An Asymptotically-Optimal Sampling-Based Algorithm for Bi-directional Motion Planning

Joseph A. Starek*, Javier V. Gomez[†], Edward Schmerling[‡], Lucas Janson[§], Luis Moreno[†], Marco Pavone*

Abstract-Bi-directional search is a widely used strategy to increase the success and convergence rates of sampling-based motion planning algorithms. Yet, few results are available that merge both bi-directional search and asymptotic optimality into existing optimal planners, such as PRM*, RRT*, and FMT*. The objective of this paper is to fill this gap. Specifically, this paper presents a bi-directional, sampling-based, asymptoticallyoptimal algorithm named Bi-directional FMT* (BFMT*) that extends the Fast Marching Tree (FMT*) algorithm to bidirectional search while preserving its key properties, chiefly lazy search and asymptotic optimality through convergence in probability. BFMT* performs a two-source, lazy dynamic programming recursion over a set of randomly-drawn samples, correspondingly generating two search trees: one in cost-tocome space from the initial configuration and another in costto-go space from the goal configuration. Numerical experiments illustrate the advantages of BFMT* over its unidirectional counterpart, as well as a number of other state-of-the-art planners.

I INTRODUCTION

Motion planning is the computation of paths that guide systems from an initial configuration to a set of goal configuration (s) around nearby obstacles, while possibly optimizing an objective function. The problem has a long and rich history in the field of robotics, and many algorithmic tools have been developed; we refer the interested reader to [1] and references therein. Arguably, sampling-based algorithms are among the most pervasive, widespread planners available in robotics, including the Probabilistic Roadmap algorithm (PRM) [2], the Expansive Space Trees algorithm (EST) [3], [4], and the Rapidly-Exploring Random Tree algorithm (RRT) [5]. Since their development, efforts to improve the "quality" of paths led to asymptotically-optimal (AO) variants of RRT and PRM, named RRT* and PRM*, respectively, whereby the cost of the returned solution converges almost surely to the optimum as the number of samples approaches infinity [6], [7]. Many other planners followed, including BIT^{*} [8] and RRT[#] [9] to name a few. Recently, a conceptually different asymptoticallyoptimal, sampling-based motion planning algorithm, called the Fast Marching Tree (FMT*) algorithm, has been presented in [10], [11]. Numerical experiments suggested that FMT* converges to an optimal solution faster than PRM* or RRT*. especially in high-dimensional configuration spaces and in scenarios where collision-checking is expensive.

It is a well-known fact that *bi-directional* search can dramatically increase the convergence rate of planning

[‡]Inst. for Computational & Mathematical Engineering, Stanford University, Stanford, CA 94305, schmrlng@stanford.edu.

algorithms, prompting some authors [12] to advocate its use for accelerating essentially any motion planning query. This was first rigorously studied in [13] and later investigated, for example, in [14], [15]. Collectively, the algorithms presented in [12]–[15] belong to the family of *non*-sampling-based approaches and are more or less closely related to a bidirectional implementation of the Dijkstra Method. More recently, and not surprisingly in light of these performance gains, bi-directional search has been merged with the sampling-based approach, with RRT-Connect and SBL representing the most notable examples [16], [17].

Though such bi-directional versions of RRT and PRM are probabilistically complete, they do not enjoy optimality guarantees. The next logical step in the quest for fast planning algorithms is the design of *bi-directional*, sampling-based, asymptotically-optimal algorithms. To the best of our knowledge, the only available results in this context are [18] and the unpublished work [19], both of which discuss bi-directional implementations of RRT*. Neither work, however, provides a mathematically-rigorous proof of asymptotic optimality starting from first principles. Accordingly, the objective of this paper is to propose and *rigorously* analyze such an algorithm.

Statement of Contributions: This paper introduces the Bidirectional Fast Marching Tree (BFMT*) algorithm.¹ To the best of the authors' knowledge, this is the first treebased, asymptotically-optimal bi-directional sampling-based planner. BFMT* extends FMT* to bi-directional search and essentially performs a "lazy," bi-directional dynamic programming recursion over a set of probabilistically-drawn samples in the free configuration space. The contribution of this paper is threefold. First, we present the BFMT* algorithm in Section III. Second, we rigorously prove the asymptotic optimality of BFMT* (under the notion of convergence in probability) and characterize its convergence rate in Section IV. We note that the convergence rate of FMT* in [11] is proved only for obstacle-free configuration spaces, while we generalize that result to allow for the presence of obstacles. Finally, we perform numerical experiments in Section V across a number of planning spaces that suggest BFMT* converges to an optimal solution at least as fast as FMT*, PRM*, and RRT*, and sometimes significantly faster.

II PROBLEM DEFINITION

Let \mathcal{X} be a *d*-dimensional configuration space, and let \mathcal{X}_{obs} be the obstacle region, such that $\mathcal{X} \setminus \mathcal{X}_{obs}$ is an open set (we consider $\partial \mathcal{X} \subset \mathcal{X}_{obs}$). Denote the obstacle-free space as $\mathcal{X}_{free} = cl(\mathcal{X} \setminus \mathcal{X}_{obs})$, where $cl(\cdot)$ denotes the closure of a set. A path planning problem, denoted

^{*}Dept. of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305, {jstarek, pavone}@stanford.edu.

[†]Dept. of Systems Engineering & Automation, Carlos III University of Madrid, Madrid, Spain, 28911, {jvgomez, moreno}@ing.uc3m.es.

[§]Dept. of Statistics, Stanford University, Stanford, CA 94305, ljanson@stanford.edu.

¹The asterisk *, pronounced "star", is intended to represent asymptotic optimality much like for the RRT* and PRM* algorithms.

by a triplet $(\mathcal{X}_{\text{free}}, \mathbf{x}_{\text{init}}, \mathbf{x}_{\text{goal}})$, seeks to maneuver from an initial configuration \mathbf{x}_{init} to a goal configuration \mathbf{x}_{goal} through $\mathcal{X}_{\text{free}}$. Let a continuous function of *bounded variation* $\sigma : [0, 1] \rightarrow \mathcal{X}$, called a *path*, be *collision-free* if $\sigma(\tau) \in \mathcal{X}_{\text{free}}$ for all $\tau \in [0, 1]$. A path is called a *feasible* solution to the planning problem $(\mathcal{X}_{\text{free}}, \mathbf{x}_{\text{init}}, \mathbf{x}_{\text{goal}})$ if it is collision-free, $\sigma(0) = \mathbf{x}_{\text{init}}$, and $\sigma(1) = \mathbf{x}_{\text{goal}}$.

Let Σ be the set of all paths. A cost function for the planning problem $(\mathcal{X}_{\text{free}}, \mathbf{x}_{\text{init}}, \mathbf{x}_{\text{goal}})$ is a function $J : \Sigma \to \mathbb{R}_{\geq 0}$ from Σ to the nonnegative real numbers; in this paper, we consider as $J(\sigma)$ the *arc length* of σ with respect to the Euclidean metric in \mathcal{X} (the extension to general cost functions will be briefly discussed in Section IV-C).

Optimal path planning problem: Given a path planning problem $(\mathcal{X}_{\text{free}}, \mathbf{x}_{\text{init}}, \mathbf{x}_{\text{goal}})$ and an arc length function $J : \Sigma \to \mathbb{R}_{\geq 0}$, find a feasible path σ^* such that $J(\sigma^*) = \min\{J(\sigma) \mid \sigma \text{ is feasible}\}$. If no such path exists, report failure.

Finally, we introduce some definitions concerning the *clearance* of a path, *i.e.*, its "distance" from \mathcal{X}_{obs} [11]. For a given $\delta > 0$, the δ -interior of \mathcal{X}_{free} is defined as the set of all points that are at least a distance δ away from any point in \mathcal{X}_{obs} . A collision-free path σ is said to have strong δ -clearance if it lies entirely inside the δ -interior of \mathcal{X}_{free} . A path planning problem with optimal path cost J^* is called δ -*robustly feasible* if there exists a strictly positive sequence $\delta_n \to 0$, with $\delta_n \leq \delta \quad \forall n \in \mathbb{N}$, and a sequence $\{\sigma_n\}_{n=1}^{\infty}$ of feasible paths such that $\lim_{n\to\infty} J(\sigma_n) = J^*$ and for all $n \in \mathbb{N}$, σ_n has strong δ_n -clearance, $\sigma_n(1) = \mathbf{x}_{goal}$, $\sigma_n(\tau) \neq \mathbf{x}_{goal}$ for all $\tau \in (0, 1)$, and $\sigma_n(0) = \mathbf{x}_{init}$.

III THE BFMT* ALGORITHM

In this section, we present the Bi-Directional Fast Marching Tree algorithm, BFMT*, represented in pseudocode as Algorithm 1. To begin, we provide a high-level description of FMT* in Section III-A, on which BFMT* is based. We follow in Section III-B with BFMT*'s own high-level description, and then provide additional details in Section III-C.

III-A FMT* – High-level description

The FMT* algorithm, introduced in [10], [11], is a unidirectional algorithm that essentially performs a forward dynamic programming recursion over a set of sampled points and correspondingly generates a *tree of paths* that grow steadily outward in cost-to-come space. The recursion performed by FMT* is characterized by three key features: (1) It is tailored to disk-connected graphs, where two samples are considered neighbors (hence connectable) if their distance is below a given bound, referred to as the connection radius; (2) It performs graph construction and graph search concurrently; and (3) For the evaluation of the immediate cost in the dynamic programming recursion, one "lazily" ignores the presence of obstacles, and whenever a locally-optimal (assuming no obstacles) connection to a new sample intersects an obstacle, that sample is simply skipped and left for later (as opposed to looking for other locally-optimal connections in the neighborhood).

The last feature, which makes the algorithm "lazy," may cause *suboptimal* connections. A central property of FMT^{*} is that the cases where a suboptimal connection is made become vanishingly rare as the number of samples goes to infinity, which helps maintain the algorithm's asymptotically optimality. This manifests itself into a key computational advantage—by restricting collision detection to only locally-optimal connections, FMT* (as opposed to, *e.g.*, PRM* [6]) avoids a large number of costly collision-check computations, at the price of a vanishingly small "degree" of suboptimality. We refer the reader to [10], [11] for a detailed description of the algorithm and its advantages.

III-B BFMT^{*} – High-level description

At its core, BFMT* implements a *bi-directional* version of the FMT* algorithm by simultaneously propagating two wavefronts (henceforth, the leaves of an expanding tree will be referred to as the wavefront of the tree) through the free configuration space. BFMT*, therefore, performs a *two-source* dynamic programming recursion over a set of sampled points, and correspondingly generates a *pair* of search trees: one in cost-to-come space from the initial configuration and another in cost-to-go space from the goal configuration (see Fig. 1). Throughout the remainder of the paper, we refer to the former as the *forward tree*, and to the latter as the *backward tree*.



Fig. 1: The BFMT* algorithm generates a *pair* of search trees: one in cost-to-come space from the initial configuration (blue) and another in cost-to-go space from the goal configuration (purple). The path found by the algorithm is in green color.

The dynamic programming recursion performed by BFMT* is characterized by the same lazy feature of FMT* (see Section III-A). However, the time it takes to run BFMT* on a given number of samples can be substantially smaller than for FMT*. Indeed, for uncluttered configuration spaces, the search trees grow hyperspherically, and hence BFMT* only has to expand about half as far (in both trees) as FMT* in order to return a solution. This is made clear in Fig. 1(a), in which FMT* would have to expand the forward tree twice as far to find a solution. Since runtime scales approximately with edge number, which scales as the linear distance covered by the tree raised to the dimension of the state space, we may expect in loosely cluttered configuration spaces an approximate speedup of a factor 2^{d-1} over FMT* in d-dimensional space (the -1 in the exponent is because BFMT^{*} has to expand 2 trees, so it loses one factor of 2 advantage).

III-C BFMT* – Detailed description

To understand the BFMT^{*} algorithm, some background notation must first be introduced. Let S be a set of points sampled independently and identically from the uniform distribution on $\mathcal{X}_{\text{free}}$, to which \mathbf{x}_{init} and \mathbf{x}_{goal} are added. (The extension to non-uniform sampling distributions is addressed in Section IV-C.) Let tree \mathcal{T} be the quadruple $(\mathcal{V}, \mathcal{E}, \mathcal{V}_{\text{unvisited}}, \mathcal{V}_{\text{open}})$, where \mathcal{V} is the set of tree nodes, \mathcal{E}

is the set of tree edges, and $\mathcal{V}_{unvisited}$ and \mathcal{V}_{open} are mutually exclusive sets containing the *unvisited* samples in S and the *wavefront* nodes in \mathcal{V} , correspondingly. To be precise, the unvisited set $\mathcal{V}_{unvisited}$ stores all samples in the sample set S that have not yet been considered for addition to the tree of paths. The wavefront set \mathcal{V}_{open} , on the other hand, tracks in sorted order (by cost from the root) only those nodes which have already been added to the tree that are near enough to tree leaves to actually form better connections. These sets play the same role as their counterparts in FMT^{*}, see [10], [11]. However, in this case BFMT^{*} "grows" two such trees, referred to as $\mathcal{T} = (\mathcal{V}, \mathcal{E}, \mathcal{V}_{unvisited}, \mathcal{V}_{open})$ and $\mathcal{T}' = (\mathcal{V}', \mathcal{E}', \mathcal{V}'_{unvisited}, \mathcal{V}'_{open})$. Initially, \mathcal{T} is the tree rooted at \mathbf{x}_{init} , while \mathcal{T}' is the tree rooted at \mathbf{x}_{goal} . Note, however, that the trees are exchanged during the execution of BFMT^{*}, so \mathcal{T} in Algorithm 1 is not always the tree that contains \mathbf{x}_{init} .

The BFMT* algorithm is represented in Algorithm 1. Before describing BFMT* in detail, we list briefly the basic planning functions employed by the algorithm. Let SAMPLEFREE(n) be a function that returns a set of $n \in$ \mathbb{N} points sampled independently and identically from the uniform distribution on $\mathcal{X}_{\text{free}}$. Let $\text{COST}(\overline{\mathbf{x}'\mathbf{x}})$ be the cost of the straight-line path between configurations x' and x. Let PATH(\mathbf{z}, \mathcal{T}) return the unique path in tree \mathcal{T} from its root to node z. Also, with a slight abuse of notation, let COST(x, T)return the cost of the unique path in tree \mathcal{T} from its root to node x, and let COLLISIONFREE(x, y) be a boolean function returning true if the straight-line path between configurations x and y is collision free. Given a set of samples \mathcal{A} , let NEAR($\mathcal{A}, \mathbf{z}, r$) return the subset of \mathcal{A} within a ball of radius r centered at sample z (*i.e.*, the set { $\mathbf{x} \in \mathcal{A} \mid ||\mathbf{x} - \mathbf{z}|| < r$ }). Let the TERMINATE function represent an external termination criterion (*i.e.*, timeout, maximum number of samples, etc.) which can be used to force early termination (or prevent infinite runtime for infeasible problems). Finally, regarding tree expansion, let SWAP($\mathcal{T}, \mathcal{T}'$) be a function that swaps the two trees \mathcal{T} and \mathcal{T}' . and let COMPANION(\mathcal{T}) return the companion tree \mathcal{T}' to \mathcal{T} (or vice versa).

We are now in position to describe the BFMT* algorithm. First, a set of n configurations in $\mathcal{X}_{\text{free}}$ is determined by drawing samples uniformly. Two trees are then initialized using INITIALIZE as shown in Algorithm 2, with a forward tree rooted at \mathbf{x}_{init} and a reverse tree rooted at \mathbf{x}_{goal} . Once complete, tree expansion begins starting with tree \mathcal{T} rooted at \mathbf{x}_{init} using the EXPAND procedure in Algorithm 3. In the following, the node selected for expansion will be consistently denoted by \mathbf{z} , while $\mathbf{x}_{\mathrm{meet}}$ will denote the lowestcost candidate node for tree connection (*i.e.*, for joining the two trees). The EXPAND procedure requires the specification of a connection radius parameter, r_n , whose selection will be discussed in Section IV. EXPAND implements the "lazy" dynamic programming recursion described (at a high level) in Section III-B, making locally-optimal collision-free connections from nodes x near z unvisited by tree \mathcal{T} (those in set $\mathcal{V}_{\text{unvisited}}$ within search radius r_n of \mathbf{z}) to wavefront nodes \mathbf{x}' near each \mathbf{x} (those in set \mathcal{V}_{open} within search radius r_n of x). Any collision-free edges and newly-connected nodes found are then added to \mathcal{T} , the connection candidate node \mathbf{x}_{meet} is updated, and \mathbf{z} is dropped from the list of wavefront nodes. The key feature of the EXPAND function is that in the execution of the dynamic programming recursion it "lazily"

ignores the presence of obstacles (see line 6) – as discussed in Section IV this comes at no loss of (asymptotic) optimality (see also [10], [11]). Note the EXPAND function is identical to that of unidirectional FMT^{*}, with the exception here of additional lines for tracking the connection candidate x_{meet} .

After expansion, the algorithm checks whether a feasible path is found on line 8. If unsuccessful so far, TERMINATE (which reports failure upon early termination) is checked before proceeding. If the algorithm has not terminated, it checks whether the wavefront of the companion tree is empty (line 14). If this is the case, the INSERT function shown in Algorithm 4 samples a new configuration s uniformly from $\mathcal{X}_{\mathrm{free}}$ and tries to connect it to a nearest neighbor in the companion tree within radius r_n . This way, the expanding tree is ensured to have at least one configuration in its wavefront available for expansion on subsequent iterations (the alternative would be to report failure). This mimics anytime behavior, and by forcing samples to lie close to tree nodes we effectively "reopen" closed nodes for expansion again. Uniform resampling may require many attempts before finding a configuration s which can be successfully connected to \mathcal{V}'_{open} , though this appeared to have a negligible impact on running time for our path planning studies. On the other hand, a more effective strategy might bias resampling towards areas requiring expansion (e.g., bottlenecks, traps) rather than uniformly within tree coverage.

The algorithm then proceeds on lines 16–17 with the selection of the next node (and corresponding tree) for expansion. As shown, BFMT^{*} "swaps" the forward and backward trees on each iteration, each being expanded in turns. As INSERT ensures the companion tree \mathcal{T}' always has at least one node in its frontier \mathcal{V}'_{open} , a node is always available for subsequent expansion as the next z. After selection, the entire process is iterated.

Algorithm 1 The Bi-directional Fast Marching Tree Algorithm (BFMT*)

Require: Motion query $(\mathbf{x}_{init}, \mathbf{x}_{goal})$, Connection radius r_n

```
1: \mathcal{S} \leftarrow {\mathbf{x}_{\text{init}}, \mathbf{x}_{\text{goal}}} \cup \text{SampleFree}(n)
```

```
2: \mathcal{T} \leftarrow \text{INITIALIZE}(\mathcal{S}, \mathbf{x}_{\text{init}})
```

- 3: $\mathcal{T}' \leftarrow \text{INITIALIZE}(\mathcal{S}, \mathbf{x}_{\text{goal}})$
- 4: $\mathbf{z} \leftarrow \mathbf{x}_{init}, \mathbf{x}_{meet} \leftarrow \emptyset$ 5: success := false 6: while success = false 7: EXPAND($\mathcal{T}, \mathbf{z}, \mathbf{x}_{meet}$) 8: if $\mathbf{x}_{meet} \neq \emptyset$ 9: $\sigma^* = PATH(\mathbf{x}_{meet}, \mathcal{T}) \cup PATH(\mathbf{x}_{meet}, \mathcal{T}')$ 10: success = true 11: else 12: if TERMINATE()

13: **return** Failure

- 14: else if $\mathcal{V}'_{\mathrm{open}} = \varnothing$
- 15: $INSERT(\mathcal{T}')$
- 16: $\mathbf{z} \leftarrow \arg\min\{\text{Cost}(\mathbf{x}', \mathcal{T}')\}$
- $\mathbf{x}' \in \mathcal{V}'_{\text{open}}$

```
17: SWAP(\mathcal{T}, \mathcal{T}')
18: return \sigma^*
```

Algorithm 2 Initializes a Fast Marching Tree

1: function INITIALIZE(S, \mathbf{x}_0) 2: $\mathcal{V} \leftarrow {\mathbf{x}_0}$ 3: $\mathcal{E} \leftarrow \varnothing$ 4: $\mathcal{V}_{\text{unvisited}} \leftarrow S \setminus {\mathbf{x}_0}$ 5: $\mathcal{V}_{\text{open}} \leftarrow {\mathbf{x}_0}$ 6: return $\mathcal{T} = (\mathcal{V}, \mathcal{E}, \mathcal{V}_{\text{unvisited}}, \mathcal{V}_{\text{open}})$

Algorithm 3 Fast Marching Tree Expansion Step

Require: Connection radius r_n 1: function EXPAND($\mathcal{T} = (\mathcal{V}, \mathcal{E}, \mathcal{V}_{unvisited}, \mathcal{V}_{open}), \mathbf{z}, \mathbf{x}_{meet})$ 2: $\mathcal{V}_{\mathrm{open,new}} \leftarrow \varnothing$ $Z_{\text{near}} \leftarrow \text{NEAR}(\mathcal{V}_{\text{unvisited}}, \mathbf{z}, r_n)$ 3: for $\mathbf{x} \in Z_{\text{near}}$ 4: $X_{\text{near}} \leftarrow \text{NEAR}(\mathcal{V}_{\text{open}}, \mathbf{x}, r_n)$ 5: $\mathbf{x}_{\min} \gets \arg\min\{\text{COST}(\mathbf{x}', \mathcal{T}) + \text{COST}(\overline{\mathbf{x}'\mathbf{x}})\}$ 6: $\mathbf{x}' \in X_{\text{near}}$ if CollisionFree($\mathbf{x}_{min}, \mathbf{x}$) 7: $\mathcal{V} \leftarrow \mathcal{V} \cup \{\mathbf{x}\}$ \triangleright Add x to tree 8: $\mathcal{E} \leftarrow \mathcal{E} \, \cup \, \{(\mathbf{x}_{\min}, \mathbf{x})\}$ ▷ Add edge to tree 9: $\mathcal{V}_{\mathrm{unvisited}} \leftarrow \mathcal{V}_{\mathrm{unvisited}} \backslash \{ \mathbf{x} \} \triangleright Mark \ \mathbf{x} \ visited$ 10: $\mathcal{V}_{\text{open,new}} \leftarrow \mathcal{V}_{\text{open,new}} \cup \{\mathbf{x}\}$ \triangleright Save x 11: if { $\mathbf{x} \in \mathcal{V}'$ and $\text{COST}(\mathbf{x}, \mathcal{T}) + \text{COST}(\mathbf{x}, \mathcal{T}') <$ 12: $COST(\mathbf{x}_{meet}, \mathcal{T}) + COST(\mathbf{x}_{meet}, \mathcal{T}') \}$ $\mathbf{x}_{meet} \leftarrow \mathbf{x} \ \triangleright$ Save \mathbf{x} as best connection 13:

14:
$$\mathcal{V}_{\text{open}} \leftarrow (\mathcal{V}_{\text{open}} \cup \mathcal{V}_{\text{open},\text{new}}) \setminus \{\mathbf{z}\} \rightarrow \text{Add new}$$

nodes to the wavefront; drop z from the wavefront
15: **return** $\mathcal{T} = (\mathcal{V}, \mathcal{E}, \mathcal{V}_{\text{unvisited}}, \mathcal{V}_{\text{open}})$

III-C.1 BFMT^{*} – *Variations:* As for any bi-directional planner, the correctness and computational efficiency of BFMT^{*} hinge upon two key aspects: (i) how computation is interleaved among the two trees (in other words, which wavefront at each step should be chosen for expansion), and (ii) when the algorithm should terminate. For instance, as an alternative tree expansion strategy (*i.e.*, item (i)), one could replace lines 16–17 with the "balanced trees" condition which enforces more of a balanced search, maintaining equal costs from the root within each wavefront such that the two wavefronts propagate and meet roughly equidistantly in cost-to-go from their roots:

16:
$$\mathbf{z}_{1} \leftarrow \underset{\mathbf{x} \in \mathcal{V}_{\text{open}}}{\operatorname{arg\,min}} \{ \operatorname{COST}(\mathbf{x}, \mathcal{T}) \}$$

17: $\mathbf{z}_{2} \leftarrow \underset{\mathbf{x}' \in \mathcal{V}'_{\text{open}}}{\operatorname{arg\,min}} \{ \operatorname{COST}(\mathbf{x}', \mathcal{T}') \}$
18: $(\mathbf{z}, \mathcal{T}) \leftarrow \underset{(\mathbf{z}_{1}, \mathcal{T}), (\mathbf{z}_{2}, \mathcal{T}')}{\operatorname{arg\,min}} \{ \operatorname{COST}(\mathbf{z}_{i}, \mathcal{T}_{i}) \}$
19: $\mathcal{T}' = \operatorname{COMPANION}(\mathcal{T})$

Similarly, as an alternative termination condition (*i.e.*, item (ii)), one might replace line 8 with the "best path" criterion:

8:
$$\mathbf{z} \in (\mathcal{V}' \setminus \mathcal{V}'_{\text{open}})$$

Currently line 8 returns the first available path discovered, at the moment that the two wavefronts touch at x_{meet} (which

Algorithm 4 Insertion of New Samples
Require: Connection radius r_n
1: function INSERT($\mathcal{T} = (\mathcal{V}, \mathcal{E}, \mathcal{V}_{unvisited}, \mathcal{V}_{open})$)
2: while $V_{open} = \emptyset$ and not TERMINATE()
3: $\mathbf{s} \leftarrow \mathbf{S} \text{AMPLEFREE}(1)$
4: $V_{\text{near}} \leftarrow \text{NEAR}(\mathcal{V}, \mathbf{s}, r_n)$
5: while $V_{\text{near}} \neq \emptyset$
6: $\mathbf{x}_{\min} \leftarrow \arg\min\{\operatorname{COST}(\mathbf{x}, \mathcal{T}) + \operatorname{COST}(\overline{\mathbf{xs}})\}$
$\mathbf{x} \in V_{\text{near}}$
8: $\mathcal{V} \leftarrow \mathcal{V} \cup \{s\} \qquad \triangleright \text{ Add } s \text{ to tree}$
9: $\mathcal{E} \leftarrow \mathcal{E} \cup \{(\mathbf{x}_{\min}, \mathbf{s})\} \triangleright \text{Add edge to tree}$
10: $\mathcal{V}_{\text{open}} \leftarrow \mathcal{V}_{\text{open}} \cup \{\mathbf{s}\} \triangleright \text{ Add to wavefrom } $
11: break
12: else
13: $V_{\text{near}} \leftarrow V_{\text{near}} \setminus \{\mathbf{x}_{\min}\}$
14: return $\mathcal{T} = (\mathcal{V}, \mathcal{E}, \mathcal{V}_{\text{unvisited}}, \mathcal{V}_{\text{open}})$

is not, in general, the lowest cost path). This alternative condition, on the other hand, returns the *exact optimal path* from \mathbf{x}_{init} to \mathbf{x}_{goal} through the given set S of n samples. This change terminates BFMT* when the two wavefronts have propagated sufficiently far through each other that no better solution can be discovered. Intuitively-speaking, this occurs at the first moment where the two trees have both selected, at the current iteration or previously, the same node as the minimum cost node \mathbf{z} from their respective roots.

Though seemingly promising ideas, no appreciable differences in performance were found using the above criteria in combination or otherwise; hence we report only the simplest version of our planner as Algorithm 1.

IV ASYMPTOTIC OPTIMALITY OF BFMT*

In this section, we prove the asymptotic optimality of BFMT*. The claim is proven assuming BFMT* acts without the INSERT procedure (see Algorithm 4) on line 15, in place of which "Failure" is reported instead. The proof for the full algorithm then follows immediately by *a fortiori* argument. We begin with a result essentially stating that any path in $\mathcal{X}_{\text{free}}$ may be "traced" arbitrarily well by connecting randomly-distributed points from a sufficiently large sample set covering the configuration space. We refer to this property as *probabilistic exhaustivity* [20]. We then provide the (asymptotic) optimality proof for BFMT* by showing that the algorithm recovers a solution with cost no greater than that of any tracing path. In the following, let ζ_d denote the volume of the unit ball in *d*-dimensional Euclidean space. Also, let the complement of a probabilistic event \mathcal{A} be denoted by \mathcal{A}^c .

IV-A Probabilistic exhaustivity

Let $\sigma : [0, 1] \to \mathcal{X}$ be a path. Given a set of samples (referred to as waypoints) $\{\mathbf{y}_m\}_{m=1}^M \subset \mathcal{X}$, we associate a path $y : [0, 1] \to \mathcal{X}$ that sequentially connects the nodes $\{\mathbf{y}_m\}$ to (ϵ, r) -trace the path σ if: (i) $||\mathbf{y}_m - \mathbf{y}_{m+1}|| \leq r$ for all m, (ii) the cost of y is bounded as $J(y) \leq (1+\epsilon)J(\sigma)$, and (iii) the distance from any point of y to σ is no more than r, *i.e.*, $\min_{t \in [0, 1]} ||y(s) - \sigma(t)|| \leq r$ for all $s \in [0, 1]$. In the context of sampling-based motion planning, we may expect to find closely-tracing $\{\mathbf{y}_m\}$ as a subset of the sampled points, provided the sample size is large. This notion is formalized in the following theorem (Theorem 4.1), proved as Theorem IV.5 in [20] for the general case of driftless control-affine control systems, a special case of which is path planning without differential constraints (as addressed in this paper).

Theorem 4.1 (Probabilistic exhaustivity): Let $(\mathcal{X}_{\text{free}}, \mathbf{x}_{\text{init}}, \mathbf{x}_{\text{goal}})$ be a path planning problem and let σ : $[0,1] \rightarrow \mathcal{X}_{\text{free}}$ be a feasible path. Let $S = { \mathbf{x}_{init}, \mathbf{x}_{goal} } \cup SAMPLEFREE(n), \epsilon > 0, \text{ and for fixed}$ *n* consider the event \mathcal{A}_n that there exist $\{\mathbf{y}_m\}_{m=1}^M \subset \mathcal{S}$, $y_1 = \mathbf{x}_{\text{init}}, y_M = \mathbf{x}_{\text{goal}}$ which (ϵ, r_n) -trace σ , where

$$r_n = 4\left(1+\eta\right)^{\frac{1}{d}} \left(\frac{1}{d}\right)^{\frac{1}{d}} \left(\frac{\mu(\mathcal{X}_{\text{free}})}{\zeta_d}\right)^{\frac{1}{d}} \left(\frac{\log n}{n}\right)^{\frac{1}{d}}$$

for a parameter $\eta \geq 0$. Then, as $n \to \infty$, the probability that \mathcal{A}_n does not occur is asymptotically bounded as $P[\mathcal{A}_n^c] =$ $O\left(n^{-\frac{\eta}{d}}\log^{-\frac{1}{d}}n\right).$

IV-B Asymptotic optimality (AO)

We are now in a position to prove the asymptotic optimality of BFMT^{*}, which represents the main result of this section. We start with an important lemma, which relates the cost of the path returned by BFMT* to that of any feasible path.

Lemma 4.2 (Bi-directional FMT^{*} cost comparison): Let $\sigma: [0, 1] \to \mathcal{X}_{\text{free}}$ be a feasible path with strong δ -clearance. Consider running BFMT^{*} to completion with n samples and a connection radius:

$$r_n = 4 \left(1+\eta\right)^{\frac{1}{d}} \left(\frac{1}{d}\right)^{\frac{1}{d}} \left(\frac{\mu(\mathcal{X}_{\text{free}})}{\zeta_d}\right)^{\frac{1}{d}} \left(\frac{\log n}{n}\right)^{\frac{1}{d}}$$

for a parameter $\eta \geq 0$. Let J_n denote the cost of the path returned by BFMT^{*}. Then for fixed $\epsilon > 0$:

$$\mathbf{P}[J_n > (1+\epsilon)J(\sigma)] = \mathbf{O}\left(n^{-\frac{\eta}{d}}\log^{-\frac{1}{d}}n\right).$$

Proof: If $\mathbf{x}_{init} = \mathbf{x}_{goal}$, then BFMT^{*} immediately terminates with $J_n = 0$, trivially satisfying the claim. Thus we assume that $\mathbf{x}_{init} \neq \mathbf{x}_{goal}$. Consider *n* sufficiently large so that $r_n \leq \min\{\delta/2, \epsilon ||\mathbf{x}_{init} - \mathbf{x}_{goal}||/2\}$, and apply Theorem 4.1 to produce, with probability at least $1 - O\left(n^{-\frac{\eta}{d}}\log^{-\frac{1}{d}}n\right)$, a sequence of waypoints $\{\mathbf{y}_m\}_{m=1}^M \subset$ $\mathcal{S}, \mathbf{y}_1 = \mathbf{x}_{\text{init}}, \mathbf{y}_M = \mathbf{x}_{\text{goal}}$ which $(\epsilon/2, r_n)$ -trace σ . We claim that in the event that such $\{\mathbf{y}_m\}$ exists, the BFMT^{*} algorithm returns a path with cost upper bounded as $J_n \leq$ $J(y) + r_n \le (1 + \epsilon/2)J(\sigma) + (\epsilon/2)J(\sigma) = (1 + \epsilon)J(\sigma)$. It is clear that the desired result follows from this claim.

Assume the existence of an $(\epsilon/2, r_n)$ -tracing $\{\mathbf{y}_m\}$. Let $\mathcal{B}(\mathbf{x},r)$ represent a ball of radius r centered at a sample **x**. Note that our upper bound on r_n implies that $\mathcal{B}(\mathbf{y}_m, r_n)$ intersects no obstacles. This follows from our choice of r_n and the distance bound

$$\inf_{\mathbf{s}\in\mathcal{X}_{\text{obs}}} ||\mathbf{y}_m - \mathbf{s}|| \ge \inf_{\mathbf{s}\in\mathcal{X}_{\text{obs}}} ||\boldsymbol{\sigma}_m - \mathbf{s}|| - ||\mathbf{y}_m - \boldsymbol{\sigma}_m|| \\
\ge 2r_n - r_n \ge r_n.$$

where σ_m is the closest point of σ to \mathbf{y}_m . This fact, along with $||\mathbf{y}_m - \mathbf{y}_{m+1}|| \leq r_n$ for all *m*, implies that when a connection is attempted for y_m , both y_{m-1} and y_{m+1} will be in the search radius and no obstacles will lie within that search radius. Running BFMT* to completion generates one cost-to-come tree $\mathcal{T}_i(\mathcal{V}_i, \mathcal{E}_i, \mathcal{V}_{\text{open},i}, \mathcal{V}_{\text{unvisited},i})$ and one costto-go tree $\mathcal{T}_g(\mathcal{V}_g, \mathcal{E}_g, \mathcal{V}_{\text{open},g}, \mathcal{V}_{\text{unvisited},g})$ rooted at \mathbf{x}_{init} and \mathbf{x}_{goal} , respectively (the subscripts *i* and *g* are used to identify the root of a tree without ambiguity). The above discussion ensures that the trees will meet and the algorithm will return a feasible path when it terminates – the path outlined by the waypoints $\{\mathbf{y}_m\}$ disallows the possibility of failure.

For each sample point $\mathbf{x} \in S$, let $J_i(\mathbf{x}) := \text{COST}(\mathbf{x}, \mathcal{T}_i)$ denote the cost-to-come of x from x_{init} in T_i , and let $J_g(\mathbf{x}) := \operatorname{COST}(\mathbf{x},\mathcal{T}_g)$ denote the cost-to-go from \mathbf{x} to \mathbf{x}_{goal} in \mathcal{T}_{g} . If \mathbf{x} is not contained in a tree \mathcal{T}_{k} , $k = \{i, g\}$, we set $J_k(\mathbf{x}) = \infty$. When the algorithm terminates, we know there exists a sample point $\mathbf{x}_{meet} \in \mathcal{V}_i \cap \mathcal{V}_g$ where the two trees meet; indeed we select the particular meeting point $\mathbf{x}_{meet} = \arg\min_{\mathbf{x}\in\mathcal{V}_i\cap\mathcal{V}_g} J_i(\mathbf{x}) + J_g(\mathbf{x})$. Then $J_n =$ $J_i(\mathbf{x}_{meet}) + J_g(\mathbf{x}_{meet})$. We now note a lemma bounding the costs-to-come of the $\{y_m\}$, the proof of which may be found as an inductive hypothesis (Eq. 5) in Theorem VI.1 of [20].

Lemma 4.3: Let $m \in \{1, \ldots, M\}$. If $J_i(\mathbf{y}_m) < \infty$, Lemma 4.3: Let $m \in \{1, ..., M\}$. If $J_i(\mathbf{y}_m) < \infty$, then $J_i(\mathbf{y}_m) \leq \sum_{k=1}^{m-1} ||\mathbf{y}_k - \mathbf{y}_{k+1}||$. Otherwise if $\mathbf{y}_m \notin \mathcal{V}_i$, then $J_i(\mathbf{x}_{meet}) \leq \sum_{k=1}^{m-1} ||\mathbf{y}_k - \mathbf{y}_{k+1}||$. Similarly if $J_g(\mathbf{y}_m) < \infty$, then $J_g(\mathbf{y}_m) \leq \sum_{k=m}^{M-1} ||\mathbf{y}_k - \mathbf{y}_{k+1}||$; otherwise $J_g(\mathbf{x}_{meet}) \leq \sum_{k=m}^{M-1} ||\mathbf{y}_k - \mathbf{y}_{k+1}||$. To bound the performance J_n of BFMT*, there are two

cases to consider. Note in either case we find that $J_n \leq$ $J(\mathbf{y}) + r_n$, thus completing the proof.

<u>Case 1</u>: There exists some $\mathbf{y}_m \in \mathcal{V}_i \cap \mathcal{V}_g$. In this case, $J_n = J_i(\mathbf{x}_{\text{meet}}) + J_g(\mathbf{x}_{\text{meet}}) \leq$ $\begin{array}{l} J_i(\mathbf{y}_m) + J_g(\mathbf{y}_m) < \infty \text{ by our choice of } \mathbf{x}_{\text{meet}}.\\ \text{Then applying Lemma 4.3 we see that } J_n \leq J_i(\mathbf{y}_m) + J_g(\mathbf{y}_m) \leq \sum_{k=1}^{M-1} ||\mathbf{y}_k - \mathbf{y}_{k+1}|| = I_{(\mathbf{x})}. \end{array}$ $J(\mathbf{y}).$

 $\begin{array}{lll} \underline{\text{Case 2}} \text{: There are no } \mathbf{y}_m \in \mathcal{V}_i \cap \mathcal{V}_g.\\ \text{Consider } \widetilde{m} &=& \max\{m \mid J_i(\mathbf{y}_m) < \infty\}. \end{array} \text{ Then} \end{array}$ $\mathbf{y}_{\widetilde{m}} \in \mathcal{V}_i$ and $\mathbf{y}_{\widetilde{m}}$ can not have been the minimum cost element of $\mathcal{V}_{\text{open},i}$ at any point during algorithm execution or else we would have connected $\mathbf{y}_{\widetilde{m}+1} \in \mathcal{V}_i$. Let \mathbf{z} denote the minimum cost element of $\mathcal{V}_{\mathrm{open},i}$ when $\mathbf{x}_{\mathrm{meet}}$ was added to $\mathcal{V}_i.$ We have the bound:

$$J_{i}(\mathbf{x}_{\text{meet}}) \leq J_{i}(\mathbf{z}) + r_{n} \leq J_{i}(\mathbf{y}_{\widetilde{m}}) + r_{n}$$
$$\leq \sum_{k=1}^{m-1} ||\mathbf{y}_{k} - \mathbf{y}_{k+1}|| + r_{n}.$$
(1)

By our assumption for this case, $\mathbf{y}_{\widetilde{m}} \notin \mathcal{V}_g$. Then by Lemma 4.3 we know that $J_g(\mathbf{x}_{\text{meet}}) \leq \sum_{k=m}^{M-1} ||\mathbf{y}_k - \mathbf{y}_{k+1}||$. Combining with the previous inequality yields $J_n = J_i(\mathbf{x}_{meet}) + J_g(\mathbf{x}_{meet}) \le \sum_{k=1}^{M-1} ||\mathbf{y}_k - \mathbf{y}_{k+1}|| + r_n = J(y) + J_g(\mathbf{x}_{meet}) \le \sum_{k=1}^{M-1} ||\mathbf{y}_k - \mathbf{y}_{k+1}|| + r_n = J(y) + J_g(\mathbf{x}_{meet}) \le J_g(\mathbf{x}_{meet}) \le J_g(\mathbf{x}_{meet}) + J_g(\mathbf{x}_{meet}) \le J_g(\mathbf{x}$ r_n .

Remark 4.4 (Tightened bound for connection radius): As discussed in [20], for the sake of clarity the constant term 4 in the expression for r_n is greater than is necessary for Theorem 4.1 to hold. A more careful argument along the lines of the original FMT^{*} AO proof [10] would suffice to show that $r_n = 2(1 + \eta)^{1/d} \cdot \left(\frac{1}{d}\right)^{1/d} \left(\frac{\mu(\mathcal{X}_{\text{free}})}{\zeta_d}\right)^{1/d} \left(\frac{\log(n)}{n}\right)^{1/d}$ satisfies the theorem as well.

Remark 4.5 (Alternative termination criteria): The proof holds as well for the different expansion and termination criteria discussed in Section III-C.1. However, due to space constraints the details are omitted.

We now have everything we need to show that BFMT^{*} is asymptotically-optimal. The next theorem defines this formally.

Theorem 4.6 (BFMT* asymptotic optimality): Assume a δ -robustly feasible path planning problem as defined in Section II with optimal path σ^* of cost J^* . Then BFMT* converges in probability to σ^* as the number of samples $n \to \infty$. Specifically, for any $\epsilon > 0$,

$$\lim \mathbf{P}[J_n > (1+\epsilon)J^*] = 0$$

Proof: The proof follows as a corollary to Lemma 4.2. By our δ -robustly feasible assumption, we can find a strong δ -clearance feasible path $\sigma : [0, 1] \rightarrow \mathcal{X}_{\text{free}}$ that approximates σ^* with cost $J(\sigma) < (1 + \epsilon/3)J^*$ (*i.e.*, less than factor $\epsilon/3$ from J^*), for any $\epsilon > 0$. By Lemma 4.2, we can choose n sufficiently large such that BFMT^{*} returns an $\epsilon/3$ cost approximation to the approximant:

$$\mathbf{P}\left[J_n > (1+\epsilon/3)^2 J^*\right] < \mathbf{P}\left[J_n > (1+\epsilon/3)J(\sigma)\right]$$
$$= \mathbf{O}\left(n^{-\frac{\eta}{d}}\log^{-\frac{1}{d}}n\right)$$

To approach the optimal path, let the number of samples $n \to \infty$. It follows that, for any $\eta \ge 0$:

$$\lim_{n \to \infty} P\left[J_n > (1 + \epsilon/3)^2 J^*\right] < \lim_{n \to \infty} O\left(n^{-\frac{\eta}{d}} \log^{-\frac{1}{d}} n\right) = 0$$

Now we relate this to the original claim. First suppose that

 $\epsilon \leq 3$. From $(1 + \epsilon/3)^2 \leq 1 + \epsilon$, the event $\{J_n > (1 + \epsilon)J^*\}$ is a subset of the event $\{J_n > (1 + \epsilon/3)^2J^*\}$, hence:

$$\lim_{n \to \infty} \mathbf{P}[J_n > (1+\epsilon)J^*] \le \lim_{n \to \infty} \mathbf{P}\left[J_n > (1+\epsilon/3)^2 J^*\right] = 0$$

Because the probability is monotone-decreasing in ϵ as ϵ increases, the statement holds for all $\epsilon > 3$ as well (to see this, apply Lemma 4.2 again for m sufficiently large to handle $\epsilon = 3$; then by similar argument as above $P[J_m > (1 + \epsilon)J^*] < P[J_m > (1 + 3)J^*] = O\left(m^{-\frac{\eta}{d}}\log^{-\frac{1}{d}}m\right)$ and take the limit as $m \to \infty$). Hence $\lim_{n\to\infty} P[J_n > (1 + \epsilon)J^*] = 0$ holds for arbitrary ϵ , and we see that BFMT^{*} converges in probability to the optimal path, as claimed.

Remark 4.7 (Convergence rate): Note that we can also translate the convergence rate from Lemma 4.2 to the setup of Theorem 4.6, which does not require strong δ -clearance. For any $\epsilon > 0$, the optimal path can be approximated by a strong- δ -clear path with cost less than $(1+\epsilon)J(\sigma)$ and we can focus on approximating that path to high-enough precision to still approximate the optimal path to within $(1 + \epsilon)$. Since the convergence rate in Lemma 4.2 only contains ϵ in the rate's constant, the big-O convergence rate remains the same. This generalizes the convergence rate result in [11], which only applied to a specific obstacle-free configuration space, initial configuration, and goal region.

IV-C Sampling and cost generalizations

It is worth mentioning that the asymptotic optimality (AO) properties of BFMT^{*} are not limited to uniform configuration sampling and arc-length cost functions. For example, if one

has prior information about areas that the optimal path is unlikely to pass through, it may be advantageous to consider a non-uniform sampling strategy that downsamples these regions. As long as the sampling density is lower-bounded by a *positive* number over the configuration space, BFMT* can be slightly altered (by merely increasing r_n by a constant factor) to ensure it stays AO. The argument is analogous to that made in [11], and essentially proceeds by making the search radius wide enough to balance out the detrimental effect of the lower sampling density (in some areas). An additional common concern is when the cost is not arc-length, but some other metric or line integral cost. In either case, BFMT* need only consider cost balls instead of Euclidean balls when making connections. Details for adjusting the algorithm and why the AO proof still holds can be derived from [11]. The argument basically shows that the triangle inequality either holds exactly (for metric costs) or approximately, and that this approximation goes away in the limit as $n \to \infty$.

V SIMULATIONS

In this section, we provide numerical path-planning experiments that compare the performance of BFMT* with other sampling-based, asymptotically-optimal planning algorithms $(namely, FMT^*, RRT^*, and PRM^*)^2$. Given a planning workspace and query, we aim to observe the quality of the solution returned as a function of the execution time allotted to the algorithm. Here dynamic constraints are neglected and arc-length is used as path cost. As a basis for quality comparison between incremental or "anytime" planners (such as RRT^{*}) and non-incremental planners (such as BFMT^{*}, which generate solutions via sample batches), we vary the number of samples drawn by the planners during the planning process (which in essence serves as a proxy to execution time). Note sample count has a different connotation depending on the planner that will not necessarily be the number of nodes stored in the constructed solution graph - for RRT* (with one sample drawn per iteration), this is the number of iterations, while for FMT*, PRM*, and BFMT*, this is the number of free space samples taken during initialization.

V-A Simulation Setup

To generate simulation data for a given experiment, we queried the planning algorithms once each for a series of sample counts, recorded the cost of the solution returned, the planner execution time³, and whether the planner succeeded or not, then repeated this process over 50 trials. To ensure a fair comparison, each planning algorithm was tested using the Open Motion Planning Library (OMPL) v1.0.0 [21], which provides high-quality implementations of many state-of-the-art planners and a common framework for executing motion plans. In this way, we could ensure that all algorithms employed the *exact same* primitive routines (*e.g.*, nearest-neighbor search, collision-checking, data handling, *etc.*), and measure their performances fairly. Regarding implementation,

²Existing state-of-the-art sampling-based, bi-directional algorithms (namely, RRT-Connect and SBL) were initially also included. However, average costs for RRT-Connect and SBL were roughly 2-4x greater, which occluded the details of other curves; they were thus omitted for clarity

³Code for all experiments was written in C++. Corresponding programs were compiled and run on a Linux-operated PC, clocked at 2.4 GHz and equipped with 7.5 GB of RAM.

BFMT^{*}, FMT^{*}, and PRM^{*} used $\eta = 0$ from Lemma 4.2 for the nearest-neighbor radius r_n in order to satisfy the theoretical bounds provided in Section IV and [6]. For RRT^{*}, we used the default OMPL settings; namely, a 5% goal bias and a steering parameter equal to 20% of the maximum extent of the configuration space (except for the α -puzzle, in which case a value of 1.1 was found to work much better). For FMT^{*}, we included the same INSERT routine as BFMT^{*} for configuration resampling upon failure. For all algorithms, early termination (*e.g.*, using TERMINATE for BFMT^{*}) was suppressed by defining a 1000 second time limit, well above each planner's worst-case execution time.

Before proceeding, note that each marker shown on the plots throughout this section represents a single simulation at a fixed sample count. The points on the curves, however, represent the mean cost/time of *successful* algorithm runs *only* for a particular sample count, with error bars corresponding to one standard deviation of the 50 run sample mean.⁴ Sample counts varied from the order of 200 to 2000 points for 2D problems, from 1000 to 30000 points for 3D problems, and 500 to 4000 points for the hypercube examples.

V-B Results and Discussion

Here we present benchmarking results (average solution cost versus average execution times and success rates) comparing BFMT* to other state-of-the-art sampling-based planners. Three benchmarking test scenarios were considered: (1) a 2D "bug trap" and (2) a 2D "maze" problem for a convex polyhedral robot in the SE(2) configuration space, as well as (3) a challenging 3D problem called the " α -puzzle" in which we seek to untangle two loops of metal (non-convex) in the SE(3) configuration space. All problems were drawn directly from OMPL's bank of tests, and are illustrated in Fig. 2. In each case, collision-checks relied on OMPL's built-in collision-checking library, FCL. Additionally, to tease out the performance of BFMT* relative to FMT* in high-dimensional environments, we also studied a point mass robot moving in cluttered unit hypercubes of 5 and 10 dimensions.⁵



(a) $\mathbb{SE}(2)$ bug trap

(b) $\mathbb{SE}(2)$ maze (c) $\mathbb{SE}(2)$

(c) $\mathbb{SE}(3) \alpha$ -puzzle

Fig. 2: Depictions of the three OMPL rigid-body planning problems

Figure 3 shows the results for each BFMT*, FMT*, RRT*, and PRM*. Performance here is measured by execution time on the x-axis and solution cost on the y-axis—high quality data points are therefore located in the lower-left corner

(low-cost solutions obtained quickly). The plots reveal that both FMT* and BFMT* for the most part outperform RRT* as well as PRM*. In particular, BFMT* and FMT* achieve higher success rates (always a flat 100% for the cases studied) in shorter time. To extract further information, we need to examine each test in detail.

In the Bug Trap and Maze problems, BFMT* notably generates the same cost-time curve as FMT* (meaning they return solutions of very similar cost for a given sample count), but with data points shifted to the left (indicating they were obtained in shorter execution time). Though not shown due to slow running times for PRM* (whose results had to be truncated to clarify detail), all planners appear to tend towards similar low-cost solutions as more execution time was allocated. However BFMT* and FMT* seem to converge to an optimum much faster, particularly for the Maze problem (on the order of 1.5 and 2.0 seconds respectively, compared to 3-4 seconds for RRT* and 5-7 seconds for PRM*). This contrast becomes even more evident for the α -puzzle. Here we see an unusual spread of solutions - one in a band at around 500 cost and another at around 275. These indicate the presence of two solution types, or homotopy classes: one corresponding to the true α -puzzle solution, and another less-efficient path. This appears to have yielded a "bump" in the BFMT* cost-curve, where increasing the sample count momentarily gives an increased average cost. We believe this is a result of how BFMT* trees interconnect; at this count, by unlucky circumstance, the longer homotopy seems to be found first more often than usual. But as proved in Section IV, the behavior disappears as $n \to \infty$. Note RRT^{*} seems to avoid this issue through goal biasing. Despite the difficult problem structure, BFMT* finds the cheaper homotopy faster than other planners, with many more of its data points clustered in the lower-left corner, generally at lower costs and times than RRT* and of equal quality but faster times than FMT*.

These results suggest that BFMT* tends to an optimal cost at least as fast as the other planners, and sometimes much faster. To shed light on the relative performance of FMT* and BFMT* further, we compare them in higher dimensions. Results for the 5D and 10D hypercube are shown in Fig. 4 (success rates were again at 100%, and were thus omitted). Here BFMT^{*} substantially outperforms FMT^{*}, particularly as dimension increases, with convergence in roughly 0.5 and 1.4 seconds (5D), and 5 and 20 seconds (10D) on average. This suggests that *reachable volumes* play a significant role in their execution time. The relatively small volume of reachable configurations around the goal at the corner implies that the reverse tree of BFMT* expands its wavefront through many fewer states than the forward tree of FMT* (which in fact needlessly expands towards the zero-vector); tree interconnection in the bi-directional case prevents its forward tree from growing too large compared to unidirectional search. This is pronounced exponentially as the dimension increases. In trap or maze-like scenarios, however, bi-directionality does not seem to change significantly the number of states explored by the marching trees, leading to comparable performance for the $\mathbb{SE}(2)$ bug-trap and maze. Note we expect a greater contrast in execution times in favor of BFMT* as the cost of collision-checking increases, such as with many non-convex obstacles or in time-varying environments.

⁴Standard deviation of the mean indicates where we expect with one- σ confidence the distribution mean to lie based on the 50-run sample mean, and is related to the standard deviation of the distribution by $\sigma_{\mu} = \sigma/\sqrt{50}$.

⁵We populated the space to 50% obstacle coverage with randomly-sized, axis-oriented hyperrectangles. \mathbf{x}_{init} was set to the center at $[0.5, \ldots, 0.5]$, with the goal \mathbf{x}_{goal} at the ones-vector (*i.e.*, $[1, \ldots, 1]$).



Fig. 3: Simulation results for the three OMPL scenarios.



Fig. 4: FMT* and BFMT* results for 5D and 10D cluttered hypercubes (50% coverage; all success rates were 100%).

VI CONCLUSION

In this paper, we presented a bi-directional, sampling-based, asymptotically-optimal motion planning algorithm named BFMT*, for which we rigorously proved its optimality and characterized its convergence rate – arguably a first in the field of bi-directional sampling-based planning. Numerical experiments in \mathbb{R}^d , $\mathbb{SE}(2)$, and $\mathbb{SE}(3)$ revealed that BFMT* tends to an optimal solution at least as fast as its state-ofthe-art counterparts, and in some cases significantly faster. Convergence rates are expected to improve with parallelization, in which each tree is grown using a separate CPU.

Future research will examine BFMT*'s interaction with more advanced techniques, such as adaptive sampling near narrow passages or sample biasing in INSERT (Algorithm 4) towards failed wavefronts. We also plan to extend BFMT* to dynamic environments through lazy re-evaluation (leveraging its tree-like forward and reverse path structures) in a way that reuses previous results as much as possible. Maintaining bounds on run-time performance and solution quality in this context will remain the greatest challenges. Ultimately, we hope that BFMT* will enable fast, easy-to-implement replanning with proven performance guarantees, analogous to planning in static environments as we have shown here.

VII ACKNOWLEDGEMENTS

This work was supported by an Early Career Faculty grant from NASA's Space Technology Research Grants Program (Grant NNX12AQ43G). The authors would like to thank technical representative Brent Barbee for his support.

REFERENCES

- S. M. LaValle, *Planning Algorithms*. Cambridge University Press, 2006, available at http://planning.cs.uiuc.edu.
- [2] L. E. Kavraki, P. Švestka, J. C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional spaces," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [3] D. Hsu, J.-C. Latombe, and R. Motwani, "Path planning in expansive configuration spaces," *International Journal of Computational Geometry & Applications*, vol. 9, no. 4, pp. 495–512, 1999.
- [4] J. M. Phillips, N. Bedrossian, and L. E. Kavraki, "Guided expansive spaces trees: A search strategy for motion-and cost-constrained state spaces," in *Proc. IEEE Conf. on Robotics and Automation*, vol. 4, New Orleans, LA, May 2004, pp. 3968–3973.
- [5] S. M. LaValle and J. J. Kuffner, "Randomized kinodynamic planning," *International Journal of Robotics Research*, vol. 20, no. 5, pp. 378–400, 2001.
- [6] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [7] J. Luo and K. Hauser, "An empirical study of optimal motion planning," in *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, Sep. 2014, pp. 1761–1768.
- [8] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot, "Bit*: Batch informed trees for optimal sampling-based planning via dynamic programming on implicit random geometric graphs," 2014, available at http://arxiv. org/abs/1405.5848.
- [9] O. Arslan and P. Tsiotras, "Use of relaxation methods in samplingbased algorithms for optimal motion planning," in *Proc. IEEE Conf.* on Robotics and Automation, Karlsruhe, Germany, May 2013, pp. 2421–2428.
- [10] L. Janson and M. Pavone, "Fast Marching Trees: A fast marching sampling-based method for optimal motion planning in many dimensions," in *International Symposium on Robotics Research*, 2013.
- [11] L. Janson, E. Schmerling, A. Clark, and M. Pavone, "Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions," *International Journal of Robotics Research*, 2015, in Press.
- [12] Y. K. Hwang and N. Ahuja, "Gross Motion Planning a Survey," ACM Comput. Surv., vol. 24, no. 3, pp. 219–291, Sep. 1992.
- [13] I. Pohl, Bi-directional and heuristic search in path problems. Department of Computer Science, Stanford University., 1969, no. 104.
- [14] M. Luby and P. Ragde, "A bidirectional shortest-path algorithm with good average-case behavior," *Algorithmica*, vol. 4, no. 1-4, pp. 551–567, 1989.
- [15] A. Goldberg, H. Kaplan, and R. Werneck, *Reach for A*: Efficient Point-to-Point Shortest Path Algorithms*. Springer, 2006, ch. 12, pp. 129–143.
- [16] J. J. Kuffner and S. M. LaValle, "RRT-Connect: An efficient approach to single-query path planning," in *Proc. IEEE Conf. on Robotics and Automation*, San Francisco, CA, Apr. 2000, pp. 995–1001.
- [17] G. Sánchez and J.-C. Latombe, "A single-query bi-directional probabilistic roadmap planner with lazy collision checking," in *International Journal of Robotics Research*. Springer, 2003, pp. 403–417.
- Journal of Robotics Research. Springer, 2003, pp. 403–417.
 [18] B. Akgun and M. Stilman, "Sampling heuristics for optimal motion planning in high dimensions," in *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems.* IEEE, 2011, pp. 2640–2645.
- [19] M. Jordan and A. Perez, "Optimal Bidirectional Rapidly-Exploring Random Trees," Aug. 2013, http://people.csail.mit.edu/aperez/obirrt/ csailtech.pdf.
- [20] E. Schmerling, L. Janson, and M. Pavone, "Optimal sampling-based motion planning under differential constraints: the driftless case," in *Proc. IEEE Conf. on Robotics and Automation*, 2015, extended version available at http://arxiv.org/abs/1403.2483/.
- [21] I. A. Şucan, M. Moll, and L. E. Kavraki, "The Open Motion Planning Library," *IEEE Robotics and Automation Magazine*, vol. 19, no. 4, pp. 72–82, Dec. 2012, http://ompl.kavrakilab.org.