

# Bidirectional Sampling-Based Motion Planning Without Two-Point Boundary Value Solution

Sharan Nayak, *Student Member, IEEE*, and Michael W. Otte, *Member, IEEE*

**Abstract**—Bidirectional path and motion planning approaches decrease planning time, on average, compared to their unidirectional counterparts. In single-query feasible motion planning, using bidirectional search to find a *continuous* motion plan requires an edge connection between the forward search tree and the reverse search tree. Such a tree-tree connection requires solving a two-point Boundary Value Problem (BVP). However, obtaining a closed-form two-point BVP solution can be difficult or impossible for many systems. While numerical methods can provide a reasonable solution in many cases, they are often computationally expensive, numerically unstable, or sensitive (to an initial guess) for the purposes of single-query sampling-based motion planning. To overcome this challenge, we present a novel bidirectional search strategy that does not require solving the two-point BVP. Instead of connecting the forward and reverse trees directly, the reverse tree’s cost information is used as a guiding heuristic for forward search. This enables the forward search to quickly grow down the reverse tree—converging to a fully feasible solution *without* a direct tree-tree connection and *without* the solution to a two-point BVP. We propose two algorithms that use this strategy for single-query feasible motion planning for various dynamical systems, performing experiments in both simulation and hardware testbeds. We find that these algorithms perform better than or comparable to existing state-of-the-art methods with respect to quickly finding an initial feasible solution.

**Index Terms**—Motion and Path Planning, Autonomous agents and Dynamics

## I. INTRODUCTION

SAMPLING based motion planning approaches have become popular in robotics due to their simplicity and rapid exploration of high-dimensional state spaces. They involve randomly sampling collision-free feasible states of a robot in the configuration space and moving between them in a way that respects a robot’s kinematics and dynamics.

Feasible planning [1], [2] and any-time asymptotically optimal (AO) motion planning [3], [4] are two problems in single-query [5] motion planning that have garnered interest in the motion planning community. Traditionally, feasible planning is used to find solutions to difficult problems (e.g. Piano mover’s problem [6], Alpha puzzle [7]) where finding any solution is viewed as an achievement. For this reason, feasible planning is concerned with finding any valid solution, and all solutions are considered equivalent—there is no

The authors Sharan Nayak (Graduate Student Researcher) and Michael Otte (Assistant Professor) are with the Department of Aerospace Engineering, University of Maryland, College Park (UMD), MD 20742 USA. Email: snayak18@umd.edu, otte@umd.edu.

### Different types of searches




Fig. 1: Comparison of unidirectional (left) and bidirectional (middle) searches commonly used in single-query feasible planning, and our bidirectional search approach (right) where the forward search uses heuristic from reverse search to reach the goal. Our approach is applicable when two-point BVP cannot be solved.

### Availability and absence of two-point BVP

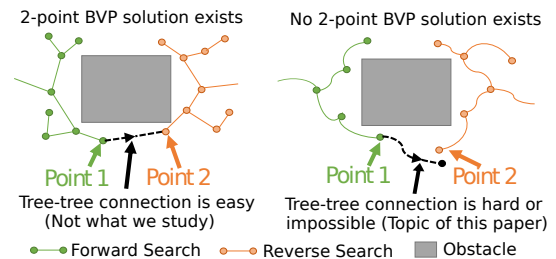


Fig. 2: Comparison of instances where the two-point BVP can (left) and cannot (right) be solved for connecting the forward and reverse trees in a bidirectional search.

requirement of minimizing cost or maximizing reward. More recently, feasible planning has also been used as the first step in any-time AO motion planning, where it is preferable to have both an actionable solution as soon as possible but then also get better solutions (with respect to a metric) whenever there is additional planning time to refine the solution toward optimal. In this article, we focus on solving the feasible motion planning problem as quickly as possible.

Most single-query sampling-based feasible motion planning algorithms use unidirectional or bidirectional search (Fig. 1) for their planning. Unidirectional search has a single search-tree built from the start state forwards to the goal state or vice-versa, whereas bidirectional search has two search trees: one expanding forward from the start or state and the other expanding backward from the goal state. Typically, the forward and reverse trees are connected if/when they grow “close” to each other. However, as we shall describe shortly, this is not the approach taken in our work.

Bidirectional searches are popular due to their improved speed [2] in finding an initial feasible solution in higher-dimensional search spaces and spaces with narrow passages [8]. Existing bidirectional search approaches [8]–[10] have focused on solving the planning problem when a steering function is available, i.e., when the two-point Boundary Value Problem (BVP) [5], [11] is solvable (Fig. 2). However, it is difficult to derive closed-form solutions for many kinodynamic robots, especially with non-holonomic constraints. Practical solutions of two-point BVPs for many systems do not exist in a closed form, and may require the use of numerical methods like shooting approaches [12], which are often computationally expensive for real-world planning time constraints. On the other hand, existing bidirectional approaches that do not require solving the two-point-BVP [13] suffer from discontinuities when connecting the two trees, while post-processing methods like curve-smoothing [2] have been employed to remove the discontinuities, which may not be trivial.

Maneuver/Trajectory libraries [14], [15] have been used to speed up execution of motion planners. The speed-up is achieved by storing the pre-computed trajectories in memory rather than calculating them online during the generation of a motion plan. However, while there are (typically) an unaccountably infinite number of trajectories that a system could potentially include in a maneuver library, only a finite subset of such trajectories can be included in practice due to memory constraints. Therefore, maneuver libraries do not remove the requirement of solving a two-point-BVP when linking two sampling-based motion planning trees rooted at arbitrary start and goal coordinates<sup>1</sup>.

**The main contribution of this article is a novel bidirectional single-query feasible sampling-based motion planning algorithm called Generalized Bidirectional Rapidly-exploring Random Tree (GBRRT) that does not require solving the two-point BVP.** The term ‘Generalized’ implies that this algorithm can be applied to many systems including systems that are: holonomic, non-holonomic, and/or kinodynamic. Our method (Fig. 3) differs from the bidirectional RRT-Connect [13] approach in that GBRRT does not try to connect the two trees when they are close; instead, it uses the heuristic information provided by the reverse search tree to help the forward tree quickly grow towards the goal configuration.

Another contribution of this article is an asymmetric variant of GBRRT called Generalized Asymmetric Bidirectional Rapidly-exploring Random Tree (GABRRT) that uses a *naive reverse search* with no dynamical trajectories in the reverse tree (Fig. 4) to provide heuristic information to the forward tree; GABRRT performs better than GBRRT in certain conditions (Section V-C). We prove that GBRRT and GABRRT are

<sup>1</sup>Special cases of symmetry can be used to create a lattice of trajectories compatible with graph search algorithms, which can be used as an alternative to sampling-based motion planning; however, such lattices present a different set trade-offs: they are not probabilistically complete and do not scale to high dimensional spaces.

## GBRRT applied to an example motion planning problem

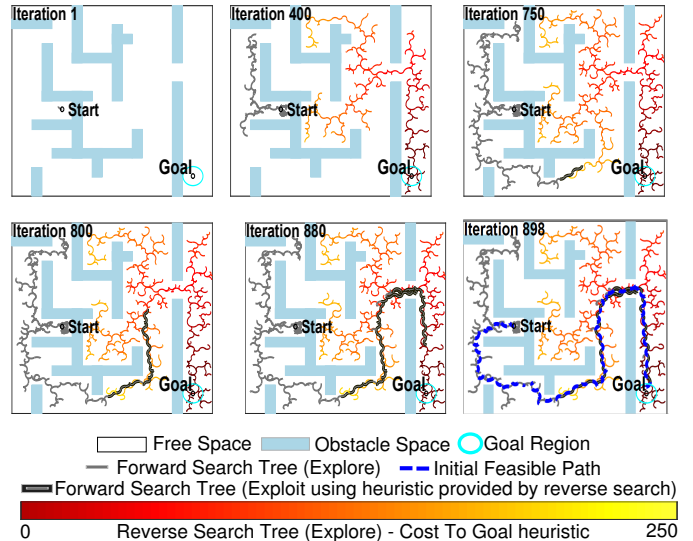


Fig. 3: Proposed Generalized Bidirectional RRT (GBRRT) algorithm, employing the forward search tree to utilize a combination of exploration and exploitation (heuristic provided by the reverse search tree) search to quickly find a feasible path in a maze.

probabilistically complete [5]. We run multiple experiments in simulation to evaluate the performance of GBRRT and GABRRT versus seven other algorithms across six different dynamical systems. We also run hardware experiments in a quadrotor test-bed to evaluate GABRRT in a more realistic physical setting.

The rest of the article is organized as follows. Section II discusses related work. Section III contains the preliminaries, including the definition of terms used throughout the article. Section IV provides a formal definition of our work. Section V discusses our proposed algorithms. Section VI presents additional details related to the specific implementations used in our experiments. Section VII provides a proof of probabilistic completeness. Section VIII describes the setup used for running experiments. Section IX presents the result of the experiments and limitations of our algorithms. Section X concludes by summarizing the contributions and main results. An appendix contains additional experiments and runtime analysis.

## II. RELATED WORK

In the following subsections, we discuss related work on existing bidirectional sampling-based motion planning algorithms, algorithms addressing the lack of two-point BVP, and algorithms that use heuristics to improve convergence.

### A. Bidirectional sampling-based algorithms

There have been a variety of bidirectional sampling-based motion planning algorithms proposed in the literature. A seminal work is bidirectional RRT-Connect [13] which grows

two trees toward each other by having each tree grow toward the nearest vertex of the opposite tree.

Jordan et al. [8] develop a bidirectional variant of RRT\* [3] that provides asymptotic optimality utilizing various heuristics and procedures to improve the convergence rate. Tahir et al. [16] and Xinyu et al. [17] use artificial potential fields [18] to generate intelligent samples using gradient descent methods to improve speed of convergence of bidirectional RRT\*. Devaurs et al. [9] update the unidirectional Transition-RRT (T-RRT) algorithm to create a bidirectional variant that achieves fast convergence and better quality paths over configuration spaces where cost functions are defined. Starek et al. [10] propose a bidirectional approach to Fast Marching Trees (FMT\*), which they show has better convergence rates to high-quality paths than many existing unidirectional planners. Strub et al. [19] use an asymmetric bidirectional search in Adaptively Informed Trees (AIT\*) with the reverse tree created using Lifelong Planning A\* (LPA\*). Although the algorithms mentioned above provide fast convergence rates, they cannot be used (in their current form) when the 2-point BVP cannot be solved, or used in practice when doing so is computationally expensive.

Bidirectional algorithms [2], [13] that do not require solving 2-point BVP, instead require extra computation to connect the discontinuity that exists between the two trees, e.g., using methods like point perturbation [13] or Bézier curves [2]. In contrast, our approach avoids the tree-tree connection problem entirely by using the heuristic provided by the reverse search tree to grow the forward tree towards the goal.

### B. Solving two-point BVP

Solving a 2-point BVP usually involves solving a differential equation constrained to the given start and end boundary conditions. This is non-trivial for robots with complex dynamics. Hence, researchers have explored solving the motion planning problem either by generating approximations to the two-point BVP [20], using shooting approaches [12], or by simplifying the dynamics by linearization [21] and then solving the resulting linear 2-point BVP. More recently, a number of random forward-propagation-based algorithms have been proposed that do not require solving the 2-point BVP. These include the Stable Sparse-RRT (SST) [22], Informed SST (iSST) [23], Dominance-Informed Region Trees (DIRT) [24], Asymptotically Optimal- RRT (AO-RRT) [25] [4] and the algorithm using bundles of edges (BOE) [26]. These are anytime AO (near) algorithms that provide fast initial solutions (specifically iSST, DIRT, BOE due to the use of heuristics) and then continue to improve their solutions over time with respect to a user-defined metric. A common theme in these random-propagation-based algorithms is that they are unidirectional. However, in expectation, bidirectional searches provide improved efficiency and faster convergence rates [5], [8] than unidirectional searches. Therefore, we present a new type of bidirectional search that does not require solving 2-

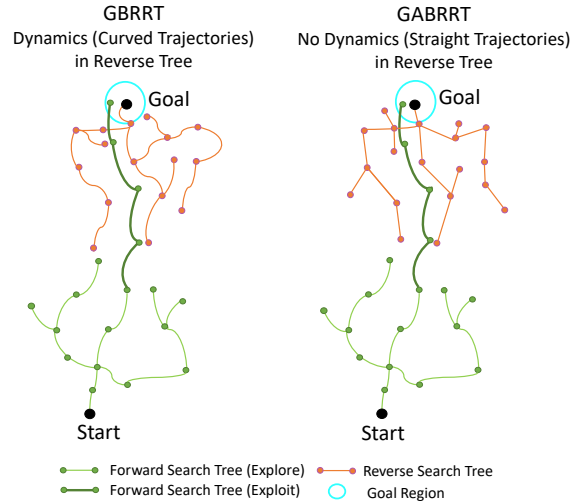


Fig. 4: Left: GBRRT (with dynamics in the reverse tree). Right: GABRRRT (with no dynamics in the reverse tree).

point BVP. Some of the comparison algorithms included in our experiments (Section VIII) are AO algorithms; however, we consider only their feasible planning phase of operation—the discovery of the first feasible solution (the phase of operation for which GBRRT could potentially be used as a drop-in replacement).

### C. Use of heuristics to improve solution convergence

A variety of heuristics have previously been used in the context of sampling-based motion planning because a well-chosen heuristic can often improve search efficiency. One common approach is sampling the goal node with a certain probability, also known as goal biasing [13]. Akgun et al. [27] use sampling heuristics to sample more nodes near a found path to increase the rate of finding a more optimal one. Urmsen et al. [28] propose the heuristically guided RRT (hRRT) where a quality measure is used to balance exploration vs. exploitation when selecting the next node to expand, leading to lower-cost solutions. Anytime-RRT [29] presents a general heuristic for rejecting samples in the context of anytime planning whenever the heuristic cost from the start to the new sample plus the heuristic cost from the sample to the goal is greater than the length of the best known path. C-Forest [30] extends this idea to AO sampling-based motion planning on both single-process and multi-process implementations. In  $D$ -dimensional Euclidean spaces ( $\mathbb{R}^D$ ), the sample region from [29], [30] takes the special form of a hyperellipse; Informed RRT\* extends the idea from [30] by showing how to generate new samples from the hyperellipse embedded in  $\mathbb{R}^D$  directly (*without* using rejection sampling).

DIRT [24] and iSST [23] use user-defined heuristics to provide high rates of convergence towards optimal solution. However, sometimes it is difficult to come up with a heuristic that provides a benefit in a particular problem instance, even

though heuristic for the entire problem domain may be easy to formulate [19]. To alleviate this issue, Strub et al. propose a geometric planner [19] that uses the problem-specific cost-to-go heuristic provided by the ever-improving reverse search tree obtained from LPA\* to converge towards the optimum. Our methods use a similar approach to [19] in using the reverse search tree to get the cost-to-goal heuristic but differ in that our algorithms do not require solving the two-point BVP.<sup>2</sup>

A technical report describing our preliminary research on bidirectional search without solving the two-point boundary value problem was released in October 2020 [31].

### III. PRELIMINARIES

Let  $\mathcal{X}$  be a smooth  $n$ -dimensional (manifold) state-space of a robot. The state-space may contain geometry states like position and angles, and higher-order dynamical states such as velocities, accelerations, etc. Let the obstacle space  $\mathcal{X}_{obs}$  be an open subset of  $\mathcal{X}$ .  $\mathcal{X}_{obs}$  is the set of all states where the system is in collision with obstacles, e.g. in the workspace. The free space  $\mathcal{X}_{free}$  is defined as  $\mathcal{X}_{free} = \mathcal{X} \setminus \mathcal{X}_{obs}$ .  $\mathcal{X}_{free}$  is the closed subset of  $\mathcal{X}$  and set of all states that are not defined by obstacle collisions. Let the start state be denoted by  $x_{start}$ , goal state  $x_{goal}$  and the goal region be  $\mathcal{X}_{goal}$  with  $x_{start} \in \mathcal{X}_{free}$ , and  $x_{goal} \in \mathcal{X}_{goal}$  and  $\mathcal{X}_{goal} \subset \mathcal{X}_{free}$ . Let  $\mathcal{U}$  be the  $m$ -dimensional input control set. The system dynamics satisfies the differential equation of the form:

$$\dot{x}(t) = f(x(t), u(t)), \quad x(t) \in \mathcal{X}, \quad u(t) \in \mathcal{U}. \quad (1)$$

System (1) is assumed to be Lipschitz continuous in both state and input arguments if there exists  $K_x, K_u > 0$  such that for all  $x_1, x_2 \in \mathcal{X}, u_1, u_2 \in \mathcal{U}$ :

$$\begin{aligned} \|f(x_1, u_1) - f(x_2, u_1)\| &\leq K_x \|x_1 - x_2\|, \\ \|f(x_1, u_1) - f(x_1, u_2)\| &\leq K_u \|u_1 - u_2\|, \end{aligned} \quad (2)$$

where  $\|\cdot\|$  is the Euclidean norm. This assumption is required for proving the probabilistic completeness of our proposed algorithms (Section VII).

**Definition 1.** A valid trajectory (also called edge)  $\mathcal{E}$  of duration  $t_{\mathcal{E}}$  is a continuous function that satisfies  $\mathcal{E} : [0, t_{\mathcal{E}}] \rightarrow \mathcal{X}_{free}$ , where  $\mathcal{E}$  is generated by forward integrating system (1) using a control function  $\Upsilon : [0, t_{\mathcal{E}}] \rightarrow \mathcal{U}$ .

This definition is valid for trajectories generated during forward search. We generate reverse tree trajectories by performing backwards integration of system (1).

In the specific implementations used in our experiments,  $t_{\mathcal{E}} \in [0, T_{max}]$  where  $T_{max}$  is a user-defined maximum propagation time. We randomly vary  $t_{\mathcal{E}}$  but apply a constant control input for the entire duration of  $t_{\mathcal{E}}$  as done in [22], [32]. Although constant control is useful in practice due to the discretization inherent in digital controllers [22], the

<sup>2</sup>We note that [13] and [28] use heuristics for feasible (initial solution) planners. [19], [23], [24], [27], [29], and [30] use heuristics for AO (initial and subsequent solutions) planners.

application of constant control is not a general requirement of our proposed algorithms.

**Definition 2.** A distance metric  $d_M$  is a function that satisfies  $d_M : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+$  where  $\times$  denotes the cartesian product and  $\mathbb{R}^+$  is the set of non-negative real numbers. Assuming  $x_1, x_2, x_3 \in \mathcal{X}$ ,  $d_M$  satisfies the following properties.

- *Non-negativity:*  $d_M(x_1, x_2) \geq 0$
- *Identity of indiscernibles:*  $d_M(x_1, x_2) = 0 \iff x_1 = x_2$
- *Symmetry:*  $d_M(x_1, x_2) = d_M(x_2, x_1)$
- *Triangle inequality:*  $d_M(x_1, x_2) + d_M(x_2, x_3) \geq d_M(x_1, x_3)$

GBRRT and GABRRT only require a distance function that satisfies the triangle inequality property, which we define next.

**Definition 3.** A distance function  $d_X$  is a function that satisfies  $d_X : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  where  $\mathbb{R}$  is the set of real numbers.

Unlike  $d_M$ ,  $d_X$  must only satisfy the triangle inequality property. Thus  $d_X$  has less restrictive requirements than  $d_M$  and this makes  $d_X$  more general than  $d_M$ .

We also define  $d_X^{ND}$  as the distance function on a lower-dimensional subspace of  $\mathcal{X}$  where the higher-order dynamics (velocity or accelerations) are not included. The superscript ND refers to “No Dynamics”. GABRRT (Algorithm 17) uses this distance function to generate the cost-to-go-heuristic from a non-dynamical reverse tree (GBRRT, Algorithm 1, does not).

**Definition 4.** The cost  $\mathcal{C}$  of a trajectory is a function that satisfies  $\mathcal{C} : \mathcal{E} \rightarrow \mathbb{R}$  where  $\mathcal{E}$  is the space of all trajectories.

The cost can be defined as the distance traveled along a trajectory or the energy or time required to execute the trajectory. In our experiments, we define  $\mathcal{C}$  as the distance along a shortest trajectory,  $\mathcal{C}(\mathcal{E}) = \limsup \sum_{i=1}^N d_X(\mathcal{E}(t_{i-1}), \mathcal{E}(t_i))$ , where  $t_0 \leq t_1 \leq t_2 \dots \leq t_N$ . By construction,  $d_X$  satisfies the triangle inequality, and can thus be used to define the “distance” between two points for the purposes of our algorithm. The trajectory cost has previously been used in the context of AO planners, but in our work, we use it to define the cost-from-start (Definition 6) and cost-to-goal (Definition 7) heuristics.

**Definition 5.** Let  $\mathcal{G} \triangleq (\mathbf{V}, \mathbf{E})$  be a graph data structure embedded in  $\mathcal{X}$  representing a motion search tree. Let  $\mathbf{V}$  and  $\mathbf{E}$  represent the node set and edge set respectively. Let  $x_a$  and  $x_b$  represent any two nodes  $\in \mathbf{V}$ . Then a feasible path  $\Gamma_{a,b}$  between  $x_a$  and  $x_b$  is defined as  $\Gamma_{a,b} \triangleq \{\mathcal{E}_1, \mathcal{E}_2 \dots \mathcal{E}_m\}$ , where  $\{\mathcal{E}_1, \mathcal{E}_2 \dots \mathcal{E}_m\}$  is an ordered set of  $m$  valid trajectories with  $\mathcal{E}_i \in \mathbf{E}$  and  $\mathcal{E}_1(0) = x_a$  and  $\mathcal{E}_m(t_{\mathcal{E}_m}) = x_b$  and  $\mathcal{E}_i(t_{\mathcal{E}_i}) = \mathcal{E}_{i+1}(0)$ .

**Definition 6.** Let  $\mathcal{G}_{for} \triangleq (\mathbf{V}_{for}, \mathbf{E}_{for})$  be a graph data structure embedded in  $\mathcal{X}$  representing a forward search tree rooted at  $x_{start}$ .  $\mathbf{V}_{for}$  and  $\mathbf{E}_{for}$  represent the node set and edge set respectively. Let  $\Gamma_{start,a}$  be a feasible path connecting  $x_{start}$  and  $x_a \in \mathbf{V}_{for}$ . The cost-from-start

heuristic  $g$  of the state  $x_a$  is defined as  $g(x_a) = \sum_{i=1}^m \mathcal{C}(\mathcal{E}_i)$ , where  $\mathcal{E}_i \in \Gamma_{start,a}$ ,  $\mathcal{E}_1(0) = x_{start}$ ,  $\mathcal{E}_m(t_{\mathcal{E}_m}) = x_a$ , and  $m$  is the number of edges connecting  $x_{start}$  to  $x_a$ .

**Definition 7.** Let  $\mathcal{G}_{rev} \triangleq (\mathbf{V}_{rev}, \mathbf{E}_{rev})$  be a graph data structure embedded in  $\mathcal{X}$  representing a reverse search tree rooted at  $x_{goal}$ .  $\mathbf{V}_{rev}$  and  $\mathbf{E}_{rev}$  represent the node set and edge set respectively. Let  $\Gamma_{a,goal}$  be a feasible path connecting  $x \in \mathbf{V}_{rev}$  and  $x_{goal}$ . The cost-to-goal heuristic  $h$  of a state  $x_a \in \mathbf{V}_{rev}$  is defined as  $h(x_a) = \sum_{i=1}^n \mathcal{C}(\mathcal{E}_i)$ , where  $\mathcal{E}_i \in \Gamma_{a,goal}$ ,  $\mathcal{E}_0(0) = x_a$ ,  $\mathcal{E}_n(t_{\mathcal{E}_n}) = x_{goal}$ , and  $n$  is the number of edges connecting  $x_a$  to  $x_{goal}$ .

**Definition 8.** The clearance of a trajectory  $\mathcal{E}$  is the maximal value  $\delta$  such that  $\mathcal{B}_\delta^m(\mathcal{E}(t)) \subset \mathcal{X}_{free} \forall t \in [0, t_{\mathcal{E}}]$ , where  $\mathcal{B}_\delta^m$  is the  $m$ -dimensional hyper-ball of radius  $\delta$  centered along a point in  $\mathcal{E}(t)$ .

Definition 8 is used in the proof of probabilistic completeness (Section VII).

#### IV. PROBLEM DEFINITION

Let  $\Gamma_{start,g}$  be a feasible path connecting  $x_{start}$  and some  $x_g \in \mathcal{X}_{goal}$ . Note that  $x_g$  is different from  $x_{goal}$  (the starting point for the reverse tree as defined in Definition 7). While  $x_{goal}$  is a valid candidate for  $x_g$  (because  $x_{goal} \in \mathcal{X}_{goal}$  by construction), the definition of  $x_g$  is deliberately more general.  $x_g$  is allowed to be any valid end state in the goal region  $\mathcal{X}_{goal}$ . Even though both  $x_g, x_{goal} \in \mathcal{X}_{goal}$ , they are usually not the same—especially in scenarios where we cannot solve the two-point BVP.

The problem definition stated below is for the feasible path planning problem, in general. It does not include a specific  $x_{goal}$ , because the problem can be solved with either a unidirectional or bidirectional algorithm.

**Problem:** Single-query feasible motion planning.

Given  $(x_{start}, \mathcal{X}_{free}, \mathcal{X}_{goal})$ , find a  $\Gamma_{start,g}$  if one exists.

#### V. ALGORITHM DESCRIPTION

The main idea behind our proposed algorithms is to use heuristic information from the reverse search tree to guide the forward search tree's advancement towards the goal. The search begins by using the forward and reverse search trees to explore different parts of the search space starting from both the start and goal states. Once the two trees encounter each other, the forward tree combines ongoing exploration with an exploitation strategy that leverages the cost-to-goal values stored in the reverse tree to direct the search. This causes the forward search tree to maneuver around obstacles and grow towards the goal quickly.

A design choice is made in our algorithms to maintain a feasible but *not* a (near-) optimal reverse tree, even though refining the reverse tree toward optimality could arguably provide a more focused guiding heuristic after the two trees become intertwined. Instead, because our algorithms are designed to solve the feasible motion planning problem,

**Algorithm 1:** GBRRT( $x_{start}, x_{goal}, \mathcal{X}_{goal}, \mathcal{U}, T_{max}, \delta_{hr}, \mathcal{P}$ )

---

```

1  $\mathcal{G}_{for} \leftarrow \{\mathbf{V}_{for} \leftarrow \{x_{start}\}, \mathbf{E}_{for} \leftarrow \{\}\}$ 
2  $\mathcal{G}_{rev} \leftarrow \{\mathbf{V}_{rev} \leftarrow \{x_{goal}\}, \mathbf{E}_{rev} \leftarrow \{\}\}$ 
3  $\mathbf{Q} \leftarrow \{\}$ 
4 for  $k \leftarrow 1$  to Miter do
5    $r_k \leftarrow \min(\gamma \left( \frac{\log|\mathbf{V}_{rev}|}{|\mathbf{V}_{rev}|} \right)^{\frac{1}{\mathcal{P}}}, \delta_{hr})$ 
6   // Reverse Tree Expansion (lines 6-11)
7    $\mathcal{E}_{rev} \leftarrow \text{RevSrchFastExplore}(\mathcal{G}_{rev}, \mathcal{U}, T_{max})$ 
8   if not CollisionCheck( $\mathcal{E}_{rev}$ ) then
9      $x_{rev} \leftarrow \mathcal{E}_{rev}.\text{initialNode}()$ 
10     $\mathbf{E}_{rev} \leftarrow \mathbf{E}_{rev} \cup \{\mathcal{E}_{rev}\}$ 
11     $\mathbf{V}_{rev} \leftarrow \mathbf{V}_{rev} \cup \{x_{rev}\}$ 
12    updatePriorityQueue( $\mathcal{G}_{for}, \mathbf{Q}, x_{rev}, r_k$ )
13   // Forward Tree Expansion (lines 12-30)
14    $q \leftarrow \mathcal{P}(k)$ 
15    $c_{rand} \sim U([0, 1])$ 
16   if  $c_{rand} < q$  then
17      $\mathcal{E}_{for} = \text{NULL}$ 
18     // Pop node from  $\mathbf{Q}$ 
19      $x_{pop} \leftarrow \text{Pop}(\mathbf{Q})$ 
20     if  $x_{pop} \neq \text{NULL}$  then
21        $\mathcal{E}_{for} \leftarrow \text{ForSrchExploit}(\mathcal{G}_{rev}, \mathcal{U}, T_{max}, x_{pop}, r_k)$ 
22     if  $\mathcal{E}_{for} = \text{NULL}$  then
23        $\mathcal{E}_{for} \leftarrow \text{ForSrchFastExplore}(\mathcal{G}_{for}, \mathcal{U}, T_{max})$ 
24   if  $c_{rand} \geq q$  or  $\mathcal{E}_{for} = \text{NULL}$  then
25     // For Probabilistic completeness
26      $\mathcal{E}_{for} \leftarrow \text{ForSrchRandomExplore}(\mathcal{G}_{for}, \mathcal{U}, T_{max})$ 
27   if  $\mathcal{E}_{for} \neq \text{NULL}$  then
28     if not CollisionCheck( $\mathcal{E}_{for}$ ) then
29        $x_{for} \leftarrow \mathcal{E}_{for}.\text{finalNode}()$ 
30        $\mathbf{E}_{for} \leftarrow \mathbf{E}_{for} \cup \{\mathcal{E}_{for}\}$ 
31        $\mathbf{V}_{for} \leftarrow \mathbf{V}_{for} \cup \{x_{for}\}$ 
32       if goalRegionReached( $\mathcal{X}_{goal}, x_{for}$ ) then
33         return Path( $\mathcal{G}_{for}, x_{start}, x_{for}$ )
34       insertToPriorityQueue( $\mathcal{G}_{rev}, \mathbf{Q}, x_{for}, r_k$ )
35 return NULL

```

---

computational effort is focused on expediting the trees' mutual encounters by exploring the free space—so that the guiding heuristic can be used as soon as possible.

In Section V-A, we present the proposed GBRRT algorithm. In Section V-B, we provide the general outline for the exploration and exploitation searches utilized in GBRRT. In Section V-C, we give a high-level description of GBRRT and then provide the algorithm details in Appendix D.

##### A. Proposed Generalized Bidirectional RRT (GBRRT)

**Inputs:** The inputs to the GBRRT (Algorithm 1) include the scenario specific start node  $x_{start}$ , goal node  $x_{goal}$ , and goal region  $\mathcal{X}_{goal}$ ; the system specific control input set  $\mathcal{U}$ ; as well as a max propagation time  $T_{max}$  (maximum duration of a trajectory); and two user defined parameters  $\delta_{hr}$  and  $\mathcal{P}$ . The parameter  $\delta_{hr}$  defines the maximum distance within which nodes from the reverse tree have a focusing effect on the growth of the forward tree, where  $0 < \delta_{hr} < \infty$  ( $hr$  stands for heuristic radius). The parameter  $\mathcal{P}$  defines the algorithm's balance between performing exploration and exploitation during forward search, where  $0 \leq \mathcal{P} \leq 1$  and, with a slight abuse of notation, 0 and 1 refer to the zero and one function, respectively.  $\delta_{hr}$  and  $\mathcal{P}$  are discussed more in Section VI-B.

**Initialization:** The forward search tree  $\mathcal{G}_{for}$  is initialized (line 1) using the node set  $\mathbf{V}_{for}$  (initialized to  $x_{start}$ ) and

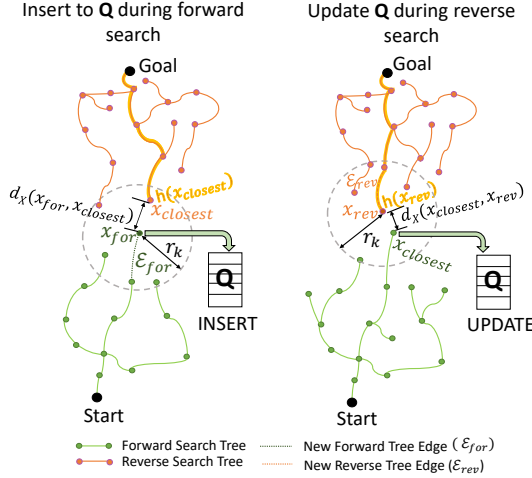


Fig. 5: Left: A new forward tree node  $x_{for}$  pushed to priority queue  $\mathbf{Q}$  with key  $d_X(x_{for}, x_{closest}) + h(x_{closest})$  during forward tree expansion. Right: the cost of the same forward tree node being updated in  $\mathbf{Q}$  during a reverse search expansion. The cost is updated only when the new cost  $d_X(x_{closest}, x_{rev}) + h(x_{rev})$  is less than its current key in  $\mathbf{Q}$ . The trees in the right sub-figure differ from those in the left to indicate that the cost update of  $x_{for}$  (right) may occur a few iterations after the insertion operation (left) — during which the trees may have been updated.

edge set  $\mathbf{E}_{for}$ . The reverse tree  $\mathcal{G}_{rev}$  is initialized (line 2) using the node set  $\mathbf{V}_{rev}$  (initialized to  $x_{goal}$ ) and edge set  $\mathbf{E}_{rev}$  (line 2). We use a priority queue  $\mathbf{Q}$  to maintain a list of potential forward tree nodes for quick expansion (line 3).

**Shrinking neighborhood radius  $r_k$ :** Shrinking the radius  $r_k$  of the neighborhood in which the two trees influence each other is necessary to balance the competing desires for a quick expected per-iteration runtime (requiring smaller  $r_k$ ) with almost sure (A.S.) asymptotic graph connectivity between  $x_{start}$  and  $x_g \in \mathcal{X}_{goal}$  (requiring larger  $r_k$ ) [3], [33]. Functions for  $r_k$  that have previously appeared in the sampling based motion planning literature include  $r_k = \min(\gamma(\frac{\log |\mathbf{V}_{rev}|}{|\mathbf{V}_{rev}|})^{1/d}, \delta_{hr})$  from [3] and  $r_k = \min(\gamma(\frac{\log |\mathbf{V}_{rev}|}{|\mathbf{V}_{rev}|})^{1/(d+1)}, \delta_{hr})$  from [33], where  $d$  is the dimension of the system and  $\gamma$  is a scenario specific percolation-theoretic parameter<sup>3</sup>. In Appendix E, we analyze the expected per-iteration runtime that results from using either function as well from using a constant  $r_k$ .

In our experiments, we calculate  $r_k$  using the expression from [33], we discuss why this rate was chosen in Ap-

<sup>3</sup>Regarding  $\gamma$ , if  $r_k$  is defined as in [3], then A.S. graph connectivity requires  $\gamma \geq (2((1 + \frac{1}{d})^{1/d} \frac{|\mathcal{X}_{free}|}{\zeta_d}))^{1/d}$ , where  $|\mathcal{X}_{free}|$  is the volume of free space and  $\zeta_d$  is the volume of the unit  $d$ -ball [3]. A new proof in [33] uses  $\gamma \geq (2 + \theta)(\frac{(1+\epsilon/4)c^*}{(d+1)\theta(1-\mu)} \frac{|\mathcal{X}_{free}|}{\zeta_d})^{\frac{1}{d+1}}$ , where  $\epsilon \in (0, 1)$  and  $\theta \in (0, 1/4)$ , and  $\mu > 0$ , and  $c^*$  is the A.S. convergent length of the path as  $n \rightarrow \infty$ . Both bounds on  $\gamma$  are of a theoretical interest rather than prescriptive because we typically do not know  $|\mathcal{X}_{free}|$  or  $c^*$ , though we may have bounds on them. Practical selection of  $\gamma$  requires making an educated guess and refining if necessary. Anecdotally, we observe that the range of  $\gamma$  producing decent results is large in practice.

pendix E.  $r_k$  is used in the insert and update  $\mathbf{Q}$  operations (line 11 and line 30) and exploitation search (line 18).

**Reverse tree expansion:** We expand the reverse search tree in lines 6-11. The reverse tree uses a pure exploration strategy for expansion using `RevSrchFastExplore` (line 6). Its purpose is to explore the state space as fast as possible and get close to and provide a focusing heuristic to the forward tree. The edge  $\mathcal{E}_{rev}$  returned by `RevSrchFastExplore` (line 6) is checked for collision using `CollisionCheck` (line 24). If no collision exists, then  $\mathcal{E}_{rev}$  and its corresponding final node  $x_{rev}$  are added to  $\mathbf{E}_{rev}$  (line 9) and  $\mathbf{V}_{rev}$  (line 10) respectively.  $x_{rev}$  is utilized to perform a heuristic update in  $\mathbf{Q}$  using `updatePriorityQueue` function (line 11). We describe `updatePriorityQueue` (Algorithm 3) as well as `insertToPriorityQueue` (Algorithm 2) later in this section.

**Forward tree expansion:** We expand the forward search tree in lines 12-30. The forward search uses a combination of exploitation (line 18) and exploration (line 20, line 22) strategies to expand the tree. The exploitation strategy uses the reverse tree’s cost information for focused expansion, while the exploration strategy uses both random exploration (for probabilistic completeness) and fast exploration (Fig. 6). These strategies are discussed in detail in Section V-B.

The choice of whether exploitation is performed in the current iteration is determined by  $\mathcal{P}$  which outputs the probability  $q$  based on  $k$ ’s value (line 12).  $q$  is compared to a value  $c_{rand}$  generated uniformly at random on  $[1, 0]$  (line 13). If  $c_{rand} < q$ , the edge is generated by an exploitation process by popping the forward tree node  $x_{pop}$  from  $\mathbf{Q}$  (line 16) and then using `ForSrchExploit` (line 18). This routine chooses the ‘best’ reverse tree node and generates a trajectory towards  $x_{pop}$  to make focused progress towards the goal.

The condition of  $x_{pop} = \text{NULL}$  implies  $\mathbf{Q}$  is empty. This may happen, e.g., near the beginning of the search when the forward tree has not yet encountered the reverse tree. If  $x_{pop} = \text{NULL}$ , the algorithm defaults to performing fast exploration using `ForSrchFastExplore` (line 20). This function is designed to quickly explore the state space (see Fig. 6). If  $c_{rand} \geq q$  or  $\mathcal{E}_{for} = \text{NULL}$  (where  $\mathcal{E}_{for} = \text{NULL}$  implies both exploitation and fast exploration were unsuccessful), the algorithm performs a random exploration using `ForSrchRandomExplore` (line 22). The use of `ForSrchRandomExplore` provides probabilistic completeness. Using only the exploitation and *fast* exploration may produce an algorithm that is not probabilistically complete. Probabilistic completeness is further discussed in Section VII.

The edge  $\mathcal{E}_{for}$  is checked for collision (line 24) and if no collision exists,  $\mathcal{E}_{for}$  and its corresponding final node  $x_{for}$  are added to  $\mathbf{E}_{for}$  (line 26) and  $\mathbf{V}_{for}$  (line 27) respectively. If  $x_{for} \in \mathcal{X}_{goal}$ , checked using `goalRegionReached` (line 28), then the output path is returned using `Path` (line 29).

**Priority queue  $\mathbf{Q}$  insert:** `GBRRT` (Algorithm 1, line 30) adds the new forward tree node  $x_{for}$  to  $\mathbf{Q}$  (Fig. 5) using `insertPriorityQueue` (Algorithm 2). This algorithm first determines the closest node  $x_{closest} \in \mathcal{G}_{rev}$  to  $x_{for}$  within

### Search types used in forward expansion of GBRRT

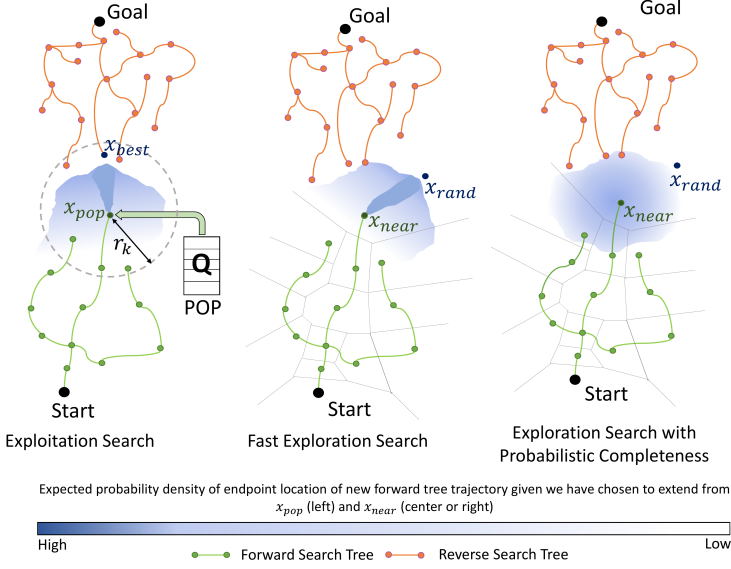


Fig. 6: **Left:** In exploitation search a forward tree node  $x_{pop}$  is popped from  $\mathbf{Q}$  and extension is biased in the direction of  $x_{best}$ , the best reverse tree node according to a user-defined cost function. The relative bias over the new endpoint's location is depicted as a tear shape oriented in the direction of  $x_{best}$ . **Center:** In fast exploration search a random node  $x_{rand}$  is sampled from the entire state-space and extension is biased in the direction of  $x_{near}$ , the nearest neighbor of  $x_{rand}$  in the forward search tree. The relative bias over the new endpoint's location is depicted as a tear shape oriented in the direction of  $x_{near}$ . **Right:** In random exploration search a random node  $x_{rand}$  is sampled from the entire state-space and extension from  $x_{near}$  uses constant random control. The probability density of endpoint location is depicted as being spread out around  $x_{near}$ . Random exploration guarantees probabilistic completeness, but is not biased toward the search-tree frontier. **All:** The new edges/trajectories are limited in duration, and duration may be random. Probability density blobs are depicted as being non-uniform to highlight the fact that the endpoint depends on  $f(x(t), u(t))$ ,  $\mathcal{U}$  and  $T_{max}$ . The embedded Voronoi diagram depicts which regions of space are nearest to each node in the forward tree. The Voronoi diagram is for illustrative purposes and is not calculated in GBRRT.

### Specific edge-generation types used in our version of GBRRT

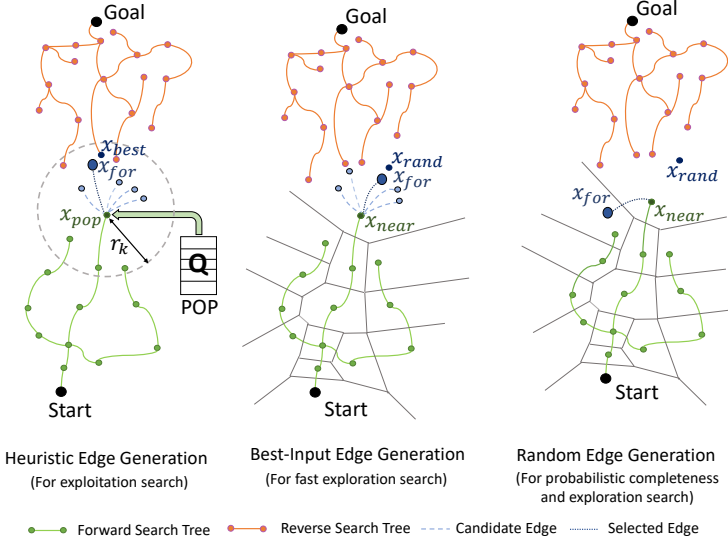


Fig. 7: These figures illustrate the specific implementation of the searches outlined in Fig. 6. The left figure shows heuristic edge-generation where a series of random edges are propagated from  $x_{pop}$  (node popped from  $\mathbf{Q}$ ) and the best edge is chosen as the one whose final node  $x_{for}$  get closest to  $x_{best}$ . The node  $x_{best}$  is chosen using the cost function  $g(x_{pop}) + d_X(x_{pop}, x) + h(x)$  where  $x$  is the node in the reverse tree. The middle figure presents best-input edge generation where a series of random edges are propagated from  $x_{near}$  and the best edge is chosen as the one that gets closest to  $x_{rand}$ . Observe that the selected edges in the left and middle figures lie in the darker tear-shaped region of probability density in Fig. 6. The right figure shows random edge-generation where an edge of random duration is propagated from  $x_{near}$  to maintain probabilistic completeness.

**Algorithm 2:** insertToPriorityQueue( $\mathcal{G}_{rev}, \mathbf{Q}, x_{for}, r_k$ )

```

1  $x_{closest} \leftarrow \mathcal{G}_{rev}.NearestNeighbor(x_{for})$ 
2  $\delta \leftarrow d_X(x_{for}, x_{closest})$ 
3 if  $\delta \leq r_k$  then
4    $\mathbf{Q}.insert(x_{for}, \delta + h(x_{closest}))$ 

```

**Algorithm 3:** updatePriorityQueue( $\mathcal{G}_{for}, \mathbf{Q}, x_{rev}, r_k$ )

```

1  $x_{closest} \leftarrow \mathcal{G}_{for}.NearestNeighbor(x_{rev})$ 
2  $\delta \leftarrow d_X(x_{closest}, x_{rev})$ 
3 if  $\delta \leq r_k$  then
4    $\mathbf{Q}.update(x_{closest}, \delta + h(x_{rev}))$ 

```

a ball of radius  $r_k$  (line 1). If  $x_{closest} \neq \text{NULL}$ , we use the key value  $d_X(x_{for}, x_{closest}) + h(x_{closest})$  to push  $x_{for}$  to  $\mathbf{Q}$  (line 4). The cost includes the actual cost of reaching  $x_{for}$  plus the heuristic estimate of reaching the goal from  $x_{for}$ . The condition of  $x_{closest} = \text{NULL}$  is true if the forward tree has not encountered the reverse tree.

Note that  $x_{for}$  is pushed to  $\mathbf{Q}$  during the forward tree's exploration and exploitation step but only popped from  $\mathbf{Q}$  during the exploitation step. This necessitates the use of a priority queue  $\mathbf{Q}$  in GBRRT to store the set of potential forward tree nodes obtained from both these steps and pop

the best one for the current exploitation step.<sup>4</sup>

**Priority queue  $\mathbf{Q}$  update:** The heuristic update in  $\mathbf{Q}$  during the reverse search (Fig. 5) occurs in updatePriorityQueue (Algorithm 3). The new reverse tree node  $x_{rev}$  obtained from expanding the reverse tree is used to select the closest forward tree node  $x_{closest}$

<sup>4</sup>A design choice worthy of mention is that we do not reinsert the node back into  $\mathbf{Q}$  once it is popped (line 16 in Algorithm 1). We found through experiments that although reinserting the node helps in some cases, there may be cases where it causes the performance to decrease. An example is reinserting a node that is close to the goal but blocked by an obstacle. There may be intelligent ways of re-pushing such a node when certain conditions are satisfied, rather than naively re-pushing them all the time. This could be a direction for future research.

within a ball of radius  $r_k$  (line 1). If  $x_{closest} \neq \text{NULL}$ , we update the heuristic of  $x_{closest}$  and corresponding key value  $d_X(x_{closest}, x_{rev}) + h(x_{rev})$  (line 4) only if the new value is less than the old key value in  $\mathbf{Q}$ .

While `insertPriorityQueue` adds a new (forward tree) node to  $\mathbf{Q}$ , `updatePriorityQueue` updates the key value and position of a (forward tree) node already in  $\mathbf{Q}$ .

**Naive goal biasing vs. Cost-To-Goal heuristic:** Naive goal biasing [13], a common approach where the goal node is sampled a certain percentage of times to improve motion planners’ performance, is not incorporated in GBRRT. GBRRT uses the cost-to-goal heuristic from the reverse search tree nodes, which is significantly more informed than traditional naive goal biasing. However, a naive goal biasing is used to potentially speed up the execution of some of the comparison algorithms specified in Section VIII (Experiment setup).

### B. General outline of exploration and exploitation searches

GBRRT uses an exploration strategy for the reverse tree expansion and a blend of exploration and exploitation strategies for the forward tree expansion. In this section, we discuss the general requirements of these strategies (Fig. 6), while specific details of the particular edge-generation implementations used in our experiments appear in Section VI.

1) **Exploitation search:** Exploitation search uses the heuristic values from the reverse tree to bias the growth of the forward tree. It begins by popping the forward tree node  $x_{pop}$  from  $\mathbf{Q}$ . Next,  $x_{pop}$  is used to select the best node  $x_{best}$  (within a radius  $r_k$ ) in the reverse tree according to a user-defined cost function. Finally, it performs forward propagation from  $x_{pop}$  (from  $\mathbf{Q}$ ) such that the endpoint of the resulting trajectory/edge gets close to  $x_{best}$  (Fig. 6). This search is visualized in Fig. 6-Left, where the probability density of the new edge’s endpoint is depicted as being higher in the region between  $x_{pop}$  and  $x_{best}$  using a tear shape). GBRRT implements this search in function `ForSrChExploit`.

2) **Fast exploration search:** A fast exploration search iteration is attempted in the event that the exploitation search iteration fails (when  $x_{pop} = \text{NULL}$ ). It involves sampling a random node  $x_{rand}$  and selecting the nearest neighbor  $x_{near}$  in the forward tree. Forward propagation from  $x_{near}$  is performed such that the endpoint of the resulting trajectory gets close to  $x_{rand}$  (Fig. 6). This has been shown to explore the state space relatively quickly [6] because, given the Voronoi partitioning generated by forward tree nodes, search is biased into the relatively large Voronoi regions associated with the search frontier. Fig. 6-Center contains a visualization in which a probability density of the new edge’s endpoint being higher in the region between  $x_{near}$  and  $x_{rand}$  is depicted using a tear shape. GBRRT implements this search in functions `ForSrChFastExplore` and `RevSrChFastExplore`. In Section VI, we discuss a method of performing fast exploration that worked well in our

experiments; faster methods may exist depending on problem scenario and system dynamics.

3) **Search with probabilistic completeness property:** This search provides probabilistic completeness for GBRRT, and also helps with exploration. It involves randomly sampling a node  $x_{rand}$  and then selecting the nearest neighbor  $x_{near}$  in the forward tree using  $d_X$ . Next, forward propagation from  $x_{near}$  is performed using a randomly selected constant control (randomness helps to ensure probabilistic completeness). This search is not biased toward unexplored regions, and so the resulting trajectory may propagate in any direction. This exploration method is visualized in Fig. 6-Right, where the probability density of endpoint of the new edge is depicted as being spread out around  $x_{near}$ . GBRRT implements this search in function `ForSrChRandomExplore`.

The Fig. 6 depictions of expected probability density being ‘tear-shaped’ and ‘spread out’ are used here for illustrative purposes and to help build intuition, they are not explicitly calculated as part of the algorithm. The actual shape of these regions will depend on many factors including system dynamics, control space, max propagation time, etc.

### C. Generalized Asymmetric Bidirectional RRT (GABRRT)

We present a variant of GBRRT called GABRRT (Algorithm 17), where non-dynamical trajectories are utilized to build the reverse tree rather than reverse trajectories based on the system’s dynamics. GABRRT generally performs well when no dynamics are considered in the planning problem (or dynamics create ‘straight-line’ like trajectories). The non-dynamical distance function  $d_X^{\text{ND}}$  is applied when updating  $\mathbf{Q}$  (Algs. 13, 11), building the reverse tree (Algorithm 14) and generating the heuristic values from the reverse tree (Algs. 16, 15). The advantage of using GABRRT is that the reverse tree is computationally inexpensive to build because it does not require performing online integration. On the other hand, omitting the effects of higher-order dynamics in the guiding heuristic may lead to poor performance in scenarios where the dynamical constraints play a relatively important role versus the geometric constraints. The algorithm for GABRRT is presented in detail in Appendix D.

## VI. GBRRT SPECIFIC IMPLEMENTATION OF SUBROUTINES USED IN OUR EXPERIMENTS

In Section VI-A, we describe how the search subroutines outlined in Section V-B have been implemented in our experiments. In Section VI-B, we explain the parameters of our specific implementations of GBRRT and GABRRT.

### A. Specific edge-generation types

Here we provide the details of the edge-generation methods for forward search (Fig. 7) used in our experiments (Section VIII). It is important to note that, while these particular implementations have the edge-generation properties required by GBRRT, many alternative strategies may exist

---

**Algorithm 4:** ForSrchExploit( $\mathcal{G}_{rev}, \mathcal{U}, T_{max}, x_{pop}, r_k$ )

```
// Heuristic edge-generation
1  $\mathbf{V}_{rev\_near} \leftarrow \{x \in \mathcal{G}_{rev}.Nodes() \mid d_X(x_{pop}, x) \leq r_k\}$ 
2  $x_{best} \leftarrow \arg \min_{x \in \mathbf{V}_{rev\_near}} (g(x_{pop}) + d_X(x_{pop}, x) + h(x))$ 
3  $\mathcal{E}_{new} \leftarrow \text{BestInputProp}(x_{pop}, x_{best}, \mathcal{U}, T_{max})$ 
4 return  $\mathcal{E}_{new}$ 
```

---

**Algorithm 5:** ForSrchFastExplore( $\mathcal{G}_{for}, \mathcal{U}, T_{max}$ )

```
// Best-input edge-generation
1  $x_{rand} \leftarrow \text{RandomState}()$ 
2  $x_{near} \leftarrow \text{NearestNeighbor}(x_{rand}, \mathcal{G}_{for})$ 
3  $\mathcal{E}_{new} \leftarrow \text{BestInputProp}(x_{near}, x_{rand}, \mathcal{U}, T_{max})$ 
4 return  $\mathcal{E}_{new}$ 
```

---

**Algorithm 6:** BestInputProp( $x_{near}, x_{rand}, \mathcal{U}, T_{max}$ )

```
// Best-input prop. for fwd. search
1  $\mathbf{E}_{fin} \leftarrow \{\}$ 
2 for  $k \leftarrow 1$  to  $N_B$  do
3  $\lfloor \mathbf{E}_{fin} \leftarrow \mathbf{E}_{fin} \cup \{\text{MonteCarloProp}(x_{near}, \mathcal{U}, T_{max})\}$ 
4  $\mathcal{E}_{new} \leftarrow \arg \min_{\mathcal{E} \in \mathbf{E}_{fin}} (d_X(\mathcal{E}.finalNode(), x_{rand}))$ 
5 return  $\mathcal{E}_{new}$ 
```

---

**Algorithm 7:** ForSrchRandomExplore( $\mathcal{G}_{for}, \mathcal{U}, T_{max}$ )

```
// Random edge-generation
1  $x_{rand} \leftarrow \text{RandomState}()$ 
2  $x_{near} \leftarrow \text{NearestNeighbor}(x_{rand}, \mathcal{G}_{for})$ 
3  $\mathcal{E}_{new} \leftarrow \text{MonteCarloProp}(x_{near}, \mathcal{U}, T_{max})$ 
4 return  $\mathcal{E}_{new}$ 
```

---

**Algorithm 8:** MonteCarloProp( $x_{initial}, \mathcal{U}, T_{max}$ )

```
// Monte Carlo propagation for forward search
1  $t_{\mathcal{E}} \leftarrow \text{Random}(0, T_{max})$ 
2  $u \leftarrow \text{Random}(\mathcal{U})$ 
3  $\mathcal{E}_{new} = \int_0^{t_{\mathcal{E}}} f(x, u) dt$  where  $\mathcal{E}_{new}(0) = x_{initial}$ 
4 return  $\mathcal{E}_{new}$ 
```

---

**Algorithm 9:** RevSrchFastExplore( $\mathcal{G}_{rev}, \mathcal{U}, T_{max}$ )

```
// Best-input edge-gen. for reverse search
1  $x_{rand} \leftarrow \text{RandomState}()$ 
2  $x_{near} \leftarrow \text{NearestNeighbor}(x_{rand}, \mathcal{G}_{rev})$ 
3  $\mathcal{E}_{new} \leftarrow \text{BestInputProp}(x_{near}, x_{rand}, \mathcal{U}, T_{max})$ 
4 return  $\mathcal{E}_{new}$ 
```

---

that also have the required properties. We use these versions in our experiments because they are easy to implement and straightforward to analyze.

1) **Heuristic edge-generation:** This edge-generation method (Algorithm 4) exploits the heuristic provided by the reverse tree. The algorithm determines the nearest vertices set  $\mathbf{V}_{rev\_near}$  to  $x_{pop}$  in  $\mathcal{G}_{rev}$  (line 1) within a ball of radius  $r_k$ . It then selects  $x_{best} \in \mathbf{V}_{rev\_near}$  (line 2) having the minimum cost  $g(x_{pop}) + d_X(x_{pop}, x) + h(x)$  (Fig. 8). Best-input propagation (Algorithm 6), explained in the next subsection, generates a new potential edge  $\mathcal{E}_{new}$  (line 3) that starts at  $x_{pop}$  and gets close to  $x_{best}$  (Fig. 7-Left).

2) **Best-Input edge-generation:** This edge-generation method (Algorithm 5) performs a fast exploration search (Fig. 7). It samples a random node  $x_{rand}$  in the state space using RandomState (line 1), then selecting its nearest neighbor in  $\mathcal{G}_{for}$  according to the distance function  $d_X$  using NearestNeighbor (line 2), and finally performs best-input propagation ([32]) using BestInputProp (line 3). Best-Input propagation (Algorithm 6) entails generating a list of  $N_B$  Monte Carlo propagation trajectories (explained in next subsection) and then selects the trajectory that gets closest to the random sampled state  $x_{rand}$  as specified by the distance function  $d_X$  (Fig. 7-Center). Although this algorithm performs fast exploration, in practice, its probabilistic completeness is still an open question [32] [34]. The reverse search version – RevSrchFastExplore, is quite similar, and is presented in Algorithm 9 and its sub-functions in Appendix C.

3) **Random edge-generation:** This edge-generation method (Algorithm 7) is required to perform search in a way that facilitates probabilistic completeness (Fig. 7). This method involves sampling a random node  $x_{rand}$  in the state space, selecting its nearest neighbor in  $\mathcal{G}_{for}$  (line 2) and then performing Monte Carlo propagation [32] using MonteCarloProp (line 3). Monte Carlo propagation (Algorithm 8) requires randomly selecting an input  $u \in \mathcal{U}$  of random duration  $t$  and then performing integration of system (1) to generate a new trajectory. We note that using this type of edge-generation within the Kinodynamic RRT algorithm as been proven to be probabilistic complete [32].

## B. Parameters of GBRRT/GABRRT

Two user-chosen parameters affect the performance of GBRRT and GABRRT: heuristic radius ( $\delta_{hr}$ ) and exploitation ratio function ( $\mathcal{P}$ ). A third parameter is required for the particular implementation of BestInputProp that we use in our experiments: the number of best-input propagation trajectories ( $N_B$ ). We now outline the role of these parameters and then perform a simulation study to evaluate the sensitivity of the performance of the algorithm to these (changing) parameters.

The heuristic radius  $\delta_{hr}$  performs two roles. First, it defines the maximum distance to the closest reverse/forward tree node that affects the cost value used to insert/update a forward tree node to the priority queue  $\mathbf{Q}$ . Second, it limits the list of the potential reverse tree nodes considered for heuristic edge-generation (line 1 in Algorithm 4). Choosing a low value of  $\delta_{hr}$  will both decrease the likelihood that a forward tree node is added to the priority queue and decrease the number of reverse nodes considered for directed propagation. On the other hand, a high value of  $\delta_{hr}$  will reduce the accuracy of the cost-to-goal heuristic and increase the chances that directed propagation is blocked by an obstacle collision.

The user-defined exploitation ratio  $\mathcal{P}$  dictates the trade-off between exploration and exploitation during forward search. We choose  $\mathcal{P}$  to be a constant probability  $q$  in our experiments. We pick  $q > 0.5$  to prioritize exploitation over exploration to quickly get to an initial feasible solution. However, a very high  $q$  can cause insufficient exploration, especially in higher-dimensional systems leading to poor performance.

The number of trajectories  $N_B$  of best-input propagation affects the per iteration run-time of best-input and heuristic edge-generation. As the value of  $N_B$  increases, the closer the set approximates the true reachability set of that node [24], at the expense of potentially increasing both the per-iteration run-time and the time to generate a feasible solution. A low value of  $N_B$  decreases the per-iteration run-time of the algorithm but potentially increases the time to generate a feasible solution. The effect of increasing  $N_B$  has less impact when using pre-computed maneuver libraries than using online integration.

To understand the effect of input parameters ( $\delta_{hr}$ ,  $q$ ,

### Pop operation from Q for heuristic edge-generation




Fig. 8: A forward tree node  $x_{pop}$  popped from  $Q$  and  $x_{best}$  chosen from reverse tree using the cost  $g(x_{pop}) + d_X(x_{pop}, x) + h(x)$  for heuristic edge generation.  $x_{best}$  is within a ball of radius  $r_k$  centered at  $x_{pop}$ .

### Covering ball sequence





Fig. 9: Covering ball sequence on reference path  $\beta$ . Only forward search tree is shown. Figure inspired from Fig. 1 in [25].

$N_B$ ) on the initial solution time, we conduct a simulation study (Fig. 10) using unicycle dynamics in the environment specified in Fig. 13. Because we have three input parameters and one output parameter (initial solution time), we fix one input parameter while plotting the other two parameters on the x-y axis and the initial solution time on the z-axis. Due to space constraints, we choose Low, Medium, and High values to plot for each fixed input parameter (see Fig. 10).


The effect of  $\delta_{hr}$  on the initial solution time is shown in Fig. 10-Top Row. The initial solution time is higher for low (left) and high (right) values of  $\delta_{hr}$  and best around the medium value (middle). The initial solution time did not vary much as  $\delta_{hr}$  changed from the medium to high value.

The effect of  $q$  on the initial solution time is shown in Fig. 10-Middle Row. The initial solution time is lower for the medium value of  $q$  (middle) when  $\delta_{hr}$  is in the medium range

### Initial solution time vs. parameters ( $N_B$ and $q$ ) with $\delta_{hr}$ fixed



### Initial solution time vs. parameters ( $\delta_{hr}$ and $N_B$ ) with $q$ fixed



### Initial solution time vs. parameters ( $\delta_{hr}$ and $q$ ) with $N_B$ fixed




Fig. 10: Variation of initial solution time vs. parameters ( $\delta_{hr}$ ,  $q$ ,  $N_B$ ) of GABRRT.  $\delta_{hr}$  and  $N_B$  axis are log scales while  $q$  axis is linear scale. We generate the figures using unicycle dynamics with the environment shown in Fig. 12. Each point is a mean of 70 trials. We set  $r_k = \delta_{hr}$  (constant radius) in all trials to highlight effect of  $\delta_{hr}$  on initial solution time.

8-14. However, low  $q$  does better when  $\delta_{hr}$  is in the low range (2-4). This happens because when  $\delta_{hr}$  is low, the exploitation search is not that effective, and hence experiments using a low value of  $q$  (more exploration) produces faster initial solutions.

The effect of  $N_B$  on the initial solution time is shown in Fig. 10-Bottom Row. The initial solution time is low and relatively constant across the tested range of  $N_B$  for medium values of  $\delta_{hr}$  around 8-14. However, the initial solution time is higher for low (and high)  $N_B$  values with low  $\delta_{hr}$  and high  $q$  because not enough (*in excess*) trajectories are considered for performing a productive exploitation search.

## VII. ANALYSIS

In this section, we prove the probabilistic completeness of GBRRT and GABRRT. Computational complexity is discussed in Appendix E. We begin by restating the definition

of probabilistic completeness referenced from [3]. Next, we prove that GBRRT is probabilistic complete and the same for GABRRT.

**Definition 9.** Let  $\mathcal{A}$  be an algorithm that solves the single-query feasible motion planning problem  $(x_{start}, \mathcal{X}_{free}, \mathcal{X}_{goal})$ . Let  $\{\mathbf{V}_n^{\mathcal{A}}\}_{n \in \mathbb{N}}$  and  $\{\mathbf{E}_n^{\mathcal{A}}\}_{n \in \mathbb{N}}$  be the sequences of vertices and edge sets returned by algorithm  $\mathcal{A}$  at each iteration  $n$ . Let the corresponding graph at iteration  $n$  be  $\mathcal{G}_n^{\mathcal{A}} = (\mathbf{V}_n^{\mathcal{A}}, \mathbf{E}_n^{\mathcal{A}})$ . The algorithm  $\mathcal{A}$  is probabilistic complete if

$$\liminf_{n \rightarrow \infty} \mathbb{P}(\exists x_g \in \mathcal{X}_{goal} \cap \mathbf{V}_n^{\mathcal{A}} \text{ such that } x_{start} \text{ is connected to } x_g \text{ in } \mathcal{G}_n^{\mathcal{A}}) = 1.$$

We leverage the proof of probabilistic completeness of RRT with random edge-generation provided by Kleinbort et al. [32] to prove the probabilistic completeness of GBRRT. For this proof, we assume the system represented by (1) resides in smooth  $n$ -dimensional Euclidean manifolds and is Lipschitz continuous (Section III). We start with a reference solution path  $\beta$  that is covered by  $(r + 1)$  hyper-balls of maximal clearance radius  $\delta$  (Fig. 9) connecting  $x_{start}$  and  $x_{goal}$ . To prove completeness property of GBRRT, we need to show that given a GBRRT vertex exists in the  $i$ th ball, the probability  $p_{GBRRT}$ , that GBRRT in the next iteration will create a vertex in the  $(i + 1)$ th ball when propagating from the  $i$ th ball is bounded below by a positive constant. To prove this statement, we state the below propositions provided by Kleinbort et al. [32] for RRT with random edge-generation. We then use these propositions to show  $p_{GBRRT} > 0$  (Lemma 1) and finally GBRRT is probabilistically complete (Theorem 1).

**Proposition 1. (From [32])** For system represented by equation (1) with Lipschitz continuity (Section III), the probability  $p_{RRT}$  that RRT will generate a vertex in the  $(i + 1)$ th ball when propagating from the  $i$ th ball in the next iteration satisfies the condition  $p_{RRT} > (|\mathcal{B}_{\delta/5}| \cdot \rho) / |\mathcal{X}|$  where  $\rho$  is a positive constant,  $\delta$  is the maximal clearance radius (Definition 8),  $\mathcal{B}_{\delta/5}$  is the hyper-ball of radius  $\delta/5$  centered at  $x_i$ ,  $|\mathcal{B}_{\delta/5}|$  and  $|\mathcal{X}|$  are the Lebesgue measures of  $\mathcal{B}_{\delta/5}$  and  $\mathcal{X}$  respectively.

**Proposition 2. (From [32])** Given Proposition 1, RRT is probabilistically complete, i.e. the probability that RRT fails to reach  $\mathcal{X}_{goal}$  from  $x_{start}$  after  $k$  iterations is at most  $a e^{-bk}$  where  $a, b > 0$ .

**Lemma 1.**  $p_{GBRRT} \geq (1 - q)p_{RRT}$  where  $q$  is the exploitation ratio and  $0 \leq q < 1$ .

*Proof.* For GBRRT, the probability of performing exploration using random edge-generation (having probabilistic completeness) during forward search is at least  $1 - q > 0$ . Using multiplication rule of probability for independent events, it directly follows that  $p_{GBRRT} \geq (1 - q)p_{RRT} > 0$ .  $\square$

**Theorem 1.** GBRRT is probabilistically complete

## Maneuver libraries




Fig. 11: Maneuver libraries for forward searches for unicycle (left) and quadrotor (right) in simulation experiments. Different shades indicate different reference distances from origin (quadrotor) and different peak velocities (unicycle) to achieve the maneuvers. Maneuvers at periodic intervals are shaded black for clarity.

*Proof.* Using Lemma 1 and Proposition 2, it follows that GBRRT is probabilistically complete.  $\square$

Like GBRRT, the probability of performing exploration using random edge-generation (having probabilistic completeness) for GABRRT is at least  $1 - q > 0$ . This leads to the following corollary.

**Corollary 1.** GABRRT is probabilistically complete

Note that the proof of Theorem 1 relies on the edge-generation method having the probabilistic completeness property in  $(1 - q) > 0$  fraction of the iterations, in expectation. It is independent of the other edge-generation implementations that are used to provide fast exploration and exploitation in the remaining  $q < 1$  proportion of iterations. Second, this analysis is only applicable to versions of GBRRT and GABRRT that generate asymptotically dense coverings of the sample space. It may *not* apply when using pre-computed maneuver libraries.

## VIII. EXPERIMENT SETUP

We run multiple software simulations and hardware experiments to test the performance of our proposed algorithms. We use the initial solution's computation time and the solution success rate as the performance metrics for comparing algorithms. We also record the initial solution cost (Fig. 15) and provide it in the appendix for completeness of analysis but note that our algorithms are not designed to produce low cost solutions. We also report the total mission time for hardware experiments, which is the sum of the computation time and flight time.

### A. Simulation experiments

We run simulation experiments using six different dynamical systems - a kinematic unicycle (3D), a kinematic quadrotor (3D), cart-pole (4D), treaded (skid-steer) vehicle (5D), second-order car with trailer (6D), and fixed-wing airplane (9D). The distance function  $d_X$  and GBRRT's parameter values for each system are specified in Table I. We run 100 trials for each dynamical system, with each trial having randomized start and goal states. At each goal state, we set the corresponding velocities to zero so that the systems comes to rest (except for the fixed-wing aircraft which is assumed

TABLE I: A summary of the dynamical systems used in experiments and the parameters used for evaluating GBRRT and GABRRT

Dynamical System	Type	State Space Dim.	Control Space Dim.	$\delta_{hr}$	$N_B$	$q$	$\gamma$	Distance Function ( $d_X$ )
Kinematic Unicycle	Non-holonomic, Kinodynamic	3	-	7	40	0.8	14	$\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$
Kinematic Quadrotor	Non-holonomic, Kinodynamic	3	-	8	90	0.8	16	$\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2}$
Cart-Pole	Holonomic [35], Kinodynamic	4	1	6	7	0.7	10	$\sqrt{(x_1 - x_2)^2 + (1.5)^2 (\Delta(\theta_1, \theta_2))^2 + (v_1 - v_2)^2 + (\omega_1 - \omega_2)^2}$
Treaded-Vehicle	Non-holonomic, Kinodynamic	5	2	3	7	0.7	7	$\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (0.25)^2 ((v_{L1} - v_{L2})^2 + (v_{R1} - v_{R2})^2)}$
Car with Trailer	Non-holonomic, Kinodynamic	6	2	4	7	0.7	8	$\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (0.25)^2 (v_1 - v_2)^2}$
Fixed-Wing Airplane	Non-holonomic, Kinodynamic	9	3	6	7	0.7	10	$\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2 + (0.9)^2 ((\dot{x}_1 - \dot{x}_2)^2 + (\dot{y}_1 - \dot{y}_2)^2 + (\dot{z}_1 - \dot{z}_2)^2)}$

The  $\Delta(\theta_1, \theta_2)$  in cart-pole  $d_X$  represents the ‘‘shortest distance’’ [36] between two angles accounting for wrap-around with  $\Delta(\theta_1, \theta_2) \in [-\pi, +\pi]$ . The  $\Delta(\theta_1, \theta_2)^2$  term is multiplied by 2.25 to give more weight to angular position of the pole. The velocity components of the treaded, car-trailer and fixed-wing systems are multiplied by values  $< 1$  to provide more priority to position than velocity. For GABRRT, we use the same parameters as GBRRT and determine  $d_X^{ND}$  from  $d_X$  by dropping the velocity terms. Finally, we choose the value of  $\gamma$  such that  $r_k < \delta_{hr}$  for approximately 80% of the initial solution planning time on average (Fig. 13).

to require a minimum velocity). We use  $S_{max}$  to represent the maximum time allotted to each algorithm to generate a system solution. We set  $S_{max} = 10$  minutes for all algorithms for unicycle, quadrotor and cart-pole and  $S_{max} = 15$  minutes for treaded, car-trailer and fixed-wing systems.

For trajectory generation, we use pre-built trajectory libraries for the unicycle and quadrotor and online integration using Runge-Kutta order 4 (RK4) [37] for the other four systems (for all algorithms). We use these different methods of generating trajectories to show that our algorithms work with distinct types of trajectory generation methods. The advantage of using trajectory libraries is that there is little overhead in considering many trajectories for the ‘‘best-input’’ propagation. We limit the number of trajectories in the considered libraries (Fig. 11) by designing them such that linear and angular velocities are zero at the start and endpoints. This velocity constraint is not assumed while performing online integration for the remaining systems. The backward integration for the reverse tree trajectories is carried out by flipping the time resolution sign and reversing the integrator’s time limits.

The dynamical systems are explained in detailed below.

1) *Unicycle*: The unicycle is 3-dimensional ( $x, y, \theta$ ) where  $x$  and  $y$  correspond to Cartesian positions and  $\theta$  is the angle from the horizontal. The zero terminal (linear and angular) velocities are achieved by using a time-varying non-linear controller [38]. The maneuvers are generated by varying the maximum linear velocity along the trajectory from 1 to 5 m/s in 1 m/s increments and the angular velocities varying from 0 to  $\pi/2$  rad/s in  $\pi/18$  rad/s increments in all four quadrants. Its kinematics is represented by  $\dot{x} = u \cos(\theta)$ ,  $\dot{y} = u \sin(\theta)$ ,  $\dot{\theta} = \omega$  where  $u$  is linear speed and  $\omega$  is angular velocity.

2) *Quadrotor*: The quadrotor system is 3-dimensional ( $x, y, z$ ) where  $x, y$ , and  $z$  correspond to Cartesian positions. The zero terminal velocities are achieved using a feedback linearization controller [39]. The maneuver library is generated by providing final positions 1 to 5 meters distance away from

the initial position (origin) in  $10^\circ$  increments in the heading. Its dynamics is represented by  $M\dot{v} + C(v)v + Dv + g(\eta) = u$ , where  $M$  is the mass matrix,  $C(v)$  is the Coriolis-centripetal matrix,  $D$  is the damping matrix,  $g$  is the gravity vector,  $\eta, v$  and  $u$  are the state, velocity and control input vectors respectively. The definition of these terms are given in [39].

3) *Cart-Pole*: The cart-pole system [22] [40] consists of a pendulum attached to a cart which is restricted to move along a line. The state space is 4-dimensional ( $x, \theta, v, \omega$ ) where  $x$  and  $v$  corresponds to the position and linear velocity of the cart and  $\theta$  and  $\omega$  corresponds to the angular position (from downward vertical) and angular velocity of the pendulum. The control space is 1-dimensional which is the force  $F$  on the cart. Its dynamics is represented by  $\dot{x} = v$ ,  $\dot{\theta} = \omega$ ,  $\dot{v} = \frac{(I+mL^2)(F+mL\omega^2 \sin(\theta)) + (mL)^2 \cos(\theta) \sin(\theta)g}{(M+m)(I+mL^2) - (mL)^2 \cos^2 \theta}$ ,  $\dot{\omega} = \frac{(-mL \cos(\theta))(F+mL\omega^2 \sin(\theta)) + (M+m)(-mgL \sin(\theta))}{(M+m)(I+mL^2) - (mL)^2 \cos^2 \theta}$ , where  $I$  is the moment of inertia,  $L$  is the length of the pendulum,  $g$  is the acceleration due to gravity,  $M$  and  $m$  are the masses of the cart and pendulum respectively.

4) *Treaded (Skid-Steer) Vehicle*: The skid steer vehicle [41] [42] is 5-dimensional ( $x, y, \theta, v_L, v_R$ ) where  $x$  and  $y$  are Cartesian positions,  $\theta$  is the angle of vehicle from horizontal, and  $v_L$  and  $v_R$  are the velocities of left and right wheels respectively. The control space is 2-dimensional ( $a_L, a_R$ ) which corresponds to accelerations of left and right wheels. Its dynamics is represented by  $\dot{x} = v_x \cos(\theta) - v_y \sin(\theta)$ ,  $\dot{y} = v_x \sin(\theta) + v_y \cos(\theta)$ ,  $\dot{\theta} = w_z$ ,  $\dot{v}_L = a_L$ ,  $\dot{v}_R = a_R$ , where  $v_x, v_y$  and  $w_z$  are functions of  $v_L$  and  $v_R$  and their definitions are provided in [41].

5) *Second-Order Car with Trailer*: The second order car with trailer [24] [43] is a 6-dimensional ( $x, y, \theta, v, \omega, \theta_1$ ) system where  $x$  and  $y$  are Cartesian positions,  $v$  is the linear velocity,  $\theta$  is the angle of car from the horizontal,  $\omega$  is the angular velocity, and  $\theta_1$  is the angle of the trailer from the horizontal. The control space is 2-dimensional ( $a, \alpha$ ) which corresponds to the desired linear and angular accelerations. Its dynamics is represented by  $\dot{x} = v \cos(\theta) \cos(\omega)$ ,

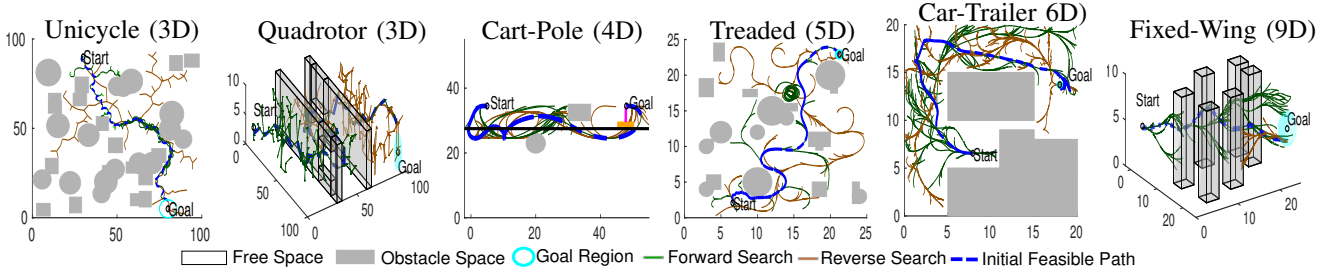


Fig. 12: Initial feasible solution found by GABRRT (unicycle and quadrotor) and GBRRT (cart-pole, treaded, car-trailer, and fixed-wing).

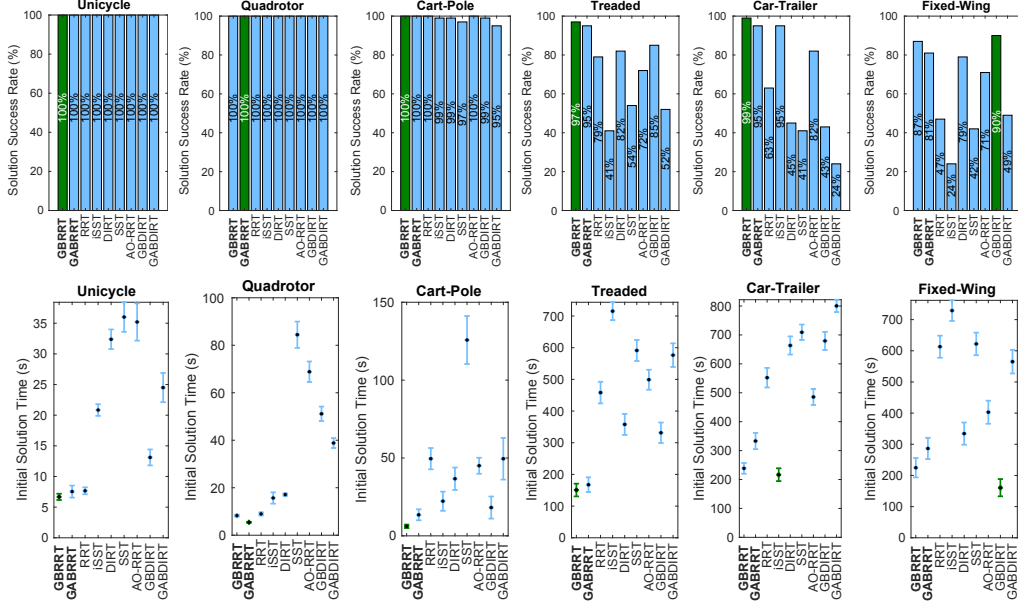


Fig. 13: Top row shows the percentage solution success rate, and the bottom row shows the mean and standard error of the initial solution time for 100 trials of running algorithms using various systems. The best performing algorithm is highlighted in green. The solution success rate is less than 100 % for algorithms because they are not able to find a solution within the max time  $S_{max}$ . The mean and standard error are calculated using results from all 100 trials (successful and unsuccessful) with  $S_{max}$  value used for unsuccessful trials.

$$\dot{y} = v \sin(\theta) \cos(\omega), \quad \dot{\theta} = v \sin(\omega), \quad \dot{v} = a, \quad \dot{\omega} = \alpha, \\ \dot{\theta}_1 = v \sin(\theta - \theta_1) \text{ respectively.}$$

6) *Fixed-Wing Airplane*: The fixed-wing airplane [22] [44] is 9-dimensional  $(x, y, z, v, \omega, \beta, T, \alpha, \mu)$  where  $x, y, z$  are the Cartesian positions,  $v$  - speed,  $\omega$  - flight path angle,  $\beta$  - heading angle,  $T$  - thrust per unit mass,  $\alpha$  - angle of attack and  $\mu$  - wind axis roll angle. The control space is 3-dimensional  $(T_c, \alpha_c, \mu_c)$  which are the commanded thrust per unit mass, commanded angle of attack and commanded wind axis roll angle respectively. Its dynamics is represented by  $\dot{x} = v \cos(\omega) \cos(\beta)$ ,  $\dot{y} = v \cos(\omega) \sin(\beta)$ ,  $\dot{z} = v \sin(\omega)$ ,  $\dot{T} = c_T (T_c - T)$ ,  $\dot{\alpha} = c_\alpha (\alpha_c - \alpha)$ ,  $\dot{\mu} = c_\mu (\mu_c - \mu)$ ,  $\dot{v} = (T \cos(\alpha) - kv^2 C_D(\alpha)) - g \sin(\omega)$ ,  $\dot{\omega} = \left( \frac{T \sin(\alpha)}{v} + kv C_L(\alpha) \right) \cos(\mu) - \frac{g \cos(\omega)}{v}$ ,  $\dot{\beta} = \left( \frac{T \sin(\alpha)}{v} + kv C_L(\alpha) \right) \frac{\sin(\mu)}{\cos(\omega)}$ , where  $g$  is the acceleration due to gravity,  $k$  is the scaled inverse of the wing loading,  $C_D(\alpha)$  and  $C_L(\alpha)$  are the coefficients of lift and drag respectively.

TABLE II: Experimental statistics of GBRRT for different systems

Informative Statistic	Uni.	Quad.	C.Pole	Tread.	C.Trail.	F.Wing
$R_{r_k}$	24.95	13.27	58.41	61.28	54.96	58.79
$R_{edge}$	2.48	2.07	1.33	3.33	10.65	3.90
$R_x$	0.07	0.08	0.11	0.12	0.2	0.24
$R_{max}$	0.049	0.057	0.095	0.084	0.14	0.158

$R_{r_k}$  and  $R_{edge}$  are averaged over 100 trials while  $R_x$  and  $R_{max}$  remain same across all trials.

### B. Comparison of algorithms

We compare our proposed algorithms with five existing state-of-the-art modern algorithms that do not require a two-point BVP solver. This includes RRT using best-input propagation (5% goal bias), SST [22] (5% goal bias), iSST [23], DIRT [24], and AO-RRT [25] [4] (5% goal bias). Even though SST, iSST, DIRT, and AO-RRT are AO algorithms which improve solution over time, we only run these algorithms until an initial solution is found because we are solving the quick initial feasible planning problem. DIRT and

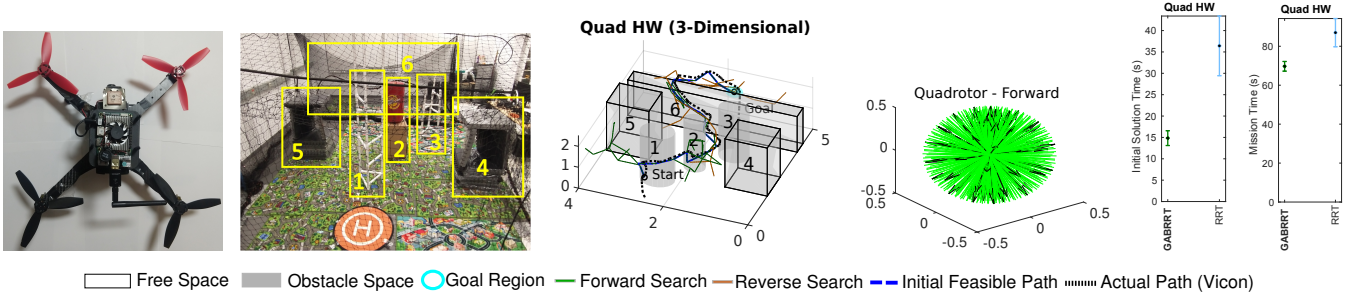


Fig. 14: Left image shows the quadrotor used in our hardware experiments. The second image shows the actual view of the obstacle course and the third image shows the model view with actual path (Vicon) overlaid. The fourth image shows the maneuver library used for performing forward search. The right figures show the initial solution time and mission time for 100 trials of hardware experiments.

iSST are heuristic-based algorithms that require a heuristic during the execution of the algorithms. As suggested in [23] (the original iSST paper) we generate the heuristics for iSST and DIRT using a lower dimensional roadmap built using PRM [45]. We use the same random seed to generate the roadmap for both iSST and DIRT in each trial. Because we are interested in single-query feasible planning, we factor in the time needed to generate the roadmap into the execution time that we report for these algorithms. We do not compare with AIT\* [19] because it is a geometric planner and requires solving the two-point BVP for its operations.

We create two new bidirectional variants of DIRT called GBDIRT (Generalized Bidirectional DIRT) and GABDIRT (Generalized Asymmetric Bidirectional DIRT) to include in our comparison. These variants use our continually growing reverse search tree to generate the heuristic instead of using a roadmap. Similar to GBRRT and GABRRT, GBDIRT grows a dynamical reverse tree whereas GABDIRT uses a non-dynamical reverse tree to obtain the heuristic. Since the original DIRT algorithm (line 4 in Algorithm 1 in [24]) requires a guiding heuristic during the entire execution of the algorithm, we choose the closest reverse tree node to the current forward tree to obtain the heuristic for GBDIRT and GABDIRT.

Finally, we also compare our proposed algorithms with a numerical bidirectional variant of RRT (NBRRT for brevity) that uses exploration to grow the forward and reverse trees and a numerical boundary value solver to connect the two trees when they are within  $\delta_{hr}$  distance of each other. This algorithm requires solving only a single two-point BVP to connect the forward and reverse trees, unlike algorithms like AIT\* that require solving many two-point BVPs for their operations. We use Python 2.7 scipy library’s `solve_bvp` [46] solver that uses a collocation formulation to solve the two point BVP. This solver requires the number of boundary conditions be equal to the number of equations describing the system’s dynamics. We have  $2n$  boundary conditions ( $n$  for both the beginning and end of the connecting trajectory) but only  $n$  equations describing the system’s dynamics. Therefore, we obtain  $n$  additional equations by differentiating the original  $n$  equations with respect to time. We verify the solution returned by the solver is valid (kinodynamically

feasible) by forward integrating the system using the input solution returned by the solver, and then accept the trajectory (Table III). For the unicycle and quadrotor, we select the best trajectory from their respective trajectory library that reduces the error on the final intended point. We implement the algorithms in Python 2.7 and perform the experiments on a system with an Intel i7-7700 4-core CPU with 32GB RAM.

### C. Statistics chosen for GBRRT insights

We report various experimental statistics for GBRRT for all systems considered (Table II). These statistics provide insights into the inner workings of GBRRT across the different systems, and include:

- $R_{r_k}$  - Average percentage of times the forward and reverse trees are within  $r_k$  distance away.
- $R_{edge}$  - Average ratio of  $\delta_{hr}$  to the average length/cost of kinodynamic edges in the forward and reverse tree.
- $R_x$  - Ratio of  $\delta_{hr}$  to the length of the ‘ $x$ ’ dimension of state space common to all systems.
- $R_{max}$  - Ratio of  $\delta_{hr}$  to the maximum distance function value possible across the state space.

We get these statistics using the same set of 100 GBRRT trials used in the comparison experiments (Fig. 13).

### D. Hardware experiments

We perform the hardware experiments in an obstacle course of size  $5\text{m} \times 4\text{m} \times 2.3\text{m}$  (Fig. 14) with start and end locations set at  $(0.2, 2.5, 1.0)$  and  $(4.5, 1.7, 1.5)$  respectively. The obstacle course is traversed using a modified Bebop 2 quadrotor [47] in 3D space. The quadrotor is equipped with an Intel UP board (Fig. 14) having Intel Atom x5-z8350 4-core, 1.44 GHz CPU with 4GB RAM for running the algorithms. We run 100 experimental trials, with each trial having a different seed. We use a Proportional-Derivative (PD) controller for position tracking and a Vicon motion capture system for state information. The obstacles in the course are modeled using 3D cylinders and rectangularoids. The quadrotor is modeled as a sphere of radius 20 cm for easy collision detection. We compare GABRRT with RRT

TABLE III: Numerical boundary value solver results

Metric	Uni.	Quad.	C.Pole	Tread.	C.Trail.	F.Wing
Solution Time	15.18	35.93	171.29	171.33	528.91	796.32
Total Success Rate	100	100	91	99	82	23
Conn. Success Rate	67	1	58	92	70	0
Tolerance	0.01	0.01	0.75	0.2	0.3	0.75

Results averaged over 100 trials. The connected success rate refers to the percentage of trials the solution was found using the numerical boundary value solver. It is less than the total success rate because there were some trials where the forward tree reached the goal region without connection to the reverse tree (attempts to connect using 2-point BVP solver were unsuccessful).

in our experiments. We choose GABRRT because it is best-performing proposed algorithm in our quadrotor simulation experiments and RRT is the best non-proposed algorithm (Fig. 13). We use a maneuver library with each trajectory having a length of 0.5 m with the start and end velocities set to zero (Fig. 14). We use the same distance function  $d_X$  as for the quadrotor simulation. We set GABRRT parameters  $\delta_{hr} = 1.0$ ,  $N_B = 90$  and  $q = 0.8$ .

## IX. RESULTS

This section provides the simulation experiment results (Section IX-A), NBRRT comparison results (Section IX-B), GBRRT experimental statistics interpretation (Section IX-C), the hardware experiment results (Section IX-D) and finally the limitations of our algorithms.

### A. Simulation experiment results

Example feasible paths generated by GBRRT and GABRRT for the systems we consider are shown in Fig. 12. The plots for the solution success rate percentage and initial solution time metrics are presented in Fig. 13. Considering these metrics, GBRRT performs best for the unicycle, cart-pole, threaded and, car-trailer (success rate metric) systems, GABRRT performs the best for the quadrotor systems, and GBDIRT performs the best for the fixed-wing system.

GABRRT uses an inexpensive non-dynamical reverse tree to generate the cost-to-go heuristic. For this reason, it performs better for systems having ‘straight-line’ like trajectories (quadrotor). On the other hand, although iSST and DIRT also use heuristics from a non-dynamical PRM roadmap, they take longer to generate a feasible solution. One reason for this is that roadmaps usually takes longer to construct than trees containing the same number of nodes (due to the fact that they typically involve more edges per node).

GBRRT uses the cost-to-go heuristic from a reverse tree with dynamics. This makes GBRRT perform better for higher-dimensional systems (cart-pole, threaded, car-trailer, and fixed-wing) that have velocity constraints in addition to position constraints at their goal states, and the distance functions have higher-order dynamics. We believe this also

causes GBRRT to perform better than DIRT and iSST for these systems (DIRT and iSST use a heuristic generated from a lower-dimensional non-dynamical roadmap). In such cases we also observe that if the goal region is made large for these systems, then GABRRT performs on par or better than GBRRT. We believe this happens because the size of the state space’s positional dimensions is relatively large compared to the size of the velocity dimensions.

GBDIRT performs the best for the fixed-wing system and performs better than DIRT for majority of the systems. We believe this demonstrates that, at least in the case of single query feasible planning, trees are usually better suited to the task of generating a cost-to-go heuristic than roadmaps. We note that this parallels the broader use of trees (vs. roadmaps) for single query planning more generally.

### B. NBRRT comparison results

NBRRT (Table III) does not perform the best in terms of the initial solution time metric among the algorithms tested. The results for the solution success rate were mixed, with NBRRT performing the best for the threaded system and the worst for the fixed-wing system. We found the connection success rate is zero for all systems if a low enough tolerance value is used. Although NBRRT’s performance is mixed, we state that if the problem can be solved efficiently using numerical methods or has special symmetries such that solving the two-point BVP is trivial, then it is advisable to use one of the pre-existing numerical bidirectional algorithms. However, if numerical methods are not easy to apply, then our approach provides a general solution method that can be used.

### C. Interpretation of GBRRT experimental statistics

Table II presents the experimental statistics for GBRRT. The interpretation of the various statistic results follows:

$R_{rk}$  - The percent overlap between the forward and reverse trees is lower for lower-dimensional systems (unicycle, quadrotor). We observe that once the trees meet each other, the reverse tree heuristic directs the forward tree towards the goal. However, the percent overlap is much higher for higher-dimensional systems (sometimes called ‘curse of dimensionality’ [48]). This causes the feasible solution to be found much later after the initial encounter of the forward and reverse trees.

$R_{edge}$  - This value is  $> 1$  for all systems (because  $\delta_{hr}$  chosen is always greater than the average length/cost of the kinodynamic edge). This is necessary to make the forward tree consider the cost-to-go heuristic from reverse tree nodes sufficiently far from its current location to make quick progress towards the goal.

$R_x$  and  $R_{max}$  - These values are less than one for all the systems because the  $\delta_{hr}$  value chosen is much less than the diameter of the systems’ state space. These values are higher for the higher-dimensional systems (car-trailer, fixed-wing) because their state-spaces have smaller positional dimensions compared to the lower-dimensional systems (Fig. 12).

#### D. Hardware experiment results

The initial feasible path generated by GABRRT in a hardware experiment trial with the actual path (obtained using Vicon motion capture system) and the results for 100 trials of hardware experiments are shown in Fig. 14. GABRRT, on average, performs better than RRT with respect to initial solution time and total mission time (solution time + flight time). A Student’s T-test [49] and Kolmogorov Smirnov (KS) [50] tests verify that the initial solution times for GABRRT and RRT are significantly different. However, we did not observe statistical significance for the mission-time (KS test). This is not surprising due to the fact that flight time is proportional to the path length, and GABRRT does not optimize for path length since it is a feasible motion planning algorithm.

#### E. Limitations of our proposed algorithms

Our proposed algorithms are parameter-based, and their performance depends on the appropriate selection of parameters, including: heuristic radius  $\delta_{hr}$ , exploitation ratio function  $\mathcal{P}$  and number of best-input trajectories  $N_B$ . The effects of using different  $\delta_{hr}$ ,  $\mathcal{P}$  and  $N_B$  values are discussed in Section VI-B. We provide guidelines for parameters selection as follows:

1) *Number of Maneuvers ( $N_B$ ):* The solution time is largely unaffected for systems using maneuver libraries when we increase  $N_B$  within an extensive range (Fig. 10). However, if we choose  $N_B$  to be a small value, the solution time does increase. Hence for systems using maneuver libraries, we recommend a high value dependent on the number of maneuvers in the library. For systems using online integration, a value between 5-10 worked well in our experiments. Higher-dimensional systems generally performed well in our experiments when  $N_B$  was lower.

2) *Exploitation ratio ( $q$ ):* We recommend using a value between 0.5-1.0 to prioritize exploitation over exploration to find the initial solution quickly. For lower-dimensional systems, a value between 0.8-0.95 worked well in our experiments. For higher-dimensional systems, the range 0.5-0.7 worked well.

3) *Heuristic radius ( $\delta_{hr}$ ):* The heuristic radius is dependent on the system and the environment used. A general rule of thumb for  $\delta_{hr}$  is that it should be chosen less than the “average size” of obstacles in the environment (assuming positional coordinates used in  $d_X$ ) so that heuristic from reverse tree nodes blocked by obstacles are not considered. For example,  $\delta_{hr} > 10$  may not work well for the quadrotor environment (Fig. 12) since the large  $\delta_{hr}$  may cause our algorithms to consider the heuristic from reverse tree nodes blocked by the walls (whose thickness is 10). In addition,  $\delta_{hr}$  should be chosen, so it is greater than the average size of the kinodynamic edge (Table II) so that it considers reverse tree nodes sufficiently far away from its current point to make progress towards the goal. In our experiments, a value between 3-10 worked well. Finally, if none of these guidelines

work, the last resort is to perform a simple parameter sweep for the system in question or a similar problem.

GBRRT (GABRRT), like RRT [13] [32], is a feasible motion planning algorithm. As such, it is not designed to be asymptotically optimal. That said, it is of potential relevance to any-time planning with asymptotic algorithms because one potential application of GBRRT (GABRRT) is that it can be used to provide an initial solution and then be improved using another method. This is because feasible motion planning algorithms tend to find an initial solution more quickly than AO motion planning algorithms. Thus if an any-time solution is desired, it is possible to start planning with a feasible algorithm until the first solution is found and then switch to an AO algorithm. In practice, both algorithms often use the same underlying graph data structure and subroutines. GBRRT (GABRRT) could potentially be made asymptotically optimal by replacing the node selection process with nearest neighbors to random selection [3] [32] at the expense of reduced convergence rates. In addition, the cost function used for inserting the node in the priority queue  $\mathbf{Q}$  has to be updated to use  $g(x_{for}) + d_X(x_{for}, x_{closest}) + h(x_{closest})$  (see Fig. 5) instead of  $d_X(x_{for}, x_{closest}) + h(x_{closest})$  to prevent the algorithm from only selecting nodes close to the goal for expansion and thus consider the total cost from start to goal.

Applying GBRRT to systems involving interactions with the environment like rigid body systems having frictional contacts is challenging. This is because computing the reverse propagation primitive is hard for these systems since we cannot uniquely solve the backward simulation problem [51]. A partial solution based on Linear Complementary Problem (LCP) formulation is provided in [51] to alleviate the issue, but future research challenges remain.

## X. CONCLUSION

We present Generalized Bidirectional RRT (GBRRT), a new bidirectional single-query sampling-based motion planning algorithm that does not require a two-point BVP solver. GBRRT produces a continuous motion plan for solving the initial feasible motion planning problem. Instead of connecting the forward and reverse trees, the forward tree to uses a cost-to-goal heuristic provided by the reverse search tree to quickly find an initial feasible solution. The cost-to-goal heuristic calculated on the fly applies to the specific problem instance and alleviates the user from coming up with a heuristic.

We provide a variant of GBRRT called Generalized Asymmetric Bidirectional RRT (GABRRT) that uses a computationally inexpensive non-dynamical reverse search to generate the cost-to-go heuristic. In experimental trials GABRRT performs better than GBRRT in scenarios without higher-order dynamics. We also combine the idea used in GBRRT/GABRRT (reverse tree provided heuristic) to an existing state-of-the-art unidirectional algorithm called DIRT [24] to create Generalized Bidirectional DIRT (GBDIRT) and Generalized Asymmetric Bidirectional (GABDIRT). GB-

DIRT outperforms DIRT in the majority of the systems tested.

We prove that GBRRT and GABRRT are probabilistically complete. Finally, we run multiple simulation experiments using systems of varying dimensions and hardware experiments using a quadrotor. These show that our algorithms quickly find a feasible solution in many different scenarios.

Considering solution success rate and average initial solution time, GBRRT performs the best for the unicycle, cart-pole, threaded and car-trailer systems (solution success rate) and GBDIRT for the fixed-wing system. This is because these systems generally consider higher-order dynamics in their planning and distance functions. GABRRT performs the best for the quadrotor system where higher-order dynamics ('straight-line' like trajectories) are not considered. GABRRT also outperforms RRT in the hardware experiments we run in initial solution time and mission time metrics. These results shows the power of performing a bidirectional search and using an on-the-fly calculated problem-specific heuristic in improving the performance of finding an initial solution.

A potential avenue for future research is developing an AO version of our proposed algorithms. This version should be able to provide good convergence rates while still maintaining asymptotically optimality. Other possible directions include testing the performance of GBRRT and GABRRT using different definitions of cost and distance functions.

#### ACKNOWLEDGMENTS

We thank Chahat Deep Singh and Nitin Sanket from the Perception and Robotics Group at the University of Maryland (UMD) for providing the modified bebop 2 quadrotor used in our experiments, and Mohamed Khalid M. Jaffar in the Motion and Teaming Lab at UMD for his valuable feedback during manuscript preparation. We genuinely appreciate the reviewers' and editor's feedback and comments, many of which helped us to improve the structure and performance of the proposed algorithms, as well as helping us to substantially improve the manuscript from its preliminary version. This research was supported by the Minta Martin Research Fellowship Fund, the Clark Doctoral Fellowship, and U.S. Office of Naval Research (ONR) award #N00014-20-1-2712.

#### APPENDIX

##### A. Initial solution cost

We present the figures for the initial solution cost in Fig. 15.

##### B. Comparison with variations of GBRRT

We compare GBRRT with algorithms created from modifications to its operations to show the effect of different operations on the performance metrics (Fig. 16). The algorithms include GBRRT-NF (No fast exploration (Alg. 5), replaced by random exploration), GBRRT-NU (No priority Q-update, (Alg. 3)), GBRRT-NE (No Exploitation, (Algs. 2, 3, 4)) and

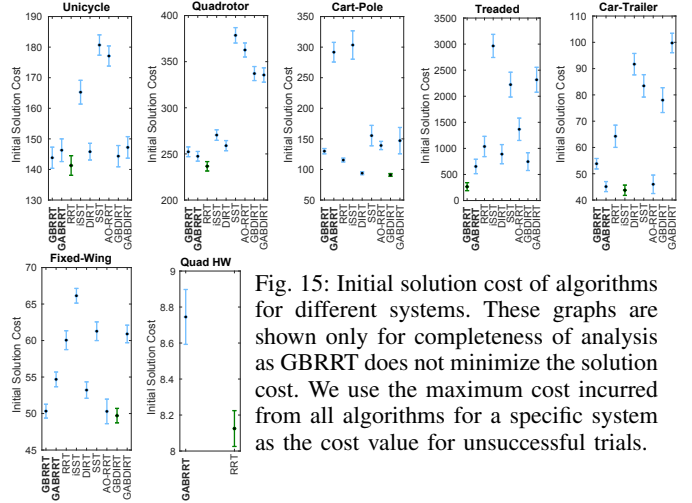


Fig. 15: Initial solution cost of algorithms for different systems. These graphs are shown only for completeness of analysis as GBRRT does not minimize the solution cost. We use the maximum cost incurred from all algorithms for a specific system as the cost value for unsuccessful trials.

GBRRT-RU (Range-based update instead of nearest update in Alg. 3). Fig. 16 shows that baseline GBRRT performs at least as good, on average and in most cases, as these other variations.

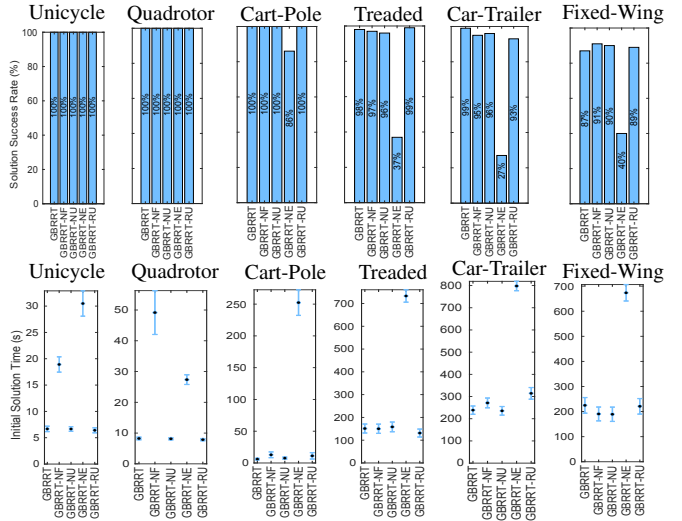


Fig. 16: Comparison of success rate and solution time metrics for different modifications of GBRRT — GBRRT-NF (No Fast exploration), GBRRT-NU (No Priority Q-Update), GBRRT-NE (No Exploitation) and GBRRT-RU (Range-Based Update).

##### C. Reverse search algorithms

Algorithms 10 and 12 present reverse search subroutines.

##### D. Generalized Asymmetric Bidirectional RRT (GABRRT)

GABRRT appears in Algorithm 17 and its subroutines in Algorithms 11, 13-16. Differences between GBRRT and GABRRT appear brown.  $\mathcal{G}_{rev}^{ND}$ ,  $\mathcal{E}_{rev}^{ND}$ ,  $\mathbf{E}_{rev}^{ND}$ ,  $d_X^{ND}$  denote the non-dynamical reverse tree, reverse trajectory, edge set, and distance functions, respectively. Extend returns a straight line trajectory  $\mathcal{E}_{new}^{ND}$  starting at  $x_{near}$  and ending at a point on the line between  $x_{near}$  and  $x_{rand}$ . The length  $\|\mathcal{E}_{new}^{ND}\| \leq \epsilon$ ,

**Algorithm 10: BestInputProp** $(x_{near}, x_{rand}, \mathcal{U}, T_{max})$ 

```

// Rev. search Best-input prop.
1  $\mathbf{E}_{initial} \leftarrow \{\}$ 
2 for  $k \leftarrow 1$  to  $N_B$  do
3    $\mathbf{E}_{initial} \leftarrow \mathbf{E}_{initial} \cup$ 
   {MonteCarloProp( $x_{near}, \mathcal{U}, T_{max}$ )}
4  $\mathcal{E}_{new} \leftarrow \arg \min_{\mathcal{E} \in \mathbf{E}_{initial}} (d_X(x_{rand}, \mathcal{E}.initialNode()))$ 
5 return  $\mathcal{E}_{new}$ 

```

**Algorithm 11: UpdatePQueueND** $(\mathcal{G}_{for}, \mathbf{Q}, x_{rev}, r_k)$ 

```

1  $x_{closest} \leftarrow \mathcal{G}_{for}.NearestNeighbor(x_{rev})$ 
2 if  $x_{closest} \leq r_k$  then
3    $\mathbf{Q}.update(x_{closest}, d_X^{ND}(x_{closest}, x_{rev}) + h(x_{rev}))$ 

```

**Algorithm 16: ForSrchExploitND** $(\mathcal{G}_{rev}^{ND}, \mathcal{U}, T_{max}, x_{pop}, r_k)$ 

```

1  $x_{best} \leftarrow \min_{g(x_{pop}) + d_X^{ND}(x_{pop}, x) + h(x)} \{x \in \mathbf{V}_{rev\_near}\}$ 
2  $\mathcal{E}_{new} \leftarrow \text{BestInputPropND}(x_{pop}, x_{best}, \mathcal{U}, T_{max})$ 
3 return  $\mathcal{E}_{new}$ 

```

**Algorithm 17: GABRR** $(x_{start}, x_{goal}, \mathcal{X}_{goal}, \mathcal{U}, T_{max}, \delta_{hr}, \mathcal{P})$ 

```

1  $\mathcal{G}_{for} \leftarrow \{\mathbf{V}_{for} \leftarrow \{x_{start}\}, \mathbf{E}_{for} \leftarrow \{\}\}$ 
2  $\mathcal{G}_{rev}^{ND} \leftarrow \{\mathbf{V}_{rev} \leftarrow \{x_{goal}\}, \mathbf{E}_{rev}^{ND} \leftarrow \{\}\}$ 
3  $\mathbf{Q} \leftarrow \{\}$ 
4 for  $k \leftarrow 1$  to  $M_{iter}$  do
   // Reverse Tree Expansion (lines 12-30)
5    $r_k \leftarrow \min(\gamma(\frac{\log |\mathbf{V}_{rev}|}{|\mathbf{V}_{rev}|})^{\frac{1}{d+1}}, \delta_{hr})$ 
6    $\mathcal{E}_{rev}^{ND} \leftarrow \text{RevSrchFastExploreND}(\mathcal{G}_{rev}^{ND}, \mathcal{U}, T_{max})$ 
7   if not CollisionCheck( $\mathcal{E}_{rev}^{ND}$ ) then
8      $x_{rev} \leftarrow \mathcal{E}_{rev}^{ND}.initialNode()$ 
9      $\mathbf{E}_{rev}^{ND} \leftarrow \mathbf{E}_{rev}^{ND} \cup \{\mathcal{E}_{rev}^{ND}\}$ 
10     $\mathbf{V}_{rev} \leftarrow \mathbf{V}_{rev} \cup \{x_{rev}\}$ 
11    UpdatePQueueND( $\mathcal{G}_{for}, \mathbf{Q}, x_{rev}, r_k$ )
   // Forward Tree Expansion (lines 10-28)
12    $q \leftarrow \mathcal{P}(k)$ 
13    $c_{rand} \sim U([0, 1])$ 
14   if  $c_{rand} < q$  then
15      $\mathcal{E}_{for} = \text{NULL}$ 
16     // Pop node from  $\mathbf{Q}$ 
17      $x_{pop} \leftarrow \text{Pop}(\mathbf{Q})$ 
18     if  $x_{pop} \neq \text{NULL}$  then
19        $\mathcal{E}_{for} \leftarrow \text{ForSrchExploitND}(\mathcal{G}_{rev}^{ND}, \mathcal{U}, T_{max}, x_{pop}, r_k)$ 
20       if  $\mathcal{E}_{for} = \text{NULL}$  then
21          $\mathcal{E}_{for} \leftarrow \text{ForSrchFastExplore}(\mathcal{G}_{for}, \mathcal{U}, T_{max})$ 
22   if  $c_{rand} \geq q$  or  $\mathcal{E}_{for} = \text{NULL}$  then
23     // For Probabilistic completeness
24      $\mathcal{E}_{for} \leftarrow \text{ForSrchRandomExplore}(\mathcal{G}_{for}, \mathcal{U}, T_{max})$ 
25   if  $\mathcal{E}_{for} \neq \text{NULL}$  then
26     if not CollisionCheck( $\mathcal{E}_{for}$ ) then
27        $x_{for} \leftarrow \mathcal{E}_{for}.finalNode()$ 
28        $\mathbf{E}_{for} \leftarrow \mathbf{E}_{for} \cup \{\mathcal{E}_{for}\}$ 
29        $\mathbf{V}_{for} \leftarrow \mathbf{V}_{for} \cup \{x_{for}\}$ 
30       if goalRegionReached( $\mathcal{X}_{goal}, x_{for}$ ) then
31          $\text{return Path}(\mathcal{G}_{for}, x_{start}, x_{for})$ 
32       insertToPQueueND( $\mathcal{G}_{rev}^{ND}, \mathbf{Q}, x_{for}, r_k$ )
33 return NULL

```

where  $\epsilon$  is a user-defined value. ForSrchFastExplore (Algorithm 5) and ForSrchRandomExplore (Algorithm 7) appear in Section VI.

### E. Computational Complexity and Shrinking $d$ -Ball Rate

In this section, we derive a worst case runtime bound and then discuss how the function used to shrink the neighborhood  $d$ -Ball affects the expected runtime of iteration  $i$ . GBRRT and GABRR require two data structures. A dynamic space partitioning tree is used to perform dynamic insertion (DI), nearest neighbor (NN), and range queries (RS), while a priority heap stores nodes based on priority key values and supports Push, Pop, and Update operations

**Algorithm 12: MonteCarloProp** $(x_{final}, \mathcal{U}, T_{max})$ 

```

// Monte Carlo propagation for
reverse search
1  $t_{\mathcal{E}} \leftarrow \text{Random}(0, T_{max})$ 
2  $u \leftarrow \text{Random}(\mathcal{U})$ 
3  $\mathcal{E}_{new} = - \int_{t_{\mathcal{E}}}^0 f(x, u) dt$  where  $\mathcal{E}_{new}(t_{\mathcal{E}}) = x_{final}$ 
4 return  $\mathcal{E}_{new}$ 

```

**Algorithm 13: insertToPQueueND** $(\mathcal{G}_{rev}^{ND}, \mathbf{Q}, x_{for}, r_k)$ 

```

1  $x_{closest} \leftarrow \mathcal{G}_{rev}^{ND}.NearestNeighbor(x_{for})$ 
2 if  $x_{closest} \leq r_k$  then
3    $\mathbf{Q}.insert(x_{for}, d_X^{ND}(x_{for}, x_{closest}) + h(x_{closest}))$ 

```

**Algorithm 14: RevSrchFastExploreND** $(\mathcal{G}_{rev}^{ND}, \mathcal{U}, T_{max})$ 

```

1  $x_{rand} \leftarrow \text{RandomState}()$ 
2  $x_{near} \leftarrow \text{NearestNeighbor}(x_{rand}, \mathcal{G}_{rev}^{ND})$ 
3  $\mathcal{E}_{new}^{ND} \leftarrow \text{Extend}(x_{near}, x_{rand})$ 
4 return  $\mathcal{E}_{new}^{ND}$ 

```

**Algorithm 15: BestInputPropND** $(x_{near}, x_{rand}, \mathcal{U}, T_{max})$ 

```

// Fwd. search for Best-input prop.
1  $\mathbf{E}_{fin} \leftarrow \{\}$ 
2 for  $k \leftarrow 1$  to  $N_B$  do
3    $\mathbf{E}_{fin} \leftarrow$ 
    $\mathbf{E}_{fin} \cup \{\text{MonteCarloProp}(x_{near}, \mathcal{U}, T_{max})\}$ 
4  $\mathcal{E}_{new} \leftarrow \arg \min_{\{\mathcal{E} \in \mathbf{E}_{fin}\}} (d_X^{ND}(\mathcal{E}.finalNode(), x_{rand}))$ 
5 return  $\mathcal{E}_{new}$ 

```

[52], [53]. Our runtime analysis assume these subroutines have the following runtime bounds, where  $N_T$  is the number of points in the tree.

**Assumption 1.** *DI takes expected time  $\mathcal{O}(\log N_T)$ .*

**Assumption 2.** *NN takes expected time  $\mathcal{O}(\log N_T)$ .*

These bounds are met by many algorithms [53]. GBRRT does not require exact NN and RS queries and, similar to other sampling based motion planning algorithms [3], can also work with ‘approximate’ solutions to these operations [54]–[56]. For example, the Skip Quadtree [56] supports a special kind of approximate RS query in which all points within range  $r_k$  are returned plus *possibly some additional points* that are further than  $r_k$  but within  $r_k(1 + \epsilon)$  for a user defined  $\epsilon$ . The Skip Quadtree is dynamic, allowing for both new point insertions and look-ups in expected time  $\log N_T$ . The approximate range query returns  $\nu$  points in expected time  $\mathcal{O}(\epsilon^{1-d} \log(N_T) + \nu)$ . The Skip Quadtree also meets Assumptions 1 and 2.

**Assumption 3.** *RS query (or approximate RS query) in dynamic space partitioning tree takes  $\mathcal{O}(c_d \log N_T + \nu)$  expected runtime and returns  $\nu$  neighbors in expectation, where  $c_d$  is a (possibly user chosen) parameter that may depend on  $d$  but is otherwise constant.*

While various implementations of priority heaps exist [53], many perform push and pop operations in  $\mathcal{O}(\log N_Q)$  runtime, where  $N_Q$  is the number of points in the queue [53].

**Assumption 4.** *Push, Pop, and Update operations of the priority queue each take time  $\mathcal{O}(\log N_Q)$ .*

A trivial worst case runtime for GBRRT and GABRR is given below in Theorem 2 for the case when Assumptions 1-4 hold. We denote  $|\mathbf{V}_{for}^i|$  and  $|\mathbf{V}_{rev}^i|$  as the number of forward and reverse tree nodes at the start of iteration  $i$ , respectively.  $n$  is the total number of nodes stored in both trees at the start of iteration  $i$ . Note that  $n = |\mathbf{V}_{for}^i| + |\mathbf{V}_{rev}^i| \leq 2i$ .

**Theorem 2.** *The worst case runtime of GBRRT and GABRR is  $\mathcal{O}(n^2)$ .*

*Proof.* In a (worse than) worst case  $\delta_{hr} = \gamma = \infty$  such that  $r_k = \infty$  and every query for forward (resp. reverse) tree neighbors returns the entire forward (resp. reverse) tree, which contains  $|\mathbf{V}_{\text{for}}^i|$  (resp.  $|\mathbf{V}_{\text{rev}}^i|$ ) nodes at iteration  $i$ . An upper bound on the time required per iteration is obtained by observing  $\nu \leq |\mathbf{V}_{\text{for}}^i| + |\mathbf{V}_{\text{rev}}^i| = n$ . Summing over iterations  $i = 1, \dots, n$  leads to a worst case bound of  $\mathcal{O}(n^2)$ .  $\square$

**Remark 1.** *The worst case  $\mathcal{O}(n^2)$  time is identical to that of many other sampling based motion planning algorithms. Such a worst case is typically only realized in degenerative cases, and average runtimes tend to be much lower in practice.*

Before deriving *expected* per iteration runtime bounds, it is useful to discuss what we do *not* know about the distribution of nodes stored within the forward and backward trees.

Expected runtime analysis of sampling based motion planning algorithms often focuses on the geometric case (without system dynamics) where the set of potential neighbors already in the tree can be assumed to be generated from a uniform random distribution over the free space, at least in the limiting case as  $n \rightarrow \infty$ . In contrast, the nodes stored in the forward and reverse trees of GBRRT and the forward tree of GABRRT are selected using a combination of randomness, forward or backward integration of system dynamics (or, alternatively, maneuver libraries). Moreover, a portion of all nodes' position in the forward tree are influenced by the reverse tree focusing heuristic. Whether or not the nodes within either tree resemble a set of points drawn uniformly at random and i.i.d. is, *at best*, system/scenario dependent. It remains an open question whether or not the sampling distributions realized by GBRRT and GABRRT converge toward the uniform distribution (even for the special case of geometric planning).

When calculating expected runtime, *if* points can be assumed to be drawn i.i.d from a uniform distribution, *then* many existing proofs [3], [57] take advantage of the fact that the expected number of nodes  $\mathbb{E}(|V_{\text{near}}|)$  in a  $d$ -ball is proportional to the ball's volume  $V_d(r)$ . Given  $N_T$  nodes drawn uniformly random and i.i.d. from  $X_{\text{free}}$  then  $\mathbb{E}(|V_{\text{near},r}|) \sim N_T \frac{V_d(r)}{\|X_{\text{free}}\|}$ , where  $\|X_{\text{free}}\|$  is the volume of the free space, and where the free space is assumed to have "nice" boundaries, e.g., locally Lipschitz continuous boundaries for some positive finite Lipschitz constant. In a static environment,  $\|X_{\text{free}}\|$  is constant. This discussion leads to the following Proposition 3.

**Proposition 3.** *If points are sampled uniformly at random and i.i.d. from  $X_{\text{free}}$ , then  $\mathbb{E}(|V_{\text{near},r}|) = \Theta(N_T V_d(r))$ .*

**Lemma 2.** *Given range  $r$  and  $N_T$  points sampled uniformly at random i.i.d., and assuming  $\varepsilon$  and  $d$  are both constant, the expected value of  $\nu$  for the approximate RS query described above is  $\mathbb{E}(\nu(r, N_T)) = \Theta(N_T V_d(r))$ .*

*Proof.* It is a useful fact that the ratio between the volume of a  $d$ -ball of radius  $r(1 + \varepsilon)$  and the volume of a  $d$ -ball of

radius  $r$  is:  $\left(\frac{\pi^{d/2}}{\Gamma(1+d/2)}(r_k(1 + \varepsilon))^d\right) / \left(\frac{\pi^{d/2}}{\Gamma(1+d/2)}r_k^d\right) = (1 + \varepsilon)^d$ , where  $\Gamma(\cdot)$  is the gamma function. The volume ratio  $(1 + \varepsilon)^d$  does not depend on  $r_k$  and, once  $d$  is chosen, is a constant.  $\square$

If the point sampling process *cannot* be modeled as uniformly at random and i.i.d., then it is possible that  $\mathbb{E}(|V_{\text{near},r}|) = \Omega(N_T V_d(r))$ . For example, if new points are more likely to be sampled in the vicinity of existing points such that, in expectation, greater than half of all points have greater than  $N_T \frac{V_d(r)}{\|X_{\text{free}}\|}$  neighbors. Since the value of  $\mathbb{E}(|V_{\text{near},r}|)$  achieved in GBRRT/GABRRT is an open question, we shall derive lower bounds on runtime for the  $\mathbb{E}(|V_{\text{near},r}|) = \Omega(N_T V_d(r))$  case. How  $\mathbb{E}(|V_{\text{near},r}|) = \Omega(N_T V_d(r))$  may affect  $\nu$  is important for our analysis.

**Lemma 3.** *If  $\mathbb{E}(|V_{\text{near},r}|) = \Omega(N_T V_d(r))$  then  $\mathbb{E}(\nu(r, N_T)) = \Omega(N_T V_d(r))$ .*

*Proof.* By definition, using Bachmann–Landau notation,  $\mathbb{E}(|V_{\text{near},r}|) = \Omega(N_T V_d(r)) \iff \liminf_{N_T \rightarrow \infty} \frac{\mathbb{E}(|V_{\text{near},r}|)}{N_T V_d(r)} > 0$ .

By construction of the approximate range search (RS) query  $\nu(r, N_T) \geq |V_{\text{near},r}|$  and so  $\mathbb{E}(\nu(r, N_T)) \geq \mathbb{E}(|V_{\text{near},r}|)$ . Thus,  $\liminf_{N_T \rightarrow \infty} \frac{\mathbb{E}(\nu(r, N_T))}{N_T V_d(r)} > 0$  and so by definition, using Bachmann–Landau notation,  $\mathbb{E}(\nu(r, N_T)) = \Omega(N_T V_d(r))$ .  $\square$

We now characterize the expected runtime required for the forward and reverse search portions of GBRRT and GABRRT in iteration  $i$ . Note  $|\mathbf{Q}| \leq |\mathbf{V}_{\text{for}}^i|$  at iteration  $i$ , where  $|\mathbf{Q}|$  is the numbers of nodes in the priority queue  $\mathbf{Q}$ .

**Lemma 4.** *Given Assumptions 1-4, if  $\mathbb{E}(\nu(r, |\mathbf{V}_{\text{rev}}^i|)) = \mathcal{O}(\log |\mathbf{V}_{\text{rev}}^i|)$ , then the expected runtime of forward tree operations during iteration  $i$  is  $\mathcal{O}(\log |\mathbf{V}_{\text{for}}^i| + \log |\mathbf{V}_{\text{rev}}^i|)$ ; and alternatively, if  $\mathbb{E}(\nu(r, |\mathbf{V}_{\text{rev}}^i|)) = \Omega(\log |\mathbf{V}_{\text{rev}}^i|)$ , then the expected runtime of forward tree operations during iteration  $i$  is  $\Theta(\nu(r, |\mathbf{V}_{\text{rev}}^i|))$ .*

*Proof.* Forward search involves a NN query and a DI operation on the forward tree and push/pop of a forward tree node from  $\mathbf{Q}$ ; all take expected time  $\mathcal{O}(\log |\mathbf{V}_{\text{for}}^i|)$ . It also involves a NN query and RS query on the reverse tree. For constant  $d$ , these operations take expected time  $\mathcal{O}(\log |\mathbf{V}_{\text{rev}}^i|)$  and  $\mathcal{O}(\log |\mathbf{V}_{\text{rev}}^i| + \nu(r, |\mathbf{V}_{\text{rev}}^i|))$ , respectively. The expected time involved in processing the  $\mathbb{E}(\nu(r, |\mathbf{V}_{\text{rev}}^i|))$  nodes returned by the RS query is  $\Theta(\nu(r, |\mathbf{V}_{\text{rev}}^i|))$ . The proof is finished by considering the case where  $\mathbb{E}(\nu(r, |\mathbf{V}_{\text{rev}}^i|)) = \mathcal{O}(\log |\mathbf{V}_{\text{rev}}^i|)$ ; and alternatively, the case where  $\mathbb{E}(\nu(r, |\mathbf{V}_{\text{rev}}^i|)) = \Omega(\log |\mathbf{V}_{\text{rev}}^i|)$ .  $\square$

**Lemma 5.** *Given Assumptions 1-4, the expected runtime of reverse search operations during iteration  $i$  is  $\mathcal{O}(\log |\mathbf{V}_{\text{for}}^i| + \log |\mathbf{V}_{\text{rev}}^i|)$ .*

*Proof.* The computationally dominant operations performed during reverse search are NN query and insertion in reverse tree which both take expected time  $\mathcal{O}(\log |\mathbf{V}_{\text{rev}}^i|)$  and NN query in forward tree and update operation in  $\mathbf{Q}$  which both take expected time  $\mathcal{O}(\log |\mathbf{V}_{\text{for}}^i|)$ .  $\square$

**Theorem 3.** Given Assumptions 1-4, if  $\mathbb{E}(\nu(r, n)) = \mathcal{O}(\log n)$ , then the expected runtime of iteration  $i$  is  $\mathcal{O}(\log n)$ , alternatively, if  $\mathbb{E}(\nu(r, n)) = \Omega(\log n)$ , then the expected runtime of iteration  $i$  is  $\Theta(\nu(r, n))$ .

*Proof.* Combine Lemma 4 and Lemma 5 with the fact that  $2 \leq |\mathbf{V}_{\text{for}}^i| + |\mathbf{V}_{\text{rev}}^i| = n \leq 2i$ , and we alternate between attempting to grow the forward and reverse trees, so we expect both trees to have a nonzero proportion of the  $n$  nodes in the limit, i.e.,  $\lim_{i \rightarrow \infty} \frac{|\mathbf{V}_{\text{for}}^i|}{|\mathbf{V}_{\text{rev}}^i| + |\mathbf{V}_{\text{for}}^i|} = c$  such that  $0 < c < 1$ . Because neither tree's cardinality is expected to overwhelm the other, the most conservative runtime associated with either tree determines the runtime of the overall algorithm.  $\square$

Theorem 3 shows that  $\mathbb{E}(\nu(r, n))$  plays a key role in determining the expected runtime of iteration  $i$  in GBRRT and GABRRT.  $\mathbb{E}(\nu(r, n))$  is affected by the radius  $r$  of the  $d$ -ball as well as the sampling distribution of the  $n$  points in  $X_{\text{free}}$ . As already mentioned, the sampling distribution that results when running GBRRT or GABRRT is an open question. We shall now discuss the radius of the ball, looking at three cases of interest — using a constant ball, and shrinking at the ball according to two functions that have appeared in the literature. We consider the uniform i.i.d. case followed by the case where  $\mathbb{E}(\nu(r, N_T)) = \Omega(N_T V_d(r))$ .

**Theorem 4.** Given Assumptions 1-4, and assuming  $d$  is constant, nodes are drawn uniformly at random and i.i.d., and  $r_k = c$  for a positive constant  $c < \infty$ , the expected runtime of GBRRT/GABRRT iteration  $i$  is  $\Theta(n)$ .

*Proof.*  $\mathbb{E}(\nu(r, n)) = \Theta(nV_d(r))$  by Proposition 3 and Lemma 2. By construction,  $V_d(r)$  is constant for constant  $r$  and so  $\mathbb{E}(\nu(r, n)) = \Theta(n)$ . By definition  $n = \Omega(\log n)$ . Thus, by Theorem 3, the expected runtime of iteration  $i$  is  $\Theta(\nu(r, n))$ . Because  $\mathbb{E}(\nu(r, n)) = \Theta(n)$  when  $r_k$  is constant, the expected runtime of iteration  $i$  is  $\Theta(n)$ .  $\square$

**Corollary 2.** Given Assumptions 1-4 and constant  $d$ , and assuming nodes are sampled such that  $\mathbb{E}(|V_{\text{near}, r}|) = \Omega(N_T V_d(r))$ , and assuming  $r_k = c$  for some positive constant  $c < \infty$ , then the expected runtime of GBRRT/GABRRT iteration  $i$  is  $\Omega(n)$ .

Corollary 2 results from combining Theorem 4 with Lemma 3.

Many sampling based motion planning algorithms shrink the neighborhood  $d$ -ball to reduce expected runtime. For example, RRT\* [3] shrinks the neighborhood ball according to  $r_k = \gamma \left(\frac{\log n}{n}\right)^{1/d}$  to achieve an expected per iteration runtime of  $\mathcal{O}(\log n)$ . We analyze the effects of using this strategy for GBRRT and GABRRT. While  $r_k = \min(\gamma \left(\frac{\log n}{n}\right)^{1/d}, \delta_{hr})$ , the use of  $\delta_{hr}$  ceases after a finite number of iterations because  $\left(\frac{\log n}{n}\right)^{1/d} = \mathcal{O}(\delta_{hr})$ . So, our asymptotic analysis focuses on the  $r_k = \gamma \left(\frac{\log n}{n}\right)^{1/d}$  case.

**Theorem 5.** Given Assumptions 1-4 and constant  $d$ , and assuming nodes are drawn uniformly at random and i.i.d.

and assuming  $r_k = \gamma \left(\frac{\log n}{n}\right)^{1/d}$  then the expected runtime of GBRRT/GABRRT iteration  $i$  is  $\Theta(\log n)$ .

*Proof.* Combining Proposition 3 with Lemma 2  $\mathbb{E}(\nu(r, n)) = \Theta(nV_d(r))$ . Given constant  $d$ , by construction  $nV_d(r) = n \left(\frac{\pi^{d/2}}{\Gamma(1+d/2)} \left(\gamma \left(\frac{\log n}{n}\right)^{1/d}\right)^d\right) = \Theta(\log n)$ , where  $\gamma$  is constant. Theorem 3 completes the proof.  $\square$

**Corollary 3.** Given Assumptions 1-4 and constant  $d$ , and assuming nodes are sampled such that  $\mathbb{E}(|V_{\text{near}, r}|) = \Omega(N_T V_d(r))$ , and assuming  $r_k = \gamma \left(\frac{\log n}{n}\right)^{1/d}$  then the expected runtime of GBRRT/GABRRT iteration  $i$  is  $\Omega(\log n)$ .

Corollary 3 results from combining Theorem 5 with Lemma 3.

A recent result from Solovey et al. [33] has highlighted a small error in Karaman and Frazzoli's proof of almost sure asymptotic optimality for RRT\* [3]. (Note: the proof of probabilistic completeness in [3] is still valid.) The original proof from [3] has motivated the use of  $r_k = \gamma \left(\frac{\log n}{n}\right)^{1/d}$  in a number of algorithms. Solovey et al. show that using a larger  $r_k = \gamma \left(\frac{\log n}{n}\right)^{1/(d+1)}$  is sufficient to achieve almost sure asymptotic optimality. Whether or not using  $r_k = \gamma \left(\frac{\log n}{n}\right)^{1/d}$  is also sufficient is now an open problem [33]. The main difference between [33] and [3] is a consideration of how sampling order affects local rewiring direction. While GBRRT and GABRRT are feasible (and not asymptotically optimal) planning algorithms, the reverse tree focusing heuristic is concerned with local wiring direction. While further study of the relationship between asymptotic optimality and focusing heuristics is outside the scope of this paper, we both (i) use  $r_k = \gamma \left(\frac{\log n}{n}\right)^{1/(d+1)}$  in our experiments and (ii) acknowledge that future results may determine that  $r_k \leq \gamma \left(\frac{\log n}{n}\right)^{1/d}$  is sufficient for focusing heuristic purposes in GBRRT and/or GABRRT. We now derive results for  $r_k = \gamma \left(\frac{\log n}{n}\right)^{1/(d+1)}$ .

**Theorem 6.** Given Assumptions 1-4 and constant  $d$ , and assuming nodes are drawn uniformly at random and i.i.d., and assuming  $r_k = \gamma \left(\frac{\log n}{n}\right)^{1/(d+1)}$ , then the expected runtime of GBRRT/GABRRT iteration  $i$  is  $\Theta((\log n)^{d/(d+1)} a^{+1} \sqrt{n})$ .

*Proof.* Combining Proposition 3 with Lemma 2  $\mathbb{E}(\nu(r, n)) = \Theta(nV_d(r))$ . Given constant  $d$ , by construction  $nV_d(r) = n \left(\frac{\pi^{d/2}}{\Gamma(1+d/2)} \left(\gamma \left(\frac{\log n}{n}\right)^{1/(d+1)}\right)^d\right) = \Theta((\log n)^{d/(d+1)} a^{+1} \sqrt{n})$ , where  $\gamma$  is a constant. Applying Theorem 3 completes the proof.  $\square$

**Corollary 4.** Given Assumptions 1-4 and constant  $d$ , and assuming nodes are sampled such that  $\mathbb{E}(|V_{\text{near}, r}|) = \Omega(N_T V_d(r))$ , and assuming  $r_k = \gamma \left(\frac{\log n}{n}\right)^{1/(d+1)}$ , then the expected runtime of GBRRT/GABRRT iteration  $i$  is  $\Omega((\log n)^{d/(d+1)} a^{+1} \sqrt{n})$ .

Corollary 4 results from combining Lemma 5 with Lemma 3.

## REFERENCES

- [1] S. M. LaValle, "Rapidly-exploring random trees: A new tool for path planning," University of Illinois at Urbana-Champaign, Tech. Rep., 1998.
- [2] H. Liu, X. Zhang, J. Wen, R. Wang, and X. Chen, "Goal-biased bidirectional rrt based on curve-smoothing," *IFAC-PapersOnLine*, vol. 52, no. 24, pp. 255–260, 2019.
- [3] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The international journal of robotics research*, vol. 30, no. 7, pp. 846–894, 2011.
- [4] K. Hauser and Y. Zhou, "Asymptotically optimal planning by feasible kinodynamic planning in a state–cost space," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1431–1443, 2016.
- [5] S. M. LaValle, *Planning algorithms*. Cambridge university press, 2006.
- [6] J. J. Kuffner and S. M. LaValle, "Rrt-connect: An efficient approach to single-query path planning," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, vol. 2. IEEE, 2000, pp. 995–1001.
- [7] L. Zhang, X. Huang, Y. J. Kim, and D. Manocha, "D-plan: Efficient collision-free path computation for part removal and disassembly," *Computer-Aided Design and Applications*, vol. 5, no. 6, pp. 774–786, 2008.
- [8] M. Jordan and A. Perez, "Optimal bidirectional rapidly-exploring random trees," 2013.
- [9] D. Devaurs, T. Siméon, and J. Cortés, "Enhancing the transition-based rrt to deal with complex cost spaces," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 4120–4125.
- [10] J. Starek, E. Schmerling, L. Janson, and M. Pavone, "Bidirectional fast marching trees: An optimal sampling-based algorithm for bidirectional motion planning," in *Workshop on Algorithmic Foundations of Robotics*, 2014.
- [11] H. B. Keller, *Numerical methods for two-point boundary-value problems*. Courier Dover Publications, 2018.
- [12] M. R. Osborne, "On shooting methods for boundary value problems," *Journal of mathematical analysis and applications*, vol. 27, no. 2, pp. 417–433, 1969.
- [13] S. M. LaValle and J. J. Kuffner Jr, "Randomized kinodynamic planning," *The international journal of robotics research*, vol. 20, no. 5, pp. 378–400, 2001.
- [14] M. S. Branicky, R. A. Knepper, and J. J. Kuffner, "Path and trajectory diversity: Theory and algorithms," in *2008 IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 1359–1364.
- [15] J. Go, T. D. Vu, and J. J. Kuffner, "Autonomous behaviors for interactive vehicle animations," *Graphical Models*, vol. 68, no. 2, pp. 90–112, 2006.
- [16] Z. Tahir, A. H. Qureshi, Y. Ayaz, and R. Nawaz, "Potentially guided bidirectionalized rrt\* for fast optimal path planning in cluttered environments," *Robotics and Autonomous Systems*, vol. 108, pp. 13–27, 2018.
- [17] W. Xinyu, L. Xiaojuan, G. Yong, S. Jiadong, and W. Rui, "Bidirectional potential guided rrt\* for motion planning," *IEEE Access*, vol. 7, pp. 95 034–95 045, 2019.
- [18] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," in *Autonomous robot vehicles*. Springer, 1986, pp. 396–404.
- [19] M. P. Strub and J. D. Gammell, "Adaptively informed trees (ait\*): Fast asymptotically optimal path planning through adaptive heuristics," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 3191–3198.
- [20] J. hwan Jeon, S. Karaman, and E. Frazzoli, "Anytime computation of time-optimal off-road vehicle maneuvers using the rrt," in *2011 50th IEEE Conference on Decision and Control and European Control Conference*. IEEE, 2011, pp. 3276–3282.
- [21] D. J. Webb and J. Van Den Berg, "Kinodynamic rrt\*: Asymptotically optimal motion planning for robots with linear dynamics," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 5054–5061.
- [22] Y. Li, Z. Littlefield, and K. E. Bekris, "Asymptotically optimal sampling-based kinodynamic planning," *The International Journal of Robotics Research*, vol. 35, no. 5, pp. 528–564, 2016.
- [23] Z. Littlefield and K. E. Bekris, "Informed asymptotically near-optimal planning for field robots with dynamics," in *Field and Service Robotics*. Springer, 2018, pp. 449–463.
- [24] Z. Littlefield and K. Bekris, "Efficient and asymptotically optimal kinodynamic motion planning via dominance-informed regions," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1–9.
- [25] M. Kleinbort, E. Granados, K. Solovey, R. Bonalli, K. E. Bekris, and D. Halperin, "Refined analysis of asymptotically-optimal kinodynamic planning in the state-cost space," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 6344–6350.
- [26] R. Shome and L. E. Kavraki, "Asymptotically optimal kinodynamic planning using bundles of edges," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 9988–9994.
- [27] B. Akgun and M. Stilman, "Sampling heuristics for optimal motion planning in high dimensions," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 2640–2645.
- [28] C. Urmson and R. Simmons, "Approaches for heuristically biasing rrt growth," in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, vol. 2. IEEE, 2003, pp. 1178–1183.
- [29] D. Ferguson and A. Stentz, "Anytime rrts," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2006, pp. 5369–5375.
- [30] M. Otte and N. Correll, "C-forest: Parallel shortest path planning with superlinear speedup," *IEEE Transactions on Robotics*, vol. 29, no. 3, pp. 798–806, 2013.
- [31] S. Nayak and M. W. Otte, "Bidirectional sampling based search without two point boundary value solution," <https://arxiv.org/abs/2010.14692>, 2020.
- [32] M. Kleinbort, K. Solovey, Z. Littlefield, K. E. Bekris, and D. Halperin, "Probabilistic completeness of rrt for geometric and kinodynamic planning with forward propagation," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. x–xvi, 2018.
- [33] K. Solovey, L. Janson, E. Schmerling, E. Frazzoli, and M. Pavone, "Revisiting the asymptotic optimality of rrt," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2189–2195.
- [34] T. Kunz and M. Stilman, "Kinodynamic rrts with fixed time step and best-input extension are not probabilistically complete," in *Algorithmic foundations of robotics XI*. Springer, 2015, pp. 233–244.
- [35] L.-W. Cai, *Fundamentals of mechanical vibrations*. John Wiley & Sons, 2016.
- [36] J. J. Kuffner, "Effective sampling and distance metrics for 3d rigid body path planning," in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, vol. 4. IEEE, 2004, pp. 3993–3998.
- [37] J. D. Lambert *et al.*, *Numerical methods for ordinary differential systems*. Wiley New York, 1991, vol. 146.
- [38] R. Carona, A. P. Aguiar, and J. Gaspar, "Control of unicycle type robots tracking, path following and point stabilization," 2008.
- [39] I. Palunko and R. Fierro, "Adaptive control of a quadrotor with dynamic changes in the center of gravity," *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 2626–2631, 2011.
- [40] G. Papadopoulos, H. Kurniawati, and N. M. Patrikalakis, "Analysis of asymptotically optimal sampling-based motion planning algorithms for lipschitz continuous dynamical systems," *arXiv preprint arXiv:1405.2872*, 2014.
- [41] J. Pentzer, S. Brennan, and K. Reichard, "Model-based prediction of skid-steer robot kinematics using online estimation of track instantaneous centers of rotation," *Journal of Field Robotics*, vol. 31, no. 3, pp. 455–476, 2014.
- [42] A. Sivaramakrishnan, Z. Littlefield, and K. E. Bekris, "Towards learning efficient maneuver sets for kinodynamic motion planning," *arXiv preprint arXiv:1907.07876*, 2019.
- [43] R. M. Murray and S. S. Sastry, "Nonholonomic motion planning: Steering using sinusoids," *IEEE transactions on Automatic Control*, vol. 38, no. 5, pp. 700–716, 1993.
- [44] A. A. Paranjape, K. C. Meier, X. Shi, S.-J. Chung, and S. Hutchinson, "Motion primitives and 3d path planning for fast flight through a forest," *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 357–377, 2015.

- [45] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [46] J. Kierzenka and L. F. Shampine, "A bvp solver based on residual control and the matlab pse," *ACM Transactions on Mathematical Software (TOMS)*, vol. 27, no. 3, pp. 299–316, 2001.
- [47] N. J. Sanket, C. D. Singh, C. Fermüller, and Y. Aloimonos, "Prgfly: Ai based indoor quadrotor autonomy framework for navigation and interaction tasks," <https://github.com/prgumd/PRGFlyt/wiki>, 2019.
- [48] M. Köppen, "The curse of dimensionality," in *5th Online World Conference on Soft Computing in Industrial Applications (WSC5)*, vol. 1, 2000, pp. 4–8.
- [49] S. Boslaugh, *Statistics in a nutshell: A desktop quick reference.* " O'Reilly Media, Inc.", 2012.
- [50] F. J. Massey Jr, "The kolmogorov-smirnov test for goodness of fit," *Journal of the American statistical Association*, vol. 46, no. 253, pp. 68–78, 1951.
- [51] C. D. Twigg and D. L. James, "Backward steps in rigid body simulation," in *ACM SIGGRAPH 2008 papers*, 2008, pp. 1–10.
- [52] H. Samet, *The design and analysis of spatial data structures.* Addison-Wesley Reading, MA, 1990, vol. 85.
- [53] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to algorithms.* MIT press, 2009.
- [54] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Y. Wu, "An optimal algorithm for approximate nearest neighbor searching fixed dimensions," *Journal of the ACM (JACM)*, vol. 45, no. 6, pp. 891–923, 1998.
- [55] S. Arya and D. M. Mount, "Approximate range searching," *Computational Geometry*, vol. 17, no. 3-4, pp. 135–152, 2000.
- [56] D. Eppstein, M. T. Goodrich, and J. Z. Sun, "The skip quadtree: a simple dynamic data structure for multidimensional data," in *Proceedings of the twenty-first annual symposium on Computational geometry*, 2005, pp. 296–305.
- [57] M. Otte and E. Frazzoli, "Rrtx: Asymptotically optimal single-query sampling-based motion planning with quick replanning," *The International Journal of Robotics Research*, vol. 35, no. 7, pp. 797–822, 2016.



**Sharan Nayak** received his Bachelor of Engineering (B.E.) degree in electronics and communications engineering from Visveswaraya Technological University, India, in 2009; his M.S. degree in electrical and computer engineering from Georgia Institute of Technology, Atlanta, GA, USA, in 2011 and M.S. degree in aerospace engineering from University of Maryland (UMD), College Park, MD in 2020. He is currently pursuing a Ph.D. degree in the department of Aerospace Engineering at UMD under the supervision of Dr. Michael Otte.

His research interests are in motion planning of single and multi-agent autonomous systems. He is a student member of IEEE and a recipient of the Clark Doctoral Fellowship.



**Michael Otte** (M'07) received the B.S. degrees in aeronautical engineering and computer science from Clarkson University, Potsdam, New York, USA, in 2005, and the M.S. and Ph.D. degrees in computer science from the University of Colorado Boulder, Boulder, CO, USA, in 2007 and 2011, respectively. From 2011 to 2014, he was a Postdoctoral Associate at the Massachusetts Institute of Technology. From 2014 to 2015, he was a Visiting Scholar at the U.S. Air Force Research Lab. From 2016 to 2018, he was a National Research Council

RAP Postdoctoral Associate at the U.S. Naval Research Lab. He has been with the Department of Aerospace Engineering, at the University of Maryland, College Park, MD, USA, since 2018. He is the author of over 30 articles, and his research interests include autonomous robotics, motion planning, and multi-agent systems.