Asymptotically near-optimal RRT
for fast, high-quality, motion planning

Oren Salzman and Dan Halperin
Blavatnik School of Computer Science, Tel-Aviv University, Israel

Abstract—We present *Lower Bound Tree-RRT* (LBT-RRT), a single-query sampling-based algorithm that is *asymptotically near-optimal*. Namely, the solution extracted from LBT-RRT converges to a solution that is within an approximation factor of $1 + \varepsilon$ of the optimal solution. Our algorithm allows for a continuous interpolation between the fast RRT algorithm and the asymptotically optimal RRT* and RRG algorithms. When the approximation factor is 1 (i.e., no approximation is allowed), LBT-RRT behaves like RRG. When the approximation factor is unbounded, LBT-RRT behaves like RRT. In between, LBT-RRT is shown to produce paths that have higher quality than RRT would produce and run faster than RRT* would run. This is done by maintaining a tree which is a sub-graph of the RRG roadmap and a second, auxiliary graph, which we call the *lower-bound* graph. The combination of the two roadmaps, which is faster to maintain than the roadmap maintained by RRT*, efficiently guarantees asymptotic near-optimality. We suggest to use LBT-RRT for high-quality, anytime motion planning. We demonstrate the performance of the algorithm for scenarios ranging from 3 to 12 degrees of freedom and show that even for small approximation factors, the algorithm produces high-quality solutions (comparable to RRG and RRT*) with little running-time overhead when compared to RRT.

I. INTRODUCTION AND RELATED WORK

Motion planning is a fundamental research topic in robotics with applications in diverse domains such as surgical planning, computational biology, autonomous exploration, search-and-rescue, and warehouse management. Sampling-based planners such as PRM [1], RRT [2] and their many variants enabled solving motion-planning problems that had been previously considered infeasible [3, C.7]. Recently, there is growing interest in the robotics community in finding *high-quality* paths, which turns out to be a non-trivial problem [4], [5]. Quality can be measured in terms of, for example, length, clearance, smoothness, energy, to mention a few criteria, or some combination of the above.

A. High-quality planning with sampling-based algorithms

Unfortunately, planners such as RRT and PRM produce solutions that may be far from optimal [4], [5]. Thus, many

variants of these algorithms and heuristics were proposed in order to produce high-quality paths.

Post-processing existing paths: Post-processing an existing path by applying *shortcutting* is a common, effective, approach to increase path quality; see, e.g., [6]. Typically, two non-consecutive configurations are chosen randomly along the path. If the two configurations can be connected using a straight-line segment in the configuration space and this connection improves the quality of the original path, the segment replaces the original path that connected the two configurations. The process is continued iteratively until a termination condition holds.

Path hybridization: An inherent problem with path post-processing is that it is local in nature. A path that was post-processed using shortcutting often remains in the same homotopy class of the original path. Carefully combining even a small number of different paths (that may be of low quality) often enables the construction of a higher-quality path [7].

Online optimization: Changing the sampling strategy [8], [9], [10], [11], or the connection scheme to a new milestone [10], [12] are examples of heuristics proposed to create higher-quality solutions. Additional approaches include, among others, useful cycles [6] and random restarts [13].

Asymptotically optimal and near-optimal solutions: In their seminal work, Karaman and Frazzoli [4] give a rigorous analysis of the performance of the RRT and PRM algorithms. They show that with probability one, the algorithms will not produce the optimal path. By modifying the connection scheme of a new sample to the existing data structure, they propose the PRM* and the RRG and RRT* algorithms (variants of the PRM and RRT algorithms, respectively) all of which are shown to be *asymptotically optimal*. Namely, as the number of samples tends to infinity, the solution obtained by these algorithms converges to the optimal solution with probability one. To ensure asymptotic optimality, the number of nodes each new sample is connected to is proportional to $\log(n)$ (here n is the number of free samples).

As PRM* may produce prohibitively large graphs, recent work has focused on sparsifying these graphs. This can be done as a post-processing stage of the PRM* [14], [15], or as a modification of PRM* [16], [17], [18].

The performance of RRT* can be improved using several heuristics that bear resemblance to the lazy approach

used in this work [19]. Additional heuristics to speed up the convergence rate of RRT* were presented in RRT*-SMART [20]. Recently, RRT$^{\#}$ [21] was suggested as an asymptotically-optimal algorithm with a faster convergence rate when compared to RRT*. RRT$^{\#}$ extends its roadmap in a similar fashion to RRT* but adds a replanning procedure. This procedure ensures that the tree rooted at the initial state contains lowest-cost path information for vertices which have the potential to be part of the optimal solution. Thus, in contrast to RRT* which only performs *local* rewiring of the search tree, RRT$^{\#}$ efficiently propagates changes to *all* the relevant parts of the roadmap. Janson and Pavone [22] introduced the asymptotically-optimal Fast Marching Tree algorithm (FMT*). The single-query asymptotically-optimal algorithm maintains a tree as its roadmap. Similarly to PRM*, FMT* samples n collision-free nodes. It then builds a minimum-cost spanning tree rooted at the initial configuration over this set of nodes (see Section E for further details). Lazy variants have been proposed both for PRM* and RRG [23] and for FMT* [24].

An alternative approach to improve the running times of these algorithms is to relax the asymptotic optimality to *asymptotic near-optimality*. An algorithm is said to be asymptotically near-optimal if, given an *approximation factor* ε, the solution obtained by the algorithm converges to within a factor of $(1 + \varepsilon)$ of the optimal solution with probability one, as the number of samples tends to infinity. Similar to this work, yet using different methods, Littlefield et al. [25] recently presented an asymptotic near-optimal variant of RRT* for systems with dynamics. Their approach however, requires setting different parameters used by their algorithm.

Anytime and online solutions: An interesting variant of the basic motion-planning problem is anytime motion-planning: In this problem, the time to plan is not known in advance, and the algorithm may be terminated at any stage. Clearly, any solution should be found as fast as possible and if time permits, it should be refined to yield a higher-quality solution.

Ferguson and Stentz [26] suggest iteratively running RRT while considering only areas that may potentially improve the existing solution. Alterovitz et al. [27] suggest the Rapidly-exploring Roadmap Algorithm (RRM), which finds an initial path similar to RRT. Once such a path is found, RRM either explores further the configuration space or refines the explored space. Luna et al. [28] suggest alternating between path shortcutting and path hybridization in an anytime fashion.

RRT* was also adapted for online motion planning [29]. Here, an initial path is computed and the robot begins its execution. While the robot moves along this path, the algorithm refines the part that the robot has not yet moved along.

B. Contribution

We present LBT-RRT, a single-query sampling-based algorithm that is *asymptotically near-optimal*. Namely, the solution extracted from LBT-RRT converges to a solution that is within a factor of $(1+\varepsilon)$ of the optimal solution. LBT-RRT allows for interpolating between the fast, yet sub-optimal, RRT algorithm

and the asymptotically-optimal RRG algorithm. By choosing $\varepsilon = 0$ no approximation is allowed and LBT-RRT maintains a roadmap identical to the one maintained by RRG. Choosing $\varepsilon = \infty$ allows for any approximation and LBT-RRT maintains a tree identical to the tree maintained by RRT.

The asymptotic near-optimality of LBT-RRT is achieved by simultaneously maintaining two roadmaps. Both roadmaps are defined over the same set of vertices but each consists of a different set of edges. On the one hand, a path in the first roadmaps may not be feasible, but its cost is always a *lower bound* on the cost of paths extracted from RRG (using the same sequence of random nodes). On the other hand, a path extracted from the second roadmap is always feasible and its cost is within a factor of $(1+\varepsilon)$ from the lower bound provided by the first roadmap.

We suggest to use LBT-RRT for high-quality, anytime motion planning. We demonstrate its performance on scenarios ranging from 3 to 12 degrees of freedom (DoF) and show that the algorithm produces high-quality solutions (comparable to RRG and RRT*) with little running-time overhead when compared to RRT.

This paper is a modified and extended version of a publication presented at the 2014 IEEE International Conference on Robotics and Automation [30]. In this paper we present additional experiments and extensions of the original algorithmic framework. Finally, we note that the conference version of this paper contained an oversight with regard to the roadmap that is used for the lower bound. We explain the problem and its fix in detail in Section III after providing all the necessary technical background.

C. Outline

In Section II we review the RRT, RRG and RRT* algorithms. In Section III we present our algorithm LBT-RRT and a proof of its asymptotic near-optimality. We continue in Section IV to demonstrate in simulations its favorable characteristics on several scenarios. In Section V we discuss a modification of the framework to further speed up the convergence to high-quality solutions. We conclude in Section VI by describing possible directions for future work. In the appendix we list several applications where either RRT or RRT* were used and argue that LBT-RRT may serve as a superior alternative with no fundamental modification to the underlying algorithms using RRT or RRT*. Moreover, we discuss alternative implementations of LBT-RRT using tools developed for either RRT or RRT* that can enhance LBT-RRT. Finally, we demonstrate how the framework presented in this paper for relaxing the optimality of RRG can be used to have a similar effect on another asymptotically-optimal sampling-based algorithm, FMT* [22].

II. TERMINOLOGY AND ALGORITHMIC BACKGROUND

We begin this section by formally stating the motion-planning problem and introducing several standard procedures used by sampling-based algorithms. We continue by reviewing the RRT, RRG and RRT* algorithms.

Algorithm 1 RRT (x_{init})

1: $\mathcal{T}.V \leftarrow \{x_{\text{init}}\}$
2: **while** `construct_roadmap()` **do**

3:    $x_{\text{rand}} \leftarrow$ `sample_free()`
4:    $x_{\text{nearest}} \leftarrow$ `nearest_neighbor`$(x_{\text{rand}}, \mathcal{T}.V)$
5:    $x_{\text{new}} \leftarrow$ `steer`$(x_{\text{nearest}}, x_{\text{rand}})$

6:    **if** (!`collision_free`$(x_{\text{nearest}}, x_{\text{new}})$) **then**
7:       CONTINUE

8:    $\mathcal{T}.V \leftarrow \mathcal{T}.V \cup \{x_{\text{new}}\}$
9:    $\mathcal{T}$.`parent`$(x_{\text{new}}) \leftarrow x_{\text{nearest}}$

Algorithm 2 RRG (x_{init})

1: $\mathcal{G}.V \leftarrow \{x_{\text{init}}\}$ $\mathcal{G}.E \leftarrow \emptyset$
2: **while** `construct_roadmap()` **do**

3:    $x_{\text{rand}} \leftarrow$ `sample_free()`
4:    $x_{\text{nearest}} \leftarrow$ `nearest_neighbor`$(x_{\text{rand}}, \mathcal{G}.V)$
5:    $x_{\text{new}} \leftarrow$ `steer`$(x_{\text{nearest}}, x_{\text{rand}})$

6:    **if** (!`collision_free`$(x_{\text{nearest}}, x_{\text{new}})$) **then**
7:       CONTINUE

8:    $\mathcal{G}.V \leftarrow \mathcal{G}.V \cup \{x_{\text{new}}\}$
9:    $\mathcal{G}.E \leftarrow \{(x_{\text{nearest}}, x_{\text{new}}), (x_{\text{new}}, x_{\text{nearest}})\}$
10:   $X_{\text{near}} \leftarrow$ `nearest_neighbors`$(x_{\text{new}},$
               $\mathcal{G}.V, k_{RRG} \log(|\mathcal{G}.V|))$
11:   **for all** $(x_{\text{near}}, X_{\text{near}})$ **do**
12:      **if** (`collision_free`$(x_{\text{near}}, x_{\text{new}})$) **then**
13:         $\mathcal{G}.E \leftarrow \{(x_{\text{near}}, x_{\text{new}}), (x_{\text{new}}, x_{\text{near}})\}$

Algorithm 3 RRT* (x_{init})

1: $\mathcal{T}.V \leftarrow \{x_{\text{init}}\}$
2: **while** `construct_roadmap()` **do**

3:    $x_{\text{rand}} \leftarrow$ `sample_free()`
4:    $x_{\text{nearest}} \leftarrow$ `nearest_neighbor`$(x_{\text{rand}}, \mathcal{G}.V)$
5:    $x_{\text{new}} \leftarrow$ `steer`$(x_{nearest}, x_{rand})$

6:    **if** (!`collision_free`$(x_{\text{nearest}}, x_{\text{new}})$) **then**
7:       CONTINUE

8:    $\mathcal{T}.V \leftarrow \mathcal{T}.V \cup \{x_{\text{new}}\}$
9:    $\mathcal{T}$.`parent`$(x_{\text{new}}) \leftarrow x_{\text{nearest}}$

10:   $X_{\text{near}} \leftarrow$ `nearest_neighbors`$(x_{\text{new}},$
               $\mathcal{T}.V, k_{RRG} \log(|\mathcal{T}.V|))$

11:   **for all** $(x_{\text{near}}, X_{\text{near}})$ **do**
12:      `rewire_RRT*`$(x_{\text{near}}, x_{\text{new}}$)

13:   **for all** $(x_{\text{near}}, X_{\text{near}})$ **do**
14:      `rewire_RRT*`$(x_{\text{new}}, x_{\text{near}}$)

A. Problem definition and terminology

We follow the formulation of the motion-planning problem as presented by Karaman and Frazzoli [4]. Let $\mathcal{X}$ denote the configuration space (C-space), $\mathcal{X}_{\text{free}}$ and $\mathcal{X}_{\text{forb}}$ denote the free and forbidden spaces, respectively. Let $(\mathcal{X}_{\text{free}}, x_{\text{init}}, \mathcal{X}_{goal})$ be the motion-planning problem where: $x_{\text{init}} \in \mathcal{X}_{\text{free}}$ is an initial free configuration and $\mathcal{X}_{goal} \subseteq \mathcal{X}_{\text{free}}$ is the goal region. A *collision-free path* $\sigma : [0,1] \to \mathcal{X}_{\text{free}}$ is a continuous mapping to the free space. It is *feasible* if $\sigma(0) = x_{\text{init}}$ and $\sigma(1) \in X_{goal}$.

We will make use of the following procedures throughout the paper: `sample_free`, a procedure returning a random free configuration; `nearest_neighbor`(x, V) and `nearest_neighbors`(x, V, k) are procedures returning the nearest neighbor and k nearest neighbors of x within the set V, respectively. Let `steer`(x, y) return a configuration z that is closer to y than x is, `collision_free`(x, y) tests if the straight line segment connecting x and y is contained in $\mathcal{X}_{\text{free}}$ and let `cost`(x, y) be a procedure returning the cost of the straight-line path connecting x and y. Let us denote by $\text{cost}_{\mathcal{G}}(x)$ the minimal cost of reaching a node x from x_{init} using a roadmap $\mathcal{G}$. These are standard procedures used by the RRT or RRT* algorithms. Finally, we use the (generic) predicate `construct_roadmap` to assess if a stopping criterion has been reached to terminate the algorithm[1].

B. Algorithmic background

The RRT, RRG and RRT* algorithms share the same high-level structure. They maintain a roadmap as the underlying data structure which is a directed tree for RRT and RRT* and a directed graph for RRG. At each iteration a configuration x_{rand} is sampled at random. Then, x_{nearest}, the nearest configuration to x_{rand} in the roadmap is found and extended in the direction of x_{rand} to a new configuration x_{new}. If the path between x_{nearest} and x_{new} is collision-free, x_{new} is added to the roadmap (see Alg. 1, 2 and 3, lines 3-9).

The algorithms differ in the connections added to the roadmap. In RRT, only the edge $(x_{\text{nearest}}, x_{\text{new}})$ is added. In RRG and RRT*, a set X_{near} of $k_{RRG} \log(|V|)$ nearest neighbors of x_{new} is considered. Here, k_{RRG} is a constant

---

[1]A stopping criterion can be, for example, reaching a certain number of samples or exceeding a fixed time budget.

---

ensuring that the cost of paths produced by RRG and RRT* indeed converges to the optimal cost almost surely as the number of samples grows. A valid choice for all problem instances is $k_{RRG} = 2e$ [4]. For each neighbor $x_{\text{near}} \in X_{\text{near}}$ of x_{new}, RRG checks if the path between x_{near} and x_{new} is collision-free and if so, $(x_{\text{near}}, x_{\text{new}})$ and $(x_{\text{new}}, x_{\text{near}})$ are added

---

Algorithm 4 `rewire_RRT*`$(x_{\text{potential\_parent}}, x_{\text{child}})$

1: **if** (`collision_free`$(x_{\text{potential\_parent}}, x_{\text{child}})$) **then**
2:    $c \leftarrow$ `cost`$(x_{\text{potential\_parent}}, x_{\text{child}})$
3:    **if** $(\text{cost}_{\mathcal{T}}(x_{\text{potential\_parent}}) + c < \text{cost}_{\mathcal{T}}(x_{\text{child}}))$ **then**
4:       $\mathcal{T}$.`parent`$(x_{\text{child}}) \leftarrow x_{\text{potential\_parent}}$

to the roadmap (lines 10-13). RRT* maintains a sub-graph of the RRG roadmap. This is done by an additional rewiring procedure (Alg. 4) which is invoked twice: The first time, it is used to find the node $x_{\text{near}} \in X_{\text{near}}$ which will minimize the cost to reach x_{new} (Alg. 3, lines 11-12). The second time, the procedure is used to to minimize the cost to reach every node $x_{\text{near}} \in X_{\text{near}}$ by considering x_{new} as its parent (Alg. 3, lines 13-14). Thus, at all time, RRT* maintains a tree which, as mentioned, is a subgraph of the RRG roadmap.

Given a sequence of n random samples, the cost of the path obtained using the RRG algorithm is a lower bound on the cost of the path obtained using the RRT* algorithm. However, RRG requires both additional memory (to explicitly store the set of $O(\log n)$ neighbours) and exhibits longer running times (due to the additional calls to the local planner). In practice, this excess in running time is far from negligible (see Section IV), making RRT* a more suitable algorithm for asymptotically-optimal motion planning.

III. ASYMPTOTICALLY NEAR-OPTIMAL MOTION-PLANNING

Clearly the asymptotic optimality of the RRT* and RRG algorithms comes at the cost of the additional $O(k_{RRG} \log(|V|))$ calls to the local planner at each stage (and some additional overhead). If we are not concerned with *asymptotically optimal* solutions, we do not have to consider all of the $k_{RRG} \log(|V|)$ neighbors when a node is added. Our idea is to initially only *estimate* the quality of each edge. We use this estimate of the quality of the edge to decide if to discard it, use it *without* checking if it is collision-free or use it after validating that it is indeed collision-free. Thus, many calls to the local planner can be avoided, though we still need to estimate the quality of many edges. Our approach is viable in cases where such an assessment can be carried out efficiently. Namely, more efficiently than deciding if an edge is collision-free. This condition holds naturally when the quality measure is *path length* which is the cost function considered in this paper; for a discussion on different cost functions, see Section VI.

A. Single-sink shortest-path problem

As we will see, our algorithm needs to maintain the shortest path from x_{init} to any node in a graph. Moreover, this graph undergoes a series of edge insertions and edge deletions. This problem is referred to as the fully dynamic *single-source shortest-path problem* or SSSP for short. Efficient algorithms [31], [32] exist that can store the minimal cost to reach each node (and the corresponding path) in such settings from a source node. In our setting, this source node is x_{init}. We make use of the following procedures which are provided by SSSP algorithms: $\texttt{delete\_edge}_{\text{SSSP}}(\mathcal{G}, (x_1, x_2))$ and $\texttt{insert\_edge}_{\text{SSSP}}(\mathcal{G}, (x_1, x_2))$ which delete and insert, respectively, the edge (x_1, x_2) from/into the graph $\mathcal{G}$ while maintaining $\text{cost}_{\mathcal{G}}$ for each node. We assume that these procedures return the set of nodes whose cost has changed due to the edge deletion or edge insertion. Furthermore, let

$\texttt{parent}_{\text{SSSP}}(\mathcal{G}, x)$ be a procedure returning the parent of x in the shortest path from the source to x in $\mathcal{G}$.

B. LBT-RRT

We propose a modification to the RRG algorithm by maintaining two roadmaps $\mathcal{G}_{lb}, \mathcal{T}_{apx}$ simultaneously. Both roadmaps have the same set of vertices but differ in their edge set. $\mathcal{G}_{lb}$ is a graph and $\mathcal{T}_{apx}$ is a tree rooted at x_{init}[2].

Let $\mathcal{G}_{RRG}$ be the roadmap constructed by RRG if run on the same sequence of samples used for LBT-RRT. The following invariants are maintained by the LBT-RRT algorithm:

> **Bounded approximation invariant** - For every node $x \in \mathcal{T}_{apx}, \mathcal{G}_{lb}$, $\text{cost}_{\mathcal{T}_{apx}}(x) \leq (1 + \varepsilon) \cdot \text{cost}_{\mathcal{G}_{lb}}(x)$.

and

> **Lower bound invariant** - For every node $x \in \mathcal{G}_{lb}$, $\text{cost}_{\mathcal{G}_{lb}}(x) \leq \text{cost}_{\mathcal{G}_{RRG}}(x)$.

The lower bound invariant is maintained by ensuring that the edges of $\mathcal{G}_{RRG}$ are a subset of the edges of $\mathcal{G}_{lb}$. As we will see, $\mathcal{G}_{lb}$ may possibly contain some edges that $\mathcal{G}_{RRG}$ considered but found to be in collision.

The main body of the algorithm (see Alg. 5) follows the structure of the RRT, RRT* and RRG algorithms with respect to adding a new milestone (lines 3-7) but differs in the connections added. If a path between the new node x_{new} and its nearest neighbor x_{nearest} is indeed collision-free, it is added to both roadmaps together with an edge from x_{nearest} to x_{new} (lines 8-11).

Similar to RRG and RRT*, LBT-RRT locates the set X_{near} of $k_{RRG} \log(|V|)$ nearest neighbors of x_{new} (line 12). Then, for each edge connecting a node from X_{near} to x_{new} and for each edge connecting x_{new} to a node from X_{near}, it uses a procedure $\texttt{consider\_edge}$ (Alg. 6) to assess if the edge should be inserted to either roadmaps. The edge is first lazily inserted into $\mathcal{G}_{lb}$ without checking if it is collision-free. This *may* cause the bounded approximation invariant to be violated, which in turn will induce a call to the local planner for a set of edges. Each such edge might either be inserted into $\mathcal{T}_{apx}$ or removed from $\mathcal{G}_{lb}$.

This is done as follows, first, the edge considered is inserted to $\mathcal{G}_{lb}$ while updating the shortest path to reach each vertex in $\mathcal{G}_{lb}$ (Alg. 6, line 1). Denote by I the set of updated vertices after the edge insertion. Namely, for every $x \in I$, $\text{cost}_{\mathcal{G}_{lb}}(x)$ has decreased due to the edge insertion. This cost decrease may, in turn, cause the bounded approximation invariant to be violated for some nodes in U. All such nodes are collected and inserted into a priority queue Q (line 2) ordered according to $\text{cost}_{\mathcal{G}_{lb}}$ from low to high. Now, the algorithm proceeds

---

[2]The subscript of $\mathcal{G}_{lb}$ is an abbreviation for lower bound and the subscript of $\mathcal{T}_{apx}$ is an abbreviation for approximation.

Algorithm 5 LBT-RRT ($x_{\text{init}}, \varepsilon$)
1: $\mathcal{T}_{lb}.G \leftarrow \{x_{\text{init}}\} \quad \mathcal{T}_{apx}.V \leftarrow \{x_{\text{init}}\}$
2: **while** `construct_roadmap()` **do**
3:     $x_{\text{rand}} \leftarrow$ `sample_free`$()$
4:     $x_{\text{nearest}} \leftarrow$ `nearest_neighbor`$(x_{rand}, \mathcal{T}_{lb}.V)$
5:     $x_{\text{new}} \leftarrow$ `steer`$(x_{\text{nearest}}, x_{\text{rand}})$
6:     **if** (!`collision_free`$(x_{\text{nearest}}, x_{\text{new}})$) **then**
7:       CONTINUE
8:     $\mathcal{T}_{apx}.V \leftarrow \mathcal{T}_{apx}.V \cup \{x_{\text{new}}\}$
9:     $\mathcal{T}_{apx}.\text{parent}(x_{\text{new}}) \leftarrow x_{\text{nearest}}$
10:    $\mathcal{G}_{lb}.V \leftarrow \mathcal{G}_{lb}.V \cup \{x_{\text{new}}\}$
11:    `insert_edge`$_{\text{SSSP}}(\mathcal{G}_{lb}, (x_{\text{nearest}}, x_{\text{new}}))$
12:    $X_{\text{near}} \leftarrow$ `nearest_neighbors`$(x_{\text{new}},$
$\mathcal{G}_{lb}.V, k_{RRG} \log(|\mathcal{G}_{lb}.V|))$
13:    **for all** $(x_{\text{near}}, X_{\text{near}})$ **do**
14:       `consider_edge`$(x_{\text{near}}, x_{\text{new}})$
15:    **for all** $(x_{\text{near}}, X_{\text{near}})$ **do**
16:       `consider_edge`$(x_{\text{new}}, x_{\text{near}})$

Algorithm 6 `consider_edge`(x_1, x_2)
1: $I \leftarrow$ `insert_edge`$_{\text{SSSP}}(\mathcal{G}_{lb}, (x_1, x_2))$
2: $Q \leftarrow \{x \in I \mid \text{cost}_{\mathcal{T}_{apx}}(x) > (1+\varepsilon) \cdot \text{cost}_{\mathcal{G}_{lb}}(x)\}$
3: **while** $Q \neq \emptyset$ **do**
4:     $x \leftarrow Q.\text{top}();$
5:     **if** $\text{cost}_{\mathcal{T}_{apx}}(x) > (1+\varepsilon) \cdot \text{cost}_{\mathcal{G}_{lb}}(x)$ **then**
6:       $x_{parent} \leftarrow$ `parent`$_{\text{SSSP}}(\mathcal{G}_{lb}, x)$
7:       **if** (`collision_free` (x_{parent}, x)) **then**
8:         $\mathcal{T}_{apx}.\text{parent}(x) \leftarrow x_{\text{parent}}$
9:         $Q.\text{pop}()$
10:      **else**
11:        $D \leftarrow$ `delete_edge`$_{\text{SSSP}}(\mathcal{G}_{lb}, (x_{\text{parent}}, x))$
12:        **for all** $y \in D \cap Q$ **do**
13:          $Q.\text{update\_cost}(y)$
14:     **else**
15:       $Q.\text{pop}()$

in iterations until the queue is empty (lines 3-15). At each iteration, the head of the queue x is considered (line 4). If the bounded approximation invariant does not hold (line 5), the algorithm checks if the edge in $\mathcal{G}_{lb}$ connecting the node x to its parent along the shortest path to x_{init} is collision free (lines 6-7). If this is the case, the approximation tree is updated (line 8) and the head of the queue is removed (line 9). If not, the edge is removed from $\mathcal{G}_{lb}$ (line 11). This causes an increase in $\text{cost}_{\mathcal{G}_{lb}}$ for a set D of nodes, some of which are already in the priority queue. Clearly, the bounded approximation invariant holds for the nodes $x \in D$ that are not in the priority queue. Thus, we take only the nodes $x \in D$ that are already in Q and update their location in Q according to their new cost (lines 12-13) . Finally, if the bounded approximation invariant holds for x then it is removed from the queue (lines 15).

C. Analysis

In this section we show that Alg. 5 maintains the lower bound invariant (Corollary III.5) and that after every iteration of the algorithm the bounded approximation invariant is maintained (Lemma III.8). We then report on the time complexity of the algorithm (Corollary III.10).

We note the following straightforward, yet helpful observations comparing LBT-RRT and RRG when run on the same sequence of random samples:

Observation III.1. *A node x is added to $\mathcal{G}_{lb}$ and to $\mathcal{T}_{apx}$ if and only if x is added to $\mathcal{G}_{RRG}$ (Alg. 2 lines 3-8 and 5, lines 3-11).*

Observation III.2. *Both LBT-RRT and RRG consider the same set of $k_{RRG} \log(|V|)$ nearest neighbors of x_{new} (Alg. 2,*

line 10 and Alg. 5, line 12).

Observation III.3. *Every edge added to the RRG roadmap (Alg. 2 line 13) is added to $\mathcal{G}_{lb}$ (Alg. 5 lines 14, 16 and Alg. 6 line 1).*

Note that some additional edges may be added to $\mathcal{G}_{lb}$ which are not added to the RRG roadmap as they are not collision-free.

Observation III.4. *Every edge of $\mathcal{T}_{apx}$ is collision free (Alg. 5, line 9 and Alg. 6, line 8).*

Thus, the following corollary trivially holds:

Corollary III.5. *After every iteration of LBT-RRT (Alg. 5, lines 3-16) the lower bound invariant is maintained.*

We continue with the following observations relevant to the analysis of the procedure `consider_edge`(x_1, x_2):

Observation III.6. *The only place where $\text{cost}_{\mathcal{G}_{lb}}$ is decreased is during a call to `insert_edge`$_{\text{SSSP}}(\mathcal{G}_{lb}, (x_1, x_2))$ (Alg. 6, line 1).*

Observation III.7. *A node x is removed from the queue Q (Alg 6, lines 9,15) only if the bounded approximation invariant holds for x.*

Showing that the bounded approximation invariant is maintained is done by induction on the number of calls to `consider_edge`(x_1, x_2). Using Obs. III.6, prior to the first call to `consider_edge`(x_1, x_2) the bounded approximation invariant is maintained. Thus, we need to show that:

Lemma III.8. *If the bounded approximation invariant holds prior to a call to the procedure `consider_edge`(x_1, x_2) (Alg. 6), then the procedure will terminate with the invariant maintained.*

    Proof: Assume that the bounded approximation invariant was maintained prior to a call to `consider_edge`(x_1, x_2). By Observation III.6 inserting a new edge (line 1) may cause

the bounded approximation invariant to be violated for a set of nodes. Moreover, it is the *only* place where such an event can occur. Observation III.7 implies that the bounded approximation invariant holds for every vertex *not* in Q.

Recall that in the priority queue we order the nodes according to $\text{cost}_{\mathcal{G}_{lb}}$ (from low to high) and at each iteration of $\texttt{consider\_edge}(x_1, x_2)$ the top of the priority queue x is considered. The parent x_{parent} of x, that has a smaller cost value, cannot be in the priority queue. Thus, the bounded approximation invariant holds for x_{parent}. Namely,

$$\text{cost}_{\mathcal{T}_{apx}}(x_{\text{parent}}) \leq (1 + \varepsilon) \cdot \text{cost}_{\mathcal{G}_{lb}}(x_{\text{parent}}).$$

Now, if the edge between x_{parent} and x is found to be free (line 7), we update the approximation tree (line 8). It follows that after such an event,

$$
\begin{aligned}
\text{cost}_{\mathcal{T}_{apx}}(x) &= \text{cost}_{\mathcal{T}_{apx}}(x_{parent}) + \\
&\quad \text{cost}(x_{\text{parent}}, x) \\
&\leq (1 + \varepsilon) \cdot \text{cost}_{\mathcal{G}_{lb}}(x_{\text{parent}}) + \\
&\quad \text{cost}(x_{\text{parent}}, x) \\
&\leq (1 + \varepsilon) \cdot \text{cost}_{\mathcal{G}_{lb}}(x).
\end{aligned}
$$

Namely, after updating the approximation tree, the bounded approximation invariant holds for the node x.

To summarize, at each iteration of Alg. 6 (lines 3-16), either: (i) we remove a node x from Q (line 9 or line 15) or (ii) we remove an incoming edge to the node x from the lower bound graph (line 11). If the node x was removed from Q (case (i)), the bounded approximation invariant holds—either it was not violated to begin with (line 15) or it holds after updating the approximation tree (lines 8-9).

To finish the proof we need to show that the main loop (lines 3-15) in Alg. 6 indeed terminates. Recall that the degree of each node is $O(\log n)$. Thus, a node x cannot be at the head of the queue more than $O(\log n)$ times (after each time we either remove an incoming edge or remove x from the queue). This in turn implies that after at most $O(n \log n)$ iterations Q is empty and the main loop terminates. ∎

From Corollary III.5, Lemma III.8 and using the asymptotic optimality of RRG we conclude,

Theorem III.9. *LBT-RRT is asymptotically near-optimal with an approximation factor of* $(1 + \varepsilon)$.

Namely, the cost of the path computed by LBT-RRT converges to a cost at most $(1 + \varepsilon)$ times the cost of the optimal path almost surely.

We continue now to discuss the time complexity of the algorithm. If δ is the number of nodes updated during a call to an SSSP procedure[3] (namely, $\texttt{insert\_edge}_{\text{SSSP}}$ or $\texttt{delete\_edge}_{\text{SSSP}}$), then the complexity of the procedure is $O(\delta \log n)$ when using the algorithm of Ramalingam et

---

[3]The number of nodes δ updated during an SSSP procedure depends on the topology of the graph and the edge weights. Theoretically, in the worst case $\delta = O(n)$ and a dynamic SSSP algorithm cannot perform better than recomputing shortest paths from scratch. However, in practice this value is much smaller.

al. [32]. Set $\hat{\delta}$ to be the maximum value of δ over all calls to SSSP procedures (Alg 5 line 11 and Alg 6, lines 1 and 11) and let n denote the final number of samples used by LBT-RRT.

We have $O(n \log n)$ edges and each edge will be inserted to $\mathcal{G}_{lb}$ once (Alg 5 line 11 or Alg 6 line 1) and possibly be removed from $\mathcal{G}_{lb}$ once (Alg 6 line 11). Therefor, the total complexity due to the SSSP procedures is $O(\hat{\delta} \cdot n \log^2 n)$. The time-complexity of all the other operations (nearest neighbours, collision detection etc.) is similar to RRG which runs in time $O(n \log n)$.

Corollary III.10. *LBT-RRT runs in time* $O(\hat{\delta} \cdot n \log^2 n)$, *where n is the number of samples and $\hat{\delta}$ is the maximal number of nodes updated over all SSSP procedures* .

While this running time may seem discouraging, we note that in practice, the local planning dominates the actual running time of the algorithm in practice. As we demonstrate in Section IV through various simulations, LBT-RRT produces high-quality results in an efficient manner.

D. Implementation details

We describe the following optimizations that we use in order to speed up the running-time of the algorithm. The first is that the set X_{near} is ordered according to the cost to reach x_{new} from x_{init} through an element x of X_{near}. Hence, the set X_{near} will be traversed from the node that yields the smallest lower bound to reach x_{new} to the node that will yield the highest lower bound. After the first edge that does *not* violate the bounded approximation invariant, no subsequent node can improve the cost to reach x_{new} and $\texttt{insert\_edge}_{\text{SSSP}}$ will not need to perform any updates. This ordering was previously used to speed up RRT* (see, e.g., [19], [33]).

The second optimization comes to avoid the situation where $\texttt{insert\_edge}_{\text{SSSP}}$ is called and immediately afterwards the same edge is removed. Hence, given an edge, we first check if the bounded approximation invariant will be violated had the edge been inserted. If this is indeed the case, the local planner is invoked and only if the edge is collision free $\texttt{insert\_edge}_{\text{SSSP}}$ is called.

E. Discussion

Let T^{ω}_{ALG} denote the time needed for an algorithm ALG to find a feasible solution on a sequence ω of random samples. Clearly, $\text{T}^{\omega}_{\text{RRT}} \leq \text{T}^{\omega}_{\text{RRG}}$ (as RRG may require more calls to the collision detector than the RRT algorithm). Moreover, for every $\varepsilon_1 \leq \varepsilon_2$ it holds that

$$\text{T}^{\omega}_{\text{RRT}} \leq \text{T}^{\omega}_{\text{LBT-RRT}(\varepsilon_2)} \leq \text{T}^{\omega}_{\text{LBT-RRT}(\varepsilon_1)} \leq \text{T}^{\omega}_{\text{RRG}}.$$

Thus, given a limited amount of time, RRG may fail to construct any solution. On the other hand, RRT may find a solution fast but will not improve its quality (if the goal is a single configuration). LBT-RRT allows to find a feasible path quickly while continuing to search for a path of higher quality.

Remark The conference version of this paper contained an oversight with regard to how the bounded approximation invariant was maintained. Specifically, instead of storing $\mathcal{G}_{lb}$

as a graph, a tree was stored which was rewired locally. When the algorithm tested if the bounded approximation invariant was violated for a node x, it only considered the *children* of x in the tree. This local test did not take into account the fact that changing the cost of x in the tree could also change the cost of nodes y that are descendants of x (but not its children). The implications of the oversight is that the algorithm was not asymptotically near optimal. The experimental results presented in the conference version of this paper suggest that in certain scenarios this oversight did not have a significant effect on the convergence to high quality solutions. Having said that, LBT-RRT as presented in this paper is both asymptotically near optimal and converges to high quality solutions faster than the original algorithm.

IV. EVALUATION

We present an experimental evaluation of the performance of LBT-RRT as an anytime algorithm on different scenarios consisting of 3,6 and 12 DoFs (Fig. 1). The algorithm was implemented using the Open Motion Planning Library (OMPL 0.10.2) [34] and our implementation is currently distributed with the OMPL release. All experiments were run on a 2.8GHz Intel Core i7 processor with 8GB of memory. RRT* was implemented by using the ordering optimization described in Section III and [19]).

The Maze scenario (Fig. 1a) consists of a planar polygonal robot that can translate and rotate. The Alternating barriers scenario (Fig. 1b) consists of a robot with three perpendicular rods free-flying in space. The robot needs to pass through a series of barriers each containing a large and a small hole. For an
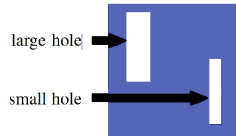
Fig. 2. One barrier of the Alternating barriers scenario.

illustration of one such barrier, see Fig. 2. The large holes are located at alternating sides of consecutive barriers. Thus, an easy path to find would be to cross each barrier through a large hole. A high-quality path would require passing through a small hole after each large hole. Finally, the cubicles scenario consists of two L-shaped robots free-flying in space that need to exchange locations amidst a sparse collection of obstacles[4].

We compare the performance of LBT-RRT with RRT, RRG and RRT* when a fixed time budget is given. We add another algorithm which we call RRT+RRT* which initially runs RRT and once a solution is found runs RRT*. RRT+RRT* will find a solution as fast as RRT and is asymptotically-optimal. For LBT-RRT we consider $(1+\varepsilon)$ values of $1.2, 1.4, 1.8$ and report on the success rate of each algorithm (Fig. 3). Additionally, we report on the path length after applying shortcuts (Fig. 5). Each result is averaged over 100 different runs.

Fig. 3 depicts similar behaviour for all scenarios: As one would expect, the success rate for all algorithms has a monotonically increasing trend as the time budget increases. For a

[4]The Maze Scenario and the Cubicles Scenario are provided as part of the OMPL distribution.

Algorithm 7 `consider_edge_goal_biased`(x_1, x_2)

1: $c_{\min}^{\text{apx}} \leftarrow \text{cost}_{\text{LPA*}}(\mathcal{T}_{apx})$
2: `insert_edge`$_{\text{LPA*}}(\mathcal{G}_{lb}, x_1, x_2)$
3: $x \leftarrow \text{shortest\_path}_{\text{LPA*}}(\mathcal{G}_{lb})$
4: $x_{parent} \leftarrow \text{parent}_{\text{LPA*}}(\mathcal{G}_{lb}, x)$
5: $c_{\min}^{\text{lb}} \leftarrow \text{cost}_{\text{LPA*}}(\mathcal{G}_{lb})$
6: **while** $c_{\min}^{\text{apx}} > (1 + \varepsilon) \cdot c_{\min}^{\text{lb}}$ **do**
7:    **if** (`collision_free`(x_{parent}, x)) **then**
8:       `insert_edge`$_{\text{LPA*}}(\mathcal{T}_{apx}, x_{parent}, x)$
9:       `shortest_path`$_{\text{LPA*}}(\mathcal{T}_{apx})$
10:      $c_{\min}^{\text{apx}} \leftarrow \text{cost}_{\text{LPA*}}(\mathcal{T}_{apx})$
11:      $x \leftarrow \text{parent}_{\text{LPA*}}(x)$
12:    **else**
13:       `delete_edge`$_{\text{LPA*}}(\mathcal{G}_{lb}, x_{parent}, x)$
14:      **GoTo** line 3

---

specific time budget, the success rate for RRT and RRT+RRT* is typically highest while that of the RRT* and RRG is lowest. The success rate for LBT-RRT for a specific time budget, typically increases as the value of ε increases. Fig. 5 also depicts similar behavior for all scenarios: the average path length decreases for all algorithms (except for RRT). The average path length for LBT-RRT typically decreases as the value of ε decreases and is comparable to that of RRT* for low values of ε. RRT+RRT* behaves similarly to RRT* but with a "shift" along the time-axis which is due to the initial run of RRT. We note that although RRG and RRT+RRT* are asymptotically-optimal, their overhead makes them poor algorithms when one desires a *high-quality* solution very fast.

Thus, Fig. 3 and 5 should be looked at simultaneously as they encompass the tradeoff between speed to find *any* solution and the quality of the solution found. Let us demonstrate this on the alternating barriers scenario: If we look at the success rate of each algorithm to find *any* solution (Fig. 3b), one can see that RRT manages to achieve a success rate of 70% after 30 seconds. RRT*, on the other hand, requires 70 seconds to achieve the same success rate (more than double the time). For all different values of ε, LBT-RRT manages to achieve a success rate of 70% after 50 seconds (around 60% overhead when compared to RRT). Now, considering the path length at 50 seconds, typically the paths extracted from LBT-RRT yield the same quality when compared to RRT* while ensuring a high success rate.

The same behavior of finding paths of high-quality (similar to the quality that RRT* produces) within the time-frames that RRT requires in order to find *any* solution has been observed for both the Maze scenario and the Cubicles scenario. Results omitted in this text. For supplementary material the reader is referred to http://acg.cs.tau.ac.il/projects/LBT-RRT.

V. LAZY, GOAL-BIASED LBT-RRT

In this section we show to further reduce the number of calls to the local planner by incorporating a lazy approach together with a goal bias.
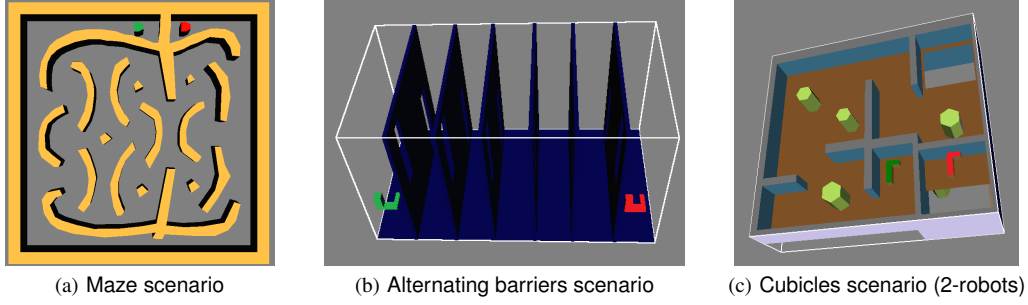
(a) Maze scenario   (b) Alternating barriers scenario   (c) Cubicles scenario (2-robots)

Fig. 1. Benchmark scenarios. The start and goal configuration are depicted in green and red, respectively.

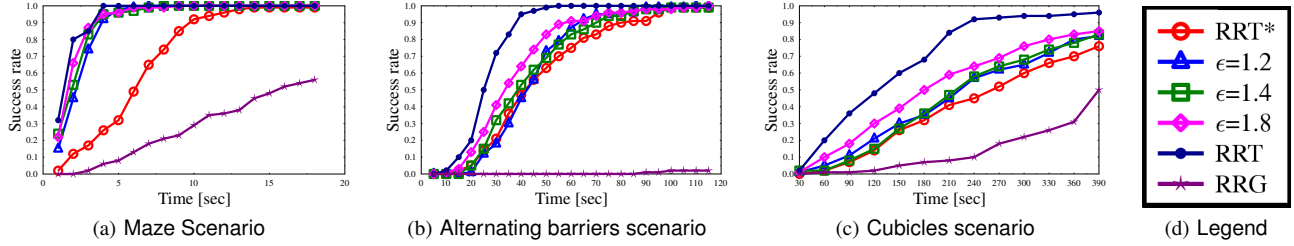(a) Maze Scenario   (b) Alternating barriers scenario   (c) Cubicles scenario   (d) Legend

Fig. 3. Success rate for algorithms on different scenarios (RRT and RRT+RRT* have almost identical success rates; the plot for RRT+RRT* is omitted to avoid cluttering of the graph).

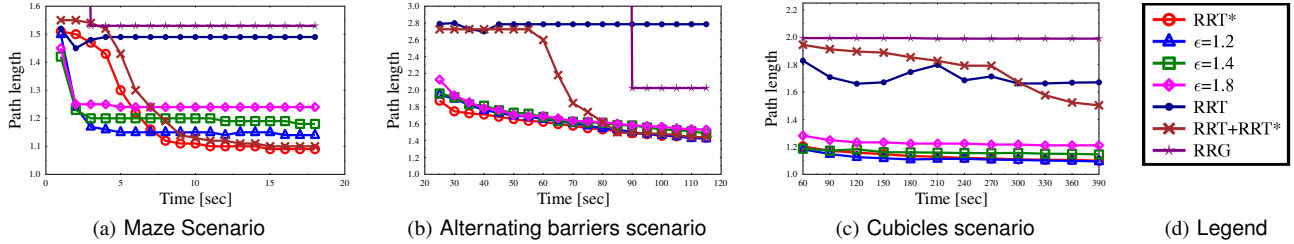(a) Maze Scenario   (b) Alternating barriers scenario   (c) Cubicles scenario   (d) Legend

Fig. 4. Path lengths for algorithms on different scenarios. Length values are normalized such that a length of one represents the length of an optimal path.

LBT-RRT maintains the lower bound invariant to *every* node. This is desirable in settings where a high-quality path to every point in the configuration space is required. However, when only a high-quality path to the goal is needed, this may lead to unnecessary time-consuming calls to the local planner.

Therefore, we suggest the following variant of LBT-RRT where we relax the bounded approximation invariant such that it holds only for nodes $x \in X_{goal}$. This variant is similar to LBT-RRT but differs with respect to the calls to the local planner and with respect to the dynamic shortest-path algorithm used. As we only maintain the bounded approximation invariant to the goal nodes, we do not need to continuously update the (approximate) shortest path to every node in $\mathcal{G}_{lb}$. We replace the SSSP algorithm, which allows to compute the shortest paths to *every* node in a dynamic graph, with Lifelong Planning A* (LPA*) [35]. LPA* allows to repeatedly find shortest paths from a given start to a given goal while allowing for edge insertions and deletions. Similiar to A* [36], this is done by using heuristic function h such that for every node x, $h(x)$ is an estimate of the cost to reach the goal from x.

Given a start vertex x_{init}, a goal region X_{Goal}, we will use the following functions which are provided when implementing LPA*: $\texttt{shortest\_path}_{\text{LPA}\star}(G)$, recomputes the shortest path to reach X_{Goal} from x_{init} on the graph G and returns the node $x \in X_{Goal}$ such that $x \in X_{Goal}$ and $\texttt{cost}_{\mathcal{G}_{lb}}(x)$ is minimal among all $x' \in X_{Goal}$. Once the function has been called, the following functions take constant running time: $\texttt{cost}_{\text{LPA}\star}(G)$ returns the minimal cost to reach X_{Goal} from x_{init} on the graph G and for every node x lying on a shortest path to the goal, $\texttt{parent}_{\text{LPA}\star}(G, x)$ returns the predecessor of the node x along this path. Additionally, $\texttt{insert\_edge}_{\text{LPA}\star}(G, x, y)$ and $\texttt{delete\_edge}_{\text{LPA}\star}(G, x, y)$ inserts (deletes) the edge (x, y) to (from) the graph G, respectively.

We are now ready to describe Lazy, goal-biased LBT-RRT which is similar to LBT-RRT except for the way new edges are considered. Instead of the function $\texttt{consider\_edge}$ called in lines 14 and 16 of Alg. 5, the function $\texttt{consider\_edge\_goal\_biased}$ is called.

$\texttt{consider\_edge\_goal\_biased}(x_1, x_2)$, outlined in Alg. 7, begins by computing the cost to reach the goal in $\mathcal{T}_{apx}$ (line 1) and in $\mathcal{G}_{lb}$ after adding the edge (x_1, x_2) lazily

to $\mathcal{G}_{lb}$ (lines 2-5). Namely, the edge is added with no call to the local planner and without checking if the bounded approximation invariant is violated. Note that the *relaxed* bounded approximation invariant is violated (line 6) only if a path to the goal is found. Clearly, if all edges along the shortest path to the goal are found to be collision free, then the invariant holds. Thus, the algorithm attempts to follow the edges along the path (starting at the last edge and backtracking towards x_{init}) one by one and test if they are indeed collision-free. If an edge is collision free (line 7), it is inserted to $\mathcal{T}_{apx}$ (line 8), and a path to the goal in $\mathcal{T}_{apx}$ is recomputed (line 9). This is repeated as long as the relaxed bounded approximation invariant is violated. If the edge is found to be in collision (line 12), it is removed from $\mathcal{G}_{lb}$ (line 13) and the process is repeated (line 14).

Following similar arguments as described in Section III, one can show the correctness of the algorithm. We note that as long as no path has been found, the algorithm performs no more calls to the local planner than RRT. Additionally, it is worth noting that the planner bares resemblance with Lazy-RRG* [37].

We compared Lazy, goal-biased LBT-RRT with LBT-RRT, RRT* and RRG on the Home scenario (Fig. 5a). In this scenario, a low quality solution is typically easy to find and all algorithms (except RRG) find a solution with roughly the same success rate as RRT (results omitted). Converging to the optimal solution requires longer running times as low-quality paths are easy to find yet high-quality ones pass through narrow passages. Fig. 5b depicts the path length obtained by the algorithms as a function of time. The convergence to the optimal solution of RRG is significantly slower than all other algorithms. Both LBT-RRT and RRT* find a low quality solution (between five and six times longer than the optimal solution) within the allowed time frame and manage to slightly improve upon its cost (with RRT* obtaining slightly shorter solutions than LBT-RRT). When enhancing LBT-RRT with a lazy approach together with goal-biasing, one can observe that the convergence rate improves substantially.

VI. CONCLUSION AND FUTURE WORK

In this work we presented an asymptotically near-optimal motion planning algorithm. Using an approximation factor allows the algorithm to avoid calling the computationally-expensive local planner when no substantially better solution may be obtained. LBT-RRT, together with the lazy, goal-biased variant, make use of *dynamic shortest path algorithms*. This is an active research topic in many communities such as artificial intelligence and communication networks.

Hence, the algorithms we proposed in this work may benefit from any advances made for dynamic shortest path algorithms. For example, recently D'Andrea et al. [38] presented an algorithm that allows for dynamically maintaining shortest path trees under *batches* of updates which can be used by LBT-RRT instead of the SSSP algorithm.

Looking to further extend our framework, we seek natural stopping criteria for LBT-RRT. Such criteria could possibly be

related to the rate at which the quality is increased as additional samples are introduced. Once such a criterion is established, one can think of the following framework: Run LBT-RRT with a large approximation factor (large ε) , once the stopping criterion has been met, decrease the approximation factor and continue running. This may allow an even quicker convergence to find any feasible path while allowing for refinement as time permits (similar to [27]). While changing the approximation factor in LBT-RRT may possibly require a massive rewiring of $\mathcal{G}_{lb}$ (to maintain the bounded approximation invariant) this is not the case in Lazy, goal-biased LBT-RRT. In this variant of LBT-RRT the approximation factor can change at any stage of the algorithm without any modifications at all.

An interesting question to be further studied is can our framework be applied to different quality measures. For certain measures, such as bottleneck clearance of a path, this is unlikely, as bounding the quality of an edge already identifies if it is collision-free. However, for some other measures such as energy consumption, we believe that the framework could be effectively used.

## VII. ACKNOWLEDGEMENTS

APPENDIX

RRT has been used in numerous applications and various efficient implementations and heuristics have been suggested for it. Even the relatively recent RRT* has already gained many applications and various implementations. Typically, the applications rely on the efficiency of RRT or the asymptotic optimality of RRT*. We list two such applications (Sections A and B below) and discuss the possible advantage of replacing the existing planner (either RRT or RRT*) with LBT-RRT.

Efficient implementations and heuristics typically take into account the primitive operations used by the RRT and the RRT* algorithms (such as collision detection, nearest neighbor computation, sampling procedure etc.). Thus, techniques suggested for efficient implementations of RRT and RRT* may be applied to LBT-RRT with little effort as the latter relies on the same primitive operations. We give two examples in Sections C and D below.

Finally, we show how to apply our approach to a different asymptotically-optimal sampling-based algorithm—Fast Marching Trees (FMT*) [22].

A. Re-planning using RRTs:

Many real-world applications involve a C-space that undergoes changes (such as moving obstacles or partial initial information of the workspace). A common approach to plan in such dynamic environments is to run RRT, and re-plan when a change in the environment occurs. Re-planning may be done from scratch although this can be unnecessary and time consuming as the assumption is that only part of the
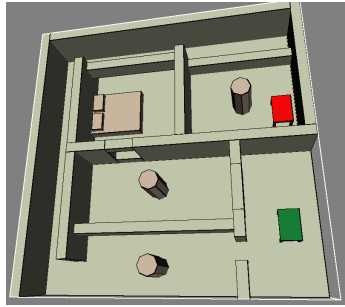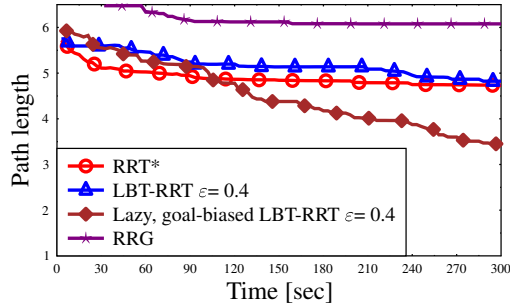
Fig. 5. Simulation results comparing Lazy, goal-biased LBT-RRT with LBT-RRT, RRT* and RRG. (a) Home scenario (provided by the OMPL distribution). Start and target table-shaped robots are depicted in green and red, respectively. (b) Path lengths as a function of computation time. Length values are normalized such that a length of one represents the length of an optimal path

environment changes. Ferguson et al. [26] suggest to (i) plan an initial path using RRT, (ii) when a change in the configuration space is detected, nodes in the existing tree may be invalidated and a "trimming" procedure is applied where invalid parts are removed and (iii) the trimmed tree is grown until a new solution is generated.

Obviously LBT-RRT can replace RRT in the suggested scheme. If the overhead of running LBT-RRT when compared to RRT is acceptable (which may indeed be the case as the experimental results in Section IV suggest), then the algorithm will be able to produce *high-quality* paths in dynamic environments.

B. High-quality planning on implicitly-defined manifolds:

Certain motion-planning problems, such as grasping with a multi-fingered hand, involve planning on implicitly-defined manifolds. Jaillet and Porta [39] address the central challenges of applying RRT* to such cases. The challenges include sampling uniformly on the manifold, locating the nearest neighbors using the metric induced by the manifold, computing the shortest path between two points and more. They suggest AtlasRRT*, an adaptation of RRT* that operates on manifolds. It follows the same structure as RRT* but maintains an atlas by iteratively adding charts to the atlas to facilitate the primitive operations of RRT* on the manifold (i.e., sampling, nearest-neighbor queries, local planning etc.).

If one is concerned with *fast convergence* to a high quality solution, LBT-RRT can be used seamlessly, replacing the guarantee for optimality with a weaker near-optimality guarantee.

C. Sampling Heuristics:

Following the exposition of RRT*, Akgun and Stilman [40] suggested a sampling bias for the RRT* algorithm. This bias accelerates cost decrease of the path to the goal in the RRT* tree. Additionally, they suggest a simple node-rejection criterion to increase efficiency. These heuristics may be applied to the LBT-RRT by simply changing the procedure `sample_free` (Alg. 5, line 3).

D. Parallel RRTs:

In recent years, hardware allowing for parallel implementation of existing algorithms has become widespread both in the form of multi-core Central Processing Units (CPUs) and in the form of Graphics Processing Units (GPUs). Parallel implementations for sampling based algorithms have already been proposed in the late 1990s [41]. Since then, a multitude of such implementations emerged (see, e.g., [42], [43] for a detailed literature review).

We review two approaches to parallel implementation of RRT and RRT* and claim that both approaches may be used for parallel implementation of LBT-RRT. The first approach, by Ichnowski et al. [42] suggests parallel variants of RRT and RRT* on multi-core CPUs that achieve superlinear speedup. By using CPU-based implementation, their approach retains the ability to integrate the planners with existing CPU-based libraries and algorithms. They achieve superlinear speedup by: (i) lock-free parallelism using atomic operations to reduce slowdowns caused by contention, (ii) partition-based sampling to reduce the size of each processor core's working data set and to improve cache efficiency and (iii) parallel backtracking in the rewiring phase of RRT*. LBT-RRT may benefit from all three key algorithmic features.

Bialkowski et al. [43] present a second approach for parallel implementation of RRT and RRT*. They suggest a massively parallel, GPU-based implementation of the collision-checking procedures of RRT and RRT*. Again, this approach may be applied to the collision-checking procedure of LBT-RRT without any need for modification.

E. Framework Extensions

FMT*, proposed by Janson and Pavone, is a recently introduced asymptotically-optimal algorithm which is shown to converge to an optimal solution faster than PRM* or RRT*. It uses a set of probabilistically-drawn configurations to construct a tree, which grows in cost-to-come space. Unlike RRT*, it is a batch algorithm that works with a predefined number of nodes n.

We first describe the FMT* algorithm (outlined in Alg. 8), we continue to describe how to apply it in an anytime fashion and conclude by describing how to apply our framework to this anytime variant. FMT* samples n collision-free nodes[5] V

---

[5]By slight abuse of notation, `sample_free(n)` is a procedure returning n collision-free samples.

(line 1) and builds a minimum-cost spanning tree rooted at the initial configuration by maintaining two sets of nodes H, W such that H is the set of nodes added to the tree that may be expanded and W is the set of nodes not in the tree (line 2). It then computes for each node the set of nearest neighbors[6] of radius $r(n)$ (line 4). The algorithm repeats the following process: the node z with the lowest cost-to-come value is chosen from H (lines 5 and 17). For each neighbor x of z that is not already in H, the algorithm finds its neighbor $y \in H$ such that the cost-to-come of y added to the distance between y and x is minimal (lines 7-10). If the local path between y and x is free, x is added to H with y as its parent (lines 11-13). At the end of each iteration z is removed from H (line 14). The algorithm runs until a solution is found or there are no more nodes to process.

We now outline a straightforward enhancement to FMT* to make it anytime. As noted in previous work (see, e.g., [44]) one can turn a batch algorithm into an anytime one by the following (general) approach: choose an initial (small) number of samples $n = n_0$ and apply the algorithm. As long as time permits, double n and repeat the process. We call this version anytime FMT* or aFMT*, for short.

We can further speed up this method by reusing both existing samples and connections from previous iterations. Assume the algorithm was run with n samples and now we wish to re-run it with $2n$ samples. In order to obtain the $2n$ random samples, we take the n random samples from the previous iteration together with n new additional random samples. For each node that was used in iteration $i - 1$, on average half of its neighbors in iteration i are nodes from iteration $i - 1$ and half of its neighbors are newly-sampled nodes. Thus, if we cache the results of calls to the local planner, we can use them in future iterations using the framework presented in this paper. We call this algorithm LBT-aFMT*.

Alg. 9 outlines one iteration of LBT-aFMT*, differences between FMT* and LBT-aFMT* are colored in red. Similar to LBT-RRT, LBT-aFMT* constructs two trees $\mathcal{T}_{lb}$ and $\mathcal{T}_{ub}$ and maintains the bounded approximation invariant and the lower bound invariant. The two invariants are maintained by using a cache that can efficiently answer if the local path between two configurations is collision-free for a subset of the nodes used. The proof of near-optimality of LBT-aFMT*, which is omitted, follows the same lines as the analysis presented in Section III-C.

In order to demonstrate the effectiveness of applying our lower-bound framework to aFMT*, we compared the two algorithms, namely aFMT* and LBT-aFMT*, in the case of a three-dimensional configuration space which we call the Corridors scenario (Fig. 6a). It consists of a planar hexagon robot that can translate and rotate amidst a collection of small obstacles. There are two main corridors from the start to the goal position, a wide one and a narrow one. A relatively small number of samples suffices to find a path through

---

[6]The nearest-neighbor computation can be delayed and performed only when needed but we present the batched mode of computation to simplify the exposition.

Algorithm 8 `fast_marching_tree`
1: $V \leftarrow \{x_{\text{init}}\} \cup \texttt{sample\_free}(n); \quad E \leftarrow \emptyset; \quad \mathcal{T} \leftarrow (V, E)$

2: $W \leftarrow V \setminus \{x_{\text{init}}\}; \quad H \leftarrow \{x_{\text{init}}\}$
3: **for all** $v \in V$ **do**
4: $\quad N_v \leftarrow \texttt{nearest\_neighbors}(V \setminus \{v\}, v, r(n))$
5: $z \leftarrow x_{\text{init}}$
6: **while** $z \notin \mathcal{X}_{\text{Goal}}$ **do**
7: $\quad H_{\text{new}} \leftarrow \emptyset; \quad X_{\text{near}} \leftarrow W \cap N_z$
8: $\quad$ **for** $x \in X_{\text{near}}$ **do**
9: $\quad\quad Y_{\text{near}} \leftarrow H \cap N_x$
10: $\quad\quad y_{\text{min}} \leftarrow \arg\min_{y \in Y_{\text{near}}} \{\texttt{cost}_\mathcal{T}(y) + \texttt{cost}(y, x)\}$
11: $\quad\quad$ **if** $\texttt{collision\_free}(y_{\text{min}}, x)$ **then**
12: $\quad\quad\quad \mathcal{T}.\texttt{parent}(x) \leftarrow y_{\text{min}}$
13: $\quad\quad\quad H_{\text{new}} \leftarrow H_{\text{new}} \cup \{x\}; \quad W \leftarrow W \setminus \{x\}$
14: $\quad H \leftarrow (H \cup H_{\text{new}}) \setminus \{z\}$
15: $\quad$ **if** $H = \emptyset$ **then**
16: $\quad\quad$ **return** FAILURE
17: $\quad z \leftarrow \arg\min_{y \in H} \{\texttt{cost}_\mathcal{T}(y)\}$
18: **return** PATH

---

the wide corridor, yet in order to compute a low-cost path through the narrow corridor, many more samples are needed. This demonstrates how LBT-aFMT* refrains from refining an existing solution when no significant advantage can be obtained. Fig. 6b depicts the success rate of finding a path through the narrow corridor as time progresses. Clearly, for these types of settings LBT-aFMT* performs favorably over aFMT* even for large approximation factors. For example, to reach a 60% succes rate in finding a path through the narrow corridor, LBT-aFMT*, run with $\varepsilon = 0.5$, needs half of the time needed by aFMT*.

REFERENCES

[1] L. E. Kavraki, P. Švestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high dimensional configuration spaces," *IEEE Trans. Robot.*, vol. 12, no. 4, pp. 566–580, 1996.
[2] J. J. Kuffner and S. M. LaValle, "RRT-Connect: An efficient approach to single-query path planning," in *ICRA*, 2000, pp. 995–1001.
[3] H. Choset, K. M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun, *Principles of Robot Motion: Theory, Algorithms, and Implementation*. MIT Press, June 2005.
[4] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *I. J. Robotic Res.*, vol. 30, no. 7, pp. 846–894, 2011.
[5] O. Nechushtan, B. Raveh, and D. Halperin, "Sampling-diagram automata: A tool for analyzing path quality in tree planners," in *WAFR*, 2010, pp. 285–301.
[6] R. Geraerts and M. H. Overmars, "Creating high-quality paths for motion planning," *I. J. Robotic Res.*, vol. 26, no. 8, pp. 845–863, 2007.
[7] B. Raveh, A. Enosh, and D. Halperin, "A little more, a lot better: Improving path quality by a path-merging algorithm," *IEEE Trans. Robot.*, vol. 27, no. 2, pp. 365–371, 2011.
[8] N. M. Amato, O. B. Bayazit, L. K. Dale, C. Jones, and D. Vallejo, "OBPRM: an obstacle-based PRM for 3D workspaces," in *WAFR*, 1998, pp. 155–168.
[9] J.-M. Lien, S. L. Thomas, and N. M. Amato, "A general framework for sampling on the medial axis of the free space," in *ICRA*, 2003, pp. 4439–4444.
[10] C. Urmson and R. G. Simmons, "Approaches for heuristically biasing RRT growth," in *IROS*, 2003, pp. 1178–1183.
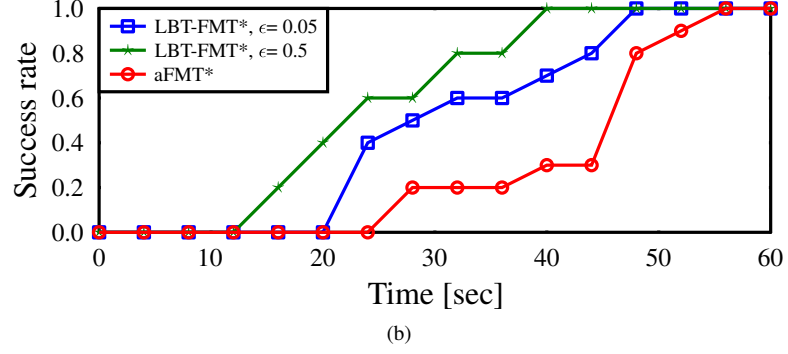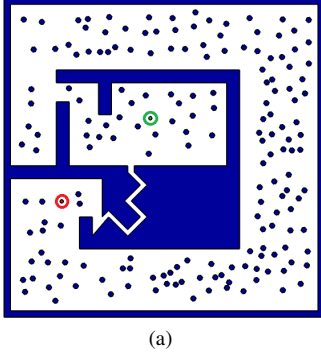
Fig. 6. (a) Corridors scenario—Start and target configurations are depicted by green and red circles, respectively. The narrow corridor appears in the lower half of the workspace. (b) Success rate to find a path through the narrow corridor.

Algorithm 9 `LBT-aFMT*` `(Cache)`

1: $V \leftarrow \{x_{\mathrm{init}}\} \cup \texttt{sample\_free}(n); \quad E \leftarrow \emptyset;$
2: $\mathcal{T}_{lb} \leftarrow (V, E); \quad \mathcal{T}_{apx} \leftarrow (V, E)$
3: $W \leftarrow V \setminus \{x_{\mathrm{init}}\}; \quad H \leftarrow \{x_{\mathrm{init}}\}$
4: **for all** $v \in V$ **do**
5: $\quad N_v \leftarrow \texttt{nearest\_neighbors}(V \setminus \{v\}, v, r(n))$
6: $z \leftarrow x_{init}$
7: **while** $z \notin \mathcal{X}_{Goal}$ **do**
8: $\quad H_{\mathrm{new}} \leftarrow \emptyset; \quad X_{\mathrm{near}} \leftarrow W \cap N_z$
9: $\quad$ **for** $x \in X_{near}$ **do**
10: $\quad\quad Y_{near} \leftarrow H \cap N_x$
11: $\quad\quad y_{lb} \leftarrow \arg\min_{y \in Y_{near}} \{\texttt{cost}_{\mathcal{T}_{lb}}(y) + \texttt{cost}(y, x)\}$
12: $\quad\quad Y_{apx} \leftarrow \{y \in Y_{near} \mid (y, x) \in \texttt{Cache}\}$
13: $\quad\quad y_{apx} \leftarrow \arg\min_{y \in Y_{apx}} \{\texttt{cost}_{\mathcal{T}_{apx}}(y) + \texttt{cost}(y, x)\}$
14: $\quad\quad c_{lb} \leftarrow \texttt{cost}_{\mathcal{T}_{lb}}(y_{lb}) + \texttt{cost}(y_{lb}, x)$
15: $\quad\quad c_{apx} \leftarrow \texttt{cost}_{\mathcal{T}_{apx}}(y_{apx}) + \texttt{cost}(y_{apx}, x)$
16: $\quad\quad$ **if** $c_{apx} \leq (1 + \varepsilon) \cdot c_{lb}$ **then**
17: $\quad\quad\quad \mathcal{T}_{lb}.\texttt{parent}(x) \leftarrow y_{lb}$
18: $\quad\quad\quad \mathcal{T}_{apx}.\texttt{parent}(x) \leftarrow y_{apx}$
19: $\quad\quad\quad H_{new} \leftarrow H_{new} \cup \{x\}; \quad W \leftarrow W \setminus \{x\}$
20: $\quad\quad$ **else**
21: $\quad\quad\quad$ **if** $\texttt{collision\_free}(y_{lb}, x)$ **then**
22: $\quad\quad\quad\quad \mathcal{T}_{lb}.\texttt{parent}(x) \leftarrow y_{lb}$
23: $\quad\quad\quad\quad \mathcal{T}_{apx}.\texttt{parent}(x) \leftarrow y_{lb}$
24: $\quad\quad\quad\quad H_{new} \leftarrow H_{new} \cup \{x\}; \quad W \leftarrow W \setminus \{x\}$
25: $\quad H \leftarrow (H \cup H_{new}) \setminus \{z\}$
26: $\quad$ **if** $H = \emptyset$ **then**
27: $\quad\quad$ **return** FAILURE
28: $\quad z \leftarrow \arg\min_{y \in H} \{\texttt{cost}(y)\}$
29: **return** PATH

[11] A. C. Shkolnik, M. Walter, and R. Tedrake, "Reachability-guided sampling for planning under differential constraints," in *ICRA*, 2009, pp. 2859–2865.

[12] T. Siméon, J.-P. Laumond, and C. Nissoux, "Visibility-based probabilistic roadmaps for motion planning," *Advanced Robotics*, vol. 14, no. 6, pp. 477–493, 2000.

[13] N. A. Wedge and M. S. Branicky, "On heavy-tailed runtimes and restarts in rapidly-exploring random trees," in *AAAI*, 2008, pp. 127–133.

[14] O. Salzman, D. Shaharabani, P. K. Agarwal, and D. Halperin, "Sparsification of motion-planning roadmaps by edge contraction," *I. J. Robotic Res.*, vol. 33, no. 14, pp. 1711–1725, 2014.

[15] J. D. Marble and K. E. Bekris, "Computing spanners of asymptotically optimal probabilistic roadmaps," in *IROS*, 2011, pp. 4292–4298.

[16] ——, "Asymptotically near-optimal is good enough for motion planning," in *ISRR*, 2011.

[17] ——, "Towards small asymptotically near-optimal roadmaps," in *ICRA*, 2012, pp. 2557–2562.

[18] A. Dobson and K. E. Bekris, "Sparse roadmap spanners for asymptotically near-optimal motion planning," *I. J. Robotic Res.*, vol. 33, no. 1, pp. 18–47, 2014.

[19] A. Perez, S. Karaman, A. Shkolnik, E. Frazzoli, S. Teller, and M. Walter, "Asymptotically-optimal path planning for manipulation using incremental sampling-based algorithms," in *IROS*, 2011, pp. 4307–4313.

[20] F. Islam, J. Nasir, U. Malik, Y. Ayaz, and O. Hasan, "RRT*-Smart: Rapid convergence implementation of RRT* towards optimal solution," *Int. J. Adv. Rob. Sys.*, vol. 10, pp. 1–12, 2013.

[21] O. Arslan and P. Tsiotras, "Use of relaxation methods in sampling-based algorithms for optimal motion planning," in *ICRA*, 2013, pp. 2413–2420.

[22] L. Janson and M. Pavone, "Fast marching trees: a fast marching sampling-based method for optimal motion planning in many dimensions," *CoRR*, vol. abs/1306.3532, 2013.

[23] J. Luo and K. Hauser, "An empirical study of optimal motion planning," in *IROS*, 2014, pp. 1761 – 1768.

[24] O. Salzman and D. Halperin, "Asymptotically-optimal motion planning using lower bounds on cost," *CoRR*, vol. abs/1403.7714, 2014.

[25] Z. Littlefield, Y. Li, and K. E. Bekris, "Efficient sampling-based motion planning with asymptotic near-optimality guarantees for systems with dynamics," in *IROS*, 2013, pp. 1779–1785.

[26] D. Ferguson and A. Stentz, "Anytime RRTs," in *IROS*, 2006, pp. 5369 – 5375.

[27] R. Alterovitz, S. Patil, and A. Derbakova, "Rapidly-exploring roadmaps: Weighing exploration vs. refinement in optimal motion planning," in *ICRA*, 2011, pp. 3706–3712.

[28] R. Luna, I. A. Şucan, M. Moll, and L. E. Kavraki, "Anytime solution optimization for sampling-based motion planning," in *ICRA*, 2013, pp. 5053–5059.

[29] S. Karaman, M. Walter, A. Perez, E. Frazzoli, and S. Teller, "Anytime motion planning using the RRT," in *ICRA*, 2011, pp. 1478–1483.

[30] O. Salzman and D. Halperin, "Asymptotically near-optimal RRT for fast, high-quality, motion planning," in *ICRA*, 2014, pp. 4680–4685.

[31] D. Frigioni, A. Marchetti-Spaccamela, and U. Nanni, "Fully dynamic algorithms for maintaining shortest paths trees," *J. Algorithms*, vol. 34, no. 2, pp. 251–281, 2000.

[32] G. Ramalingam and T. W. Reps, "On the computational complexity of dynamic graph problems," *Theor. Comput. Sci.*, vol. 158, no. 1&2, pp. 233–277, 1996.

[33] D. J. Webb and J. van den Berg, "Kinodynamic RRT*: Asymptotically optimal motion planning for robots with linear dynamics," in *ICRA*, 2013, pp. 5054–5061.

[34] I. A. Şucan, M. Moll, and L. E. Kavraki, "The Open Motion Planning Library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.

[35] S. Koenig, M. Likhachev, and D. Furcy, "Lifelong planning A," *Artificial Intelligence*, vol. 155, no. 1, pp. 93–146, 2004.

[36] J. Pearl, *Heuristics: Intelligent Search Strategies for Computer Problem Solving*. Addison-Wesley, 1984.

[37] K. Hauser, "Lazy collision checking in asymptotically-optimal motion planning," in *ICRA*, 2015, to appear.

[38] A. D'Andrea, M. D'Emidio, D. Frigioni, S. Leucci, and G. Proietti, "Dynamically maintaining shortest path trees under batches of updates," in *Structural Information and Communication Complexity*, ser. Lecture Notes in Computer Science, T. Moscibroda and A. Rescigno, Eds. Springer International Publishing, 2013, vol. 8179, pp. 286–297.

[39] L. Jaillet and J. M. Porta, "Asymptotically-optimal path planning on manifolds," in *RSS*, 2012.

[40] B. Akgun and M. Stilman, "Sampling heuristics for optimal motion planning in high dimensions," in *IROS*, 2011, pp. 2640–2645.

[41] N. M. Amato and L. K. Dale, "Probabilistic roadmap methods are embarrassingly parallel," in *ICRA*, 1999, pp. 688–694.

[42] J. Ichnowski and R. Alterovitz, "Parallel sampling-based motion planning with superlinear speedup," in *IROS*, 2012, pp. 1206–1212.

[43] J. Bialkowski, S. Karaman, and E. Frazzoli, "Massively parallelizing the RRT and the RRT*," in *IROS*, 2011, pp. 3513–3518.

[44] W. Wang, D. Balkcom, and A. Chakrabarti, "A fast streaming spanner algorithm for incrementally constructing sparse roadmaps," *IROS*, pp. 1257–1263, 2013.

Oren Salzman is a PhD-student at the School for Computer Science, Tel-Aviv University, Tel Aviv 69978, ISRAEL.

Dan Halperin is a Professor at the School for Computer Science, Tel-Aviv University, Tel Aviv 69978, ISRAEL.