

Planning Long Dynamically-Feasible Maneuvers for Autonomous Vehicles

Maxim Likhachev
School of Computer Science
University of Pennsylvania
Philadelphia, PA
maximl@seas.upenn.edu

Dave Ferguson
Intel Research Pittsburgh
4720 Forbes Ave
Pittsburgh, PA
dave.ferguson@intel.com

Abstract—In this paper, we present an algorithm for generating complex dynamically-feasible maneuvers for autonomous vehicles traveling at high speeds over large distances. Our approach is based on performing anytime incremental search on a multi-resolution, dynamically-feasible lattice state space. The resulting planner provides real-time performance and guarantees on and control of the suboptimality of its solution. We provide theoretical properties and experimental results from an implementation on an autonomous passenger vehicle that competed in, and won, the Urban Challenge competition.

I. INTRODUCTION

Autonomous vehicles navigating through cluttered, unstructured environments or parking in parking lots often need to perform complex maneuvers and reason over large distances. Furthermore, this reasoning usually needs to be performed very quickly so that the resulting maneuvers can be executed in a timely manner, particularly if the environment is inhabited, dynamic, or dangerous. In particular, our current focus is planning for autonomous urban driving including both off-road scenarios and large unstructured parking lots such as the ones in front of malls and large stores (on the order of 200×200 meters). Maneuvering at human driving speeds (~ 15 mph) through such areas requires very efficient planning, especially if they contain static obstacles or other moving vehicles.

Roboticians have concentrated on the problem of mobile robot navigation for several decades, providing a large body of research. Early approaches concentrated on *local* planning, where very short term reasoning is performed to generate the next action for the vehicle. These include potential field-based techniques, where obstacles exert repulsive forces on the vehicle while the goal exerts an attractive force [1], and the curvature velocity [2] and dynamic window [3] approaches, where planning is performed in control space to generate dynamically-feasible actions. One major limitation of these purely local approaches was their capacity to get the vehicle stuck in local minima en route to the goal (for instance, cul-de-sacs). Further, these approaches are unable to perform complex multi-stage maneuvers, such as three-point turns, as these maneuvers are not within the set of local actions considered by the planner.

To reduce the susceptibility to local minima of these approaches, algorithms were developed that incorporated *global*

as well as local information [4, 5, 6, 7]. Typically, these approaches generate a set of candidate simple local actions and evaluate each based on both their local traversability cost and the desirability of their endpoints based on a global value function (e.g. the expected distance to the goal based on known obstacle information). Although these approaches perform better with respect to local minima, their simple local planning can still cause the vehicle to get stuck or take highly suboptimal paths. Subsequent approaches have focused on improving this local planning by using more sophisticated local action sets that better follow the global value function [8, 9], and by generating sequences of actions to perform more complex local maneuvers [10, 11, 12]. The most complex of these approaches are able to perform very precise local maneuvering but are limited by the mismatch between their powerful local planning and their approximate global planning, resulting once more in a susceptibility to local minima.

Recognizing this mismatch, other researchers have concentrated on improving the quality of global planning, so that a global path can be easily tracked by the vehicle [13, 14, 15, 16, 17]. However, the computational expense of generating complex global plans over large distances has remained very challenging, and these approaches are restricted to either small distances, fairly simple environments, or highly suboptimal solutions.

In this paper, we present an efficient, global planning approach that attempts to overcome these challenges. First, we employ a multi-resolution lattice search space to reduce the complexity of the global search while still providing extremely high-quality solutions. Second, we use an efficient anytime, incremental search to quickly generate bounded suboptimal solutions, then improve these solutions while deliberation time allows and repair them when new information is received. The resulting approach is able to plan complex, dynamically-feasible maneuvers over hundreds of meters and improve and repair them in real-time for vehicles traveling at high (~ 15 mph) speeds.

We first describe the key ideas and components of our approach, then provide key theoretical properties and results from both simulation and the Urban Challenge competition.

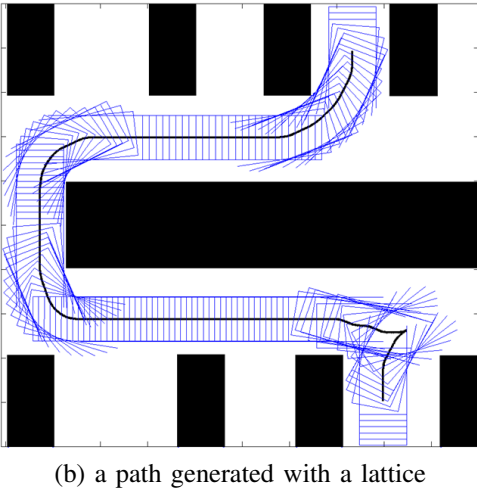
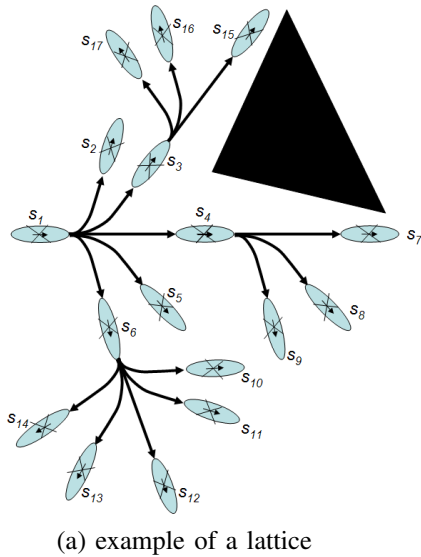


Fig. 1. (a) shows a 3D (x, y, θ) lattice with at most five forward actions for each state and no backward actions. A full set of actions is shown for state s_1 . For every state, this set of actions is translated and rotated appropriately, and all actions intersecting obstacles are removed. (b) shows an example of a complex 3D path for a vehicle with a large rectangular footprint (obstacles are shown in black).

II. MULTI-RESOLUTION LATTICE STATE-SPACE

A state lattice [18] is a discretization of the configuration space into a set of states, representing configurations, and connections between these states, where every connection represents a feasible path (see Figure 1 for an example of a lattice). As such, lattices provide a method for motion planning problems to be formulated as graph searches. However, in contrast to many graph-based representations (such as 4-connected or 8-connected grids), the feasibility requirement of lattice connections guarantees that any solutions found using a lattice will also be feasible. This makes them very well suited to planning for non-holonomic and highly-constrained robotic systems, such as passenger vehicles.

State-space. The two key considerations in constructing a lattice are the discretization (or sampling) strategy used for




Fig. 2. High- and low-resolution action spaces.

representing the states in the lattice, and the action space (or control set) used for the inter-state connections. For our application we employ a four dimensional (x, y, θ, v) state representation, where (x, y) represent the position of the center of the vehicle in the world, θ represents the orientation of the vehicle, and v represents its translational velocity. The (x, y, θ) coordinates are important for computing the validity of the poses of the vehicle in the world and making sure that no path in the lattice requires an instantaneous change in the orientation of the vehicle. For the velocity v we use two possible values: maximum forwards velocity and maximum reverse velocity. We take velocity into account because the time involved in switching between forward and backward directions is substantial so reasoning about this cost is important for generating fast, smooth paths¹.

Action Space. The action space for each state in the lattice is intended to be dense enough that every possible feasible path through the lattice can be constructed by combining sequences of these actions. However, because this action space represents the branching factor of the subsequent graph search, in practice it must be carefully constructed to provide flexibility in path selection while maintaining computational tractability.

The offline construction of our action space is based on work by Pivtoraiko and Kelly [18] that attempts to create near-minimal spanning action spaces. Given a state s , we compute the action space by first calculating a subset of states within a distance d of s that are reachable via some feasible action. To generate the feasible actions we use a trajectory generation algorithm originally developed by Howard and Kelly [9]. This algorithm employs an accurate vehicle model to produce feasible, directly-executable actions and an optimization technique to minimize the endpoint error of these actions with respect to a desired endpoint state. We use this approach to ‘snap’ the actions to the lattice so that the endpoint of each action lands on a lattice state. Next, we look at this set of actions and calculate whether any single action can be approximately recomposed out of a combination of other, shorter actions. If so, these longer actions are discarded from our set. This provides us with a compact set of actions that approximate the full reachable space. However, in contrast to the approach in [18], we maintain multiple straight segments of varying lengths to improve the speed of the subsequent search, as we

¹We do not reason about curvature (the orientation of wheels) because we found this to be less critical for the speeds we are interested in traveling at, as discussed in the results section.

will discuss in Section III-A. Figure 2(a) illustrates the action space for a single state (oriented to the right) in our lattice.

Multi-resolution Lattice. Even with a compact action space, planning long complex maneuvers over lattices can be expensive in terms of both computation and memory. An important observation, however, is that usually, there exists a wide spectrum of smooth, dynamically-feasible paths between the vehicle and goal configurations and it is waste of time and memory to explore all of them. On the other hand, all of these paths start and end at the exact same configurations, and the challenge is in finding a path that satisfies the current vehicle configuration and the specific goal configuration precisely.

This motivated us to take a novel, multi-resolution approach, where we use a high-resolution action space in the vicinity of the robot and the goal, and a low-resolution action space elsewhere. We call the resulting combination a multi-resolution lattice. With this approach, we can harness most of the benefit of the high-resolution representation without paying anything near the full computational cost. The trick is making sure that the high-resolution and low-resolution lattices connect together smoothly.

Our multi-resolution approach maintains the same dimensionality (x, y, θ, v) for both resolutions, but the action space for the low-resolution lattice is a strict subset of the action space for the high-resolution lattice. Figure 2(a) shows the action space used in the high-resolution lattice and Figure 2(b) shows the action space used in the low-resolution lattice². Using this method ensures that the low-resolution lattice is utilized fully and that paths in the multi-resolution lattice are guaranteed to be feasible, which is a strong advantage over existing combined local and global approaches for navigation.

Theorem 1: Every path in a lattice that uses only a low-resolution action space is also a valid path in our multi-resolution lattice. Further, every path in the multi-resolution lattice is a valid path in a lattice that uses only the high-resolution action space.

Proof. The proof of the first claim follows trivially from the fact that any action in the low-resolution lattice is a valid action in both the low-resolution and high-resolution lattices, and therefore is a valid action in the multi-resolution lattice. A similar argument applies for the second claim. ■

Enforcing the low-resolution action space to be a subset of the high-resolution action space decreases the *branching factor* of the graph constructed by the search, which is certainly important, but it does not necessarily decrease the *size* of the graph. However, it is also possible to decrease the size of the graph as follows. Suppose A_h is an action space used in the high-resolution space, and A_l is an action space used in the low-resolution space. Thus, $A_l \subset A_h$. Then, we can construct A_l by picking only the actions from A_h that end at states with a coarser discretization than the end states of actions in A_h . For example, we can choose for A_l only those actions whose end

states have θ equal to one of 16 possible angles, while actions in A_h can connect states with 32 possible values of θ . (This is precisely what we used in our system.) Mathematically, the construction of the action space A_l can be expressed as follows in terms of a high-resolution discretization Q_h and a lower-resolution discretization Q_l of variables x, y, θ, v : an action a connecting states $s_1 = (x_1, y_1, \theta_1, v_1)$ and $s_2 = (x_2, y_2, \theta_2, v_2)$ belongs to A_l if and only if $a \in A_h$ and $(x_2, y_2, \theta_2, v_2) \in Q_l$.

Restricting Q_l to a coarser discretization for (x, y) or θ corresponds to using a discretization that adapts based on the vehicle and goal configurations. This technique can also be used to explicitly constrain the behavior of the vehicle in the different areas. For instance, restricting Q_l to contain only positive v -values prevents the vehicle from moving backward when far from the initial and goal configurations. This general approach allows for an arbitrarily-reduced state and action space in the low-resolution portion of the lattice, and can also be trivially extended to more than two levels of resolution if desired.

III. ANYTIME, INCREMENTAL SEARCH

Given a search space (in our case, in the form of a multi-resolution lattice) and a cost function associated with each action, we need an efficient method for searching through this space for a solution path. A* search is perhaps one of the most popular methods for doing this [19]. It utilises a heuristic to focus the search towards the most promising areas of the search space. While highly efficient, A* aims to find an optimal path which may not be feasible given time constraints and the size of environments autonomous vehicles need to operate in. To cope with very limited deliberation time, anytime variants of A* search have been developed [20, 15]. These algorithms generate an initial, possibly highly-suboptimal solution very quickly and then concentrate on improving this solution while deliberation time allows. Furthermore, these anytime algorithms are able to provide bounds on the suboptimality of the solution at any point of time during the search.

A* and its anytime variants work best when the search space, and thus environment, is mostly known a priori. In robotic path planning this is rarely the case, and the robot typically receives updated environmental information through onboard and/or offboard sensors during execution. To cope with imperfect initial information and dynamic environments, efficient incremental variants of A* search have been developed that update previous solutions based on new information (e.g. from sensors) [21, 22, 23]. These algorithms repair existing solutions for a fraction of the computation required to generate such solutions from scratch.

When faced with limited deliberation time *and* imperfectly-known or dynamic environments, it is extremely useful to have a search algorithm that is both anytime and incremental. The Anytime Dynamic A* algorithm developed by Likhachev et al. is a version of A* search that combines these two properties into a single approach and has been shown to be very effective

²In practice, choosing the appropriate set can be achieved with a basic check: if the (x, y) location of a state is not within some distance d of the vehicle or goal, its action set is the low-resolution set.

for a range of robotic planning tasks [16]. We employ this algorithm for planning and re-planning paths in our multi-resolution lattice.

A. Anytime Dynamic A*

Anytime Dynamic A* (AD*) exploits a property of A* that can result in much faster generation of solutions, namely that if consistent heuristics are used and multiplied by an inflation factor $\epsilon > 1$, then A* can often generate a solution much faster than if no inflation factor is used [24], and the cost of the solution generated by A* will be at most ϵ times the cost of an optimal solution [25]. AD* operates by performing a series of these inflated A* searches with decreasing inflation factors, where each search reuses information from previous searches. By doing so, it is able to provide suboptimality bounds on all solutions generated and allows for control of these bounds, since the user can decide how much the inflation factor is decreased between searches. To cope with updated information, AD* also borrows ideas from the D* and D* Lite algorithms [21, 22] and only propagates updated information through the affected and relevant (given the current search) portions of the search space.

To enable efficient anytime planning and replanning as the vehicle moves, we use AD* to search backwards from the goal configuration towards the current configuration of the vehicle. The heuristic used thus needs to estimate the cost of a shortest path from the vehicle configuration (rather than goal) to each state in question.

The effectiveness of Anytime Dynamic A* is highly dependent on its use of an informed heuristic to focus its search. An accurate heuristic can reduce the time and memory required to generate a solution by orders of magnitude, while a poor heuristic can diminish the benefits of the algorithm. It is thus important to devote careful consideration to the heuristic used for a given search space. Further, because we are inflating heuristic values, it is useful to have long actions that can skip over several nodes and reduce the number of states in the search. It is for this reason we add several straight line actions of varying length in both the forwards and backwards directions to our action set (see Section II).

B. Informative Heuristics

The purpose of a heuristic is to improve the efficiency of the search by guiding it in promising directions. A common approach for constructing a heuristic is to use the results from a simplified search problem (e.g. from a lower-dimensional search problem where some of the original constraints have been relaxed). In selecting appropriate heuristics, it is important to analyze the original search problem and determine the key factors contributing to its complexity. In robotic path planning these are typically the complexity inherited from the constraints of the mechanism and the complexity inherited from the nature of the environment.

To cope with the complexity inherited from the mechanism constraints, a very useful general heuristic is the cost of an optimal solution through the search space assuming a




Fig. 3. (a) Inner (r) and outer (R) radii of the robot. (b) Example where the 2D heuristic function may overestimate the cost of a path derived purely from convolution.

completely empty environment. This can be computed offline and stored as a heuristic lookup table, and several efficiencies can be used to reduce the required memory for this table [17]. This is a very well informed heuristic function when operating in sparse environments and is guaranteed to be an optimistic (or admissible) approximation of the actual path cost. In obstacle-laden environments, however, this heuristic function can grossly underestimate the true costs and mislead the search into the exploration of wrong regions of the state-space.

To cope with the complexity inherited from the nature of the environment, it is not practical to pre-compute heuristic values for all possible environment configurations, as there are an effectively infinite number of possibilities for any reasonably-sized environment. However, in this case it is beneficial to solve online a simplified search problem given the actual environment and use the result of this search as a heuristic to guide the original, complex search. In particular, we solve a 2D $((x, y))$ version of the problem by running a single Dijkstra's search starting at the cell that corresponds to the center of the current vehicle position. The search computes the costs of shortest paths from the cell the robot is in to all other cells in the environment. This search is therefore re-run every time the vehicle pose is changed. Even though the search is very fast, we still restrict this search to only compute the states that are no more than twice as far (in terms of path cost) from the vehicle cell as the goal cell. In other words, the search is terminated whenever the cost of the cell it is about to compute is larger than or equal to the cost of the goal cell. To make the heuristic function more informative, the cost of each cell in the 2D grid used for computing this 2D heuristic is set to infinity whenever the distance from it to the closest obstacle is less than the inner radius of the robot (shown in Figure 3(a)). Distances from cells to obstacles can be efficiently computed using a distance transform (see Section IV-A).

AD* requires the heuristics to be admissible and consistent. This holds if $h(s_{start}) = 0$ and for every pair of states s, s' such that s' is an end state of a single action executed at state s , $h(s) + c(s, s') \geq h(s')$, where $h(s)$ is a heuristic of state s , s_{start} is a state that corresponds to the vehicle configuration and $c(s, s')$ is the cost of the action that connects s to s' .




Fig. 4. Mechanism-constrained (solid) and environment-constrained (dashed) heuristic paths. In each case, the initial and desired vehicle poses are shown as blue and red rectangles, respectively (with the interior triangles specifying the headings). (a) The mechanism-constrained heuristic is perfectly informed when no obstacles are present in the environment. (b) The environment-constrained 2D heuristic can provide significant benefit when obstacles exist. Here, an obstacle (shown in black) resides over the direct path to the desired pose.

The cost $c(s, s')$ of the action is typically computed as the length of the action times the average of the costs of the cells covered by the vehicle when moving from state s to state s' . The heuristic based on the 2D search, however, may overestimate these costs since it estimates the cost of moving the center of the vehicle only. To demonstrate this, imagine a path that involves the vehicle moving through a narrow corridor with a high-cost strip going exactly along the center of this corridor (Figure 3(b)). The cost of the 2D path from the initial coordinates of the vehicle to the goal coordinates corresponds to the summation of the costs of the transitions going along the high-cost strip. Cells on either side of the strip are impassable (have infinite cost) since they lie closer to the obstacles than the inner radius of the vehicle, and therefore the center of the vehicle can not reside in any of them. The cost of the actual path, on the other hand, is lower than the cost of the path along the high-cost strip because the cost of each actual action is computed as an average of the cells covered by the vehicle.

To resolve this inadmissibility of the heuristic function, the cost of each transition $c(s, s')$ is computed as the length of the transition times the maximum of two quantities: (a) the average value of the costs of the cells covered by the vehicle when moving from state s to state s' (same as before), and (b) the maximum of the 2D grid cell costs, used to compute heuristics, traversed through by the center of the vehicle when moving from s to s' . Intuitively, this cost function penalizes slightly more those actions for which the vehicle traverses high-cost areas (e.g., obstacles) that reside right under its center. In addition, the heuristics are scaled down by a factor of 1.08 to compensate for the suboptimality of optimal paths in 8-connected grids. It can be then shown that our 2D heuristic function is admissible and consistent with respect to this cost function.

Each of these heuristic functions, mechanism-relative and environment-relative, have strong and complementary benefits (see Figure 4). Rather than selecting one, it is possible to combine the two. We do this by constructing a new

heuristic that, for each state s , returns the value $h(s) = \max(h_{fsh}(s), h_{2D}(s))$, where $h_{fsh}(s)$ is the heuristic value of state s according to the mechanism-constrained heuristic (freespace heuristic), and $h_{2D}(s)$ is the value according to the environment-constrained heuristic (2D heuristic). As shown in the experimental results, this combined heuristic function can be an order of magnitude more effective than either of the component heuristic functions. Since both $h_{fsh}(s)$ and $h_{2D}(s)$ are admissible and consistent, the combined heuristic is also admissible and consistent [26]. This property satisfies the conditions required by AD* to guarantee ϵ -bound on the suboptimality of the paths it returns (see Theorem 1 in [16]):

Theorem 2: The cost of a path returned by Anytime Dynamic A* is no more than ϵ times the cost of a least-cost path from the vehicle configuration to the goal configuration using actions in the multi-resolution lattice, where ϵ is the current value by which Anytime Dynamic A* inflates heuristics.

IV. OPTIMIZATIONS

A. Efficient Convolution

Typically, one of the most computationally expensive parts of planning for vehicles is computing the cost of actions, as this involves convolving the geometric footprint of the vehicle for a given action with a map from perception. In our application, we used a 0.25m resolution 2D perception map and the (x, y) dimensions of our vehicle were $5.5m \times 2.25m$. Thus, even a short 1m action requires collision checking roughly 300 cells. Further, the specific cells need to be calculated based on the action and the initial pose of the vehicle.

To reduce the processing required for this convolution, we performed two optimization steps. First, for each action a we pre-computed the cells covered by the vehicle when executing this action. During online planning, these cells are quickly extracted and translated to the appropriate position when needed. Second, we generated two configuration space maps to be used by the planner to avoid performing convolutions. The first of these maps expanded all obstacles in the perception map by the inner radius (see Figure 3(a)) of the robot; this map corresponded to an optimistic approximation of the actual configuration space. Given a specific action a , if any of the cells through which the center of the robot executing action a passes are obstacles in this inner map, then a is guaranteed to collide with an obstacle. The second map expanded all obstacles in the perception map by the outer radius (see Figure 3(a)) of the robot and therefore corresponded to a pessimistic approximation of the configuration space. If all of the cells through which the center of the vehicle passes when executing action a are obstacle-free in this map, then a is guaranteed to be collision-free. Only those actions that do not produce a conclusive result from these simple tests need to be convolved with the perception map. Typically, this is a severely reduced percentage, thus saving considerable computation. To create these auxiliary maps efficiently, we performed a single distance transform on the perception map and then thresholded the distances using the corresponding radii of the robot for each map.

With incremental planning algorithms such as AD*, when changes are observed in the cost map, they must be propagated through the relevant portions of the search space. However, detecting which actions and states in the search space are directly affected by these changes in the cost map can be expensive. For example, if the cost of the cell (x_c, y_c) changes, then the costs of all actions that involve vehicle going over that cell may change. Typically, there could be thousands of such actions. AD* needs to iterate and update the values of all the states $((x, y, \theta, v)$ poses) from which these actions can be executed. Given the large number of affected actions, this iteration can be very expensive. However, AD* needs to update the values of only those states that have actually been computed in the previous planning iterations. If a state has not been computed, then there is no need to update its value. We exploit this property to decrease the amount of computational efforts involved in iterating over the states that may possibly be affected by changes in the cost map. In the following we explain this optimization.

First, we pre-compute offline all the states that have actions whose costs depend on the cost of the cell $(0, 0)$. These states are grouped into mutually disjoint sets, where each i^{th} set $\mathfrak{R}^{x_i \dots x_i+d, y_i \dots y_i+d}$ contains all those states (x, y, θ, v) , whose $x_i \leq x < x_i + d$ and $y_i \leq y < y_i + d$, where d is a (small) positive integer. We used $d = 5$. In other words, all the states whose values need to be updated by AD* whenever the cost of the cell $(0, 0)$ is modified are pre-computed and stored in a low-resolution grid map. Let us denote this map by \mathfrak{R} . Each cell in this low-resolution grid map is d times wider and d times longer than a cell in the perception map.

Second, during online operations, we maintain a low-resolution replanning map of the same discretization as \mathfrak{R} . The value of each cell in this replanning map is true whenever at least one state whose (x, y) coordinates fall into this cell has been generated (computed) by AD*. Thus, while planning, whenever AD* generates (computes a value of) a state (x, y, θ, v) , then it also sets the corresponding cell in the replanning map to be true.

Finally, whenever the cost of a cell (x_c, y_c) is modified, for each non-empty cell $\mathfrak{R}^{x_i \dots x_i+d, y_i \dots y_i+d}$ in \mathfrak{R} we look up if any one of the following four cells in the replanning map are set to true:

$$\begin{aligned} &(((x_i + x_c) \bmod d), ((y_i + y_c) \bmod d)) \\ &(((x_i + x_c) \bmod d) + 1, ((y_i + y_c) \bmod d)) \\ &(((x_i + x_c) \bmod d), ((y_i + y_c) \bmod d) + 1) \\ &(((x_i + x_c) \bmod d) + 1, ((y_i + y_c) \bmod d) + 1) \end{aligned}$$

If so, then we make AD* to update the value of every state stored in $\mathfrak{R}^{x_i \dots x_i+d, y_i \dots y_i+d}$ translated by (x_c, y_c) . No other states need to be updated since it is guaranteed that they have not been generated by AD* previously. This optimization can save a tremendous amount of replanning efforts.

We have implemented our approach on an autonomous passenger vehicle (lower-left image in Figure 5) where it has been used to drive over 3000 kilometers in urban environments, including competing in the DARPA Urban Challenge. The multi-resolution lattice planner was used for planning through parking lots and into parking spots, as well as for geometric road following in off-road areas, and in error recovery scenarios. During these scenarios, the vehicle traveled speeds of up to 15 miles per hour while performing complex maneuvers and avoiding static and dynamic obstacles.

In all cases, the multi-resolution lattice planner searches backwards out from the goal pose (or set of goal poses) and generates a path consisting of a sequence of feasible high-fidelity maneuvers that are collision-free with respect to the static obstacles observed in the environment. This path is also biased away using cost function from undesirable areas such as curbs and locations in the vicinity of dynamic obstacles.

When new information concerning the environment is received (for instance, a new static or dynamic obstacle is observed), the planner is able to incrementally repair its existing solution to account for the new information. This repair process is expedited by performing the search in a backwards direction, as in such a scenario updated information in the vicinity of the vehicle affects a smaller portion of the search space and so less repair is required. The lattice plan is typically updated once per second, however in trivial or very difficult scenarios this time may vary.

As mentioned earlier, the lattice used in this application does not explicitly represent curvature. Theoretically, this means that the paths produced over this lattice are guaranteed feasible only if we allow the vehicle to stop at each lattice state and re-orient its steering wheel. However, in practice we reduce (by a small fraction) the maximum curvature used in generating connections between states and we reduce the maximum speed at which we execute higher-curvature sections of lattice paths (from 5 m/s down to 2 m/s) so that this curvature discontinuity is not a critical issue. We also use a lookahead during execution to slow down and stop when switching velocity directions³. As a result, we don't need to stop during execution unless the path contains velocity sign changes.

The lattice path is tracked using a local planner that employs the same trajectory generation algorithm used to provide the action space for the lattice. Although a simple, single-trajectory tracker would suffice given the feasibility of the lattice plan, multiple trajectories are produced to account for dynamic obstacles and new observations that could require immediate reaction (the local planner runs at 10 Hz). The trajectories generated by the trajectory generation algorithm are shown as short multi-color trajectories in the top row images in Figure 5. The local planner always chooses the

³A maximum lookahead of $2m$ is required given our vehicle's maximum deceleration and the top speed used for following lattice paths, but we use a slightly higher lookahead for smooth deceleration.




Fig. 5. Replanning when new information is received

trajectory that stays as close to the path returned by the lattice-based planner as possible while avoiding obstacles.

To ensure that a high-quality path is available for the vehicle as soon as it enters a parking lot, the lattice planner begins planning for the desired goal pose while the vehicle is still approaching the lot. By planning a path from the entry point of the parking lot in advance, the vehicle can seamlessly transition into the lot without needing to stop, even for very large and complex lots. Further, the anytime property of the search enables the solution to be improved during the pre-planning stage and, depending on how much time is available for pre-planning, the resulting path for the vehicle can converge to a (provably) optimal solution.

As well as providing smooth navigation amongst partially-known static objects, the efficiency of the multi-resolution lattice planner makes it possible to intelligently interact with several dynamic obstacles in the environment. In our application, we were able to not only avoid such obstacles but through updating regions of high cost as the obstacles moved, we could stay well clear of them unless necessary and also exhibit intelligent yielding behavior in unstructured areas (e.g. keeping to the right when approaching oncoming vehicles).

The multi-resolution lattice planner was also used for performing complex maneuvers in error recovery scenarios during on-road driving, such as when a lane or intersection is partially blocked with vehicles or obstacles, or a road is fully blocked and a u-turn is required. It was also used when there was some uncertainty as to where the road was; in these scenarios it uses the geometric perceptual information to bias the vehicle towards the center of the road (when there are perceivable curbs or berms).

We have included here a number of examples from the Urban Challenge and our testing to illustrate key characteristics of the approach.

a) Pre-planning: Figure 6 illustrates the pre-planning used by the lattice planner, as well as its anytime performance. The left image shows our vehicle approaching a parking lot (parking lot boundary shown in green, road lanes shown in


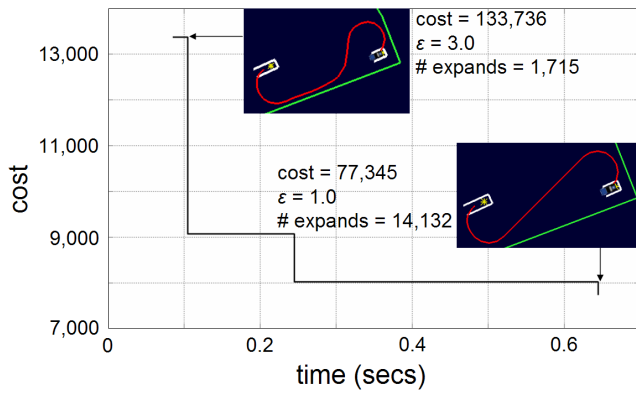




Fig. 6. Pre-planning a path into a parking spot and improving this path in an anytime fashion.

blue), with its intended parking spot indicated by the white triangle. While the vehicle is still outside the parking lot it begins planning a path from one of the parking lot entries to the desired spot (path shown in red). Although the initial path shown in this left image is feasible, it is not ideal as it involves more turning than necessary. The right image shows how this path is improved over time as the vehicle approaches. This path is optimal with respect to our cost function and is generated well before the vehicle enters the parking lot.

b) Anytime Planning: Figure 7 is intended to provide insights into the benefits provided by each of the main components of our approach. Figure 7(a) illustrates the anytime behavior of the approach when planning between two parking spots. We have included a plot of the cost of the solution produced by AD* as a function of computation time. Here, the initial suboptimality bound ϵ was set to 3. The upper image shows the first path AD* finds. This path was found in less than 100 msecs (and after 1, 715 state expansions). The cost of the path was 133, 736. Given additional deliberation time, AD* improves upon this solution, and after 650 msecs, the search converges to an optimal solution. This solution is significantly



(a) anytime behavior



lattice	states expanded	time (secs)
high-res	2,933	0.19
multi-res	1,228	0.06

(b) effect of multi-res lattice

heuristic	states expanded	time (secs)
h	2,019	0.06
h_{2D}	26,108	1.30
h_{fsh}	124,794	3.49

(c) effect of heuristic

Fig. 7. An example highlighting anytime behavior of our approach and the benefits of the multi-resolution lattice and the combined heuristic function.

shorter than the initial path (as seen in the bottom image) and has a cost of 77,345.

c) Multi-resolution Planning: Figure 7(b) shows the benefits of using our multi-resolution lattice approach on the same simple example. The top row in the table represents a uniform high-resolution lattice, while the bottom row represents our multi-resolution lattice (in both cases, $\epsilon = 2$). Planning with the multi-resolution lattice is more than three times faster. Note that the improvement in states expanded is less than a factor of three. This is because using a multi-resolution lattice decreases not only the number of states expanded but also the time spent expanding each state, since the number of possible actions from each state is decreased.

Figure 8 shows the comparison of planning with multi-resolution lattice versus planning with uniform high-resolution lattice in more details. In particular, the table shows the time spent by the algorithm as it decreases its suboptimality bound ϵ . (The time is shown as cumulative.) The table also gives the solution cost generated by the planner in each of the cases (second column in each table). As expected, the data shows that planning with multi-resolution lattice is about three times faster than planning with high-resolution lattice both in producing the first solution as well as in time it takes to converge to a provably optimal solution ($\epsilon = 1.0$). The solution quality, on the other hand, at the time of convergence

ϵ	sol. cost	# of expands	cumul. time (secs)	ϵ	sol. cost	# of expands	cumul. time (secs)
2.0	83850	1228	0.062	2.0	102462	2933	0.191
1.9	83850	0	0.063	1.9	102462	157	0.205
1.8	83850	0	0.064	1.8	102462	269	0.225
1.7	83850	0	0.065	1.7	102462	757	0.264
1.6	83850	0	0.065	1.6	83840	1839	0.375
1.5	83850	0	0.066	1.5	83840	0	0.377
1.4	83850	0	0.066	1.4	83840	0	0.378
1.3	83850	0	0.067	1.3	83840	0	0.379
1.2	83850	0	0.068	1.2	83840	0	0.381
1.1	78168	2937	0.221	1.1	77416	5027	0.786
1.0	77345	10034	0.534	1.0	76740	16912	1.618

(a) multi-res. lattice planning (b) high-res. lattice planning

Fig. 8. The performances of planning with multi-resolution lattice and planning with uniform high-resolution lattice on the example shown in Figure 7(b)

is roughly the same (just slightly better for the planner that uses a high-resolution lattice).

d) Combining Mechanism-relative and Environment-relative Heuristics: Figure 7(c) demonstrates the benefits of using our combined heuristic function on a simple example. The first row in the table represents our combined heuristic function. It combines the 2D environment-constrained heuristic (2nd row) and freespace mechanism-constrained heuristic (3rd row). Using this combination is over 21 times faster than using the 2D heuristic alone and over 58 times faster than using the freespace heuristic alone.

e) Replanning: Figure 5 illustrates the replanning capability of the lattice planner. These images were taken from a parking task performed during the National Qualification Event. The top-left image shows the initial path planned for the vehicle to enter the parking spot indicated by the white triangle. Several of the other spots were occupied by other vehicles (shown as rectangles of varying colors), with detected obstacles shown as red areas. The trajectories generated to follow the path are shown emanating from our vehicle (the selected trajectory is shown in blue). As the vehicle gets closer to its intended spot, it observes more of the vehicle parked in the right-most parking spot (top, second image from left). At this point, it realizes its current path is infeasible and replans a new path that has the vehicle perform a loop and pull in smoothly. This path was favored in terms of time over stopping and backing up to re-position. The three right-most photographs on the bottom row were taken by an onboard camera during the run.

f) Long-range Planning: As with other teams participating in the Urban Challenge, our vehicle underwent extensive testing before and during the competition. During the competition, the planner was able to continuously plan and replan without having the vehicle ever stop to wait for a plan. The scenarios we used for testing before the competition were numerous and included expansive obstacle-laden parking lots

as well as narrow, highly-constrained parking lots. An example of the former is shown in Figure 9(a-b). This parking lot is 200m by 200m. Initially, it is unknown and as the robot traverses the lot, it discovers a series of obstacles (shown as white dots in the image on the right). The robot has to replan in real-time to account for these obstacles. The time for replanning in this scenario varied from a few milliseconds for small re-planning adjustments to the path to a few seconds for finding drastically different trajectories, such as the one shown in Figure 9(b).

g) *Complex Maneuvering*: An example of a testing scenario involving a highly-constrained parking lot is shown in Figure 9(c-d). The trajectory planned involves the robot making an initial narrow U-turn and then making another one immediately before pulling into the final parking spot. While executing the trajectory, the robot discovers a series of obstacles and has to re-plan as shown in Figure 9(d). The new trajectory now requires the robot to backup a number of times. Moreover, it requires the robot to enter the desired spot in reverse since the discovered obstacles prohibit the robot from pulling in.

h) *Coping with Dynamic Obstacles*: Figure 10 shows the lattice planner being used to plan amongst several other moving vehicles in simulation. In these images, the current goal is shown as the white triangle and the inferred short-term trajectories of the other vehicles are included as fading polygons.

i) *Coping with Static Obstacles*: Figure 11 provides an example testing scenario for our physical vehicle. The left image shows the layout of the parking lot, the static obstacles (initially unknown to the vehicle), and the parking spots to be visited in order (1 through 5). The vehicle entered the lot through the left entrance between spots 3 and 4. The other images show snapshots from an onboard camera during the vehicle's traverse through this difficult environment.

VI. CONCLUSIONS

We have presented a general approach for complex planning involving large, high-dimensional search spaces. Our approach employs a novel multi-resolution action and state space that significantly reduces complexity while providing a seamless interface between the resolutions, as well as guarantees of solution feasibility. The approach also relies on an anytime, incremental search algorithm for generating solutions in partially-known or dynamic environments when deliberation time is limited. This search exploits a low-dimensional environment-dependent heuristic coupled with a full-dimensional freespace heuristic for efficient focusing, a powerful technique applicable to many high-dimensional planning problems. The resulting approach provides global, feasible solutions to challenging navigation tasks, and all the core techniques presented are applicable to a wide range of complex planning problems.

In the future, it is important to investigate several directions in which our approach to planning may possibly be extended. Perhaps the most pressing direction would be to research an

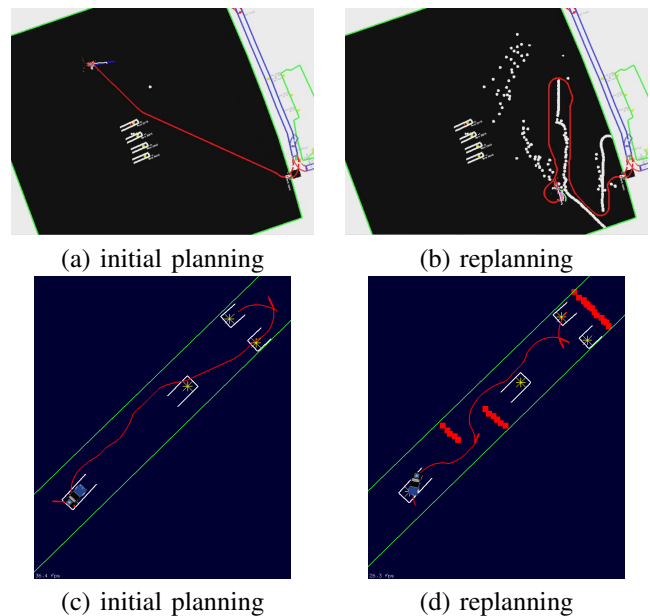


Fig. 9. Planning and replanning in large (a,b) and highly-constrained (c,d) environments

automatic way of building a multi-resolution lattice. In our current implementation, we use high-resolution lattices around the robot pose and around its goal. In some domains however, it may be required to use high-resolution lattices in additional areas. One example of this is planning for relatively large robots operating indoors. The size of these robots often does not leave much freedom for how to traverse through doorways. As a result, the areas around the doorways would have to be modeled with high-resolution lattices. In these domains, it is therefore important to explore methods for building multi-resolution lattices that change resolutions adaptively based on some properties of the local map (for instance, obstacle density).

Another challenge that our approach would not be able to deal well with is operating in environments full of dynamic obstacles (e.g., lots of cars and pedestrians). In such environments, it is necessary to plan collision-free time-parameterized trajectories, as opposed to planning purely spatial trajectories that are produced by our current approach. In addition, estimating the trajectories of dynamic obstacles accurately is very difficult, and it is important to plan collision-free trajectories by taking into account the uncertainty in these estimates. Researching how to extend multi-resolution lattice-based planning to highly-dynamic environments is therefore another topic of future research. Finally, we are currently working on applying our approach to other robots, including several different indoor ground robots and large outdoor unmanned vehicles.

REFERENCES

- [1] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *International Journal of Robotics Research*, vol. 5, no. 1, pp. 90–98, 1986.
- [2] R. Simmons, "The curvature velocity method for local obstacle avoidance," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1996.




Fig. 10. Planning amongst moving obstacles



Fig. 11. Planning in complex obstacle environments

- [3] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance." *IEEE Robotics and Automation*, vol. 4, no. 1, 1997.
- [4] S. Thrun *et al.*, "Map learning and high-speed navigation in RHINO," in *AI-based Mobile Robots: Case Studies of Successful Robot Systems*, D. Kortenkamp, R. Bonasso, and R. Murphy, Eds. MIT Press, 1998.
- [5] O. Brock and O. Khatib, "High-speed navigation using the global dynamic window approach," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1999.
- [6] A. Kelly, "An intelligent predictive control approach to the high speed cross country autonomous navigation problem," Ph.D. dissertation, Carnegie Mellon University, 1995.
- [7] R. Philippsen and R. Siegwart, "Smooth and efficient obstacle avoidance for a tour guide robot," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2003.
- [8] S. Thrun *et al.*, "Stanley: The robot that won the DARPA Grand Challenge," *Journal of Field Robotics*, vol. 23, no. 9, pp. 661–692, August 2006.
- [9] T. Howard and A. Kelly, "Optimal rough terrain trajectory generation for wheeled mobile robots," *International Journal of Robotics Research*, vol. 26, no. 2, pp. 141–166, 2007.
- [10] C. Stachniss and W. Burgard, "An integrated approach to goal-directed obstacle avoidance under dynamic constraints for dynamic environments," in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2002.
- [11] C. Urmson *et al.*, "A robust approach to high-speed navigation for unrehearsed desert terrain," *Journal of Field Robotics*, vol. 23, no. 8, pp. 467–508, August 2006.
- [12] D. Braid, A. Broggi, and G. Schmiedel, "The TerraMax autonomous vehicle," *Journal of Field Robotics*, vol. 23, no. 9, pp. 693–708, August 2006.
- [13] S. LaValle and J. Kuffner, "Rapidly-exploring Random Trees: Progress and prospects," *Algorithmic and Computational Robotics: New Directions*, pp. 293–308, 2001.
- [14] G. Song and N. Amato, "Randomized motion planning for car-like robots with C-PRM," in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2001.
- [15] M. Likhachev, G. Gordon, and S. Thrun, "ARA*: Anytime A* with provable bounds on sub-optimality," in *Advances in Neural Information Processing Systems*. MIT Press, 2003.
- [16] M. Likhachev, D. Ferguson, G. Gordon, A. Stentz, and S. Thrun, "Anytime Dynamic A*: An Anytime, Replanning Algorithm," in *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, 2005.
- [17] R. Knepper and A. Kelly, "High performance state lattice planning using heuristic look-up tables," in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2006.
- [18] M. Pivtoraiko and A. Kelly, "Generating near minimal spanning control sets for constrained motion planning in discrete state spaces," in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2005.
- [19] N. Nilsson, *Principles of Artificial Intelligence*. Tioga Publishing Company, 1980.
- [20] R. Zhou and E. Hansen, "Multiple sequence alignment using A*," in *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2002, Student abstract.
- [21] A. Stentz, "The Focussed D* Algorithm for Real-Time Replanning," in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 1995.
- [22] S. Koenig and M. Likhachev, "Improved fast replanning for robot navigation in unknown terrain," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2002.
- [23] M. Barbehenn and S. Hutchinson, "Efficient search and hierarchical motion planning by dynamically maintaining single-source shortest path trees," *IEEE Transactions on Robotics and Automation*, vol. 11, no. 2, pp. 198–214, 1995.
- [24] J. G. Gaschnig, "Performance measurement and analysis of certain search algorithms," Ph.D. dissertation, Carnegie Mellon University, 1979.
- [25] H. W. Davis, A. Bramanti-Gregor, and J. Wang, "The advantages of using depth and breadth components in heuristic search," in *Methodologies for Intelligent Systems*, 3, Z. W. Ras and L. Saitta, Eds. New York: North-Holland, 1988, pp. 19–28.
- [26] J. Pearl, *Heuristics: Intelligent Search Strategies for Computer Problem Solving*. Addison-Wesley, 1984.