

CSE-571

Deterministic Path Planning in Robotics

*Courtesy of Maxim Likhachev
University of Pennsylvania*

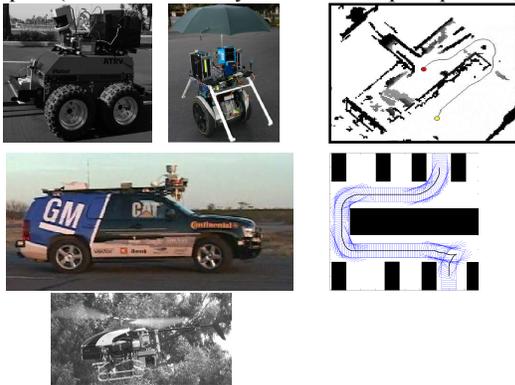
Motion/Path Planning

- 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, ...

CSE-571: Courtesy of Maxim Likhachev, CMU

Motion/Path Planning

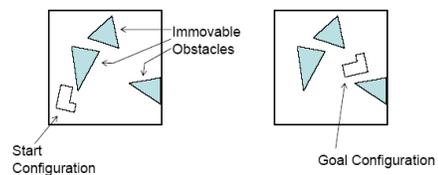
Examples (of what is usually referred to as path planning):



CSE-571: Courtesy of Maxim Likhachev, CMU

Motion/Path Planning

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



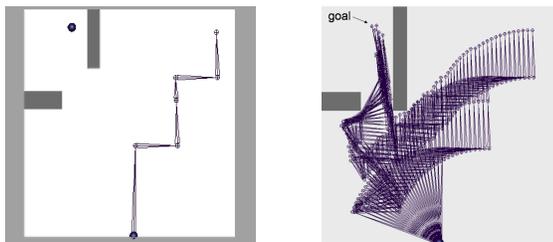
Piano Movers' problem

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

CSE-571: Courtesy of Maxim Likhachev, CMU

Motion/Path Planning

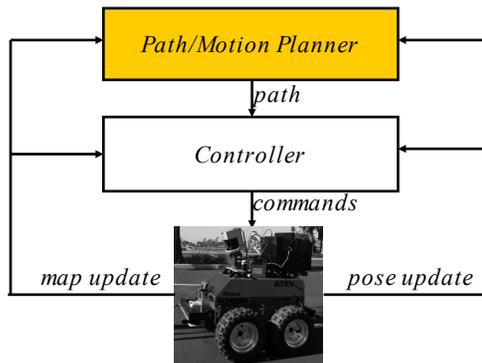
Examples (of what is usually referred to as motion planning)



Planned motion for a 6DOF robot arm

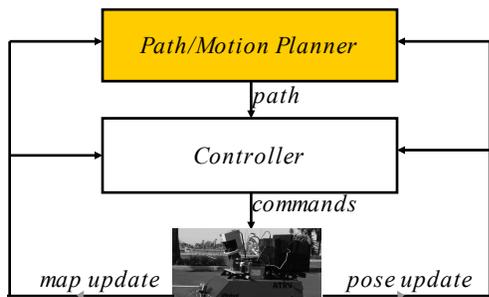
CSE-571: Courtesy of Maxim Likhachev, CMU

Motion/Path Planning



CSE-571: Courtesy of Maxim Likhachev, CMU

Motion/Path Planning



i.e., deterministic registration or Bayesian update

i.e., Bayesian update (EKF)

CSE-571: Courtesy of Maxim Likhachev, CMU

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

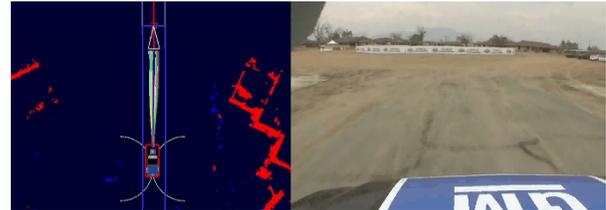
CSE-571: Courtesy of Maxim Likhachev, CMU

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

CSE-571 - Courtesy of Maxim Likhachev, CMU

Example



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

CSE-571 - Courtesy of Maxim Likhachev, CMU

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

CSE-571 - Courtesy of Maxim Likhachev, CMU

Outline

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

CSE-571 - Courtesy of Maxim Likhachev, CMU

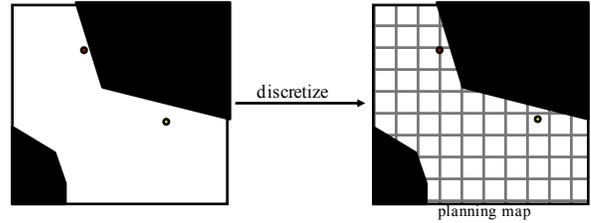
Outline

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

CSE-571: Courtesy of Maxim Likhachev, CMU

Planning via Cell Decomposition

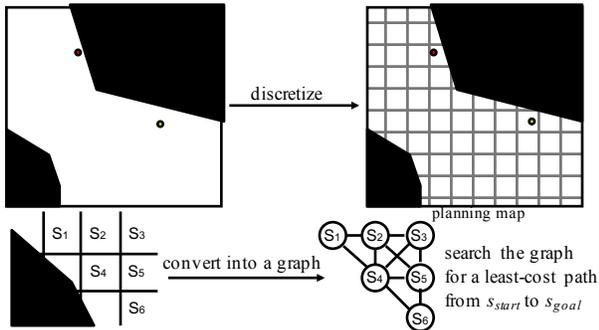
- Approximate Cell Decomposition:
 - overlay uniform grid over the C-space (discretize)



CSE-571: Courtesy of Maxim Likhachev, CMU

Planning via Cell Decomposition

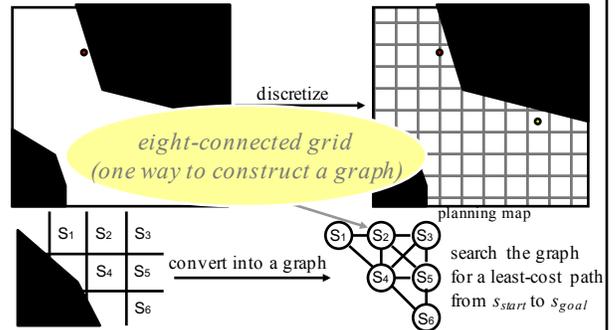
- Approximate Cell Decomposition:
 - construct a graph and search it for a least-cost path



CSE-571: Courtesy of Maxim Likhachev, CMU

Planning via Cell Decomposition

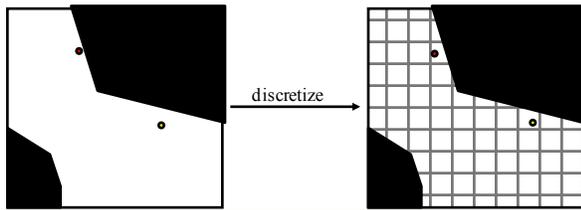
- Approximate Cell Decomposition:
 - construct a graph and search it for a least-cost path



CSE-571: Courtesy of Maxim Likhachev, CMU

Planning via Cell Decomposition

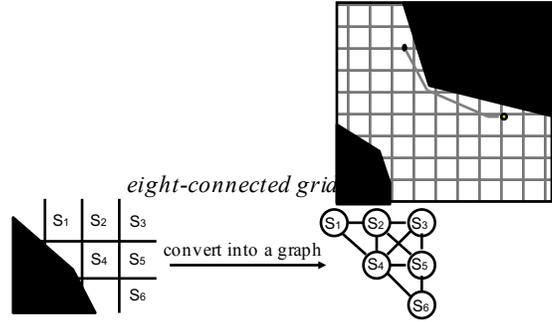
- Approximate Cell Decomposition:
 - construct a graph and search it for a least-cost path
 - VERY popular due to its simplicity and representation of arbitrary obstacles



CSE-571 - Courtesy of Maxim Likhachev, CMU

Planning via Cell Decomposition

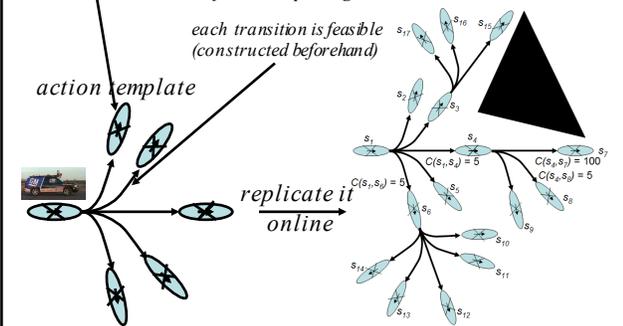
- Graph construction:
 - major problem with paths on the grid:
 - transitions difficult to execute on non-holonomic robots



CSE-571 - Courtesy of Maxim Likhachev, CMU

Planning via Cell Decomposition

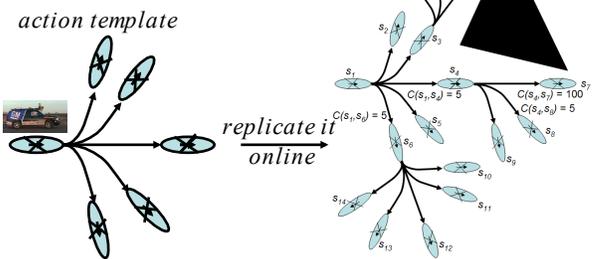
- Graph construction:
 - lattice graph
 - outcome state is the center of the corresponding cell



CSE-571 - Courtesy of Maxim Likhachev, CMU

Planning via Cell Decomposition

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



CSE-571 - Courtesy of Maxim Likhachev, CMU

Outline

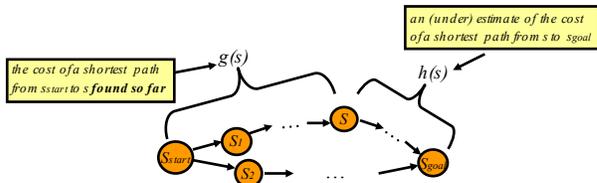
- Deterministic planning
 - constructing a graph
 - search with A*
 - search with D*
- Planning under uncertainty
 - Markov Decision Processes (MDP)
 - Partially Observable Decision Processes (POMDP)

CSE-571: Courtesy of Maxim Likhachev, CMU

A* Search

- Computes optimal g-values for relevant states

at any point of time:

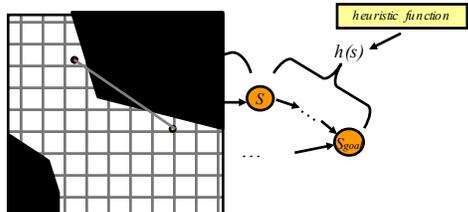


CSE-571: Courtesy of Maxim Likhachev, CMU

A* Search

- Computes optimal g-values for relevant states

at any point of time:

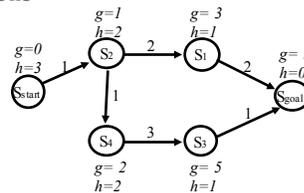


one popular heuristic function – Euclidean distance

CSE-571: Courtesy of Maxim Likhachev, CMU

A* Search

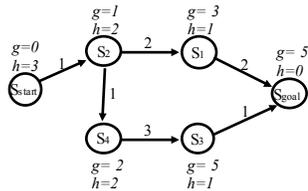
- Is guaranteed to return an optimal path (in fact, for every expanded state) – optimal in terms of the solution
- Performs provably minimal number of state expansions required to guarantee optimality – optimal in terms of the computations



CSE-571: Courtesy of Maxim Likhachev, CMU

A* Search

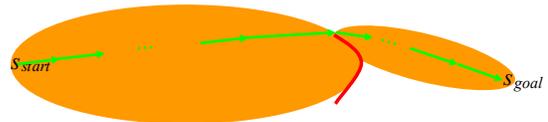
- Is guaranteed to return an optimal path (in fact, for every expanded state) – *helps with robot deviating off its path if we search with A* backwards (from goal to start)*
- Performs provably minimal number of state expansions required to guarantee optimality – optimal in terms of the computations



CSE-571: Courtesy of Maxim Likhachev, CMU

Effect of the Heuristic Function

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

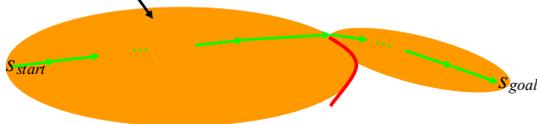


CSE-571: Courtesy of Maxim Likhachev, CMU

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: $O(n)$)*



CSE-571: Courtesy of Maxim Likhachev, CMU

Effect of the Heuristic Function

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

*solution is always ϵ -suboptimal:
 $cost(solution) \leq \epsilon cost(optimal solution)$*

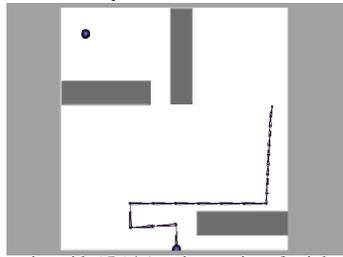


CSE-571: Courtesy of Maxim Likhachev, CMU

Effect of the Heuristic Function

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

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

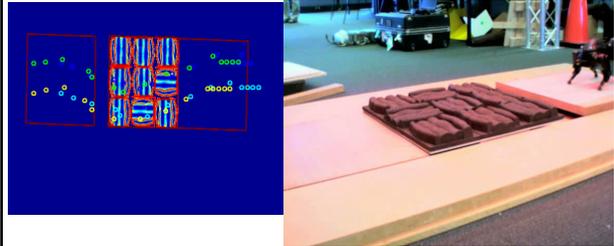


planning with ARA* (anytime version of weighted A*)

CSE-571: Courtesy of Maxim Likhachev, CMU

Effect of the Heuristic Function

- planning in 8D ($\langle x, y \rangle$ 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, Alexander Kushleyev, Paul Vernica

CSE-571: Courtesy of Maxim Likhachev, CMU

Outline

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

CSE-571: Courtesy of Maxim Likhachev, CMU

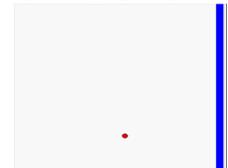
Incremental version of A* (D*/D* Lite)

- 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



CSE-571: Courtesy of Maxim Likhachev, CMU

Motivation for Incremental Version of A*

- Reuse state values from previous searches
cost of least-cost paths to s_{goal} initially

14	13	12	11	10	9	8	7	6	6	6	6	6	6	6	6
14	13	12	11	10	9	8	7	6	5	5	5	5	5	5	5
14	13	12	11	10	9	8	7	6	5	4	4	4	4	4	4
14	13	12	11	10	9	8	7	6	5	4	3	3	3	3	3
14	13	12	11	10	9	8	7	6	5	4	3	2	2	2	2
14	13	12	11	10	9	8	7	6	5	4	3	2	1	1	2
14	13	12	11	10	9	8	7	6	5	4	3	2	1	1	2
14	13	12	11	10	9	8	7	6	5	4	3	2	1	1	2
14	13	12	11	10	9	8	7	6	5	4	3	2	1	1	2
14	13	12	11	10	10	7	6	5	4	4	4	4	4	4	4
14	13	12	11	11	10	7	6	5	5	5	5	5	5	5	5
14	13	12	12	12	12	7	6	6							
18	18	16	16	15	14	14									

Would # of changed g-values be very different for forward A*?

14	13	12	11	10	9	8	7	6	6	6	6	6	6	6	6
14	13	12	11	10	9	8	7	6	5	5	5	5	5	5	5
14	13	12	11	10	9	8	7	6	5	4	4	4	4	4	4
14	13	12	11	10	9	8	7	6	5	4	3	3	3	3	3
14	13	12	11	10	9	8	7	6	5	4	3	2	2	2	2
14	13	12	11	10	9	8	7	6	5	4	3	2	1	1	2
14	13	12	11	10	9	8	7	6	5	4	3	2	1	1	2
14	13	12	11	10	9	8	7	6	5	4	3	2	1	1	2
14	13	12	11	10	10	7	6	5	4	4	4	4	4	4	4
14	13	12	11	11	10	7	6	5	5	5	5	5	5	5	5
14	13	12	12	12	12	7	6	6							
15	14	13	12	11	11	7	6	5	4	4	4	4	4	4	4
15	14	13	12	12	12	7	6	5	4	4	4	4	4	4	4
15	14	14	14	14	14	7	6	5	5	5	5	5	5	5	5
15	15	15	15	15	15	7	6	6	6	6	6	6	6	6	6
16	16	16	16	16	16	7	7	7	7	7	7	7	7	7	7
17	17	17	17	17	17	8	8	8	8	8	8	8	8	8	8
21	20	19	18	17	17										

CSE-571: Courtesy of Maxim Likhachev, CMU

Motivation for Incremental Version of A*

- Reuse state values from previous searches
cost of least-cost paths to s_{goal} initially

14	13	12	11	10	9	8	7	6	6	6	6	6	6	6	6
14	13	12	11	10	9	8	7	6	5	5	5	5	5	5	5
14	13	12	11	10	9	8	7	6	5	4	4	4	4	4	4
14	13	12	11	10	9	8	7	6	5	4	3	3	3	3	3
14	13	12	11	10	9	8	7	6	5	4	3	2	2	2	2
14	13	12	11	10	9	8	7	6	5	4	3	2	1	1	2
14	13	12	11	10	9	8	7	6	5	4	3	2	1	1	2
14	13	12	11	10	9	8	7	6	5	4	3	2	1	1	2
14	13	12	11	10	10	7	6	5	4	4	4	4	4	4	4
14	13	12	11	11	10	7	6	5	5	5	5	5	5	5	5
14	13	12	12	12	12	7	6	6							
18	18	16	16	15	14	14									

Any work needs to be done if robot deviates off its path?

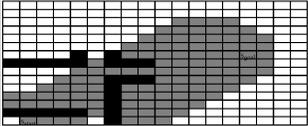
14	13	12	11	10	9	8	7	6	6	6	6	6	6	6	6
14	13	12	11	10	9	8	7	6	5	5	5	5	5	5	5
14	13	12	11	10	9	8	7	6	5	4	4	4	4	4	4
14	13	12	11	10	9	8	7	6	5	4	3	3	3	3	3
14	13	12	11	10	9	8	7	6	5	4	3	2	2	2	2
14	13	12	11	10	9	8	7	6	5	4	3	2	1	1	2
14	13	12	11	10	9	8	7	6	5	4	3	2	1	1	2
14	13	12	11	10	9	8	7	6	5	4	3	2	1	1	2
14	13	12	11	10	10	7	6	5	4	4	4	4	4	4	4
14	13	12	11	11	10	7	6	5	5	5	5	5	5	5	5
14	13	12	12	12	12	7	6	6							
15	14	13	12	11	11	7	6	5	4	4	4	4	4	4	4
15	14	13	12	12	12	7	6	5	4	4	4	4	4	4	4
15	14	14	14	14	14	7	6	5	5	5	5	5	5	5	5
15	15	15	15	15	15	7	6	6	6	6	6	6	6	6	6
16	16	16	16	16	16	7	7	7	7	7	7	7	7	7	7
17	17	17	17	17	17	8	8	8	8	8	8	8	8	8	8
21	20	19	18	17	17										

CSE-571: Courtesy of Maxim Likhachev, CMU

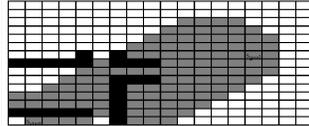
Incremental Version of A*

- Reuse state values from previous searches

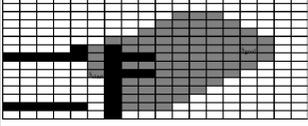
initial search by backwards A*



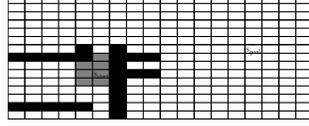
initial search by D* Lite



second search by backwards A*

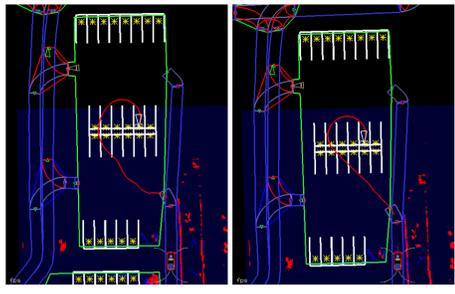


second search by D* Lite



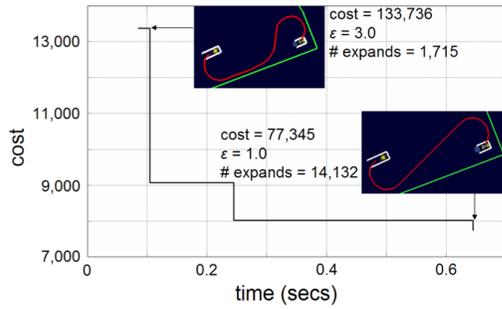
CSE-571: Courtesy of Maxim Likhachev, CMU

Anytime Aspects



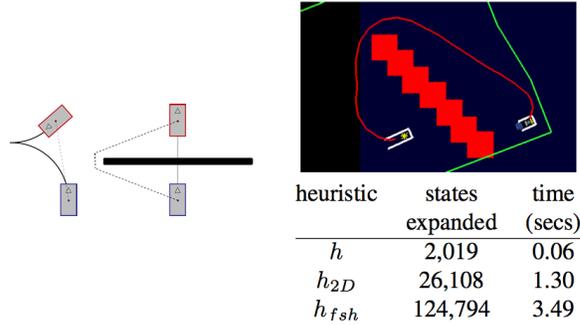
CSE-571: Courtesy of Maxim Likhachev, CMU

Anytime Aspects



CSE-571: Courtesy of Maxim Likhachev, CMU

Heuristics



CSE-571: Courtesy of Maxim Likhachev, CMU

Summary

- Deterministic planning
 - constructing a graph *used a lot in real-time*
 - search with A*
 - search with D* *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!!*

CSE-571: Courtesy of Maxim Likhachev, CMU