CSE-P571

Deterministic Path Planning in Robotics

Courtesy of Maxim Likhachev

Carnegie Mellon University

1

Motion/Path Planning

- Task:
  find a feasible (and cost-minimal) path/motion from the current configuration of the robot to its goal configuration (or one of its goal configurations)

- Two types of constraints:
  environmental constraints (e.g., obstacles)
  dynamics/kinematics constraints of the robot

- Generated motion/path should (objective):
  be any feasible path
  minimize cost such as distance, time, energy, risk, …

2

Motion/Path Planning

Examples (of what is usually referred to as path planning):

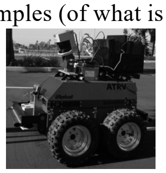

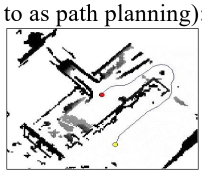

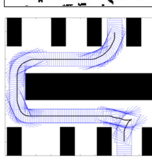

3

Motion/Path Planning

Examples (of what is usually referred to as motion planning):

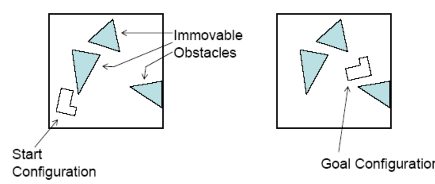

Piano Movers' problem

the example above is borrowed from www.cs.cmu.edu/~awm/tutorials

4

Motion/Path Planning

Examples (of what is usually referred to as motion planning):

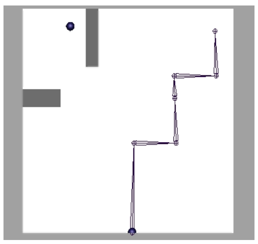

Planned motion for a 6DOF robot arm

5

Motion/Path Planning

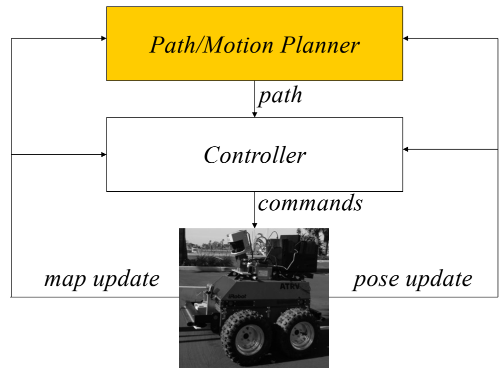

6

Motion/Path Planning

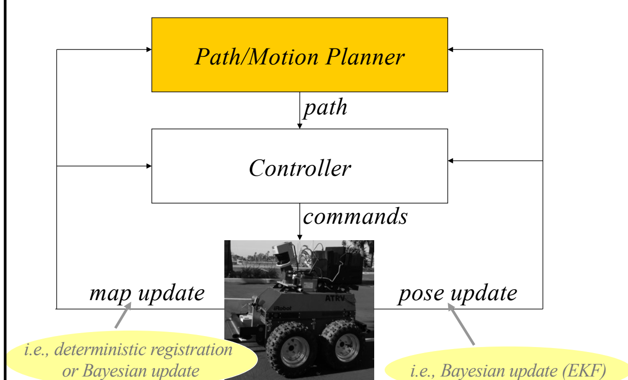

7

Uncertainty and Planning

- Uncertainty can be in:
  - prior environment (i.e., door is open or closed)
  - execution (i.e., robot may slip)
  - sensing environment (i.e., seems like an obstacle but not sure)
  - pose
- Planning approaches:
  - deterministic planning:
    - assume some (i.e., most likely) environment, execution, pose
    - plan a single least-cost trajectory under this assumption
    - re-plan as new information arrives
  - planning under uncertainty:
    - associate probabilities with some elements or everything
    - plan a policy that dictates what to do for each outcome of sensing/action and minimizes expected cost-to-goal
    - re-plan if unaccounted events happen

8

Uncertainty and Planning

- Uncertainty can be in:
  - prior environment (i.e., door is open or closed)
  - execution (i.e., robot may slip)
  - sensing environment (i.e., seems like an obstacle but not sure)
  - pose
- Planning approaches:
  - deterministic planning:
    - assume some (i.e., most likely) environment,
    - plan a single least-cost trajectory under thi
    - re-plan as new information arrives
  - planning under uncertainty:
    - associate probabilities with some elements or everything
    - plan a policy that dictates what to do for each outcome of sensing/action and minimizes expected cost-to-goal
    - re-plan if unaccounted events happen

re-plan every time sensory data arrives or robot deviates off its path

re-planning needs to be FAST

9

Uncertainty and Planning

- Uncertainty can be in:
  - prior environment (i.e., door is open or closed)
  - execution (i.e., robot may slip)
  - sensing environment (i.e., seems like an obstacle but not sure)
  - pose
- Planning approaches:
  - deterministic planning:
    - assume some (i.e., most likely) environment, execution, pose
    - plan a single least-cost trajectory under this assumption
    - re-plan as new information arrives
  - planning under uncertainty:
    - associate probabilities with some elements or everything
    - plan a policy that dictates what to do for each outcome of sensing/action and minimizes expected cost-to-goal
    - re-plan if unaccounted events happen

computationally MUCH harder

10

Example

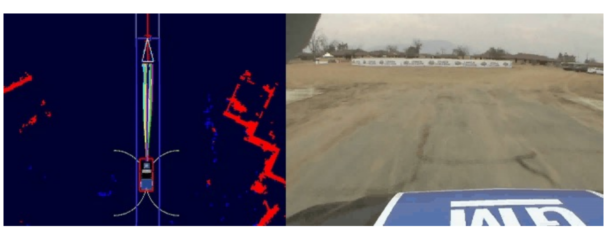

*Urban Challenge Race, CMU team, planning with Anytime D**

11

Outline

- Deterministic planning
  - constructing a graph
  - search with A*
  - search with D*

12

Outline

- Deterministic planning
  - constructing a graph
  - search with A*
  - search with D*

13

Planning via Cell Decomposition

- Approximate Cell Decomposition:
  - overlay uniform grid over the C-space (discretize)

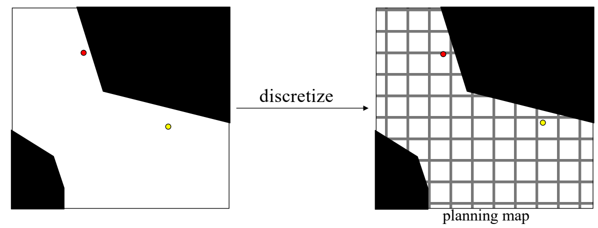

14

Planning via Cell Decomposition

- Approximate Cell Decomposition:
  - construct a graph and search it for a least-cost path

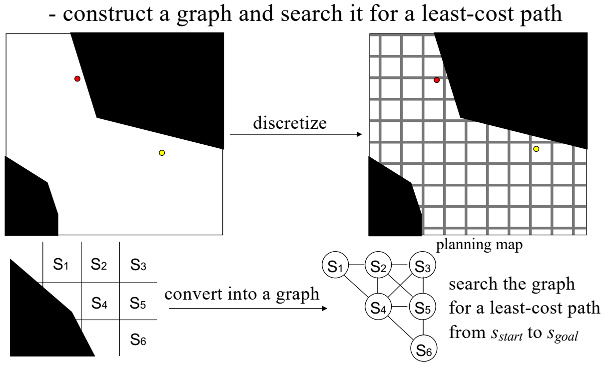

15

Planning via Cell Decomposition

- Approximate Cell Decomposition:
  - construct a graph and search it for a least-cost path

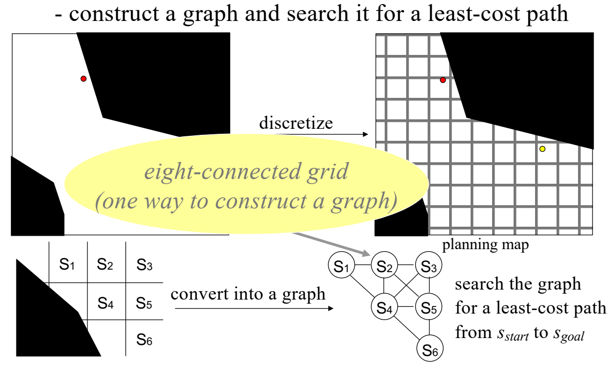

16

Planning via Cell Decomposition

- Approximate Cell Decomposition:
  - construct a graph and search it for a least-cost path
    - VERY popular due to its simplicity and representation of arbitrary obstacles
    - Problem: transitions difficult to execute on non-holonomic robots

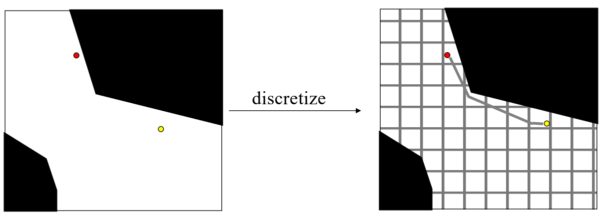

17

Planning via Cell Decomposition

- Graph construction:
  - lattice graph

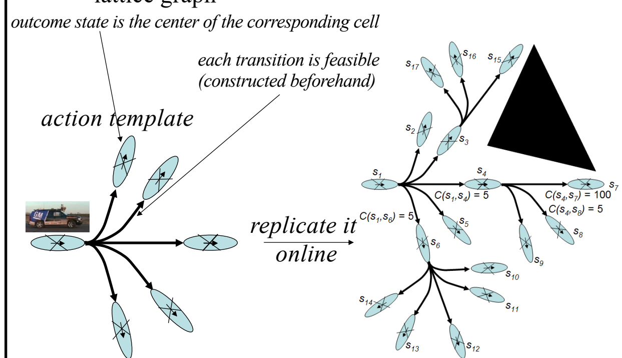

18

Planning via Cell Decomposition

- Graph construction:
  - lattice graph
  - pros: sparse graph, feasible paths
  - cons: possible incompleteness

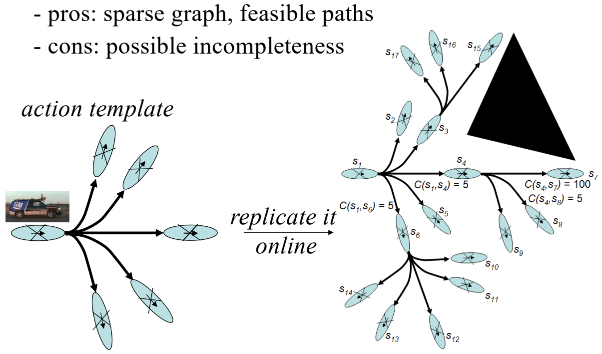

19

Outline

- Deterministic planning
  - constructing a graph
  - search with A*
  - search with D*

- Planning under uncertainty
  - Markov Decision Processes (MDP)
  - Partially Observable Decision Processes (POMDP)

20

A* Search

- Computes optimal g-values for relevant states

at any point of time:

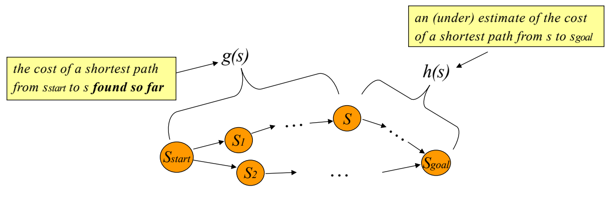

21

A* Search

- Computes optimal g-values for relevant states

at any point of time:

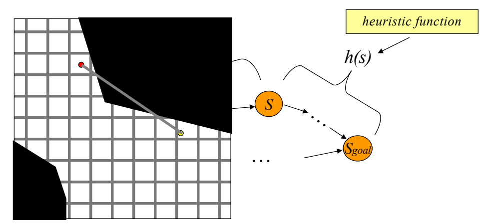

one popular heuristic function – Euclidean distance

22

A* Search

- Computes optimal g-values for relevant states

ComputePath function

while(s_{goal} is not expanded)
  remove *s* with the smallest *[f(s) = g(s)+h(s)]* from *OPEN*;
  insert *s* into *CLOSED*;
  for every successor *s'* of *s* such that *s'* not in *CLOSED*
    if *g(s') > g(s) + c(s,s')*
      g(s') = g(s) + c(s,s');
      insert *s'* into *OPEN*;

CLOSED = {}
OPEN = {s_{start}}
next state to expand: s_{start}

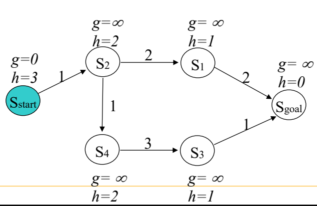

23

A* Search

- Computes optimal g-values for relevant states

ComputePath function

while(s_{goal} is not expanded)
  remove *s* with the smallest *[f(s) = g(s)+h(s)]* from *OPEN*;
  insert *s* into *CLOSED*;
  for every successor *s'* of *s* such that *s'* not in *CLOSED*
    if *g(s') > g(s) + c(s,s')*
      g(s') = g(s) + c(s,s');
      insert *s'* into *OPEN*;

CLOSED = {}
OPEN = {s_{start}}
next state to expand: s_{start}

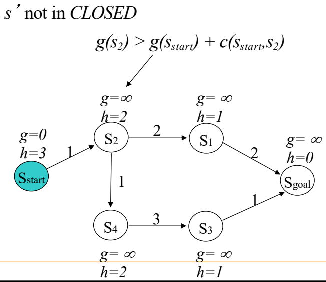

24

A* Search

- Computes optimal g-values for relevant states

ComputePath function
while(s_{goal} is not expanded)
  remove s with the smallest $[f(s) = g(s)+h(s)]$ from *OPEN*;
  insert s into *CLOSED*;
  for every successor s' of s such that s' not in *CLOSED*
    if $g(s') > g(s) + c(s,s')$
      $g(s') = g(s) + c(s,s')$;
      insert s' into *OPEN*;

$CLOSED = \{s_{start}\}$
$OPEN = \{s_2\}$
next state to expand: s_2

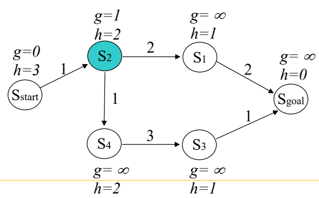

25

A* Search

- Computes optimal g-values for relevant states

ComputePath function
while(s_{goal} is not expanded)
  remove s with the smallest $[f(s) = g(s)+h(s)]$ from *OPEN*;
  insert s into *CLOSED*;
  for every successor s' of s such that s' not in *CLOSED*
    if $g(s') > g(s) + c(s,s')$
      $g(s') = g(s) + c(s,s')$;
      insert s' into *OPEN*;

$CLOSED = \{s_{start}, s_2\}$
$OPEN = \{s_1, s_4\}$
next state to expand: s_1

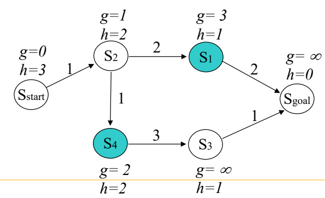

26

A* Search

- Computes optimal g-values for relevant states

ComputePath function
while(s_{goal} is not expanded)
  remove s with the smallest $[f(s) = g(s)+h(s)]$ from *OPEN*;
  insert s into *CLOSED*;
  for every successor s' of s such that s' not in *CLOSED*
    if $g(s') > g(s) + c(s,s')$
      $g(s') = g(s) + c(s,s')$;
      insert s' into *OPEN*;

$CLOSED = \{s_{start}, s_2, s_1\}$
$OPEN = \{s_4, s_{goal}\}$
next state to expand: s_4

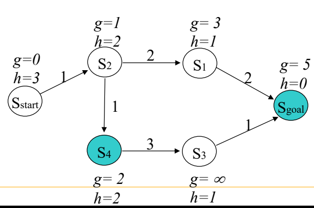

27

A* Search

- Computes optimal g-values for relevant states

ComputePath function
while(s_{goal} is not expanded)
  remove s with the smallest $[f(s) = g(s)+h(s)]$ from *OPEN*;
  insert s into *CLOSED*;
  for every successor s' of s such that s' not in *CLOSED*
    if $g(s') > g(s) + c(s,s')$
      $g(s') = g(s) + c(s,s')$;
      insert s' into *OPEN*;

$CLOSED = \{s_{start}, s_2, s_1, s_4\}$
$OPEN = \{s_3, s_{goal}\}$
next state to expand: s_{goal}

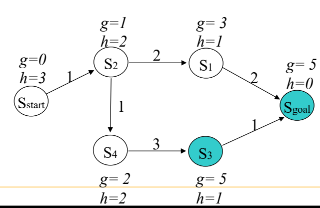

28

A* Search

- Computes optimal g-values for relevant states

ComputePath function

while(s_{goal} is not expanded)
  remove s with the smallest $[f(s) = g(s)+h(s)]$ from *OPEN*;
  insert s into *CLOSED*;
  for every successor s' of s such that s' not in *CLOSED*
    if $g(s') > g(s) + c(s,s')$
      $g(s') = g(s) + c(s,s')$;
      insert s' into *OPEN*;

CLOSED = $\{s_{start}, s_2, s_1, s_4, s_{goal}\}$
OPEN = $\{s_3\}$
done

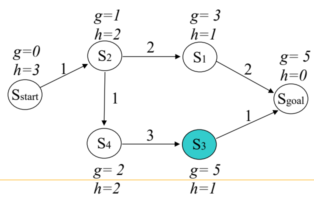

29

A* Search

- Computes optimal g-values for relevant states

ComputePath function

while(s_{goal} is not expanded)
  remove s with the smallest $[f(s) = g(s)+h(s)]$ from *OPEN*;
  insert s into *CLOSED*;
  for every successor s' of s such that s' not in *CLOSED*
    if $g(s') > g(s) + c(s,s')$
      $g(s') = g(s) + c(s,s')$;
      insert s' into *OPEN*;

for every expanded state g(s) is optimal
for every other state g(s) is an upper bound
we can now compute a least-cost path

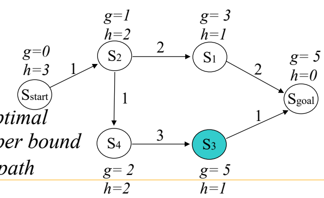

30

A* Search

- Is guaranteed to return an optimal path (in fact, for every expanded state) – optimal in terms of the solution
- Performs provably minimal number of state expansions required to guarantee optimality – optimal in terms of the computations

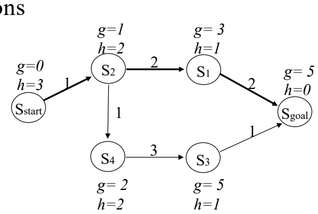

31

A* Search

- Is guaranteed to return an optimal path (in fact, for every expanded state) – optimal in terms of the solution
- Performs provably minimal number of state expansions required to guarantee optimality – optimal in terms of the computations

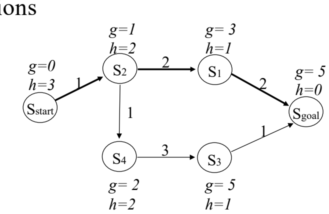

32

Effect of the Heuristic Function

- A* Search: expands states in the order of $f = g+h$ values

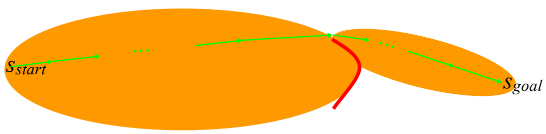

33

Effect of the Heuristic Function

- A* Search: expands states in the order of $f = g+h$ values

for large problems this results in A quickly running out of memory (memory: O(n))*

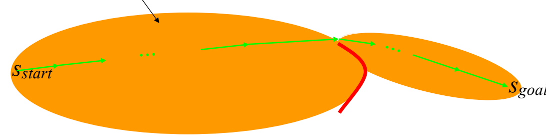

34

Effect of the Heuristic Function

- Weighted A* Search: expands states in the order of $f = g+\varepsilon h$ values, $\varepsilon > 1$ = bias towards states that are closer to goal

solution is always ε-suboptimal: cost(solution) ≤ ε·cost(optimal solution)

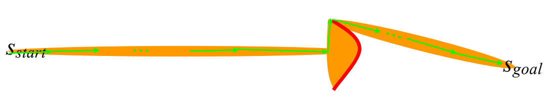

35

Effect of the Heuristic Function

- Weighted A* Search: expands states in the order of $f = g+\varepsilon h$ values, $\varepsilon > 1$ = bias towards states that are closer to goal

20DOF simulated robotic arm
state-space size: over 10^{26} states

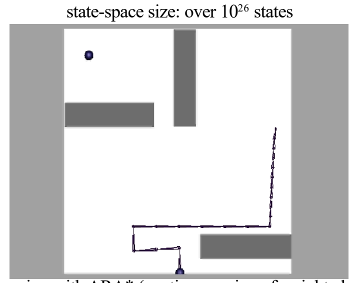

planning with ARA* (anytime version of weighted A*)

36

Effect of the Heuristic Function

- planning in 8D (<*x*,*y*> for each foothold)
- heuristic is Euclidean distance from the center of the body to the goal location
- cost of edges based on kinematic stability of the robot and quality of footholds

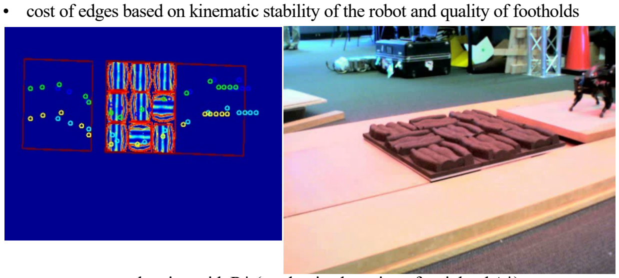

planning with R* (randomized version of weighted A*)

joint work with Subhrajit Bhattacharya, Jon Bohren, Sachin Chitta, Daniel D. Lee, Aleksandr Kushleyev, Paul Vernaza

37

Outline

- Deterministic planning
  - constructing a graph
  - search with A*
  - search with D*

38

Incremental version of A* (D*/D* Lite)

- Robot needs to re-plan whenever
  - new information arrives (partially-known environments or/and dynamic environments)
  - robot deviates off its path

ATRV navigating initially-unknown environment

planning map and path

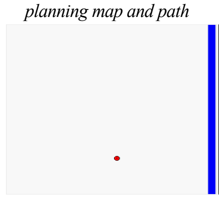

39

Motivation for Incremental Version of A*

- Reuse state values from previous searches

cost of least-cost paths to s_{goal} initially

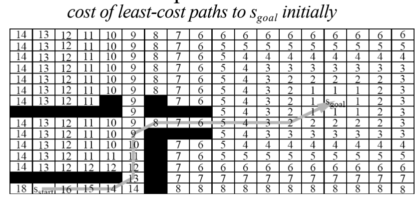

cost of least-cost paths to s_{goal} after the door turns out to be closed

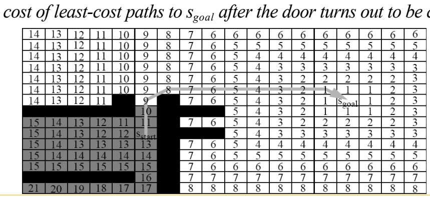

40

Motivation for Incremental Version of A*

- Reuse state values from previous searches

cost of least-cost paths to s_{goal} initially

cost of least-cost paths to s_{goal} after the door turns out to be closed

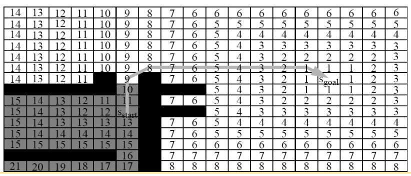

41

Motivation for Incremental Version of A*

- Reuse state values from previous searches

cost of least-cost paths to s_{goal} initially

cost of least-cost paths to s_{goal}

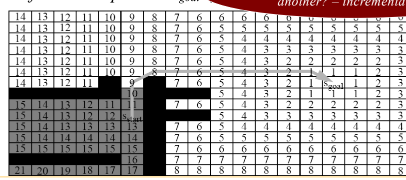

42

Motivation for Incremental Version of A*

- Reuse state values from previous searches

cost of least-cost paths to s_{goal} initially

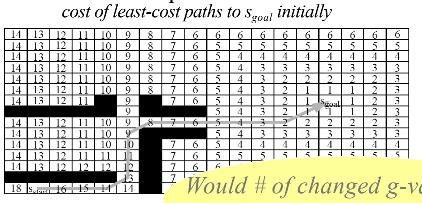

cost of least-cost paths to

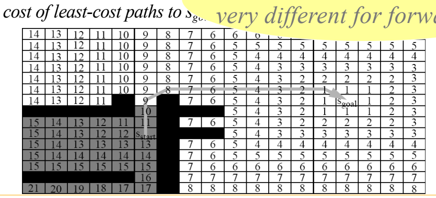

43

Motivation for Incremental Version of A*

- Reuse state values from previous searches

cost of least-cost paths to s_{goal} initially

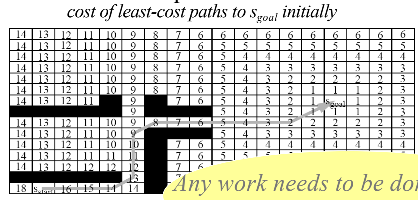

cost of least-cost paths to

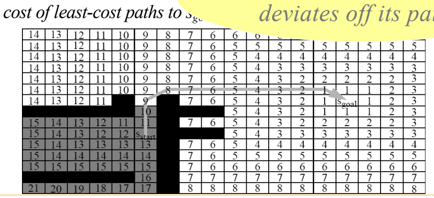

44

Incremental Version of A*

- Reuse state values from previous searches

*initial search by backwards A** *initial search by D* Lite*

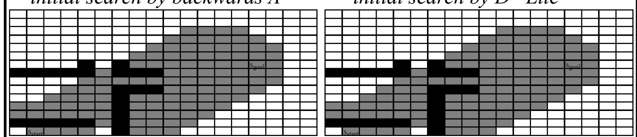

*second search by backwards A** *second search by D* Lite*

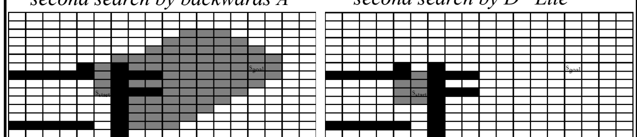

45

Anytime Aspects

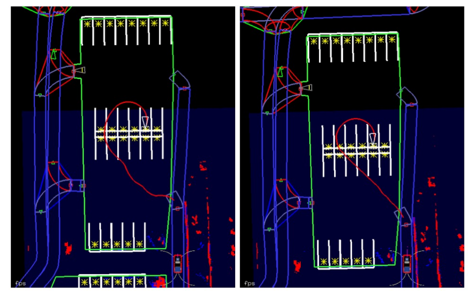

46

Anytime Aspects

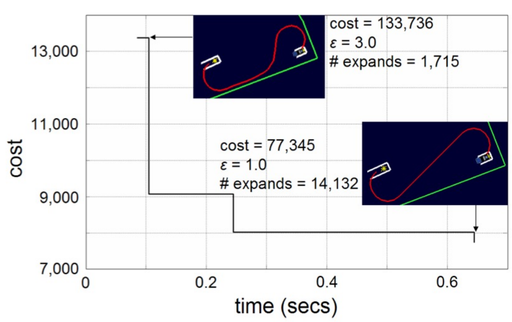

47

Searching the Graph

- Incremental behavior of Anytime D*:

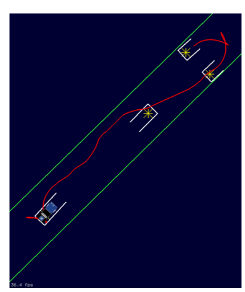

initial path

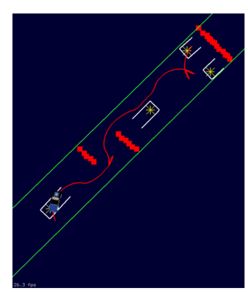

a path after re-planning

48

48

Searching the Graph

- Performance of Anytime D* depends strongly on heuristics *h(s)*: estimates of cost-to-goal

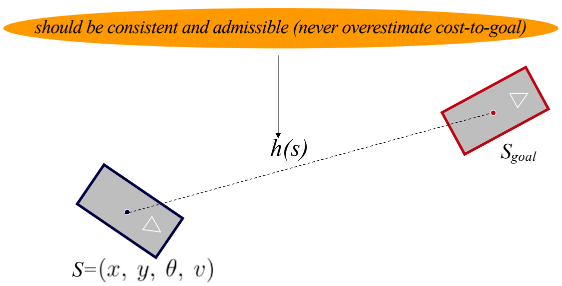

49

49

Searching the Graph

- In our planner: $h(s) = max(h_{mech}(s), h_{env}(s))$, where
  - $h_{mech}(s)$ – mechanism-constrained heuristic
  - $h_{env}(s)$ – environment-constrained heuristic

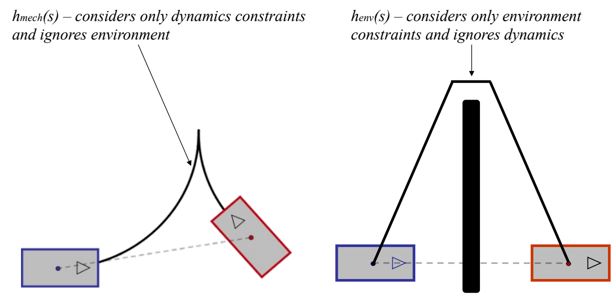

50

50

Searching the Graph

- In our planner: $h(s) = max(h_{mech}(s), h_{env}(s))$, where
  - $h_{mech}(s)$ – mechanism-constrained heuristic
  - $h_{env}(s)$ – environment-constrained heuristic

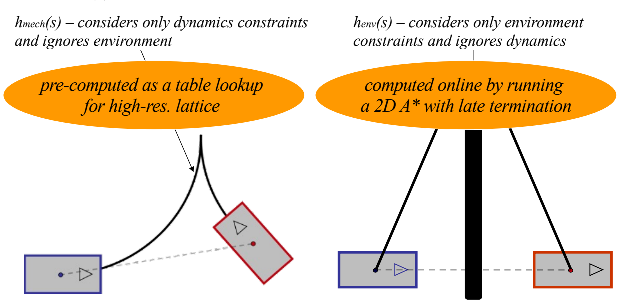

51

51

Heuristics

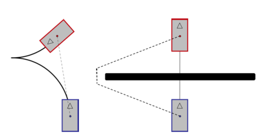

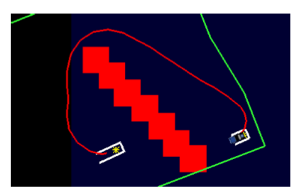

heuristic	states expanded	time (secs)
h	2,019	0.06
h_{2D}	26,108	1.30
h_{fsh}	124,794	3.49

52

Example, again

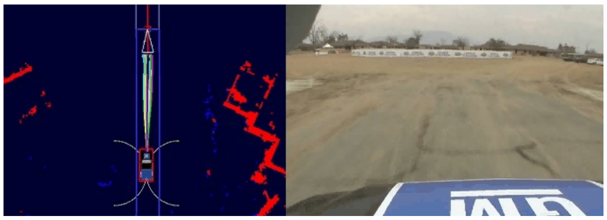

*Urban Challenge Race, CMU team, planning with Anytime D**

53

Summary

- Deterministic planning
  - constructing a graph
  - search with A*
  - search with D*

used a lot in real-time

think twice before trying to use it in real-time

- Planning under uncertainty
  - -Markov Decision Processes (MDP)
  - -Partially Observable Decision Processes (POMDP)

think three or four times before trying to use it in real-time

Many useful approximate solvers for MDP/POMDP exist!!

54

Manipulation Planning Examples

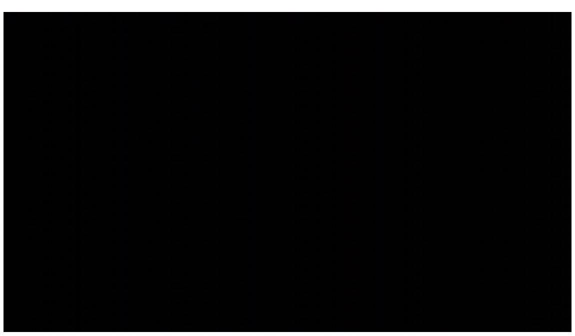

55