

Motion Planning

(It's all in the discretization)

R&N: Chap. 25 gives some background

1

Motion planning is the ability for an agent to compute its own motions in order to achieve certain goals. All autonomous robots and digital actors should eventually have this ability



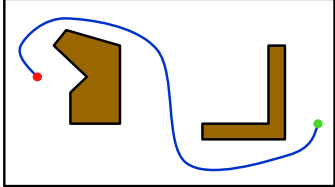
2

Digital Actors

- [video 1](#)
- [video 2](#)

3

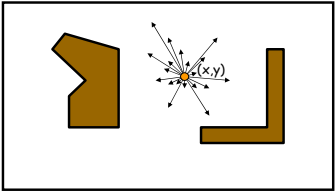
Basic problem



- Point robot in a 2-dimensional workspace with obstacles of known shape and position
- Find a collision-free path between a start and a goal position of the robot

4

Basic problem

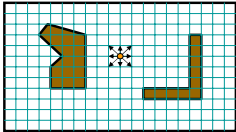
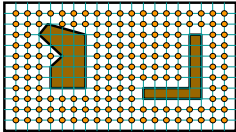


- Each robot position (x,y) can be seen as a state
- → *Continuous* state space
- Then each state has an infinity of successors
- We need to *discretize* the state space

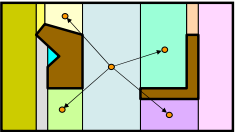
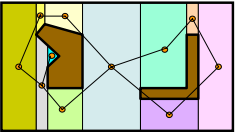
5

Two Possible Discretizations

Grid-based





Criticality-based





6

Two Possible Discretizations




Intruder Finding Problem




- A moving intruder is hiding in a 2-D workspace
- The robot must "sweep" the workspace to find the intruder
- Both the robot and the intruder are points

Does a solution always exist?

No !




Easy to test:
"Hole" in the workspace



Hard to test:
No "hole" in the workspace


9

Information State




- Example of an information state = $(x,y,a=1,b=1,c=0)$
- An **initial state** is of the form $(x,y,1,1, \dots, 1)$
- A **goal state** is any state of the form $(x,y,0,0, \dots, 0)$

Critical Line



11


Criticality-Based Discretization



Each of the regions A, B, C, D, and E consists of "equivalent" positions of the robot, so it's sufficient to consider a single position per region


12

Criticality-Based Discretization




13

Criticality-Based Discretization




14

Criticality-Based Discretization




15

Criticality-Based Discretization



16


Grid-Based Discretization



- Ignores critical lines → Visits many "equivalent" states
- Many information states per grid point
- Potentially very inefficient

17

Example of Solution



18


But ...

Criticality-based discretization does not scale well in practice when the dimensionality of the continuous space increases

(It becomes prohibitively complex to define and compute)

19


Motion Planning for an Articulated Robot



Find a path to a goal configuration that satisfies various constraints: collision avoidance, equilibrium, etc...

20


Configuration Space of an Articulated Robot



- A **configuration** of a robot is a list of non-redundant parameters that fully specify the position and orientation of each of its bodies
 - In this robot, one possible choice is: (q_1, q_2)
- The **configuration space (C-space)** has 2 dimensions

21

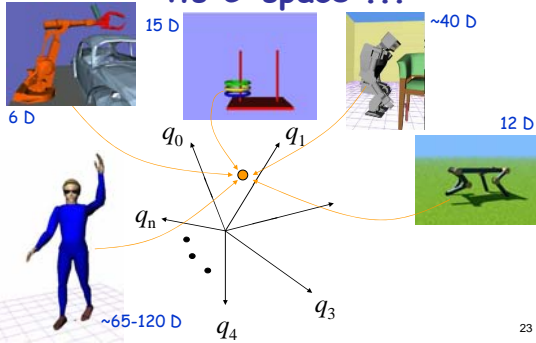
How many dimensions has the C-space of these 3 rings?



Answer:
 $3 \times 5 = 15$

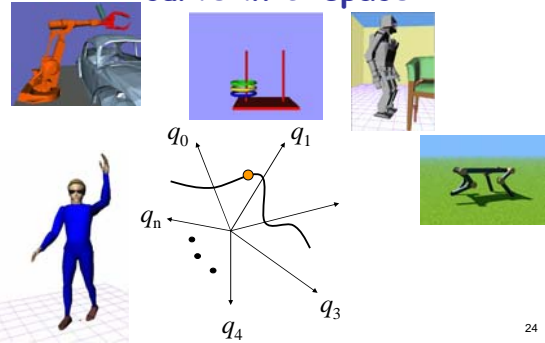
22

Every robot maps to a point in its C-space ...



23

... and every robot path is a curve in C-space



24

A robot path is a curve in C-space

So, the C-space is the continuous state space of motion planning problems

25

C-space "reduces" motion planning to finding a path for a point

But how do the obstacle constraints map into C-space ?

26

A Simple Example: Two-Joint Planar Robot Arm

Problems:

- Geometric complexity
- Space dimensionality

27

Continuous state space

↓

Discretization

↓

Search

C-space

28

About Discretization

- Dimensionality + geometric complexity
 - Criticality-based discretization turns out to be prohibitively complex
- Dimensionality
 - Grid-based discretization leads to impractically large state spaces for $\dim(C\text{-space}) > 6$
 - Each grid node has $3^n - 1$ neighbors, where $n = \dim(C\text{-space})$

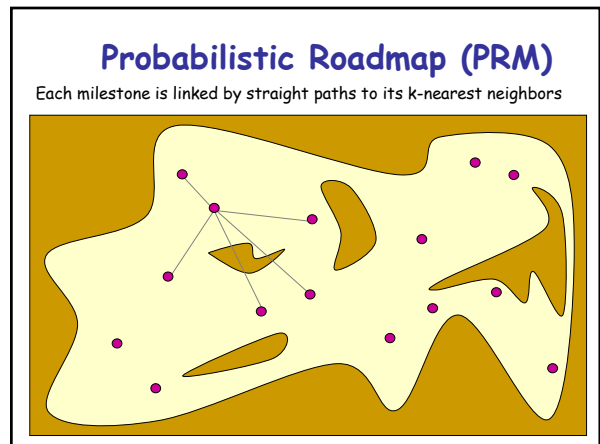
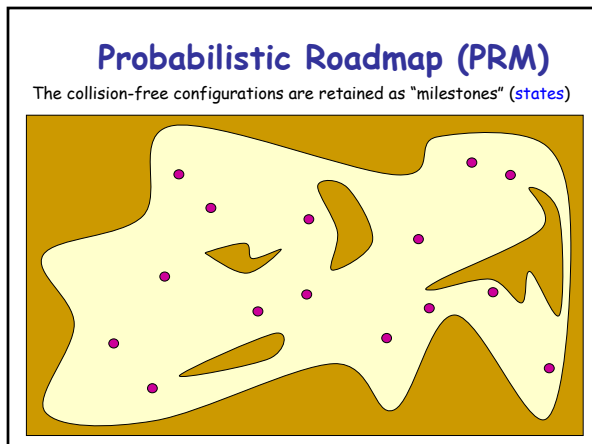
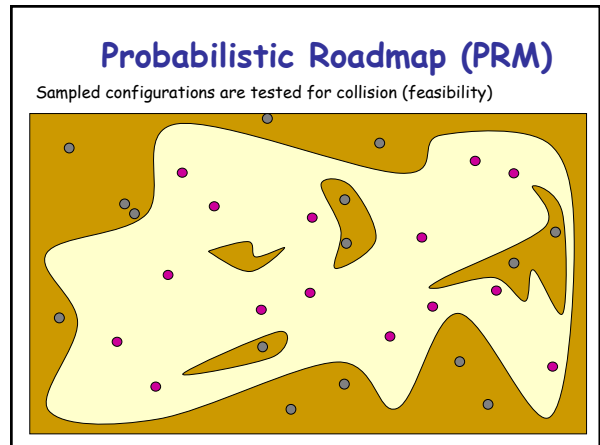
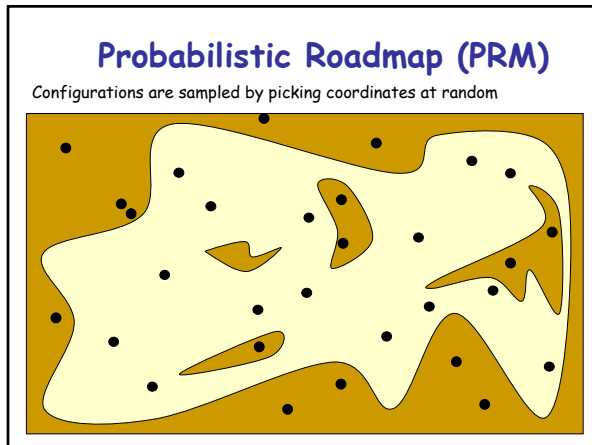
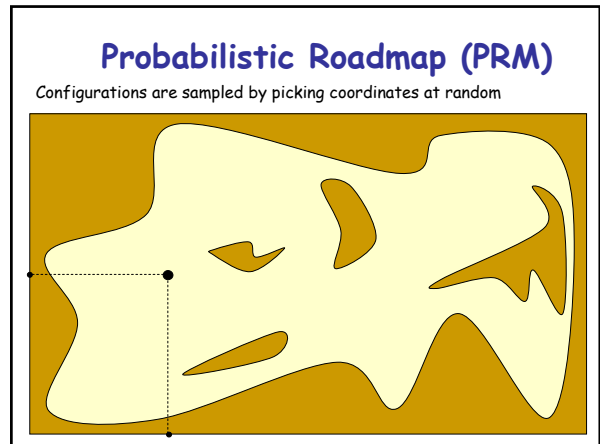
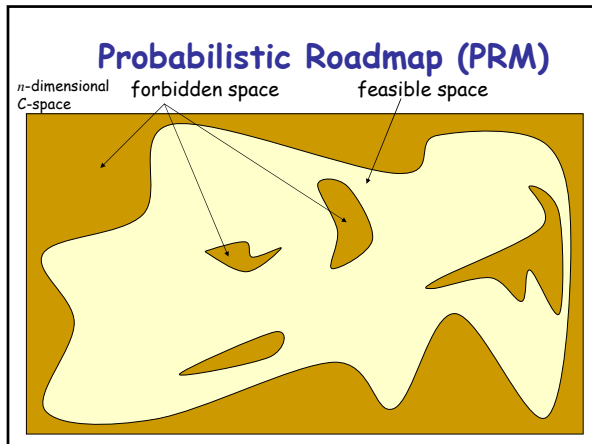
29

Robots with many joints: Modular Self-Reconfigurable Robots

Millipede-like robot with 13,000 joints


(M. Yim) (S. Redon)

30




Probabilistic Roadmap (PRM)

Each milestone is linked by straight paths to its k-nearest neighbors




Probabilistic Roadmap (PRM)

The collision-free links are retained to form the PRM (state graph)




Probabilistic Roadmap (PRM)

The start and goal configurations are connected to nodes of the PRM



Probabilistic Roadmap (PRM)


The PRM is searched for a path from s to g



Continuous state space



Discretization



Search A*

41


Why Does PRM Work?

Because most feasible spaces verifies some good geometric (visibility) properties

42

Why Does PRM Work?


In most feasible spaces, every configuration "sees" a significant fraction of the feasible space



→ A relatively small number of milestones and connections between them are sufficient to cover most feasible spaces with high probability


43

Narrow-Passage Issue



The **lookout** of a subset S of the feasible space is the set of all configurations in S from which it is possible to "see" a significant fraction of the feasible space outside S

The feasible space is **expansive** if all of its subsets have a large lookout



45

Probabilistic Completeness of a PRM Motion Planner

In an expansive feasible space, the probability that a PRM planner with uniform sampling strategy finds a solution path, if one exists, goes to 1 exponentially with the number of milestones (\sim running time)

A PRM planner can't detect that no path exists. Like A^* , it must be allocated a **time limit** beyond which it returns that no path exists. But this answer may be **incorrect**. Perhaps the planner needed more time to find one!


46

Sampling Strategies

- **Issue:** Where to sample configurations? That is, which probabilistic distribution to use?


- **Example:** Two-stage sampling strategy:

1. Construct initial PRM with uniform sampling
2. Identify milestones that have few connections to their close neighbors
3. Sample more configurations around them



→ Greater density of milestones in "difficult" regions of the feasible space

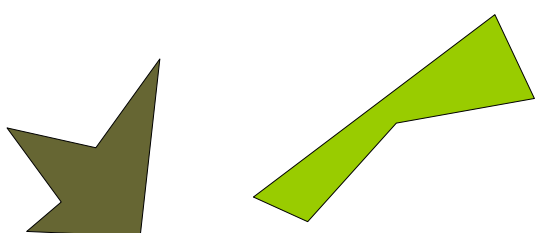
47



48

Collision Checking

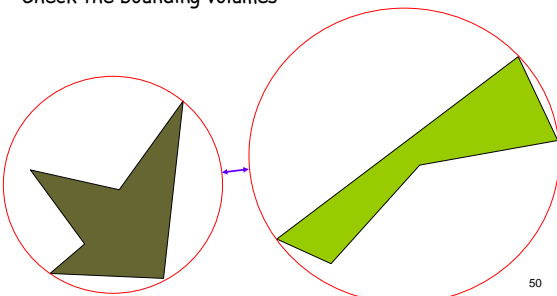
- Check whether objects overlap



49

Hierarchical Collision Checking

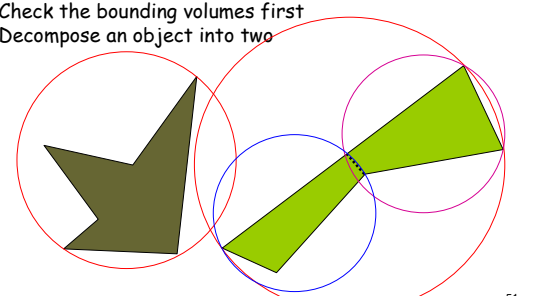
- Enclose objects into bounding volumes (spheres or boxes)
- Check the bounding volumes



50

Hierarchical Collision Checking

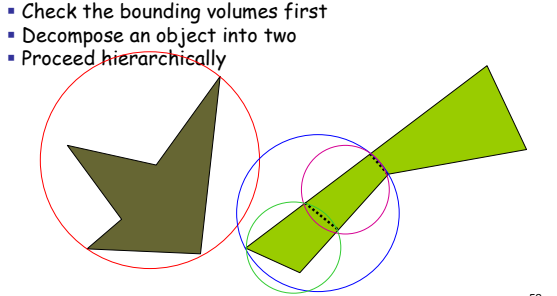
- Enclose objects into bounding volumes (spheres or boxes)
- Check the bounding volumes first
- Decompose an object into two



51

Hierarchical Collision Checking

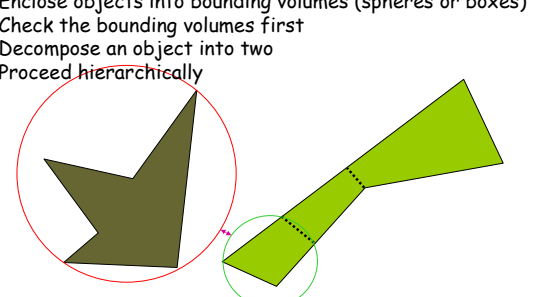
- Enclose objects into bounding volumes (spheres or boxes)
- Check the bounding volumes first
- Decompose an object into two
- Proceed hierarchically



52

Hierarchical Collision Checking

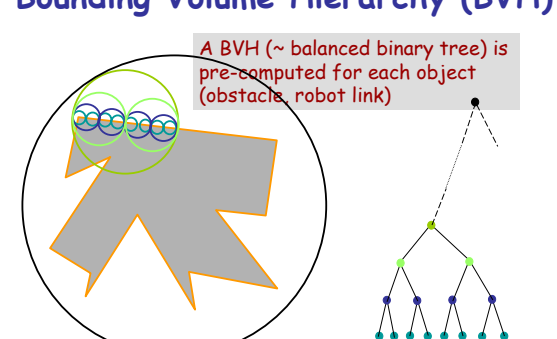
- Enclose objects into bounding volumes (spheres or boxes)
- Check the bounding volumes first
- Decompose an object into two
- Proceed hierarchically



53

Bounding Volume Hierarchy (BVH)

A BVH (~ balanced binary tree) is pre-computed for each object (obstacle, robot link)



54

BVH of a 3D Triangulated Cat



55

Collision Checking Between Two Objects



BVH of object 1

BVH of object 2


[Usually, the two trees have different sizes]

→ Search for a collision

56

Search for a Collision

Search tree



57

Search for a Collision

Search tree




Heuristic: Break the largest BV




58

Search for a Collision

Search tree




Heuristic: Break the largest BV



59

Search for a Collision


Search tree



60

Search for a Collision

Search tree




If two leaves of the BVH's overlap (here, **C** and **B**) check their content for collision

61

Search Strategy

- If there is no collision, all paths must eventually be followed down to pruning or a leaf node
- But if there is collision, one may try to detect it as quickly as possible
- → **Greedy best-first search strategy** with $f(N) = h(N) = d/(r_x+r_y)$

[Expand the node **XY** with largest relative overlap (most likely to contain a collision)]



62

So, to discretize the state space of a motion planning problem, a PRM planner performs **thousands of auxiliary searches** (sometimes even more) to detect collisions!

But from an outsider's point of view the search of the PRM looks like the main search

63

Fortunately, hierarchical collision checkers are quite fast

On average, over **10,000 collision checks per second** for two 3-D objects each described by 500,000 triangles, on a contemporary PC

Checks are much faster when the objects are either neatly separated (→ early pruning) or neatly overlapping (→ quick detection of collision)

64



65

Free-Climbing Robot




LEMUR IIb robot (created by NASA/JPL)



66

Only friction and internal degrees of freedom are used to achieve equilibrium



67



[Bretl, 2003]


68

Two Levels of Planning

- 1) **One-step planning:**
Plan a path for moving a foot/hand from one hold to another
Can be solved using a PRM planner
- 2) **Multi-step planning:**
Plan a sequence of one-step paths
Can be solved by searching a stance space


69

Multi-Step Planning



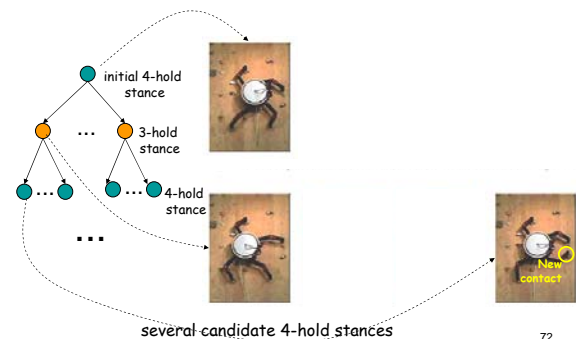
70

Multi-Step Planning



71

Multi-Step Planning



72

Multi-Step Planning

initial 4-hold stance

3-hold stance

4-hold stance

New contact

73

Multi-Step Planning

initial 4-hold stance

3-hold stance

4-hold stance

breaking contact / zero force

74

Multi-Step Planning

one-step planning

initial 4-hold stance

3-hold stance

4-hold stance

breaking contact / zero force

The one-step planner is needed to determine if a one-step path exists between two stances

75

Multi-Step Planning

initial 4-hold stance

3-hold stance

4-hold stance

breaking contact / zero force

76

Multi-Step Planning

initial 4-hold stance

3-hold stance

4-hold stance

breaking contact / zero force

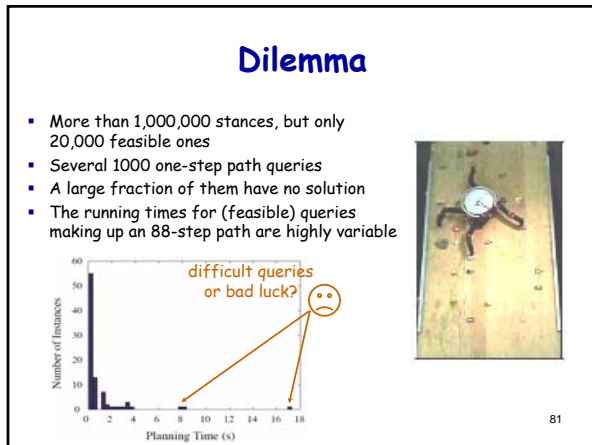
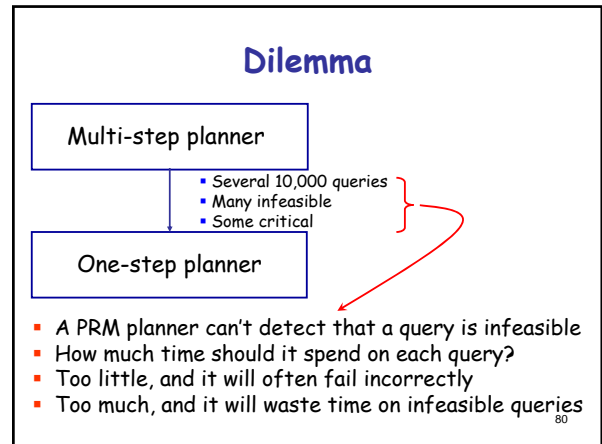
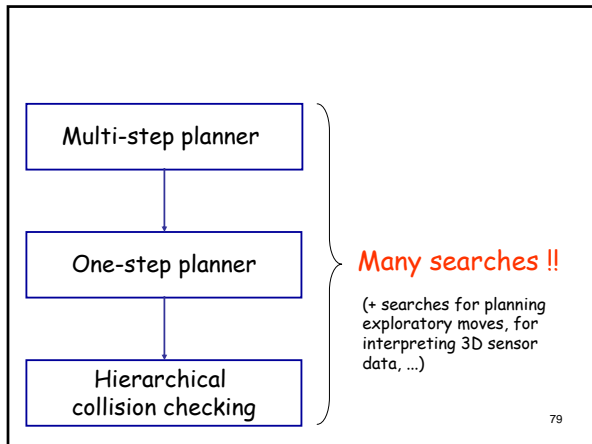
one-step planning

77

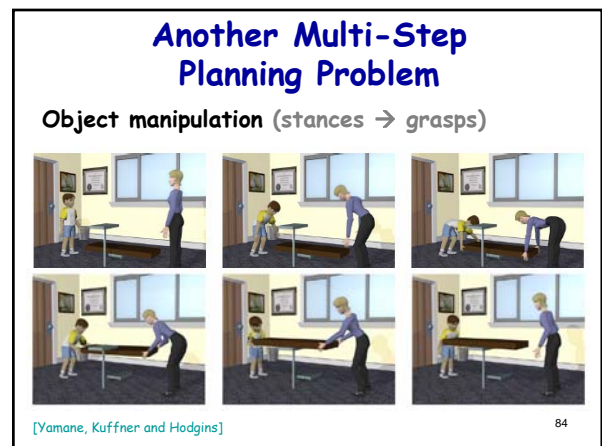
One-Step Planning

- The contact constraints define specific C-space that is easy to sample at random
- It is also easy to test (self-)collision avoidance and equilibrium constraints at sampled configurations
- PRM planning

78





- ### Possible Solution
- Use learning method to train a "feasibility" classifier
 - Use this classifier to avoid infeasible one-step queries in the multi-step search tree
 - More on this later in a lecture on Learning (if there is enough time)
- 82



Some Applications of Motion Planning


85

Design for Manufacturing and Servicing



86

Automatic Robot Programming



ABB


87

Navigation through Virtual Environments



M. Lin, UNC


Virtual Angiography



[S. Napel, 3D Medical Imaging Lab, Stanford]

89

Radiosurgery



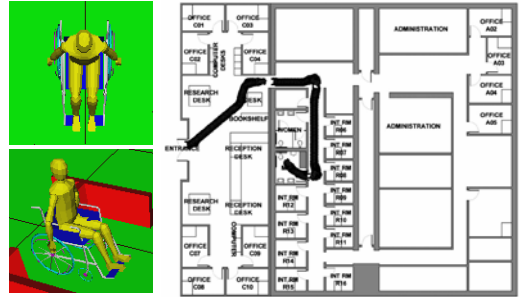
CyberKnife (Accuray)

Transportation of A380 Fuselage through Small Villages



Kineo


Architectural Design: Verification of Building Code



C. Han


92

Architectural Design: Egress Analysis




93

Planet Exploration




Autonomous Digital Actors



[Yamane, Kuffner and Hodgins]

95

Animation of Crowds



Amato

96



97