

# Lazy Receding Horizon A\* for Efficient Path Planning in Graphs with Expensive-to-Evaluate Edges

**Aditya Mandalika**  
University of Washington  
adityavk@cs.uw.edu \*

**Oren Salzman**  
Carnegie Mellon University  
osalzman@andrew.cmu.edu \*

**Siddhartha Srinivasa**  
University of Washington  
siddh@cs.uw.edu \*

## Abstract

Motion-planning problems, such as manipulation in cluttered environments, often require a collision-free shortest path to be computed quickly given a roadmap graph  $\mathcal{G}$ . Typically, the computational cost of evaluating whether an edge of  $\mathcal{G}$  is collision-free dominates the running time of search algorithms. Algorithms such as Lazy Weighted A\* (LWA\*) and LazySP have been proposed to reduce the number of edge evaluations by employing a *lazy lookahead* (one-step lookahead and infinite-step lookahead, respectively). However, this comes at the expense of additional graph operations: the larger the lookahead, the more the graph operations that are typically required. We propose Lazy Receding-Horizon A\* (LRA\*) to minimize the *total planning time* by balancing edge evaluations and graph operations. Endowed with a lazy lookahead, LRA\* represents a family of lazy shortest-path graph-search algorithms that generalizes LWA\* and LazySP. We analyze the theoretic properties of LRA\* and demonstrate empirically that, in many cases, to minimize the total planning time, the algorithm requires an intermediate lazy lookahead. Namely, using an intermediate lazy lookahead, our algorithm outperforms both LWA\* and LazySP. These experiments span simulated random worlds in  $\mathbb{R}^2$  and  $\mathbb{R}^4$ , and manipulation problems using a 7-DOF manipulator.

## 1 Introduction

Robotic motion-planning has been widely studied in the last few decades. Since the problem is computationally hard (Reif 1979; Sharir 2004), a common approach is to apply sampling-based algorithms which typically construct a graph where vertices represent robot configurations and edges represent potential movements of the robot (Choset et al. 2005; LaValle 2006). A shortest-path algorithm is then run to compute a path between two vertices on the graph.

There are numerous shortest-path algorithms, each suitable for a particular problem domain based on the computational efficiency of the algorithm. For example, A\* (Hart, Nilsson, and Raphael 1968) is optimal with respect to node expansions, and planning techniques such as partial expansions (Yoshizumi, Miura, and Ishida 2000) and iterative

deepening (Korf 1985) are well-suited for problems with large graphs and large branching factors.

However, in most robotic motion-planning problems, path validations and edge evaluations are the major source of computational cost (LaValle 2006). Our work addresses these problems of quickly producing collision-free optimal paths, when the cost of evaluating an edge for collision is a computational bottleneck in the planning process.

A common technique to reduce the computational cost of edge evaluation and consequently the planning time is to employ a *lazy* approach. Two notable search-based planners that follow this paradigm are Lazy Weighted A\* (LWA\*) (Cohen, Phillips, and Likhachev 2014) and LazySP (Dellin and Srinivasa 2016; Haghtalab et al. 2017).

Both LWA\* and LazySP assume there exists a lower bound on the weight of an edge that is efficient to compute. This lower bound is used as a lookahead (formally defined in Section 4) to guide the search without having to explicitly evaluate edges unless necessary. LazySP uses an infinite-step lookahead which can be shown to minimize the number of edge evaluations but requires a large number of graph operations (node expansions, updating the shortest-path tree, etc.). On the other hand, LWA\* uses a one-step lookahead which may result in a larger number of edge evaluations compared to LazySP but with much fewer graph operations.

Our key insight is that there should exist an optimal lookahead for a given environment, which balances the time for edge evaluations and graph operations, and minimizes the *total planning time*. We make the following contributions:

1. We present Lazy Receding-Horizon A\* (LRA\*), a family of lazy shortest-path algorithms parametrized by a *lazy lookahead* (Sections 4 and 5) which allows us to continuously interpolate between LWA\* and LazySP, balancing edge evaluations and graph operations.
2. We analyze the theoretic properties of LRA\* (Section 6). Part of our analysis proves that LazySP is optimal with respect to minimizing edge evaluations thus closing a theoretic gap left open in LazySP (Dellin and Srinivasa 2016).
3. We demonstrate in Section 7, the efficacy of our algorithm on a range of planning problems for simulated  $\mathbb{R}^n$  worlds and robot manipulators. We show that LRA\* outperforms both LWA\* and LazySP by minimizing not just edge evaluations or graph operations but the *total planning time*.

\*This work was (partially) funded by the National Science Foundation IIS (#1409003), and the Office of Naval Research. Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

## 2 Related Work

A large number of motion-planning algorithms consist of (i) constructing a graph, or a roadmap, embedded in the configuration space and (ii) finding the shortest path in this graph. The graph can be constructed in a preprocessing stage (Kavraki et al. 1996; Karaman and Frazzoli 2011) or vertices and edges can be added in an incremental fashion (Gammell, Srinivasa, and Barfoot 2015; Salzman and Halperin 2015).

In domains where edge evaluations are expensive and dominate the planning time, a *lazy approach* is often employed (Bohlin and Kavraki 2000; Hauser 2015) wherein the graph is constructed *without* testing if edges are collision-free. Instead, the search algorithm used on this graph is expected to evaluate only a subset of the edges in the roadmap and hence save computation time. While standard search algorithms such as Dijkstra’s (Dijkstra 1959) and A\* (Hart, Nilsson, and Raphael 1968) can be used, specific search algorithms (Cohen, Phillips, and Likhachev 2014; Dellin and Srinivasa 2016) were designed for exactly such problems. They aim to further reduce the number of edge evaluations and thereby the planning time.

Alternative algorithms that reduce the number of edge evaluations have been studied. One approach was by foregoing optimality and computing near-optimal paths (Salzman and Halperin 2016; Dobson and Bekris 2014). Another approach was re-using information obtained from previous edge evaluations (Bialkowski et al. 2016; Choudhury, Dellin, and Srinivasa 2016; Choudhury et al. 2017).

In this paper, we propose an algorithm that makes use of a *lazy lookahead* to guide the search and minimize the total planning time. It is worth noting that the idea of a lookahead has previously been used in algorithms such as RTA\*, LRTA\* (Korf 1990) and LSS-LRTA\* (Koenig and Sun 2009). However, these algorithms use the lookahead in a different context by interleaving planning with execution before the shortest path to the goal has been completely computed. Using a lazy lookahead has also been considered in the control literature (Kwon and Han 2006). Receding horizon optimization can be summarized as iteratively solving an optimal-control problem over a fixed future interval. Only the first step in the resulting optimal control sequence is executed and the process is repeated after measuring the state that was reached. Our lazy lookahead is analogous to the fixed horizon used by these algorithms. Additionally, the lazy lookahead can also be seen as a threshold that defines the extent to which (lazy) search is performed. This is similar to the Iterative Deepening version of A\* (IDA\*) (Korf 1985) which performs a series of depth-first searches up to a (increasing) threshold over the solution cost.

## 3 Algorithmic Background

### 3.1 Single Source Shortest Path (SSSP) Problem

Given a directed graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with a cost function  $w : \mathcal{E} \rightarrow \mathbb{R}^+$  on its edges, the Single Source Shortest Path (SSSP) problem is to find a path of minimum cost between two given vertices  $v_{\text{source}}$  and  $v_{\text{target}}$ . Here, a path  $P = (v_1, \dots, v_k)$  on the graph is a sequence of vertices where  $\forall i, (v_i, v_{i+1}) \in \mathcal{E}$ . An

edge  $e = (u, v)$  belongs to a path if  $\exists i$  s.t.  $u = v_i, v = v_{i+1}$ . The *cost* of a path is the sum of the weights of the edges along the path:

$$w(P) = \sum_{e \in P} w(e).$$

### 3.2 Solving the SSSP Problem

To solve the SSSP problem, algorithms such as Dijkstra (Dijkstra 1959) compute the shortest path by building a shortest-path tree  $\mathcal{T}$  rooted at  $v_{\text{source}}$  and terminate once  $v_{\text{target}}$  is reached. This is done by maintaining a minimal-cost priority queue  $Q$  of nodes called the OPEN list. Each node  $\tau_u$  is associated with a vertex  $u$  and a pointer to  $u$ ’s parent in  $\mathcal{T}$ . The nodes are ordered in  $Q$  according to their cost-to-come i.e., the cost to reach  $u$  from  $v_{\text{source}}$  in  $\mathcal{T}$ .

The algorithm begins with  $\tau_{v_{\text{source}}}$  (associated with  $v_{\text{source}}$ ) in  $Q$  with a cost-to-come of 0. All other nodes are initialized with a cost-to-come of  $\infty$ . At each iteration, the node  $\tau_u$  with the minimal cost-to-come is removed from  $Q$  and *expanded*, wherein the algorithm considers each of  $u$ ’s neighbours  $v$ , and evaluates if the path to reach  $v$  through  $u$  is cheaper than  $v$ ’s current cost-to-come. If so, then  $\tau_v$ ’s parent is set to be  $\tau_u$  (an operation we refer to as “rewiring”) and is inserted into  $Q$ .

The search, i.e., the growth of  $\mathcal{T}$ , can be biased towards  $v_{\text{target}}$  using a heuristic function  $h : \mathcal{V} \rightarrow \mathbb{R}$  which estimates the cost-to-go, the cost to reach  $v_{\text{target}}$  from vertex  $v \in \mathcal{V}$ . It can be shown that under mild conditions on  $h$ , if  $Q$  is ordered according to the sum of the cost-to-come and the estimated cost-to-go, this algorithm, called A\*, *expands* fewer nodes than any other search algorithm with the same heuristic (Dechter and Pearl 1983; Pearl 1984).

A key observation in the described approach for Dijkstra or A\* is that for *every* node in  $Q$ , the algorithm computed the cost  $w(e)$  of the edge to reach this node from its current parent, a process we will refer to as *evaluating* an edge. Edge evaluation occurs for all edges leading to nodes in the OPEN list  $Q$  irrespective of whether there exists a better parent to the node or whether the node will subsequently be expanded for search. In problem domains where computing the weight of an edge is an expensive operation, such as in robotic applications, this is highly inefficient. To alleviate this problem, we can apply *lazy* approaches for edge evaluations that can dramatically reduce the number of edges evaluated.

### 3.3 Computing SSSP via Lazy Computation

In problem domains where computing  $w(e)$  is expensive, we assume the existence of a function  $\hat{w} : \mathcal{E} \rightarrow \mathbb{R}^+$  which (i) is efficient to compute and (ii) provides a lower bound on the true cost of an edge i.e.,  $\forall e \in \mathcal{E}, \hat{w}(e) \leq w(e)$ . We call  $\hat{w}$  a *lazy estimate* of  $w$ .

Given such a function, Cohen, Phillips, and Likhachev proposed LWA\* which modifies A\* as follows: Each edge  $(u, v)$  is evaluated, i.e.,  $w(u, v)$  is computed, only when the algorithm believes that  $\tau_v$  should be the next node to be expanded. Specifically, each node in  $Q$  is augmented with a flag stating whether the edge leading to this vertex has been evaluated or not. Initially, every edge is given the estimated

value computed using  $\hat{w}$  for its cost and this is used to order the nodes in  $Q$ . Only when a node is selected for expansion, is the true cost of the edge leading to it evaluated. After the edge is evaluated, its cost may be found to be higher than the lazy estimate, or even  $\infty$  (if the edge is untraversable—a notion we will refer to as “in collision”). In such cases, we simply discard the node. If it is valid, we now know the true cost of the edge, as well as the true cost-to-come for this vertex from its current parent. The node is marked to have its true cost determined and is inserted into  $Q$  again. Only when this node is chosen from  $Q$  the second time will it actually be expanded to generate paths to its neighbours. The algorithm terminates when  $v_{\text{target}}$  is removed from  $Q$  for the second time.

This approach increases the size of  $Q$  as there can be multiple nodes associated with every vertex, one for each incoming edge. Since the true cost of an edge is unknown until evaluated, it is essential that all these nodes be stored. Although this causes an increase in computational complexity and in the memory footprint of the algorithm, the approach can lead to fewer edges evaluated and hence reduce the overall running time of the search.

LWA\* uses a one-step lookahead to reduce the number of edge evaluations. Namely, every path in the shortest-path tree  $\mathcal{T}$  may contain one edge (the last) which has only been evaluated lazily. Taking this idea to the limit, Dellin and Srinivasa proposed the Lazy Shortest Path, or LazySP algorithm which uses an infinite-step lookahead. Specifically, it runs a series of shortest-path searches on the graph defined using  $\hat{w}$  for all unevaluated edges. At each iteration, it chooses the shortest path to the goal and evaluates edges along this path<sup>1</sup>. When an edge is evaluated, the algorithm considers the evaluated true cost of the edge for subsequent iterations of the search. Hence, LazySP evaluates only those edges which potentially lie along the shortest path to the goal. The algorithm terminates when all the edges along the current shortest path have been evaluated to be valid (namely, not in collision).

A naïve implementation of LazySP would require running a complete shortest-path search every iteration. However, the search tree computed in the previous iterations can be reused: When an edge is found to be in collision, the search tree computed in previous iteration is locally updated using dynamic shortest-path algorithms such as LPA\* (Koenig, Likhachev, and Furcy 2004).

### 3.4 Motivation

As described in Section 3.3, LWA\* and LazySP attempt to reduce the number of edge evaluations by delaying evaluations until necessary. As we shall prove in Section 6, LazySP (with a lookahead of infinity), minimizes the number of edge evaluations, at the expense of greater graph operations. When an edge is found to be in collision, the entire subtree

<sup>1</sup>In the original exposition of LazySP, the method for which edges are evaluated along the shortest path is determined using a procedure referred to as an *edge selector*. In our work we consider the most natural edge selector, called forward edge selector. Here, the first unevaluated edge closest to the source is evaluated.

emanating from that edge needs to be updated (a process we will refer to as *rewiring*) to find the new shortest path to each node in the subtree. On the other hand, LWA\* which has a lookahead of one, evaluates a larger number of edges relative to LazySP but does not perform any rewiring or re-pairing. When an edge to a node is found to be invalid, the node is simply discarded and the algorithm continues.

Therefore, these two algorithms, LWA\* and LazySP, with a one-step and an infinite-step lookahead respectively, form two extremals to an entire spectrum of potential lazy-search algorithms based on the lookahead chosen. We aim to leverage the advantage that the lazy lookahead can provide to interpolate between LWA\* and LazySP, and strike a balance between edge evaluations and graph operations to minimize the total planning time.

## 4 Problem Formulation

In this section we formally define our problem. To make this section self contained, we repeat definitions that were mentioned in passing in the previous sections. We consider the problem of finding the shortest path between source and target vertices  $v_{\text{source}}$  and  $v_{\text{target}}$  on a given graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ . Since we are motivated by robotic applications where edge evaluation, i.e., checking if the robot collides with its environment while moving along an edge, is expensive, we do not build the graph  $\mathcal{G}$  with just feasible edges. Rather, as in the lazy motion-planning paradigm, the idea is to construct a graph with edges *assumed* to be feasible and delay the evaluation to only when absolutely necessary.

For simplicity, we assume the lazy estimate  $\hat{w}$  to tightly estimate the true cost  $w$  for edges that are collision-free<sup>2</sup>. Therefore  $\hat{w}$  is a *lazy estimate* of  $w$  such that

$$w(e) = \begin{cases} \hat{w}(e) & \text{if } e \text{ is not in collision,} \\ \infty & \text{if } e \text{ is in collision.} \end{cases} \quad (1)$$

We use the cost function  $w$  and its lazy estimate  $\hat{w}$  to define the cost of a path on the graph. The (*true*) cost of a path  $P$  is the sum of the weights of the edges along  $P$ :

$$w(P) = \sum_{e \in P} w(e).$$

Similarly, the *lazy cost* of a path is the sum of the lazy estimates of the edges along the path:

$$\hat{w}(P) = \sum_{e \in P} \hat{w}(e).$$

Our algorithm will make use of paths which are only partially evaluated. Specifically, every path  $P$  will be a concatenation of two paths  $P = P_{\text{head}} \cdot P_{\text{tail}}$  (here,  $(\cdot)$  denotes the concatenation operator). Edges belonging to  $P_{\text{head}}$  will have been evaluated and known to be collision-free while edges belonging to  $P_{\text{tail}}$  will only be lazily evaluated. Notice that  $P_{\text{tail}}$  may be empty. We also define the *estimated total cost* of a path  $P = P_{\text{head}} \cdot P_{\text{tail}}$  as:

$$\bar{w}(P) = w(P_{\text{head}}) + \hat{w}(P_{\text{tail}}).$$

<sup>2</sup>We discuss relaxing the assumption that  $\hat{w}(e)$  tightly estimates  $w(e)$  in Section 8. In the general case, we require only that it is a lower bound i.e.,  $\forall e \in \mathcal{E}, \hat{w}(e) \leq w(e)$ .

Although  $\hat{w}$  helps guide the search of a *lazy* algorithm, as in LWA\* or LazySP, as noted in Section 3.4, it can lead to additional computational overhead when the estimate is wrong, i.e., when the search algorithm encounters edges in collision. In this work we balance this computational overhead with the number of edge evaluations, by endowing our search algorithm with a *lookahead*  $\alpha$ . In essence, the *lookahead* controls the extent to which we use  $\hat{w}$  to guide our search.

As we will see later in Section 8, the lookahead  $\alpha$  can be interpreted in various ways. However, in this paper we interpret the lookahead as the *number of edges* over which we use  $\hat{w}$  to guide our search.

## 5 Lazy Receding-Horizon A\* (LRA\*)

### 5.1 Algorithmic details

Our algorithm maintains a lazy shortest-path tree  $\mathcal{T}$  over the graph  $\mathcal{G}$ . Every node in  $\mathcal{T}$  is associated with a vertex of  $\mathcal{G}$  and the tree is rooted at the node  $\tau_{\text{source}}$  associated with the vertex  $v_{\text{source}}$ . We define the node entry  $\tau \in \mathcal{T}$  as  $\tau = (u, p, c, \ell, b)$ , where  $u[\tau] = u$  is the vertex associated with  $\tau$ ,  $p[\tau] = p$  is  $\tau$ 's parent in  $\mathcal{T}$  which can be backtracked to compute a path  $P[\tau]$  from  $v_{\text{source}}$  to  $u$ . The node  $\tau$  also stores  $c[\tau] = c$  and  $\ell[\tau] = \ell$  which are the costs of the evaluated and lazily-evaluated portions of  $P[\tau]$ , respectively. Namely,  $c[\tau] = w(P[\tau]_{\text{head}})$  and  $\ell[\tau] = \hat{w}(P[\tau]_{\text{tail}})$ . Finally,  $b[\tau] = b$  is the budget of  $P[\tau]$  i.e., the number of edges that have been lazily evaluated in  $P[\tau]$  or equivalently, the number of edges in  $P[\tau]_{\text{tail}}$ . Given a lookahead  $\alpha$ , our algorithm will maintain shortest paths to a set of nodes represented by the search tree  $\mathcal{T}$ , where  $\forall \tau \in \mathcal{T}$ ,  $b[\tau] \leq \alpha$ . The budget of any node in  $\mathcal{T}$  never exceeds  $\alpha$ .

Given a node  $\tau \in \mathcal{T}$ , we call it a *frontier* node if  $b[\tau] = \alpha$  ( $P[\tau]_{\text{tail}}$  has exactly  $\alpha$  edges). Additionally,  $\tau$  is said to belong to the  $\alpha$ -band if  $b[\tau] > 0$ . We call  $\tau$  a *border* node if it does not belong to the  $\alpha$ -band but one of its children does. Finally,  $\tau$  is called a *leaf* node if it has children in  $\mathcal{G}$  but not in  $\mathcal{T}$ . Note that all frontier nodes are leaf nodes. See Fig. 1a for reference.

The algorithm maintains four priority queues to efficiently process the different kinds of nodes, each of which is ordered according to the estimated cost-to-come  $\bar{w}(P[\tau]) = c[\tau] + \ell[\tau]$ . Specifically, we will make use of the following queues:

- $\mathcal{Q}_{\text{frontier}}$  stores the frontier nodes. This queue is used to choose which path to evaluate at each iteration.
- $\mathcal{Q}_{\text{extend}}$  stores leaf nodes that have a budget smaller than  $\alpha$ .
- $\mathcal{Q}_{\text{rewire}}$  stores nodes that require rewiring. It is used to update the structure of  $\mathcal{T}$  when an edge is evaluated to be in collision.
- $\mathcal{Q}_{\text{update}}$  stores the nodes that have children in  $\mathcal{T}$  whose entries need to be updated. It is used to update the structure of  $\mathcal{T}$  after an edge is found to be collision-free.

For ease of exposition, we present a high-level description of the algorithm (Alg. 1) which uses only one of these queues. For detailed pseudo-code, see Appendix A.

---

### Algorithm 1 LRA\* ( $\mathcal{G}, v_{\text{source}}, v_{\text{target}}, \alpha$ )

---

```

1:  $\mathcal{Q}_{\text{frontier}}, \mathcal{T} := \emptyset$  ▷ Initialization
2: insert  $\tau_{v_{\text{source}}} = (v_{\text{source}}, \text{NIL}, 0, 0, 0)$  into  $\mathcal{T}$ 
3: for each leaf node  $\tau \in \mathcal{T}$  do ▷ Extend  $\alpha$ -band
4:   add all nodes at distance  $(\alpha - b[\tau])$  edges into  $\mathcal{T}$ 
5: insert all frontier nodes in  $\mathcal{T}$  into  $\mathcal{Q}_{\text{frontier}}$ 
6: while  $\mathcal{Q}_{\text{frontier}}$  is not empty do ▷ Search
7:   remove  $\tau$  with minimal key  $\bar{w}(\tau)$  from  $\mathcal{Q}_{\text{frontier}}$ 
8:   evaluate first edge  $(u, v)$  along  $P[\tau]_{\text{tail}}$  ▷ Expensive
9:   if  $(u, v)$  is collision-free then
10:    update  $\tau_v$ 
11:    if  $v = v_{\text{target}}$  then
12:      return  $P[\tau_{v_{\text{target}}}]$ 
13:    update descendants  $\tau$  of  $\tau_v$  s.t  $\tau \in \mathcal{T}$ 
14:   else ▷ Edge is in collision
15:     remove edge  $(u, v)$  from graph
16:     for each descendant  $\tau$  of  $\tau_v$  s.t  $\tau \in \mathcal{T}$  do
17:       rewire  $\tau$  to the best parent  $\tau' \in \mathcal{T}$ ,  $\tau' \neq \tau_{v_{\text{target}}}$ 
18:   repeat steps 3-5 to extend the  $\alpha$ -band
19: return failure

```

---

### 5.2 Algorithm Description

Lazy Receding-Horizon A\* (LRA\*) begins by initializing the node  $\tau_{v_{\text{source}}}$  associated with  $v_{\text{source}}$  (line 2). Our algorithm maintains the invariant that at the beginning of any iteration all leaf nodes are frontier nodes. When the algorithm starts,  $\tau_{v_{\text{source}}}$  is a leaf node with  $b[\tau_{v_{\text{source}}}] = 0 < \alpha$ . Therefore we extend the  $\alpha$ -band (lines 3-4) and consequently the search tree  $\mathcal{T}$ , adding all the frontier nodes to  $\mathcal{Q}_{\text{frontier}}$  (line 5).

The algorithm iteratively finds the frontier node  $\tau$  with minimal estimated cost (line 7) and evaluates the first edge along the lazy portion  $P[\tau]_{\text{tail}}$  of the path  $P[\tau]$  from  $\tau_{v_{\text{source}}}$  to  $\tau$  in  $\mathcal{T}$  (line 8). If a collision-free shortest path to  $v_{\text{target}}$  is found during this evaluation (line 11-12), the algorithm terminates. Every evaluation of a collision-free edge  $(u, v)$  causes the node  $\tau_v$ , that was previously in the  $\alpha$ -band, to be a border node. Consequently the node entry is updated and this update is cascaded to all the nodes in the  $\alpha$ -band belonging to the subtree rooted at  $\tau_v$  (lines 10, 13). Specifically, the new cost, lazy cost and budget of  $\tau_v$  is used to update the nodes in its subtree. However, if the edge  $(u, v)$  is found to be in collision, the edge is removed from the graph, and the entire subtree of  $\tau_v$  is rewired appropriately (lines 14-17). This can potentially lead to some of the nodes being removed from the  $\alpha$ -band. Both updating and rewiring subtrees can generate leaf nodes with budget less than  $\alpha$ . Therefore at the end of the iteration, the  $\alpha$ -band is again extended to ensure all leaf nodes have budget equal to the lookahead  $\alpha$  (line 18). See Fig. 1 for an illustration.

As we will show in Section 6, the algorithm described is guaranteed to terminate with the shortest path, if one exists, and is hence complete for all values of  $\alpha$ .

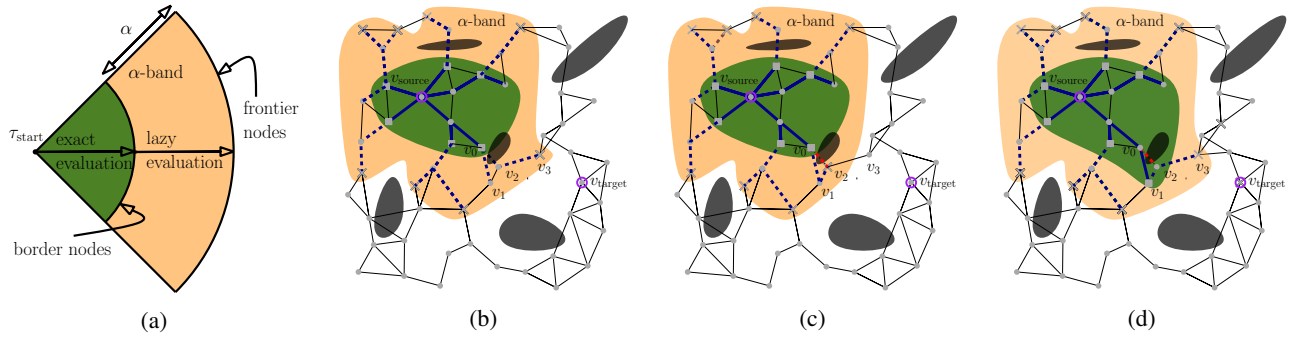


Figure 1: (a) Search space of LRA\*. Figures (b, c, d) visualize LRA\* running on  $\mathcal{G}$  embedded in a workspace cluttered with obstacles (dark grey) and  $\alpha = 2$ . The regions where edges are evaluated and lazily evaluated are depicted by green and orange regions, respectively. Shortest-path tree  $\mathcal{T}$  in the two regions is depicted by solid and dashed blue edges, respectively. Finally, vertices associated with border and frontier nodes are depicted by squares and crosses, respectively. Figure is best viewed in color. (b) Node associated with  $v_3$  has the minimal key and the path ending with nodes  $v_0, v_2, v_3$  is evaluated. Edge  $(v_0, v_2)$  is found to be in collision. (c) Node  $\tau_2$  associated with  $v_2$  is rewired and the  $\alpha$ -band is recomputed. Now  $\tau_2$  has the minimal key and the path ending with nodes  $v_0, v_1, v_2$  is evaluated and found to be collision free. (d) The  $\alpha$ -band is extended from  $v_2$ .

### 5.3 Implementation Details—Lazy computation of the $\alpha$ -band

Every time an edge  $(u, v)$  is evaluated, a series of updates is triggered (Alg. 1 lines 13 and 16-17) Specifically, let  $\tau$  be the node associated with  $v$  and  $\mathcal{T}(\tau)$  be the subtree of  $\mathcal{T}$  rooted at  $\tau$ . If the edge  $(u, v)$  is collision-free, then the budget of all the nodes  $\mathcal{T}(\tau)$  needs to be updated. Alternatively, if the edge  $(u, v)$  is in collision, then a new path to every node in  $\mathcal{T}(\tau)$  needs to be computed. These updates may be time-consuming and we would like to minimize them. To this end, we propose the following optimization which reduces the size of the  $\alpha$ -band and subsequently, potentially reduces the number of nodes in  $\mathcal{T}(\tau)$ .

We suggest that if we already know that a node  $\tau'$  in the  $\alpha$ -band will *not* be part of a path that is chosen for evaluation in an iteration, then we defer expanding the  $\alpha$ -band through this node. The key insight behind the optimization is that there is no need to expand a node  $\tau'$  in the  $\alpha$ -band if its key,  $\bar{w}(P[\tau'])$ , is larger than the key of the first node in  $\mathcal{Q}_{frontier}$ . Using this optimization may potentially reduce the size of  $\mathcal{T}(\tau)$  and save computations.

This is implemented by changing the termination criteria in Alg. 6 (line 1) to test if the key of the first node in  $\mathcal{Q}_{extend}$  is larger than the key of the head of the first node in  $\mathcal{Q}_{frontier}$ .

### 5.4 Implementation Details—Heuristically guiding the search

We described our algorithm as a lazy extension of Dijkstra’s algorithm which orders its search according to cost-to-come. In practice we will want to heuristically guide the search similar to A\*, which orders its search queues according to the sum of cost-to-come to a vertex from  $v_{source}$  and an estimate of the cost-to-go to  $v_{target}$  from the vertex, i.e., a heuristic.

We apply a similar approach by assuming that the algorithm is given a heuristic function that under estimates the cost to reach  $v_{target}$ . We add this value to the key of every node in  $\mathcal{Q}_{frontier}$  and  $\mathcal{Q}_{extend}$ . In Section 6 we state and prove

that as the heuristic is strictly more informative, the number of edge evaluations and rewires further reduce, for a given lazy lookahead.

### 5.5 Discussion—LRA\* as an approximation of optimal heuristic

In this section, we provide an intuition on the role that the lazy lookahead plays when guided by a heuristic. Given a graph  $\mathcal{G}$ , we can define the optimal heuristic  $h_{\mathcal{G}}^*(v)$  as the length of the shortest path from  $v$  to  $v_{target}$  in  $\mathcal{G}$ . Indeed, if all edges of  $\mathcal{G}$  are collision-free, an algorithm such as A\* guided by  $h_{\mathcal{G}}^*$  will only evaluate edges along the shortest path to  $v_{target}$ . To take advantage of this, LazySP proceeds by computing  $h_{\mathcal{G}}^*$ . If an edge is found to be in collision, it is removed from  $\mathcal{G}$  and  $h_{\mathcal{G}}^*$  is recomputed. This is why no other algorithm can perform fewer edge evaluations (see Section 6).

Using a finite *lookahead* and a *static* admissible heuristic, LRA\* can be seen as a method to *approximate* the optimal heuristic. Every frontier node  $\tau$  is associated with the key  $c[\tau] + \ell[\tau] + h(u[\tau])$ . The minimal of all such keys forms the approximation for the optimal heuristic  $h_{\mathcal{G}}^*(v_{source})$  i.e., if  $\tau_v$  associated with vertex  $v$  has the minimal key, we have,

$$h_{\mathcal{G}}^*(v_{source}) \geq c[\tau_v] + \ell[\tau_v] + h(v) \geq h(v_{source})$$

and the algorithm chooses to evaluate an edge along the path from  $v_{source}$  to  $v$  in  $\mathcal{T}$ . This approximation improves as the  $\alpha$ -band approaches the target. When the algorithm starts, this approximation may be crude (when a small lazy lookahead is used). However, as the algorithm proceeds and  $\alpha$ -band is expanded, this approximation dynamically converges to the optimal heuristic. The approximation can also be improved by increasing the lookahead since a larger lookahead enables the algorithm to be more informed. We formalize these ideas in Section 6 (see Lemmas 3, 4), Section 7.3 and show this phenomenon empirically in Section 7.

## 5.6 Discussion—Is greediness beneficial?

A possible extension to LRA\* is to employ greediness in edge evaluation: Given a path, we currently evaluate the first edge along this path (Alg. 1, lines 7 and 8). However, we can choose to evaluate more than one edge, hence performing an exploitative action. This introduces a second parameter  $\beta \leq \alpha$  that indicates how many edges to evaluate along the path. However, we can show that our current formulation using a minimal greediness value of  $\beta = 1$  always outperforms any other greediness value. This is only the case when we seek *optimal* paths. If we relax the algorithm to produce suboptimal paths, greediness may be of use in early termination. While this relaxation is out of the scope of the paper, we provide proofs pertaining to the superiority of no greediness in the Appendix C for the case that optimal paths are required.

## 6 Correctness, Optimality and Complexity

In this section we provide theoretical properties regarding our family of algorithms LRA\*. For brevity, we defer all proofs to Appendix B. We start in Section 6.1 with a correctness theorem stating that upon termination of the algorithm, the shortest path connecting  $v_{\text{source}}$  and  $v_{\text{target}}$  is found. We continue in Section 6.2 to detail how the lazy lookahead affects the performance of the algorithm with respect to edge evaluations. Specifically, we show that for  $\alpha = \infty$ , the algorithm is edge optimal. That is, it tests the minimal number of edges possible (this notion is formally defined). Furthermore, we examine how the lazy lookahead affects the number of edges evaluated by our algorithm. Finally, in Section 6.3 we bound the running time of the algorithm as well as its space complexity as a function of the lazy lookahead  $\alpha$ . Here, we show that the running time (governed, in this case, by graph operations) can grow exponentially with the lazy lookahead  $\alpha$ . This further backs our intuition that in order to minimize the running time in practice, an intermediate lookahead is required to balance edge evaluation and graph operations.

The following additional notation will be used throughout this section: Let  $P_v^*$  denote the shortest collision-free path from  $v_{\text{source}}$  to a vertex  $v$  and let  $w^*(v) = w(P_v^*)$  be the minimal *true* cost-to-come to reach  $v$  from  $v_{\text{source}}$ . Finally, for the special case of  $v_{\text{target}}$ , we will use  $w^* = w^*(v_{\text{target}})$ . That is,  $w^*$  denotes the minimal cost-to-come to reach  $v_{\text{target}}$  from  $v_{\text{source}}$ .

### 6.1 Correctness

**Lemma 1.** *Let  $(v_0, v)$  be an edge evaluated by LRA\* and found to be collision free. Then the shortest path to the node  $\tau_v$  associated with vertex  $v$  has been found and  $c[\tau_v] = w^*(v)$ .*

Replacing  $v$  with  $v_{\text{target}}$ , we have,

**Corollary 1.** *LRA\* is complete, i.e., if an edge  $(v_0, v_{\text{target}})$  is found to be collision-free, the shortest path to  $\tau_{v_{\text{target}}}$  associated with  $v_{\text{target}}$  has been found.*

### 6.2 Edge Optimality

We analyze how the lazy lookahead allows to balance between the number of edge evaluations and rewiring operations. We start by looking at the extreme case where there is an infinite lookahead ( $\alpha = \infty$ ). We define a general family of algorithms  $\mathcal{SP}$  that solve the shortest-path problem and show in Lemma 2 that when  $\alpha = \infty$ , no other algorithm in  $\mathcal{SP}$  can perform fewer edge evaluations. We then show in Lemma 4 that larger the lookahead, fewer the edge evaluations LRA\* will perform.

Recall that a shortest-path problem consists of a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , a lazy estimate of the weights  $\hat{w}$ , a weight function  $w$  and start ( $v_{\text{source}}$ ) and goal ( $v_{\text{target}}$ ) vertices. Given a shortest-path problem, let  $\mathcal{SP}$  be the family of shortest-path algorithms that build a shortest-path tree  $\mathcal{T}$  rooted at  $v_{\text{source}}$ . Assume that for every shortest-path problem, there are no two paths in  $\mathcal{G}$  that have the same weight<sup>3</sup>.

An algorithm  $\text{ALG} \in \mathcal{SP}$  can only call the weight function  $w$  for an edge  $e = (u, v)$  if  $u \in \mathcal{T}$ . When terminating, it must report the shortest path from  $P_{v_{\text{target}}}^*$  and validate that no shorter path exists. Thus for any other path  $P$  from  $v_{\text{source}}$  to  $v_{\text{target}}$  with  $\hat{w}(P) < w(P_{v_{\text{target}}}^*)$ , ALG must explicitly test an edge  $e \in P$  with  $w(e) = \infty$ . Since ALG constructs a shortest-path tree, this will be the first edge on  $P$  that is in collision.

Finally, an algorithm  $\text{ALG} \in \mathcal{SP}$  is said to be *edge-optimal* if for any other algorithm  $\text{ALG}' \in \mathcal{SP}$ , and any shortest-path problem, ALG will test no more edges than  $\text{ALG}'$ .

**Lemma 2.** *LRA\* with  $\alpha = \infty$  is edge-optimal.*

**Corollary 2.** *LazySP is edge-optimal.*

Considering LRA\* as an approximation of the optimal heuristic can provide a different perspective on how LazySP is edge optimal. Consider consistent heuristics  $h_1$  and  $h_2$ , such that  $h_1$  strictly dominates  $h_2$  i.e.,

$$h_{\mathcal{G}}^*(v) \geq h_1(v) > h_2(v) \quad \forall v \in \mathcal{V}, v \neq v_{\text{target}},$$

where  $h_{\mathcal{G}}^*(v)$  is the optimal heuristic for a given graph  $\mathcal{G}$ .

**Lemma 3.** *For every graph  $\mathcal{G}$  and lookahead  $\alpha$ , we have that  $E_1 \subseteq E_2$ , where  $E_i$  denotes the set of edges evaluated by LRA\* with heuristic  $h_i$ ,  $i \in \{1, 2\}$ .*

**Corollary 3.** *LazySP is edge-optimal.*

**Lemma 4.** *For every graph  $\mathcal{G}$  and every  $\alpha_1 > \alpha_2$ , we have that  $E_1 \subseteq E_2$ . Here,  $E_i$  denotes the set of edges evaluated by LRA\* with  $\alpha = i$ .*

### 6.3 Complexity

In this section we analyse LRA\* with respect to the space (Lemma 5) and running time (Lemma 6) complexity.

**Lemma 5.** *The total space complexity of our algorithm is bounded by  $O(n + m)$ , where  $n$  and  $m$  are the number of vertices and edges in  $\mathcal{G}$ , respectively.*

**Lemma 6.** *The total running time of the algorithm is bounded by  $O(nd^\alpha \cdot \log(n) + m)$ , where  $n$  and  $m$  are the number of vertices and edges,  $d$  is the maximal degree of a vertex and  $\alpha$  is the lookahead.*

<sup>3</sup>To avoid handling tie-breaking in proofs. LRA\* does not require this assumption.

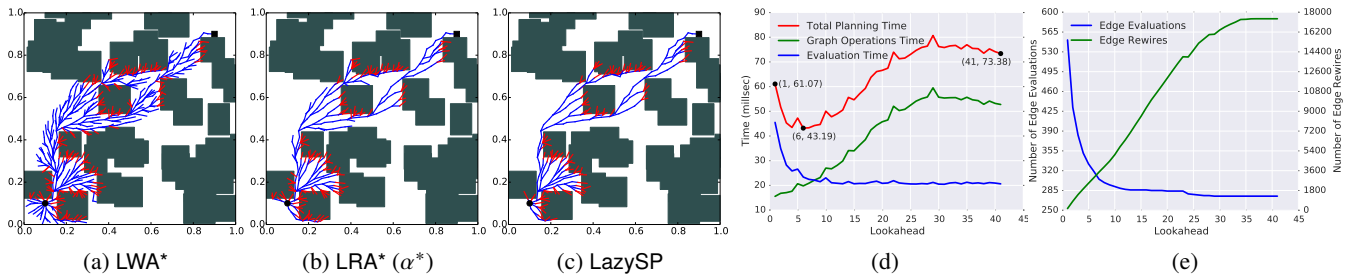


Figure 2: Visualization of edge evaluations by (a) LWA\*, (b) LRA\* with an optimal lookahead  $\alpha^*$ , and (c) LazySP. Source and target are  $(0.1, 0.1)$  and  $(0.9, 0.9)$ , respectively. Edges evaluated to be in collision and free are marked red and blue, respectively. Computation times (d) and number of operations (e) of LRA\* as a function of the lookahead  $\alpha$ .

## 7 Results

In this section we empirically evaluate LRA\*. We start by demonstrating the different properties of LRA\* as a *family* of algorithms parameterized by  $\alpha$ . Specifically, we show that to minimize the total planning time, an optimal lookahead  $\alpha^*$  exists (where  $1 < \alpha^* < \infty$ ) that allows to balance between edge evaluation and graph operations.

We then continue to evaluate properties of the optimal lookahead  $\alpha^*$ . While choosing the exact lookahead value is out of the scope of the paper (see Sec. 8), we provide general guidelines regarding this choice.

### 7.1 Experimental setup

We evaluated LRA\* on a range of planning problems in simulated random  $\mathbb{R}^2$  and  $\mathbb{R}^4$  environments as well as real-world manipulation problems on HERB (Srinivasa et al. 2009), a mobile manipulator with 7-DOF arms. We implemented the algorithm using the Open Motion Planning Library (OMPL) (Sucan, Moll, and Kavraki 2012)<sup>4</sup>. Our source code is publicly available and can be accessed at <https://github.com/personalrobotics/LRA-star>.

**Random environments** We generated 10 different random environments for  $\mathbb{R}^2$  and  $\mathbb{R}^4$ . For a given environment, we consider 10 distinct random roadmaps for a total of 100 trials for each dimension. Each roadmap was constructed as follows: The set of vertices were generated in a unit hypercube using Halton sequences (Halton 1964), which are characterized by low dispersion. The vertex positions were then offset by uniform random values to generate distinct roadmaps. An edge existed in the graph between every pair of vertices whose Euclidean distance is less than a predefined threshold  $r$ . The value  $r$  was chosen to ensure that, asymptotically, the graph can capture the shortest path connecting the start to the goal (Janson et al. 2015). The number of vertices was chosen such that the roadmap contained a solution. Specifically, it was 2000 for  $\mathbb{R}^2$  and 3000 for  $\mathbb{R}^4$ .

The source and target were set to  $(0.1, 0.1, \dots, 0.1)^d$  and  $(0.9, 0.9, \dots, 0.9)^d$ , respectively, with  $d \in \{2, 4\}$ . For the 2D environments, the obstacles were a set of axis-aligned hypercubes that occupy 70% of an environment to simulate a cluttered space. One such randomly-generated environment

is shown in Fig. 2 along with the edges evaluated by LWA\*, LazySP and LRA\* with an optimal lookahead. For the 4D environments, we chose a maze generated similar to the recursive mazes defined by Janson et al.. The choice of such a maze in  $\mathbb{R}^4$  is motivated by the fact that it is inherently a *hard* problem to solve, since many lazy shortest paths need to be invalidated before a true shortest path is determined by the planner. A detailed discussion about the complexity of the recursive maze problem is found in (Janson et al. 2015).

**Manipulation** Our manipulation problems simulate the task of reaching into a bookshelf while avoiding obstacles such as a table. We consider 10 different roadmaps, each with 30,000 vertices constructed by applying a random offset to the 7D Halton sequence. Two vertices are connected if their Euclidean distance is less than  $r = 1.3$  radians. These choices are similar to the simulated  $\mathbb{R}^n$  worlds, where we choose  $r$  using the bounds provided by Janson et al. and enough vertices such that we are ensured a solution exists on the roadmap. Fig. 3 illustrates the environment and the planning problem considered.

### 7.2 Properties of LRA\*

Figures 2 and 3 visualize the search space for our simulated  $\mathbb{R}^2$  environments as well as our manipulation environment. For both settings, we ran LRA\* with a range of lookahead values.

Notice that the number of edge evaluations as a function of the lookahead is a monotonically decreasing function (Fig. 2e and Lemma 4). However, the time spent on edge evaluations (Fig. 2d) is not monotonic. This is because the time for evaluating an edge depends on the edge length and if it is in collision. Having said that, the overall trend of this plot decreases as the lookahead increases. In addition, the time spent on rewiring (Fig. 2d and 3d) roughly increases with the lookahead. Following these two trends we find that, in both experiments, an intermediate lookahead does indeed balance edge evaluations and graph operations. This, in turn reduces the overall planning time.

### 7.3 Lazy Lookahead and Dynamic Heuristic

We consider every border node in the  $\alpha$ -band to be associated with a dynamic heuristic that extracts information about the graph structure up to  $\alpha$  edges away, and the static heuristic associated with a frontier node exactly  $\alpha$  edges away.

<sup>4</sup>Simulations were run on a desktop machine with 16GB RAM and an Intel i5-6600K processor running a 64-bit Ubuntu 14.04.

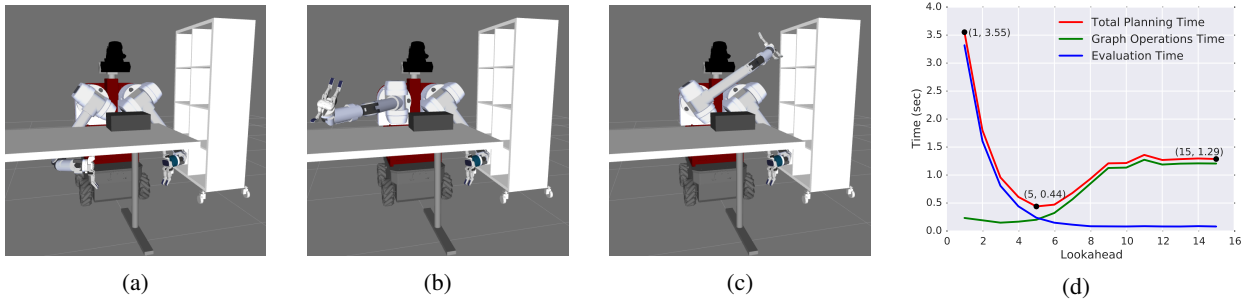


Figure 3: Manipulation experiments. (a-c) HERB is required to reach into the bookshelf while avoiding collision with the table. (d) Edge evaluation, rewiring and total planning time as a function of the lookahead.

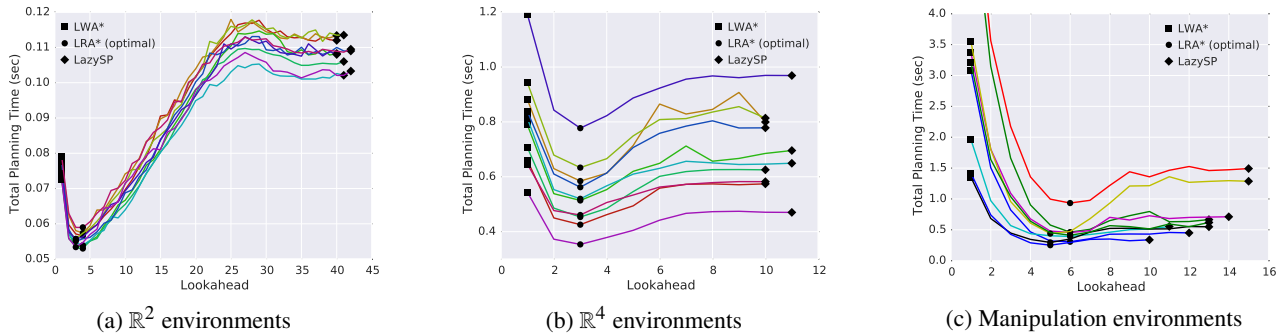


Figure 4: Planning time vs. lookahead for similar problems on different environments.




Figure 5: The  $f$ -value of the top node popped from  $Q_{\text{frontier}}$  every iteration of  $LRA^*$  for various lookaheads.

As the lookahead increases from one to infinity,  $LRA^*$  lazily obtains an increasing amount of information about the underlying graph structure. This information, encoded in the dynamic heuristic, allows  $LRA^*$  with a larger lookahead to search a smaller region of the graph and evaluate at most as many edges as  $LRA^*$  with a smaller lookahead (Lemma 4).

In Fig. 5 we plot, for a 2D problem, the  $f$ -value<sup>5</sup> of the top node in  $Q_{\text{frontier}}$  for each iteration of the algorithm. Note that each iteration of the algorithm corresponds to an edge evaluation. LazySP converges to the optimal  $f$ -value in the fewest number of iterations while  $LWA^*$  evaluates the most number of edges since it is least informed amongst the family of  $LRA^*$  algorithms.

We observe in Fig. 5 that for  $LRA^*$  with an intermediate lookahead  $1 < \alpha < \infty$ , the  $f$ -values of the nodes popped from

<sup>5</sup>Recall  $f$ -value is the key used to order nodes in  $Q_{\text{frontier}}$ . For a given node, it is the sum of estimated cost-to-come and cost-to-go.

the  $Q_{\text{frontier}}$  do not monotonically increase. This is attributed to the fact that the dynamic heuristic does not necessarily capture the true underlying graph structure when there are edges in collision. When an edge is found to be in collision, the  $\alpha$ -band can potentially shrink closer to the source based on the budget available. This generates new leaf nodes closer to the source which can possibly be characterized by a lower  $f$ -value compared to the  $f$ -value of the node popped in the earlier iteration. This does not occur in  $LWA^*$  since it has a dynamic heuristic over one step. In case of a collision, the heuristic is trivially updated after removing the edge from the graph. LazySP does not exhibit this behavior either since it always pops the goal node and the  $f$ -value of a *particular* node in the graph can only monotonically increase.

## 7.4 Properties of optimal lookahead $\alpha^*$

While determining how to choose the lookahead value for a specific problem instance is beyond the scope of this paper (see Sec. 8), we provide some insight on some properties of optimal lookahead  $\alpha^*$ . In Fig. 4 we plotted the planning time as a function of the lookahead for different random instances. We observe two phenomena: (i) the value of the optimal lookahead  $\alpha^*$  has a very small variance when considering similar environments. Thus, if we will face multiple problems on a specific type of environment, it may be beneficial to run a preprocessing phase to estimate  $\alpha^*$ . (ii) As the dimension increases, the relative speedup, when compared to LazySP diminishes. We conjecture that this is because the cost of edge evaluation increases with the complexity of the robot (namely, with the dimension).



## 8 Future Work

**Setting the lazy lookahead** Our formulation assumed that the lazy lookahead  $\alpha$  is fixed and provided by the user. In practice, we would like to automatically find the value of  $\alpha$  and, possibly, change its value through the running time of the algorithm. This is especially useful when the search algorithm is interleaved with graph construction—namely, when vertices and edges are incrementally added to  $\mathcal{G}$  (see, e.g., (Gammell, Srinivasa, and Barfoot 2015)).

**Non-tight estimates of edge weights** In this paper we assumed that  $\hat{w}$  tightly estimates the true cost  $w$  (see Eq. 1), however it can be easily extended to take into account non-tight estimates. Once an edge  $(u, v)$  is evaluated, if its true cost is larger than the estimated cost, the entire subtree rooted at  $v$  may need to be rewired to potentially better parents. Our immediate goal is to run our algorithm on such settings.

### Alternative budget definitions and optimization criteria

In this paper, we defined the budget and the optimization criteria in terms of number of unevaluated edges and path length, respectively. However, the same approach can be used for alternative definitions. For example, we can define the budget in terms of the *length* of the unevaluated path. This definition is somewhat more realistic since the computational cost of evaluating an edge is typically proportional to its length. A different optimization criteria that we wish to consider is minimizing the *expected* number of edges checked given some belief over the probability that edges are collision-free. This can be further extended to balance between path length (which is a proxy for the execution time) and number of edge evaluations (which is a proxy for the planning time). Here, we need to consider some combination of path length and probability of being collision-free as the optimization criteria.

### Implementing LRA\* using advanced priority queues

Recall that LRA\* maintains only the best path to reach each frontier node in the search tree  $\mathcal{T}$  at every point in time. This is done to avoid an exponential increase in the memory footprint of LRA\* with respect to the laziness value  $\alpha$ . Consequently, all the priority queues that we use require the ability to update the key of elements in the queue. In contrast, LWA\*, which may maintain several paths to the same node, only requires a priority queue that supports inserting elements and popping the element with the minimal key. Interestingly, such priority queues allow for a significant speedup in A\*-like algorithms (Chen et al. 2007). LRA\* can easily be modified to maintain all unevaluated paths to frontier nodes. If the laziness value  $\alpha$  is relatively small, then the use of fast priority queues that do not support key updates may significantly speed up the algorithm’s running time.

## 9 Acknowledgements

The authors would like to thank Shushman Choudhury, previously at Personal Robotics Lab, now at Stanford University, for his valuable insights and discussions in the development of this work.

---

### Algorithm 2 LRA\* ( $\mathcal{G}, v_{\text{source}}, v_{\text{target}}, \alpha$ )

---

```

1:  $\tau_{v_{\text{source}}} = (v_{\text{source}}, \text{NIL}, 0, 0, 0)$ ;  $\mathcal{T}.\text{insert}(\tau_{v_{\text{source}}})$ 
2: for all  $v \in \mathcal{V}, v \neq v_{\text{source}}$  do
3:    $\tau_v = (v, \text{NIL}, \infty, \infty, \infty)$  ▷ Initialization
4:  $Q_{\text{update}}, Q_{\text{frontier}}, Q_{\text{rewire}} \leftarrow \emptyset$ 
5:  $Q_{\text{extend}}.\text{push}(\tau_{v_{\text{source}}})$ 
6: extend_ $\alpha$ _band() ▷ populate  $Q_{\text{frontier}}$ 
7: while  $Q_{\text{frontier}} \neq \emptyset$  do
8:    $\tau \leftarrow Q_{\text{frontier}}.\text{pop}()$ 
9:    $P \leftarrow P[\tau]_{\text{tail}}$  ▷ extract path from border node to  $\tau$ 
10:  evaluate_path $(P)$  ▷ populate  $Q_{\text{update}}, Q_{\text{rewire}}$ 
11:  if  $\tau_{v_{\text{target}}} \in Q_{\text{update}}$  then
12:    return  $P[\tau_{v_{\text{target}}}]$  ▷ return path from  $v_{\text{source}}$  to  $v_{\text{target}}$ 
13:  update_ $\alpha$ _band() ▷ populate  $Q_{\text{extend}}$ 
14:  rewire_ $\alpha$ _band() ▷ populate  $Q_{\text{extend}}$ 
15:  extend_ $\alpha$ _band() ▷ populate  $Q_{\text{frontier}}$ 
16: return failure

```

---

## A Algorithm Description

In this section we provide detailed pseudo-code of LRA\*. We start in Alg. 2 which details the main loop used by LRA\*.

### A.1 Main Algorithm

We start (lines 1-3) by adding a node corresponding to the source into the search tree and initializing all other nodes. We continue (lines 4-5) by initializing all the priority queues used by the algorithm. The algorithm then extends the  $\alpha$ -band (line 6 and Alg. 6) which computes the frontier nodes stored in  $Q_{\text{frontier}}$ .

From this point, the algorithm iterates between choosing a path (defined by the frontier node with the minimal cost-to-come) and evaluating the first edge along this path (lines 8-10). If the target was reached, the algorithm terminates (lines 11-12). This evaluation also adds nodes to either  $Q_{\text{update}}$  and  $Q_{\text{rewire}}$ , depending if the edge evaluated was collision free or not. A node will be added to  $Q_{\text{update}}$  because the budget of all nodes in its subtree needs to be updated. Similarly, a node will be added to  $Q_{\text{rewire}}$  because its current parent is in collision and all nodes in its subtree need to be rewired. Thus, if the target was not reached the algorithm updates vertices in the  $\alpha$ -band (line 13 and Alg. 4), rewires vertices if needed (line 14 and Alg. 5), and re-extends the  $\alpha$ -band (line 15 and Alg. 6).

### A.2 Path Evaluation

When a path is chosen, LRA\* evaluates the first edge along the tail of the path (Alg. 3, line 2). If the edge is collision free (lines 3-6), its entries are updated and its target is pushed into  $Q_{\text{update}}$ . This will later be used in Alg. 4) to update all nodes in its subtree in a systematic manner. If the edge is in collision (lines 8-9), the corresponding edge is removed from the graph and the subtree rooted at the target vertex of the edge is set for rewiring. Similar to the previous case this will later be used in Alg. 5) to rewire all nodes in the subtree in a systematic manner.

---

**Algorithm 3** evaluate\_path ( $P = (\tau_0, \tau_1, \dots, \tau_\alpha)$ )

---

```
1:  $e \leftarrow (u[\tau_0], u[\tau_1])$ 
2: if  $w(e) = \hat{w}(e)$  then ▷ Expensive check
3:    $\tau_1 \leftarrow (u_1, u_0, c[\tau_0] + w(e), 0, 0)$  ▷ Update node
4:    $Q_{\text{update}}.\text{push}(\tau_1)$ 
5:   if  $u_1 = v_{\text{target}}$  then
6:     return
7: else ▷ Invalid edge: rewiring required
8:    $\mathcal{E}.\text{remove}(e)$  ▷ Remove edge from graph
9:    $\mathcal{T}_{\text{rewire}} \leftarrow \mathcal{T}_{\text{sub}}(\tau_1)$ 
10:  return
11: return
```

---

---

**Algorithm 4** update\_α\_band ()

---

```
1: while  $Q_{\text{update}} \neq \emptyset$  do
2:    $\tau \leftarrow Q_{\text{update}}.\text{pop}()$ 
3:    $T_{\text{succ}} \leftarrow \{\tau' \in \mathcal{T} \mid p[\tau'] = \tau\}$ 
4:   if  $T_{\text{succ}} = \emptyset$  then ▷ Leaf node with budget
5:      $Q_{\text{extend}}.\text{push}(\tau)$ 
6:     continue
7:   for all  $\tau' \in T_{\text{succ}}$  do
8:     if  $b[\tau'] = \alpha$  then ▷ Cleanup queue
9:        $Q_{\text{frontier}}.\text{remove}(\tau')$ 
10:     $\tau' \leftarrow (u[\tau'], u[\tau], c[\tau], \ell[\tau] + \hat{w}(u[\tau], u[\tau']), b[\tau] + 1)$ 
11:     $Q_{\text{update}}.\text{push}(\tau')$ 
12:  return
```

---

### A.3 Updating α-band

In Alg. 3, when an edge  $(u, v)$  has been evaluated to be collision-free, node  $\tau_v$  associated with vertex  $v$  is updated and  $Q_{\text{update}}$  is populated with the updated node. This update to  $\tau_v$  needs to be cascaded to all the vertices in its subtree in a breadth-first search manner such that the parent node is updated before the child node. In every iteration of Alg. 4, the node with minimal key is popped from  $Q_{\text{update}}$  (line 2) and its successors in the subtree are obtained (line 3). If the set of successors is empty, this implies that the node is a leaf node and is hence pushed into  $Q_{\text{extend}}$  (lines 4-6). Otherwise, each of the successor nodes is updated using the parent node entries and pushed into  $Q_{\text{update}}$  (lines 7-11). Note that leaf nodes belonging to this subtree are removed from  $Q_{\text{frontier}}$  as their budget is updated to less than  $\alpha$  (lines 8-10).

### A.4 Rewiring α-band

In Alg. 3, when an edge  $(u, v)$  is found to be in collision, the subtree rooted at  $\tau_v$  is to be rewired as in Alg. 5. Initially every node in the subtree is updated to have an infinite key, and removed from the search tree (lines 1-3). Since these nodes have their entries re-initialized, they are removed from any priority queue they might exist in, namely,  $Q_{\text{frontier}}$  (line 5). For each of these nodes, the best *valid* parent in the graph is determined. The node is updated using the new parent's node entries and pushed into  $Q_{\text{rewire}}$ . Note that valid parents do not appear in the subtree being rewired since their node entries are still unknown (lines 8-9). Nodes belonging to  $Q_{\text{extend}}$  are subsequently extended in line 15 of Alg.2 and

---

**Algorithm 5** rewire\_α\_band ( $Q_{\text{rewire}}$ )

---

```
1: for all  $\tau \in \mathcal{T}_{\text{rewire}}$  do ▷ Assign keys to nodes in subtree
2:    $\mathcal{T}.\text{remove}(\tau)$ 
3:    $\tau = (u[\tau], \text{NIL}, \infty, \infty, \infty)$ 
4:   if  $\tau \in Q_{\text{frontier}}$  then
5:      $Q_{\text{frontier}}.\text{remove}(\tau)$ 
6:    $S_{\text{parents}} \leftarrow \{\tau' \in \mathcal{T} \text{ s.t. } (u[\tau'], u[\tau]) \in \mathcal{E}, b[\tau'] < \alpha\}$ 
7:   for all  $\tau' \in S_{\text{parents}} - \{\mathcal{T}_{\text{rewire}} \cup Q_{\text{extend}} \cup \tau_{\text{target}}\}$  do
8:     if  $c[\tau] + \ell[\tau] > c[\tau'] + \ell[\tau'] + \hat{w}(u[\tau'], u[\tau])$  then
9:        $\tau \leftarrow (u[\tau], u[\tau'], c[\tau'], \ell[\tau'] + \hat{w}(u[\tau'], u[\tau]), b[\tau'] + 1)$ 
10:     $Q_{\text{rewire}}.\text{push}(\tau)$ 
11:  while  $Q_{\text{rewire}} \neq \emptyset$  do ▷ Rewire
12:     $\tau \leftarrow Q_{\text{rewire}}.\text{pop}()$ 
13:    if  $p[\tau] = \text{NIL}$  then
14:      continue
15:     $\mathcal{T}.\text{insert}(\tau)$ 
16:    if  $b[\tau] = \alpha$  or  $u[\tau] = v_{\text{target}}$  then
17:       $Q_{\text{frontier}}.\text{push}(\tau)$ 
18:      continue
19:    if  $b[\tau] < \alpha$  then ▷ Note  $u[\tau] \neq v_{\text{target}}$ 
20:       $Q_{\text{extend}}.\text{push}(\tau)$ 
21:    for all  $v \in \mathcal{V}$  s.t.  $(u[\tau], v) \in \mathcal{E}, \tau_v \in Q_{\text{rewire}}$  do
22:      if  $c[\tau] + \ell[\tau] + \hat{w}(u[\tau], v) < c[\tau_v] + \ell[\tau_v]$  then
23:         $\tau_v \leftarrow (v, u[\tau], c[\tau], \ell[\tau] + \hat{w}(u[\tau], v), b[\tau] + 1)$ 
24:         $Q_{\text{rewire}}.\text{update\_node}(\tau_v)$ 
25:     $\mathcal{T}_{\text{rewire}}.\text{clear}()$ 
26:  return
```

---

hence are considered as invalid parents (lines 10-11).

Once  $Q_{\text{rewire}}$  is populated with all the nodes in the subtree, the algorithm iteratively pops the node with the minimal key from  $Q_{\text{rewire}}$  (lines 17-18). If a valid parent has been determined for the node, it is inserted into the search tree, and priority queues  $Q_{\text{extend}}$  and  $Q_{\text{frontier}}$  depending on its budget (lines 21-25). Otherwise the node is left as initialized in line 3. Essentially, this implies that nodes can be inserted and also removed from the search tree during rewiring.

If a node has been successfully rewired and has budget less than  $\alpha$ , it is now a potential valid best parent to nodes associated with its successor vertices in the graph. Lines 27-31 verify if the node is indeed a better parent for each of its successors and updates them accordingly.

### A.5 Extending α-band

The queue  $Q_{\text{extend}}$  contains leaf nodes in the search tree  $\mathcal{T}$  that have a budget less than  $\alpha$ . In Alg. 6, the top node  $\tau$  in  $Q_{\text{extend}}$  with minimal key is popped (lines 1-2) and extended unless it is the node associated with  $v_{\text{target}}$  in which case it is pushed into  $Q_{\text{frontier}}$  (lines 3-4). For each of  $\tau$ 's successors  $\tau_v$  in  $\mathcal{G}$ , if the cost to reach  $\tau_v$  through  $\tau$  is cheaper than the current cost, the node entry for  $\tau_v$  is updated using  $\tau$  (lines 7-17) and inserted into the search tree. If  $\tau_v$  already belongs to the search tree, it needs to be removed from any of the priority queues it previously belongs to (lines 10-16). If the successor node  $\tau_v$  has been updated to have  $\tau$  as the parent in  $\mathcal{T}$ , we push  $\tau_v$  into  $Q_{\text{extend}}$  or  $Q_{\text{frontier}}$  depending on its budget.

---

**Algorithm 6** `extend_α_band()`

---

```
1: while  $Q_{\text{extend}} \neq \emptyset$  do
2:    $\tau \leftarrow Q_{\text{extend}}.\text{pop}()$ 
3:   if  $u[\tau] = v_{\text{target}}$  then  $\triangleright$  Goal node needn't be extended
4:    $Q_{\text{frontier}}.\text{push}(\tau)$ 
5:   else
6:     for all  $v \in \mathcal{V}$  s.t.  $(u[\tau], v) \in \mathcal{E}$  do
7:        $\tau'_v \leftarrow (v, u[\tau], c[\tau], \ell[\tau] + \hat{w}(u[\tau], v), b[\tau] + 1)$ 
8:       if  $c[\tau'_v] + \ell[\tau'_v] > c[\tau_v] + \ell[\tau_v]$  then
9:         continue
10:      if  $\exists \tau_v \in \mathcal{T}$  then  $\triangleright$  Cleanup queues before update
11:        for all  $\tau' \in T_{\text{subtree}}(\tau_v)$  do
12:           $\mathcal{T}.\text{remove}(\tau')$ 
13:          if  $\tau' \in Q_{\text{frontier}}$  then
14:             $Q_{\text{frontier}}.\text{remove}(\tau')$ 
15:          if  $\tau' \in Q_{\text{extend}}$  then
16:             $Q_{\text{extend}}.\text{remove}(\tau')$ 
17:           $\tau_v \leftarrow \tau'_v$   $\triangleright$  Update
18:           $\mathcal{T}.\text{insert}(\tau_v)$ 
19:          if  $b[\tau_v] = \alpha$  then
20:             $Q_{\text{frontier}}.\text{push}(\tau_v)$ 
21:          else
22:             $Q_{\text{extend}}.\text{push}(\tau_v)$ 
23: return
```

---

## B Algorithmic Properties: Proofs to Section 6

In this section we provide accompanying proofs to the lemmas presented in Sec. 6. For clarity we repeat the statements of the proofs throughout this section.

**Lemma 1.** *Let  $(v_0, v)$  be an edge evaluated by LRA\* and found to be collision free. Then the shortest path to the node  $\tau_v$  associated with vertex  $v$  has been found and  $c[\tau_v] = w^*(v)$ .*

*Proof.* Since the algorithm has evaluated the edge  $(v_0, v)$  to be collision free, there exists a path  $P = (v_{\text{start}}, \dots, v_1, v_0, v)$  for which all edges were found to be collision free by LRA\*. Assume there exists another collision-free path  $P' = (v_{\text{start}}, \dots, v'_1, v'_0, v)$  from  $v_{\text{start}}$  to  $v$  such that  $w(P') < w(P)$ .

Consider the iteration before LRA\* evaluates  $(v_0, v)$ . Since the edge  $(v_0, v)$  is evaluated, there exists a border node  $\tau_0$  associated with  $v_0$  where and a frontier node  $\tau$  that is  $\alpha$  edges away from  $\tau_0$  for which  $\bar{w}(P[\tau])$  is minimal.

Let  $\tau'_j$  be the last border node on  $P'$  (associated with vertex  $v'_j$ ) and let  $\tau'$  be the frontier node that is  $\alpha$  edges away from  $\tau'_j$  for which  $\bar{w}(P[\tau'])$  is minimal. Note that since  $(v_0, v)$  was evaluated, we have that  $\bar{w}(P[\tau])$  is minimal which implies that  $\bar{w}(P[\tau]) < \bar{w}(P[\tau'])$ .

Consider the following cases:

**C1. We have that  $j > \alpha$ .** Namely, vertex  $v'_j$  lies more than  $\alpha$  edges before  $v$ . Thus, the vertex  $v'$  associated with frontier node  $\tau'$  lies on  $P'$  before  $v$ . By the assumption that  $w(P') < w(P)$ , we have that  $\bar{w}(P[\tau']) < \bar{w}(P[\tau])$  which gives a contradiction to the fact that  $\bar{w}(P[\tau]) < \bar{w}(P[\tau'])$ . See Fig. 6a.

**C2. We have that  $\alpha \geq j$ .** Namely, vertex  $v'_j$  lies at most  $\alpha$  edges before  $v$ . Thus, the vertex  $u'$  associated with frontier node  $\tau'$  is  $v$  or a descendant of  $v$ . Consider the node  $\tau_v$  associated with vertex  $v$ . Clearly,  $\tau_v$  is at most  $\alpha$  edges from both border nodes  $\tau_i$  and  $\tau'_j$ . By the assumption that  $w(P') < w(P)$ , we have that  $\tau_v$  will have  $\tau_i$  as its ancestor in  $\mathcal{T}$  and *not*  $\tau'_j$  which contradicts the fact that  $(v_0, v)$  is chosen for evaluation. See Fig. 6b.  $\square$

**Lemma 2.** *LRA\* with  $\alpha = \infty$  is edge-optimal.*

*Proof.* Notice that LRA\* with  $\alpha = \infty$  has an infinite lookahead. Thus, at each iteration, it will take the shortest path connecting  $v_{\text{source}}$  to  $v_{\text{target}}$  according to the lazy estimate  $\hat{w}$  excluding edges that were already found to be in collision. It will then test edges on this path until one is found to be in collision or until an entire path was found to be collision-free. Moreover, any edge  $e$  tested by LRA\* lies on a path  $P$  with  $\hat{w}(P) \leq w(P^*)$ .

Assume that there exists some shortest-path problem and some algorithm ALG where LRA\* tests more edges than ALG. Let  $e = (u, v)$  be an edge tested by LRA\* and not by ALG. Since LRA\* tested  $e$ , there exists a collision-free path connecting  $v_{\text{source}}$  to  $u$ . Furthermore,  $e$  lies on a path  $P$  with  $\hat{w}(P) \leq w(P^*)$ .

If  $\hat{w}(P) = w(P^*)$  then  $P = P^*$  and ALG must validate all edges of  $P$  including  $e$ . If  $\hat{w}(P) < w(P^*)$  then there exists an edge on  $P$  after  $u$  that is in collision and ALG must indeed test the first edge on  $P$  that is in collision. This implies that ALG must test  $e$ .  $\square$

**Lemma 3.** *For every graph  $\mathcal{G}$  and lookahead  $\alpha$ , we have that  $E_1 \subseteq E_2$ , where  $E_i$  denotes the set of edges evaluated by LRA\* with heuristic  $h_i$ ,  $i \in \{1, 2\}$ .*

*Proof.* Assume that  $E_1 \setminus E_2 \neq \emptyset$  and let  $(v_0, v_1) \in E_1 \setminus E_2$  be the edge such that  $v_0$ 's cost-to-come is minimal. Let  $u$  be the parent of  $v_0$  on the shortest path from  $v_{\text{source}}$  to  $v_0$  and note that  $(u, v_0) \in E_1 \cap E_2$ . Furthermore, let  $\mathcal{T}_i$  denote the search tree of LRA\* with heuristic  $h_i$ ,  $i \in \{1, 2\}$ . Finally, recall that  $w^*$  denotes the (true) weight of the shortest path from  $v_{\text{source}}$  to vertex  $v_{\text{target}}$  and that  $\forall v \in \mathcal{V}$  with  $v \neq v_{\text{target}}$  we have that  $h_{\mathcal{G}}^*(v) \geq h_1(v) > h_2(v)$ .

The edge  $(v_0, v_1)$  was evaluated by LRA\* with heuristic  $h_1$ . Thus, a frontier node  $\tau_{v_\alpha}^1 \in \mathcal{T}_1$  associated with some vertex  $v_\alpha$  was at the head of  $Q_{\text{frontier}}$  (with  $v_\alpha$  exactly  $\alpha$  edges from  $v_0$ ). Since node  $\tau_{v_\alpha}^1$  was popped from  $Q_{\text{frontier}}$ , its key is minimal. Specifically,

$$c[\tau_{v_\alpha}^1] + \ell[\tau_{v_\alpha}^1] + h_1[\tau_{v_\alpha}^1] \leq w^*.$$

The edge  $(u, v_0)$  was evaluated by LRA\* with heuristic  $h_2$ , thus a border node  $\tau_{v_0}^2 \in \mathcal{T}_2$  associated with the vertex  $v_0$  exists. This, in turn, implies that there is a frontier node associated with the vertex  $v_\alpha$ . Let  $\tau_{v_\alpha}^2 \in \mathcal{T}_2$  be this node which was created immediately after edge  $(u, v_0)$  was evaluated. Note that  $c[\tau_{v_\alpha}^1] = c[\tau_{v_\alpha}^2]$ ,  $\ell[\tau_{v_\alpha}^1] = \ell[\tau_{v_\alpha}^2]$  and that  $h_1[\tau_{v_\alpha}^1] = h_1[\tau_{v_\alpha}^2]$ . Since  $h_1$  strictly dominates  $h_2$ , we have that

$$c[\tau_{v_\alpha}^2] + \ell[\tau_{v_\alpha}^2] + h_2[\tau_{v_\alpha}^2] < c[\tau_{v_\alpha}^1] + \ell[\tau_{v_\alpha}^1] + h_1[\tau_{v_\alpha}^1].$$




Figure 6: Different cases considered in the proof of Lemma 1.

From the above, it follows that  $c[\tau_{v_\alpha}^2] + \ell[\tau_{v_\alpha}^2] + h_2[\tau_{v_\alpha}^2] < w^*$ . Therefore,  $\tau_{v_\alpha}^2$  will be popped from  $Q_{\text{frontier}}$  by LRA\* with heuristic  $h_2$  before any node associated with  $v_{\text{target}}$  which, in turn, implies that the edge  $(v_0, v_1)$  will be evaluated by LRA\* with heuristic  $h_2$ .  $\square$

Note that the proof of Lemma 3 assumes that  $h_1$  strictly dominates  $h_2$ . It will *not* hold if  $h_1$  weakly dominates  $h_2$  (namely, if  $\forall v, h_1(v) \geq h_2(v)$ ). Interestingly, there may be cases where both A\* and IDA\* will *not* expand fewer nodes if they use  $h_2$  than if they use  $h_1$  (Holte 2010).

**Lemma 4.** *For every graph  $\mathcal{G}$  and every  $\alpha_1 > \alpha_2$ , we have that  $E_1 \subseteq E_2$ . Here,  $E_i$  denotes the set of edges evaluated by LRA\* with  $\alpha = i$ .*

*Proof.* Assume that  $E_1 \setminus E_2 \neq \emptyset$  and let  $(v_0, v_1) \in E_1 \setminus E_2$  be the edge such that  $v_0$ 's cost-to-come is minimal. Let  $u$  be the parent of  $v_0$  on the shortest path from  $v_{\text{source}}$  to  $v_0$  and note that  $(u, v_0) \in E_1 \cap E_2$ . Let  $\mathcal{T}_i$  denote the search tree of LRA\* with laziness  $\alpha_i$ . Finally, recall that  $w^*(v)$  denotes the (true) weight of the shortest path from  $v_{\text{source}}$  to vertex  $v$ .

The edge  $(v_0, v_1)$  was evaluated by LRA\* with laziness  $\alpha_1$ , thus a frontier node  $\tau_{v_{\alpha_1}}^1 \in \mathcal{T}_1$  associated with some vertex  $v_{\alpha_1}$  was at the head of  $Q_{\text{frontier}}$ . Note that  $v_{\alpha_1}$  is exactly  $\alpha_1$  edges from  $v_0$  and set  $v_1, \dots, v_{\alpha_1-1}$  to be the intermediate vertices along this path. Since node  $\tau_{v_{\alpha_1}}^1$  was popped from  $Q_{\text{frontier}}$ , its key is minimal. Specifically,

$$c[\tau_{v_{\alpha_1}}^1] + \ell[\tau_{v_{\alpha_1}}^1] = w^*(u) + \sum_{i=1}^{\alpha_1} \hat{w}(v_{i-1}, v_i) \leq w^*(v_{\text{target}}).$$

The edge  $(u, v_0)$  was evaluated by LRA\* with laziness  $\alpha_2$ , thus a border node  $\tau_{v_0}^2 \in \mathcal{T}_2$  associated with the vertex  $v_0$  exists. This, in turn, implies that there is a node associated with every vertex that is at most  $\alpha_2$  edges from  $v_0$ , including with the vertex  $v_{\alpha_2}$ . Let  $\tau_{v_{\alpha_2}}^2 \in \mathcal{T}_2$  be this node which was created immediately after edge  $(u, v_0)$  was evaluated. We have that

$$c[\tau_{v_{\alpha_2}}^2] + \ell[\tau_{v_{\alpha_2}}^2] \leq w^*(u) + \sum_{i=1}^{\alpha_2} \hat{w}(v_{i-1}, v_i).$$

Since  $\alpha_2 < \alpha_1$ , we have that  $c[\tau_{v_{\alpha_2}}^2] + \ell[\tau_{v_{\alpha_2}}^2] < w^*(v_{\text{target}})$  which implies that  $\tau_{v_{\alpha_2}}^2$  will be popped from  $Q_{\text{frontier}}$  before any node associated with  $v_{\text{target}}$ . This implies edge  $(v_0, v_1)$  will be evaluated by LRA\* with laziness  $\alpha_2$ . See Fig. 7.  $\square$




Figure 7: Construction used in Lemma 4.

**Lemma 5.** *The total space complexity of our algorithm is bounded by  $O(n + m)$ , where  $n$  and  $m$  are the number of vertices and edges in  $\mathcal{G}$ , respectively.*

*Proof.* Each vertex  $v$  has exactly one node  $\tau \in \mathcal{T}$  associated with it. It appears in a constant number of queues. In addition, the algorithm needs to store the graph  $G$ . Thus, the total space complexity of our algorithm is bounded by  $O(n + m)$ .  $\square$

**Lemma 6.** *The total running time of the algorithm is bounded by  $O(nd^\alpha \cdot \log(n) + m)$ , where  $n$  and  $m$  are the number of vertices and edges,  $d$  is the maximal degree of a vertex and  $\alpha$  is the lookahead.*

*Proof.* Let  $\text{Anc}(v)$  be the set of all vertices that are  $\alpha$  edges from  $v$  and lie on a path between  $v_{\text{source}}$  and  $v$  and note that  $|\text{Anc}(v)| = O(d^\alpha)$ . Furthermore, the number of edges connecting vertices in  $\text{Anc}(v)$  to  $v$  is bounded by  $O(d^\alpha)$ .

We wish to bound the number of times  $\tau_v$  associated with  $v$  will be updated through the algorithm's execution. We charge each update to the event that the algorithm evaluates one of the edges connecting vertices in  $\text{Anc}(v)$  to  $v$ .

Each such update involves updating queues of nodes. The total number of nodes is bounded by  $n$  and the cost of updating the queue is logarithmic in its size. Finally, note that each edge is evaluated at most once.

Thus, the algorithm's running time can be bounded by

$$\underbrace{O(n)}_{\text{\# of nodes}} \cdot \underbrace{O(d^\alpha)}_{\text{\# of node updates}} \cdot \underbrace{O(\log n)}_{\text{cost of node update}} + \underbrace{O(m)}_{\text{\# of edges}},$$

which concludes the proof.  $\square$




Figure 8: Example where given a larger lookahead results in *more* edge evaluations. Top (a-d) and bottom (e-h) depict the flow of LRA\* for lazy lookahead value of  $\alpha = 4$  and  $\alpha = 3$ , respectively and a greediness value of  $\beta = 3$ . Each figure depicts one iteration, blue edges are found to be collision free while red edges are found to be in collision. Value of frontier nodes are shown at each iteration and all the edges evaluated by each algorithm are shown in (d) and (h).

### C Is greediness beneficial?—Theorems

In this section we provide accompanying theorems and proofs to the discussion regarding greediness presented in Section 5.6. We start by noting that Lemma 1 and 2 can be easily extended to the case where greediness is employed. We next continue by fixing one parameter (either lazy lookahead  $\alpha$  or greediness  $\beta$ ) and show under what conditions of the other parameter ( $\beta$  or  $\alpha$ , respectively) LRA\* performs less edge evaluations. Recall (Lemma 4) that where the algorithm has no greediness, namely,  $\beta = 1$ . then when considering edge evaluations, the larger the lookahead, the better the algorithm. This can be seen as warm-up for Lemma 7 which gives a general relationship between different lookaheads  $\alpha$  for a fixed greediness  $\beta$ . We then move on to fix  $\alpha$  and see how varying  $\beta$  affects the algorithm. In Lemma 8 we show that for a fixed lookahead, no greediness ( $\beta = 1$ ) is always better (in terms of edge evaluations) when compared to any other value of  $\beta$ . Finally, in Lemma 9 we show the somewhat counter-intuitive result that for larger greediness values ( $\beta > 1$ ) and fixed lookahead, there is always an example where the greater the greediness, the better.

We start by noting that if  $\beta > 1$  it may be the case that the larger the lookahead, the better (when considering edge evaluations). In a nutshell, the greediness  $\beta$  may drive the algorithm to evaluate edges along paths that, at first glance, seem promising but as the algorithm evaluates edges, it becomes evident that other paths are more promising. See Fig. 8 for an example. In the next Lemma, we show under what conditions (for  $\beta > 1$ ) this natural behaviour does indeed hold.

**Lemma 7.** *For every graph  $G$  and every  $\alpha_1 > \alpha_2 \geq \beta$ , we have that  $E_1 \subseteq E_2$  if  $\alpha_1 \geq \alpha_2 + \beta - 1$ . Here,  $E_i(\beta)$  denotes the set of edges evaluated by LRA\* with laziness  $i$  and greediness  $\beta$ .*

*Proof.* Assume that  $E_1 \setminus E_2 \neq \emptyset$  and let  $(v_0, v) \in E_1 \setminus E_2$  be an edge such that  $v_0$ 's cost-to-come is minimal. Note that this implies that both algorithms compute the shortest path to  $v_0$ . Furthermore, let  $\mathcal{T}_i$  denote the search tree of LRA\* with laziness  $i$  and greediness  $\beta$ .

Consider the iteration before LRA\* with laziness  $\alpha_1$  and greediness  $\beta$  evaluates  $(v_0, v)$  and let

$(v_{\beta-1}, v_{\beta-2}, \dots, v_1, v_0, v)$  be the sequence of vertices along the  $\beta$  edges lying on the shortest path from  $v_{\text{start}}$  to  $v$ . Since the edge  $(v_0, v)$  is evaluated, there exists a border node  $\tau_i \in \mathcal{T}_1$  associated with  $v_i$  where  $0 \leq i \leq \beta - 1$ . Furthermore, the lazy path from  $\tau_i$  that passes through  $v_i$  was considered, hence there is a node  $\tau \in \mathcal{T}_1$  which is  $\alpha_1$  edges from  $\tau_i$  whose key is minimal. Namely  $\bar{w}(P[\tau]) < w^*$  with  $w^* = w^*(v_{\text{target}})$  the minimal cost to reach  $v_{\text{target}}$ . Note that this path  $P[\tau]$  contains the edge  $(v_0, v)$ .

Clearly,  $\mathcal{T}_2$  contains a border node associated with  $v_0$ . LRA\* with laziness  $\alpha_2$  and greediness  $\beta$  does not expand any path from  $\tau_0$  that contains the edge  $(v_0, v)$ , thus all paths  $\alpha_2$  edges away from  $\tau_0$  passing through  $v$  have lazy cost larger than  $w^*$ . However, the node  $\tau$  (which caused LRA\* with laziness  $\alpha_1$  and greediness  $\beta$  to evaluate  $(v_0, v)$ ) is  $\alpha_1 - i \geq \alpha_1 - (\beta - 1) \geq \alpha_2$  edges from  $v_0$ . We know that  $\bar{w}(P[\tau]) < w^*$  thus  $(v_0, v)$  should have been evaluated by LRA\* with laziness  $\alpha_2$  and greediness  $\beta$  which gives us a contradiction. For a visualization, see Fig. 9a.  $\square$

We now move to the case where the lookahead  $\alpha$  is fixed and we compare the edge evaluation of LRA\* for different values of  $\beta$ . We start with the simple case where one algorithm has a greediness value of  $\beta = 1$ .

**Lemma 8.** *For every graph  $\mathcal{G}$  and every  $\alpha \geq \beta > 1$ , we have that  $E_1 \subseteq E_\beta$ . Here  $E_x$  denotes the edges set of LRA\* with laziness  $\alpha$  and greediness  $x$ .*

*Proof.* Assume that  $E_1 \setminus E_\beta \neq \emptyset$  and let  $(v_0, v) \in E_1 \setminus E_\beta$  be an edge such that  $v_0$ 's cost-to-come is minimal. Note that this implies both algorithms compute the shortest path to  $v_0$ . Furthermore, let  $\mathcal{T}_x$  denote the search tree of LRA\* with greediness  $x$ .

Consider the iteration before LRA\* with greediness 1 (no greediness) evaluates  $(v_0, v)$ . Since the edge  $(v_0, v)$  is evaluated, the node  $\tau_0 \in \mathcal{T}_1$  associated with  $v_0$  was the border node and there exists a node  $\tau \in \mathcal{T}_1$  which is  $\alpha$  edges from  $\tau_0$  whose key is minimal. Namely,  $\bar{w}(P[\tau]) < w^*$  with  $w^* = w^*(v_{\text{target}})$  the minimal cost to reach  $v_{\text{target}}$ . Note that the path  $P[\tau]$  contains the edge  $(v_0, v)$ .

Now, consider the search tree  $\mathcal{T}_\beta$  of LRA\* with greediness  $\beta$ . Clearly,  $\mathcal{T}_\beta$  contains a border node  $\tau'_0$  with  $u[\tau'_0] = v_0$ .

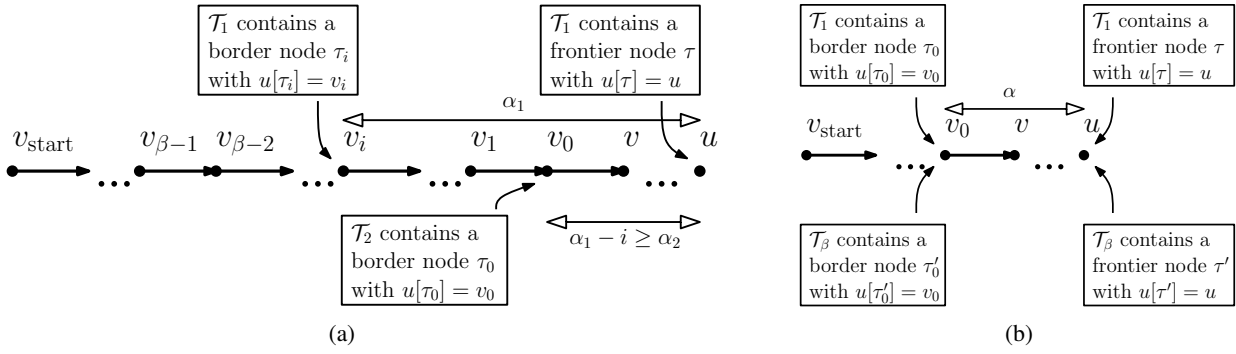


Figure 9: Constructions used in Lemma 7 (Fig. a) and Lemma 8 (Fig. b)

There exists a node  $\tau'$  with  $u[\tau'] = u[\tau]$ . Namely, the node  $\tau'$  which is exactly  $\alpha$  edges away from  $\tau'_0$  has  $\bar{w}(P[\tau']) < w^*$  where  $P[\tau']$  contains the edge  $(v_0, v)$ . Hence the node  $\tau'$  would be popped from  $Q_{\text{frontier}}$  before any node associated with  $v_{\text{target}}$  implying the edge  $(v_0, v)$  would be evaluated giving us a contradiction. For a visualization, see Fig. 9b.  $\square$

We continue to examine the general case where the lookahead  $\alpha$  is fixed and we compare the edge evaluation of LRA\* for different values of  $\beta$  for  $\beta > 1$ . Intuitively, we would expect some result stating that the smaller the greediness the better (in terms of maximal number of edge evaluations). Indeed, in Lemma 8 we showed that this is the case when  $\beta = 1$ . However, the following Lemma states that for general values of  $\beta$ , there exists cases where an algorithm with large greediness may outperform an algorithm with smaller greediness.

**Lemma 9.** *For every lookahead  $\alpha < \infty$  and every greediness  $\alpha \geq \beta_2 > \beta_1 > 1$ , there exists a graph  $\mathcal{G}$  where  $E_{\beta_1} \setminus E_{\beta_2} \neq \emptyset$ . Here,  $E_\beta$  denotes the set of edges evaluated by LRA\* with laziness  $\alpha$  and greediness  $\beta$ .*

*Proof.* We construct the graph  $\mathcal{G}$  explicitly. (Fig. 10 and 11.

We consider two following two cases (i)  $\beta_2 \bmod \beta_1 \neq 0$  and (ii)  $\beta_2 \bmod \beta_1 = 0$ . For each case we provide a different graph  $\mathcal{G}$  and show that  $E_{\beta_1} \setminus E_{\beta_2} \neq \emptyset$ . See Fig. 10 and 11 for depictions of each case described.

For case (i) where  $\beta_2 \bmod \beta_1 \neq 0$ , we have a path of length  $\beta_2$  followed by two paths of length  $\alpha$ . For LRA\* with greediness  $\beta_1$ , (Fig. 10a-10d), the algorithm starts by evaluating edges along the path of length  $\beta_2$  (Fig. 10a). Since  $\beta_2 \bmod \beta_1 \neq 0$ , at some point it will evaluate the first edge along the upper path (which is in collision) (Fig. 10b). This path is longer than the lower one, but to see this, the algorithm requires a lookahead of  $\alpha$  edges from the end of the first path. The algorithm continues to evaluate edges along the lower path until the target is reached (Fig. 10c).

For LRA\* with greediness  $\beta_2$  (Fig. 10e-10g), the algorithm starts by evaluating all edges along the path of length  $\beta_2$  (Fig. 10a). Since it can see all edges along the upper path (which is collision free) it continues to evaluate the lower path until the target is reached (Fig. 10f). The final edges evaluated by each algorithm are depicted in Fig. 10d and 10g.

For case (ii) where  $\beta_2 \bmod \beta_1 = 0$ , we have a path of length  $\beta_2$  which after one edge has a shorter path of  $\alpha$  edges. The rest of the construction is similar to case (i). Essentially, LRA\* with greediness  $\beta_1$  and LRA\* with greediness  $\beta_2$  behave similarly to case (i) except that both algorithms will evaluate the first edge along the path of length  $\beta_2$  followed by the first (in-collision edge) of the shorter path of  $\alpha$  edges. After this first iteration for both algorithms (Fig. 11a and 11e) the behaviour reduces to that of case (i).  $\square$

## References

- [Bialkowski et al. 2016] Bialkowski, J.; Otte, M. W.; Karaman, S.; and Frazzoli, E. 2016. Efficient collision checking in sampling-based motion planning via safety certificates. *I. J. Robotics Res.* 35(7):767–796.
- [Bohlin and Kavraki 2000] Bohlin, R., and Kavraki, L. E. 2000. Path planning using lazy PRM. In *ICRA*, volume 1, 521–528. IEEE.
- [Chen et al. 2007] Chen, M.; Chowdhury, R. A.; Ramachandran, V.; Roche, D. L.; and Tong, L. 2007. *Priority queues and Dijkstra’s algorithm*. UTCS Technical Report TR-07-54.
- [Choset et al. 2005] Choset, H.; Lynch, K. M.; Hutchinson, S.; Kantor, G.; Burgard, W.; Kavraki, L. E.; and Thrun, S. 2005. *Principles of Robot Motion: Theory, Algorithms, and Implementation*. MIT Press.
- [Choudhury et al. 2017] Choudhury, S.; Salzman, O.; Choudhury, S.; and Srinivasa, S. S. 2017. Densification strategies for anytime motion planning over large dense roadmaps. In *ICRA*, 3770–3777.
- [Choudhury, Dellin, and Srinivasa 2016] Choudhury, S.; Dellin, C. M.; and Srinivasa, S. S. 2016. Pareto-optimal search over configuration space beliefs for anytime motion planning. In *IROS*, 3742–3749.
- [Cohen, Phillips, and Likhachev 2014] Cohen, B. J.; Phillips, M.; and Likhachev, M. 2014. Planning single-arm manipulations with n-arm robots. In *RSS*.
- [Dechter and Pearl 1983] Dechter, R., and Pearl, J. 1983. The optimality of A\* revisited. In *AAAI*, 95–99.
- [Dellin and Srinivasa 2016] Dellin, C. M., and Srinivasa, S. S. 2016. A unifying formalism for shortest path problems

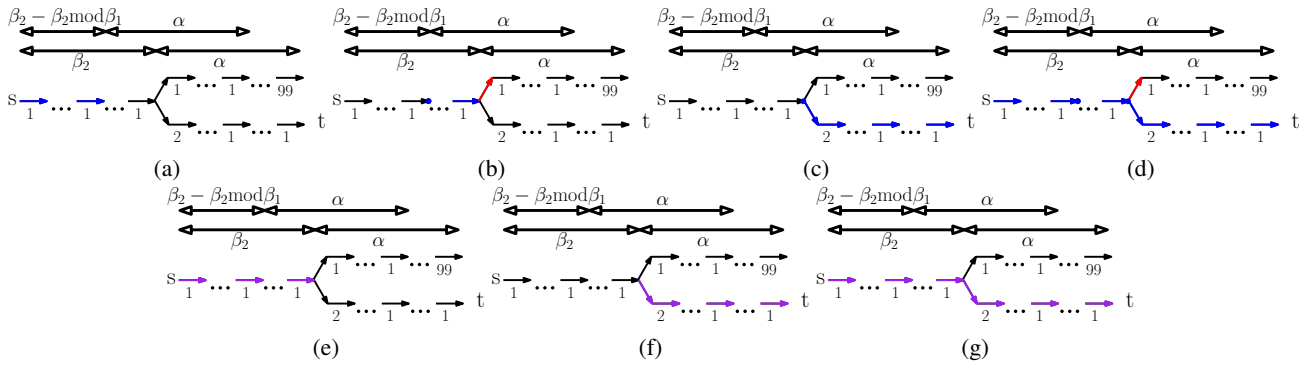


Figure 10: Construction used in Lemma 9, when  $\beta_2 \bmod \beta_1 \neq 0$ . The flow of LRA\* with greediness  $\beta_1$  is depicted in (a-d). Edges found to be collision free and in collision are depicted in blue and red, respectively. The flow of LRA\* with greediness  $\beta_2$  is depicted in (e-g). Edges found to be collision free and in collision are depicted in purple and red, respectively. All the edges evaluated by each algorithm are shown in (d) and (g).

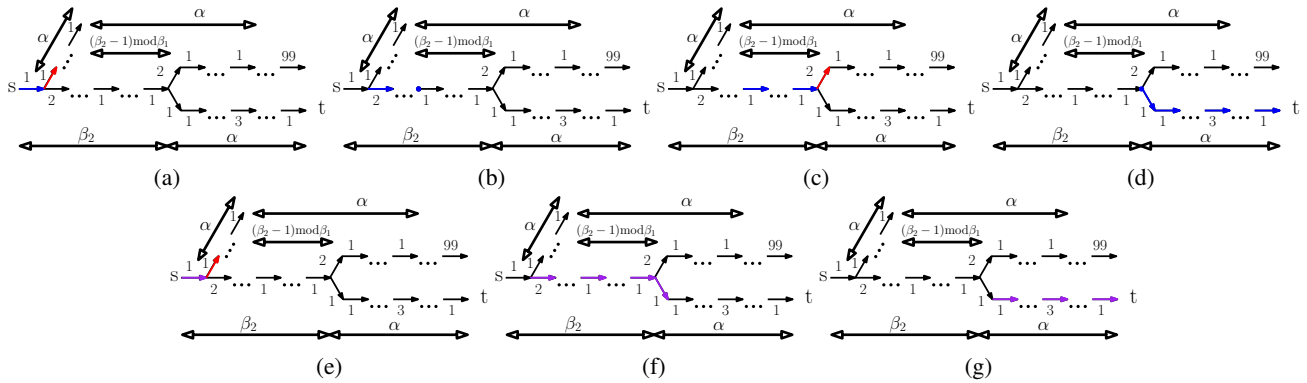


Figure 11: Construction used in Lemma 9, when  $\beta_2 \bmod \beta_1 = 0$ . The flow of LRA\* with greediness  $\beta_1$  is depicted in (a-d). Edges found to be collision free and in collision are depicted in blue and red, respectively. The flow of LRA\* with greediness  $\beta_2$  is depicted in (e-g). Edges found to be collision free and in collision are depicted in purple and red, respectively. All the edges evaluated by each algorithm are shown in (d) and (g).

with expensive edge evaluations via lazy best-first search over paths with edge selectors. In *ICAPS*, 459–467.

[Dijkstra 1959] Dijkstra, E. W. 1959. A note on two problems in connexion with graphs. *Numer. Math.* 1(1):269–271.

[Dobson and Bekris 2014] Dobson, A., and Bekris, K. E. 2014. Sparse roadmap spanners for asymptotically near-optimal motion planning. *I. J. Robotics Res.* 33(1):18–47.

[Gammell, Srinivasa, and Barfoot 2015] Gammell, J. D.; Srinivasa, S. S.; and Barfoot, T. D. 2015. Batch informed trees (BIT\*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs. In *ICRA*, 3067–3074.

[Haghtalab et al. 2017] Haghtalab, N.; Mackenzie, S.; Proccaccia, A. D.; Salzman, O.; and Srinivasa, S. S. 2017. The provable virtue of laziness in motion planning. *CoRR* abs/1710.04101.

[Halton 1964] Halton, J. H. 1964. Algorithm 247: Radical-inverse quasi-random point sequence. *Commun. ACM* 7(12):701–702.

[Hart, Nilsson, and Raphael 1968] Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968. A formal basis for the heuristic de-

termination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics* 4(2):100–107.

[Hauser 2015] Hauser, K. 2015. Lazy collision checking in asymptotically-optimal motion planning. In *ICRA*, 2951–2957.

[Holte 2010] Holte, R. C. 2010. Common misconceptions concerning heuristic search. In *SoCS*, 46–51.

[Janson et al. 2015] Janson, L.; Schmerling, E.; Clark, A. A.; and Pavone, M. 2015. Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions. *I. J. Robotics Res.* 34(7):883–921.

[Karaman and Frazzoli 2011] Karaman, S., and Frazzoli, E. 2011. Sampling-based algorithms for optimal motion planning. *I. J. Robotics Res.* 30(7):846–894.

[Kavraki et al. 1996] Kavraki, L. E.; Svestka, P.; Latombe, J.-C.; and Overmars, M. H. 1996. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robotics and Automation* 12(4):566–580.

[Koenig and Sun 2009] Koenig, S., and Sun, X. 2009. Comparing real-time and incremental heuristic search for real-

- time situated agents. *Autonomous Agents and Multi-Agent Systems* 18(3):313–341.
- [Koenig, Likhachev, and Furcy 2004] Koenig, S.; Likhachev, M.; and Furcy, D. 2004. Lifelong planning A\*. *Artif. Intell.* 155(1-2):93–146.
- [Korf 1985] Korf, R. E. 1985. Depth-first iterative-deepening: An optimal admissible tree search. *Artificial Intelligence* 27:97–109.
- [Korf 1990] Korf, R. E. 1990. Real-time heuristic search. *Artificial Intelligence* 42(2):189 – 211.
- [Kwon and Han 2006] Kwon, W. H., and Han, S. H. 2006. *Receding horizon control: model predictive control for state models*. Springer Science & Business Media.
- [LaValle 2006] LaValle, S. M. 2006. *Planning Algorithms*. Cambridge University Press.
- [Pearl 1984] Pearl, J. 1984. *Heuristics - intelligent search strategies for computer problem solving*. Addison-Wesley series in artificial intelligence.
- [Reif 1979] Reif, J. H. 1979. Complexity of the mover’s problem and generalizations. In *FOCS*, 421–427. IEEE.
- [Salzman and Halperin 2015] Salzman, O., and Halperin, D. 2015. Asymptotically-optimal motion planning using lower bounds on cost. In *ICRA*, 4167–4172.
- [Salzman and Halperin 2016] Salzman, O., and Halperin, D. 2016. Asymptotically near-optimal RRT for fast, high-quality motion planning. *IEEE Trans. Robotics* 32(3):473–483.
- [Sharir 2004] Sharir, M. 2004. Algorithmic motion planning. In *Handbook of Discrete and Computational Geometry, Second Edition*. 1037–1064.
- [Srinivasa et al. 2009] Srinivasa, S. S.; Ferguson, D.; Helfrich, C. J.; Berenson, D.; Collet, A.; Diankov, R.; Gallagher, G.; Hollinger, G.; Kuffner, J.; and Weghe, M. V. 2009. HERB: a home exploring robotic butler. *Autonomous Robots* 28(1):5.
- [Sucan, Moll, and Kavraki 2012] Sucan, I. A.; Moll, M.; and Kavraki, L. E. 2012. The open motion planning library. *IEEE Robotics Automation Magazine* 19(4):72–82.
- [Yoshizumi, Miura, and Ishida 2000] Yoshizumi, T.; Miura, T.; and Ishida, T. 2000. A\* with partial expansion for large branching factor problems. In *AAAI*, 923–929.