visibility graphs**
  - Applicable to **simple polygons** – straight sides *without intersections*
    - **Nodes** at all polygon corners
    - **Edges** wherever a pair of nodes can be connected using the local planner
  - Mainly interesting in 2D
    - Optimal in 2D, not in 3D

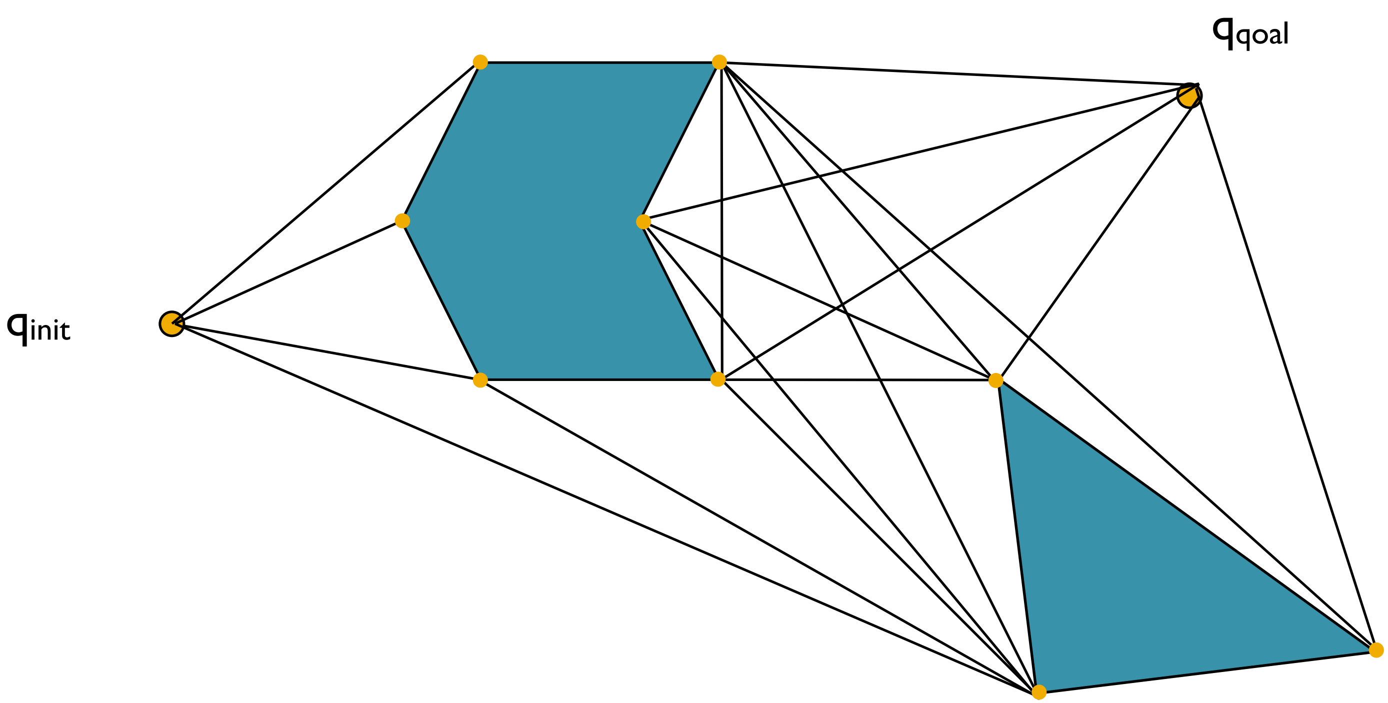

Voronoi Diagrams

- **Voronoi** diagrams
  - Find all points that have the same distance to two or more obstacles
    - Maximizes **clearance** (free distance to the nearest obstacle)
  - Creates unnecessary **detours**
  - Mainly interesting in 2D – does not scale well

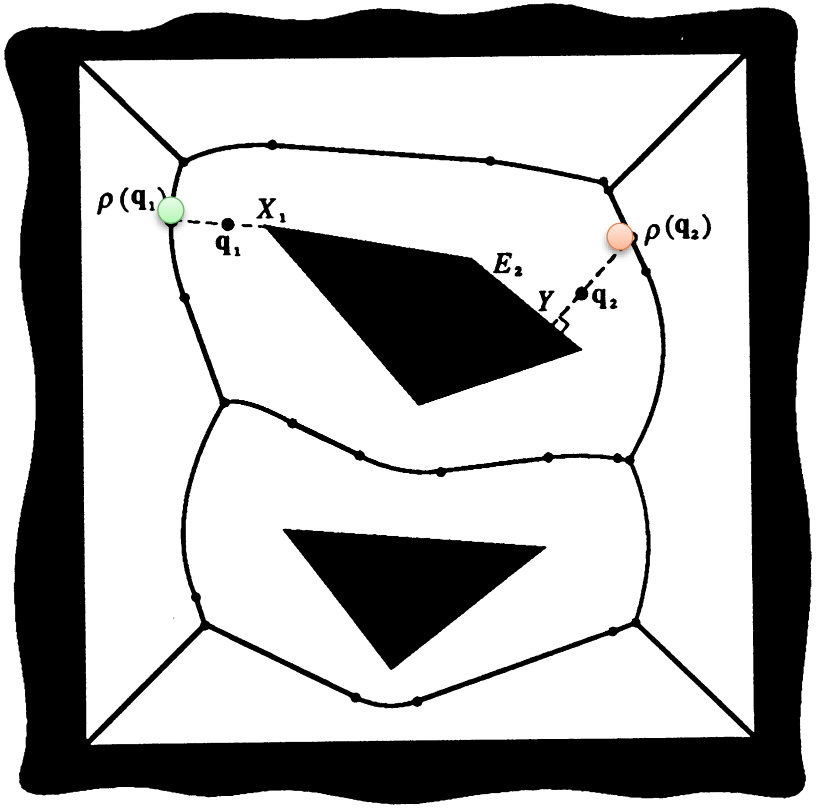

Motion Constraints and Complex Motion Planning Problems

Motion Constraints (1)

- So far, we implicitly assumed:
  - If we can **<u>draw a line</u>** between two waypoints, the robot can **<u>move</u>** between the waypoints

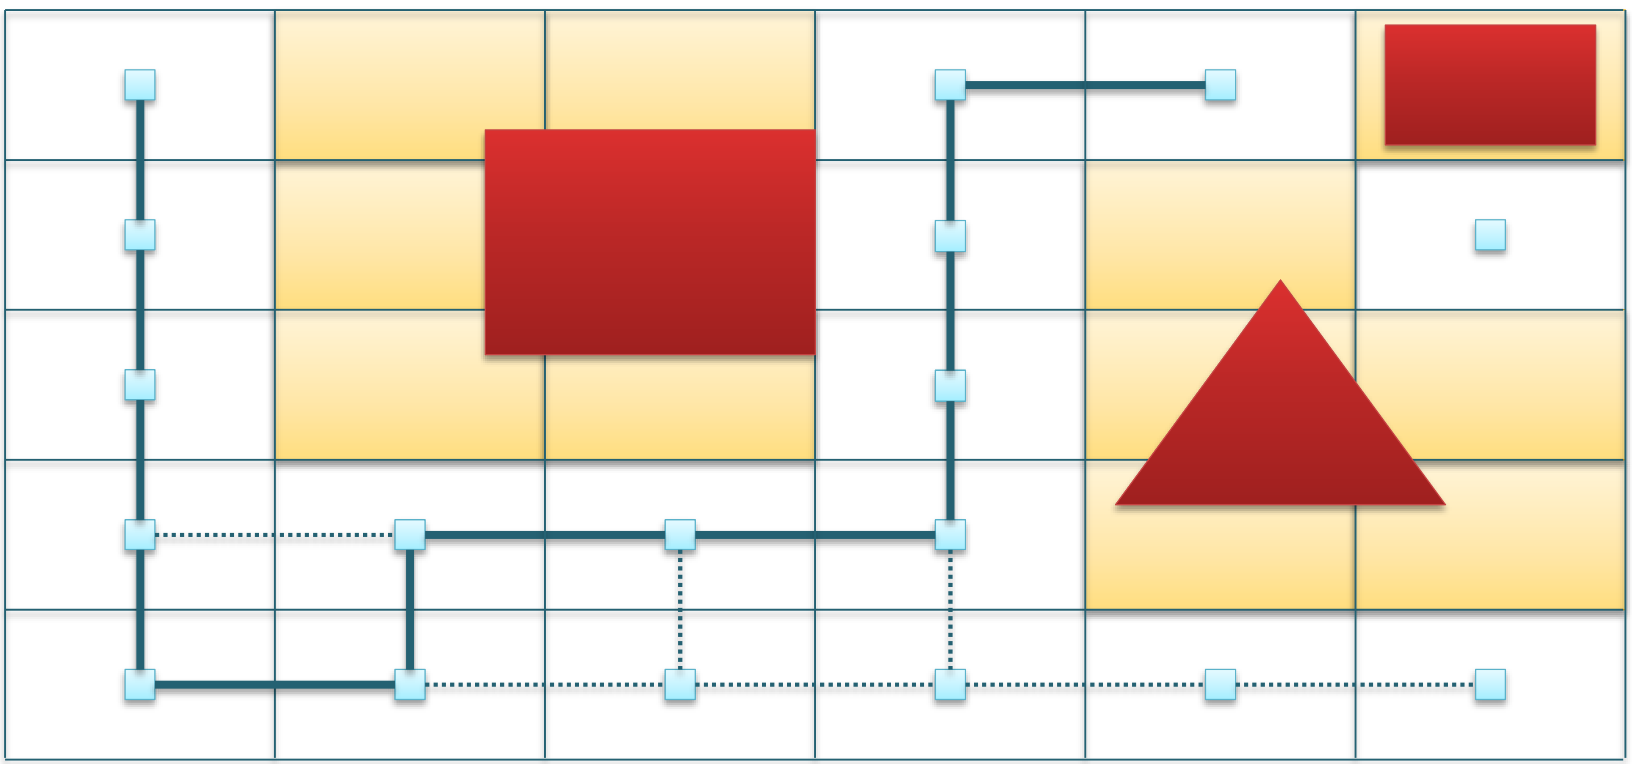

Motion Constraints (2): Holonomic

- May work if the robot is **holonomic**
  - Informally: Can move in any direction (possibly by first rotating, then moving)

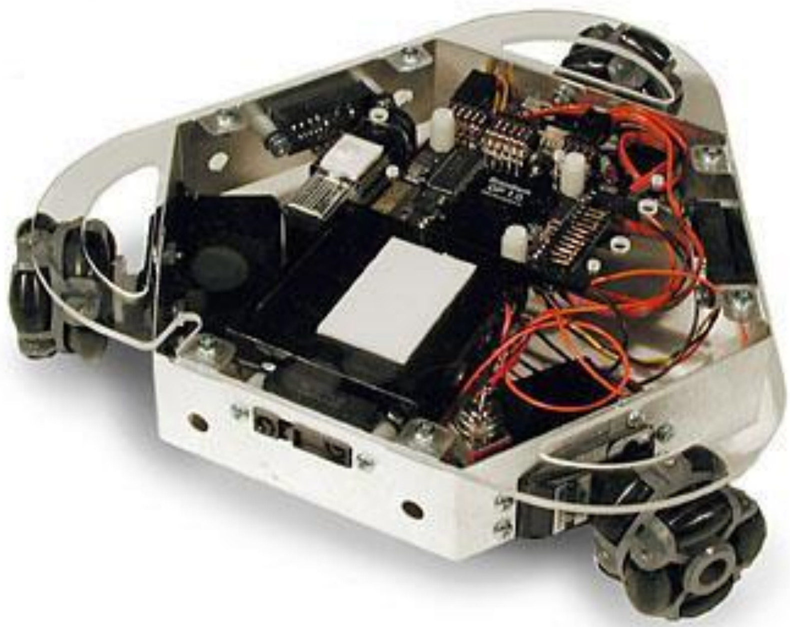

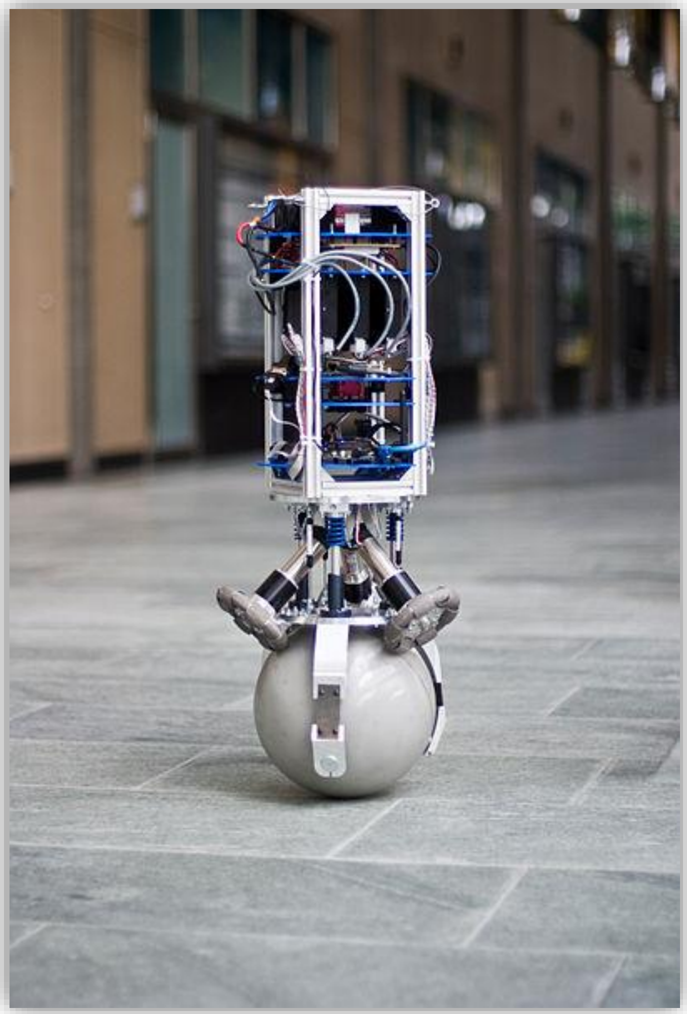

Motion Constraints (3): Non-Holonomic

jonkv@ida

- But: Can an airplane fly this path?
  - How do we know? What are the constraints?

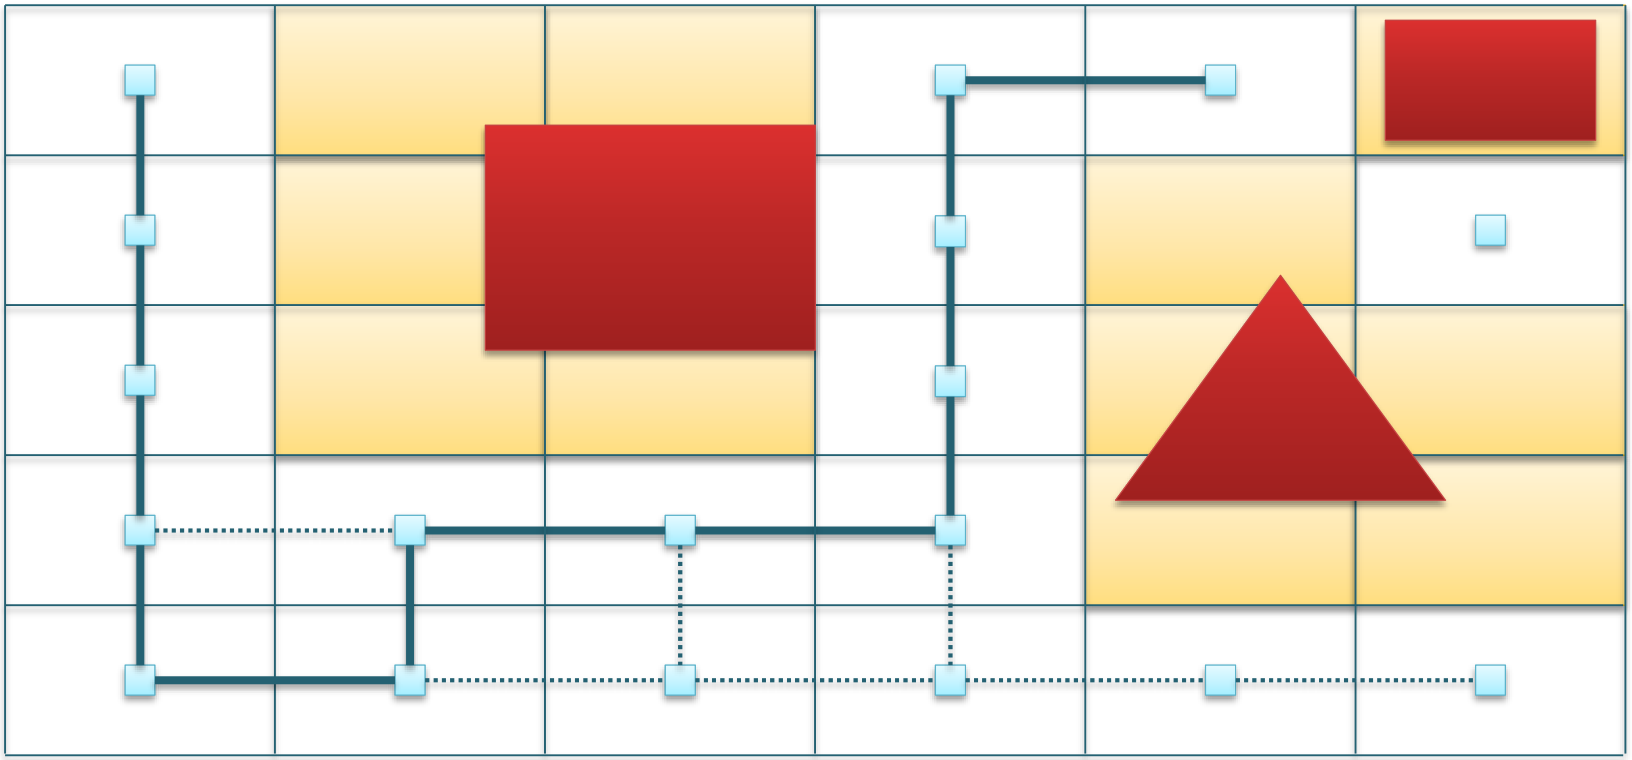

We need some new concepts…

Workspace and Configuration Space

Workspace (1)

- The **workspace** is (1) the **physical space** in which we work…
  - 3 physical dimensions, 3-dimensional coordinates, 3-dimensional obstacles
    - Need full 3D geometry to determine how the helicopter can move

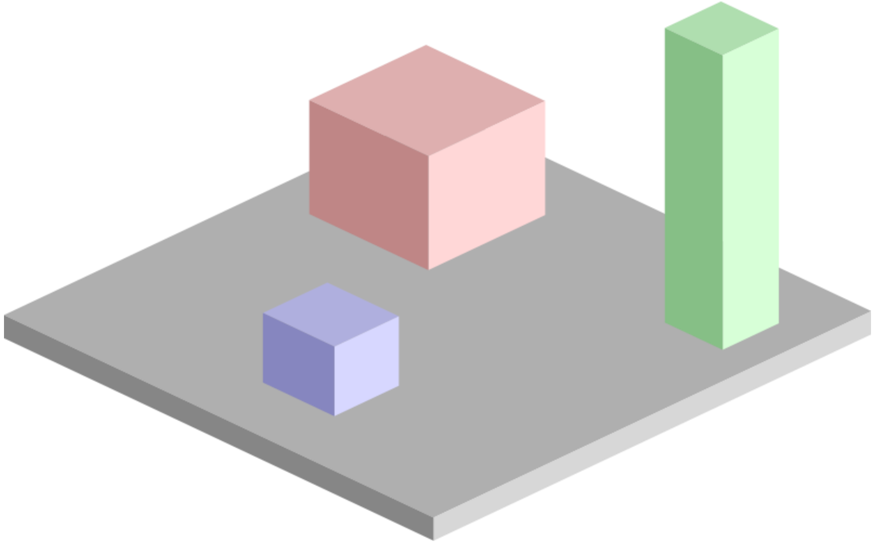

Workspace (2)

- … or (2) a **<u>2D projection</u>**, in case this is sufficient
  - For a car:
    - Can describe position, rotation in 2D (except tunnels, bridges, …)
    - Can describe obstacles in 2D
  - ➔ Workspace can be 2D
    - **<u>Still represents physical locations</u>**

Configuration Space

- Even a car has 3 **physical degrees of freedom** (DOF)!
  - The **configuration space** of the car
    - **Location** in the plane (x/y),
    - **Angle** (θ)
  - Each DOF is essential!
    - As part of the *goal* – park at the correct angle
    - As part of the *solution* – must turn the car to get through narrow passages

Motion planning takes place in **configuration space**:
How do I get from $(200, 200, 12°)$ to $(800, 400, 90°)$?

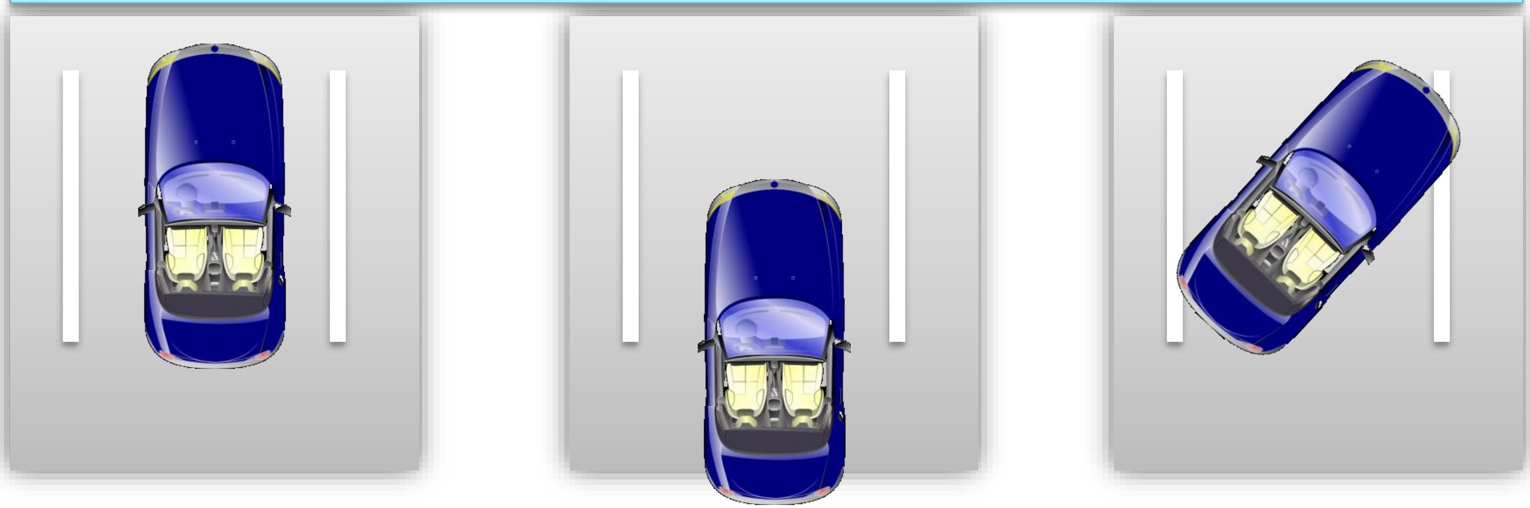

The Ladder Problem

- The **ladder problem** is similar
  - Move a ladder in a 2D workspace , with 3 **physical** DOF
  - Configuration:
    - **Location** in the plane (x/y),
    - **Angle** (θ)

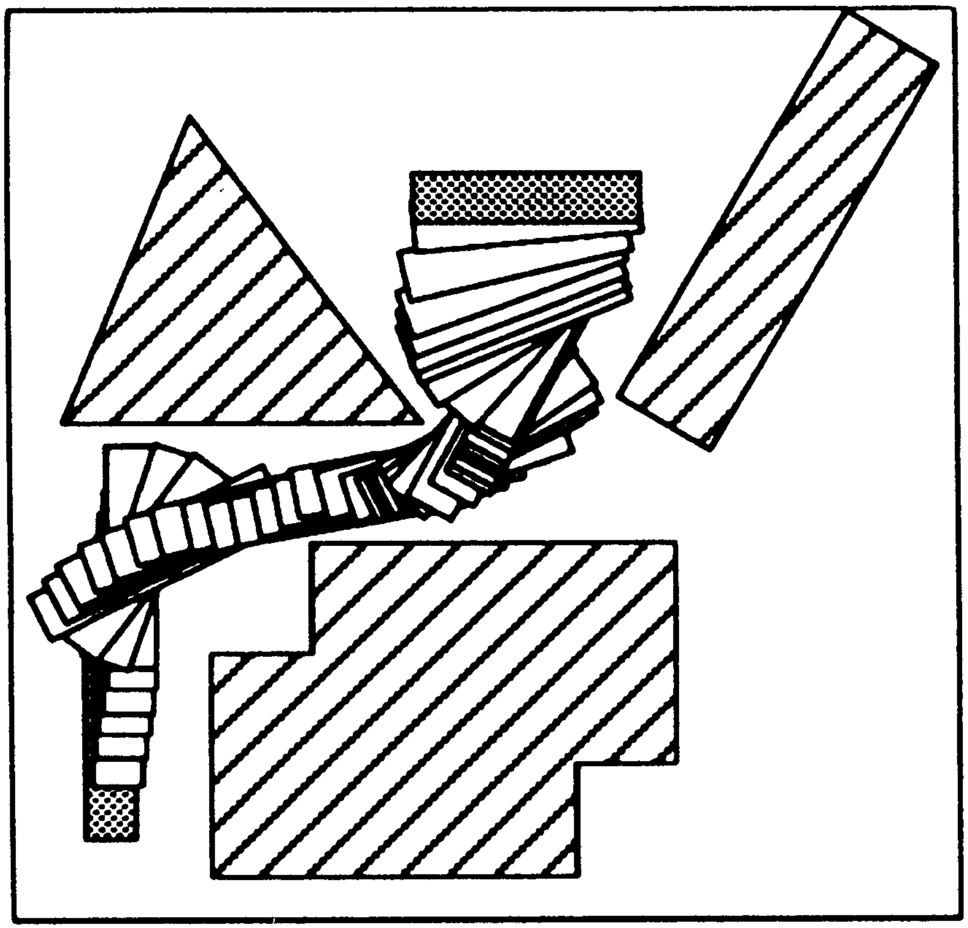

- Again, each DOF is essential:
  - As part of the ***goal***
    - We want the ladder to end up at a specific angle
  - As part of the ***solution***
    - We need to turn the ladder to get it past the obstacles

The Ladder Problem: Controllable DOF

- For ladders, each physical DOF is directly controllable!
  - You can:
    - Change x (translate sideways)
    - Change y (translate up/down)
    - Change angle (rotate in place)
  - Therefore:
    - If you want to get from $(200, 200, 12°)$ to $(800, 400, 90°)$, any path **connecting** these 3D points and **going through** free configuration space is sufficient
  - The ladder is **holonomic**!
    - Controllable DOF >= physical DOF

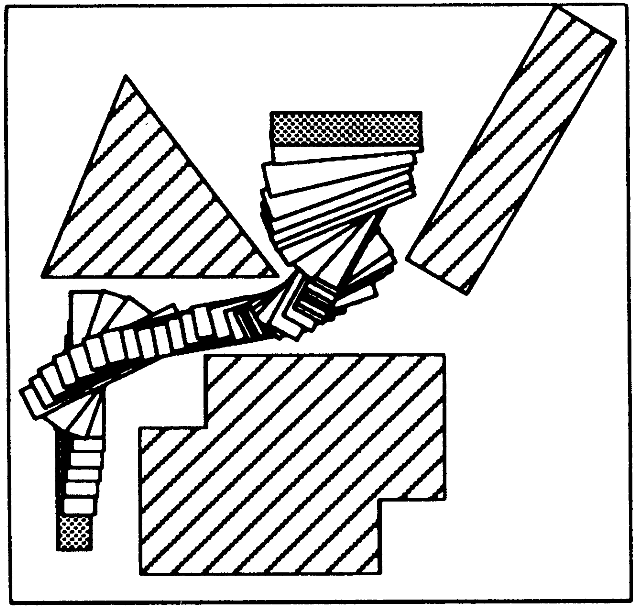

Controllable Degrees of Freedom

- Cars have 3 degrees of freedom
  - But only move **back** and forward,
    along **curves** with constrained **turning radius**
  - ➔ **constrained curves in configuration space**

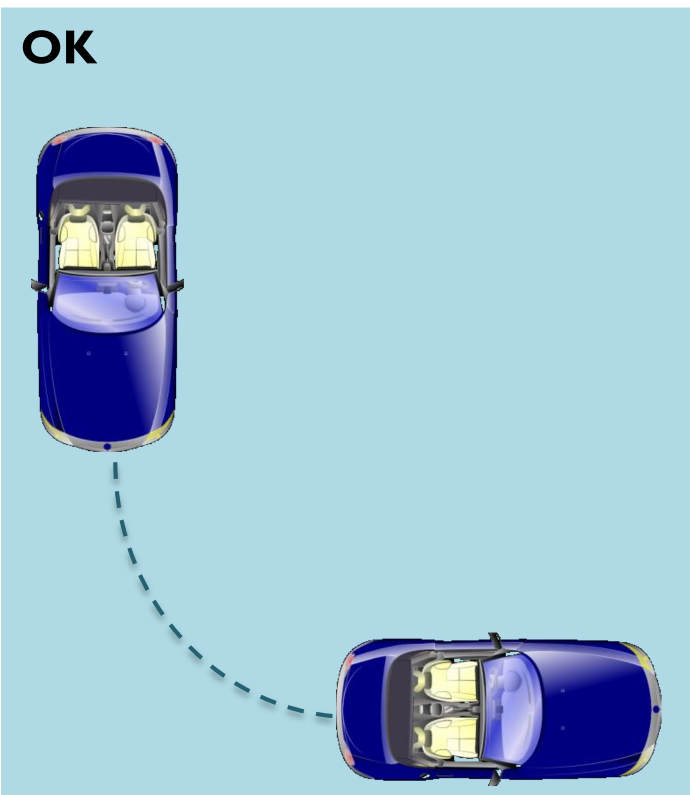

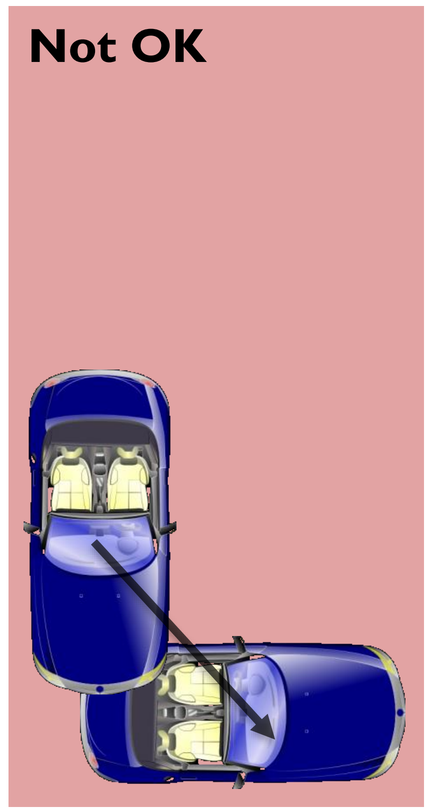

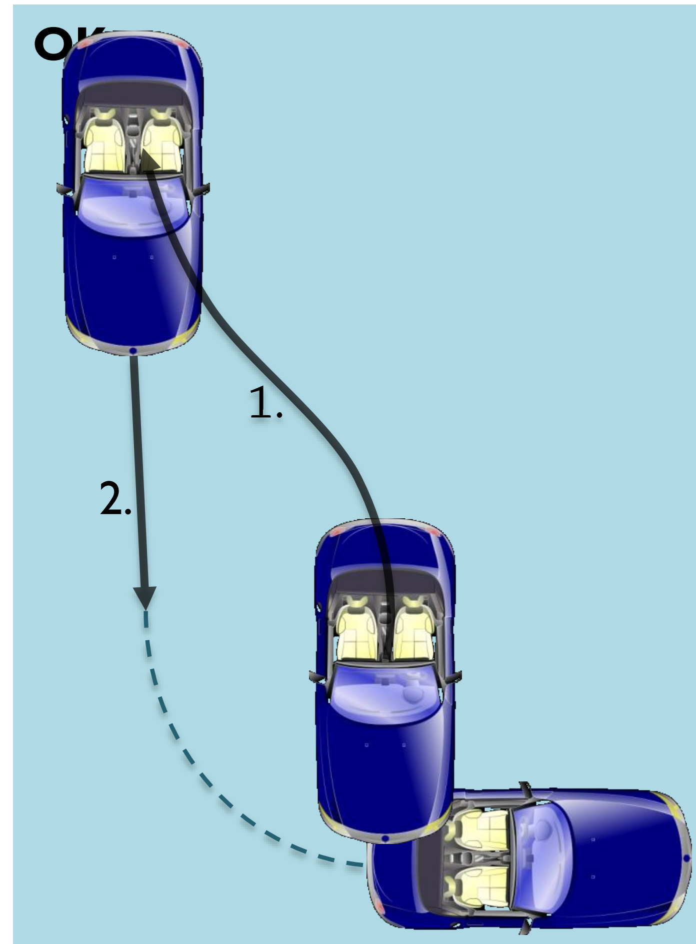

Controllable Degrees of Freedom (2)

- In this **parallel parking** example:
  - There **is** free space between current and desired configurations
    - But we can't slide in sideways!
  - Fewer controllable DOF than physical DOF!
    - ➔ non-holonomic
    - *Limits possible curves in 3D configuration space!*

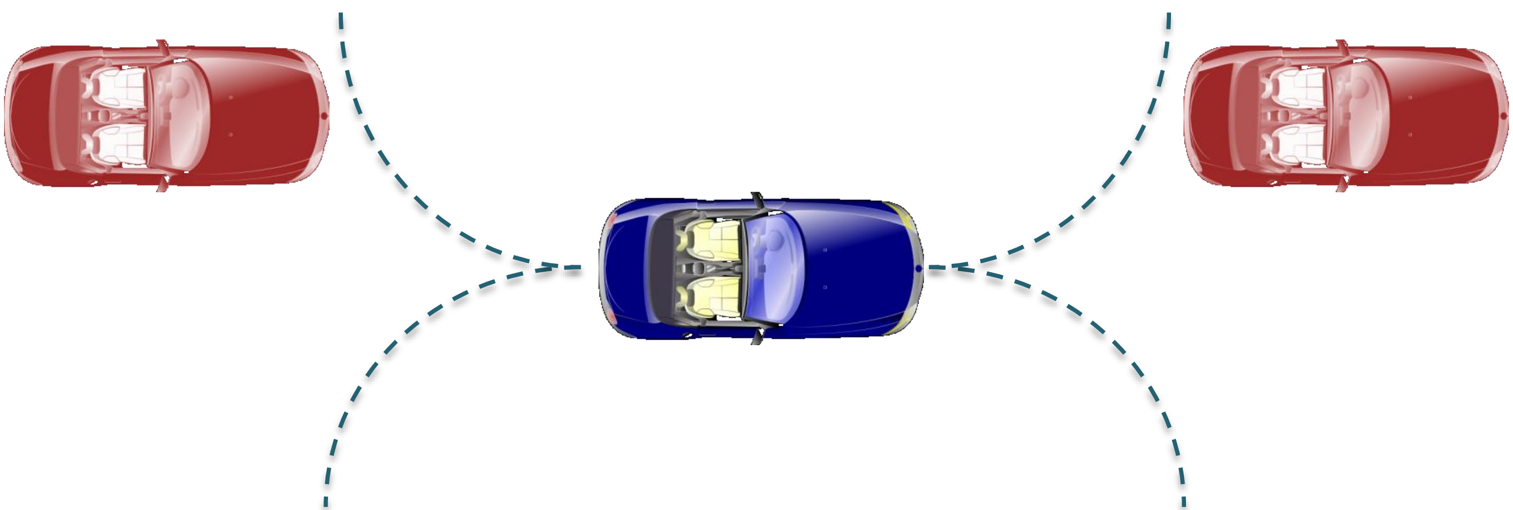

Work Space, Configuration Space

- Summary of important concepts:
  - **Work space**: The physical space in which you move
    - 3-dimensional for this robot arm
  - **Configuration space**:
    The set of possible configurations of the robot
    - Usually **continuous**
    - Often **many-dimensional**
      (one dimension per physical DOF)
    - Will often be **visualized** in 2D for clarity
  - We have to search
    in the **configuration space**!
    - Connect *configurations*, not *waypoints*

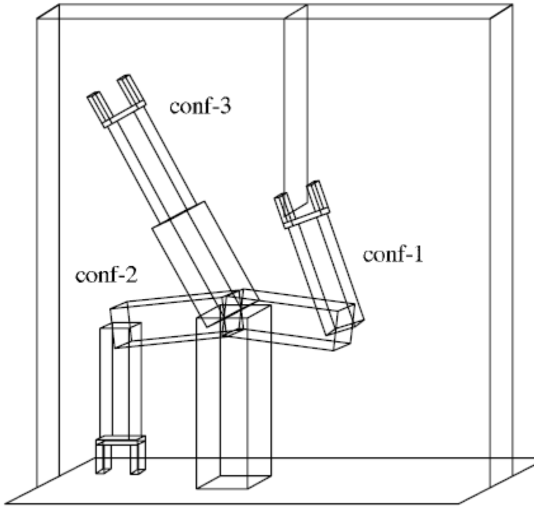

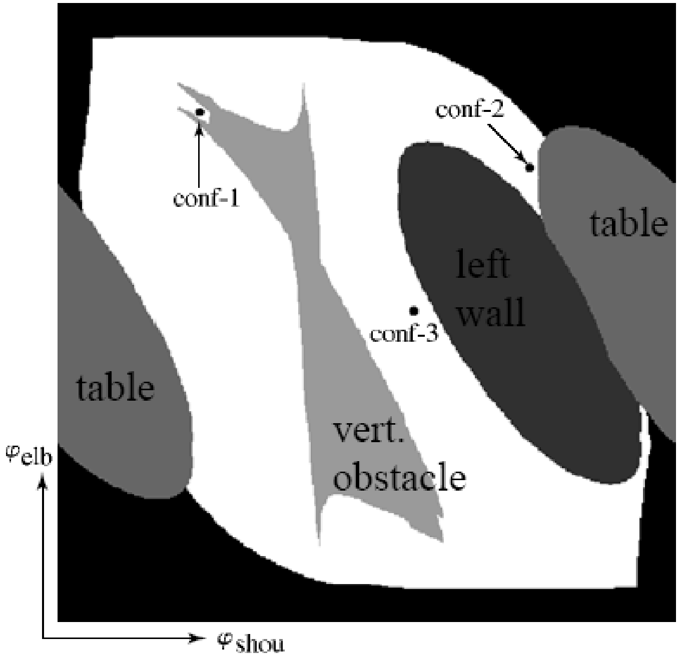

Local and Global Planners:

Divide and Conquer
in Configuration Space

Searching the Configuration Space

- Divide and Conquer!
  - **Local** path planner
    - Determines whether two **configurations** can be connected with a **path in configuration space**, and how
    - Considers vehicle-specific constraints

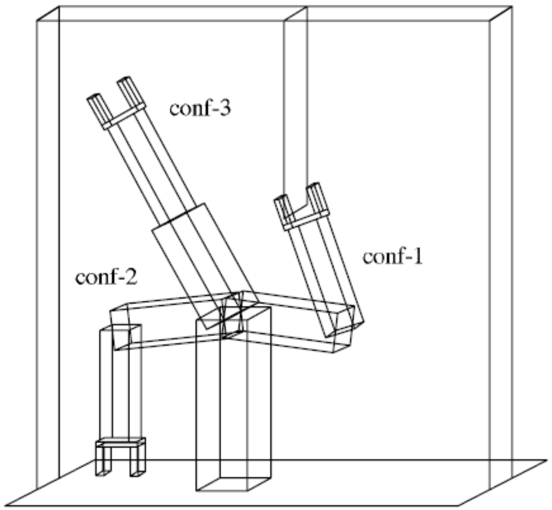

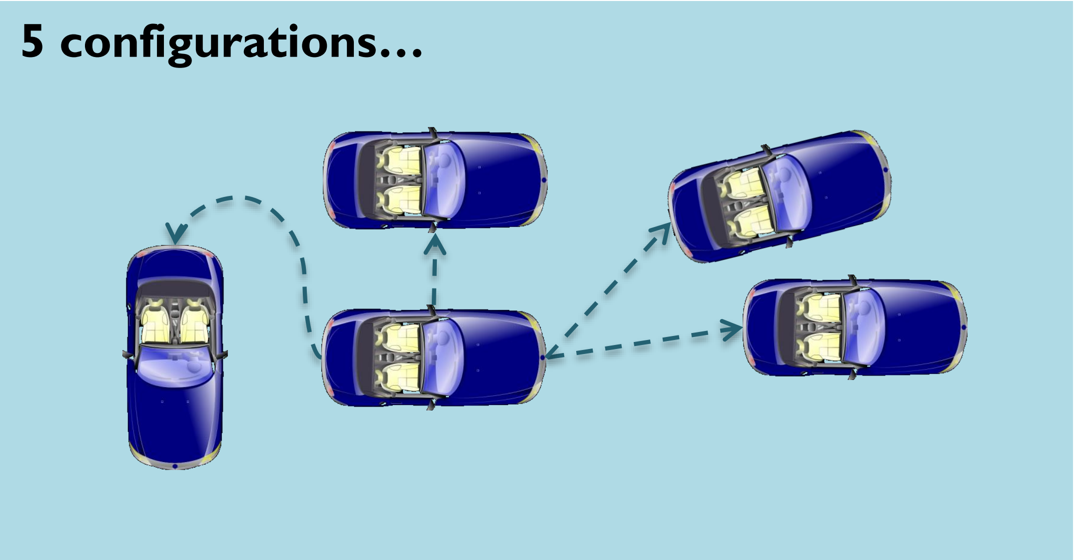

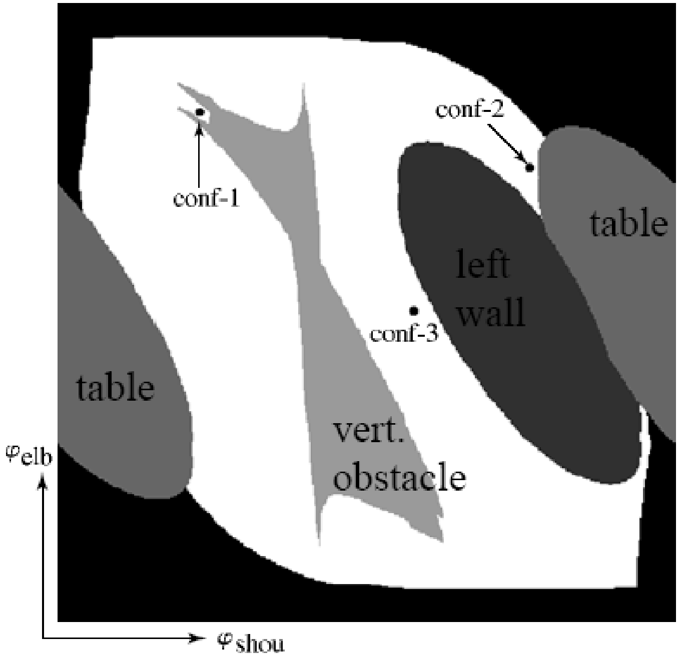

Searching the Configuration Space

- Divide and Conquer!
  - **High-level** path planner
    - **Generates** a finite set of **configurations**
    - **Calls local planner** to determine which configurations can be connected
    - Uses **discrete search** to determine a sequence of configurations to "pass through"

5 configurations…

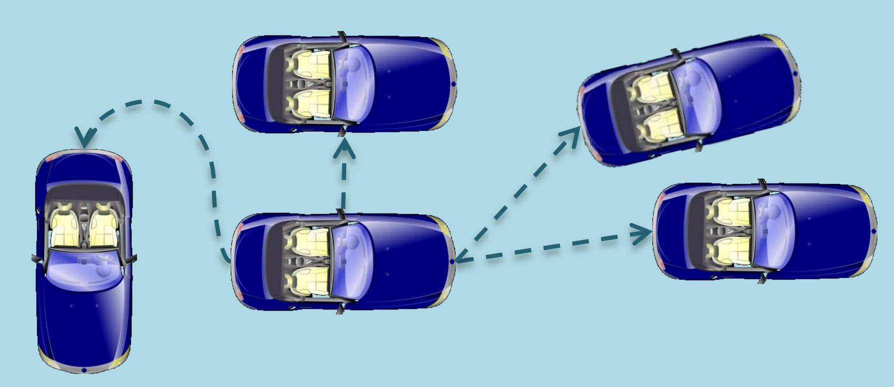

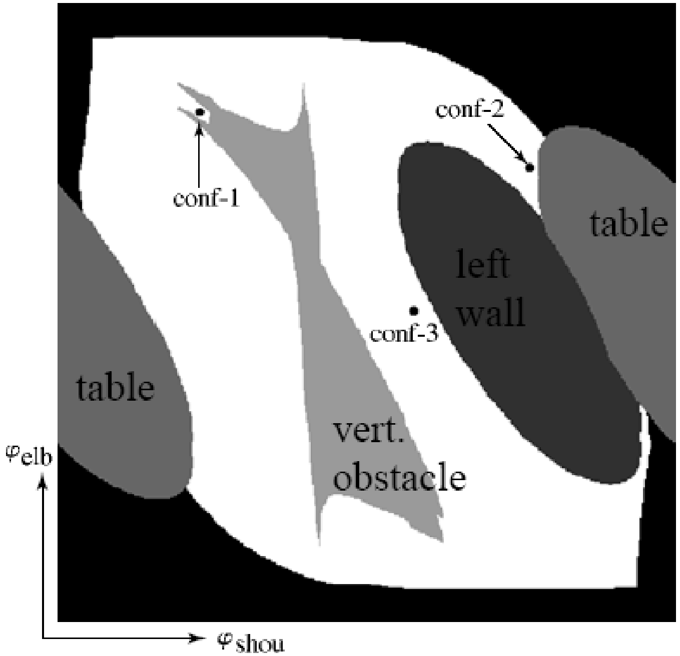

Low-Dimensional Problems

- In low-dimensional problems:
  - The high-level planner ***could*** select configurations in a grid (”equal distance”)
    - Car: 3-dim configuration space
    - Example: 6 locations, 4 angles considered per spatial location, 24 configurations

(0, 0, 0°) (0, 0, 90°) (0, 0, 180°) (0, 0, 270°)	(1, 0, 0°) (1, 0, 90°) (1, 0, 180°) (1, 0, 270°)	(2, 0, 0°) (2, 0, 90°) (2, 0, 180°) (2, 0, 270°)
(0, 1, 0°) (0, 1, 90°) (0, 1, 180°) (0, 1, 270°)	(1, 1, 0°) (1, 1, 90°) (1, 1, 180°) (1, 1, 270°)	(2, 1, 0°) (2, 1, 90°) (2, 1, 180°) (2, 1, 270°)

Low-Dimensional Problems (2)

- Let's illustrate this more **graphically**…

(0, 0, 0°)
(0, 0, 90°)
(0, 0, 180°)
(0, 0, 270°)

(1, 0, 0°)
(1, 0, 90°)
(1, 0, 180°)
(1, 0, 270°)

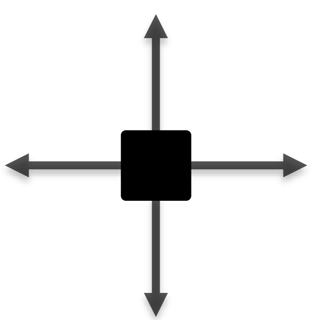

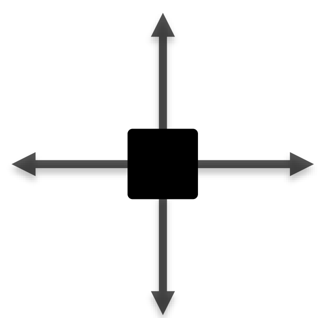

Local Planner (1)

- Ask local planner: "Can I connect these configurations"?

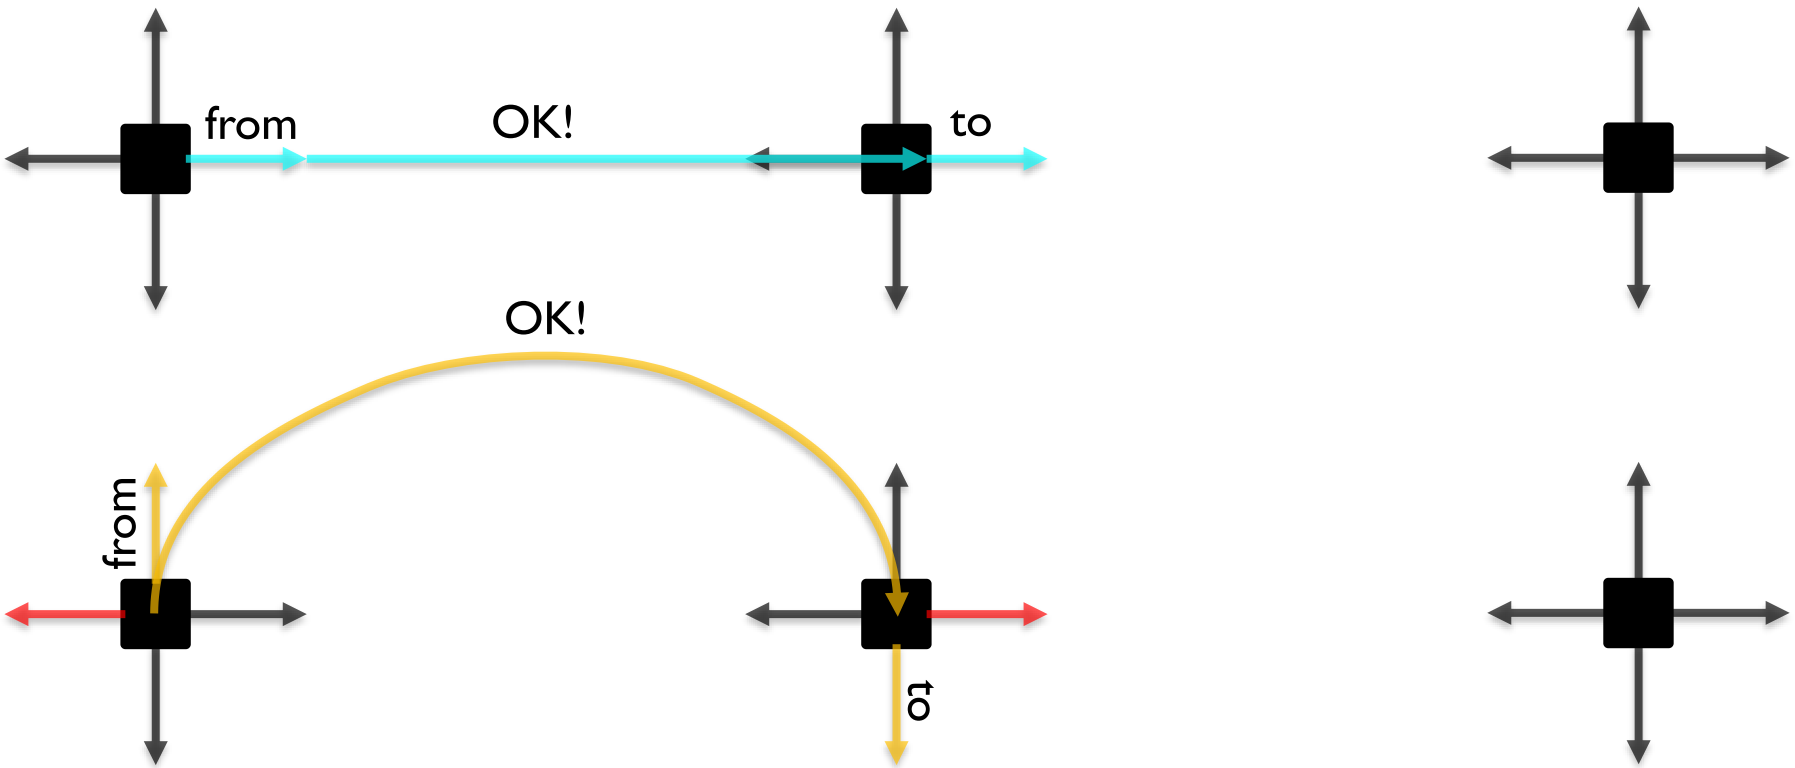

<u>Configurations</u>, not locations or points!

Can I go from **<u>here in this direction</u>** to **<u>there in that direction</u>**?
Can I go from **<u>these arm joint angles</u>** to **<u>those arm joint angles</u>**?

Local Planner (2)

- Ask local planner: "Can I connect these configurations"?

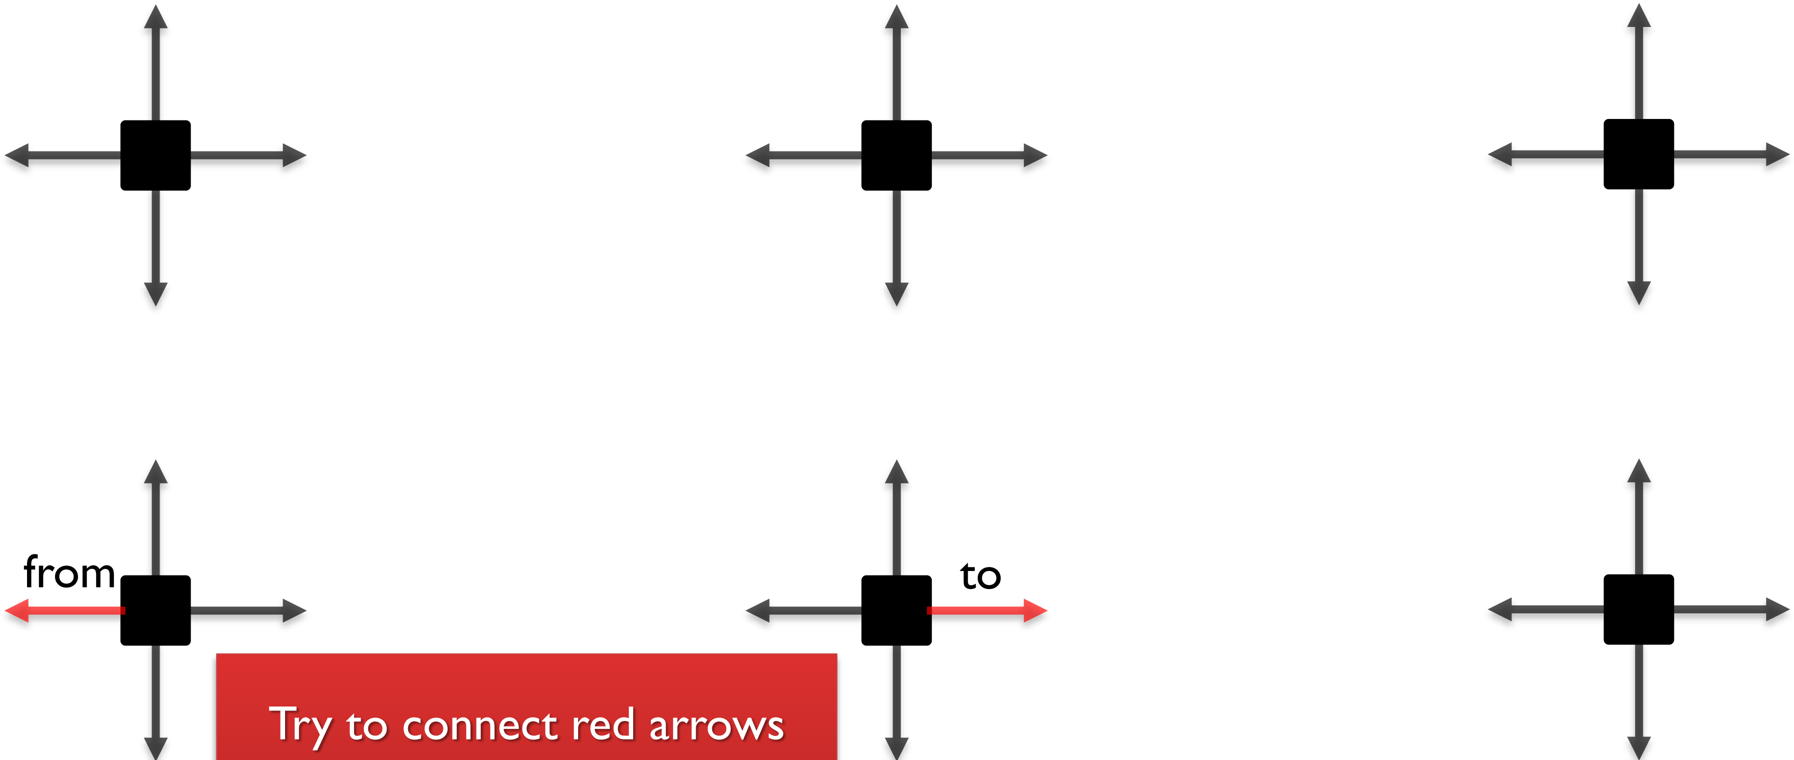

Try to connect red arrows **directly**:

The local planner might say "Sorry, too complex"

Local Planner (3): Local vs Global

- Other paths may be possible

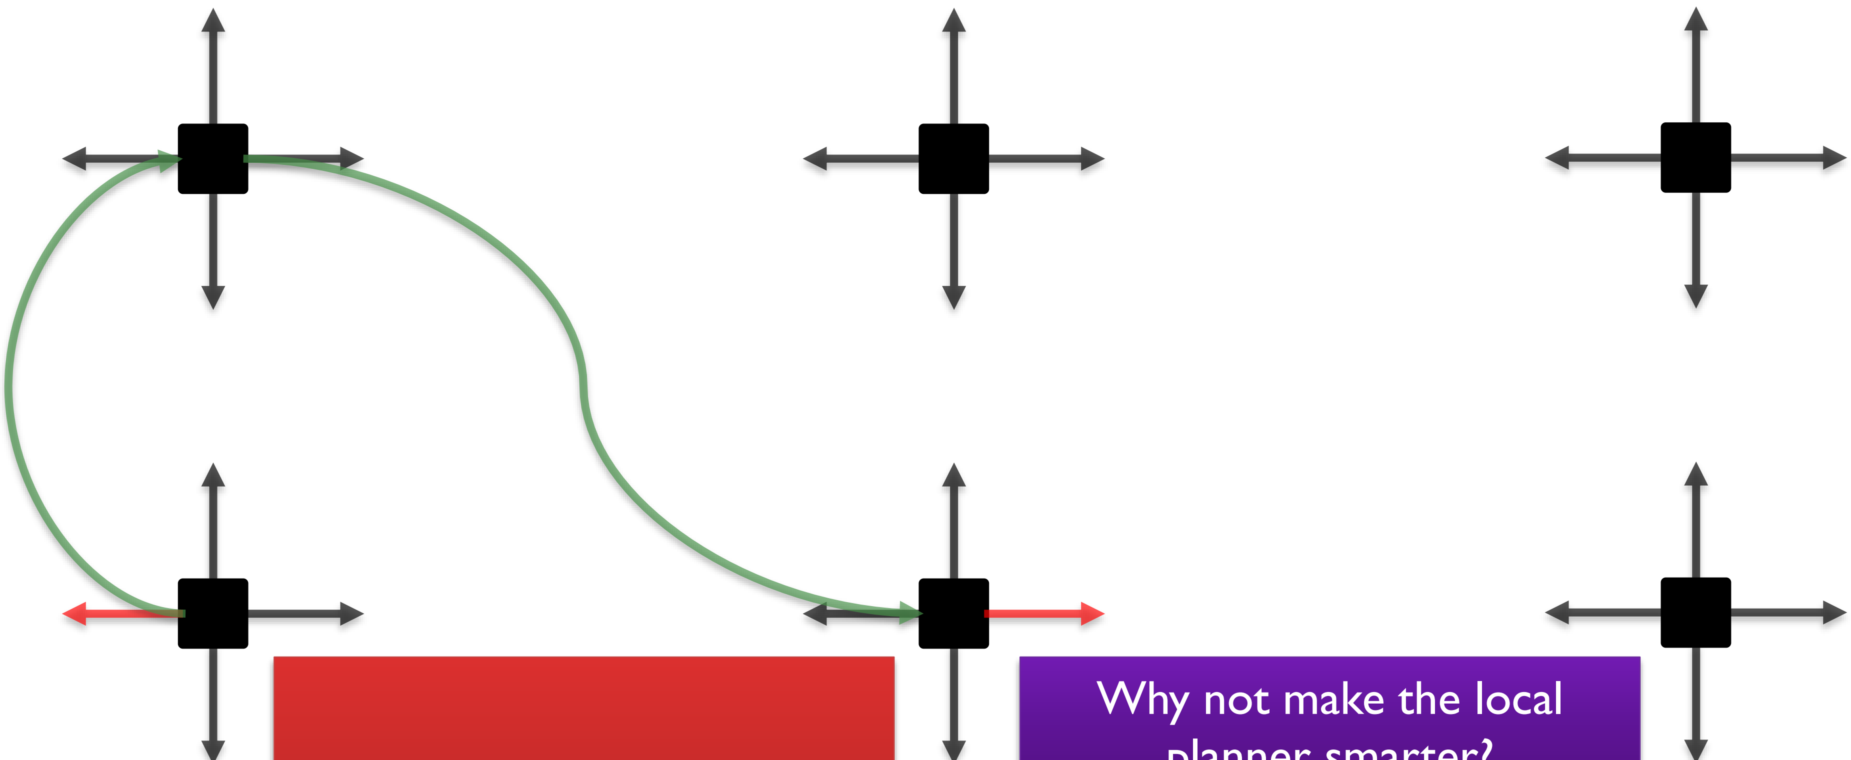

The **global** planner might be able to connect them **indirectly** through other configurations

Why not make the local planner smarter?

Divide and conquer:
Local planner should be *fast*, the rest is handled through the high-level planner

Local Planner (4)

- **<u>Local</u>** planner also considers **<u>obstacles</u>**

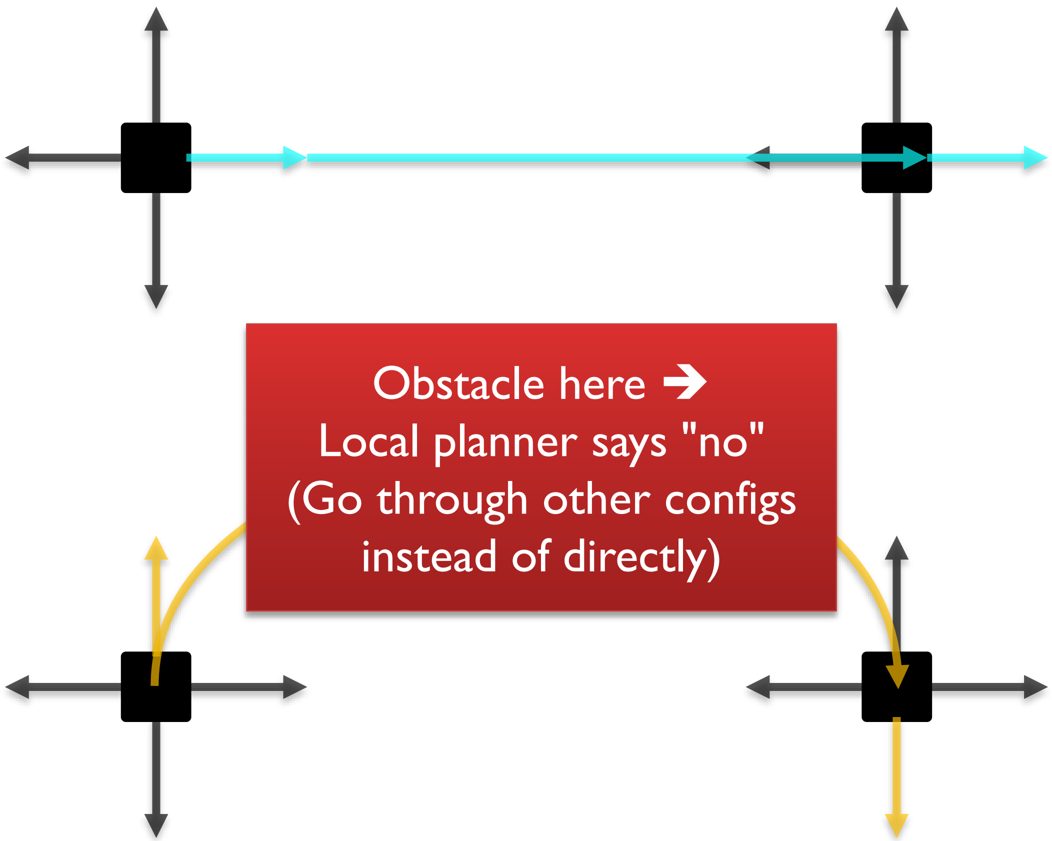

High-Dimensional Problems

High-Dimensional Problems

- For an **aircraft**, a configuration could consist of:
  - **location** in 3D space ($x/y/z$)
  - **pitch angle**
  - **yaw angle**
  - **roll angle**

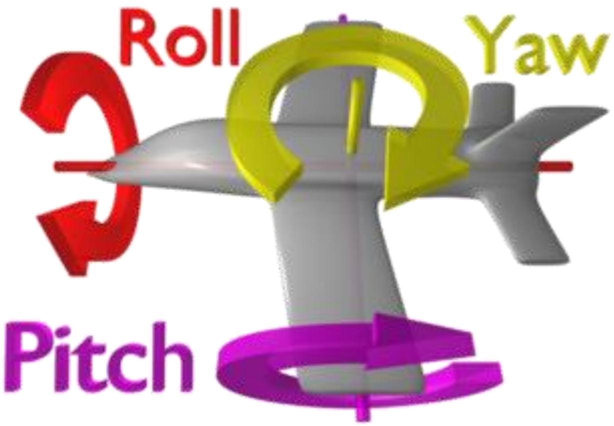

- A **path** is:
  - a continuous **curve** in 6-dimensional configuration space
    avoiding **obstacles**
    and obeying **constraints** on how the aircraft can turn
    - Can make tighter turns at low speed
    - Can't fly at arbitrary pitch angles
    - …

High-Dimensional Problems (2)

- For a **robot arm**, a configuration could consist of:
  - The position / angle of each joint
- A **path** is a continuous **curve** in n-dimensional configuration space (all joints move continuously to new positions, without "jumping"), avoiding **obstacles** and obeying **constraints** on joint endpoints etc.
- Typical goal: Reach inside the car you are painting / welding, without colliding with the car itself

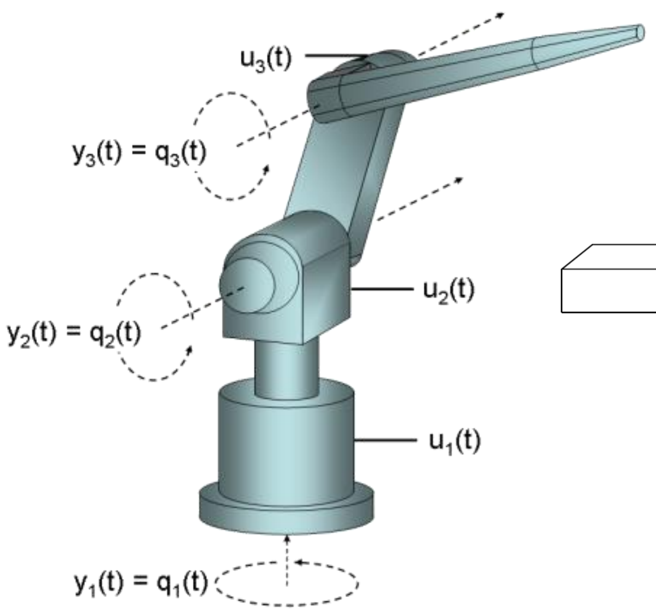

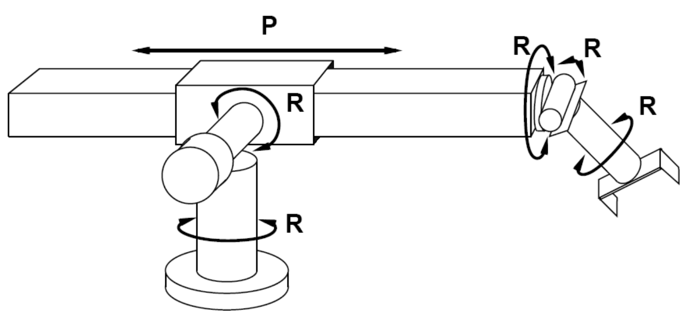

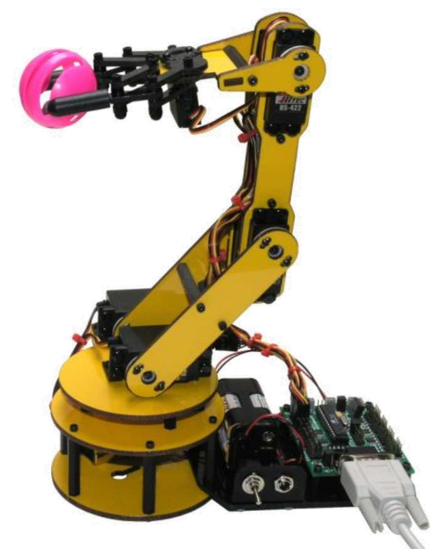

High-Dimensional Problems (3)

- Moving in tight spaces, again…

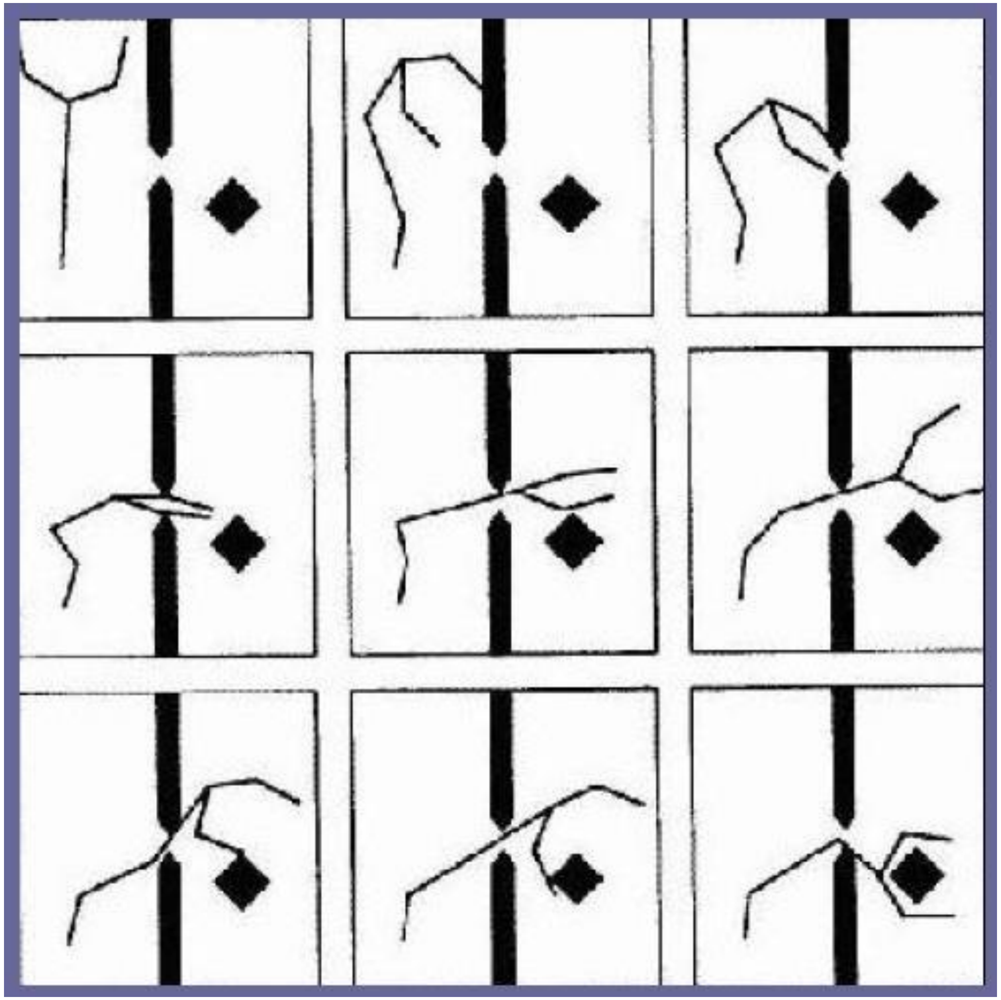

High-Dimensional Problems (4)

- For a **humanoid robot**, a configuration could consist of:
  - Position in x/y space
  - The position of each joint
- The Nao robot:
  - 14, 21 or 25 degrees of freedom depending on model
  - Up to 25-dimensional motion planning!
- Grid methods generally do not scale
  - 25-dimensional configuration space, with 1000 cells in each direction: 10^{75} cells…

High-Dimensional Problems (5)

- Honda Asimo: 57 DOF

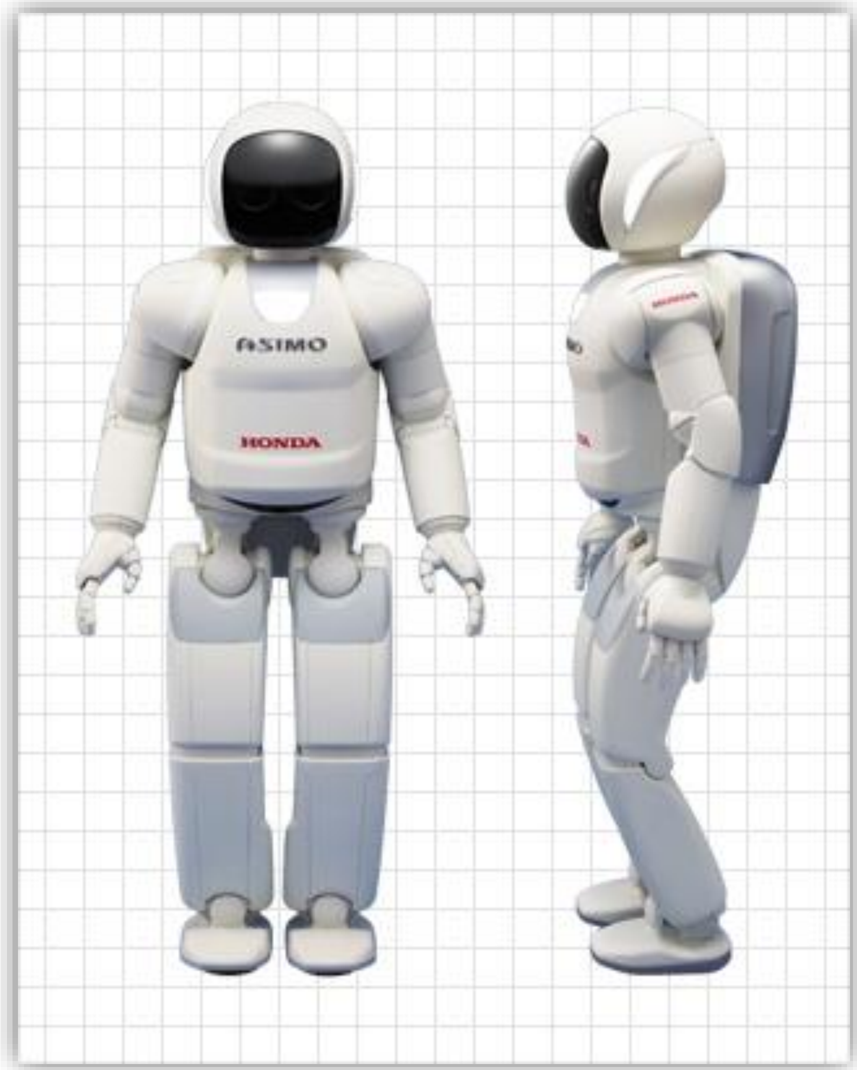

We can often omit some DOF from planning…

But then we don't use the robot's full capabilities!

Alpha Puzzle: Narrow Passages

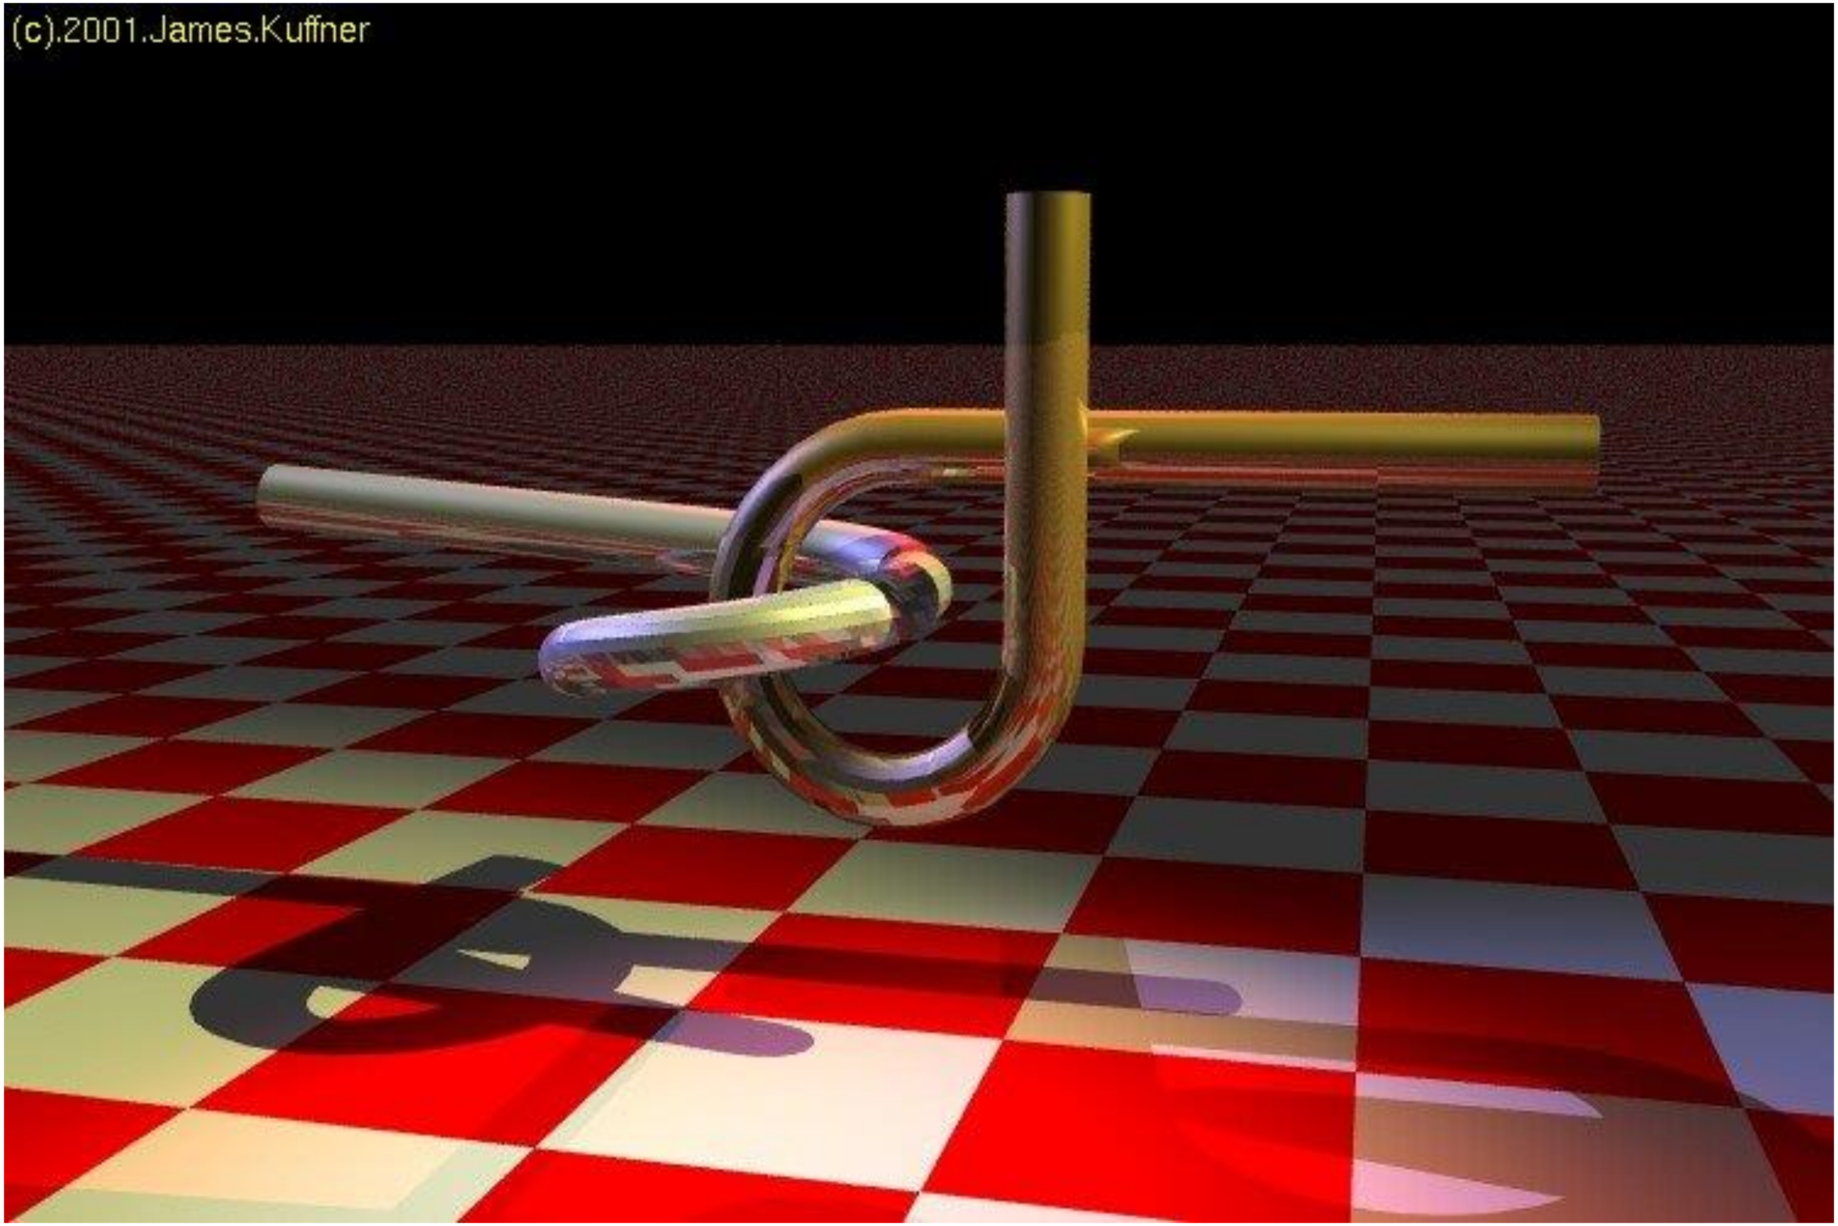

Choosing Potential Configurations: Probabilistic Methods

Preliminaries: Coverage Domain

- Given a **configuration** q in the **free config space**:
  - A particular **local planner** can connect it to a *set of other configs*
  - Called the **coverage domain** $D(q)$ – generally an infinite set

Can connect q to
any config in the green area

Can't connect q to
any other points

Example: Simple *2D planning*,
local planner uses *straight lines*…

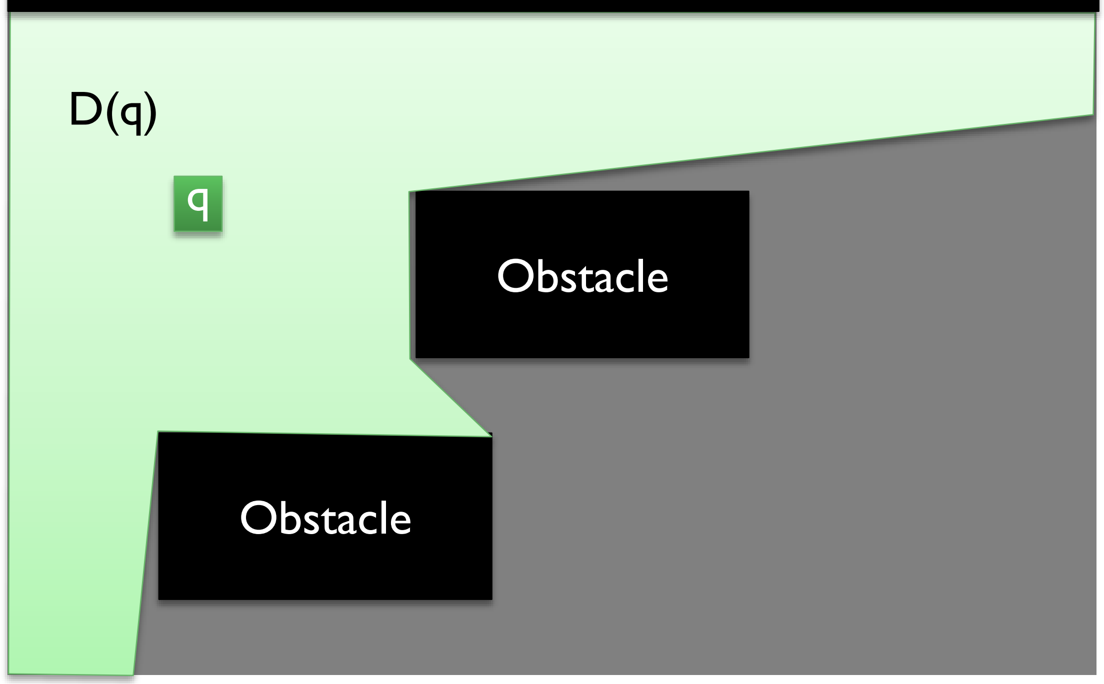

Preliminaries: Preprocessing

- **<u>Preprocessing</u>**: *<u>Suppose</u>* we can select configurations so that:
  - Their domains **<u>cover</u>** the entire config space
  - The configs can be **<u>connected</u>**

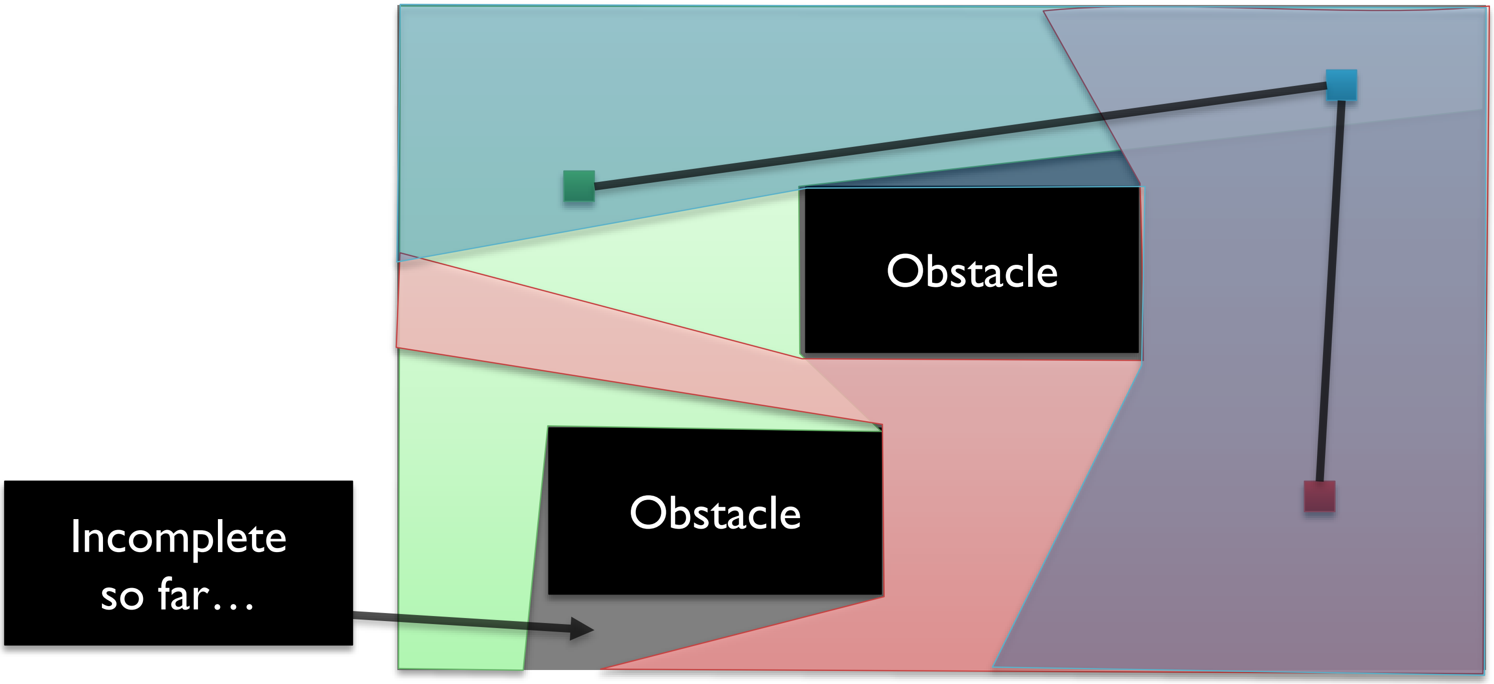

(Imagine many obstacles, hundreds or thousands of configurations, many *dimensions*…)

Preliminaries: Solving

- **Solving**: We get…
  - Start configuration q_{start}
    - **Connect** to another configuration
    - Must be possible:
      The ***domains*** of the existing configurations ***covered the entire space***
  - Goal configuration q_{goal}
    - Connect…
  - Find a path
    through the graph!

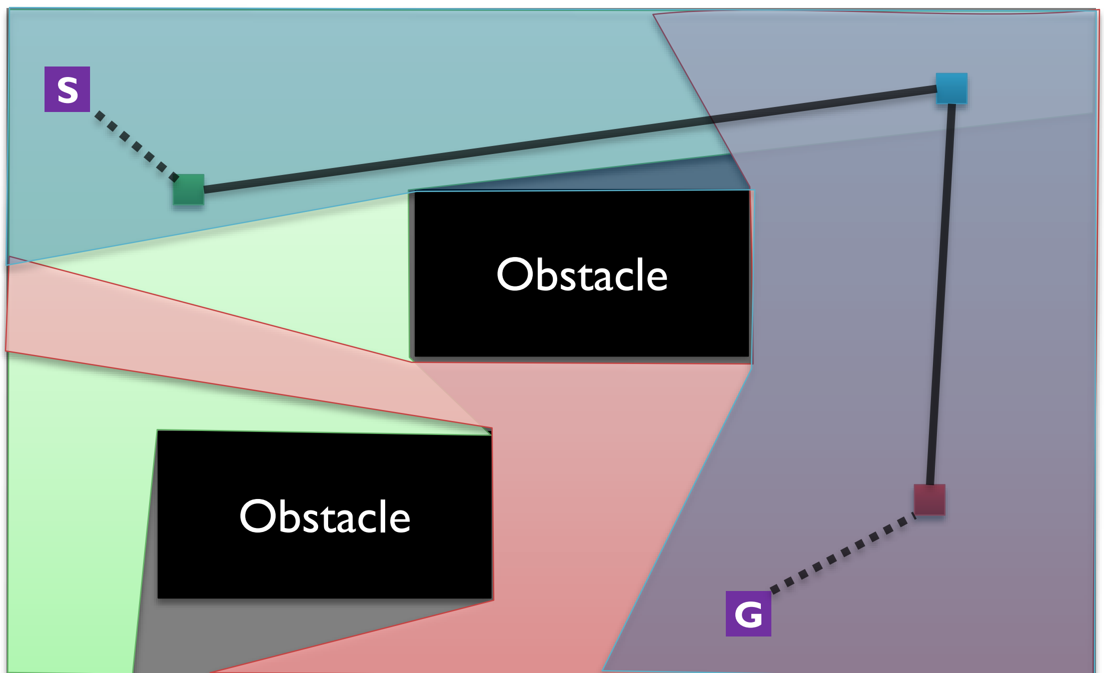

Preliminaries: Coverage Domains are Implicit

- Problem: We **can't calculate** the coverage domain $D(q)$
  - Local planner answers "can you connect q_1 with the **specific config** q_2?
  - Computing "all the configurations you can connect q_1 to":
    - High-dimensional spaces (57D???)
    - Complex motion constraints, not just physical obstacles
    - Too computationally complex, even if finite
    - Usually *infinitely* many possibilities

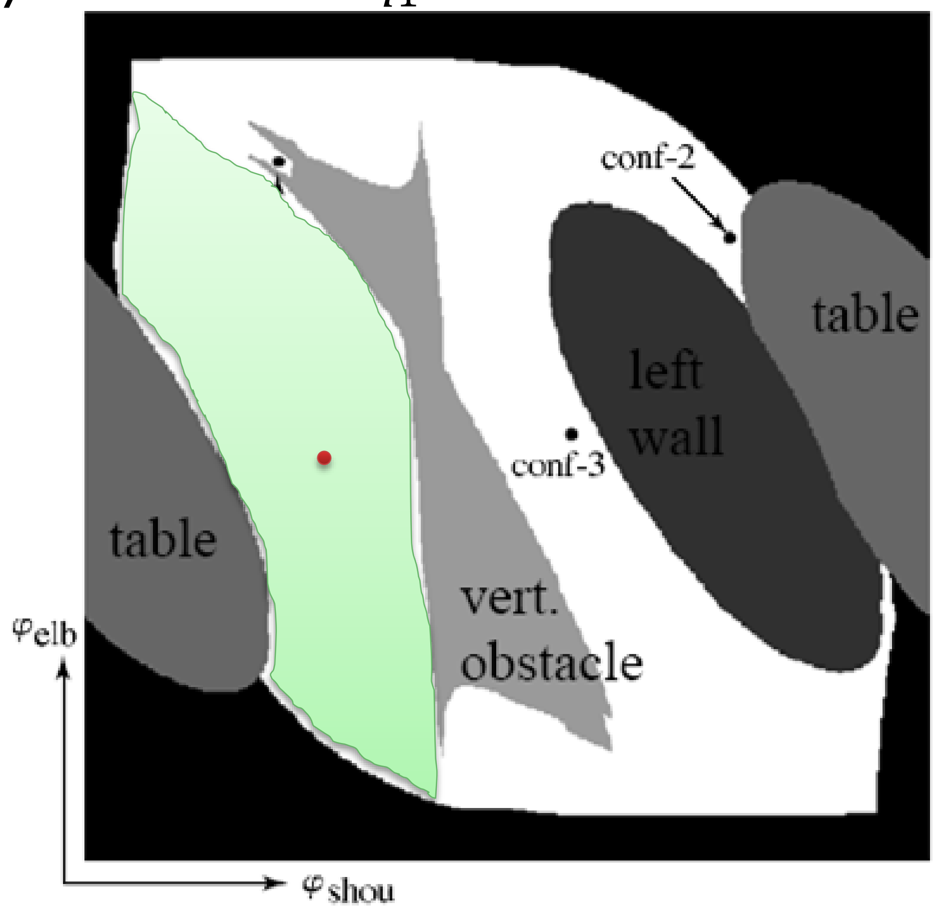

Preliminaries: Probabilistic Methods

- Solution: **Probabilistic methods**
  - Given a set of configurations $Q = \{q_1, \ldots, q_n\}$:
    - Don't compute

$$\bigcup_{q \in Q} D(q)$$

    - Directly compute **probability**:

$$P\left(\bigcup_{q \in Q} D(q) \text{ covers entire free configuration space}\right)$$

    - Or:

$$P\left(\text{if you pick a random free config, it belongs to } \bigcup_{q \in Q} D(q)\right)$$

    - Add configurations until probability is sufficiently high

Probabilistic Roadmaps

(Lydia Kavraki et al, 1996)

Probabilistic Roadmaps

- Probabilistic Roadmaps (PRM): **Construction Phase**
  - M ← empty roadmap
    - **do** {

      randomly generate configuration *q* in free config space

      if (*q* was previously unreachable, so it would extend coverage) {

      add *q* and associated edges to M

      } **else if** (*q* was reachable, but now connects two previously unconnected configs) {

      add *q* and associated edges to M

      }

      } **until** (sufficient coverage)

A new config here would *not* be added!

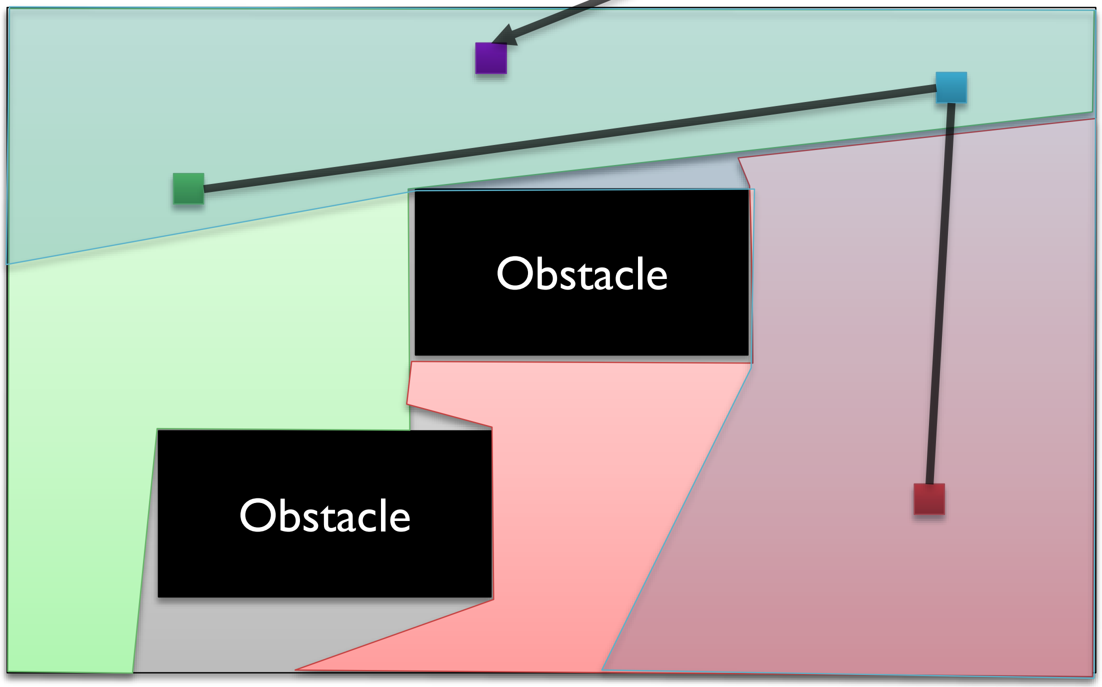

Tweaks:
Only consider points within a maximum distance;
only consider up to N neighbors;
…

PRM: Sufficient Coverage

- When do you have **sufficient coverage**?
  - Suppose you have tested n configurations in a row *without being able to add one* to the road map
  - Then the roadmap covers the free config space with probability $1 - \frac{1}{n}$
    - Example: $n = 1000$ ➔ likely that 99.9% of the free config space is covered
- Why generate *randomly*? Why don't we ***select** a non-covered config*?
  - How? Many dimensions, complex connectivity, …
  - Random ➔ no need to explicitly calculate coverage domains!
- Construction phase done **in advance**
  - In a sense, a **learning phase**
  - Road map reused for many queries

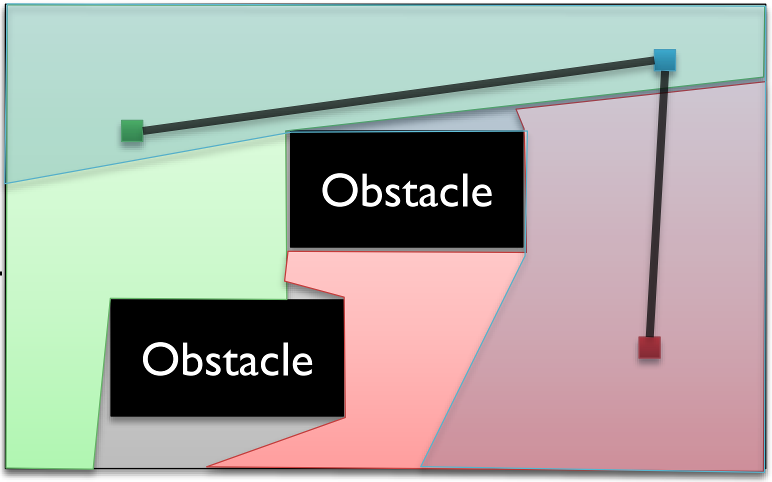

PRM: Construction in Advance

- Construction phase typically done **in advance**
  - In a sense, a **learning phase**
  - Road map reused for many queries
- But we can **improve** the road map later!
  - No solution? Add more nodes.
  - Detect new obstacles? Remove edges.
  - …

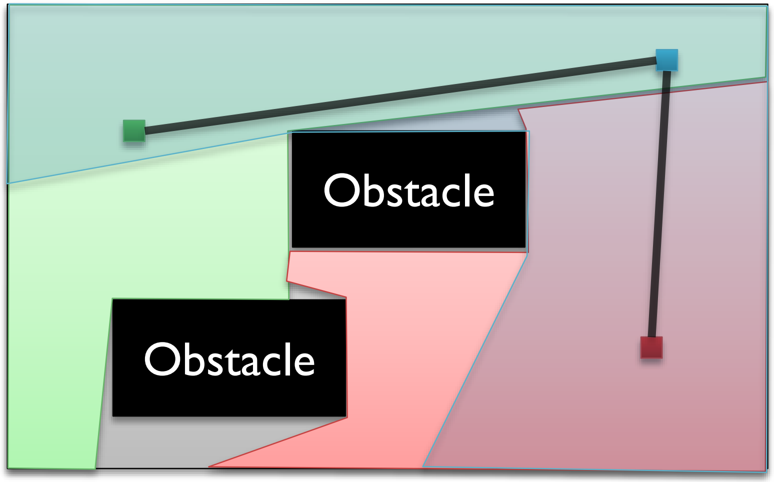

PRM: Node Placement

- Node placement is *random* but not always *uniform*
  - Can be *biased* towards difficult areas

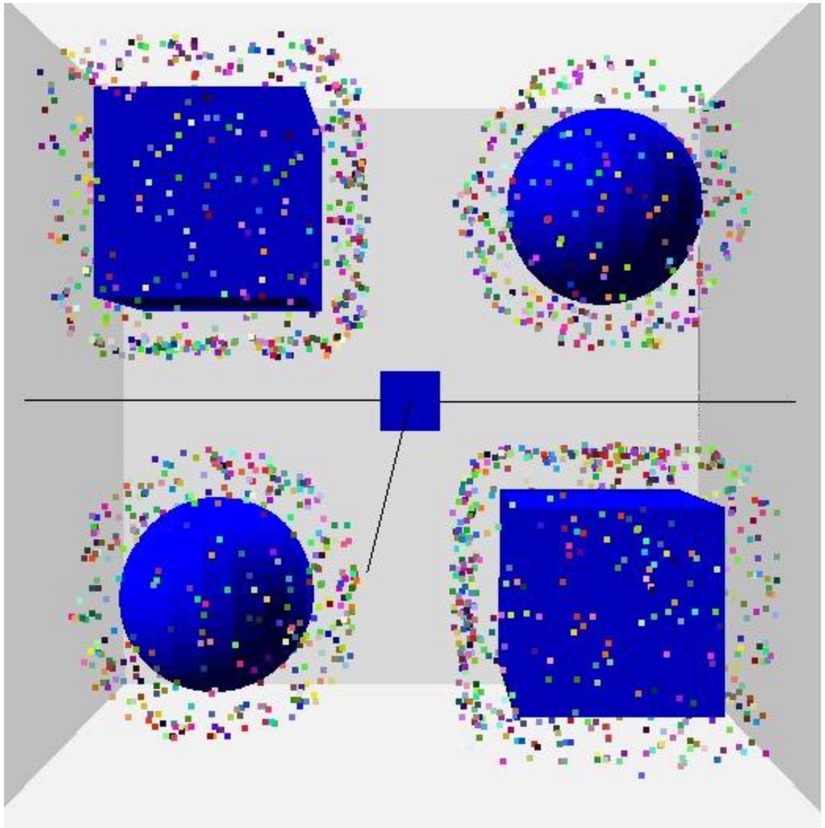

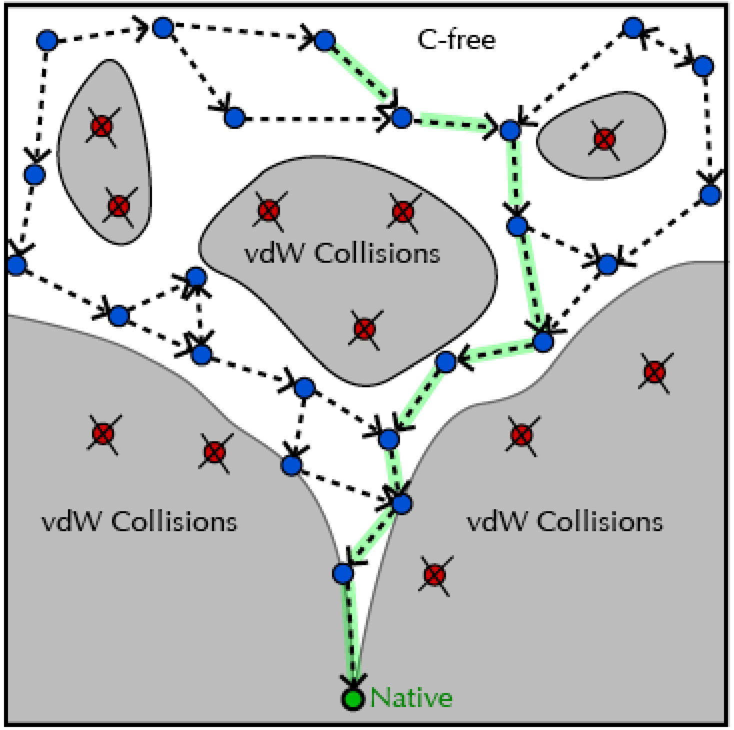

The "obstacles" above are "obstacles" in **configuration space**!

PRM: Protein Folding

- (Second example was from a protein folding application…)

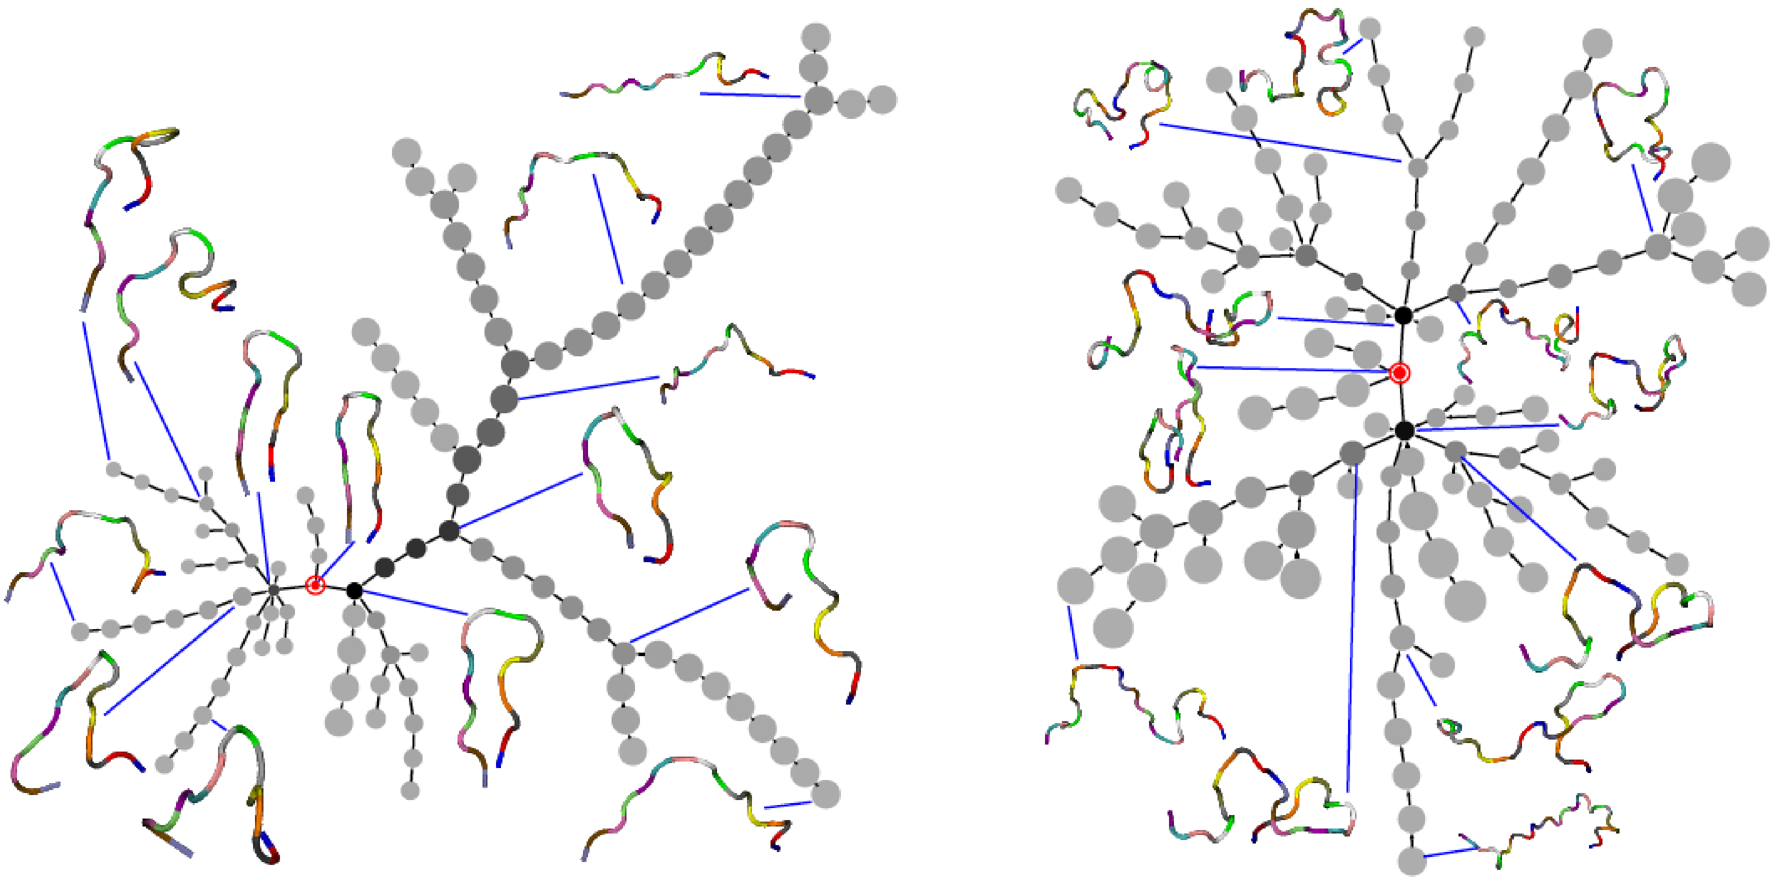

2GB1(16) 1E0L(28)

PRM: Query Phase

- Query Phase:

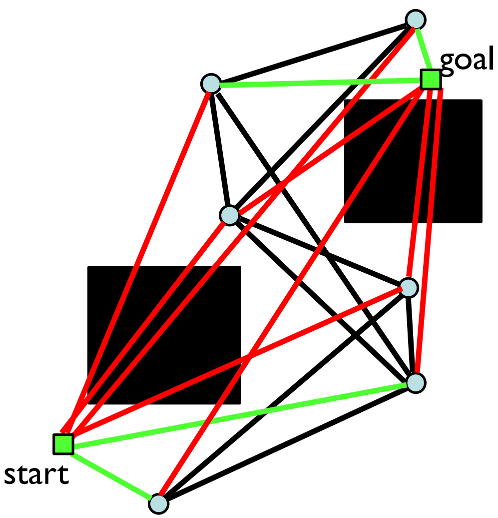

Add and connect start and goal configs to the roadmap (should be possible, as we have good coverage)

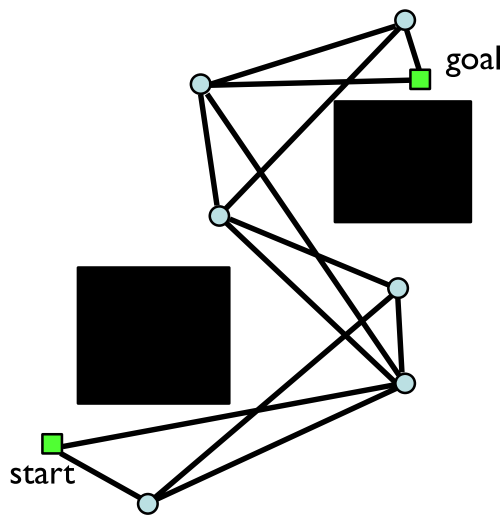

A* search

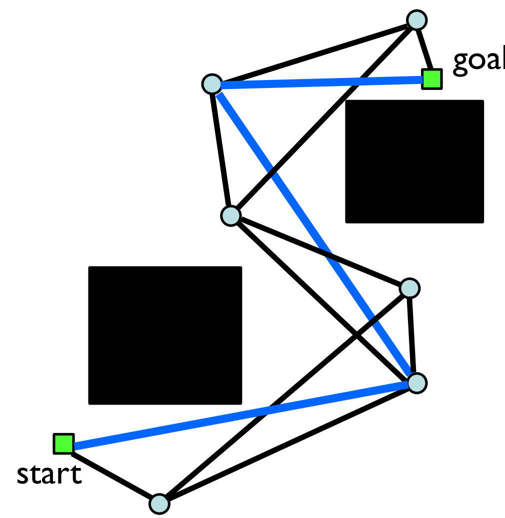

PRM: Result

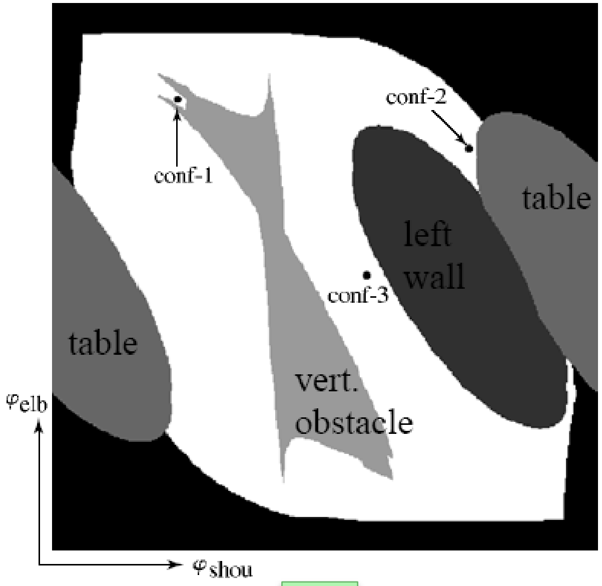

Visualized i 2D
Could be 25D
Even in 2D, we have no *closed form description* of the shape – must *sample*!

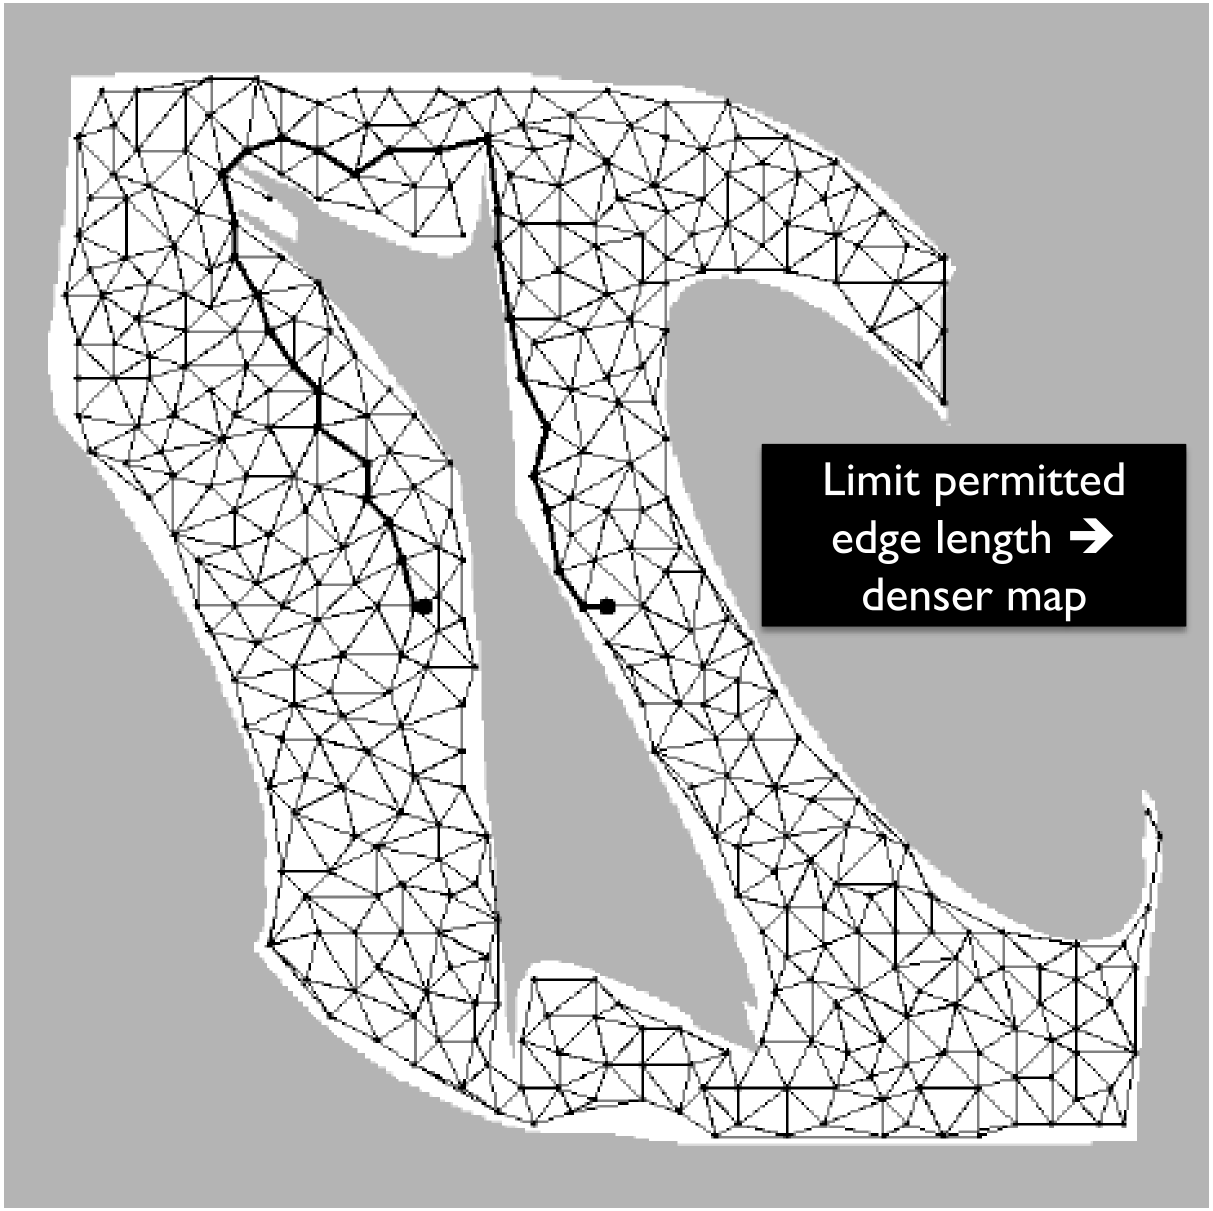

Limit permitted edge length → denser map

PRM: Properties

- **Properties**:
  - Scales better to higher dimensions
  - Deterministically incomplete, **probabilistically complete**
    - The more configurations you create,
      the greater the probability that a path can be found if one exists
      (approaching 1.0)

Graph Search

Adapting to New Obstacles

- Suppose **new obstacles** are detected during execution
  - A*: Update map and replan from scratch
    - Inefficient
  - D* (Dynamic A*): Informed **incremental** search
    - First, find a path using information about known obstacles
    - When new obstacles are detected:
      - Affected nodes are returned to the OPEN list, marked as RAISE: More expensive than before
      - Incrementally updates only those nodes whose cost change due to the new obstacles
  - Focused D*:
    - Focuses propagation towards the robot – additional speedup
  - …

Anytime Search

- **Anytime** algorithms:
  - ***Be able to answer whenever I interrupt you!***
  - In practice: ***Create*** some path quickly, then incrementally ***improve*** it
  - "Repeated weighted A*" (standard technique)
    - Run A* with $f(n) = g(n) + \boldsymbol{W} \cdot h(n)$, where $W > 1$: Faster but suboptimal

$w = 1$:
Standard A*

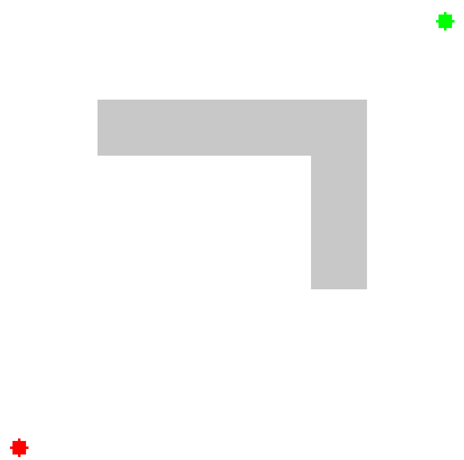

$w = 5$:
Distance from goal is exaggerated
➔
suboptimal

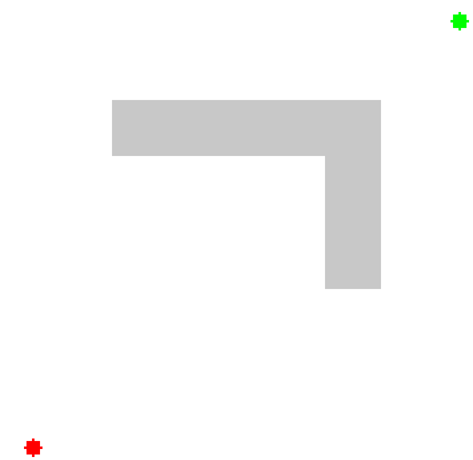

- Decrease W and **repeat**
- But: Has to **redo search** from scratch in each run!

Anytime Search (2)

- **Anytime** algorithms:
  - Anytime Repairing A*
    - Like "repeated weighted A*", but reuses search results from earlier iterations
  - Anytime Dynamic A* (AD*)
    - **Both** replanning when problems change and anytime planning
  - …

Post-Processing: Path Smoothing

Suboptimal Paths

- Paths are often **suboptimal** in the continuous space
  - Only the chosen points in the cells are used
  - In this example: The midpoints

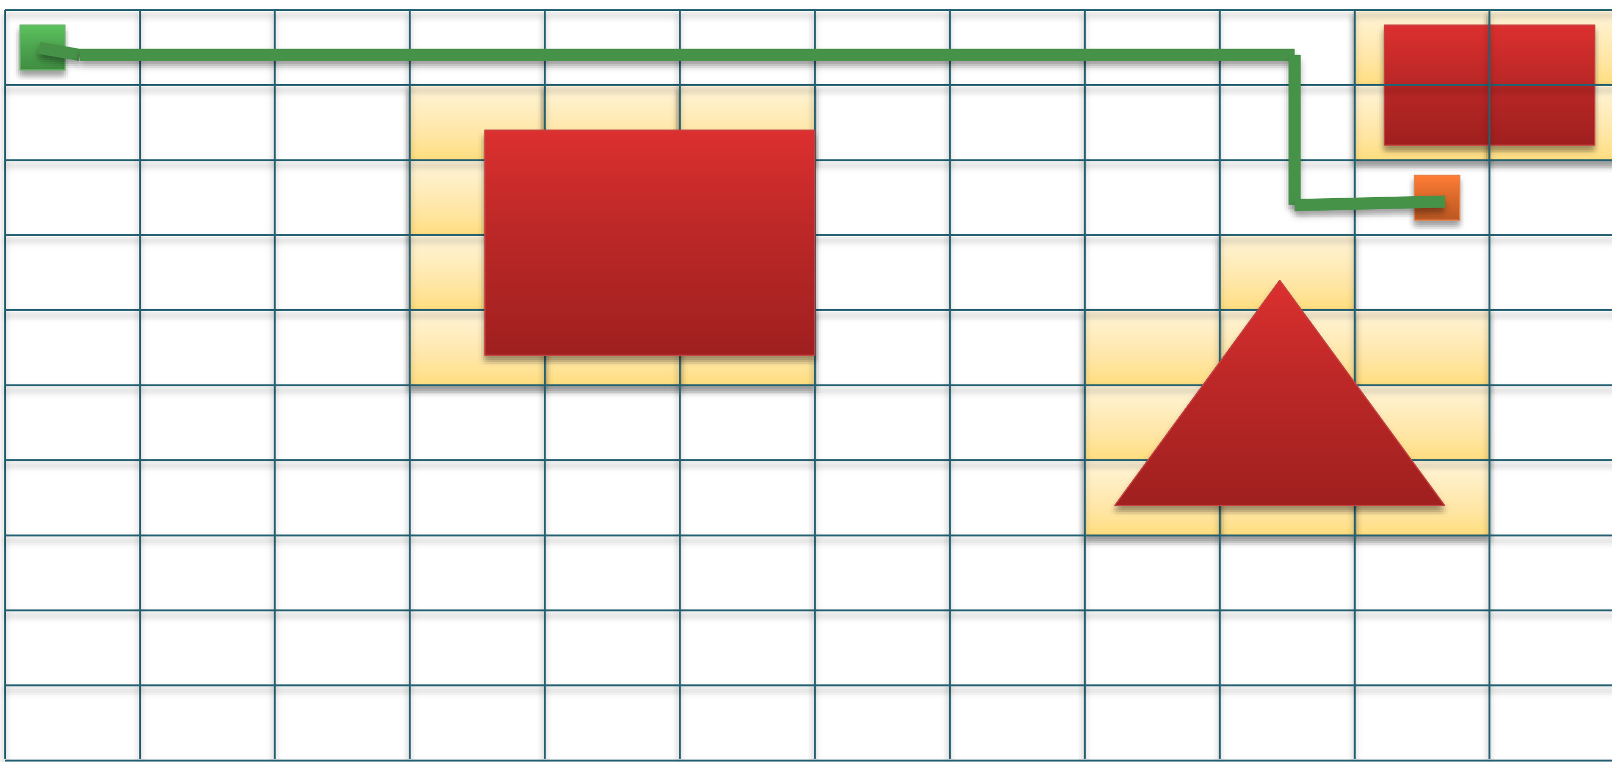

Smoothing

- Paths can be improved through **smoothing** after generation
  - Still generally does not lead to optimal paths
  - This is just a simple example, where smoothing is easy

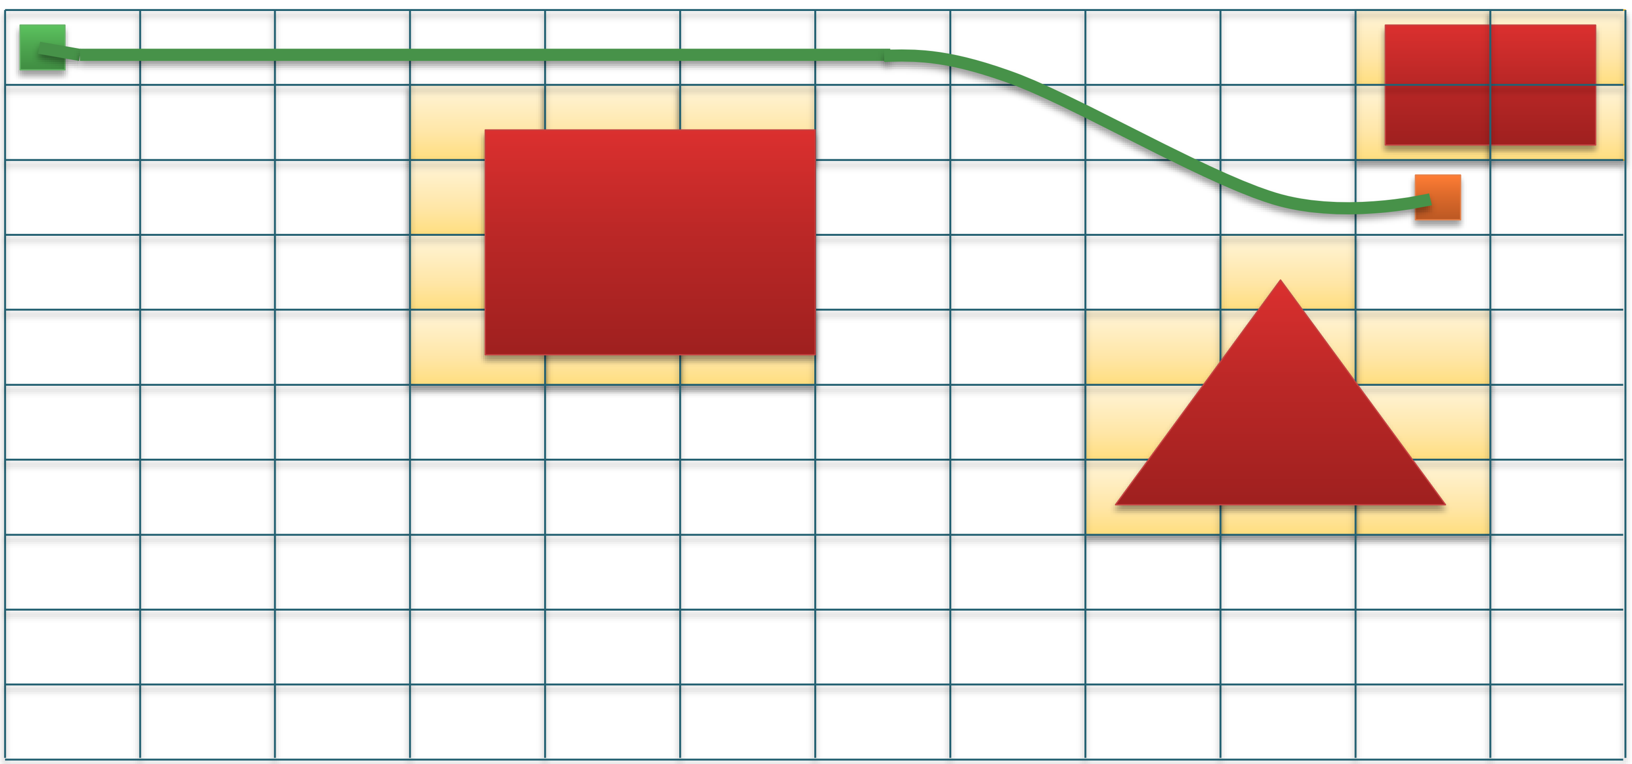

Open Motion Planning Library

- Want to experiment?
  - Open Motion Planning Library
  - http://ompl.kavrakilab.org/index.html