

Sampling-based Planning 2

Jane Li

Assistant Professor

Mechanical Engineering & Robotics Engineering


<http://users.wpi.edu/~zli11>

Problem with KD-tree


- Curse of dimension
 - N-dimensional configuration space, requires N level to halve the cell diameters
- Other (popular) trees for space partition in nearest neighbor search?
 - Random projection (RP) tree
 - Principal direction (PD) tree

Dimension Reduction


- A lot of data which superficially lie in a very high-dimensional space, actually have **low intrinsic dimension**



Dimension Reduction




Random projection



Principle Component Analysis


Random Projection Tree

- If the N -dimensional data has intrinsic dimension n , then an RP tree halves the diameter in just d levels – no dependence on N



KD-tree

Pick coordinate direction and split at median




RP-tree

Pick random direction and split at median plus noise


Principal Direction Tree

- Principle value decomposition



KD-tree

Partition with hyperplane perpendicular to an axis




PD-tree

Partition with hyperplane perpendicular to
the principle axis direction

Recap

- Last time, we discussed PRMs



- Two issues with the PRM:
 1. Uniform random sampling misses **narrow passages**
 2. Exploring **whole space**, but all we want is a path

Outline

- Sampling strategies
- RRTs

Sampling Strategies

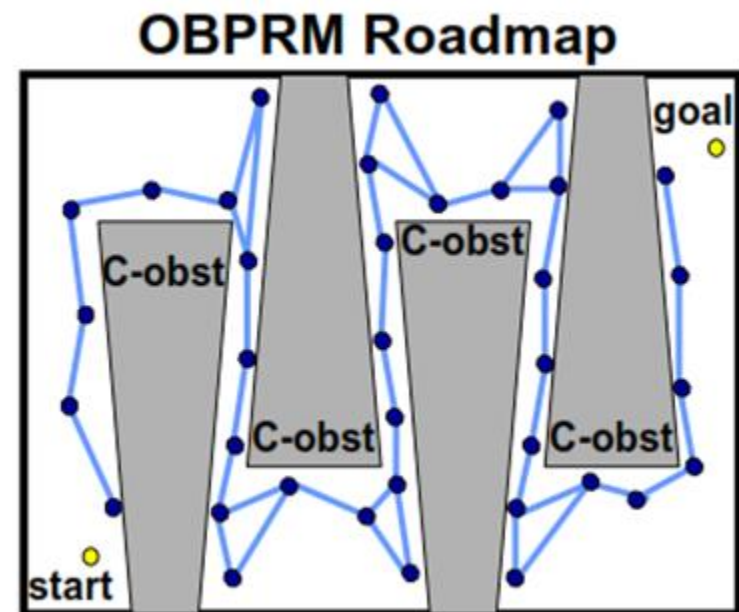

- Uniform random sampling – Most common
 - The bigger the area, the more likely it will be sampled
 - Problem – Narrow passages



- Different sampling strategies?

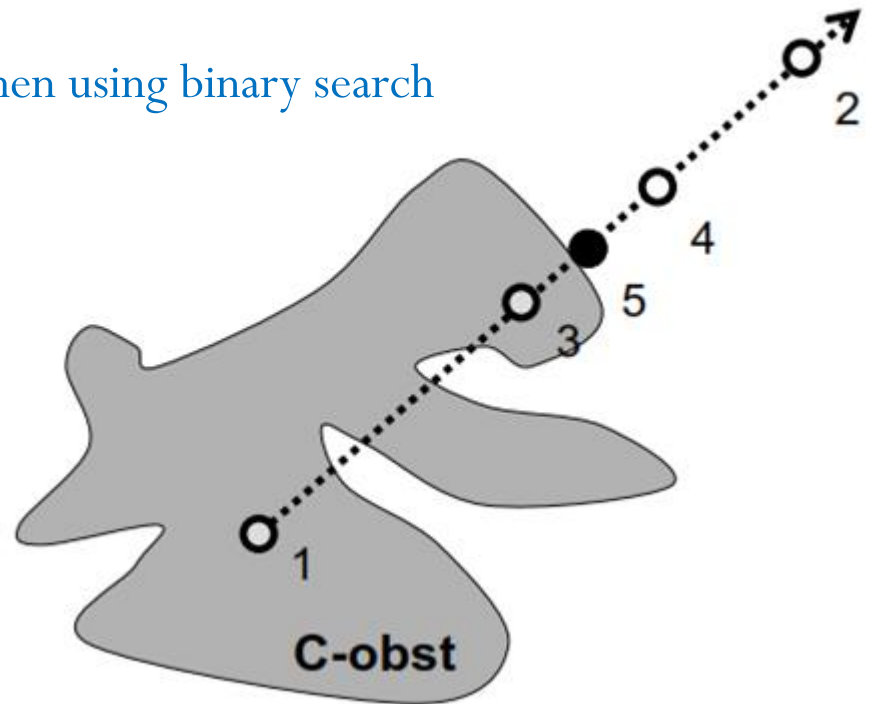
Obstacle-based PRM

- To navigate narrow passage, we must sample in them
 - Uniform sampling – Most PRM points fall in where planning is easy
 - Sample near C-obstacles?
 - But we cannot explicitly construct C-obstacles ...




Obstacle-based PRM

- How to find points on C-obstacles?
 - Find a point in the C-obstacles – a collision configuration
 - Select a random direction in C-space
 - Find a free point in that direction
 - Find the boundary point between then using binary search




PRM VS OBPRM



PRM

- 328 nodes
- 4 major CCs



OBPRM


- 161 nodes
- 2 major CCs

Gaussian Sampling

- Gaussian sampler

What is the effect?


- Find a q_1
- Pick a q_2 from a Gaussian distribution centered at q_1
- If **both** are in collision or collision-free, discard them, if one free, keep it




Sampling distribution for varying Gaussian width
(width decreasing from left to right)

Boor, Valérie, Mark H. Overmars, and A. Frank van der Stappen. "The gaussian sampling strategy for probabilistic roadmap planners." *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*. Vol. 2. IEEE, 1999.

Gaussian Sampling



Milestones (13,000) created by uniform sampling before the narrow passage was adequately sampled




Milestones (150) created by Gaussian sampling

The gain is not in sampling fewer milestones, but in connecting fewer pairs of milestones

Bridge


- Bridge sampler
 - Sample a q_1 that is in collision
 - Sample a q_2 in neighborhood of q_1 using some probability distribution (e.g. gaussian)
 - If q_2 in collision, get the midpoint of (q_1, q_2)
 - Check if midpoint is in collision, if not, add it as a node

What is the effect?




Hsu, David, et al. "The bridge test for sampling narrow passages with probabilistic roadmap planners." *Robotics and Automation, 2003. Proceedings. ICRA'03. IEEE International Conference on*. Vol. 3. IEEE, 2003.


Bridge



Bridge

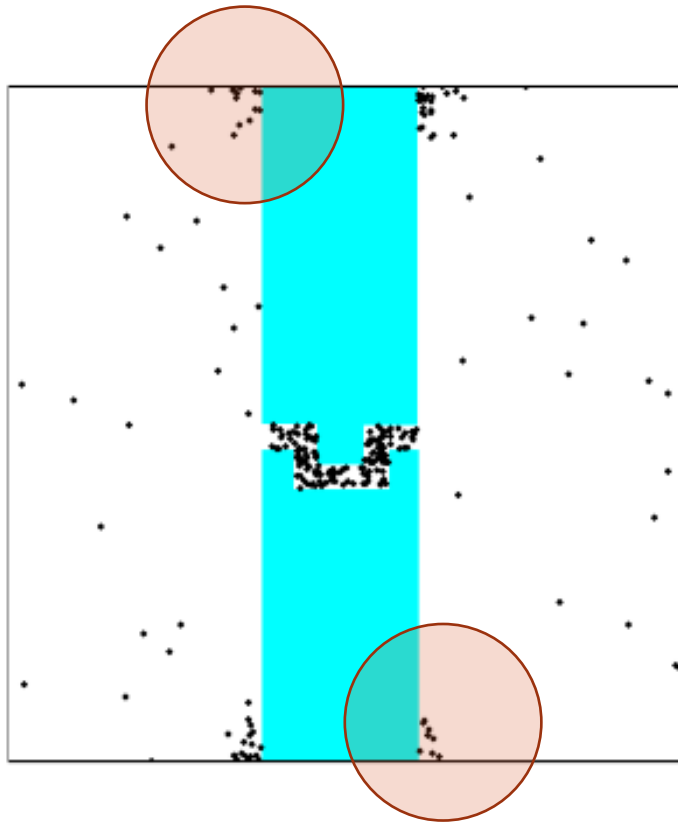


Gaussian




Bridge test

Bridge



Bridge Sampling performs well in narrow passages



What's going on at the corners?

Deterministic Sampling

What to do?

- The problem:
 - Random sampling (biased or not) can be unpredictable and irregular
 - Each time you run your algorithm you get a **different sequence of samples**, so **performance varies**
 - In the limit, space will be sampled well, but in **finite time result may be irregular**

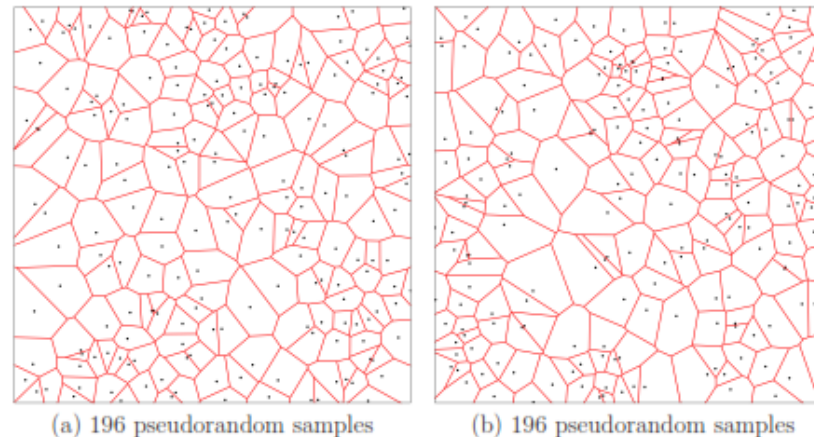


Figure 5.3: Irregularity in a collection of (pseudo)random samples can be nicely observed with Voronoi diagrams.

Deterministic Sampling


- What do we care about?
 - Dispersion

$$\delta(P) = \sup_{x \in X} \{ \min_{p \in P} \{ \rho(x, p) \} \}.$$



P is a finite set of points, (X, ρ) is a metric space (ρ is a distance metric)

In English: the radius of the largest empty ball

- What does it mean?
 - Intuitively, the dispersion quantifies how well a space is covered by a set of points S in terms of the largest open Euclidean ball that touches none of the points.



(a) L_2 dispersion



Deterministic Sampling

- Deterministic Sampling
 - Similar to discretization we saw in Discrete Motion Planning, but order of samples matters
- Sequences?
 - Van der Corput sequence – 1D
 - Halton sequence
 - **n-dimensional generalization** of van der Corput sequence
 - Hammersley sequence
 - Adaptation of Halton sequence that yields a **better distribution**. **BUT need to know number of samples** in advance.

Multi- vs. Single-Query Roadmaps

- Multi-query roadmaps
 - Pre-compute roadmap
 - Re-use roadmap for answering queries
 - The roadmap must cover the free space well
 - Why try to capture the connectivity of the whole space when all you need is one path?
- Single-query roadmaps
 - Dynamic environment
 - Compute a roadmap from scratch for each new query

Single-query Methods

- Key idea
 - Build a tree instead of a general graph.
 - The tree grows in C_{free}
 - Like PRM, captures some connectivity
 - Unlike PRM, only explores what is connected to q_{start}
- Algorithms:
 - Single-Query BiDirectional Lazy PRM (SBL-PRM)
 - Expansive Space Trees (EST)
 - **Rapidly-exploring Random Tree (RRT)**

Naïve Tree Algorithm

- Steps
 - Pick a node at random
 - Sample a new node near it
 - Grow tree from the random node to the new node

$q_{\text{node}} = q_{\text{start}}$

For $i = 1$ to NumberSamples


$q_{\text{rand}} = \text{Sample near } q_{\text{node}}$

 Add edge $e = (q_{\text{rand}}, q)$ if

collision-free

$q_{\text{node}} = \text{Pick random node of}$

tree



Grow a RRT

- Steps
 - Grow a tree rooted at the starting configuration,
 - Randomly sample from the search space
 - For each sample, try to connect it to the nearest node of the tree
 - Success – add a new node
 - Fail – discard the sample

Basic RRT Algorithm

```
BUILD_RRT ( $q_{init}$ ) {  
   $T.init(q_{init});$   
  for  $k = 1$  to  $K$  do  
     $q_{rand} = RANDOM\_CONFIG();$   
     $EXTEND(T, q_{rand})$   
}
```


STEP_LENGTH: How far to sample

1. Sample just at end point
2. Sample all along
3. Small Step

Extend returns


1. Trapped, cant make it
2. Extended, steps toward node
3. Reached, connects to node

$EXTEND(T, q_{rand})$

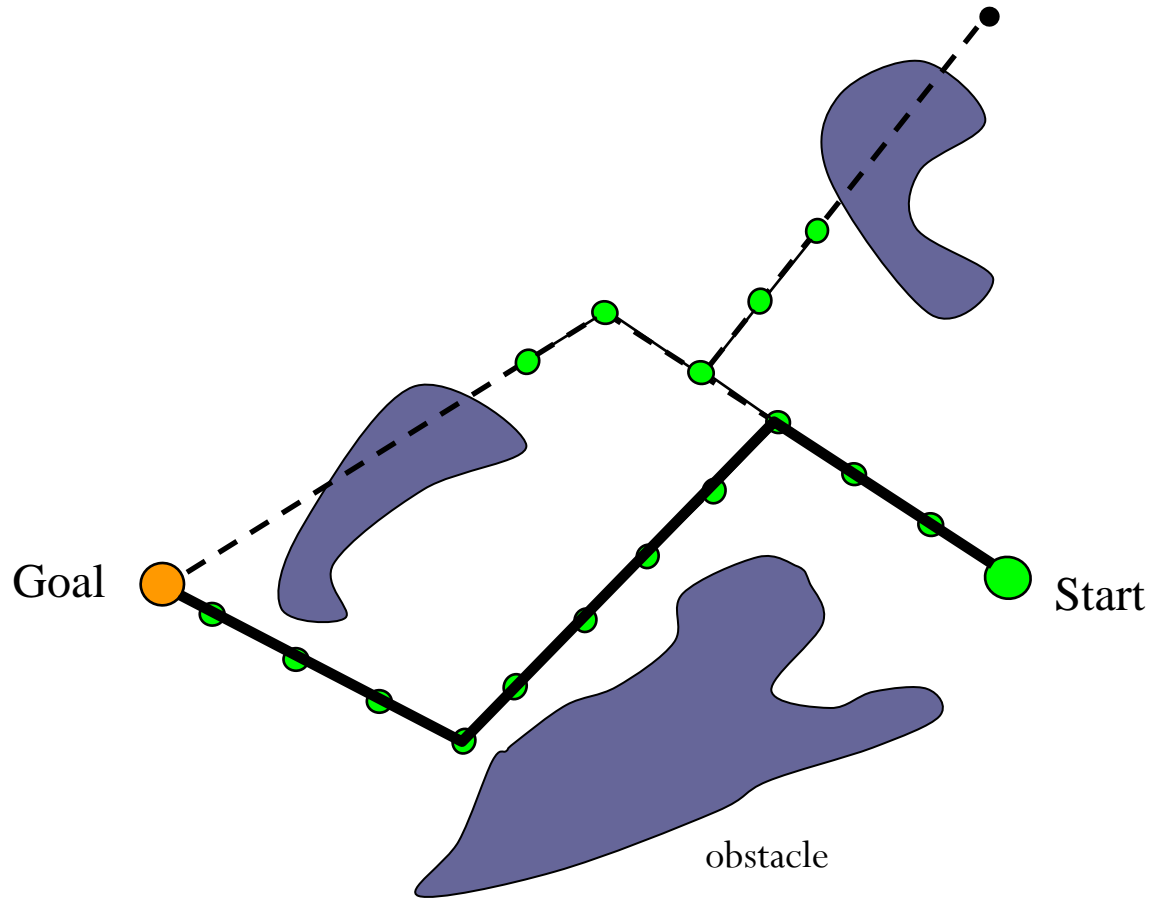


[Kuffner & LaValle , ICRA'00]

RRT Growing in Empty Space



RRT with Obstacles and Goal bias




Sample Bias


- Bias
 - Toward larger spaces
 - Toward goal
 - When generating a random sample, with some probability pick the goal instead of a random node when expanding
 - This introduces another parameter
 - 5-10% is the right choice [James Kuffner]
 - What happens if you set probability of sampling goal to 100%?

RRT Extension Types

- RRT-Extend
 - Take one step toward a random sample




- RRT-Connect
 - Step toward random sample until it is either
 - Reached
 - You hit an obstacle



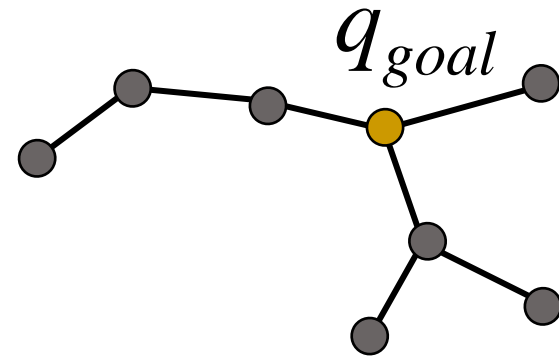
BiDirectional RRTs

- BiDirectional RRT
 - Grow trees from both start and goal
 - Try to get trees to connect to each other
 - Trees can both use Extend or both use Connect or one use Extend and one Connect
- BiDirectional RRT with Connect for both trees
 - Preferred, since this variant has only one parameter, the step size

Example of BiDirectional RRT




Connect




Extend

Example of BiDirectional RRT

1) One tree grown using random target




Connect




Extend

Example of BiDirectional RRT

2) New node becomes target for other tree



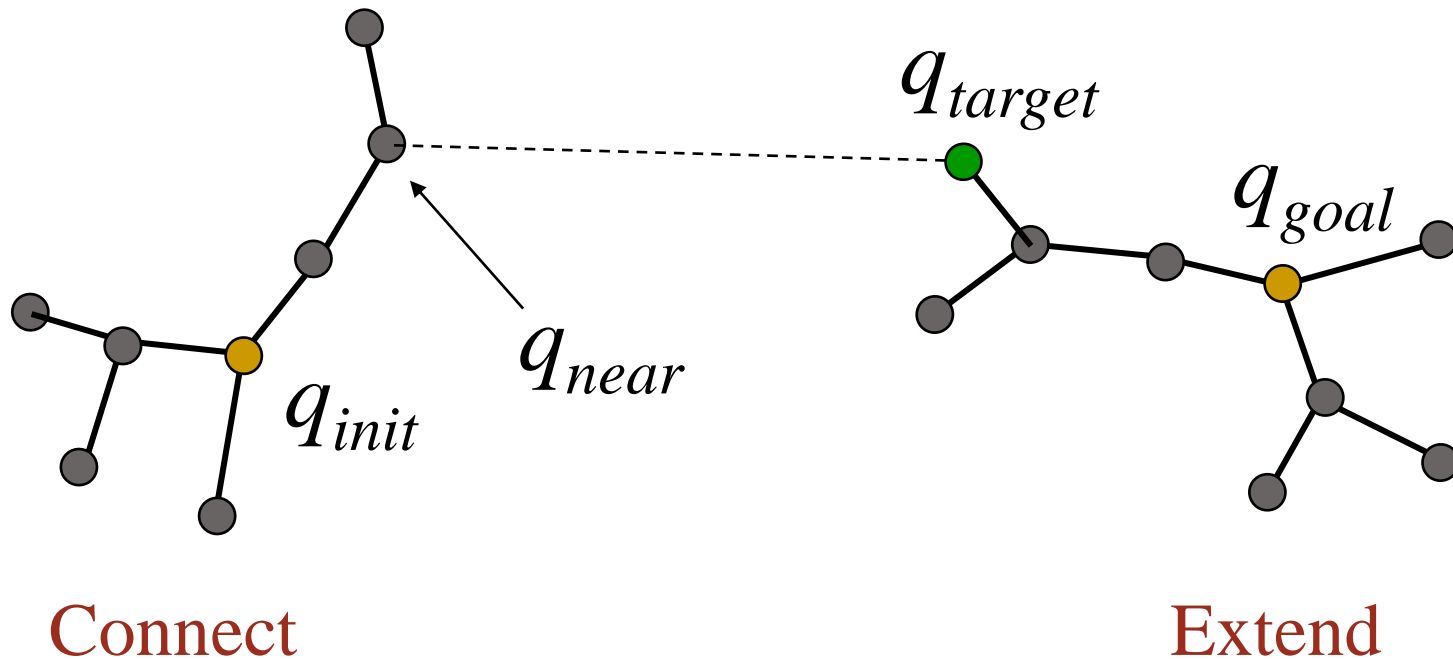
Connect



Extend


Example of BiDirectional RRT

3) Calculate node "nearest" to target




Example of BiDirectional RRT

4) Try to add new collision-free branch




Example of BiDirectional RRT

5) If successful, keep extending branch




Example of BiDirectional RRT

5) If successful, keep extending branch



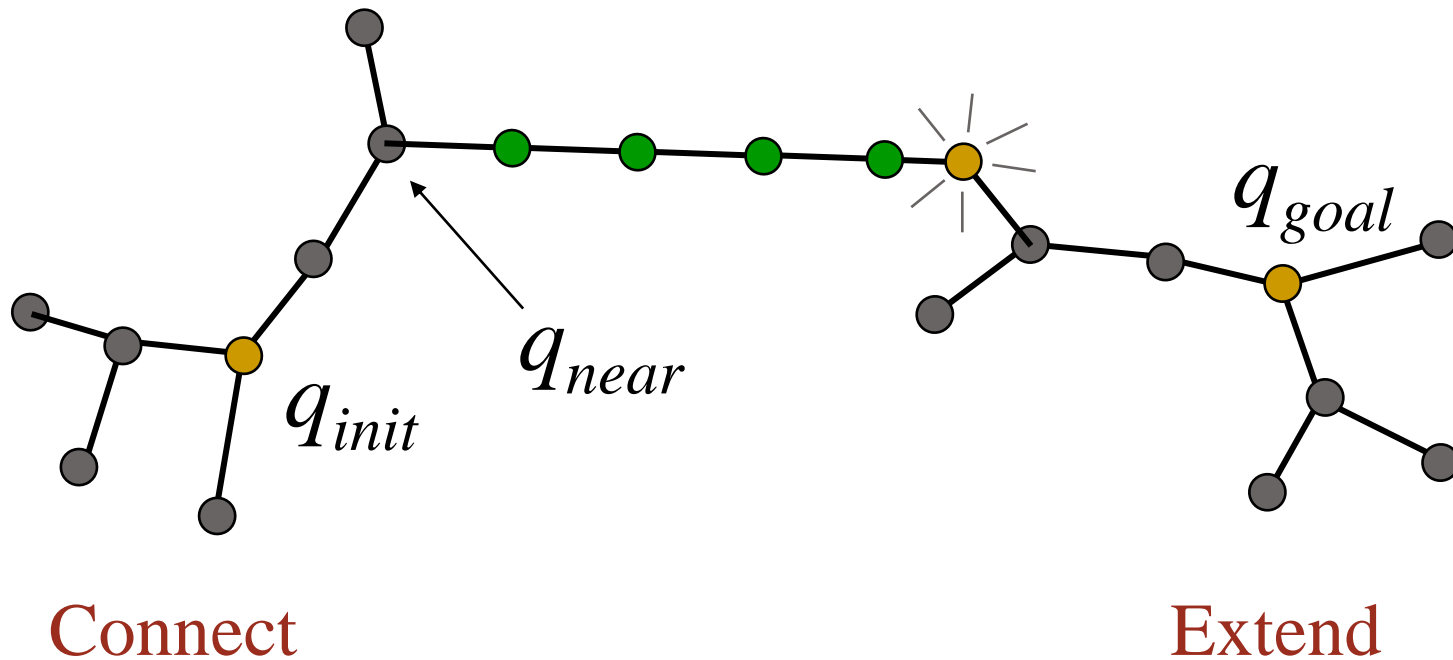
Example of BiDirectional RRT

5) If successful, keep extending branch




Example of BiDirectional RRT

6) Path found if branch reaches target



Example of BiDirectional RRT

7) Return path connecting start and goal



Tree Swapping and Balancing


- Tree Swapping and Balancing
- Some use Tree Balancing:

```
RDT BALANCED BIDIRECTIONAL( $q_I, q_G$ )
1   $T_a$ .init( $q_I$ );  $T_b$ .init( $q_G$ );
2  for  $i = 1$  to  $K$  do
3     $q_n \leftarrow$  NEAREST( $S_a, \alpha(i)$ );
4     $q_s \leftarrow$  STOPPING-CONFIGURATION( $q_n, \alpha(i)$ );
5    if  $q_s \neq q_n$  then
6       $T_a$ .add_vertex( $q_s$ );
7       $T_a$ .add_edge( $q_n, q_s$ );
8       $q'_n \leftarrow$  NEAREST( $S_b, q_s$ );
9       $q'_s \leftarrow$  STOPPING-CONFIGURATION( $q'_n, q_s$ );
10     if  $q'_s \neq q'_n$  then
11        $T_b$ .add_vertex( $q'_s$ );
12        $T_b$ .add_edge( $q'_n, q'_s$ );
13     if  $q'_s = q_s$  then return SOLUTION;
14     if  $|T_b| > |T_a|$  then SWAP( $T_a, T_b$ );
15 return FAILURE
```

- What is a situation where this would *help* performance?
- What is a situation where this would *hurt* performance?


Path Smoothing/Optimization

- RRTs produce notoriously bad paths
 - Not surprising since no consideration of path quality
- ALWAYS smooth/optimize the returned path
 - Many methods exists, e.g. shortcut smoothing (from previous lecture)



RRT Examples: The Alpha Puzzle

- VERY hard 6DOF motion planning problem (long, winding narrow passage)



- *“In 2001, it was solved by using a balanced bidirectional RRT, developed by James Kuffner and Steve LaValle. There are no special heuristics or parameters that were tuned specifically for this problem. On a current PC (circa 2003), it consistently takes a few minutes to solve”* –RRT website
- RRT became famous in large part because it was able to solve this puzzle

RRT Analysis

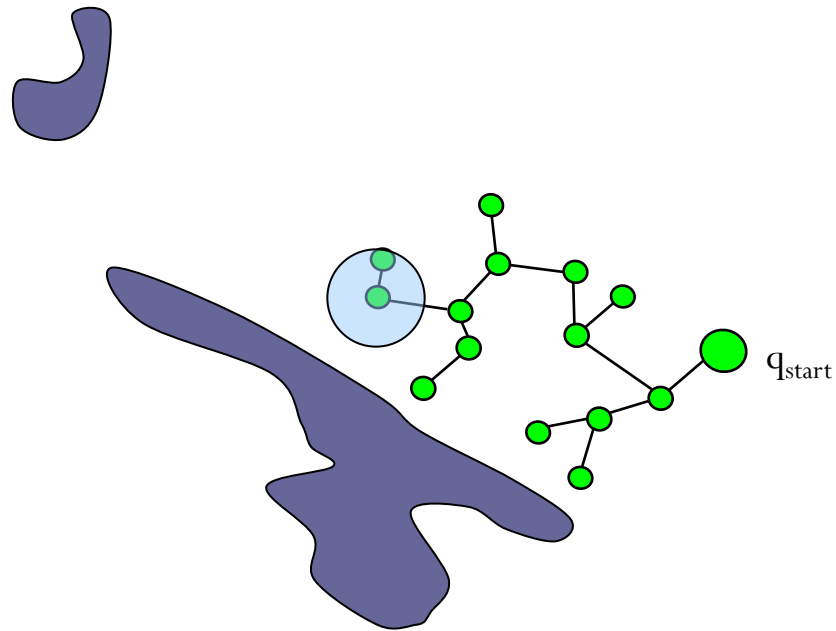
The limit of the distribution of vertices

- THEOREM: X_k converges to X with probability 1 as time goes to infinity
 - X_k : The RRT vertex distribution at iteration k
 - X : The distribution used for generating samples
- If using uniform distribution,
 - Tree nodes converge to the free space

Probabilistic Completeness


- Definition:
 - A path planner is *probabilistically complete* if, given a solvable problem, the probability that the planner solves the problem goes to 1 as time goes to infinity.
- Will RRT explore the whole space?

Proof of RRT Probabilistic Completeness




Kuffner and LaValle, ICRA, 2000

Proof of RRT Probabilistic Completeness



Kuffner and LaValle, ICRA, 2000

Proof of RRT Probabilistic Completeness



Kuffner and LaValle, ICRA, 2000

Probabilistic Completeness

- As the RRT reaches all of Q_{free} ,
 - The probability that q_{rand} immediately becomes a new vertex approaches 1.

- So, is RRT probabilistically complete?

Sampling-Based Planning

- The good:
 - Provides fast **feasible** solution
 - Popular methods have **few** parameters
 - Works on **practical** problems
 - Works in **high**-dimensions

Sampling-Based Planning

- The bad:
 - No quality guarantees on paths quality
 - In practice: smooth/optimize path afterwards
 - No termination when there is no solution
 - In practice: set an arbitrary timeout
 - Probabilistic completeness is a weak property
 - Completeness in high-dimensions is impractical

Readings

- Read “Non-holonomic motion planning guide” linked on class webpage