

Motion planning: sampling-based planners IV implementation details

Vojtěch Vonásek

Department of Cybernetics
Faculty of Electrical Engineering
Czech Technical University in Prague

- Sampling-based planning relies on low-level routines
- Efficient implementation of these routines is necessary




- Random numbers generator
- Metric
- Nearest-neighbor search
- Collision-detection

```
1 initialize tree  $\mathcal{T}$  with  $q_{init}$ 
2 for  $i = 1, \dots, l_{max}$  do
3    $q_{rand} = \text{generate randomly in } \mathcal{C}$ 
4    $q_{near} = \text{find nearest node in } \mathcal{T} \text{ towards}$ 
    $q_{rand}$ 
5    $q_{new} = \text{localPlanner from } q_{near} \text{ towards}$ 
    $q_{rand}$ 
6   if  $\text{canConnect}(q_{near}, q_{new})$  then
7      $\mathcal{T}.\text{addNode}(q_{new})$ 
8      $\mathcal{T}.\text{addEdge}(q_{near}, q_{new})$ 
9     if  $\varrho(q_{new}, q_{goal}) < d_{goal}$  then
10      return path from  $q_{init}$  to  $q_{new}$ 
```

- These routines are required also in PRM, EST and all their variants


Generation using standard `rand()`

- Many variants of random number generator (RNG)
- RNG is (usually) implemented as a Linear Congruent Generator (LCG)
- Fast, easy for usage, provides “enough” number of samples
- High dispersion (the largest empty ball according to the used metric)



Generation using standard `rand()`


- Many variants of random number generator (RNG)
- RNG is (usually) implemented as a Linear Congruent Generator (LCG)
- Fast, easy for usage, provides “enough” number of samples
- High dispersion (the largest empty ball according to the used metric)



Random samples in 2D \mathcal{C}

Generation using standard `rand()`


- Many variants of random number generator (RNG)
- RNG is (usually) implemented as a Linear Congruent Generator (LCG)
- Fast, easy for usage, provides “enough” number of samples
- High dispersion (the largest empty ball according to the used metric)



Random samples in 2D \mathcal{C} + dispersion

Generation using standard `rand()`


- Many variants of random number generator (RNG)
- RNG is (usually) implemented as a Linear Congruent Generator (LCG)
- Fast, easy for usage, provides “enough” number of samples
- High dispersion (the largest empty ball according to the used metric)



Voronoi diagram

Generation using standard `rand()`


- Many variants of random number generator (RNG)
- RNG is (usually) implemented as a Linear Congruent Generator (LCG)
- Fast, easy for usage, provides “enough” number of samples
- High dispersion (the largest empty ball according to the used metric)



PRM roadmap, note the “holes”

Generation using standard `rand()`

- Many variants of random number generator (RNG)
- RNG is (usually) implemented as a Linear Congruent Generator (LCG)
- Fast, easy for usage, provides “enough” number of samples
- High dispersion (the largest empty ball according to the used metric)



PRM roadmap, note the “holes” → due to the dispersion

Alternatives to `rand()`


- Many libraries provide various RNG
 - e.g., Boost, GSL, numpy
- GSL — GNU Scientific library, offers tens of random generators
- Most of them are based on LCG

Does RNG influence the performance of sampling-based planners?

- Test scenario: square robot, 3D \mathcal{C} -space, narrow passage




Generating random samples




- What can you conclude from this measurement?

Using low-discrepancy sequences

- Halton/Hammersley deterministic sequences
- Number of samples must be known in advance
- Slower computation in comparison to basic `rand()`
- Lower dispersion than LCG-based `rand()`



Halton points



Hammersley points

• J. M. Hammersley. Monte-Carlo methods for solving multivariable problems. Annals of the New York Academy of Science, 86:844–874, 1960.

- Random numbers generator
- Metric
- Nearest-neighbor search
- Collision-detection

```
1 initialize tree  $\mathcal{T}$  with  $q_{init}$ 
2 for  $i = 1, \dots, l_{max}$  do
3    $q_{rand} = \text{generate randomly in } \mathcal{C}$ 
4    $q_{near} = \text{find nearest node in } \mathcal{T} \text{ towards}$ 
      $q_{rand}$ 
5    $q_{new} = \text{localPlanner from } q_{near} \text{ towards}$ 
      $q_{rand}$ 
6   if canConnect( $q_{near}, q_{new}$ ) then
7      $\mathcal{T}.\text{addNode}(q_{new})$ 
8      $\mathcal{T}.\text{addEdge}(q_{near}, q_{new})$ 
9     if  $\varrho(q_{new}, q_{goal}) < d_{goal}$  then
10      return path from  $q_{init}$  to  $q_{new}$ 
```

- Sampling-based planners require a metric $\varrho(q_1, q_2)$, $q_1, q_2 \in \mathcal{C}$
- Often used are L_p metrics:

$$\varrho(x, x') = \left(\sum_{i=1}^n |x_i - x'_i|^p \right)^{1/p}$$

- L_2 is Euclidean metric
- L_1 is Manhattan metric
- Metric for 1D rotation:

$$\varrho(\theta_1, \theta_2) = \min(|\theta_1 - \theta_2|, 2\pi - |\theta_1 - \theta_2|)$$

- Metrics can be combined, let's assume that $\mathcal{C} = X \times Y$ with ϱ_X and ϱ_Y :


$$\varrho(q, q') = \left(c_x \varrho_X(x, x')^p + c_y \varrho_Y(y, y')^p \right)^{1/p}$$

- where $c_x, c_y \geq 0$ are weights


Remind that for $a, b, c \in X$ in a metric space X and metric ϱ : $\varrho(a, b) \geq 0$; $\varrho(a, b) = 0$ if and only if $a = b$; $\varrho(a, b) = \varrho(b, a)$; $\varrho(a, b) + \varrho(b, c) \geq \varrho(a, c)$

- 2D object, translation + rotation $\rightarrow q = (x, y, \varphi) \in \mathcal{C}$

$$\varrho(q, q') = \sqrt{c_1((x - x')^2 + (y - y')^2) + c_2\varrho_\theta(\varphi, \varphi')}$$



- where $\varrho_\theta(\varphi, \varphi')$ is the metric for 1D rotation
- The weights for translation c_1 is “usually” bigger than c_2 , so the effect of the rotation is suppressed
- Wrong setting of weights can worsen motion planning



Correct




Big weight on y- distance

- Random numbers generator
- Metric
- Nearest-neighbor search
- Collision-detection


```
1 initialize tree  $\mathcal{T}$  with  $q_{init}$ 
2 for  $i = 1, \dots, l_{max}$  do
3    $q_{rand} = \text{generate randomly in } \mathcal{C}$ 
4    $q_{near} = \text{find nearest node in } \mathcal{T} \text{ towards}$ 
      $q_{rand}$ 
5    $q_{new} = \text{localPlanner from } q_{near} \text{ towards}$ 
      $q_{rand}$ 
6   if canConnect( $q_{near}, q_{new}$ ) then
7      $\mathcal{T}.\text{addNode}(q_{new})$ 
8      $\mathcal{T}.\text{addEdge}(q_{near}, q_{new})$ 
9     if  $\varrho(q_{new}, q_{goal}) < d_{goal}$  then
10      return path from  $q_{init}$  to  $q_{new}$ 
```

- RRT/PRM rely on nearest-neighbor search
- Naïve implementation $\mathcal{O}(n)$
- Fast nearest-neighbor search can be realized using KD-trees
- KD-tree is a binary tree, nodes represent a decision value
- Each level (of node) is for a different dimension




- Search is $\mathcal{O}(\log n)$ for n points in the KD-tree
- Construction $\mathcal{O}(dn \log n)$, n is number of points, d is dimension

- RRT/PRM rely on nearest-neighbor search
- Naïve implementation $\mathcal{O}(n)$
- Fast nearest-neighbor search can be realized using KD-trees
- KD-tree is a binary tree, nodes represent a decision value
- Each level (of node) is for a different dimension




- Search is $\mathcal{O}(\log n)$ for n points in the KD-tree
- Construction $\mathcal{O}(dn \log n)$, n is number of points, d is dimension

- RRT/PRM rely on nearest-neighbor search
- Naïve implementation $\mathcal{O}(n)$
- Fast nearest-neighbor search can be realized using KD-trees
- KD-tree is a binary tree, nodes represent a decision value
- Each level (of node) is for a different dimension




- Search is $\mathcal{O}(\log n)$ for n points in the KD-tree
- Construction $\mathcal{O}(dn \log n)$, n is number of points, d is dimension

- RRT/PRM rely on nearest-neighbor search
- Naïve implementation $\mathcal{O}(n)$
- Fast nearest-neighbor search can be realized using KD-trees
- KD-tree is a binary tree, nodes represent a decision value
- Each level (of node) is for a different dimension



- Search is $\mathcal{O}(\log n)$ for n points in the KD-tree
- Construction $\mathcal{O}(dn \log n)$, n is number of points, d is dimension

- RRT/PRM rely on nearest-neighbor search
- Naïve implementation $\mathcal{O}(n)$
- Fast nearest-neighbor search can be realized using KD-trees
- KD-tree is a binary tree, nodes represent a decision value
- Each level (of node) is for a different dimension



- Search is $\mathcal{O}(\log n)$ for n points in the KD-tree
- Construction $\mathcal{O}(dn \log n)$, n is number of points, d is dimension

- Random numbers generator
- Metric
- Nearest-neighbor search
- Collision-detection

```
1 initialize tree  $\mathcal{T}$  with  $q_{init}$ 
2 for  $i = 1, \dots, l_{max}$  do
3    $q_{rand} = \text{generate randomly in } \mathcal{C}$ 
4    $q_{near} = \text{find nearest node in } \mathcal{T} \text{ towards}$ 
      $q_{rand}$ 
5    $q_{new} = \text{localPlanner from } q_{near} \text{ towards}$ 
      $q_{rand}$ 
6   if  $\text{canConnect}(q_{near}, q_{new})$  then
7      $\mathcal{T}.\text{addNode}(q_{new})$ 
8      $\mathcal{T}.\text{addEdge}(q_{near}, q_{new})$ 
9     if  $\varrho(q_{new}, q_{goal}) < d_{goal}$  then
10      return path from  $q_{init}$  to  $q_{new}$ 
```


- Determines if/how objects collide/overlap/intersect/touch
- “Collision detection” covers two different techniques:

Collision-detection:

- true/false answer
- fast, suitable for sampling-based planners


Collision-determination:

- report details about collisions
- identify which objects intersect
- enumerate involved primitives
- optionally computes penetration vector, point of intersection
- slower than collision-detection



Consider the manipulator at collision

- How can you react with collision-detection?
- How collision-determination helps to overcome the problem?




Geometric primitives


- Are basic shapes used to describe our object (robot, obstacle, ...)
- Objects are constructed from these primitives
- Points, Lines, Circles, Triangles, Spheres, Cylinders, Rectangles
- CD relies on intersection tests between the primitives
- Geometry of the primitives determines which CD algorithm can be chosen
- Convex shapes are often better, CD is faster with them

Collision detection between n and m primitives

- Naïve CD: $\mathcal{O}(mn)$
- This can be too slow!




- Reduce complexity of CD by replacing the original object by a simpler object that contains the original one
- Represent an object by a Bounding Box (BB)
- If two BBs do not overlap, object inside cannot collide (fast test)
- If two BBs collide, further test is made using internal objects (slow test)




- BB should be geometrically simple to enable fast BB-vs-BB tests
- Spheres/circles, ellipses, rectangles
- BB should be as tight as possible to minimize false-positives

- **AABB — Axis Aligned Bounding Box**
 - Faces of bounding box are parallel to the coordinated system
 - Very fast detection of overlap of two BBs
 - Not suitable for 'rotated' objects that lead to large BB
- **OBB — Oriented Bounding Box**
 - Faces of BB are oriented according to the object
 - Lower volume of BB, less false-positives
 - Slower detection of BBs overlap than for AABB
- **k-DOP — k Discrete Oriented Polytope**
 - Boolean intersection along k directions
 - Axes of DOPs do not have to be orthogonal
 - Generalization of AABB/OBB (e.g., AABB in 2D is 4-DOP)






- Is used to determine overlap of two convex objects
- Two convex polytopes do not overlap if there exists a line onto which the projection of the two objects do not overlap
- Separating line can be determined by testing all combinations of lines/faces of both objects
- Convex objects!



Bounding Volume Hierarchy (BVH)

- Original objects are recursively split to subsets
- BVH is a tree structure of bounding-boxes (BB) for each subset
- A Node in BVH is either a BB or a geometric object






- S. Liu, C. C. L. Wang, K. Hui, X. Jin, H. Zhao. Ellipsoid-tree construction for solid objects. ACM symposium on Solid and physical modeling, 2007.


- **Broad phase**

- Traverse BVH from the root
- At each level, evaluate overlaps between BBs
- If BBs do not overlap, return no-collision
- If BBs overlap, continue to child nodes




- **Narrow phase**


- for two overlapping BBs, perform collision detection of their internal objects
- Hierarchical CD: $\mathcal{O}(\log n)$ for n geometric primitives
- Building of BVH (depends on its type) takes at least $\mathcal{O}(n)$




- Usual representation for 2D objects:
 - Combination of boxes/spheres, polygons, triangulated polygons
- Usual representation for 3D objects:
 - Combination of 3D geometric primitives (boxes,spheres,cylinders), triangle mesh
- Note: triangle meshes are hollow → detection of 'object inside object' is not possible



No collision




Collision



No collision

- CD between objects of the same type is usually faster than between objects of different types
- It's a good practice to represent the robot by a combination of basic primitives than using a full CAD model



Collision-detection
~ 100 triangles/robot



Visualization, from CAD
~ 10k triangles + textures/robot

- Sampling-based planners rely on a “local planner”
- Given configurations $q_a \in \mathcal{C}_{\text{free}}$ and $q_b \in \mathcal{C}_{\text{free}}$, local planner attempts to find a path τ :


$$\tau : [0, 1] \rightarrow \mathcal{C}_{\text{free}}$$

such that $\tau(0) = q_a$ and $\tau(1) = q_b$, and τ must be collision free!


- Two-point boundary value problem (BVP)

Types of local planners (revision)


- Exact: analytic solution to BVP, e.g., Dubins or Reeds Shepp, straight-line (sometimes)
- Approximate: τ from q_a with q_{new} that is near-enough from q_b , e.g., straight-line
- Black-box models: physical simulation, e.g., for situations that cannot be solved analytically



Exact local planner






Approximate









Straight-line

- Let's assume a non-trivial scenario, e.g.,
 - mobile robot moving on a undulating terrain
 - or a legged robot walking on stones
- Analytic motion model is not easy to derive
- Instead, we can use a (physical) simulation
- Simulation is used as a “black-box”



- Motion model of objects based on Newton physics
- Complex objects (robots) are composed of basic primitives
 - Spheres, Boxes, Cylinders
 - Analytic collision determination
- Each object has shape, mass and mass-density
- Objects are connected using static/movable joints
- Each joint has limits/maximal moments, speed (+ internal states)
- Internal state s_i of object i : position, rotation, velocity, angular velocity
- Open-source libraries: ODE, Bullet
- User-friendly simulators (Gazebo, V-REP, ...)







- Box, four wheels,
- two one-axis hinges (rear wheels)
- two two-axis hinges (front wheels)
- Surface: either plane of triangular mesh


Particle

- Position $x(t)$, velocity $v(t)$, and mass m
- Various forces \mathbf{F}_i are applied on the particle
- Particle movement is not constrained
- $\mathbf{F} = \sum \mathbf{F}_i$ is the total (net) force
- $\mathbf{F} = m\mathbf{a}(\mathbf{t})$, $\mathbf{a}(\mathbf{t}) = \dot{\mathbf{v}}(\mathbf{t})$, $\mathbf{v}(\mathbf{t}) = \dot{\mathbf{x}}(\mathbf{t})$
- Simulator computes $\mathbf{a}(\mathbf{t}) \rightarrow \mathbf{v}(\mathbf{t}) \rightarrow \mathbf{x}(\mathbf{t})$
- Integration over time-step ε (resolution of the simulation)
- Requires integration (Euler method, Midpoint, Runge-Kutta, ...)
- Particle has no rotation




Rigid body

- Center of mass (CM) $\mathbf{x}(t)$, orientation $\mathbf{R}(t)$ (around CM)
- Translational velocity $\mathbf{v}(t)$ and angular velocity $\omega(t)$
- Net force $\mathbf{F}(t)$
- State vector is $\mathbf{y}(t) = (\mathbf{x}(t), \mathbf{R}(t), \mathbf{p}(t), \mathbf{L}(t))$
- $\mathbf{p}(t)$ is impulse, $\mathbf{L}(t)$ is angular momentum
- Constants: mass m , inertia tensor
- Unconstrained motion
- Leads to $\dot{\mathbf{y}}(t) = f(t, \mathbf{y}(t))$
- initial conditions: $\mathbf{y}(t_0) = \mathbf{y}_0$
- Solved by numerically




Colliding contact

- Particle is falling to a table
- Integration goes by step Δt : $t_0, t_0 + \Delta t, t_0 + 2\Delta t, \dots$
- Integration is terminated if collision happens
- The time of collision t_c is estimated
- Change of velocities of colliding bodies is computed
- Simulation is started from t_c with new velocities
- This ensures instant change of velocities after the collision





Detection of t_c assuming $t_0 < t_c < t_0 + \Delta t$

- Integration by intervals determined by the bisection method
- Alternatively, obtain collision depth from CD and accept t_c if the penetration depth is less than ε



- Vertex/face
 - Vertex of one object is in contact with face of the other one
 - The Normal vector of the face determine the 'normal of the contact' \hat{n}
- Edge/Edge
 - Two edges ea and eb (each from different object) are in collision
 - $\hat{n} = ea \times eb$ (ea and ev are unit vectors)



- Contacts $p_a(t_0)$ and $p_b(t_0)$, their velocity is $\dot{p}_a(t_0)$ and $\dot{p}_b(t_0)$
- $v_{rel} = \hat{n}(t_0) \cdot (\dot{p}_a(t_0) - \dot{p}_b(t_0))$
- Value of v_{rel} determines the type of collision

Separating




- $v_{rel} > 0$: bodies moving apart
- No reaction is needed

Contact/resting


- $v_{rel} = 0$

Colliding

- $v_{rel} < 0$




How many contacts we have? What is their type?




Example of contacts

How many contacts we have? What is their type?



Particle

- 1 Create objects, create joints, ...
- 2 User callback (read/set variables, display, ...)
- 3 Apply forces
- 4 Update velocities and positions
- 5 Detect collisions
- 6 Solve constraints
- 7 Goto 2





Physical engines (sw. libraries)

- Box2D, Chimpunk physics engine (2D)
- ODE, Bullet, Newton Game Physics (3D)



Robotic simulators (usually with GUI)

- They use physical engine inside, but offer more functionalities:
- Visualization, tools for interactive design of robots, import/export from URDF, sensors
- Gazebo, V-Rep (now CoppeliaSim), Webots, Player/Stage


- If wrongly set up, it can “explode” or “freeze”
- Wrongly set up hinges, unrealistic masses, no gravity
- Too complex geometries → too complex (slow) collision detection
- Wrong friction parameters
- It's better to prefer convex shapes (or composition of them) if possible



- 3D design/CAD simulation — design a machine and see how it works
- Virtual reality — e.g. for realistic object manipulation
- Computer games — realistic behavior of objects (without programming it)
- Evolving robots — evolutionary approaches to design robots or their parts, simulation serves as the fitness function evaluator



In 2013, we saw simulated robots made of soft voxel cells evolve the ability to run.





Cheney, N., MacCurdy, R., Clune, J., & Lipson, H. (2013). Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proceeding of the fifteenth annual conference on genetic and evolutionary computation* (pp. 167-174). ACM.


- Most of physical simulators (ODE, Bullet and their derivatives) assume time-linear simulation
- In motion planning, we need a non-linear simulation
- We need to “restart” simulation for each tree expansion

RRT with system simulator


- Each node contains: $x = (s_i), i = 1, \dots, n$ (simulator state)
- Tree expansion from node $x_{\text{near}} = x(0)$ using input u
 - Set simulator to state x_{near} (restart)
 - Apply control inputs u (usually joint moments)
 - Run simulation for time Δt
 - Read simulator state x
 - Add node x to the tree
- Usually several control inputs $u \in \mathcal{U}$ is tested



We need




No restarts




With restarts

- Try to minimize number of objects/joints
- Avoid using triangle mesh for collision-detection
 - Physical simulation needs collision determination (penetration vector)
 - CD may be unstable on (non-convex) meshes, simulation can “explode”
- If possible, approximate robots by boxes/spheres/cylinders → fast and stable collision detection
- Use separate models for physics and visualization



Mass
10 boxes/robot



Collision-detection
~ 100 triangles/robot



Visualization
~ 10k triangles + textures/robot

