NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

Planning with Map Uncertainty

Dave Ferguson Anthony Stentz

CMU-RI-TR-04-09

February 2004

Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

© Carnegie Mellon University

University Libraries
Carnegie Mellon University

Abstract

We describe an efficient method for planning in environments for which prior maps are plagued with uncertainty. Our approach processes the map to determine key areas whose uncertainty is crucial to the planning task. It then incorporates the uncertainty associated with these areas using the recently developed PAO* algorithm to produce a fast, robust solution to the original planning task.

Contents

1 Introduction **1**

2 Planning with Uncertainty **1**

3 Extracting Pinch Points **3**

4 Planning with Pinch Points **5**
4.1 The Adjacency Graph 5
4.2 Planning with the Adjacency Graph: PAO* 6

5 Results **8**

6 Replanning **10**

7 Conclusion **11**

8 Acknowledgments **11**

1 Introduction

There are several domains in which agents are required to plan with imperfect information. Consider a land-based robot navigating outdoors with a low-resolution overhead map generated by a helicopter. This map may be quite inaccurate: the density of the data may be low and the position estimation used to project this data onto the overhead map could contain significant error. As a result, any planning map extracted from this overhead map and used by the robot will be imperfect.

In examples like these, the navigating agent knows that it has incomplete or uncertain information. Moreover, it is often also aware of the *nature* of this uncertainty. In our example, the robot may know the density of the data used to produce the original map, as well as the uncertainty associated with the helicopter position for each data point. It would also have an idea of the error associated with any footprint convolution process used to take the original map and extract a terrain map for planning.

The combination of these sources of information can be used to derive error models for the final planning map. Given these error models, the robot can determine a reasonable approximation of the probability that a given cell in its planning grid holds a particular terrain value.

In this paper we deal with exactly the above situation: an agent is equipped with a planning grid of its environment, where every cell in the grid has some uncertainty associated with its value.

We begin by discussing the nature of the problem and describe current approaches to planning with uncertainty. We go on to introduce our novel solution and provide key results and extensions.

2 Planning with Uncertainty

Figure 1 shows a sample uncertainty distribution over the terrain associated with a particular cell in an environment. The idea is that, although we may not have perfect information regarding a cell's terrain, we can often extract such terrain distributions using the information we do have and some error models.

Dealing with distributions over terrains rather than fixed values requires some modification to classical planning techniques. Currently, there are two common methods of planning that incorporate this type of uncertainty.

The first method, known as **assumptive planning** [5], computes an approximate cost of traversing each cell (or 'assumes' a default value) and plans using these approximations. When the resulting plan is executed, if a discrepancy is found between a cell's assumed cost and actual cost, the plan can be updated to reflect the newly acquired information. This type of planning is very fast, since it can use A* techniques to focus its initial computation and D* techniques to repair previous plans [7, 4].

However, planning with approximate costs breaks down when cells have some probability of being untraversable. Typically, assumptive planners solve for a path from the start state s to a goal state g while minimizing the overall cost of the path.

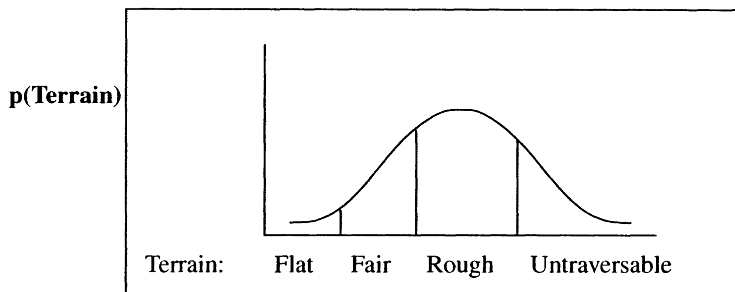

Figure 1: A sample terrain distribution

The overall cost of a cell x on the path is defined as:

$$Cost(x,g) = \min_{y \in nbrs(x)} Cost(x,y) + Cost(y,g).$$

The value $Cost(x, y)$ is computed from the approximate cost of traversing from cell x to neighboring cell y. But if there is some non-zero probability that cell x is untraversable, the equation above is no longer valid, as it does not account for this possibility (and its associated cost). In order to get the appropriate cost, an alternative path needs to be planned *around* x to the goal. The cost of this path is nontrivial to compute, as it requires making similar considerations for all other cells encountered, which makes the overall computation exponential in the number of cells in the environment. As a result, accurate cost approximations are difficult to generate and computed solution paths based on rough estimates can be highly sub-optimal.

A second method addresses the limitations of approximating costs and deals with terrain uncertainty more comprehensively. **Partially observable planning** [3] generates every possible terrain of every cell and finds the overall best plan taking into account these possibilities and their associated probabilities. Thus, rather than using an approximation of the cost of using a particular cell on a path to the goal, it computes the true expected cost by generating every possible outcome and its associated cost and probability. However, as mentioned above, computing these costs is exponential in the number of cells in the environment. Moreover, as an agent moves through its environment it learns the true terrain values of areas it encounters. Consequently, the terrain distribution associated with each cell can change and this change needs to be taken into account in the planning stage. As a result, partially observable planning is often intractable for reasonably sized environments. Even if our terrain model is reduced to two possibilities, traversable and untraversable, and the agent is equipped with a perfect contact sensor (the simplest case), planning involves dealing with $3^{m \cdot n}$ information states[1], where our planning grid is $m \cdot n$ cells in size. This is a prohibitively large state space.

[1] Each cell may be known to be traversable (t), known to be untraversable (o), or not yet seen by the agent (u), in which case it has its initial probability distribution over being traversable/untraversable. This results in 3 different states of *information* the agent may have concerning the terrain of each cell.

1. Initially, set the cost of each cell in the environment to its expected traversable cost. Use D* to plan an optimal path (relative to these costs) from the robot position to the goal. Mark the pinch points along this path (see 2(b)) and use them to construct an ordered list P.

2. While there are still pinch points in P:

   (a) *Replan path from pinch point to goal*: Remove the top pinch point from P and set the terrain of each cell belonging to a pinch point to *untraversable.* Invoke D* to replan a path to the goal from the cell previous to the pinch point (on the path which the pinch point was found on).

   (b) *Mark new pinch points along path*: Step along this new path and look for sections which have a high probability of being untraversable *and* are costly to navigate around. Mark these cells as pinch points and add them to the end of P.

3. Return all pinch points encountered.

Figure 2: Pinch Point Extraction Algorithm

It is possible to reduce our state space to only those cells which are likely to be important to an agent navigating the environment. In this way, we gain the significant advantage of partially observable planning without incurring anything like its computational cost.

3 Extracting Pinch Points

The idea is to focus our computation on areas of the environment whose traversability, and hence uncertainty, is crucial to the planning task. Thus, rather than dealing comprehensively with the uncertainty associated with every cell, as partially observable planning does, we restrict our attention to those cells which are most useful.

These are the ones which would cause a costly detour in the robot's path to the goal if they turned out to be untraversable. We would like to be able to detect these cells and incorporate their terrain uncertainty in our planner so that we could decide from the outset whether it is better to avoid these cells entirely or risk going through them.

To find these cells, coined 'pinch points' [2], we generate a set of cells which could reasonably be encountered by an agent navigating to the goal and look for key members of this set. Figure 2 outlines the process.

We first assume each cell with a reasonable probability of being traversable is in fact traversable and generate its resulting expected terrain cost[2]. We then calculate a path to the goal which is optimal relative to these costs. Given this path, labelled as consecutive cells $r_1 \dots r_n$, we can find all sections of the path which are potential blockages *and* would cause significant detours if found to be blocked.

[2]This is computed by normalising its terrain distribution to only contain the traversable range then taking an expectation.

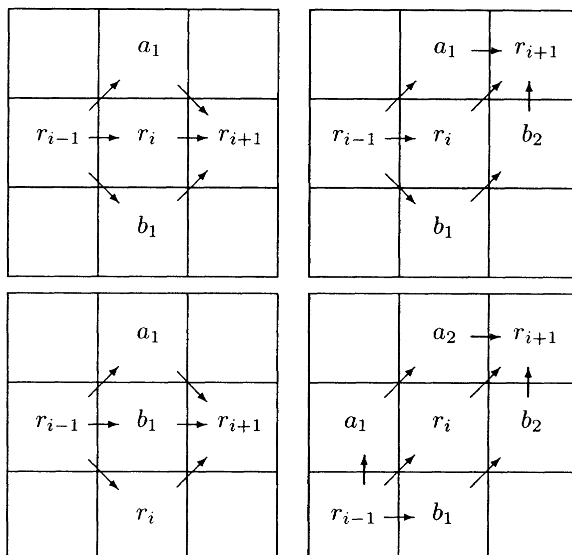

Figure 3: Finding potential path blockages in an eight-connected grid. Each diagram highlights three local paths from cell r_{i-1} to r_{i+1} for different relative positions of cells r_{i-1}, r_i, and r_{i+1}. If there is a nontrivial probability that all of these three paths are untraversable, the cells along the paths are grouped together as a potential pinch point.

A section of the path centered on cell r_i is a potential blockage if there is a nontrivial probability that an agent at cell r_{i-1} will not be able to get to cell r_{i+1} using one of the three shortest non-intersecting routes between the two cells. Here, the idea is that if r_i turns out to be untraversable, there is still a good chance r_{i+1} is reachable from r_{i-1} using one or more of the neighbors of r_i, but if r_i is untraversable and the next obvious two paths from r_{i-1} to r_{i+1} are also untraversable, then it is not so likely. To generate the probability that the path is blocked at cell r_i, we look at the three shortest paths from cell r_{i-1} to cell r_{i+1} and use the terrain distributions of the cells along these paths to determine the probability that all three paths are untraversable. If this probability is greater than some threshold, we group cell r_i and the cells along the shortest paths together as a potential pinch point with position r_i.

Figure 3 illustrates the shortest paths between r_{i-1} and r_{i+1} for four different relative positions of the consecutive path cells r_{i-1}, r_i, and r_{i+1}. The paths for all other possible relative positionings can be obtained from these four. In this figure, r_i provides one path between the two cells, while the other paths are denoted by the cells marked a_1, a_2 and b_1, b_2, respectively.

If a group of cells is marked as a potential pinch point we know that there is a nontrivial probability that an agent may not be able to get through this area. In order

to determine the consequences of this possible outcome, we then check how costly a route around the potential pinch point would be. To do this, we generate a cheapest cost path from the previous cell on the path, r_{i-1}, to the next cell on the path, r_{i+1}, without using any of the cells constituting the potential pinch point. If the cost of this path is significant we add the potential pinch point to our list of true pinch points.

If we come across a number of pinch points in a row, for instance in a narrow valley, we combine them into a single pinch point. The resulting probability of the combined pinch point being untraversable is the sum of the individual probabilities and the position of the pinch point used later for planning is taken as the mean of its constituents.

Once we have found all pinch points along the original path, we then set the terrains of all cells comprising these pinch points to untraversable and replan paths from the close side of each of these pinch points to the goal, i.e., from r_{i-1} for each pinch point r_i.This enables us to locate new pinch points which might be encountered by the agent if it found the current pinch point to be untraversable. We iterate the procedure to generate a set of pinch points which could reasonably be encountered (assuming an agent that always acts optimally given its map information). Figure 4 illustrates the approach in action. Note that, in this figure, only the cell r_i corresponding to each pinch point is shown (surrounded by a large red square to aid in illustration), rather than the collection of all cells comprising the pinch point.

In extracting pinch points, D* is used to replan each subsequent path. The efficiency of D* over A* has been widely recognized (see [7, 4]) and, as shown in the results section, this efficiency allows us to generate our set of pinch points very quickly. Once equipped with this set, we can then incorporate its members and their associated terrain uncertainties into the planning process. To do this, we make use of the recently developed PAO* algorithm.

4 Planning with Pinch Points

In [2] the algorithm PAO* was introduced, which solves planning problems involving hidden state such as pinch points. PAO* is applied to an adjacency graph containing the robot position, the goal position, and the pinch points in the environment. It uses the adjacency information to calculate an optimal solution graph, as with the AO* algorithm, in a highly efficient manner.

4.1 The Adjacency Graph

Each pinch point may provide a bridge between several different regions of the environment. For instance, a pinch point located at a Y-junction connects three different regions to each other. The collection of cells adjacent to the pinch point in each region constitute a *face* of the pinch point. The adjacency graph links up these faces by inserting arcs between every pair of faces that are reachable from one another. The cost of an arc between two faces represents the lowest cost associated with moving along a pinch-point free path between the faces and is used to propagate values from one face to another.

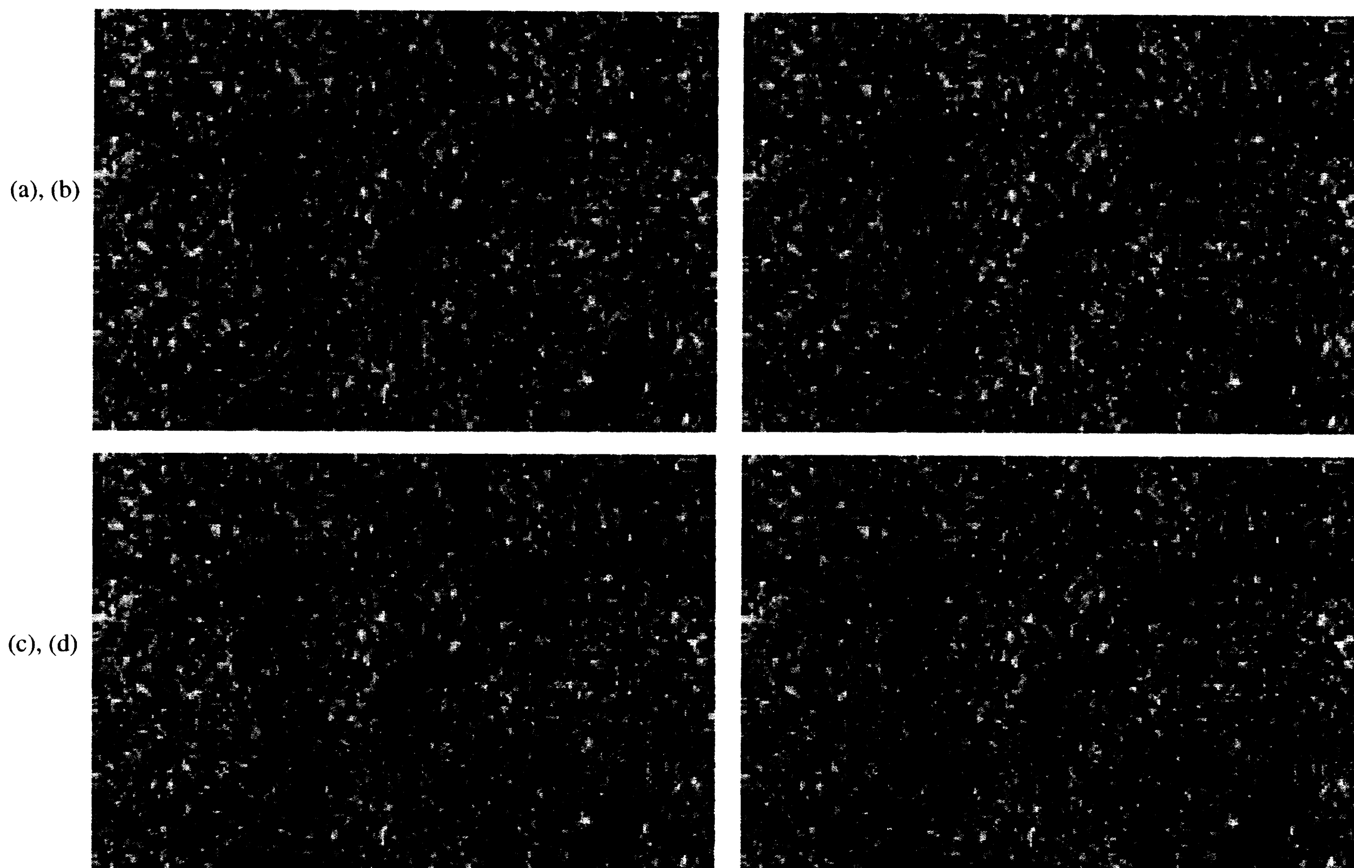

Figure 4: (a) Example of a section of a fractally generated terrain map. The darker a cell, the larger its probability of being untraversable. (b) An optimal path (in blue) from the mid-left side to the mid-right side of this map, assuming cells with a reasonable probability of being traversable are traversable. In red are the positions of pinch points found along this path. (c) An alternative path from the start-side of the first pinch point to the goal (in purple). This path is computed assuming all pinch points from (b) are untraversable. (d) The final set of all pinch points in red with all alternative paths in purple.

4.2 Planning with the Adjacency Graph: PAO*

A shortest path is planned from the robot position to the goal position using this adjacency graph. To do this, the problem is phrased as a search over an AND-OR graph [6]. An AND-OR graph contains two types of nodes: AND nodes obtain their values from combining the values of all their child nodes, while OR nodes compute their values from choosing a single child node value.

Our planning domain can be represented as an AND-OR graph as follows. Each node in the graph corresponds to a face in a particular *information state*. In our setting, an information state is the state of knowledge the agent has concerning the terrain values of each of the pinch points. Following our discussion on partially observable planning, we restrict each pinch point to be known to be traversable (t), known to be untraversable (o), or not yet seen by the agent (u).

The root of the AND-OR graph (an OR node) is the start cell s in the information state characterised by every pinch point being as yet unseen (i.e., of value u). The next level of the graph corresponds to all elements of the adjacency graph, both faces and goal, which have arcs to s. The faces are AND nodes: each has two children representing the two possible information states realizable from visiting the node. These two children each have the same face as their parent but reside in different information states (one has the pinch point associated with the face of value t, the other o). These children are OR nodes because their pinch point has a known value.

PAO*, short for *Propagating AO**, is an algorithm which searches an AND-OR graph by gradually building a solution graph from the start state through two alternating phases, as with AO* [1, 6]. First, it grows the best partial solution by expanding one of the non-terminal leaf nodes and assigning admissible heuristic costs to its children. Next, it uses the newly computed costs to propagate cost revisions throughout the partial solution graph. At each stage in this propagation, OR nodes which are part of the current partial solution update their choice of child to reflect the most recent cost values. An example partial solution graph for our domain is shown in Figure 5.

In the first phase, an initial heuristic cost for a leaf node l is obtained by solving for the cost of the 'heuristic counterpart' of l: the fully-known state characterized by the most desirable true values the pinch points in l could have. Pinch points with known values are left untouched. Pinch points not yet seen (with value u) are assigned the value t. The resulting cost is guaranteed to be admissible.

The major benefit of PAO* lies in its second phase: its propagation of cost revisions. Unlike AO*, PAO* propagates cost changes not only upwards to parents in the partial solution graph, but sideways to neighbors (in the complete AND-OR graph), and downwards to children. The resulting approach makes full use of all received information and thus allows for more informed decisions to be made at each stage of the process. The complete algorithm is given in Figure 6 and thoroughly discussed in [2]. Briefly, there are three key propagation steps that PAO* performs but AO* does not.

Firstly, when the cost of a face in a given information state changes, PAO* propagates this updated cost *across* all faces in the information state, so that dependent faces will have their costs updated as well. While AO* only propagates new information up the partial solution graph, PAO* also propagates it across the full AND-OR graph at each level.

Secondly, PAO* uses the nature of the current problem domain to propagate cost changes *down* the AND-OR graph. Given an AND node with two children corresponding to the two possible true values of the node's pinch point (traversable and untraversable), the cost of the parent node should never be greater than the cost of the untraversable child node. Taking advantage of this piece of intuition, PAO* updates the face costs of the states associated with untraversable nodes so that they are lower bounded by their parent state values. This also enables PAO* to provide a more realistic

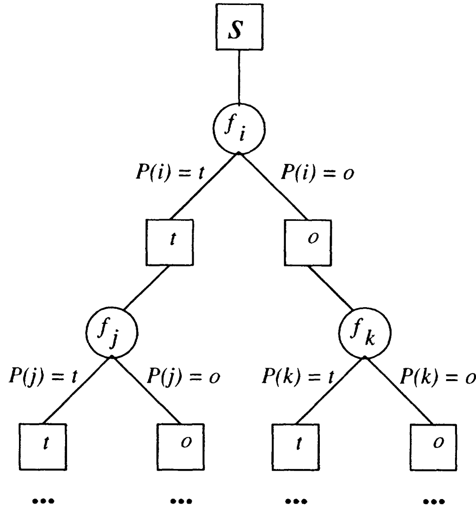

Figure 5: An example partial solution graph. Each circle corresponds to an AND node and each square to an OR node. $P(i) = o$ represents the probability that the pinch point associated with face i turns out to be untraversable.

value for the untraversable child of each newly expanded AND node.

Similarly, the cost of a parent node should never be less than the cost of its traversable child node. Thus, PAO* also updates the face costs of parent information states so that they are lower bounded by their traversable child states. It performs this update as part of its propagation of cost changes back up the solution graph.

These propagation steps combine to allow information gained at one end of the solution graph to be accessible at the other. As a result of these differences, PAO* has been shown to be orders of magnitude more efficient than AO* while still guaranteeing optimal solutions to the graph search.

5 Results

Focussing our computation on pinch points, we are able to solve the planning problem much more efficiently than the full partially observable approach yet still incorporate key areas of uncertainty to produce robust paths. In addition, the complexity of our approach, through both the extraction of pinch points and the subsequent PAO* search for a solution, is highly dependent on both the quality of the information the planning agent holds and the nature of the environment. Thus, when confronted with simple environments and reliable information, it is able to exploit these desirable attributes to produce solutions very quickly. Meanwhile, in more complex environments with less reliable information, our approach incorporates more areas of uncertainty to produce very robust solutions.

1. The initial solution graph consists solely of the start node s in the original information state.

2. While the solution graph has some nonterminal leaf node:

   (a) *Generate fringe node*: Starting from the root, traverse down the solution graph until a nonterminal leaf node is encountered. Along the way, update untraversable child states to have their face costs lower bounded by their parent states.

   (b) *Expand best partial solution*: Expand the nonterminal leaf node and compute cost values for the information states of its children. Traversable child states are given heuristic costs. Untraversable child states inherit their parents' cost values as lower bounds then perform limited value iterations over their heuristic counterparts to potentially increase these values. Add the children to the solution graph, noting whether they are terminal.

   (c) *Propagate cost changes and update solution*: Compute an updated cost of the original leaf node given the costs of its children. If the node's cost has changed, update the cost estimates for its *entire information state* and update its parent's cost to reflect these changes. If the parent is an OR node, the current node may be replaced if it no longer provides the minimum cost. If the node is a traversable child, update the costs associated with the entire parent state to be lower bounded by the current state. Continue propagating up the graph until a node is reached whose cost does not change.

3. Return the optimal solution graph.

Figure 6: The PAO* algorithm

To test the computation required to extract pinch point areas in environments of varying complexity, we generated 1000 fractal terrain grid maps. Each was of size 200 x 200. These environments ranged from very open, easy to navigate areas to very complex, cluttered areas (such as in Figure 4)[3]. With each map, we restricted the total number of pinch points to be 10, so that our results could easily be interpreted in conjunction with the planning results given in [2]. As Figure 4 shows, this still enables us to consider a large number of diverse paths when the environment is highly cluttered.

Over our 1000 terrain test cases, the average number of extracted pinch points was 2.47 and the average time taken for this extraction was 0.07 seconds. The maximum amount of time required to extract the pinch points from any map was 0.34 seconds (with the minimum being 0.003).

Combining these results with the computation required to plan with extracted pinch points (from [2]), we have the entire process conservatively taking 6.7 seconds. This time reflects the worst results reported here and in [2] for environments with 10 pinch points.

[3]For details of the fractal generation process, see [7]. We used a gain of 20 and varied the number of levels from 5 to 9.

A complete partially observable solution over such an environment would need to contend with $3^{200 \cdot 200}$ states rather than the 3^{10} used in our PAO* planning. This is currently far too large a problem to be solved. However, by restricting our attention to the most important areas of the environment, we have been able to obtain robust, realistic paths for minimal computation. The resulting approach is fast enough to be used by real systems operating with imperfect information in real environments.

6 Replanning

The approach described above allows us to deal with the uncertainty in our planning grid which is likely to be most relevant. But there are two related situations in which this approach alone will not be sufficient. The first is when an agent finds itself in a highly unlikely situation where a large number of non-pinch point cells have been found to be untraversable and, as a result, the agent is unable to traverse its computed path. The second is when the original information given to the agent is incorrect or the environment is dynamic.

However, we can deal with both of these situations by augmenting our approach to perform dynamic replanning while the agent navigates the environment, as follows.

First, the agent extracts the pinch points in the environment and computes its solution path to the goal as previously described.

Next, the agent computes which neighboring cell it should move to first, based on its neighbors' costs to each reachable face and the costs from each reachable face as returned by the solution path. Initially, this is trivial: the starting cell is part of the adjacency graph so the agent knows which face to move towards. It only needs to look for the neighbor which will take it to that face with least overall cost. However, as the robot traverses towards the desired face, it is able to update the terrain information concerning encountered cells. As a result, the face which is initially least costly may not remain so. For example, if the agent finds that a number of cells along its path to the desired face are untraversable, it might not be able to reach the face without a costly diversion. In such a situation it may be less costly to move towards a different face which was initially more expensive.

We can allow for this type of online replanning by interleaving D* and PAO*. First, we use D* to maintain cheapest cost paths to every reachable face from the agent cell. Thus, each time the agent updates the terrain of an encountered cell, the new terrain cost is propagated through these paths and the path costs to each face are updated accordingly. Given these costs (and the costs from each face returned by PAO*), the agent can update the best overall face to move towards.

However, PAO* derives its efficiency by ignoring much of the AND-OR graph and only considering promising faces. Thus, only the faces along the final solution graph are guaranteed to have their correct costs. The rest may only have admissible costs. This means we need to invoke PAO* with the current desired face each time we find we are moving towards a face which has some nonterminal leaf node in its solution graph. Thus, we take the path costs returned by D* and use these to update the arcs between the robot position and the rest of the adjacency graph. We can then use PAO* to compute the best reachable face to move towards. Because PAO* does not terminate

until the solution graph is complete, we are guaranteed that the face returned by the algorithm will have its optimal cost given the current adjacency graph information. When the agent reaches a face, it can update the adjacency graph for all altered arc costs and replan accordingly.

7 Conclusion

We have presented an algorithm for efficient planning with imperfect information. Our approach focusses computation on areas of the environment whose uncertainty is most important to the planning task. It then makes use of the recently developed PAO* algorithm to solve the subsequent planning problem. As a result, it is able to produce robust paths in a fraction of the time required by a complete solution. We have also described a replanning extension to the algorithm to cope with dynamic environments.

8 Acknowledgments

This work was sponsored in part by the U.S. Army Research Laboratory, under contract "Robotics Collaborative Technology Alliance" (contract number DAAD19-01-2-0012). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies or endorsements of the U.S. Government.

References

[1] C.L. Chang and J.R. Slagle. An admissible and optimal algorithm for searching AND-OR graphs. *Artificial Intelligence*, 2:117 – 128, 1971.

[2] D. Ferguson, A. Stentz, and S. Thrun. Pinch point planning. Technical Report CMU-RI-TR-04-06, Carnegie Mellon Robotics Institute, January 2004.

[3] L.P. Kaebling, M.L. Littman, and A.R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 1998.

[4] S. Koenig and M. Likhachev. Incremental A*. In *Advances in Neural Information Processing Systems*. MIT Press, 2002.

[5] Illah Nourbakhsh and Michael Genesereth. Assumptive planning and execution: a simple, working robot architecture. *Autonomous Robots Journal*, 3(1):49–67, 1996.

[6] E. Rich and K. Knight. *Artificial Intelligence*. McGraw-Hill, 1992.

[7] Anthony Stentz. The focussed D* algorithm for real-time replanning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 1995.