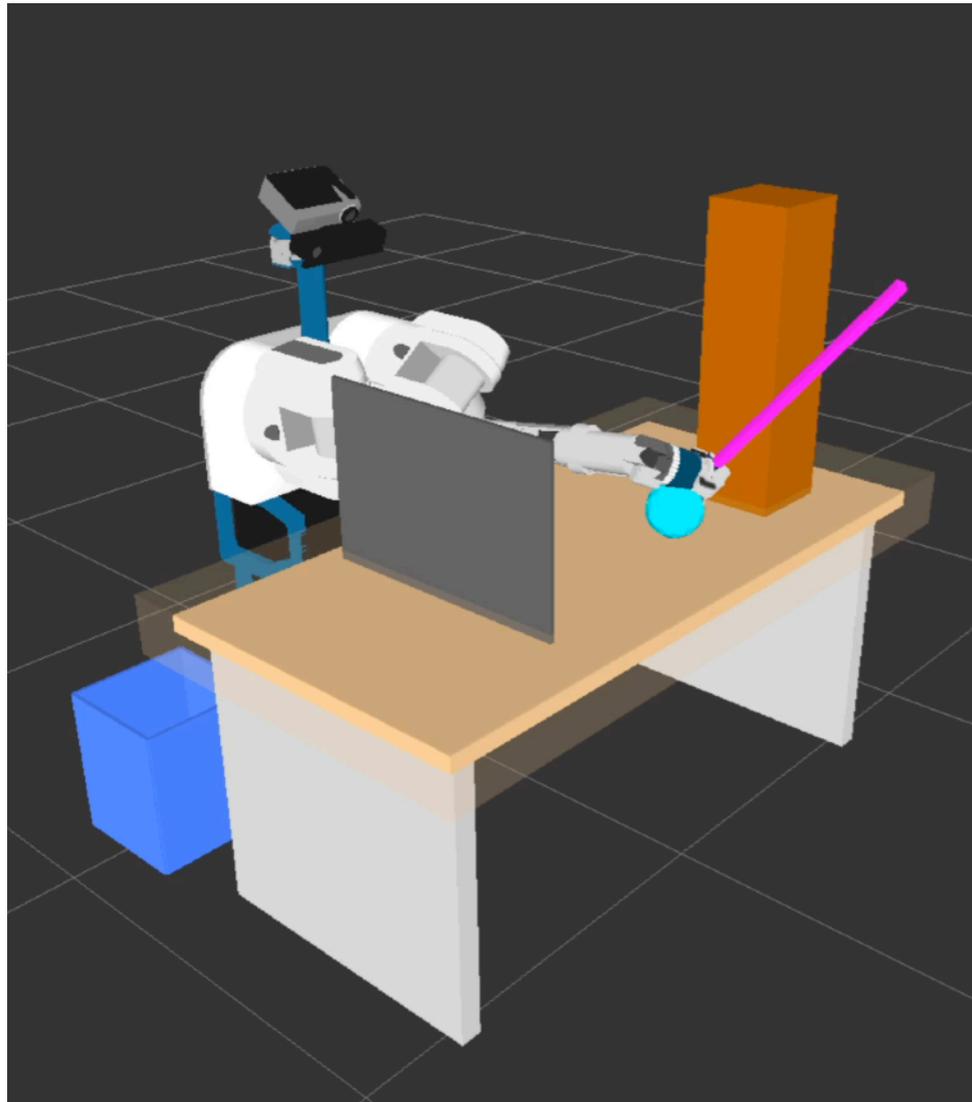


Planning on Roadmaps

Sanjiban Choudhury

TAs: Matthew Rockett, Gilwoo Lee, Matt Schmittle

Geometric Path Planning Problem



Also known as

Piano Mover's Problem (Reif 79)

Given:

1. A *workspace* \mathcal{W} , where either $\mathcal{W} = \mathbb{R}^2$ or $\mathcal{W} = \mathbb{R}^3$.
2. An *obstacle region* $\mathcal{O} \subset \mathcal{W}$.
3. A *robot* defined in \mathcal{W} . Either a rigid body \mathcal{A} or a collection of m links: $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_m$.
4. The *configuration space* \mathcal{C} (\mathcal{C}_{obs} and \mathcal{C}_{free} are then defined).
5. An *initial configuration* $\mathbf{q}_I \in \mathcal{C}_{free}$.
6. A *goal configuration* $\mathbf{q}_G \in \mathcal{C}_{free}$. The initial and goal configuration are often called a *query* $(\mathbf{q}_I, \mathbf{q}_G)$.

Compute a (continuous) path, $\tau : [0, 1] \rightarrow \mathcal{C}_{free}$, such that $\tau(0) = \mathbf{q}_I$ and $\tau(1) = \mathbf{q}_G$.

Also may want to minimize cost $c(\tau)$

But I just want to know
how to plan for my racecar!

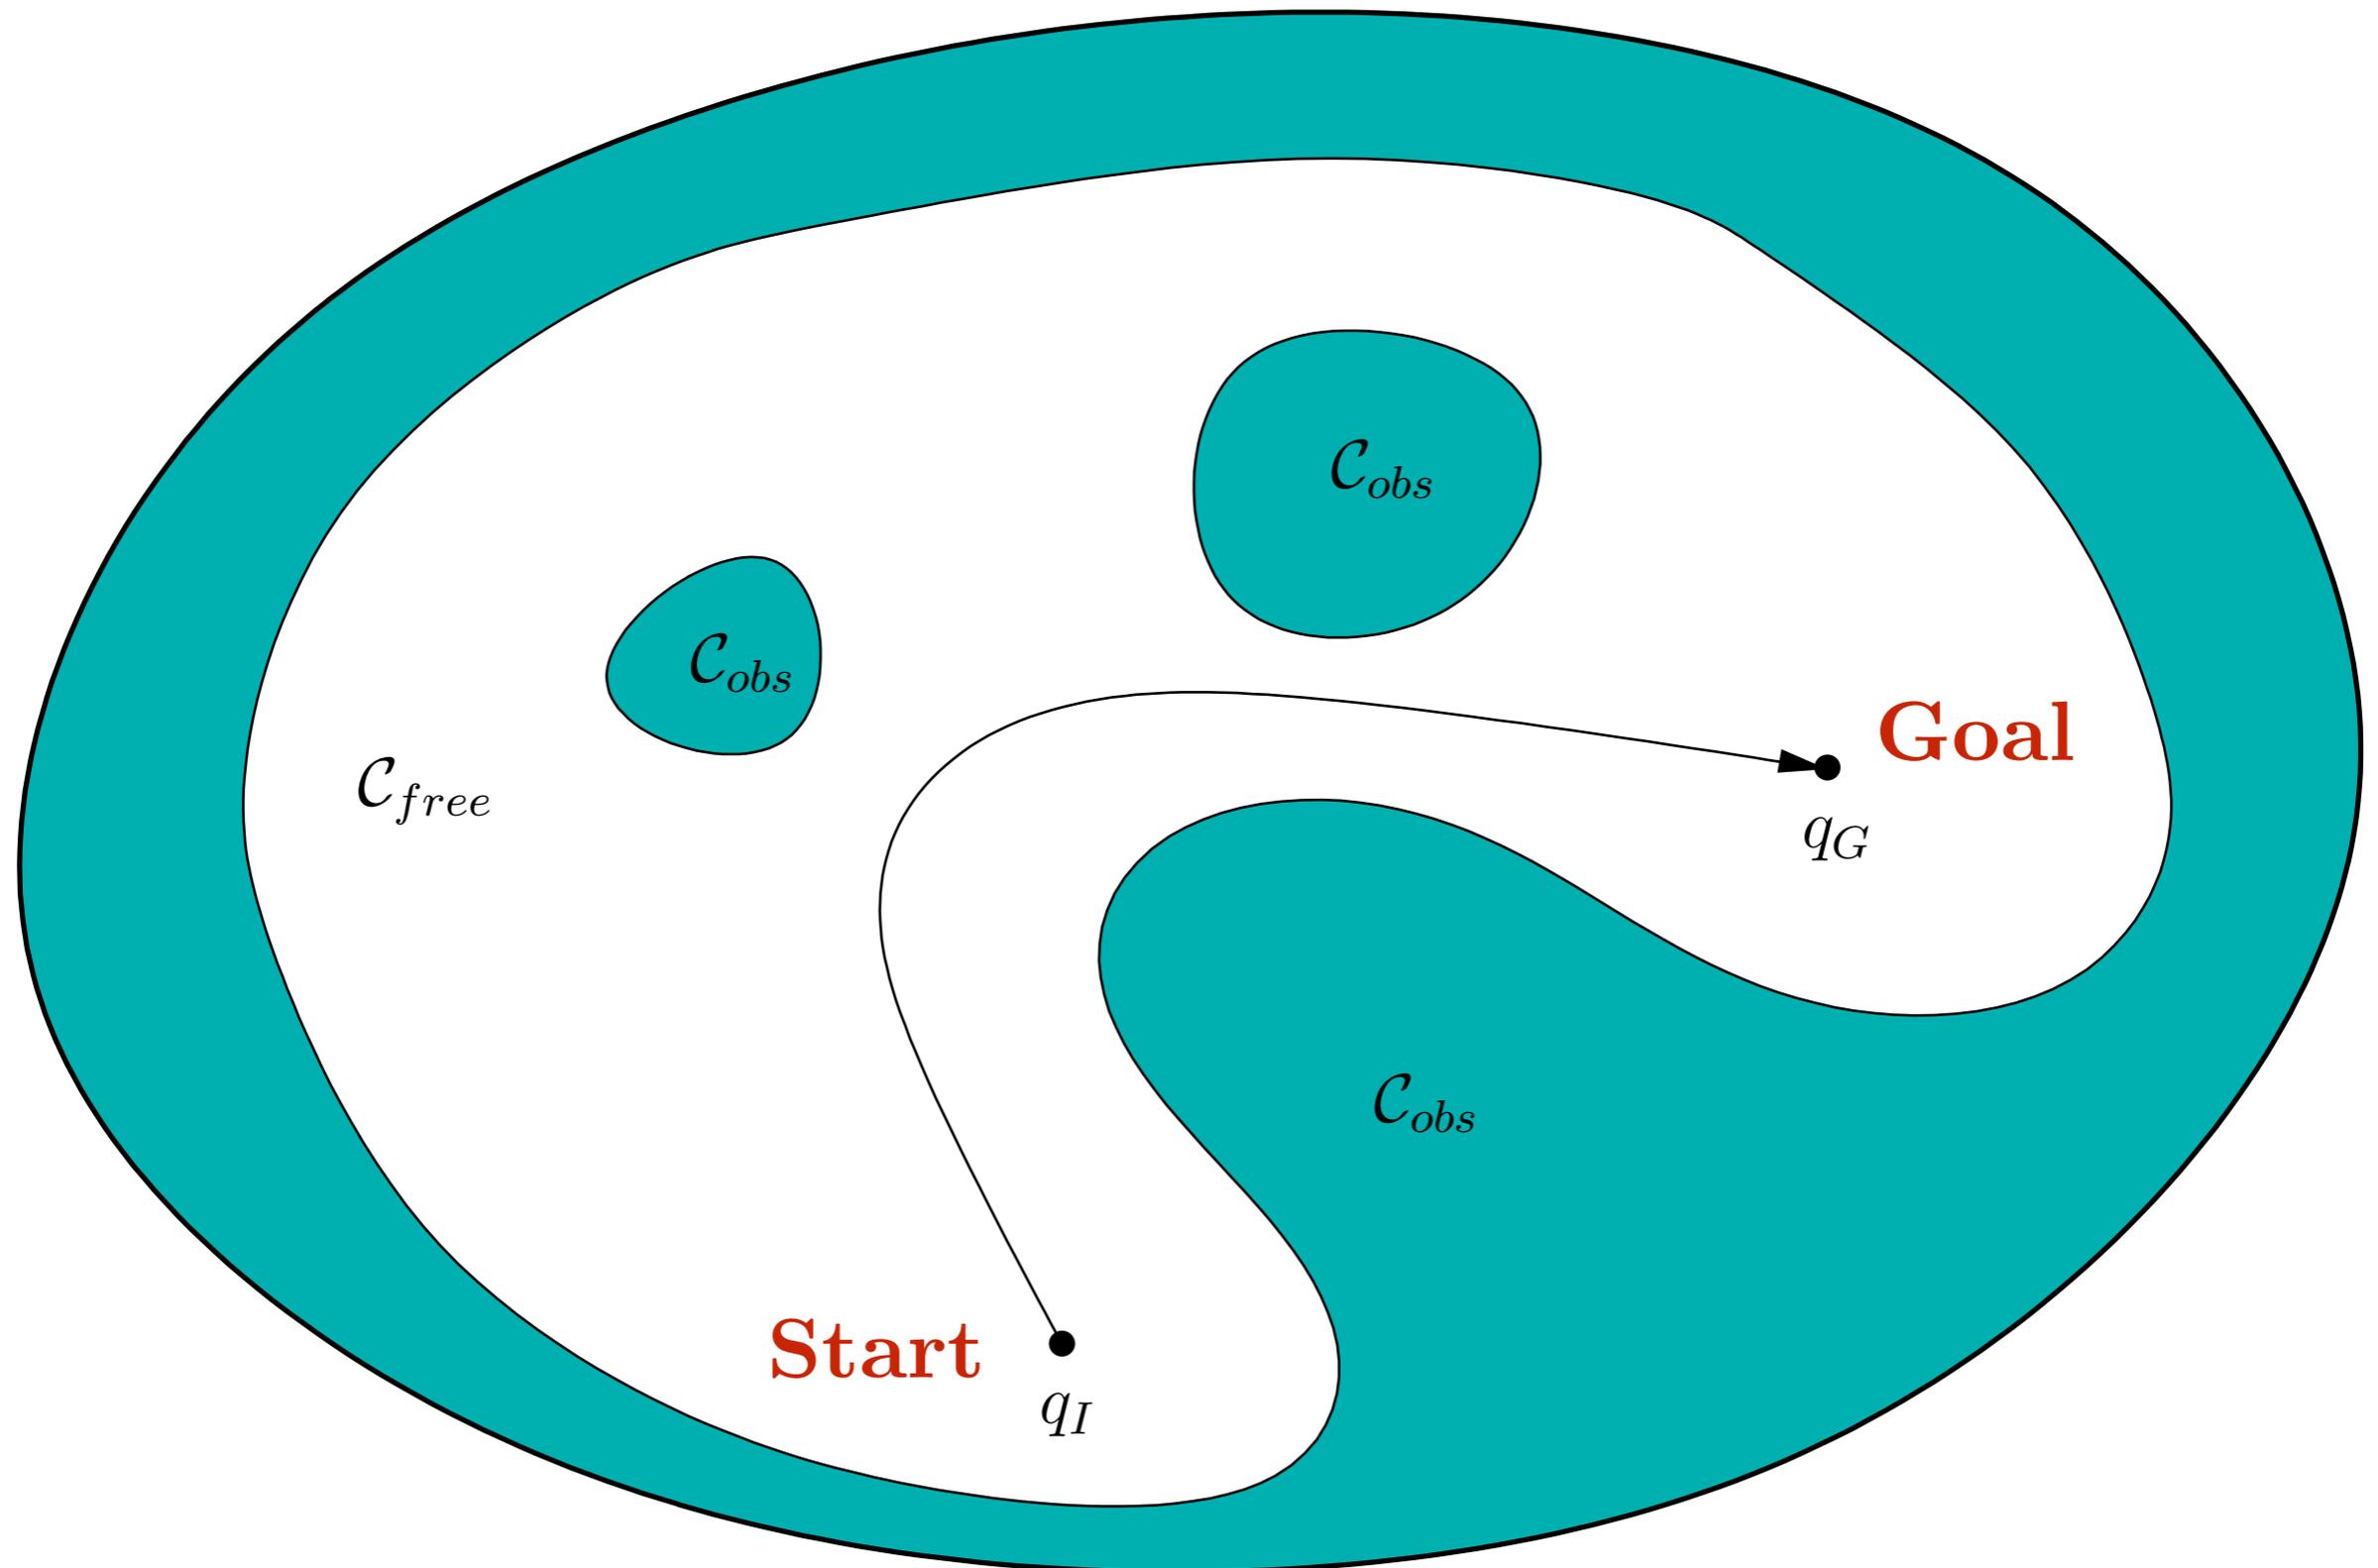


But I just want to know
how to plan for my racecar!



Patience! Upcoming lec on differential constraints

Piano Mover's Problem



Theoretical guarantees that we desire

Completeness

Optimality

Theoretical guarantees that we desire

Completeness

A planner is complete if for any input, it correctly reports whether or not a feasible path exists in **finite time**

Optimality

Theoretical guarantees that we desire

Completeness

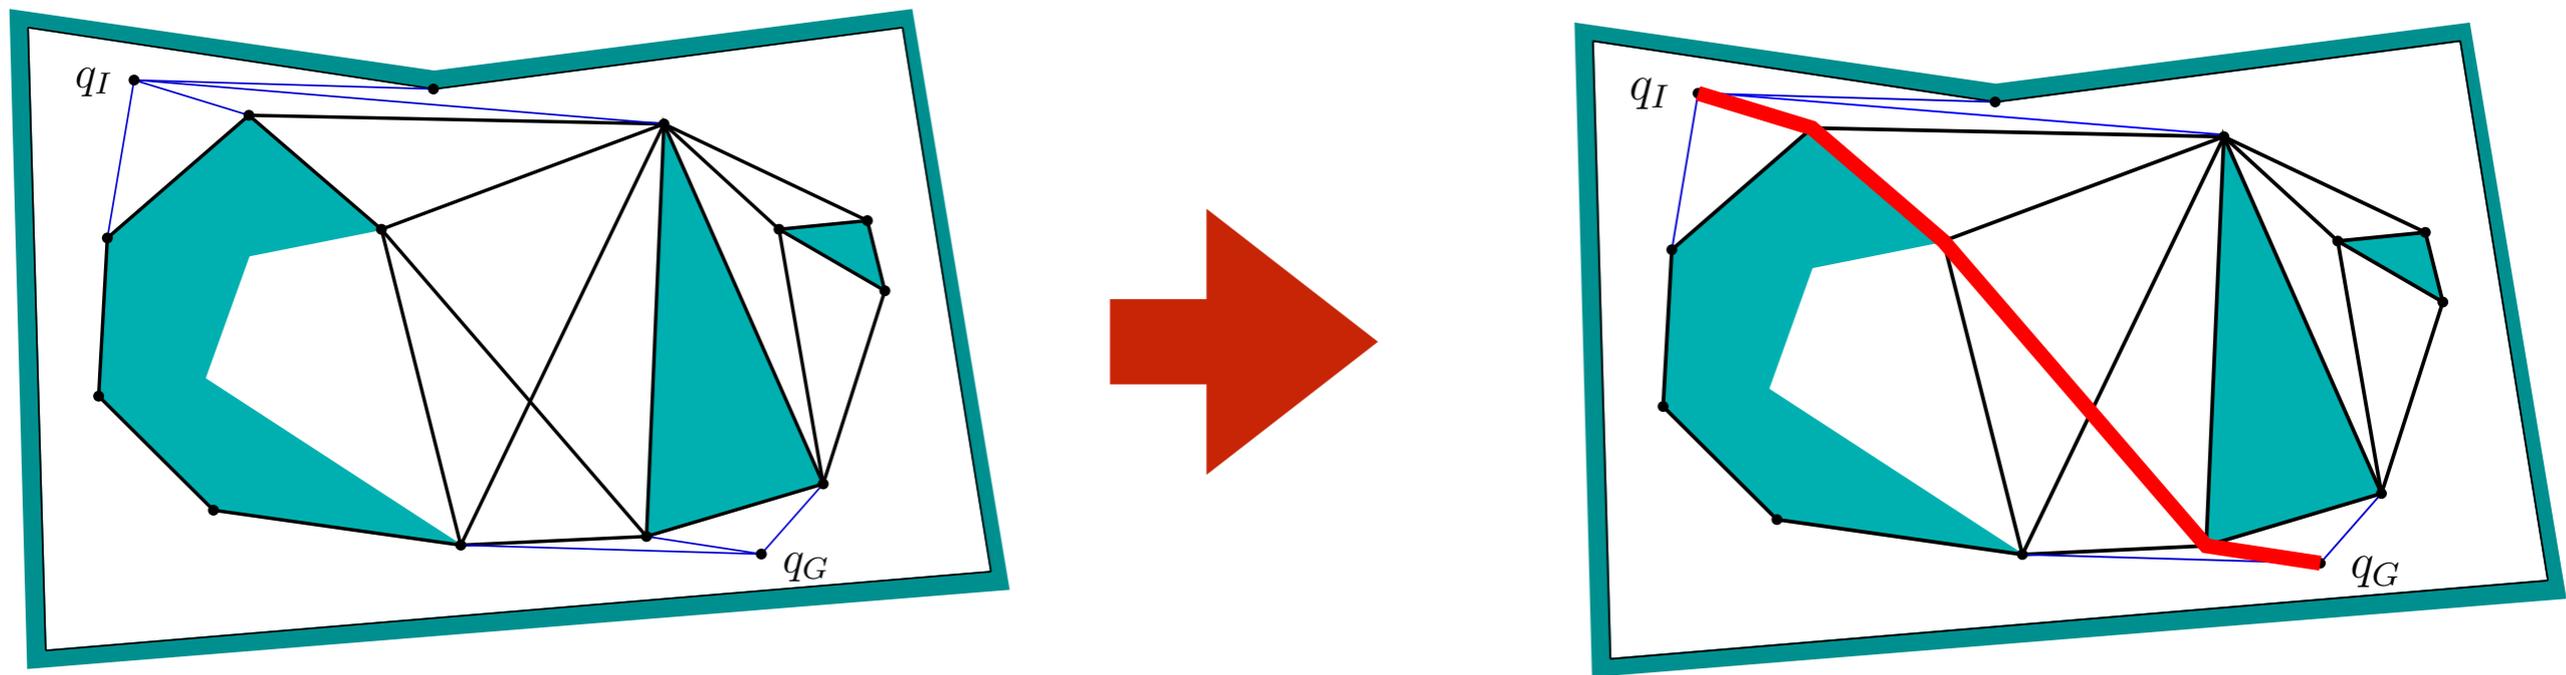
A planner is complete if for any input, it correctly reports whether or not a feasible path exists in **finite time**

Optimality

Returns the best solution in finite time.

Is there any planner that guarantees this?

Yes! 2D Visibility Graphs!



E.g. 2D polygon robots / obstacles can be solved
with visibility graphs

Typical runtime: $O(N^2 \log N)$

So, are we done ... ?

So, are we done ... ?

No! Planning in general is **hard**

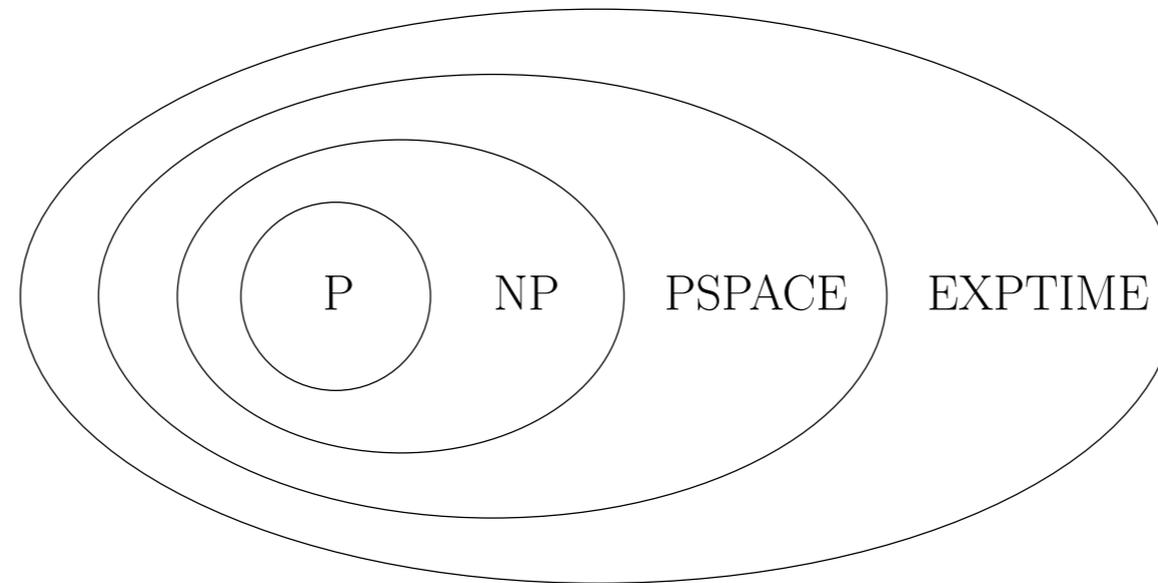
Hardness of motion planning

Hardness of motion planning

Piano Mover's problem is PSPACE-hard (Reif et al. 79)

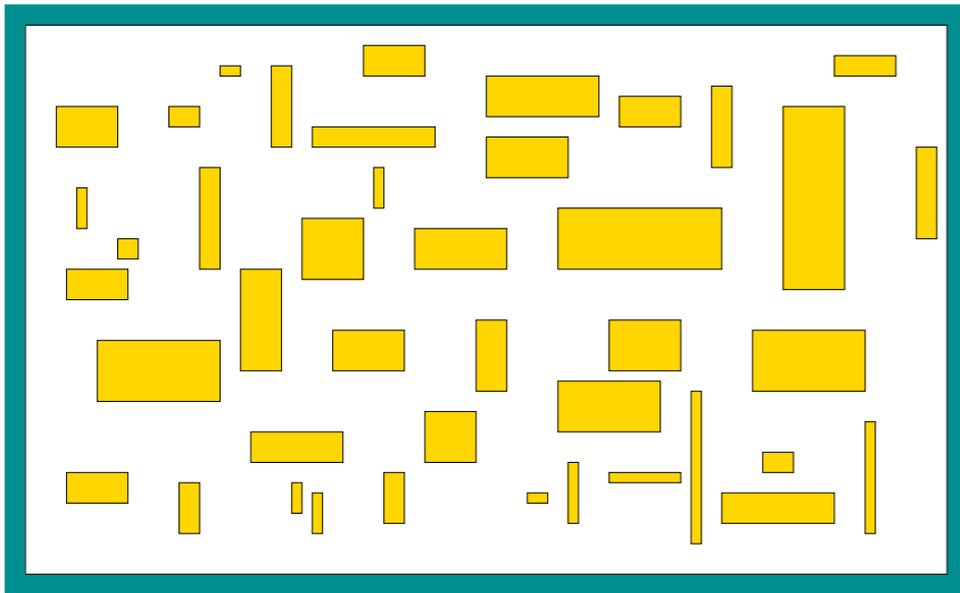
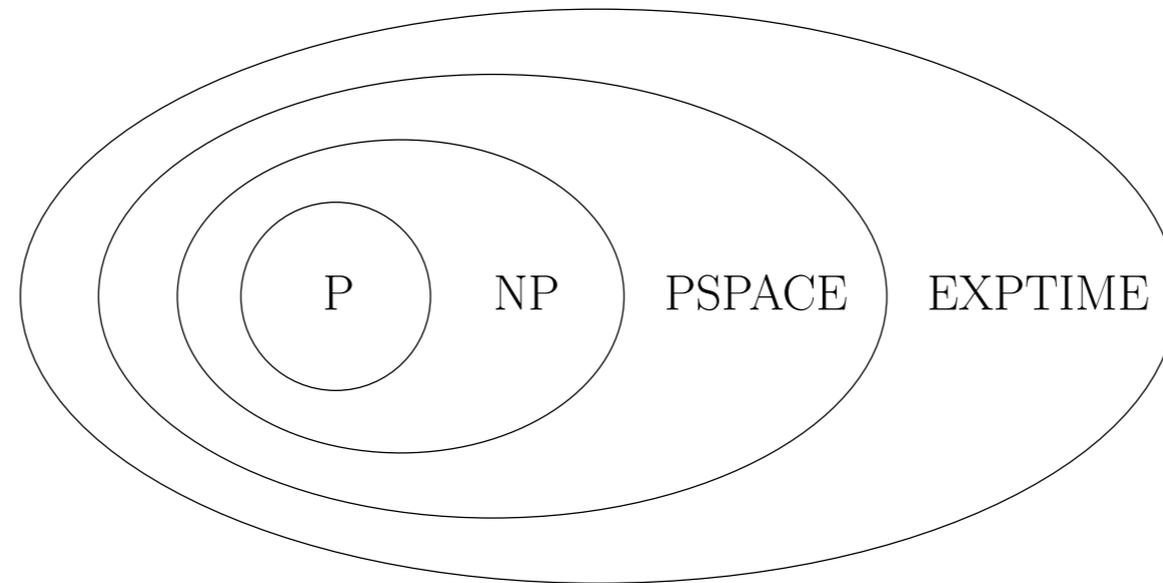
Hardness of motion planning

Piano Mover's problem is PSPACE-hard (Reif et al. 79)



Hardness of motion planning

Piano Mover's problem is PSPACE-hard (Reif et al. 79)

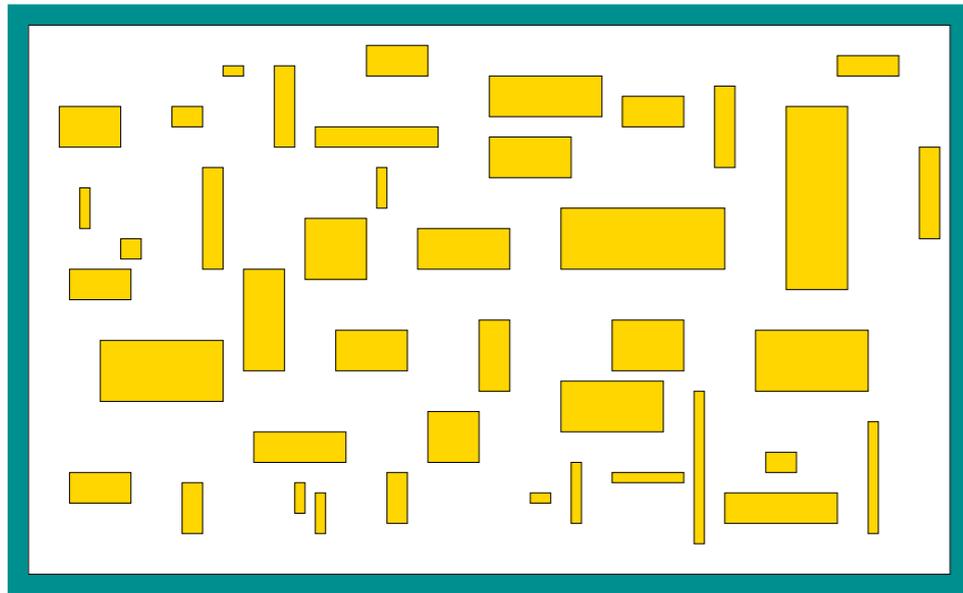
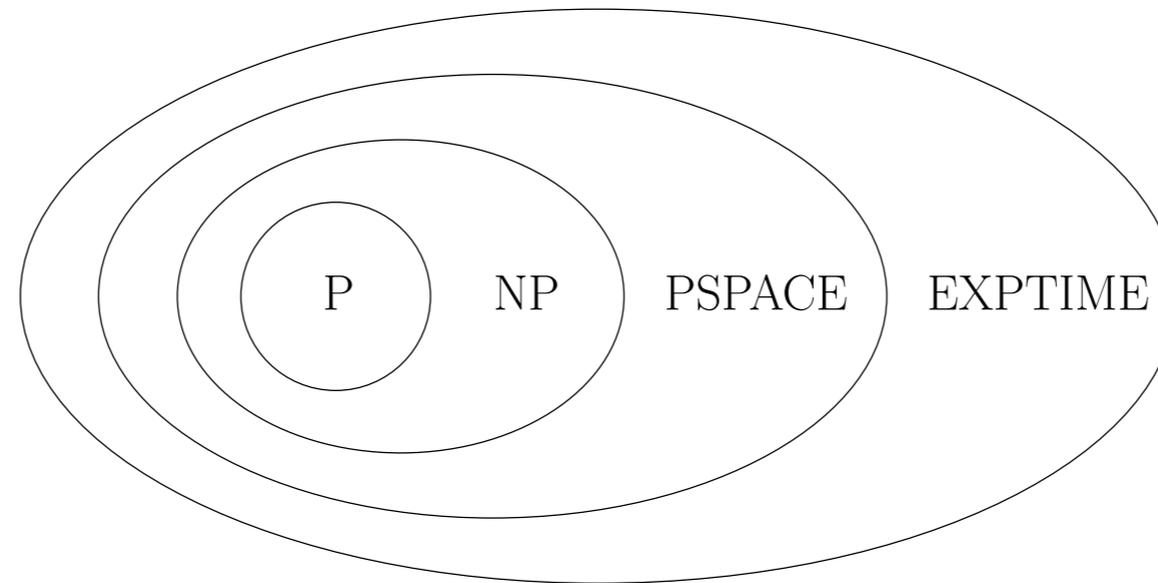


Even planning for translating rectangles is PSPACE-hard!

(Hopcroft et al. 84)

Hardness of motion planning

Piano Mover's problem is PSPACE-hard (Reif et al. 79)



Even planning for translating rectangles is PSPACE-hard!

(Hopcroft et al. 84)

Certain 3D robot planning under uncertainty is NEXPTIME-hard!

(Canny et al. 87)

Why is it so hard?

1. Computing the C-space obstacle in high dimensions is **hard**

2. Planning in continuous high-dimension space is **hard**

Exponential dependency on dimension

Why is it so hard?

1. Computing the C-space obstacle in high dimensions is **hard**
2. Planning in continuous high-dimension space is **hard**

Why is it so hard?

~~1. Computing the C space obstacle in high dimensions is hard~~

2. Planning in continuous high-dimension space is hard

Why is it so hard?

~~1. Computing the C space obstacle in high dimensions is hard~~

We won't! Instead we will use a collision checker!

2. Planning in continuous high-dimension space is hard

Why is it so hard?

~~1. Computing the C space obstacle in high dimensions is hard~~

We won't! Instead we will use a collision checker!

~~2. Planning in continuous high-dimension space is hard~~

Why is it so hard?

~~1. Computing the C space obstacle in high dimensions is hard~~

We won't! Instead we will use a collision checker!

~~2. Planning in continuous high-dimension space is hard~~

We will bring it to discrete space by sampling configurations!

Research in Motion Planning:

Make good approximations

Research in Motion Planning:

Make good approximations
(that have guarantees)

Today's objective

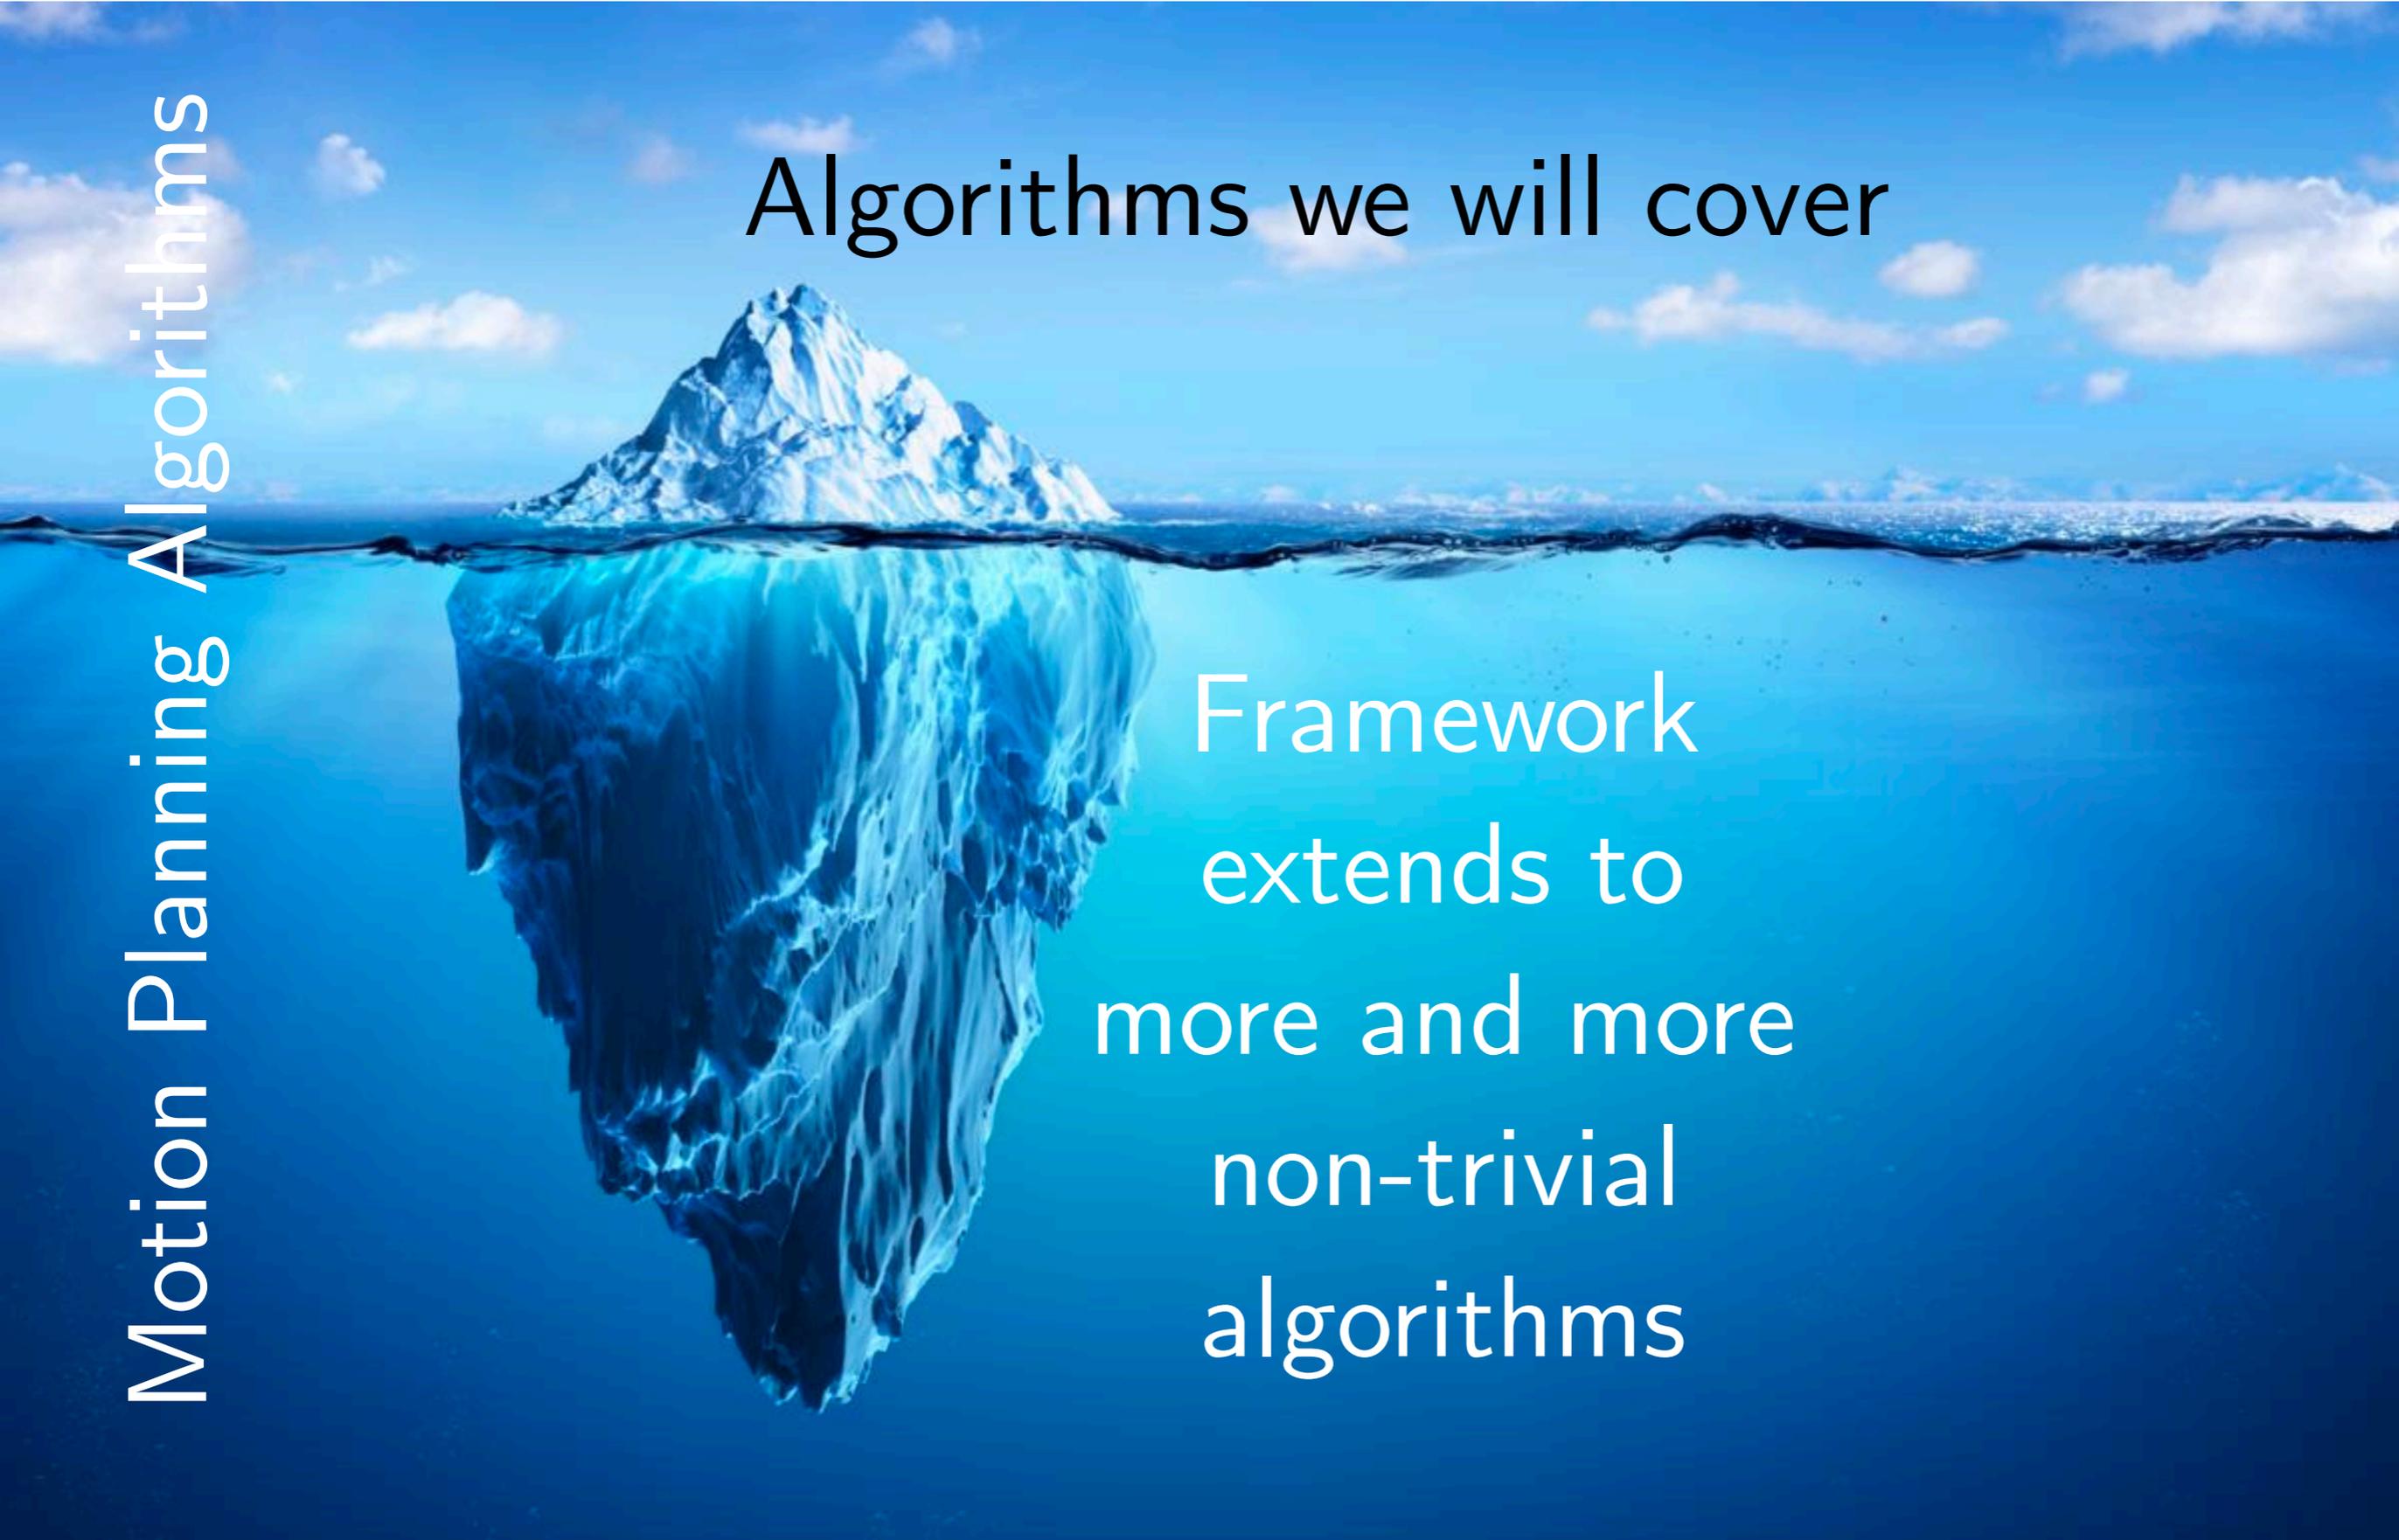
1. General framework for motion planning
2. Inputs to any planner: Collision checking and steering
3. Planning on roadmaps - one class of instantiations of the framework

Why an abstract framework?

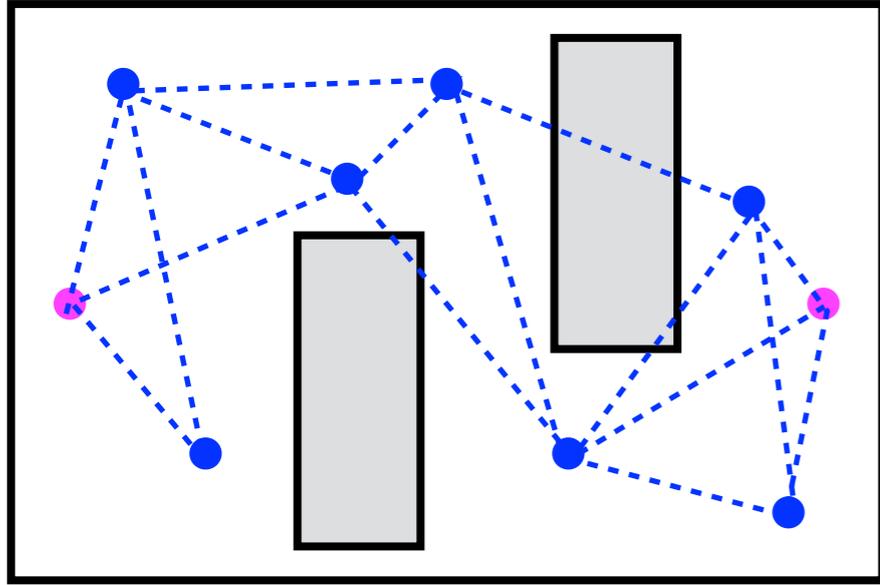
Motion Planning Algorithms

Algorithms we will cover

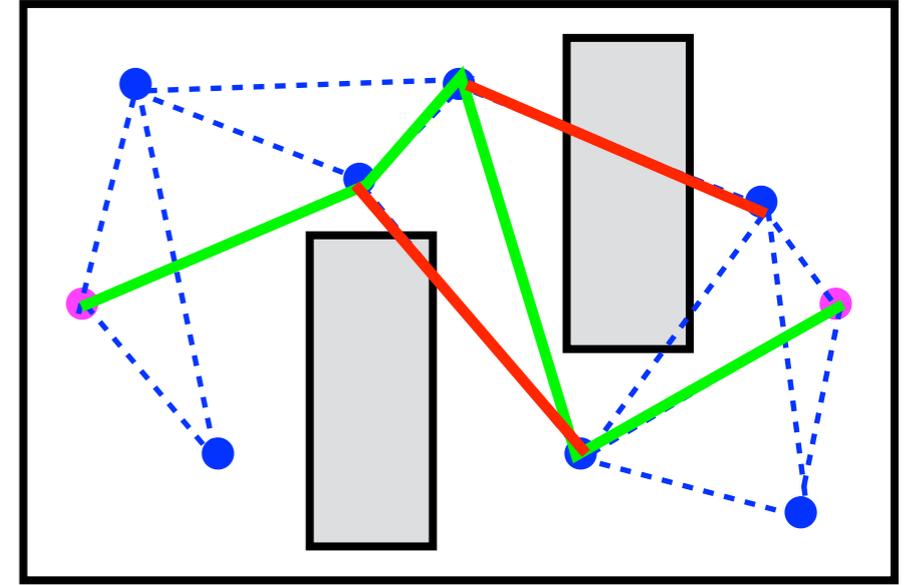
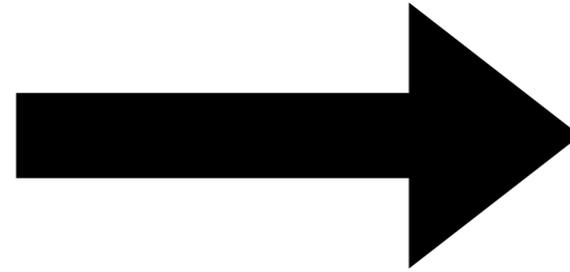
Framework
extends to
more and more
non-trivial
algorithms



General framework for motion planning

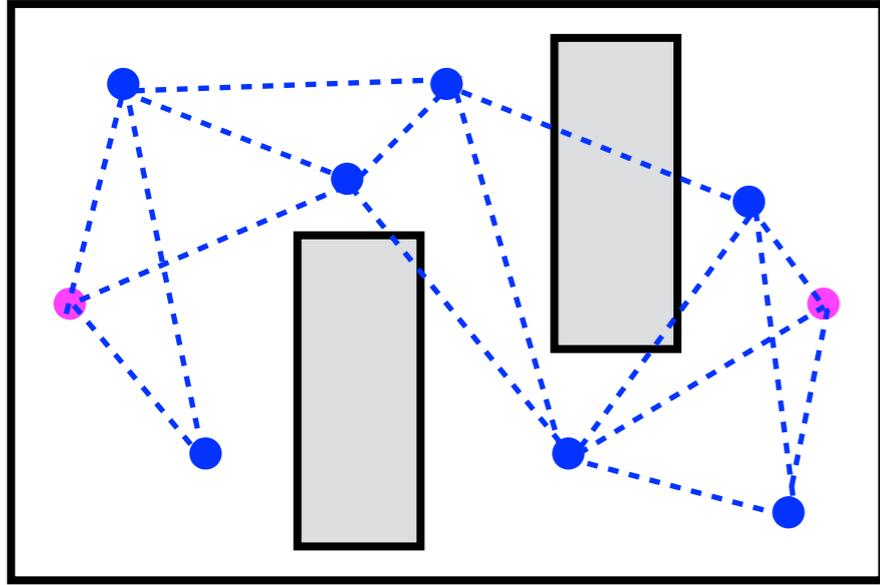


Create a graph

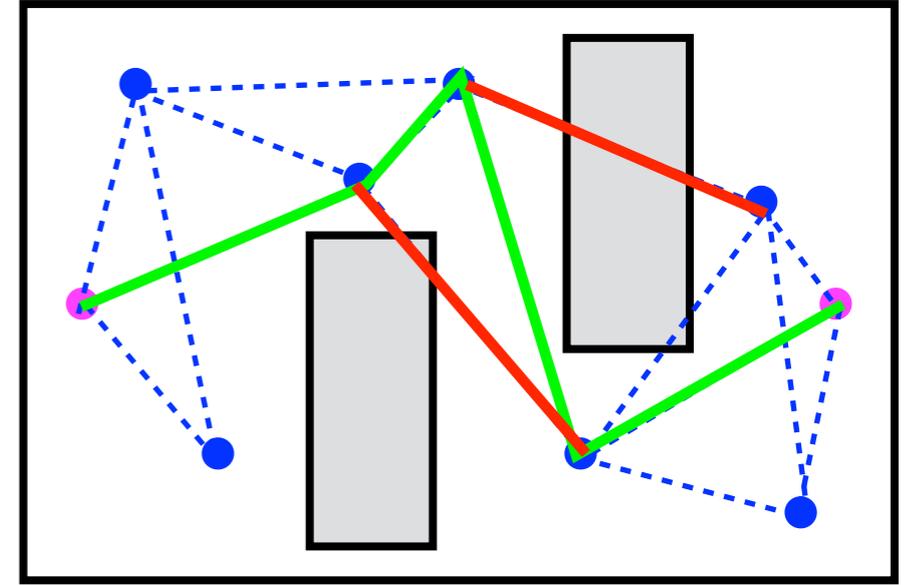
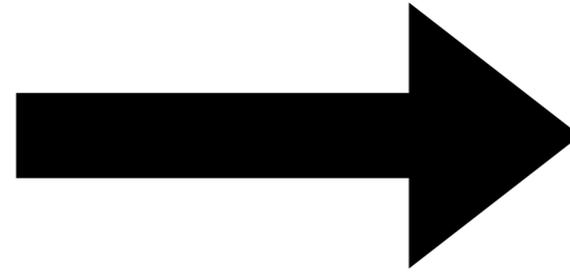


Search the graph

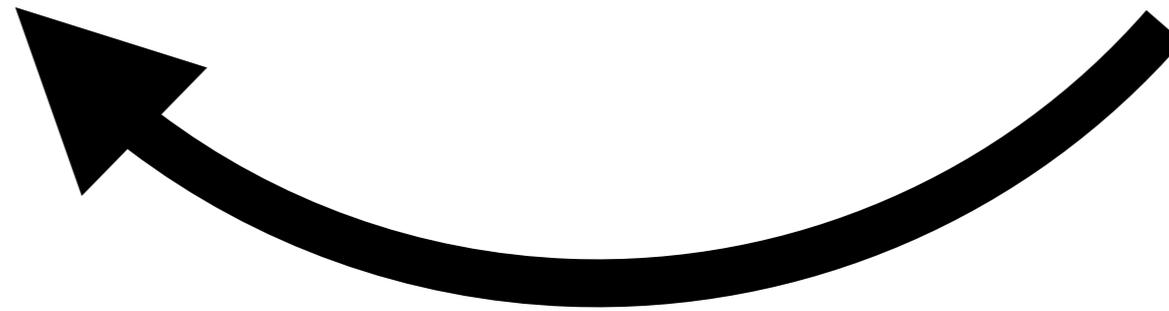
General framework for motion planning



Create a graph



Search the graph



Interleave

General framework for motion planning

Any planning
algorithm

Create graph

Search graph

Interleave

General framework for motion planning

Any planning
algorithm

Create graph

Search graph

Interleave



General framework for motion planning

Any planning
algorithm

Create graph

Search graph

Interleave



=

General framework for motion planning

Any planning algorithm

Create graph

Search graph

Interleave



=

e.g. fancy
random
sampler

×

e.g. fancy
heuristic

×

e.g. fancy
way of
densifying

General framework for motion planning

Any planning algorithm

Create graph

Search graph

Interleave



=

e.g. fancy random sampler

×

e.g. fancy heuristic

×

e.g. fancy way of densifying

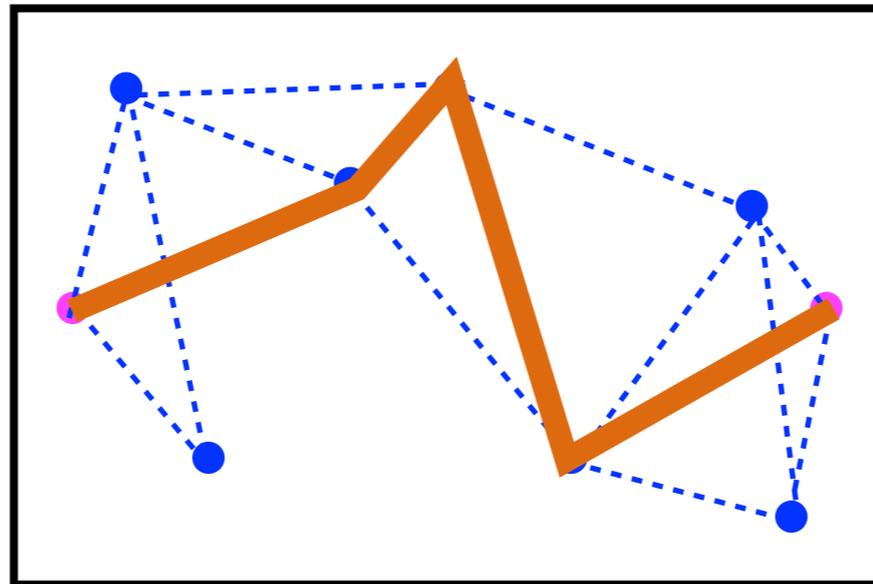
Whats the best we can do?

Whats the best we can do?

Whats the best we can do?

For this lecture....

Assume you are given a super awesome search subroutine!

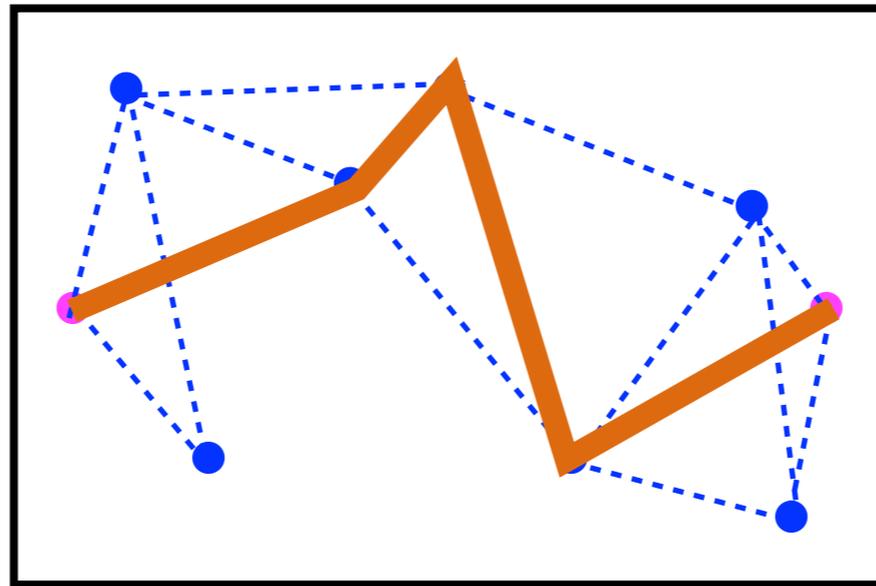


Optimal Path = $\text{SHORTESTPATH}(V, E, \text{start}, \text{goal})$

(Next lecture we will talk about how we get this)

For this lecture....

Assume you are given a super awesome search subroutine!



Optimal Path = $\text{SHORTESTPATH}(V, E, \text{start}, \text{goal})$

(Next lecture we will talk about how we get this)

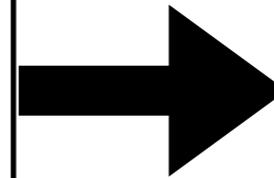
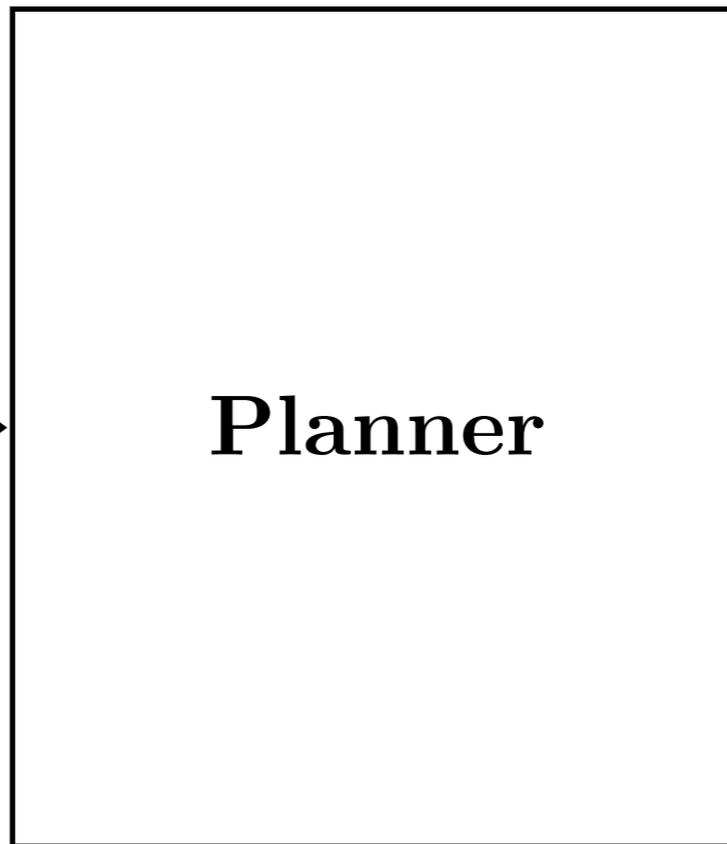
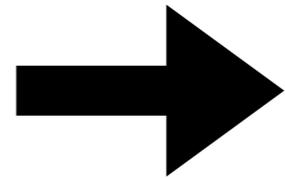
Assume complexity is $O(|V| \log |V| + |E|)$

API for motion planning

Input

1. A collision checker
 $\text{coll}(q)$

2. Steering method
 $\text{steer}(q_1, q_2)$



Output

Collision
free path
joining
start and goal

Let's take a look at the inputs

We need to give the planner a collision checker

$$\text{coll}(q) = \begin{cases} 0 & \text{in collision, i.e. } q \in \mathcal{C}_{obs} \\ 1 & \text{free, i.e. } q \in \mathcal{C}_{free} \end{cases}$$

What work does this function have to do?

Let's take a look at the inputs

We need to give the planner a collision checker

$$\text{coll}(q) = \begin{cases} 0 & \text{in collision, i.e. } q \in \mathcal{C}_{obs} \\ 1 & \text{free, i.e. } q \in \mathcal{C}_{free} \end{cases}$$

What work does this function have to do?

Collision checking is **expensive!**

Let's take a look at the inputs

We need to give the planner a steer function

steer(q_1, q_2)

A steer function tries to join two configurations with a feasible path

Computes simple path, calls `coll(q)`, and returns success if path is free

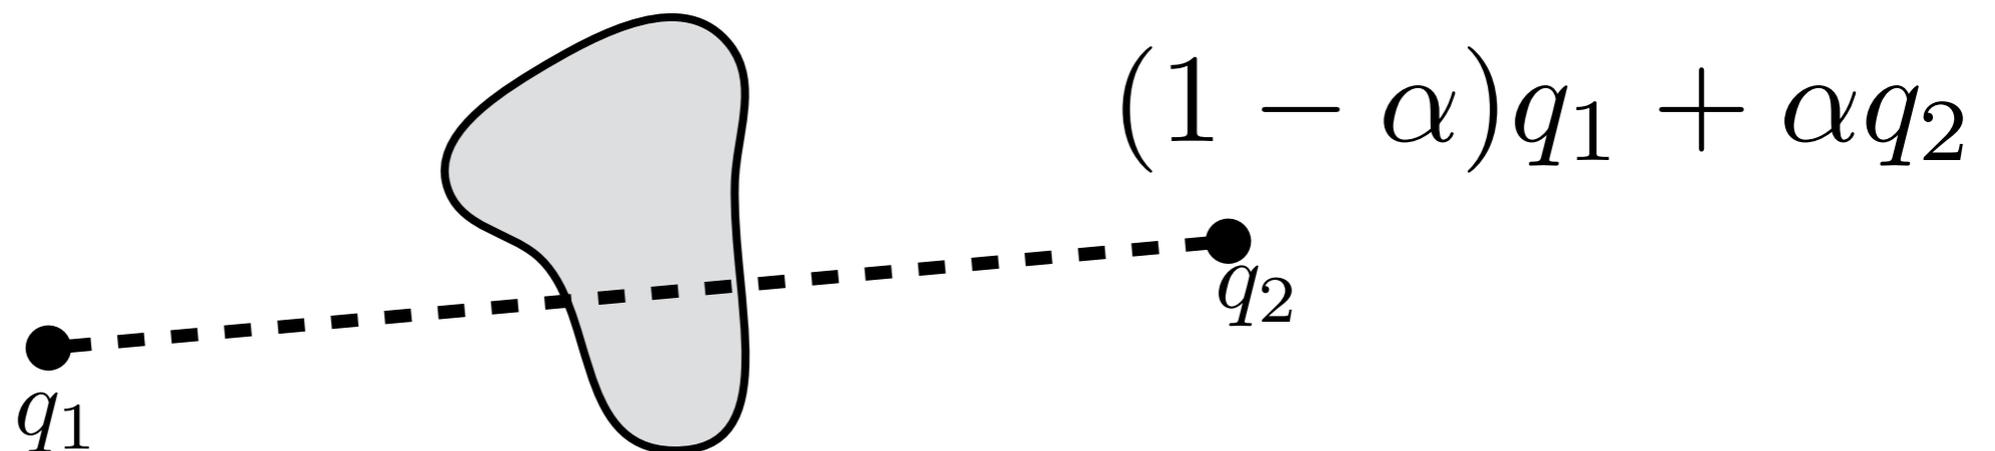
Let's take a look at the inputs

We need to give the planner a steer function

$$\text{steer}(q_1, q_2)$$

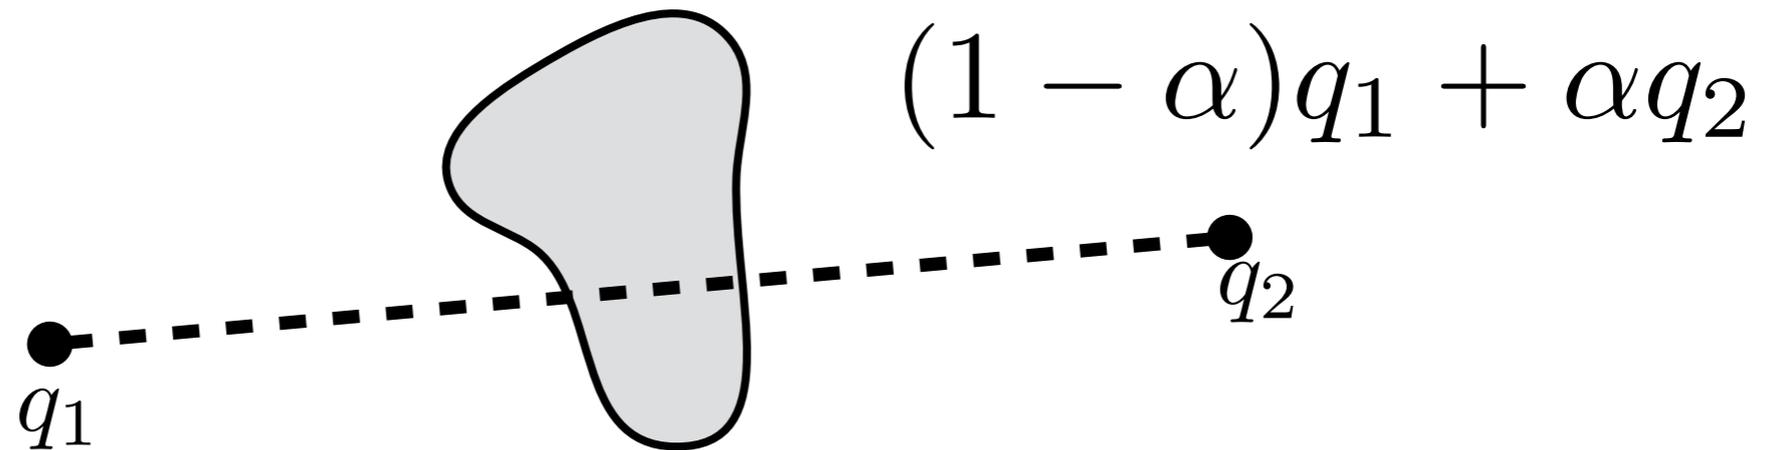
A steer function tries to join two configurations with a feasible path

Computes simple path, calls $\text{coll}(q)$, and returns success if path is free



Example: Connect them with a straight line and check for feasibility

Can steer be smart about collision checking?



$\text{steer}(q_1, q_2)$ has to assure us line is collision free (upto a resolution)

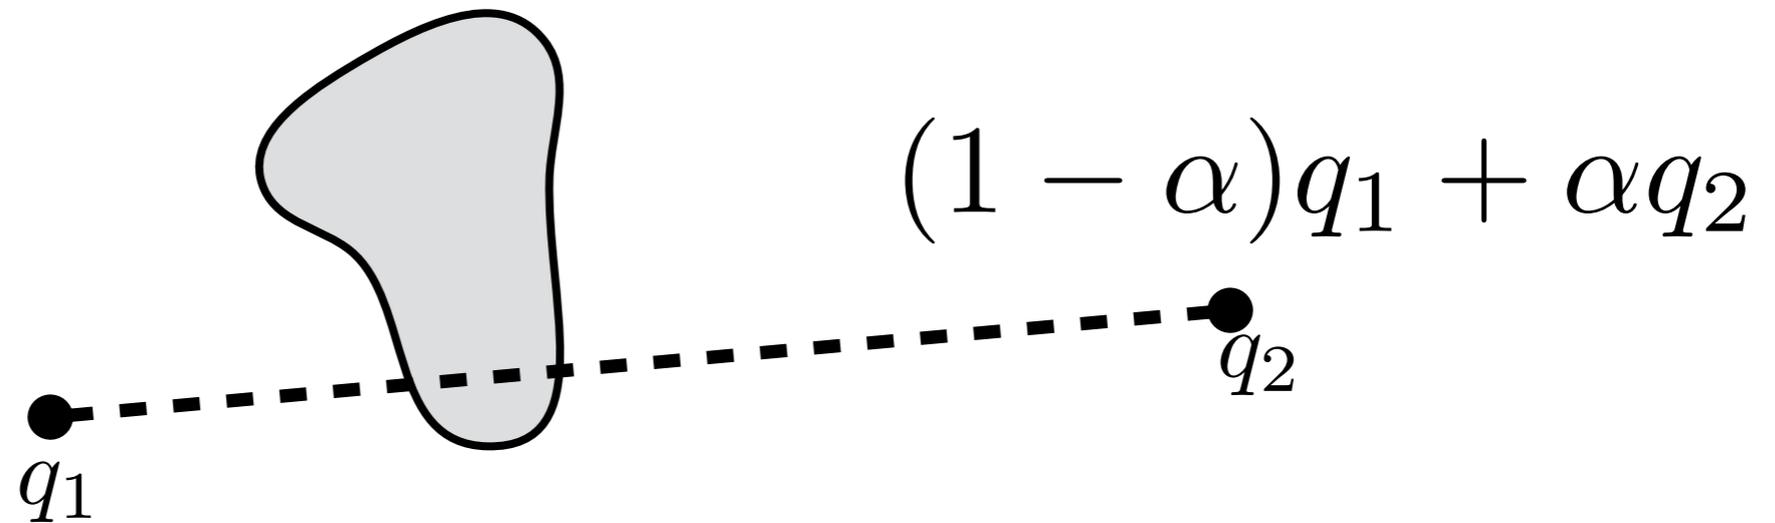
Things we can try:

1. Step forward along the line and check each point
2. Step backwards along the line and check each point

.....

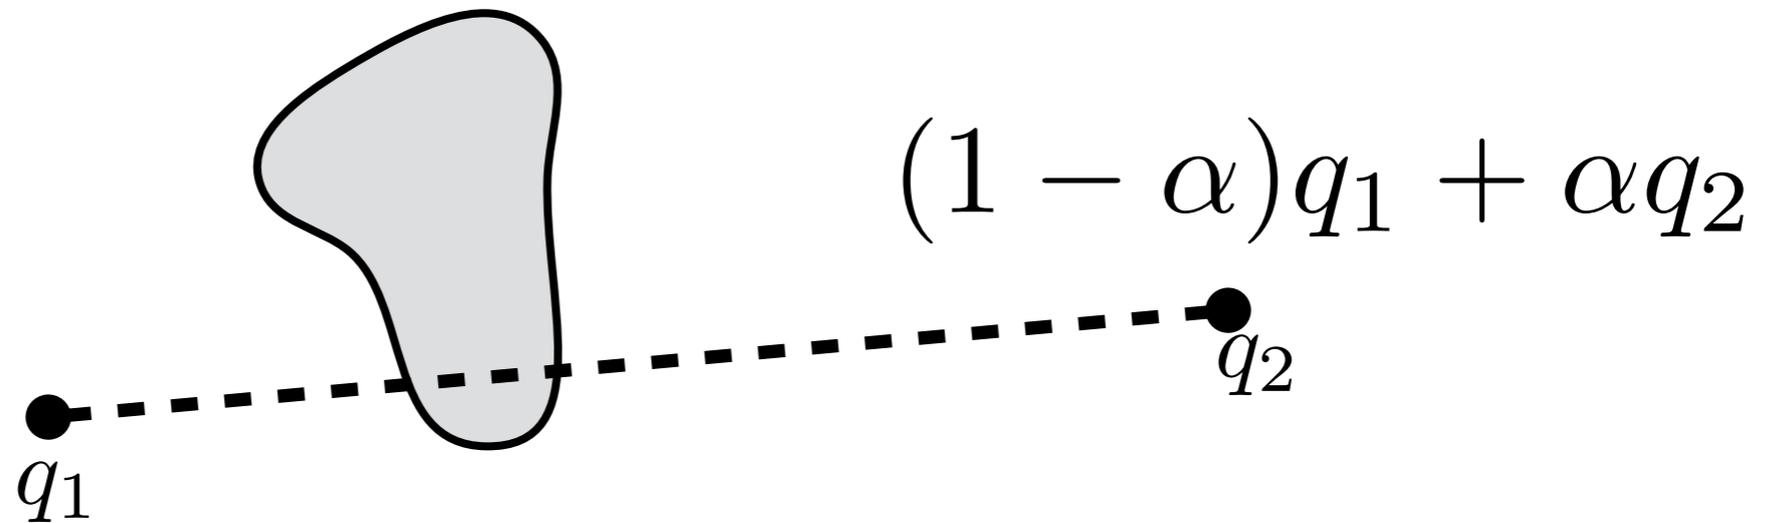
Can steer be smart about collision checking?

Say we chunk the line into 16 parts



Can steer be smart about collision checking?

Say we chunk the line into 16 parts

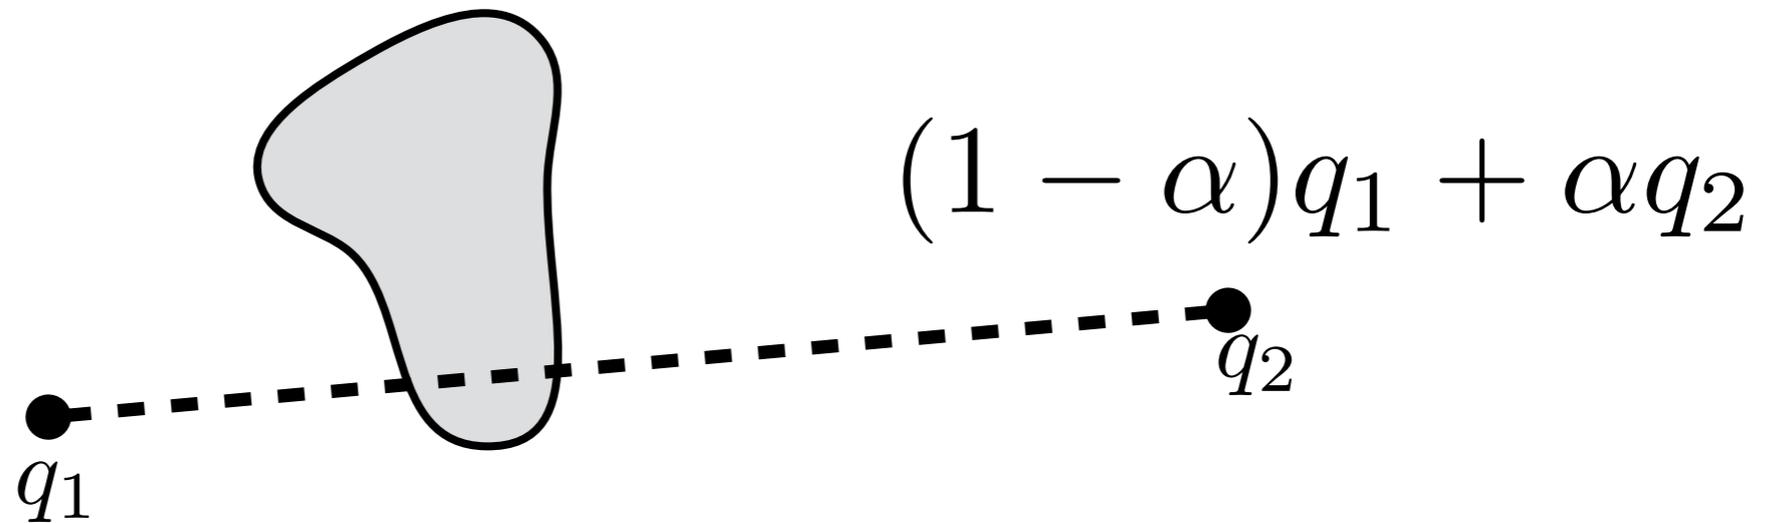


Any collision checking strategy corresponds to sequence

(Naive) $\alpha = 0, \frac{1}{16}, \frac{2}{16}, \frac{3}{16}, \dots, \frac{15}{16}$

Can steer be smart about collision checking?

Say we chunk the line into 16 parts



Any collision checking strategy corresponds to sequence

(Naive) $\alpha = 0, \frac{1}{16}, \frac{2}{16}, \frac{3}{16}, \dots, \frac{15}{16}$

(Bisection) $\alpha = 0, \frac{8}{16}, \frac{4}{16}, \frac{12}{16}, \dots, \frac{15}{16}$

Ans: Van der Corput sequence

i	Naive Sequence
1	0
2	1/16
3	1/8
4	3/16
5	1/4
6	5/16
7	3/8
8	7/16
9	1/2
10	9/16
11	5/8
12	11/16
13	3/4
14	13/16
15	7/8
16	15/16

Ans: Van der Corput sequence

i	Naive Sequence	Binary
1	0	.0000
2	1/16	.0001
3	1/8	.0010
4	3/16	.0011
5	1/4	.0100
6	5/16	.0101
7	3/8	.0110
8	7/16	.0111
9	1/2	.1000
10	9/16	.1001
11	5/8	.1010
12	11/16	.1011
13	3/4	.1100
14	13/16	.1101
15	7/8	.1110
16	15/16	.1111

Ans: Van der Corput sequence

i	Naive Sequence	Binary	Reverse Binary
1	0	.0000	.0000
2	1/16	.0001	.1000
3	1/8	.0010	.0100
4	3/16	.0011	.1100
5	1/4	.0100	.0010
6	5/16	.0101	.1010
7	3/8	.0110	.0110
8	7/16	.0111	.1110
9	1/2	.1000	.0001
10	9/16	.1001	.1001
11	5/8	.1010	.0101
12	11/16	.1011	.1101
13	3/4	.1100	.0011
14	13/16	.1101	.1011
15	7/8	.1110	.0111
16	15/16	.1111	.1111

Ans: Van der Corput sequence

i	Naive Sequence	Binary	Reverse Binary	Van der Corput
1	0	.0000	.0000	0
2	1/16	.0001	.1000	1/2
3	1/8	.0010	.0100	1/4
4	3/16	.0011	.1100	3/4
5	1/4	.0100	.0010	1/8
6	5/16	.0101	.1010	5/8
7	3/8	.0110	.0110	3/8
8	7/16	.0111	.1110	7/8
9	1/2	.1000	.0001	1/16
10	9/16	.1001	.1001	9/16
11	5/8	.1010	.0101	5/16
12	11/16	.1011	.1101	13/16
13	3/4	.1100	.0011	3/16
14	13/16	.1101	.1011	11/16
15	7/8	.1110	.0111	7/16
16	15/16	.1111	.1111	15/16

Ans: Van der Corput sequence

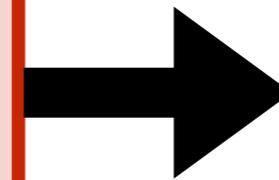
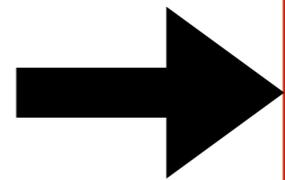
i	Naive Sequence	Binary	Reverse Binary	Van der Corput	Points in $[0, 1]/ \sim$
1	0	.0000	.0000	0	
2	1/16	.0001	.1000	1/2	
3	1/8	.0010	.0100	1/4	
4	3/16	.0011	.1100	3/4	
5	1/4	.0100	.0010	1/8	
6	5/16	.0101	.1010	5/8	
7	3/8	.0110	.0110	3/8	
8	7/16	.0111	.1110	7/8	
9	1/2	.1000	.0001	1/16	
10	9/16	.1001	.1001	9/16	
11	5/8	.1010	.0101	5/16	
12	11/16	.1011	.1101	13/16	
13	3/4	.1100	.0011	3/16	
14	13/16	.1101	.1011	11/16	
15	7/8	.1110	.0111	7/16	
16	15/16	.1111	.1111	15/16	

Now we are ready to talk about planner!

Input

1. A collision checker
 $\text{coll}(q)$

2. Steering method
 $\text{steer}(q_1, q_2)$



Output

Collision
free path
joining
start and goal

Framework for planner

1. Create a graph

(Think about what makes a good graph as we go along)

2. Search the graph (assume solved for now)

Creating a graph: Abstract algorithm

$$G = (V, E)$$

Vertices: set of configurations

Edges: paths connecting
configurations

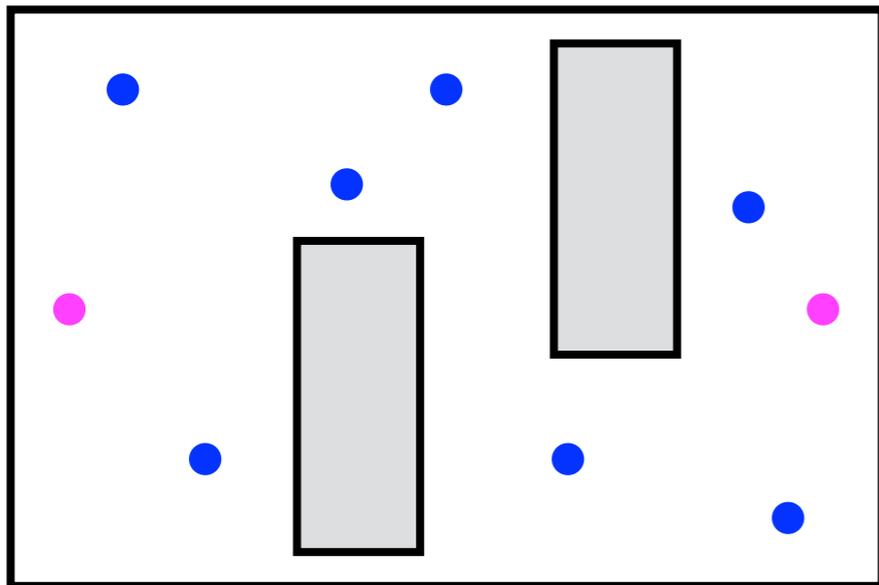
Creating a graph: Abstract algorithm

$$G = (V, E)$$

Vertices: set of configurations

Edges: paths connecting configurations

1. Sample a set of collision free vertices V (add start and goal)



Sample a configuration q

if $\text{coll}(q) = 1$

$$V \leftarrow V \cup \{q\}$$

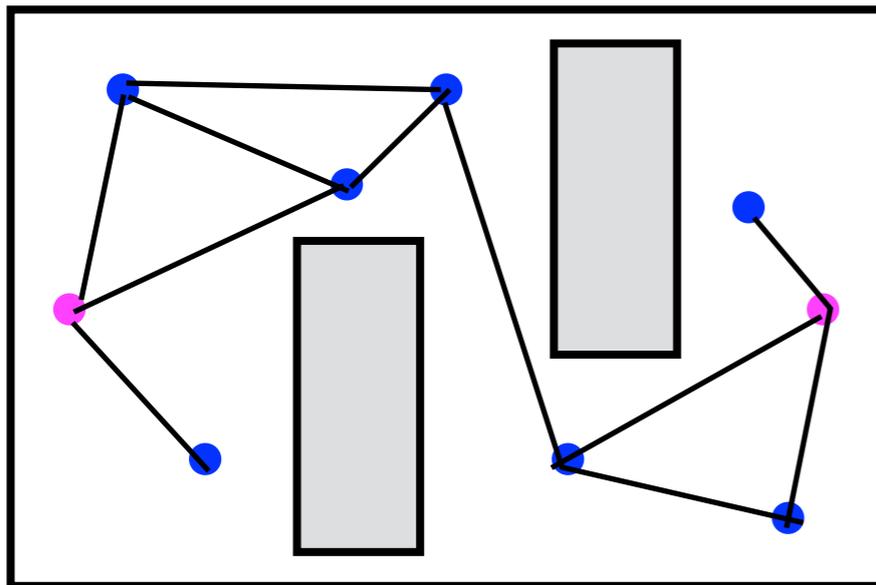
Creating a graph: Abstract algorithm

$$G = (V, E)$$

Vertices: set of configurations

Edges: paths connecting configurations

1. Sample a set of collision free vertices V (add start and goal)
2. Connect “neighboring” vertices to get edges E

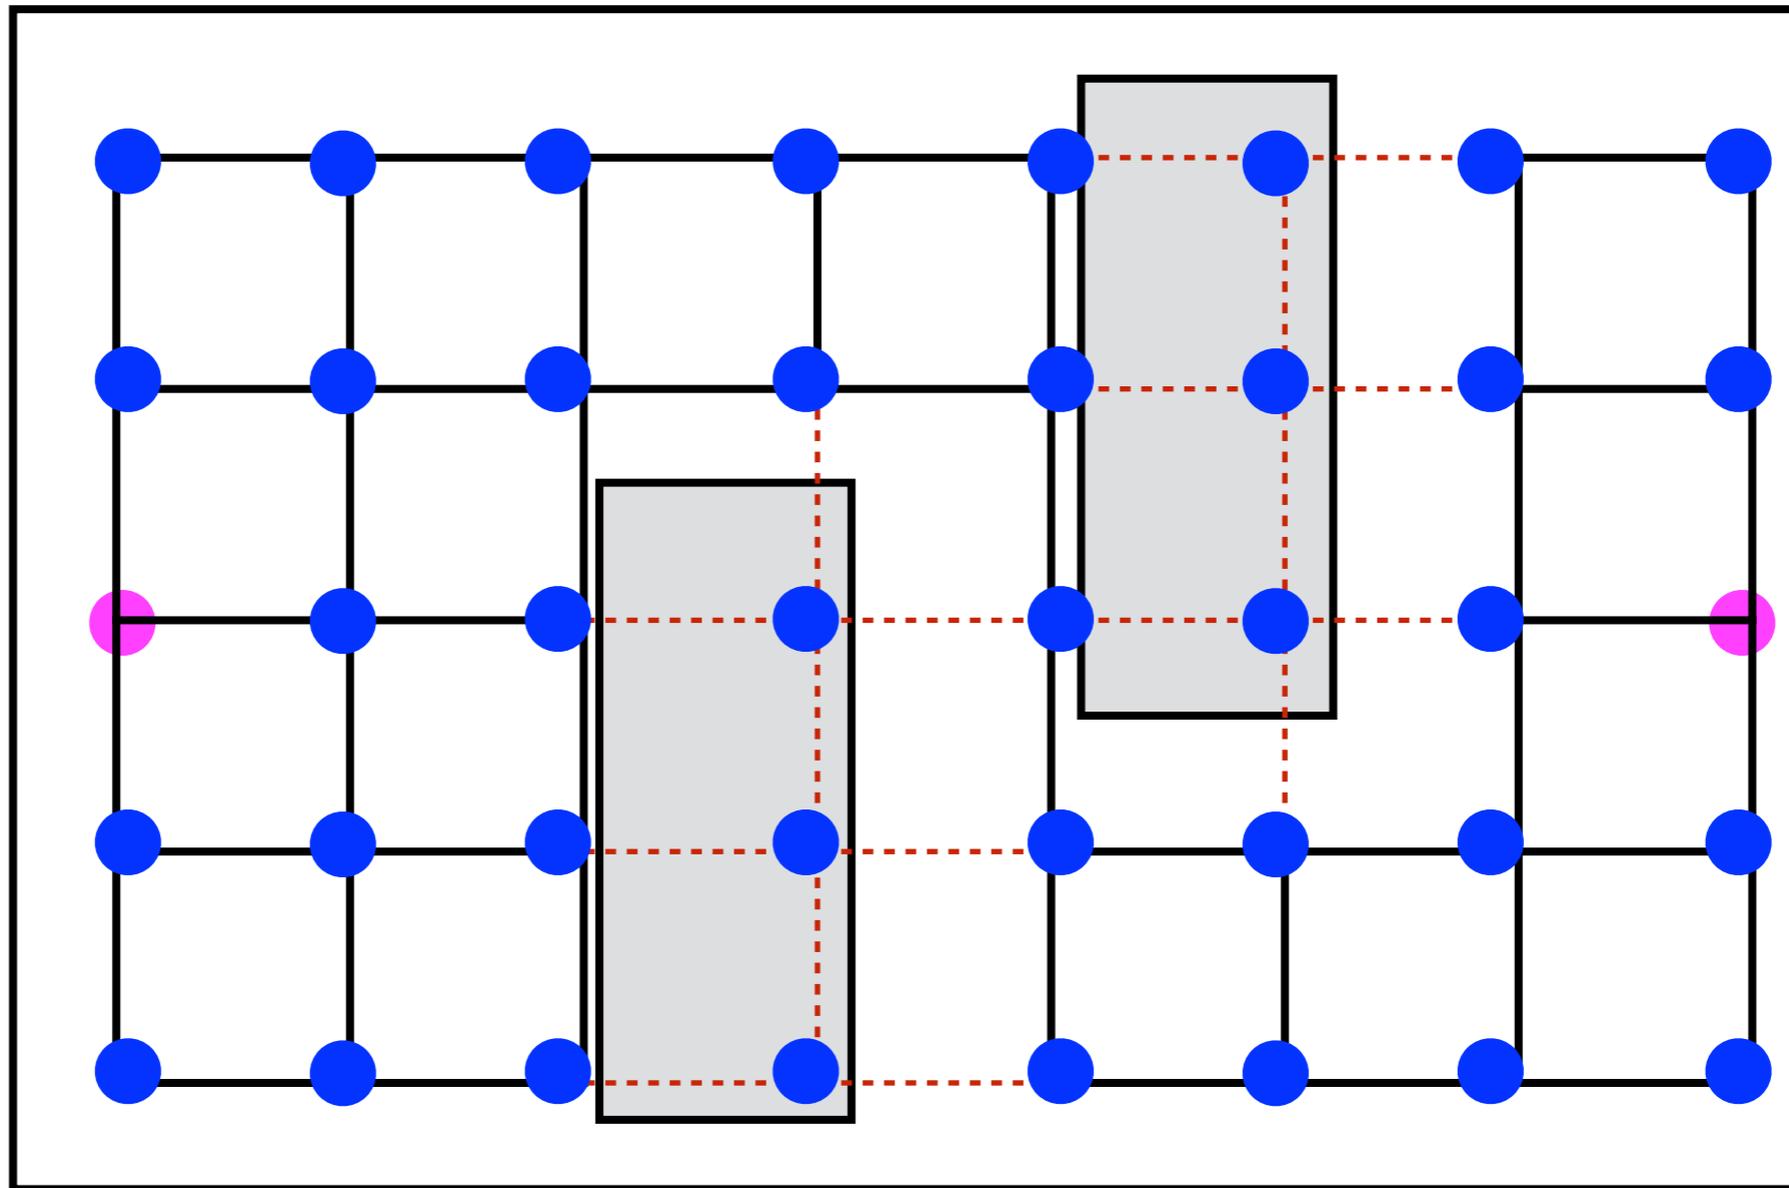


for each candidate pair (v_1, v_2)
if `steer` (v_1, v_2) succeeds

$$E \leftarrow E \cup (v_1, v_2)$$

Strategy 1: Discretize configuration space

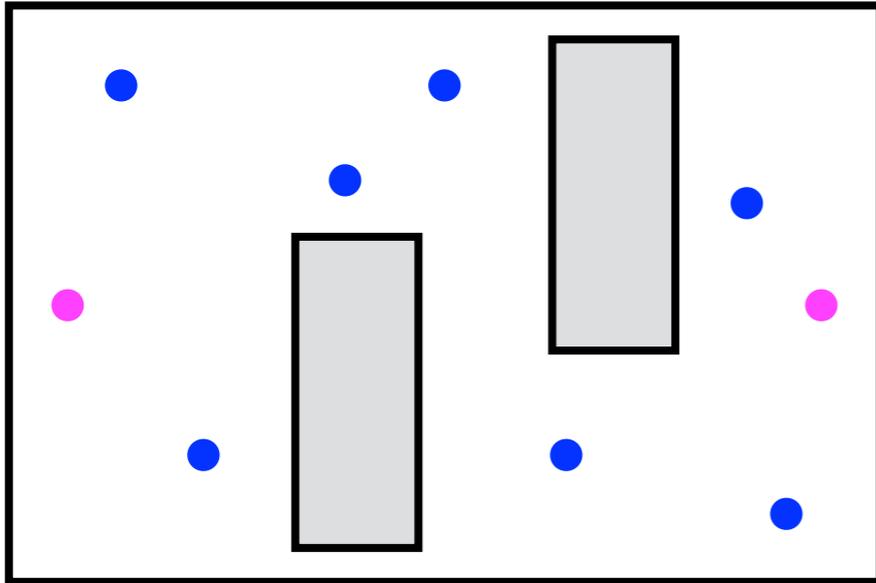
Create a lattice. Connect neighboring points (4-conn, 8-conn, ...)



Theoretical guarantees: Resolution complete

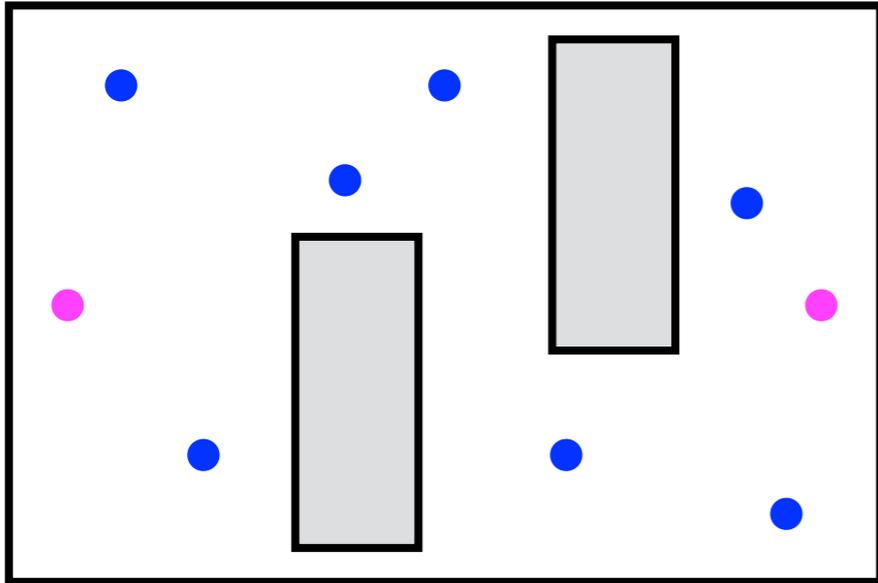
What are the pros? What are the cons?

Strategy 2: Uniformly randomly sample



If C-space is a real vector space
for each dimension i
sample $q(i) \sim [lb, ub]$

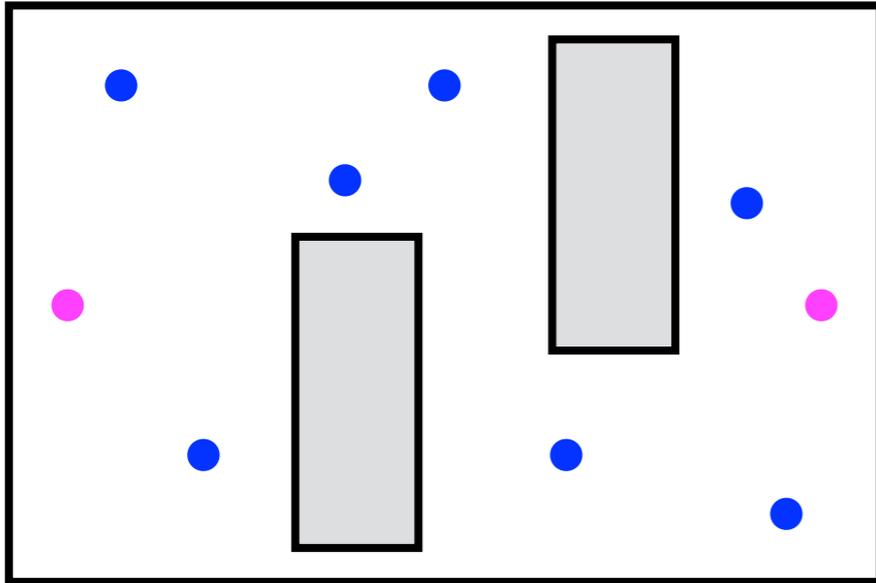
Strategy 2: Uniformly randomly sample



If C-space is a real vector space
for each dimension i
sample $q(i) \sim [lb, ub]$

What are the pros of random sampling? Cons?

Strategy 2: Uniformly randomly sample



If C-space is a real vector space
for each dimension i
sample $q(i) \sim [lb, ub]$

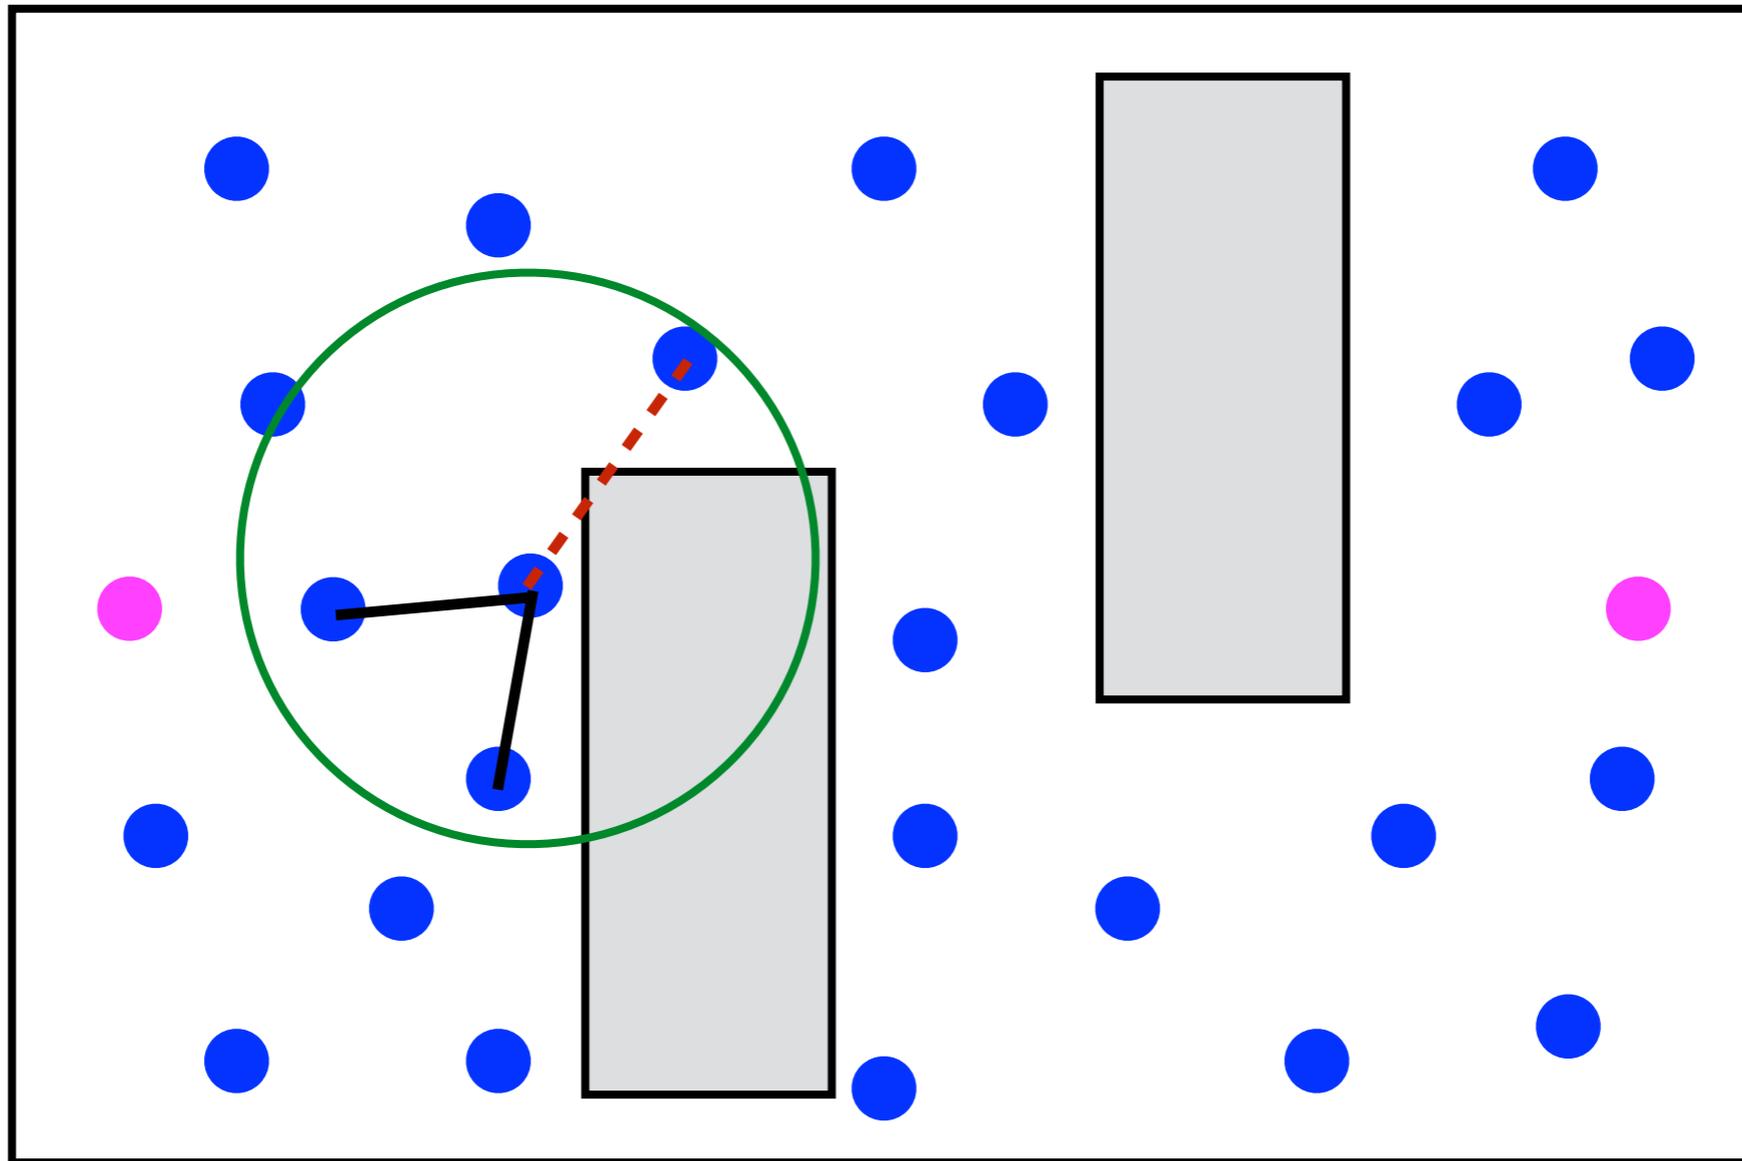
What are the pros of random sampling? Cons?

Question:

How do we decide which vertices to connect?

Strategy 2: Uniformly randomly sample

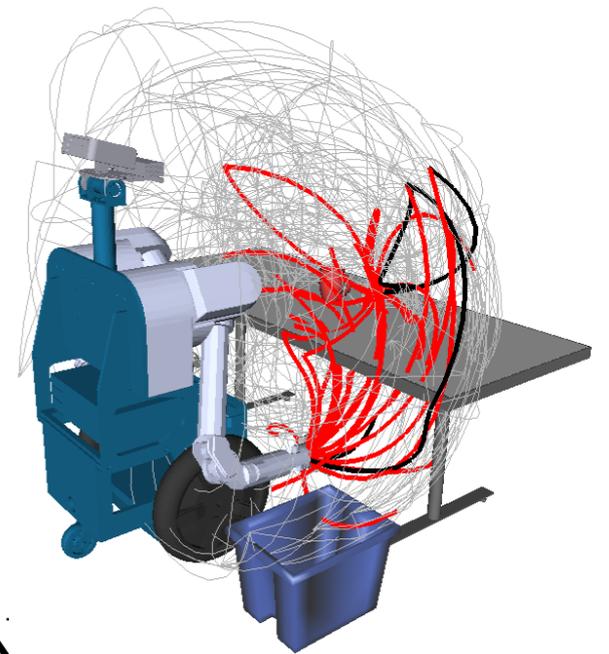
Connect vertices that are within a radius
(Alternatively can connect k-nearest neighbors)



This is the PRM Algorithm!

PRM = Probabilistic Roadmap

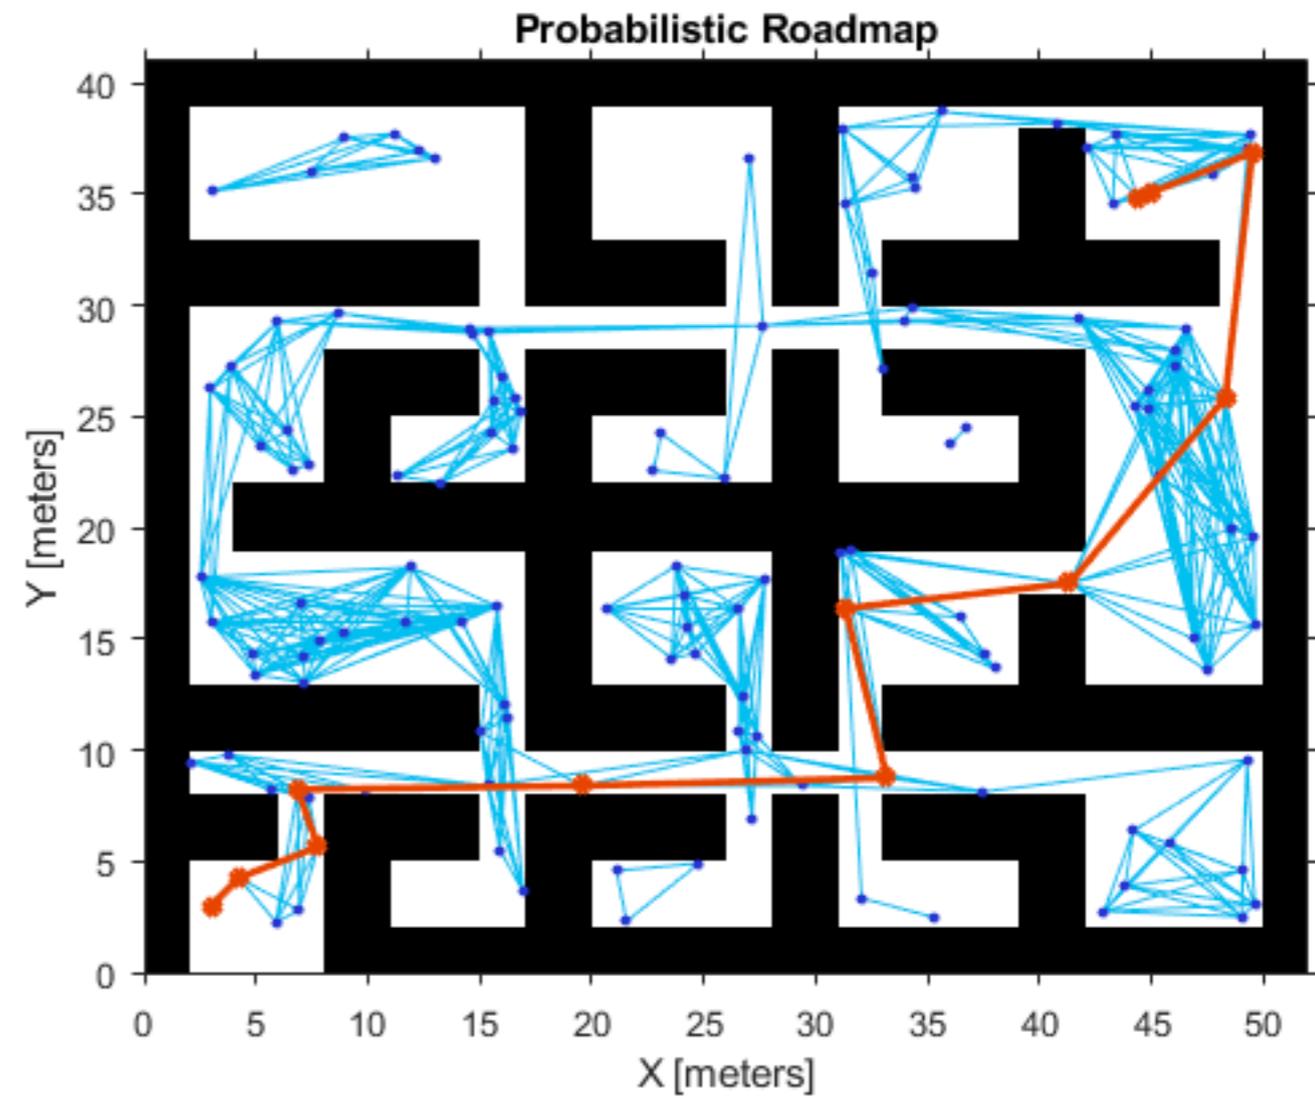
1. Sample vertices randomly
2. Connect vertices within radius (or k neighbors)
3. Search graph to find a solution



Theoretical Guarantees: It depends ...

L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics & Automation*, 12(4):566–580, June 1996.

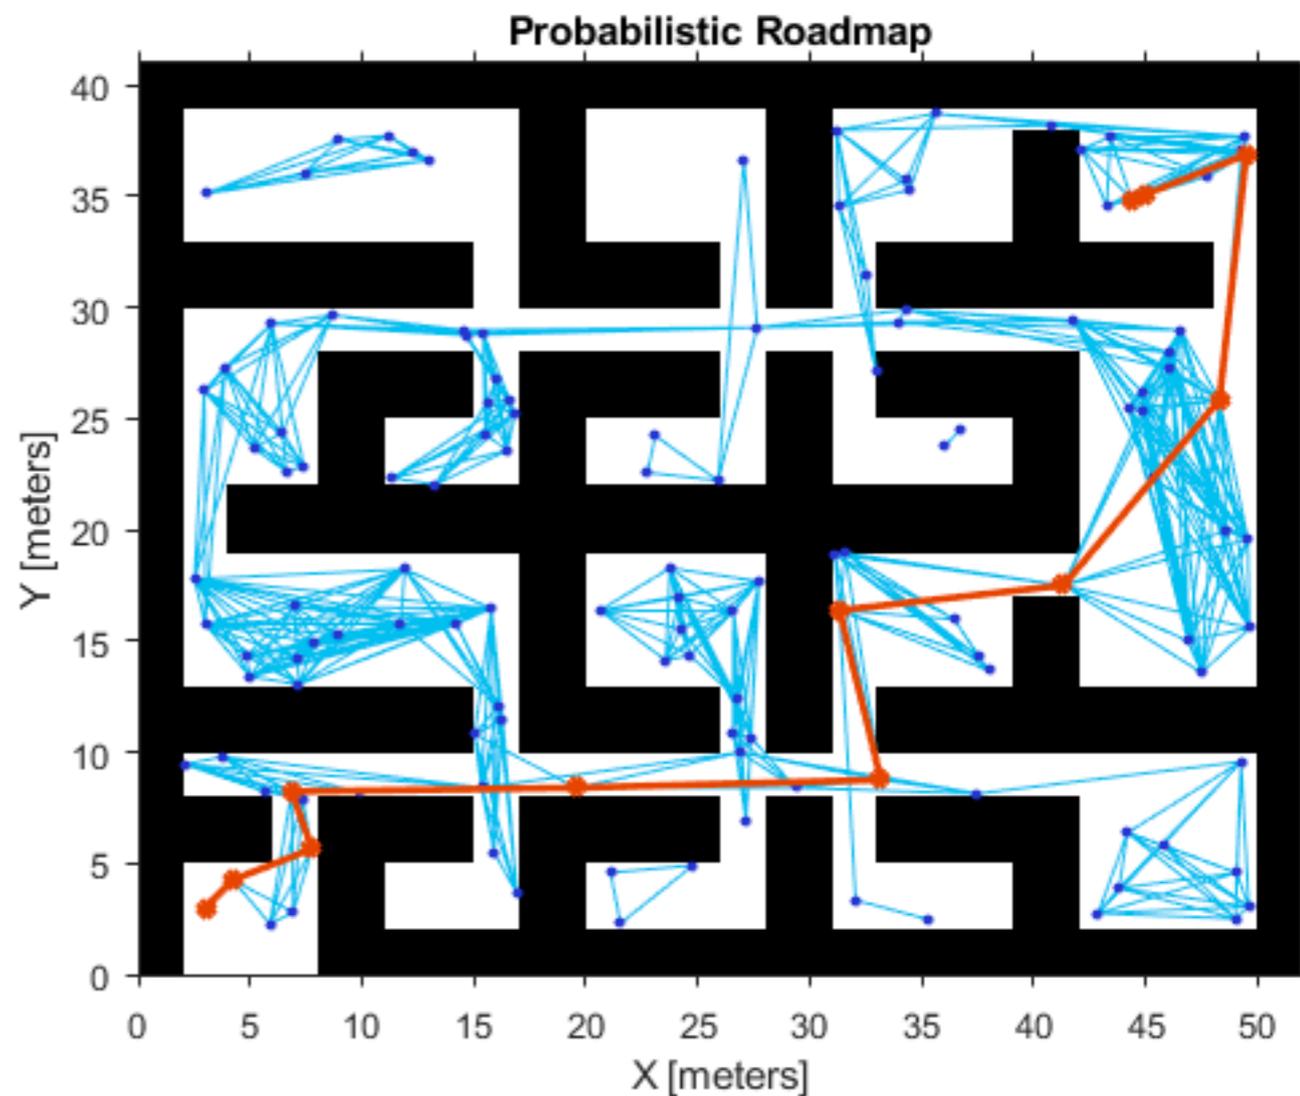
Questions we can ask PRM



Questions we can ask PRM

1. When is it a good idea to collision check every single edge?

Ans: Multi-query!



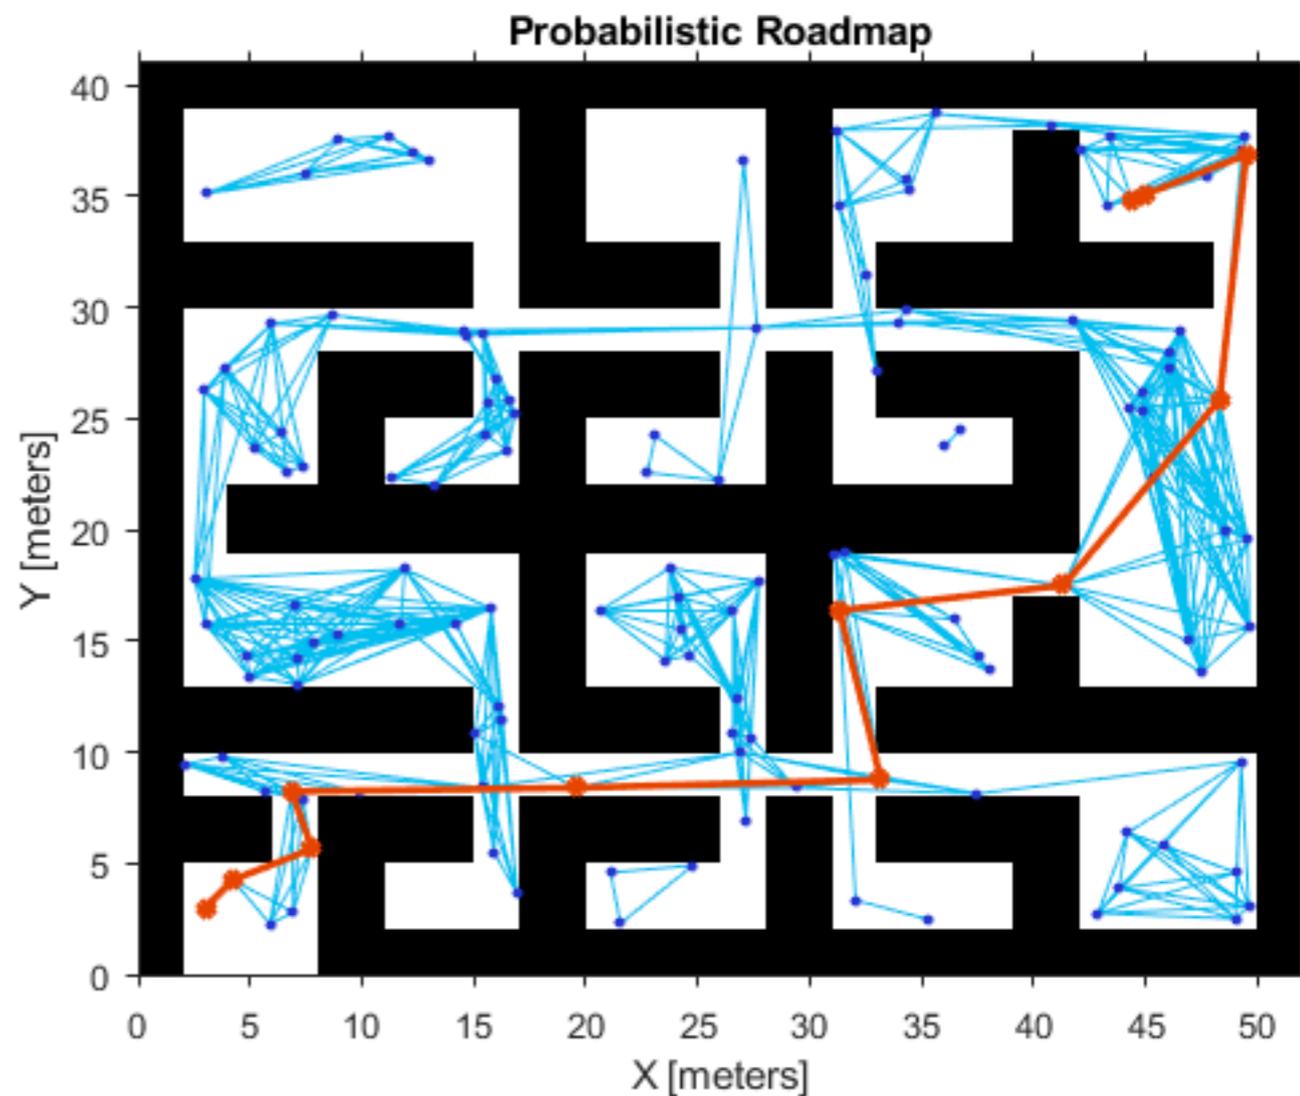
Questions we can ask PRM

1. When is it a good idea to collision check every single edge?

Ans: Multi-query!

2. How should we efficiently find nearest neighbors?

Ans: Use a KD-Tree data-structure



Questions we can ask PRM

1. When is it a good idea to collision check every single edge?

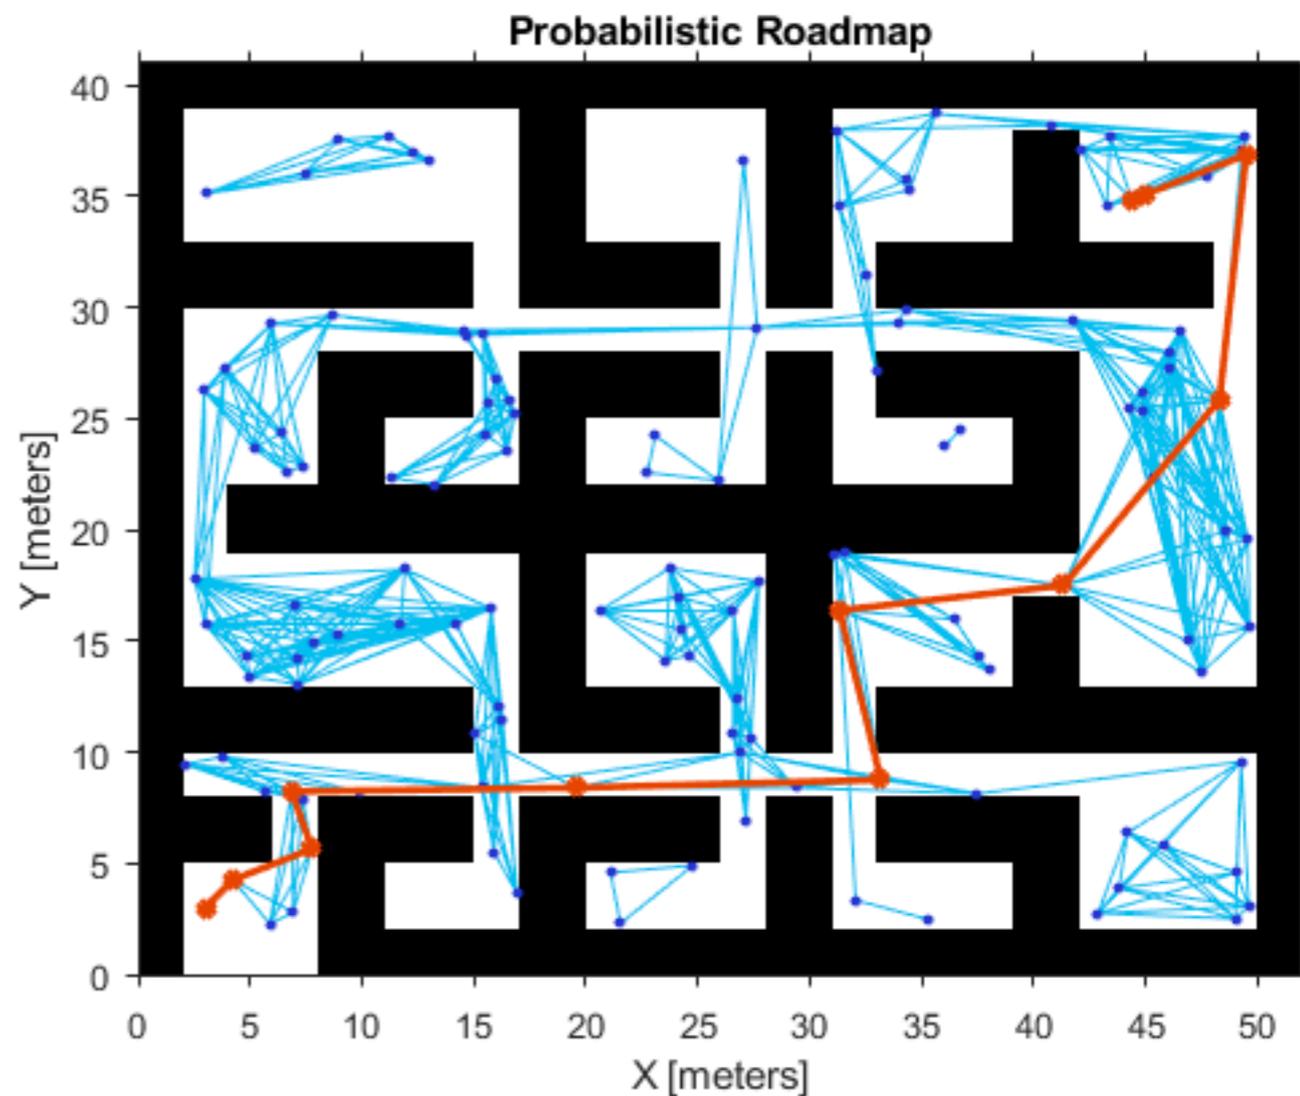
Ans: Multi-query!

2. How should we efficiently find nearest neighbors?

Ans: Use a KD-Tree data-structure

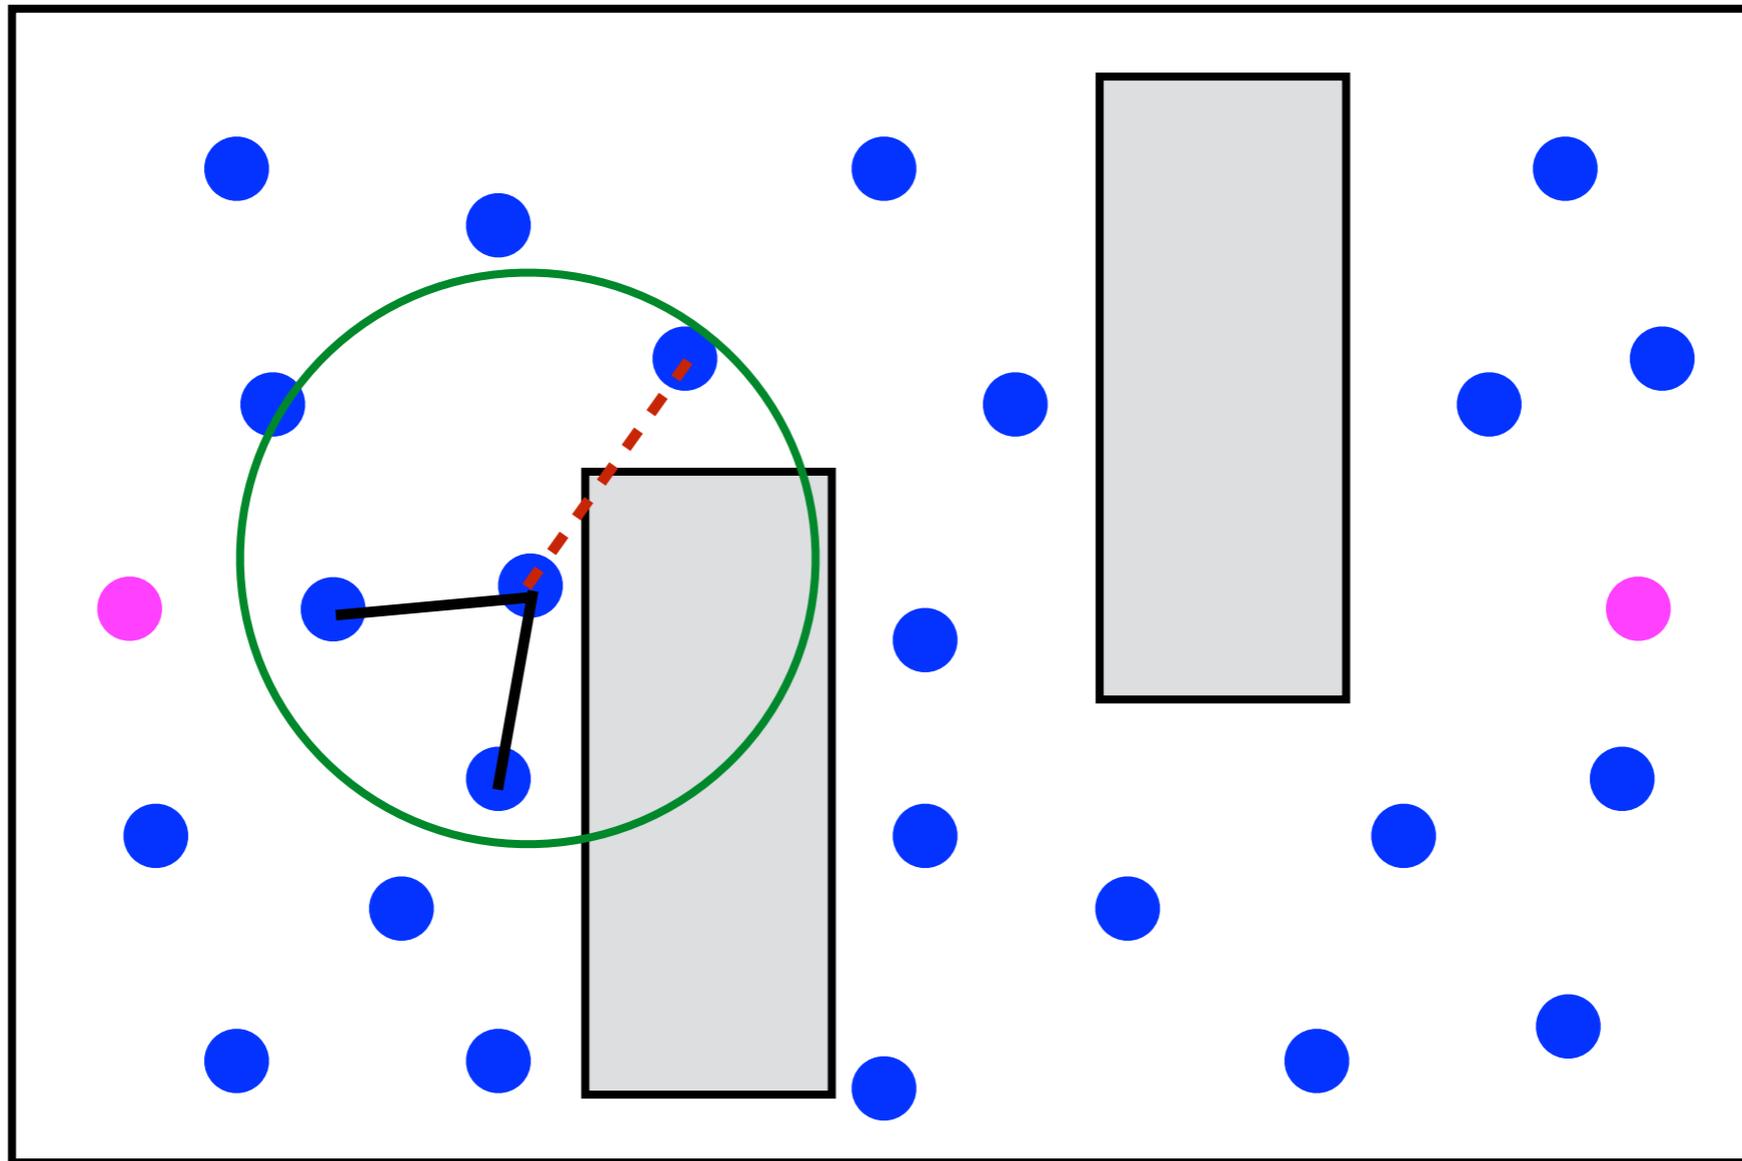
3. How should we choose which vertices to connect?

Ans: Up Next!



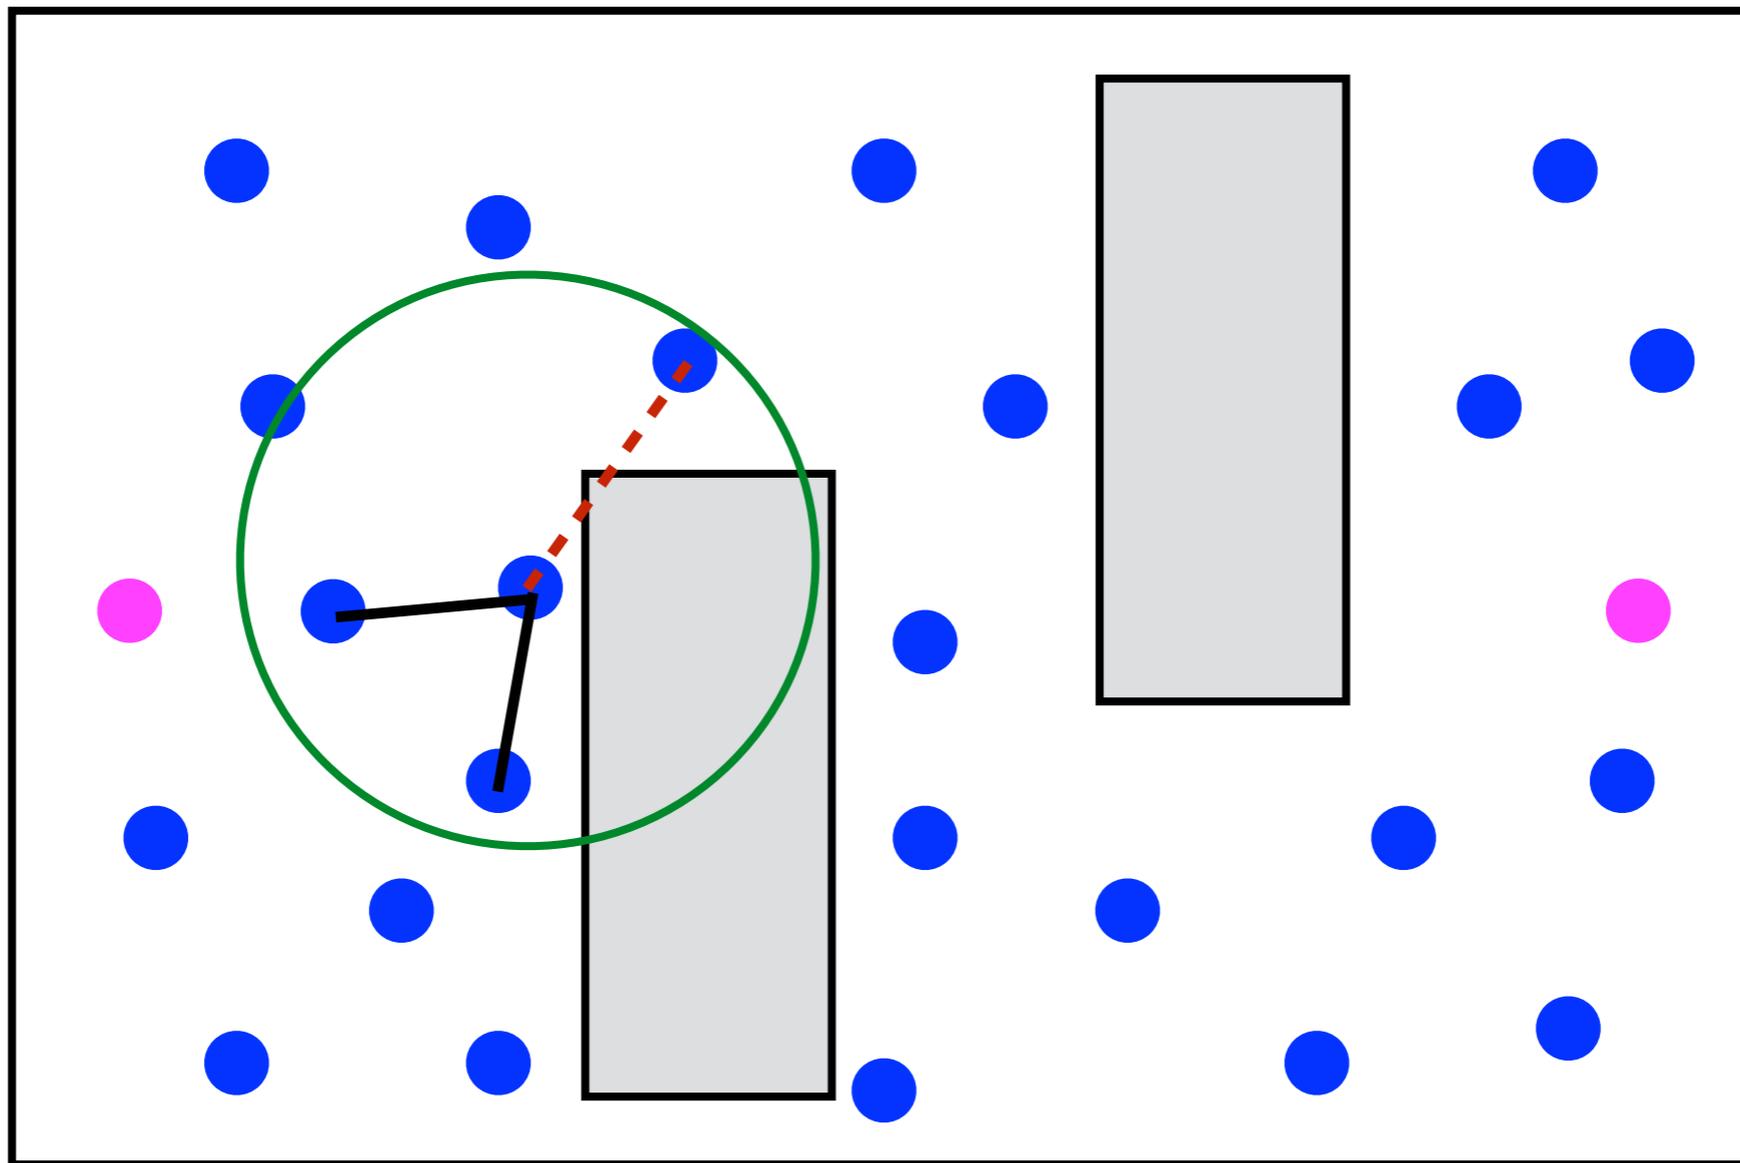
What is the optimal radius?

What happens if radius too large? too small?



What is the optimal radius?

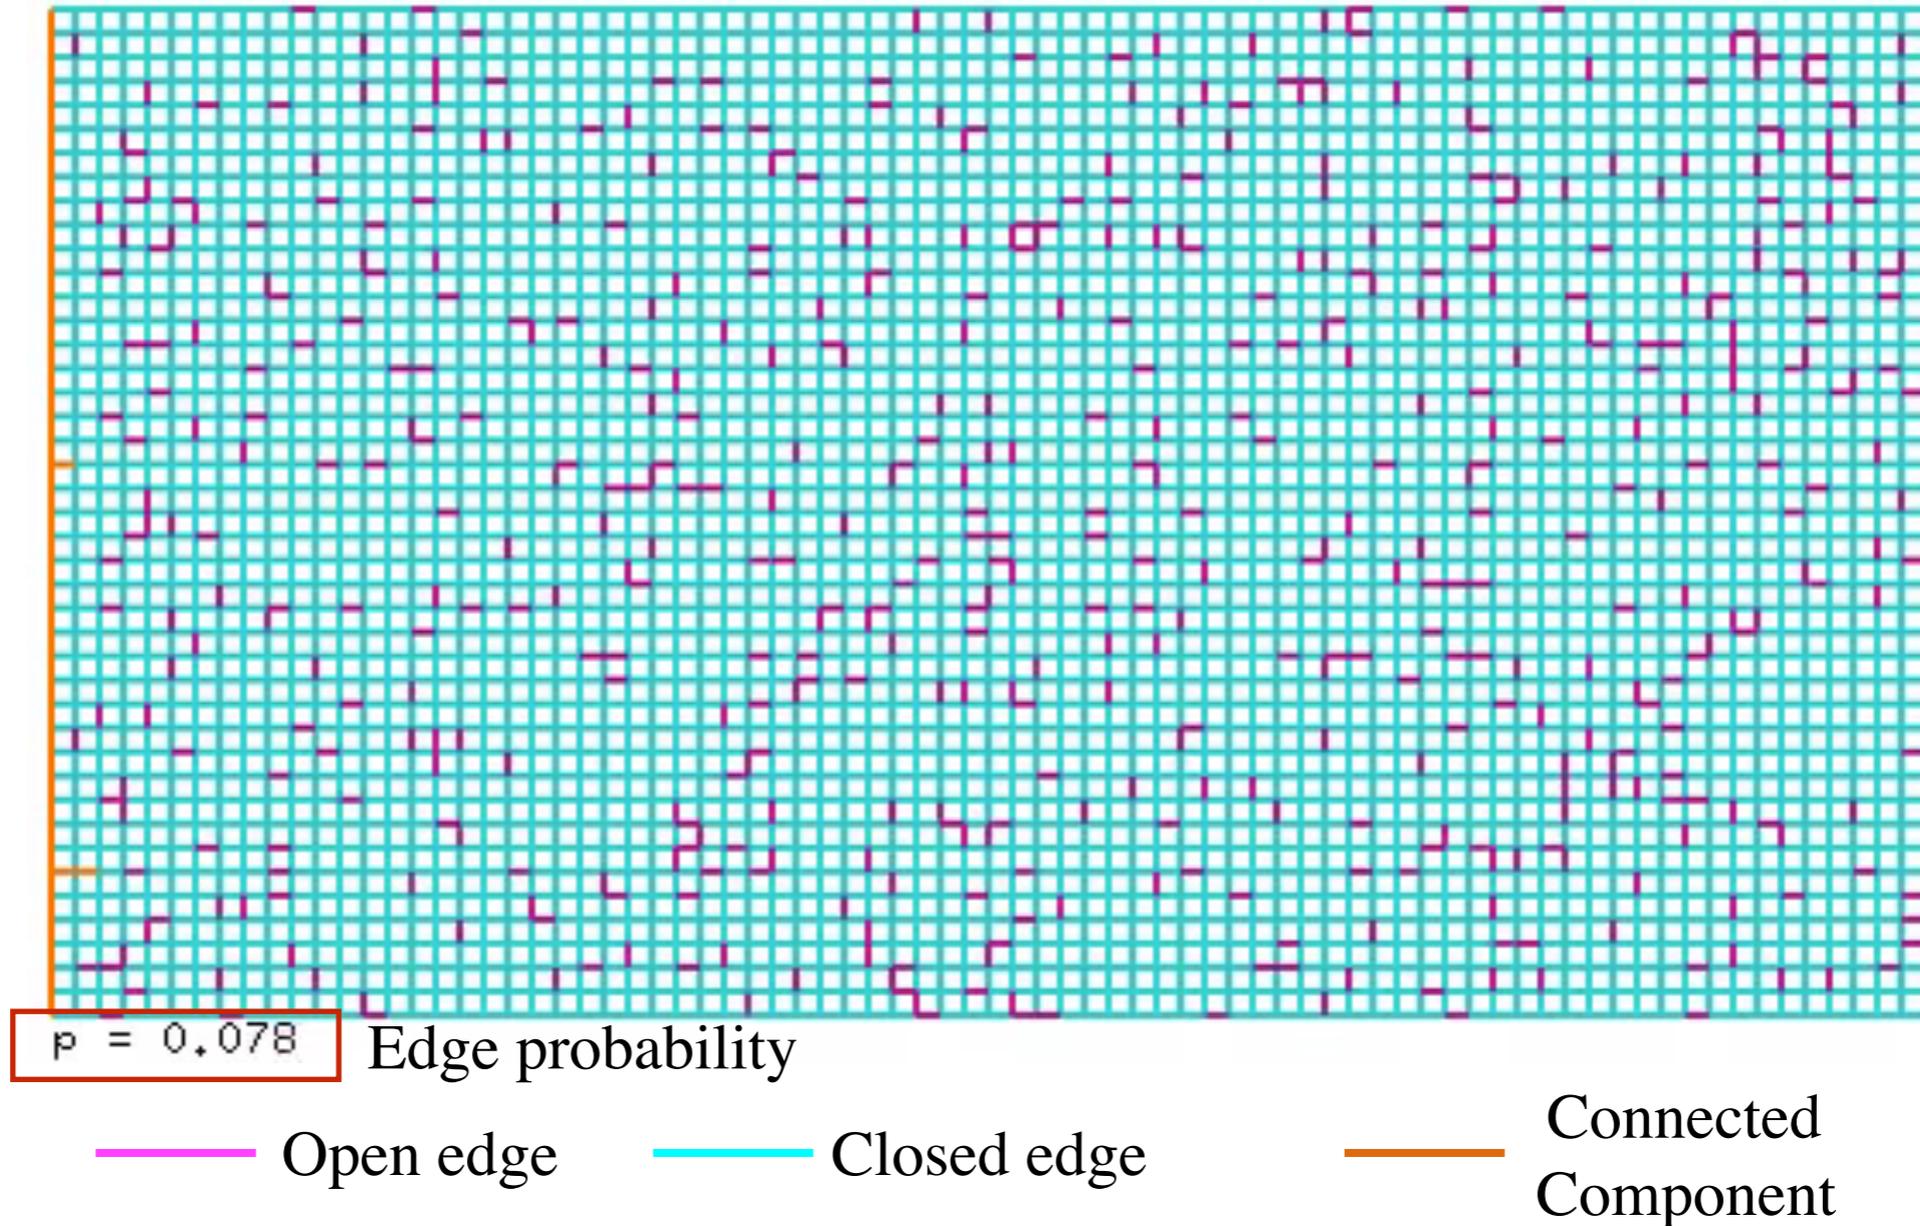
Set the radius to $r = \gamma \left(\frac{\log |V|}{|V|} \right)^{1/d}$ where magic constant!
 $\gamma \geq 2(1 + 1/d)^{1/d} \frac{\mu(\mathcal{C}_{free})}{\zeta_d}$



Also known as a Random Geometric Graph (RGG)

Aside: Percolation theory

<http://www.univ-orleans.fr/mapmo/membres/berglund/ressim.html>

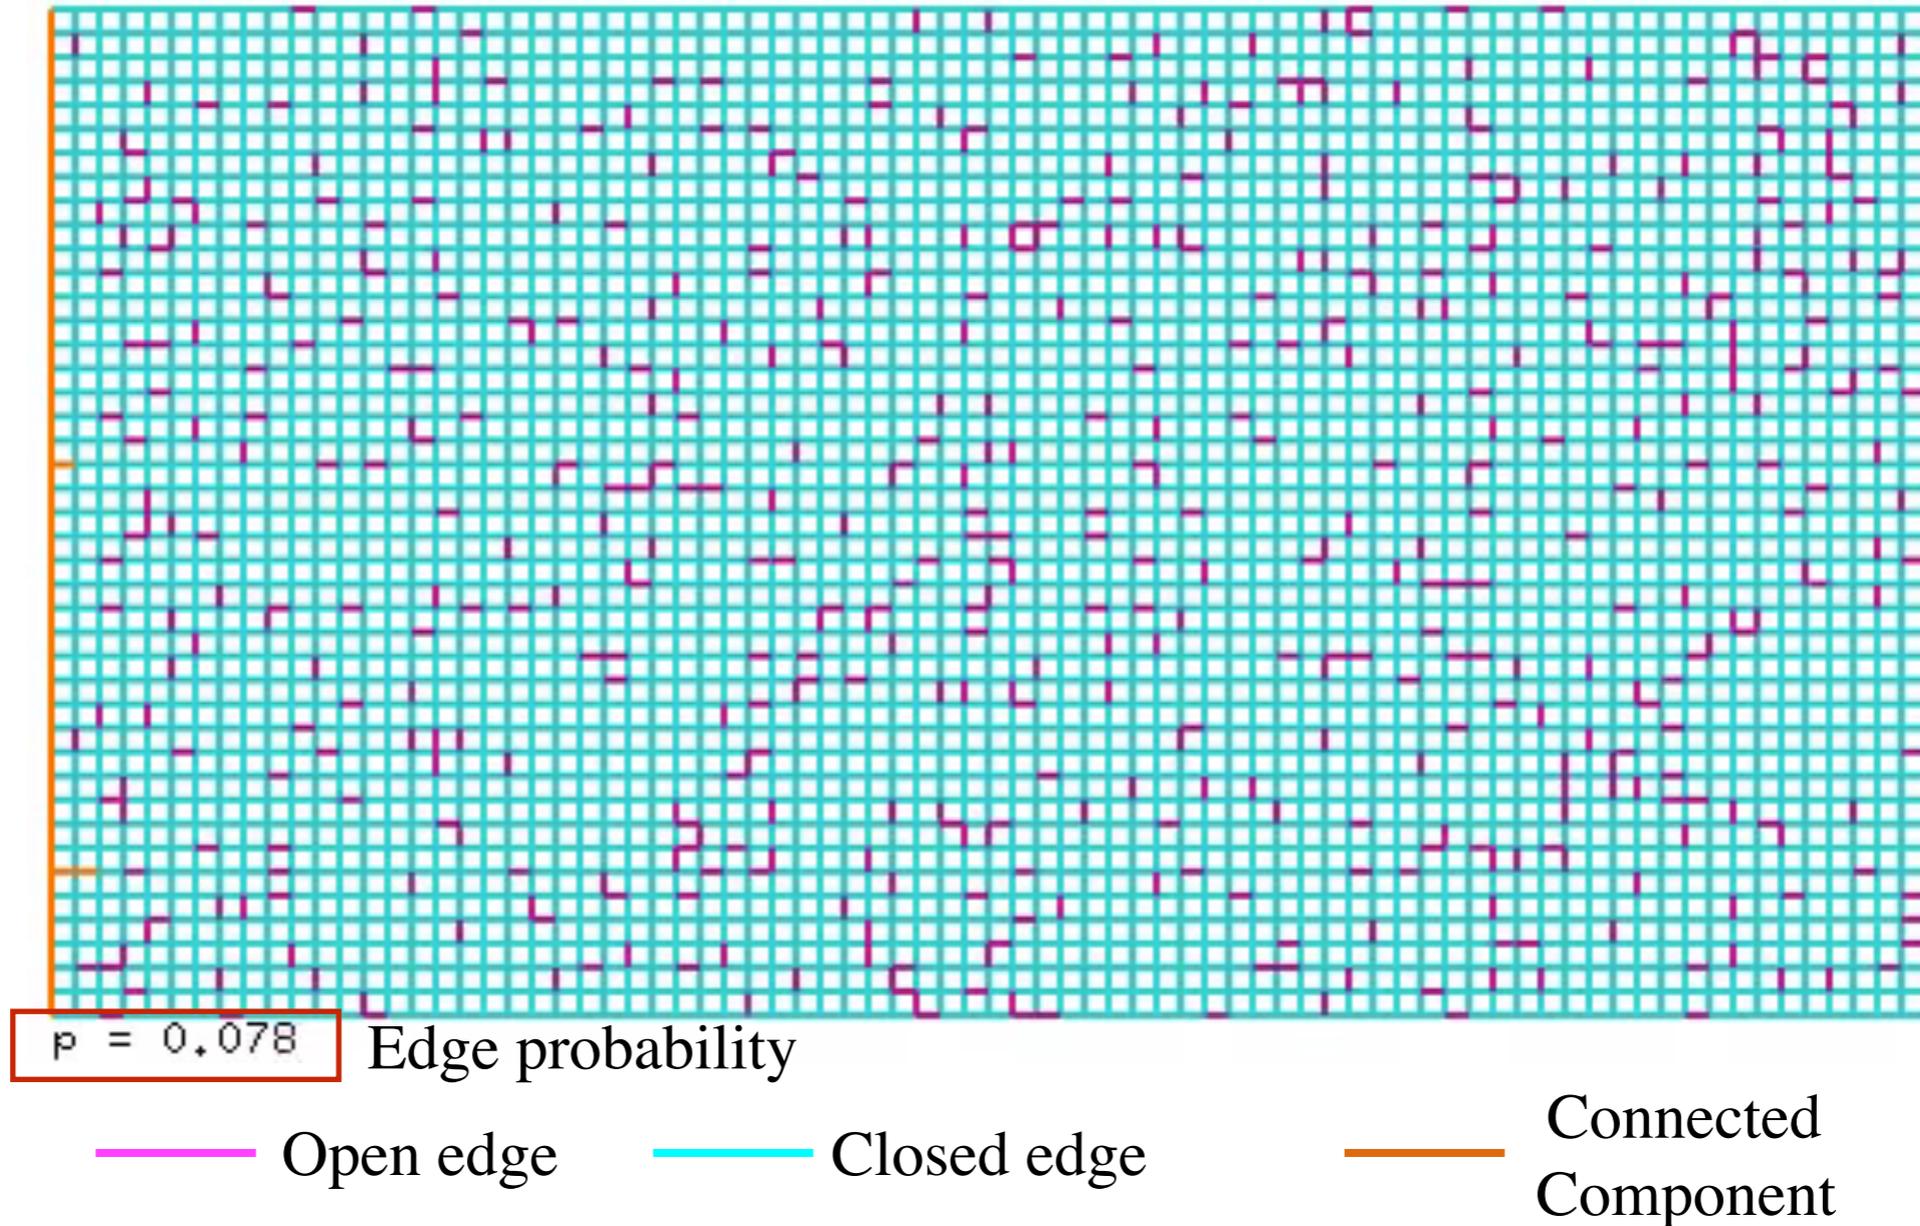


Other uses of percolation theory in planning

Theoretical Limits of Speed and Resolution for Kinodynamic Planning in a Poisson Forest
Sanjiban Choudhury, Sebastian Scherer and J. Andrew (Drew) Bagnell

Aside: Percolation theory

<http://www.univ-orleans.fr/mapmo/membres/berglund/ressim.html>



Other uses of percolation theory in planning

Theoretical Limits of Speed and Resolution for Kinodynamic Planning in a Poisson Forest
Sanjiban Choudhury, Sebastian Scherer and J. Andrew (Drew) Bagnell

This is the PRM* Algorithm!

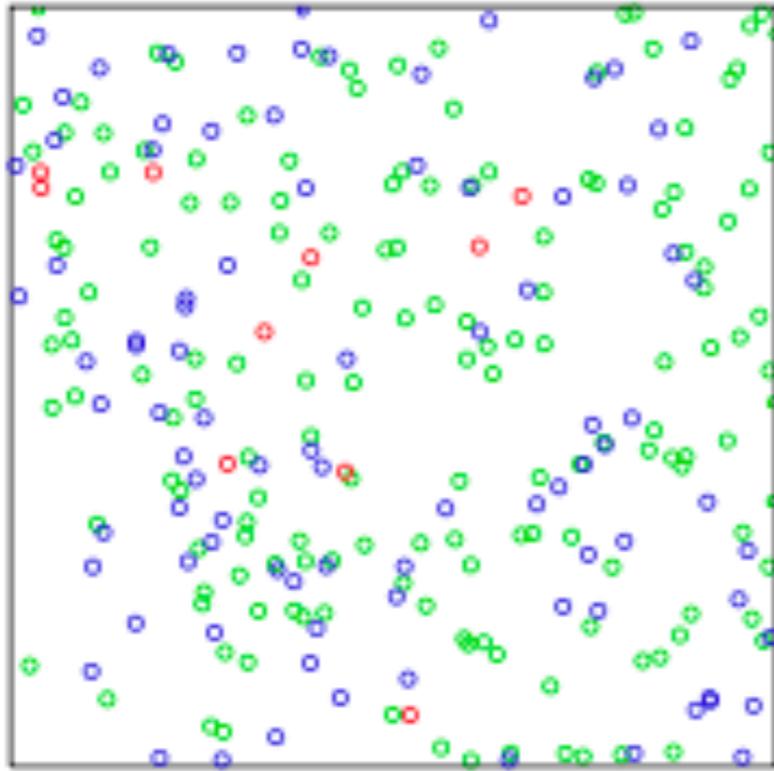
1. Sample vertices randomly
2. Use **optimal radius formula** to connect vertices
3. Search graph to find a solution

Theorem: Probabilistically complete AND Asymptotically optimal

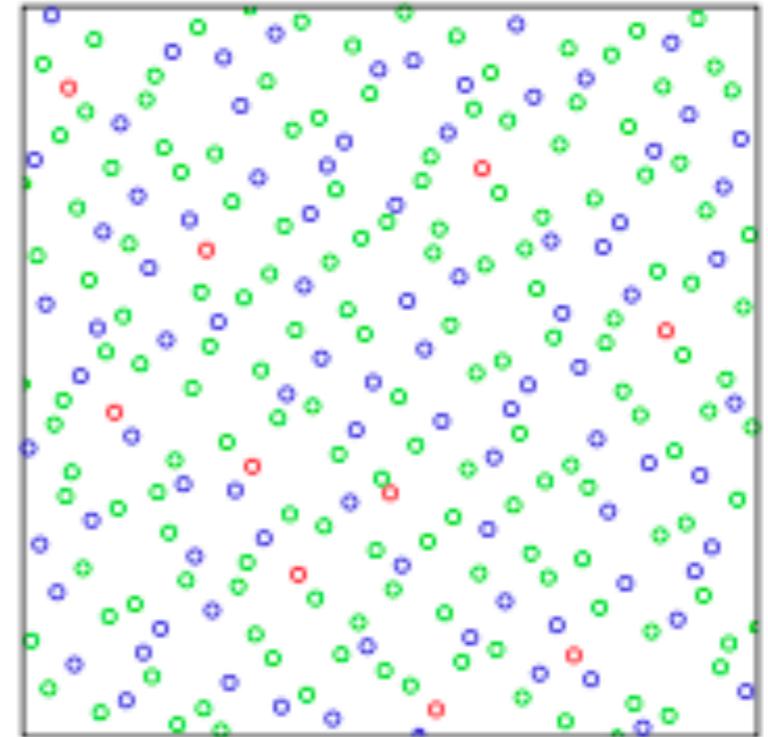
“Sampling-based Algorithms for Optimal Motion Planning”

Sertac Karaman and Emilio Frazzoli, IJRR 2011

Can we do better than random?



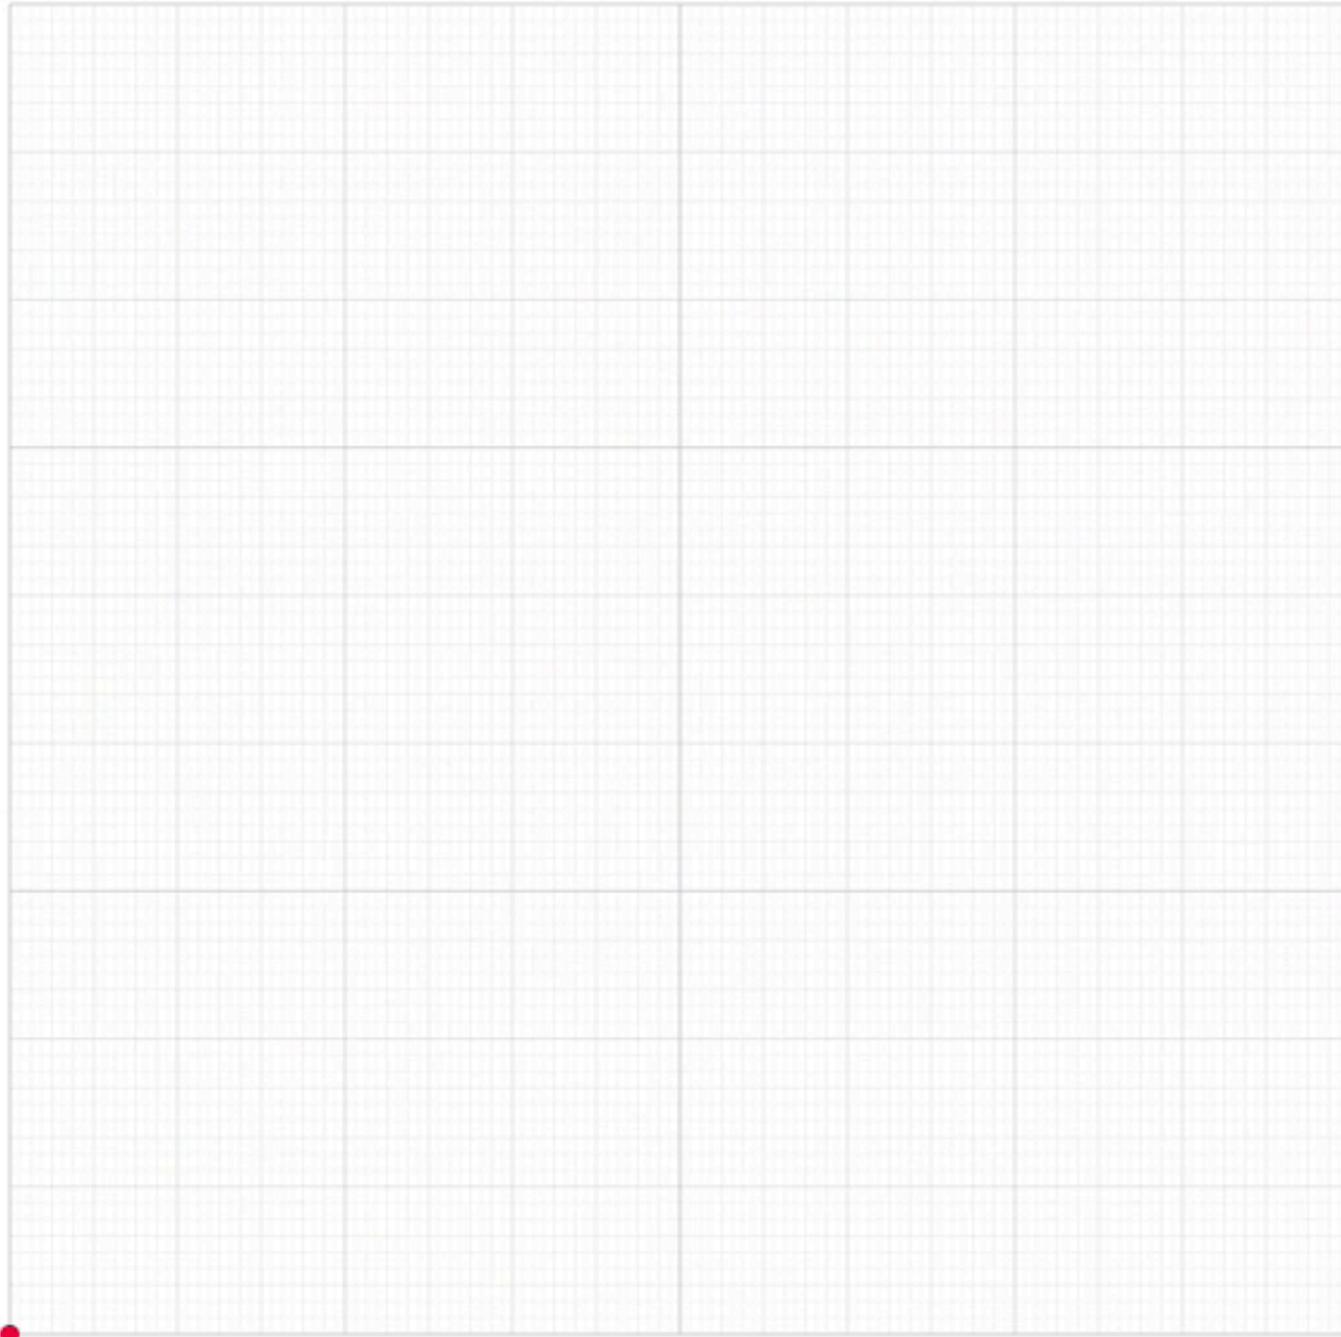
Uniform random
sampling tends to
clump



Ideally we would
want points to be
spread out evenly

Question: How do we do this without discretization?

Halton Sequence

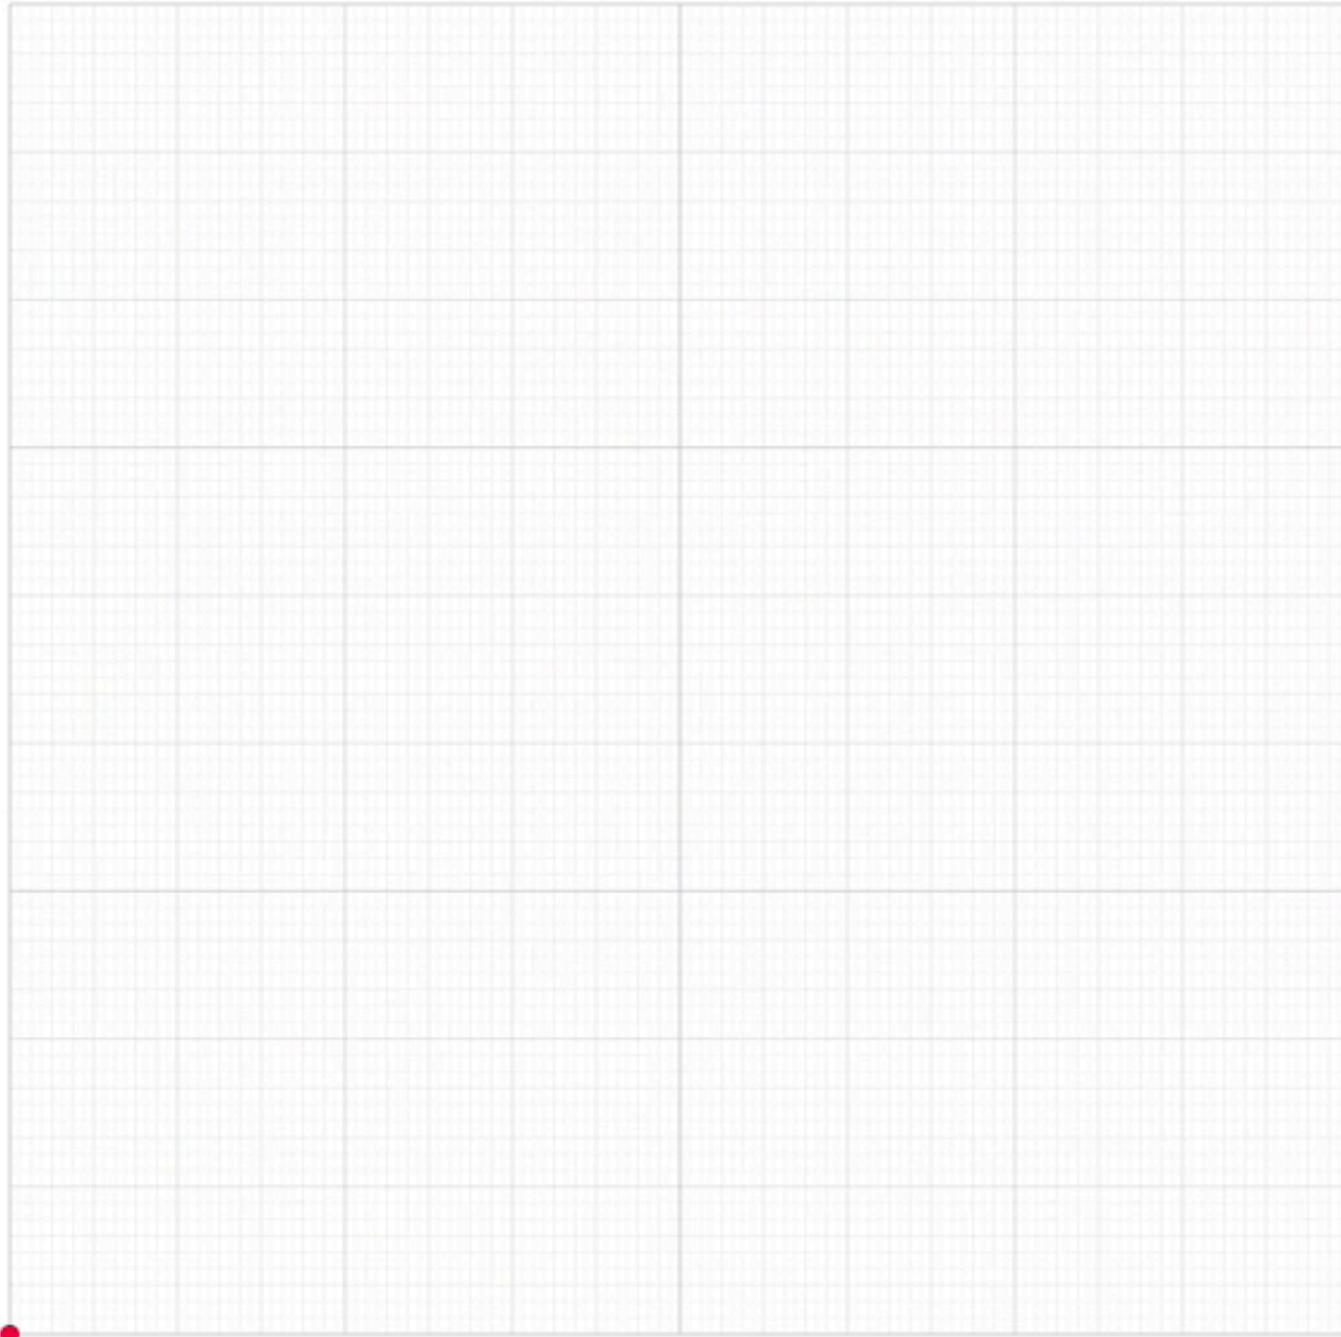


Generalization of
Van de Coruput Sequence

Intuition: Create a sequence
using prime numbers that
uniformly densify space

Link for exact algorithm:
<https://observablehq.com/@jrus/halton>

Halton Sequence

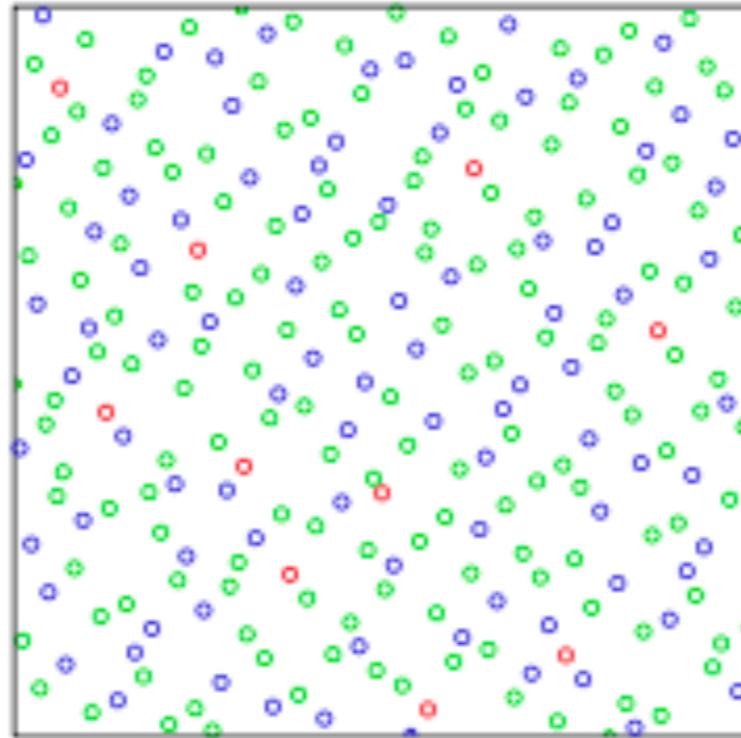


Generalization of
Van de Coruput Sequence

Intuition: Create a sequence
using prime numbers that
uniformly densify space

Link for exact algorithm:
<https://observablehq.com/@jrus/halton>

How do we connect vertices?



Halton sequences have much better coverage
(i.e. they are low dispersion)

Connect vertices that are within a radius of

$$r = \gamma \left(\frac{1}{|V|} \right)^{1/d} \quad (\text{as opposed to:})$$
$$r = \gamma \left(\frac{\log |V|}{|V|} \right)^{1/d}$$

This is the gPRM Algorithm!

1. Sample vertices randomly
2. Use **optimal radius formula** to connect vertices
3. Search graph to find a solution

Theorem: Probabilistically complete AND Asymptotically optimal
AND Asymptotic rate of convergence

“Deterministic Sampling-Based Motion Planning: Optimality, Complexity, and Performance”

Lucas Janson, Brian Ichter, Marco Pavone, IJRR 2017

What makes a good graph?

What makes a good graph?

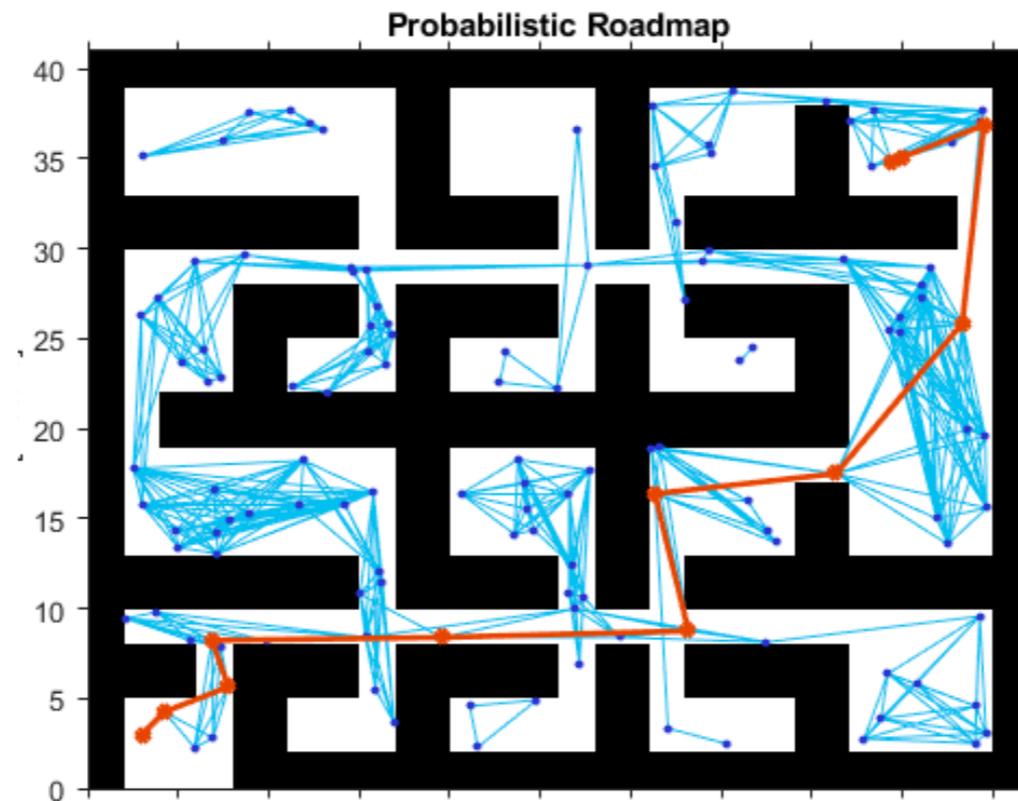
1. A good graph must be **sparse** (both in vertices and edges)

What makes a good graph?

1. A good graph must be **sparse** (both in vertices and edges)
 2. A good graph must have **good coverage**

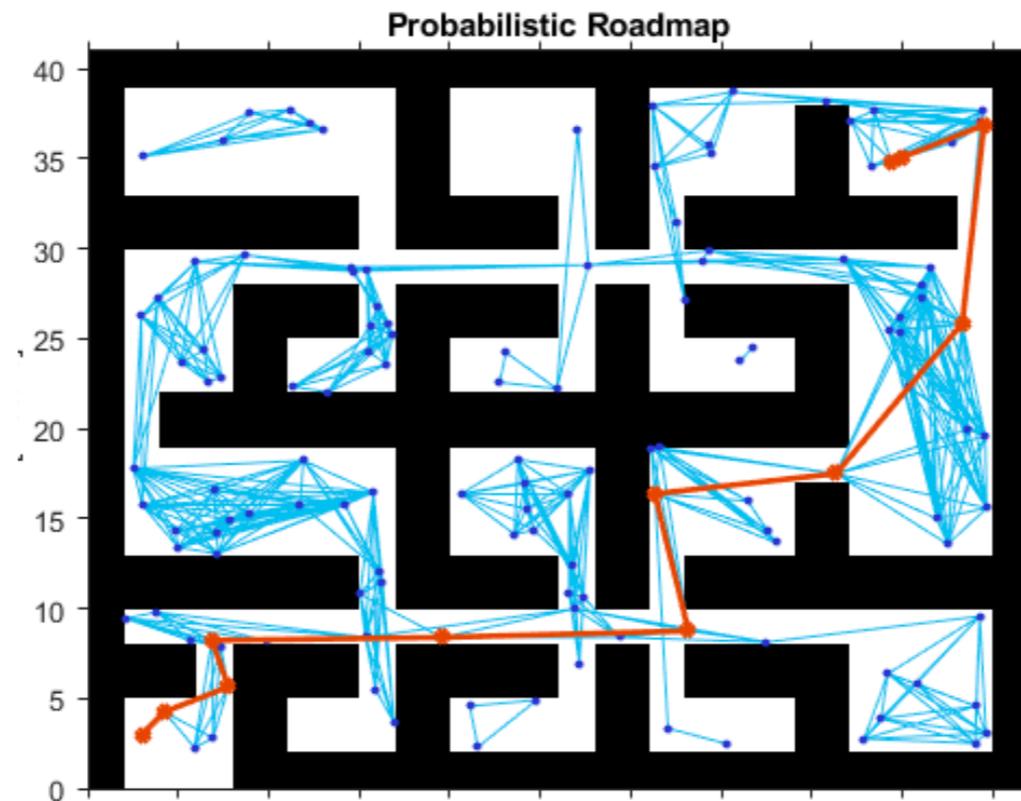
What makes a good graph?

1. A good graph must be **sparse** (both in vertices and edges)
2. A good graph must have **good coverage**



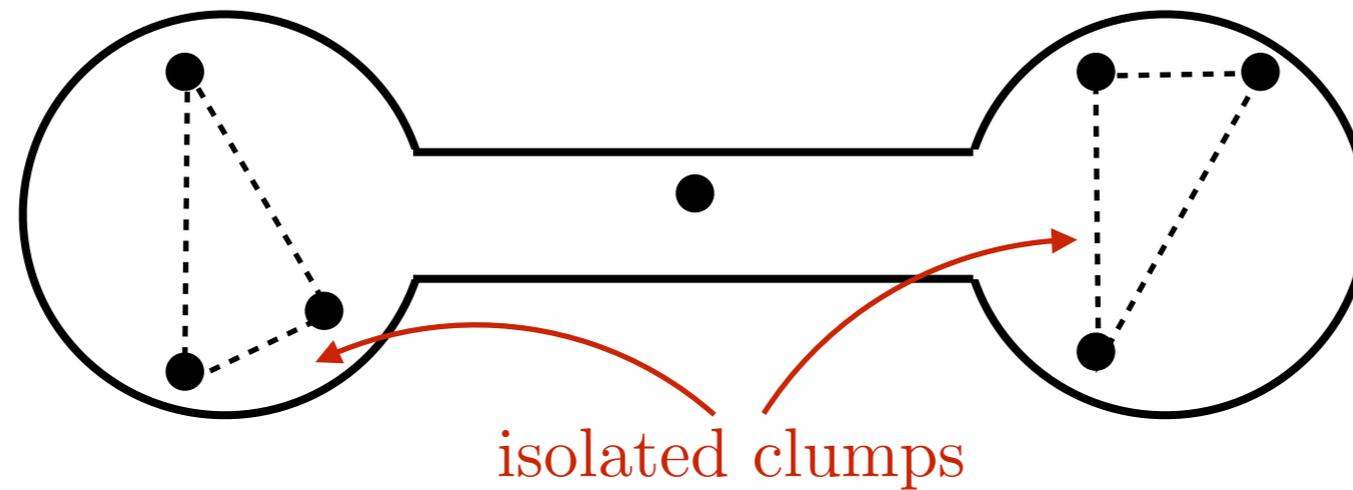
What makes a good graph?

1. A good graph must be **sparse** (both in vertices and edges)
2. A good graph must have **good coverage**



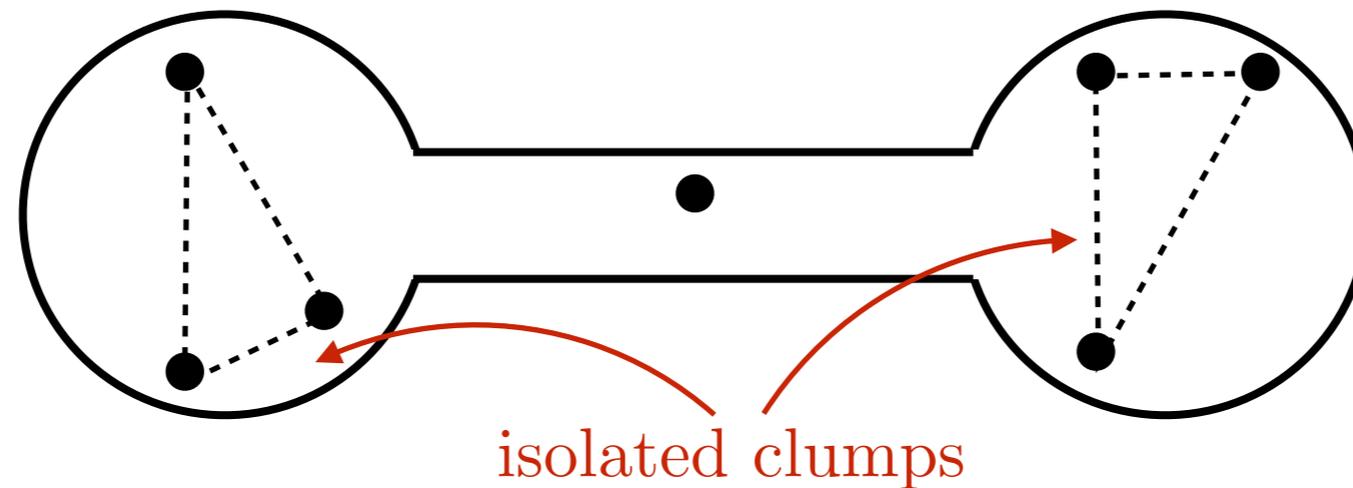
3. A good graph must have the **same connectivity** of free space

The Narrow Passage: Planning's boogie man!



Why is narrow passage mathematically hard to plan in?

The Narrow Passage: Planning's boogie man!



Why is narrow passage mathematically hard to plan in?

Mathematical Question:

How many samples do we need to connect the space
(with high probability)?

How many samples do we need?

Theorem [Hsu et al., 1999] Let $2n$ vertices be sampled from X_{free} . Then the roadmap is connected with probability at least $1 - \gamma$ if:

$$n \geq \left\lceil 8 \frac{\log\left(\frac{8}{\epsilon \alpha \gamma}\right)}{\epsilon \alpha} + \frac{3}{\beta} + 2 \right\rceil$$

The shape of free C-space is dictated by α , β , $\epsilon \in [0, 1]$

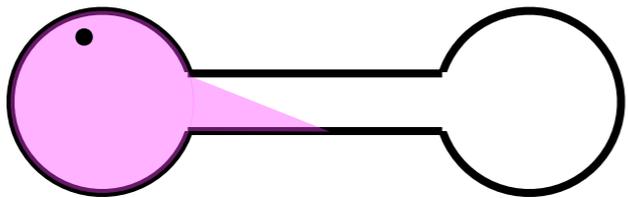
How many samples do we need?

Theorem [Hsu et al., 1999] Let $2n$ vertices be sampled from X_{free} . Then the roadmap is connected with probability at least $1 - \gamma$ if:

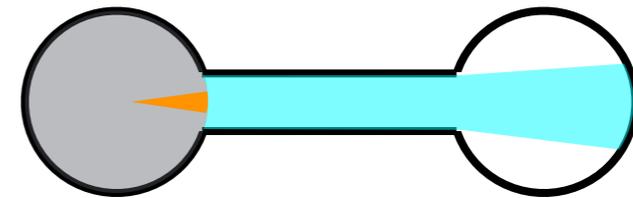
$$n \geq \left\lceil 8 \frac{\log\left(\frac{8}{\epsilon \alpha \gamma}\right)}{\epsilon \alpha} + \frac{3}{\beta} + 2 \right\rceil$$

The shape of free C-space is dictated by α , β , $\epsilon \in [0, 1]$

Visibility of free space (ϵ)



Expansion of visibility (α , β)

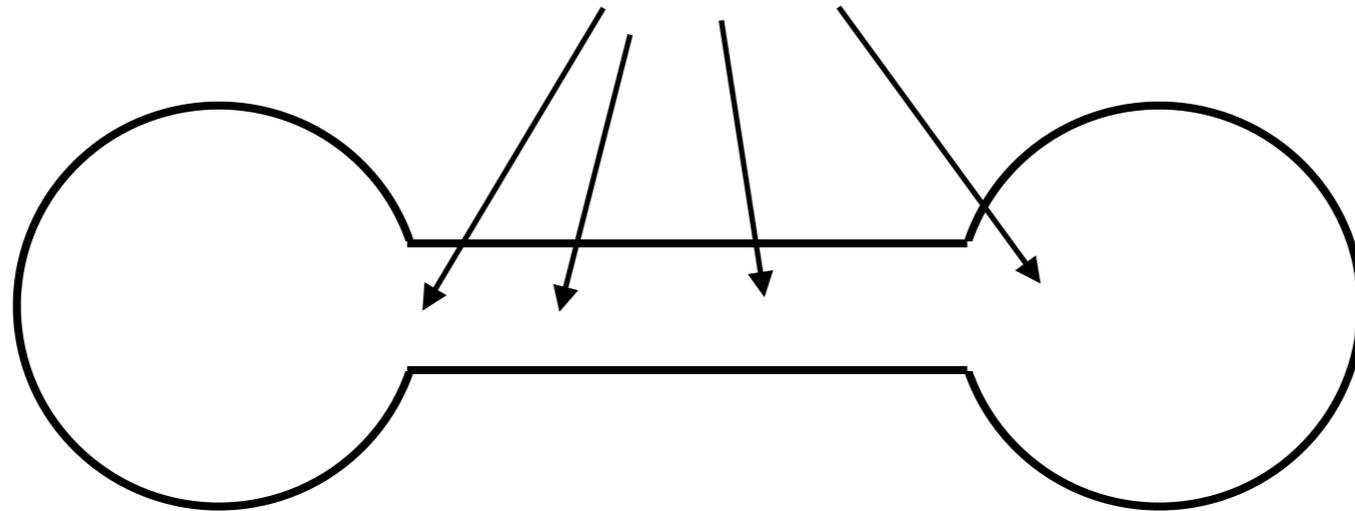


Narrow passage has small values of α , β , ϵ

Hence, **needs more samples** to find a path

How do we bias sampling?

We somehow need more samples here



1. Sample near obstacle surface?

V. Boor, M. H. Overmars, and A. F. van der Stappen. The Gaussian sampling strategy for probabilistic roadmap planners. 1999

2. Add samples that are in between two obstacles?

D. Hsu, T. Jiang, J. Reif, and Z. Sun. The bridge test for sampling narrow passages with probabilistic roadmap planners. 2003.

3. Train a learner to detect the narrow passages?

B. Ichter, J. Harrison, M. Pavone. Learning Sampling Distributions for Robot Motion Planning, 2018

Summary of ways to create graphs

	How to sample vertices?	How to connect vertices?
Lattice	Discretize	connectivity rule
PRM	Uniform random	r-disc, k-nn
PRM*	Uniform random	optimal r-disc, k-nn
gPRM	Halton sequence	optimal r-disc, k-nn
Bridge	Sample with bridge test	any visible points
Gaussian	Sample near obstacles	r-disc, k-nn
MAPRM	Sample along medial axis	r-disc, k-nn
Approx. Visibility Graph	Sample on surface of obstacles	any visible points
Learnt Sampler	Use CVAE to approximate free space	optimal r-disc, k-nn