

Real-Time Randomized Path Planning for Robot Navigation*

James Bruce and Manuela M. Veloso

Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh PA 15213, USA
{jbruce,mmv}@cs.cmu.edu

Abstract. Mobile robots often find themselves in a situation where they must find a trajectory to another position in their environment, subject to constraints posed by obstacles and the robot's capabilities. This poses the problem of planning a path through a continuous domain. Several approaches have been used to address this problem each with some limitations, including state discretizations, planning efficiency, and lack of interleaved execution. Rapidly-exploring random trees (RRTs) are a recently developed algorithm on which fast continuous domain path planners can be based. In this work, we build a path planning system based on RRTs that interleaves planning and execution, first evaluating it in simulation and then applying it to physical robots. Our algorithm, ERRT (execution extended RRT), introduces two novel extensions of previous RRT work, the waypoint cache and adaptive cost search, which improve replanning efficiency and the quality of generated paths. ERRT is successfully applied to a multi-robot system. Results demonstrate that ERRT improves efficiency and performs competitively with existing heuristic and reactive real-time path planning approaches. ERRT has shown to offer a major step with great potential for path planning in challenging continuous, highly dynamic domains.

Introduction

The path-planning problem is as old as mobile robots, but is not one that has found a universal solution. Specifically, in complicated, fast evolving environments such as RoboCup [3], currently popular approaches have their strengths, but still leave much to be desired. A recently developed tool that may help tackle the problem of real-time path planning are Rapidly-exploring random trees (RRTs) [7]. RRTs employ randomization to explore large state spaces

* This research was sponsored by Grants Nos. DABT63-99-1-0013, F30602-98-2-0135 and F30602-97-2-0250. The information in this publication does not necessarily reflect the position of the funding agencies and no official endorsement should be inferred.

efficiently without tiling, and can form the basis for a probabilistically complete though non-optimal kinodynamic path planner [8]. They can efficiently find plans in high dimensional spaces because they avoid the state explosion that discretization faces. Furthermore, due to their incremental nature, they can maintain complicated kinematic constraints if necessary.

Most current robot systems that have been developed to date are controlled by heuristic or potential field methods at the lowest level, and many extend this upward to the level of path navigation [5]. Since the time to respond must be bounded, *reactive* methods are used to build constant or bounded time heuristics for making progress toward the goal. One reactive method that has proved quite popular is motor schemas [1]. Although they meet the need for action under time constraints, most of these methods suffer from the lack of lookahead, which can lead to highly non-optimal paths and problems with oscillation. RRTs, as used in our work and presented in this paper, should provide a good compliment for very simple control heuristics, and take much of the complexity out of composing them to form a navigation system. Specifically, local minima can be reduced substantially through lookahead.

While not as popular as heuristic methods, non-reactive planning methods for interleaved planning and execution have been developed, with promising results. Among these are agent-centered A* search methods [4] and the D* variant of A* search [9]. However, using these planners requires discretization or tiling of the world in order to operate in continuous domains. This leads to a tradeoff between a higher resolution, with is higher memory and time requirements, and a low resolution with non-optimality due to discretization. Most of the features of agent-centered search methods do not rely on A* as a basis, however, so we can achieve many of their benefits using an RRT based planner which fits more naturally into domains with continuous state spaces. In addition, the base RRT system is relatively easy to extend to environments with moving obstacles, higher dimensional state spaces, and kinematic constraints. No other planners we are currently aware of offer these possible enhancements at the simplicity and speed of an RRT-based solution. As a step in that direction, this work appears to be the first successful application of an RRT planner to a real mobile robot [6].

In order to make online planning efficient enough to be practical, we introduce two novel additions to the planner, specifically the waypoint cache for replanning and adaptive cost penalty search. The second section of this paper defines the basic RRT algorithm. The next section introduces our ERRT contribution, followed by implementation results. The final sections offers concluding remarks.

RRT Planning

Basic RRT Algorithm

In essence, an RRT planner searches for a path from an initial state to a goal state by iteratively expanding a search tree. For its search, it requires the following three domain-specific function primitives:

Function *Extend* (env:environment,current:state,
target:state):state

Function *Distance* (current:state,target:state):real

Function *RandomState* ():state

The *Extend* function calculates a new state that can be reached from the target state by some incremental distance (usually a constant distance or time), which in general makes progress toward the goal. If a collision with an obstacle in the environment would occur by moving to that new state, *EmptyState* is returned. In general, any heuristic methods suitable for control of the robot can be used here, provided there is a reasonably accurate model of the results of performing its actions. The heuristic does not need to be very complicated, and does not even need to avoid obstacles (just detect when a state would hit them). However, the better the heuristic, the fewer nodes the planner will need to expand on average, since it will not need to rely as much on random exploration.

Table 1 shows the complete basic RRT planner with its stochastic decision between the search options:

- with probability p , it expands towards the goal minimizing the objective function *Distance*,
- with probability $1 - p$, it does random exploration by generating a *RandomState*.

The function *Distance* needs to provide an estimate of the time or distance (or any other objective that the algorithm is trying to minimize) that estimates how long repeated application of *Extend* would take to reach the goal. For a simple example, a holonomic point robot with no acceleration constraints can implement *Extend* simply as a step along the line from the current state to the target, and *Distance* as the Euclidean distance between the two states. Next, the function *Nearest* uses the distance function implemented for the domain to find the nearest point in the tree to some target point outside of it. The function *RandomState* simply returns a state drawn uniformly from the state space of the environment.

Finally, *ChooseTarget* chooses the goal part of the time as a directed search, and otherwise chooses a target taken uniformly from the domain as an exploration step. The main planning procedure uses these functions to iteratively pick a stochastic target and grow the nearest part of the tree towards that target. The algorithm terminates when a threshold distance to the goal has been reached, though it is also common to limit the total number of nodes that can be expanded to bound execution time.

Extended RRT Algorithm – ERRT

Some optimizations over the basic described in existing work are bidirectional search to speed planning, and encoding the tree's points in an efficient spatial data structure [2]. In this work, a KD-tree was used to speed nearest neighbor

Table 1. The basic RRT planner stochastically expands its search tree to the goal or to a random state.

```

function RRTPlan (env:environment,initial:state,
                    goal:state):rrt-tree
  var nearest,extended,target:state;
  var tree:rrt-tree;
  nearest := initial;
  rrt-tree := initial;
  while(Distance (nearest,goal) < threshold)
    target = ChooseTarget (goal);
    nearest = Nearest (tree,target);
    extended = Extend (env,nearest,target);
    if extended ≠ EmptyState then
      AddNode (tree,extended);
  return tree;

function ChooseTarget (goal:state):state
  var p:real;
  p = UniformRandom in [0.0 .. 1.0];
  if 0 < p < GoalProb then
    return goal;
  else if GoalProb < p < 1 then
    return RandomState();

function Nearest (tree:rrt-tree,target:state):state
  var nearest:state;
  nearest := EmptyState;
  foreach state s in current-tree
    if Distance (s,target) <
      Distance (nearest,target) then
      nearest := s;
  return nearest;

```

lookup, but bidirectional search was not used because it decreases the generality of the goal state specification (it must then be a specific state, and not a region of states). Additional possible optimizations include a more general biased distribution, which was explored in this work in the form of a waypoint cache. If a plan was found in a previous iteration, it is likely to yield insights into how a plan might be found at a later time when planning again; The world has changed but usually not by much, so the history from previous plans can be a guide. The waypoint cache was implemented by keeping a constant size array of states, and whenever a plan was found, all the states in the plan were placed into the cache with random replacement. This stores the knowledge of where a plan might again be found in the near future. To take advantage of this for planning, Table 2 shows the modifications to the function *ChooseTarget*.

Now, there are three probabilities in the distribution of target states. With probability $P[goal]$, the goal is chosen as the target; With probability

Table 2. The extended RRT planner stochastically expands its search tree to the goal; to a random state; or to a waypoint cache.

```

function ChooseTarget'(goal:state):state
  var p:real;
  var i:integer;
  p = UniformRandom in [0.0 .. 1.0];
  i = UniformRandom in [0 .. NumWayPoints-1];
  if 0 < p < GoalProb then
    return goal;
  else if GoalProb < p < GoalProb+WayPointProb then
    return WayPointCache[i];
  else if GoalProb+WayPointProb < p < 1 then
    return RandomState();
    
```




Fig. 1. Extended RRT with a waypoint cache for efficient replanning.

$P[waypoint]$, a random waypoint is chosen, and with the remaining probability a uniform state is chosen as before. The way the extended algorithm progresses is illustrated in Figure 1. Typical values used in this work were $P[goal] = 0.1$ and $P[waypoint] = 0.6$. Finally, A simple RRT planner is building a greedy approximate spanning tree, and thus ignores the path lengths from the initial state (the root node in the tree). The distance metric can be modified to include not only the distance from the tree to a target state, but also the cost to reach that state. A higher value of the multiplier on that history cost (beta) results in shorter paths from the root to the leaves, but also decreases the amount of exploration of the state space, biasing it to near the initial state in a “bushy” tree. A value of 1 for beta will always extend from the root node for any Euclidean metric in a continuous domain, while a value of 0 is equivalent to the original algorithm. The best value seems to vary with domain and even problem instance, and appears to be a steep tradeoff between finding an shorter plan and not finding one at all. To address this we implemented an adaptive mechanism that appears to work quite well. When the planner starts, beta is set to 0. Then on successive replans, if the previous run found a plan, beta is incremented, and decremented otherwise. In addition the value is clipped to between 0 and 0.65. This adaptive bias schedule reflects the idea that a bad plan is better than no plan initially, and once a plan is in the cache and search is biased toward the




Fig. 2. An example from the simulation-based RRT planner, shown at three times during a run. The initial state (agent) is in blue, while the fixed target state is shown in red. The best plan (the one that gets closest to the goal) is shown in white.

waypoints, nudges the system to try to shorten the plan helping to improve it over successive runs.

Domain Implementations and Results

RoboCup F180 Robot Control

We tested our RRT system in simulation and then applied it so the control of real robots. In both, planning and execution are interleaved equally; A plan is created for each step taken. Three snapshots from our simulator can be seen in Figure 2. In this simple domain, the waypoints seemed to help qualitatively in that the robot didn't tend to get stuck as long in local minima. This effect of waypoints on performance was evaluated numerically, and the results are shown in in Figure 3. Since the curves diverge at moderate difficulty, it appears that waypoints help speed planning by offering “hints” from previous solutions. When the problem becomes impossible or nearly impossible, neither performs well.

For physical robot control, the system must take input from a global vision system at 30Hz, reporting the position of all field objects detected from a fixed overhead camera, and send the output to a radio server which sends velocity commands to the robots that are being controlled. The path planning problem is to navigate quickly among other robots, while they also more around executing their own behaviors. As a simplification, we examined the more simple problem of running from one end of the field to the other, with static robots acting as obstacles in the field. The extension metric we used was hen became a model of a simple heuristic “goto-point” that had already been implemented for the robot. The motivation for this heuristic approach was that executing a bad plan immediately is often better than sitting still looking for a better one that models the robot more correctly. To increase system performance, after a path had been determined, it is post processed, replacing the head of the plan with the longest obstacle-free straight path. Not only does this smooth out the resulting plan, but the robot tends to go straight at the first “restricted” point, always trying to aim at the free space. This allowed the robot to navigate at




Fig. 3. The left side shows the effect of waypoints. The lines show the sorted planning times for 2670 plans with and without waypoints, sorted by time to show the cumulative performance distribution. The curve with waypoints is significantly lower for intermediate planning times. The right side shows the plot of planning times vs. the nodes expanded. KD-trees hold a reasonable advantage even at small numbers of nodes.

up to $1.7m/s$, performing better than any previous system we have used on our robots. Videos of the current system are available at the following address: <ftp://sponge.coral.cs.cmu.edu/pub/movies/F180-RRT>

The best combination of parameters that we were able to find, trading off physical performance and success with execution time was the following: 500 nodes, 200 waypoints, $P[goal] = 0.1$, $P[waypoint] = 0.7$, and a step size of $1/15sec$. To examine the efficiency gain from using a KD-tree, we ran the system with and without a KD-tree. The results are shown in Figure 3. Not only does the tree have better scalability to higher numbers of nodes due to its algorithmic advantage, but it provides an absolute performance advantage even with as few as 100 nodes. The planner was able to perform on average in 2.1ms, with the time rarely going above 3ms. This makes the system fast enough to use in our production RoboCup team, as it will allow 5 robots to be controlled from a reasonably powerful machine while leaving some time left over for higher level action selection and strategy.

Conclusion

In this work a robot control system was developed that used an RRT path planner to turn a simple reactive scheme into a high performance path planning system. The novel mechanisms of the waypoint cache and adaptive cost search were introduced, with the waypoint cache providing much improved performance on difficult but possible path planning problems. The real robot was able to perform better than previous fully reactive schemes, traveling 40% faster while avoiding obstacles, and drastically reducing oscillation and local minima

problems that the reactive scheme had. This is also the first application of which we are aware using an RRT-based path planner on a real mobile robot.

Several important lessons can be drawn from this work in the context of real-time path planning in highly dynamic domains:

- A heuristic may perform better than a correct model when planning time is critical. In other words, a better model may not improve the entire system even if it makes the generated plans better.
- A plan generated from an incorrect model requires post-processing for optimal performance. The system worked without post processing, but not nearly as well as when the local metric could apply its accurate model over a longer range of the plan and thus remove most of its inaccuracies over that period of the plan.
- Pre-existing reactive control methods can easily be adapted to be RRT extension and distance metrics. It can build on these to help eliminate oscillation and local minima through its lookahead mechanism. Since many existing robots already have reactive navigation systems, and the RRT core code is highly generic, We expect this to be a common adaptation in the future.

This work could not have been conducted without the many people in our group who support our RoboCup F180 small robot team. We would specifically like to thank Brett Browning and Mike Bowling, without whom we wouldn't have a team or robots with which interesting research projects could be done.

References

1. R. C. Arkin. Motor schema-based mobile robot navigation. *International Journal of Robotics Research*, August 1989, 8(4):92–112, 1989.
2. A. Atramentov and S. M. LaValle. Efficient nearest neighbor searching for motion planning. In *submitted to 2002 IEEE International Conference on Robotics and Automation*, 2002.
3. H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, and E. Osawa. Robocup: The robot world cup initiative. In *Proceedings of the IJCAI-95 Workshop on Entertainment and AI/ALife*, 1995.
4. S. Koenig. Agent-centered search. *Artificial Intelligence*, 2002, in print.
5. J.-C. Latombe. *Robot Motion Planning*. Kluwer, 1991.
6. S. M. LaValle. Rapidly-exploring random trees (<http://msl.cs.uiuc.edu/rrt/>).
7. S. M. LaValle. Rapidly-exploring random trees: A new tool for path planning. In *Technical Report No. 98-11*, October 1998.
8. S. M. LaValle and J. James J. Kuffner. Randomized kinodynamic planning. In *International Journal of Robotics Research*, Vol. 20, No. 5, pages 378–400, May 2001.
9. A. Stentz. Optimal and efficient path planning for unknown and dynamic environments. In *International Journal of Robotics and Automation*, Vol. 10, No. 3, 1995.