Efficient Sampling-Based Bottleneck Pathfinding over Cost Maps

Kiril Solovey* and Dan Halperin*

November 12, 2018

Abstract

We introduce a simple yet effective sampling-based planner that is tailored for *bottleneck pathfinding*: Given an implicitly-defined cost map $\mathcal{M} : \mathbb{R}^d \to \mathbb{R}$, which assigns to every point in space a real value, we wish to find a path connecting two given points, which minimizes the maximal value with respect to $\mathcal{M}$. We demonstrate the capabilities of our algorithm, which we call *bottleneck tree* (BTT), on several challenging instances of the problem involving multiple agents, where it outperforms the state-of-the-art cost-map planning technique T-RRT*. In addition to its efficiency, BTT requires the tuning of only a single parameter: the number of samples. On the theoretical side, we study the asymptotic properties of our method and consider the special setting where the computed trajectories must be monotone in all coordinates. This constraint arises in cases where the problem involves the coordination of multiple agents that are restricted to forward motions along predefined paths.

1 Introduction

Motion planning is a widely studied problem in robotics. In its basic form it is concerned with moving a robot between two given configurations while avoiding collisions with obstacles. Typically, we are interested in paths of high quality, which lead to lower energy consumption, shorter execution time, or safer execution for the robot and its surrounding. One of the most studied quality measures is path length, which was first studied from the combinatorial and geometric perspective [28] and has recently gained popularity with the introduction of PRM*, RRT* [23] and subsequent work. Another popular measure is the clearance (see, e.g, [50]) of a path: the *clearance* of a particular robot configuration is the distance to the closest forbidden configuration. The goal here is to find a path that maximizes the minimum clearance of the configurations along it.

The latter example is a special case of the *bottleneck pathfinding* problem: we are given an implicitly-defined cost map $\mathcal{M} : \mathcal{C} \to \mathbb{R}$ which assigns to every configuration $c \in \mathcal{C}$ of the robot a value; the goal is to find, given start and target configurations, a continuous path $\nu : [0, 1] \to \mathcal{C}$, which minimizes the expression $\max_{\tau \in [0,1]}\{\mathcal{M}(\nu(\tau))\}$.

Bottleneck pathfinding can also arise in settings involving multiple robots. However, in most cases multi-robot motion planning is computationally intractable (see, e.g., [20, 38]), even when one is only concerned with finding *a* solution. Thus, *decoupled* planners (see, e.g., [48, 37, 4]) usually construct a set of d paths—one for each robot—and attempt to coordinate between the robots along their predefined paths in order to avoid collisions. Now suppose that we also wish

to find the *safest* coordination—the one which keeps the robots in a maximal distance apart. This problem again, can be reformulated as a bottleneck-pathfinding problem, where d is the dimension of the effective search space. In particular, denote by $\sigma_i : [0,1] \to \mathcal{C}$ the path assigned for robot i, where $1 \leq i \leq d$, and define for a given d-dimensional point $x = (x_1, \ldots, x_m) \in [0,1]^d$ the cost map $\mathcal{M}(x) := 1/\max_{1 \leq i < j \leq c} \|\sigma(x_i) - \sigma(x_j)\|$. In other cases, the robots are restricted a priori to predefined paths, as in factory assembly lines [43], aviation routes [10] and traffic intersections [14].

An interesting complication that typically arises from the examples of the previous paragraph, is that each robot must move forward along its path and cannot backtrack. This restriction induces paths that are *monotone* in each of the coordinates in the search space $[0,1]^d$. This is reminiscent of the notion of *Fréchet matching*, which can be viewed as another instance of bottleneck pathfinding. Fréchet matching is a popular similarity measure between curves which has been extensively studied by the computational-geometry community (see, e.g., [3, 18]) and recently was used in robotics [49, 32, 19]. This matching takes into consideration not only the "shape" of curves but also the order in which the points are arranged along them.

Contribution. Rather than tackling each problem individually, we take a sampling-based approach and develop an effective planner that can cope with complex instances of the bottleneck-pathfinding problem. We demonstrate the capabilities of our technique termed *bottleneck tree* (BTT) on a number of challenging problems, where it outperforms the state-of-the-art technique T-RRT* by several orders of magnitude. In addition to its efficiency, BTT requires the tuning of only a single parameter: the number of samples. On the theoretical side, we study the asymptotic properties of our method and consider the special setting where the returned paths must be monotone in all coordinates.

Notice that the monotonicity requirement makes the problem harder, particularly in the sampling-based setting. Consider for example a PRM G that is constructed using a connection radius r_n and n samples. By removing all the non-monotone edges from G we obtain a much sparser graph G': informally, G' may retain only a $1/2^d$ fraction of the edges of G, where d is the dimension of the problem. Fortunately, our analysis of BTT indicates that in order to guarantee optimality in the monotone regime, the connection radius does not have to be drastically increased over its value in the non-monotone case of say PRM.

We have already considered sampling-based bottleneck pathfinding in a previous work [39]. However, there we employed a PRM-based approach which is far inferior to BTT. Already when working in a four-dimensional C-space, the PRM-base approach, which seeks to explore the *entire* C-space, becomes prohibitively expensive. In contrast, BTT quickly produces low-cost solutions even in a seven-dimensional space. We also provide here stronger theoretical analysis of our new technique for the monotone case. These differences are discussed in more detail in Sections 5 and 6.

In Section 2 we review related work. In Section 3 we provide a formal definition of bottleneck pathfinding. In Section 4 we describe our BTT algorithm for the problem. In Section 5 we provide an asymptotic analysis of the method in the monotone case. In Section 6 we report on experimental results and conclude with a discussion and future work in Section 7.

2 Related work

Particular instances of the basic motion-planning problem involving a small number of degrees of freedom can be solved efficiently in a complete manner [36]. Complete planners can even cope with scenarios involving multiple robots, if some assumptions are made about the input (see, e.g., [40, 1, 47]). However, in general motion planning is computationally intractable [20, 38, 44,

34]. Thus, most of the recent efforts in this area are aimed at the development of sampling-based planners, which attempt to capture the connectivity of the free space using random sampling.

We continue our literature review with general-purpose sampling-based planners, with an emphasis on planners that are applicable to settings involving cost maps. We then proceed to specific examples of bottleneck pathfinding that were studied from the combinatorial-algorithmic perspective.

2.1 Sampling-based motion planning

Early sampling-based planners such as PRM [25] and RRT [27] have focused on finding *a* solution. Although some efforts were made to understand the quality of paths produced by such planners [33, 30, 17] this issue remained elusive until recently. In their influential work Karaman and Frazzoli [23] introduced the planners PRM* and RRT*, which were shown to be *asymptotically optimal*, i.e., guaranteed to return a solution whose cost converges to the optimum. Several asymptotically-optimal planners that focus on the path-length cost have later emerged [5, 22, 35]. In our recent work [41] we develop a general framework for the analysis of theoretical properties of a variety of sampling-based planners through a novel connection with the theory of random geometric graphs.

We now proceed to discuss sampling-based planning in the presence of an underlying cost map $\mathcal{M}$, as in bottleneck pathfinding. RRT* theoretically guarantees to converge to the optimum even when the cost of the path is computed with respect to $\mathcal{M}$. However, in practice this convergence is very slow, since RRT* does not take into account the cost of $\mathcal{M}$ while exploring $\mathcal{C}$, and puts too much effort in the traversal of regions of high cost. On the other hand, T-RRT [21] equips RRT with a *transition test*, which biases the exploration towards low-cost regions of $\mathcal{M}$. However, T-RRT does not take path quality into consideration and cannot improve the cost of a solution that was already obtained. Consequently, this planner provides no optimality guarantees. The T-RRT* algorithm [16] combines RRT* with T-RRT, which yields an asymptotically-optimal planner that also works well in practice. The authors of [16] test T-RRT* for the case of *integral cost*, which sums the cost of the configurations along the path, and *mechanical cost*, which sums the positive cost variations along the path. We mention that T-RRT* is also applicable to the bottleneck cost, although this setting was not tested in that work. Lastly, we note that we are not aware of other techniques that consider planning in the presence of a cost map in the general setting of sampling-based planning.

2.2 Bottleneck pathfinding and other problems

Cost maps typically arise as *implicit* underlying structures in many problems involving motion and path planning. de Berg and van Kreveld [15] consider the problem of path planning over a *known* mountainous region. They describe data structures that can efficiently obtain, given two query points (which are the start and goal of the desired path), paths with various properties, such as those that strictly descend and paths that minimize the maximal height obtained, i.e., bottleneck paths.

We have already mentioned the problem of high-clearance paths. For the simple case of a disc robot moving amid polygonal obstacles this problem can be solved efficiently by constructing a Voronoi diagram over the sites that consist of the obstacles [31]. There is an obvious trade-off between clearance and path length, a short paths tend to get very close to the obstacles. The work by Wein et al. [50] describes a data structure that given a clearance threshold $\delta > 0$ returns the shortest path with that clearance. Agarwal et al. [2] consider a particular cost function suggested by Wein et al. [51] that combines path length and clearance and develop an efficient approximation algorithm for this case.

Perhaps the most popular example of bottleneck pathfinding from the recent years is *Fréchet matching*, which is concerned with quantifying the similarity of two given curves: the Fréchet distance between two given curves can be described as the shortest leash that allows a person to walk her dog, where the former is restricted to move along the first curve and latter along the other. It is generally assumed that this quantity tends to be more informative when only *forward motions* along the curves are permitted, which transform into *monotone* paths in the search space. Fréchet matching has been extensively studied for the case of two curves: several efficient techniques that solve the problem exist (see, e.g., [3, 8]). The problem can be extended to multiple curves although only algorithms that are exponential in the number of curves are known [18, 9]. Several papers consider extensions of the problem involving more complex input objects [6, 7] or cases where additional constraints such as obstacles are imposed on the "leash" [12, 11]. Finally, we mention our recent work [39] where we apply a sampling-based PRM-like technique for solving several Fréchet-type problems.

3   Preliminaries

We provide a formal definition of the bottleneck-pathfinding problem. For some fixed dimension $d \geq 2$, let $\mathcal{M} : [0,1]^d \to \mathbb{R}$ be a cost map that assigns to every point in $[0,1]^d$ a value in $\mathbb{R}$. We will refer from now on to $[0,1]^d$ as the *parameter space* in order to distinguish it from the configuration space of the underlying problem.

Notice that the definition of $\mathcal{M}$ does not imply that we restrict our discussion to the setting of a single robot whose configuration space is $[0,1]^d$. For instance, $[0,1]^d$ can represent the parameter space of d curves of the form $\sigma_i : [0,1] \to \mathcal{C}_i$, where for every $1 \leq i \leq d$, σ_i describes that motion of robot i in the configuration space $\mathcal{C}_i$. Typically in such settings involving multiple robots having predefined paths the robots are only permitted to move forward along their paths and never backtrack. This restriction induces *monotone* motions in $[0,1]^d$. Given two points $x = (x_1, \ldots, x_d), y = (y_1, \ldots, y_d) \in \mathbb{R}^d$ the notation $x \preceq y$ implies that for every $1 \leq i \leq d$ it hold that $x_i \leq y_i$. Additionally, we use the notation $x \preceq_\delta y$ for $\delta > 0$ to indicate that $x \preceq y$ and for every $1 \leq i \leq d$ it hold that $y_i - x_i \geq \delta$.

For simplicity we will assume that the goal consists of planning paths between the points $\mathbb{0}, \mathbb{1} \in [0,1]^d$, where $\mathbb{0} = (0, \ldots, 0), \mathbb{1} = (1, \ldots, 1)$.

Definition 1. *A plan $\nu : [0,1] \to [0,1]^d$ is a continuous path in $[0,1]^d$ that satisfies the following two constraints: (i) $\nu(0) = \mathbb{0}, \nu(1) = \mathbb{1}$; (ii) it is monotone in each of the d coordinates, i.e., for every $0 \leq \tau \leq \tau' \leq 1$ and $\nu(\tau) \preceq \nu(\tau')$.*

Given a path (or a plan) ν, its *bottleneck cost* is defined to be $\mathcal{M}(\nu) = \max_{\tau \in [0,1]} \mathcal{M}(\nu(\tau))$. In Section 6 we will address specific examples of the following problem:

Definition 2. *Given a cost map $\mathcal{M} : [0,1]^d \to \mathbb{R}$ the (monotone) bottleneck-pathfinding problem consists of finding a plan ν which minimizes $\mathcal{M}(\nu)$.*

4   Bottleneck-tree planner

In Algorithm 1 we describe our sampling-based technique for bottleneck pathfinding, which we call bottleneck tree (BTT). The algorithm can be viewed as a lazy version of PRM which explores the underlying graph in a Dijkstra-like fashion, according to cost-to-come, with respect to the bottleneck cost over $\mathcal{M}$. Thus, it bears some resemblance to FMT* [22] since the two implicitly explore an underlying PRM. However, FMT* behaves very differently from BTT since it minimizes the path-length cost.

Algorithm 1 BOTTLENECK-TREE(n, r_n)

1: $V := \{\mathbb{0}, \mathbb{1}\} \cup \mathtt{sample}(n)$;
2: **for all** $x \in V$ **do**
3:    $c(x) := \infty$; $p(x) := \text{null}$
4: $c(\mathbb{0}) := \mathcal{M}(\mathbb{0})$
5: **while** $\exists x \in V$ such that $c(x) < \infty$ **do**
6:    $z := \mathtt{get\_min}(V)$
7:    **if** $z == \mathbb{1}$ **then**
8:      **return** $\mathtt{path}(T, \mathbb{0}, \mathbb{1})$
9:    $V := V \setminus \{z\}$
10:   $N_z := \mathtt{near}(z, V, r_n)$
11:   **for all** $x \in N_z$ **do**
12:     **if** $z \preceq x$ **and** $c(z) < c(x)$ **then**
13:       $c_{\text{new}} := \max\{c(z), \mathcal{M}(z, x)\}$
14:       **if** $c_{\text{new}} < c(x)$ **then**
15:         $c(x) := c_{\text{new}}$; $p(x) := z$
16: **return** $\emptyset$

BTT accepts the number of samples n and the connection radius r_n as parameters. It maintains a directed minimum spanning tree $\mathcal{T}$ of an underlying *monotone* PRM or a random geometric graph (RGG) [41][1]:

Definition 3. *Given $n \in \mathbb{N}_+$ denote by $\mathcal{X}_n = \{X_1, \ldots, X_n\}$, n points chosen independently and uniformly at random from $[0,1]^d$. For a connection radius $r_n > 0$ denote by $\mathcal{G}_n = \mathcal{G}(\{\mathbb{0}, \mathbb{1}\} \cup \mathcal{X}_n; r_n)$ the monotone RGG defined over the vertex set $\{\mathbb{0}, \mathbb{1}\} \cup \mathcal{X}_n$. $\mathcal{G}_n$ has a directed edged (x, y) for $x, y \in \{\mathbb{0}, \mathbb{1}\} \cup \mathcal{X}_n$ iff $x \preceq y$ and $\|x - y\| \leq r_n$, where $\|\cdot\|$ represents the standard Euclidean distance.*

The algorithm keeps track of the unvisited vertices of $\mathcal{G}_n$ using the set V. It maintains for each vertex x the cost-to-come over the visited portion of $\mathcal{G}_n$ and its parent for this specific path, are denoted by $c(x)$ and $p(x)$, respectively.

In line 1 we initialize V with $\mathbb{0}, \mathbb{1}$, and a set of n uniform samples in $[0,1]^d$, which are generated using $\mathtt{sample}$. In lines 2-3 the cost-to-come and the parent attributes are initialized. In each iteration of the **while** loop, which begins in line 5, the algorithm extracts the vertex $z \in V$ which minimizes $c(z)$, using $\mathtt{get\_min}$ (line 6). If z turns out to be $\mathbb{1}$ (line 7), then a path from $\mathbb{0}$ to $\mathbb{1}$ is returned. This is achieved using $\mathtt{path}$, which simply follows the ancestors of $\mathbb{1}$ stored in the p attribute until reaching $\mathbb{0}$. Otherwise, z is removed from V (line 9) and a subset of its neighbors in $\mathcal{G}_n$ that are also members of V are obtained using $\mathtt{near}$ (line 10). For each such neighbor x (line 11) it is checked whether $z \preceq x$ and whether the cost-to-come of x can potentially improve by arriving to it from z (line 12). An alternative cost-to-come for x which is induced by a path that reaches from z is calculated (line 13). In particular, this cost is the maximum between the cost-to-come of z and the cost of the edge[2] from z to x with respect to $\mathcal{M}$. If the alternative option improves its cost then the parents and the cost of x are updated accordingly (line 15).

In preparation for the following section, where we study the asymptotic behavior of BTT, we mention here that given that the underlying graph $\mathcal{G}_n$ contains a path from $\mathbb{0}$ to $\mathbb{1}$, BTT finds

---

[1]PRMs and RGGs are equivalent in our context.

[2]The cost of an edge with respect to a given cost map can be approximated by dense sampling along the edge, as is customary in motion planning.

the minimum bottleneck path over $\mathcal{G}_n$, namely the path whose maximal $\mathcal{M}$ value is minimized. We omit a formal proof of this fact here. It is rather straightforward and is very similar similar to the completeness proof of the standard Dijkstra algorithm, but for the bottleneck cost; see, e.g., [13, Theorem 25.10].

5   Asymptotic analysis

In the previous section we have argued that BTT returns a solution that minimizes the bottleneck cost over a fixed graph $G \in \mathcal{G}_n$. A major question in the analysis of sampling-based motion-planning algorithms is under what conditions does a discrete graph structure captures the underlying continuous space. For this purpose, we study the asymptotic properties of $\mathcal{G}_n = \mathcal{G}(\{0, 1\} \cup \mathcal{X}_n; r_n)$ (Definition 3). To the best of our knowledge, existing proofs concerning asymptotic optimality or completeness of sampling-based planners (see, e.g., [23, 22]) only apply to the non-monotone (and standard) setting. Such results cannot be extended as-are to the analysis of $\mathcal{G}_n$. The remainder of this section is dedicated to strengthening our our analysis concerning the special and harder case that imposes the monotonicity requirement.

Previously [39], we were able to give a convergence guarantee in the monotone case by introducing a connection radius (r_n) of the form $f(n) \left(\frac{\log n}{n} \right)^{1/d}$, where $f(n)$ is any function that diminishes as n tends to ∞. Unfortunately, this statement does not indicate which specific function should be used in practice, and an uncareful selection of f could lead to undesired behavior of BTT. In particular, if f grows too slowly with n the graph becomes disconnected and BTT may not be able to find a solution at all. On the other hand, f that grows too fast with n could result in a needlessly dense graph which will take very long time to traverse. In this paper, we replace this somewhat abstract definition with a connection radius of the form $\gamma \left(\frac{\log n}{n} \right)^{1/d}$, where γ is a constant depending only on d. This is a more typical structure of bounds in the context of random geometric graphs and sampling-based planners. Furthermore, we show in Section 6 that this formulation of r_n leads to a quick convergence of BTT to the optimum.

For the purpose of the analysis, we define a supergraph of $\mathcal{G}_n$, denoted by $\mathcal{L}_n = \mathcal{L}(\{0, 1\} \cup \mathcal{X}_n; r_n)$, which corresponds to the standard and non-monotone random geometric graph (see [41]). In particular, it connects every pair of vertices $x, y \in \{0, 1\} \cup \mathcal{X}_n$ with an edge if and only if $\|x - y\| \leq r_n$, regardless of whether the monotonicity constraint is satisfied.

In order to guarantee asymptotic optimality, one typically has to make some assumptions with respect to the solution one hopes to converge to. First, we introduce basic definitions. Denote by $\mathcal{B}_r(x)$ the d-dimensional Euclidean ball of radius $r > 0$ centered at $x \in \mathbb{R}^d$ and $\mathcal{B}_r(\Gamma) = \bigcup_{x \in \Gamma} \mathcal{B}_r(x)$ for any $\Gamma \subseteq \mathbb{R}^d$. Given a curve $\nu : [0, 1] \to \mathbb{R}^d$ define $\mathcal{B}_r(\nu) = \bigcup_{\tau \in [0,1]} \mathcal{B}_r(\nu(\tau))$. Additionally, denote the image of a curve ν by $\mathrm{Im}(\nu) = \bigcup_{\tau \in [0,1]} \{\nu(\tau)\}$.

Definition 4. *Given* $\mathcal{M} : [0, 1]^d \to \mathbb{R}$, *a plan* $\nu := [0, 1] \to [0, 1]^d$ *is called* robust *if for every* $\varepsilon > 0$ *there exists* $\delta > 0$ *such that for any plan* ν' *for which* $\mathrm{Im}(\nu) \subset \mathcal{B}_\delta(\nu)$ *it follows that* $\mathcal{M}(\nu') \leq (1 + \varepsilon)\mathcal{M}(\nu)$. *A plan that attains the infimum cost, over all robust plans, is termed* robustly optimal *and is denoted by* ν^*.

The following theorem is the main theoretical contribution of this paper. The constant γ, which is defined below, was first obtained in [22], for a different problem, and involving non-monotone connections.

Theorem 1. *Suppose that* $\mathcal{L}_n = \mathcal{L}(\mathcal{X}_n \cup \{0, 1\}; r_n)$ *with* $r_n = \gamma \left(\frac{\log n}{n} \right)^{1/d}$, *where* $\gamma = (1 + \eta)2(d\theta_d)^{-1/d}$ *with* θ_d *representing the Lebesgue measure of the unit Euclidean ball, and* $\eta > 0$ *is*

*any positive constant. Then $\mathcal{L}_n$ contains a path ν_n connecting $\mathbb{0}$ to $\mathbb{1}$ which has the following properties a.s.: **(i)** $\mathcal{M}(\nu) = (1 + o(1))\mathcal{M}(\nu^*)$; **(ii)** at most $o(1)$ fraction of the edges along ν_n are non-monotone.*

This theorem implies that $\mathcal{L}_n$ contains a path ν that is *almost* entirely monotone and its cost converges to the optimum. Notice that in theory this does not necessarily mean that BTT is asymptotically optimal, as the formulation of the algorithm in Alg. 1 is concerned with paths that are *entirely* monotone. However, our experiments (Section 6) suggest that this theoretical gap can be bridged.

5.1 Proof of Theorem 1

This subsection is devoted to the proof of Theorem 1. It relies on some components that were previously described in the works of Janson et al. [22], Karaman and Frazzoli [23], and our work [39]. However, we note that several ideas employed here are brand new and so is the final result.

Our proof follows the common "ball-covering" argument [23, 22, 24], which shows that there exists a collection of samples from $\mathcal{X}_n$, that are located in the vicinity of "landmark" points that lie on some plan of interest. This guarantees that $\mathcal{L}_n$ contains a path that has a cost similar to the robustly optimal plan ν^*. Then we use a refined "(α, β) ball-covering" [22] argument, which shows that many of the samples lie very closely to the landmarks—so close that most of them form monotone subchains. This induces the existence of paths that are almost-entirely monotone and have desirable bottleneck cost.

We start by selecting a fixed $\varepsilon > 0$. As ν^* is robustly optimal, there exists $\delta > 0$ such that for every plan ν with $\mathrm{Im}(\nu) \subseteq \mathcal{B}_\delta(\nu^*)$ it follows that $\mathcal{M}(\nu) \leq (1 + \varepsilon)\mathcal{M}(\nu^*)$. Define $\ell = 2/\delta$. Similarly to the proof of Theorem 6 in [39], there exists a sequence of $\ell - 1 \leq k \leq \ell$ points $Q = \{q_1, \ldots, q_k\} \subset \mathrm{Im}(\nu^*)$ where $q_1 = \mathbb{0}, q_k = \mathbb{1}$, and $q_j \preceq_{\delta/2} q_{j+1}$ for any $1 \leq j < k$. Note that ℓ is finite and independent of n. This follows from the fact that $\mathbb{0} \preceq_1 \mathbb{1}$ and every subpath of ν^* must "make progress" in relation to at least one of the coordinates. In particular, for every $1 \leq j \leq d$ define $Q^j = \{q_1^j, \ldots, q_\ell^j\} \subset \mathrm{Im}(\nu^*)$ to be a collection of ℓ points along ν^* with the following property: for every $1 \leq i' \leq \ell$ the jth coordinate of $q_{i'}^j$ is equal to $j \cdot 2/\delta$. Observe that for any $1 \leq i' \leq \ell$ there exists $1 \leq j \leq d$ and $1 \leq i'' \leq \ell$ such that $\mathbb{0} \preceq_{(j-1)\cdot 2/\delta} q_{i''}^j$ and $\mathbb{0} \npreceq_{(j-1)\cdot 2/\delta+\Delta} q_{i''}^j$ for any $\Delta > 0$, i.e., one of the coordinates of $q_{i''}^j$ is equal to $j \cdot 2/\delta$ (see Figure 1). Thus, it follows that $Q \subset \bigcup_{j=1}^d Q^j$.

As ν^* can be arbitrarily complex we define a simplified plan, which is located in the vicinity of ν^*, and is consequently of low cost. Firstly, for every $1 \leq j < k$ denote by $\nu_{j,j+1}$ the straight-line path from q_j to q_{j+1} (see Figure 1). Now, define $\nu_{1,k}$ to be the concatenation of these $k - 1$ subpaths. Observe that $\nu_{1,k}$ is a plan.

Fix $\psi \in (0, 1)$ and define $r'_n := \frac{r_n}{2(2+\psi)}$. We generate a set of M_n "landmark" points $P = \{p_1, \ldots, p_{M_n}\}$, which are placed along $\nu_{1,k}$, i.e., $P \subset \mathrm{Im}(\nu_{1,k})$, such that for every $1 \leq i \leq M_n - 1$ it holds that $\|p_i - p_{i+1}\| \leq r'_n$. Note that $p_1 = \mathbb{0}$. For simplicity we also assume that $p_{M_n} = \mathbb{1}$.

Now define a set of balls centered at the landmarks. In particular, for every $1 \leq i \leq M_n$ define the set $B_{n,i} := \mathcal{B}_{r'_n}(p_i)$. Additionally, denote by A_n the event representing that for every $1 \leq i \leq M_n$ it holds that $\mathcal{X}_n \cap B_{n,i} \neq \emptyset$. We have the following lemma, which is proven in [22, Lemma 4.4]:

Lemma 1. *The event A_n holds a.s.*

Corollary 1. *The graph $\mathcal{L}_n$ contains a path ν_n connecting $\mathbb{0}$ to $\mathbb{1}$ such that $\mathcal{M}(\nu_n) \leq (1 + \varepsilon)\mathcal{M}(\nu^*)$ a.s..*

Figure 1: Visualization of the proof of Theorem 1 for $d = 2$. The red curve represents ν^*, on which lie the points $Q^1 \cup Q^2$ (represented by red bullets). Observe that $Q \subset Q^1 \cup Q^2$. The dashed blue straight lines represent $\nu_{1,2}, \nu_{2,3}$.

Proof. Suppose that A_n holds. Then for every $1 < i < M_n$ there exists $x_i \in \mathcal{X}_n$ such that $x_i \in B_{n,i}$. Observe that for any $1 \leq i < M_n$ it holds that $\|x_i - x_{i+1}\| \leq r_n$ and so $\mathcal{L}_n$ contains an edge from x_i to x_{i+1}. Thus, we can define ν_n to be the path consisting of the points $\mathbb{0}, x_2, \ldots, x_{M_n-1}, \mathbb{1}$.

For any $1 < i < M_n$ denote by p'_i a point along ν^* which is the closest to p_i. Observe that $\|p_i - p'_i\| \leq \delta/2$. This, combined with the fact that $\|x_i - p_i\| \leq r'_n$, implies that $\|x_i - p'_i\| \leq \delta/2 + r'_n \leq \delta$, which then guarantees that $\mathcal{M}(x_i) \leq (1+\varepsilon)\mathcal{M}(\nu^*)$, by definition of ν^*. Similar arguments can be applied to the edges of ν_n. Thus, $\mathcal{M}(\nu_n) \leq (1+\varepsilon)\mathcal{M}(\nu^*)$. $\qquad\square$

In order to reason about the monotonicity of ν_n we make the following crucial observation:

Lemma 2. *There exists a fixed constant $\beta \in (0,1)$ such that for any $1 \leq i < M_n$ it holds that $p_i \preceq_{\beta r_n} p_{i+1}$.*

Proof. First, assume that there exists $1 \leq j < k$ for which $p_i, p_{i+1} \in \mathrm{Im}(\nu_{j,j+1})$, and note that by definition the straight line from p_i to p_{i+1} is a subsegment of $\nu_{j,j+1}$. Now, recall that $q_j \preceq_{\delta/2} q_{j+1}$ and $\|p_i - p_{i+1}\| = r'_n$. Additionally, denote by $L_{j,j+1} = \|q_j - q_{j+1}\|$ and note that this value is a constant independent of n. Finally, let β_j be the maximal value for which $p_i \preceq_{\beta_j r_n} p_{i+1}$ and observe that $\frac{L_{j,j+1}}{r'_n} = \frac{\delta/2}{\beta_j r_n}$. Thus, $\beta_j = \frac{\psi \delta}{2(2+\psi)}$.

Now consider the case where $p_i \in \mathrm{Im}(\nu_{j,j+1}), p_{i+1} \in \mathrm{Im}(\nu_{j+1,j+2})$. Without loss of generality, assume that $\|p_i - q_{j+1}\| \geq \|p_{i+1} - q_{j+1}\|$. Thus, $p_i \preceq_{\beta_j r_n/2} q_{j+1}$, which implies that $p_i \preceq_{\beta_j r_n/2} p_{i+1}$. Finally, define $\beta = \min_{1 \leq j < k}\{\beta_j/2\}$. $\qquad\square$

Corollary 2. *Suppose that $p_i, p_{i+1} \in P$ for some $1 \leq i \leq M_n$. Then for any two points $x_i, x_{i+1} \in \mathcal{X}_n$ such that $x_i \in \mathcal{B}_{\beta' r_n}(p_i), x_{i+1} \in \mathcal{B}_{\beta' r_n}(p_{i+1})$, where $\beta' \leq \beta/2$, it holds that $x_i \preceq x_{i+1}$.*

This motivates the construction of another set of balls. For every $1 \leq i \leq M_n$ define the set $B^\beta_{n,i} := \mathcal{B}_{\beta r_n/2}(p_i)$. Additionally, define K^β_n to be the number of $B^\beta_{n,i}$ balls which do not

contain samples from $\mathcal{X}_n$. Formally, $K_n^{\beta} := \left|\left\{B_{n,i}^{\beta} \cap \mathcal{X}_n = \emptyset\right\}_{i=1}^{M_n}\right|$. We borrow the following lemma from [22, Lemma 4.3].

Lemma 3. *For any fixed $\alpha \in (0,1)$ it holds that $K_n^{\beta} < \alpha M_n$ a.s.*

Corollary 3. *For any fixed $\alpha \in (0,1)$ the following holds a.s.: $\mathcal{L}_n$ contains a path ν_n connecting $\mathbb{0}$ to $\mathbb{1}$ such that $\mathcal{M}(\nu_n) \leq (1+\varepsilon)\mathcal{M}(\nu^*)$, and with at most $2\alpha M_n$ non-monotone edges.*

Proof. For every $1 \leq i \leq M_n$ define $x_i \in \mathcal{X}_n$ to be some point in $\mathcal{X}_n \cap B_{n,i}^{\beta}$ if this set is non-empty, and otherwise some point in $B_{n,i}$. By Corollary 1, the points $x_1, \ldots, x_{M_n}$ induce a path ν_n such that $\mathcal{M}(\nu_m) \leq (1+\varepsilon)\mathcal{M}(\nu^*)$, and such a path exists a.s. By Corollary 2, every two consecutive points x_i, x_{i+1} such that $x_i \in B_{n,i}^{\beta}, x_{i+1} \in B_{n,i+1}^{\beta}$ contribute one monotone edge to ν_n. Conversely, each point $x_i \notin B_{n,i}^{\beta}$ contributes at most two non-monotone edges. By Lemma 3, there are at most αM_n points of the latter type a.s. $\qquad\square$

It only remains to replace the two constants ε, α with $o(1)$ in order to obtain the exact formulation of Theorem 1. Since we have made no assumptions concerning these two values, they can be replaced with the sequences $\varepsilon_i = 1/i, \alpha_{i'} = 1/i'$, where $i, i' \in \mathbb{N}$. See more details in [42, Theorem 6].

6 Experiments

In this section we report on experimental results of applying BTT to several challenging problems. We also compare its performance with T-RRT*, which is the state-of-the-art for planning on cost maps. Since BTT is concerned with optimality, we compare it against T-RRT* rather than with its non-optimal original version T-RRT. Our technique is faster than T-RRT* by several orders of magnitude. We mention that in a previous work [39] we experimented on similar problems but with a simpler PRM-flavor technique, which is significantly slower than BTT. Due to this fact we do not compare against it. We also note that the we compute fully monotone paths in all the experiments reported below, using the formulation described in Alg. 1.

We implemented BTT and T-RRT* in C++, and tested them on scenarios involving two-dimensional objects. BTT has only two parameters, which are the number of samples n and the connection radius r_n. The latter was set to the value described in Theorem 1 with $\eta = 1$. We note that even though this theorem only suggests that this value suffices for asymptotic optimality in the monotone case, the experiments support this claim also after finite, relatively short, running time. In particular, BTT obtains an initial solution fairly quickly and converges to the optimum, in scenarios where the optimum is known.

Our implementation of T-RRT* is based on its OMPL [45] implementation. T-RRT* uses a connection radius, which we set according to [23]. T-RRT* also has the parameters T, T_{rate}, which are set according to the guidelines in [16]. The length of the `extend` routine, and the rate of target bias, were set according to the OMPL implementation of T-RRT*.

Nearest-neighbor search calls in T-RRT* were performed using FLANN [29]. As the set of samples of BTT is generated in one batch, we use the efficient RTG data structure (see, e.g., [26]) for nearest-neighbor search. Geometric objects were represented with CGAL [46]. Experiments were conducted on a PC with Intel i7-2600 3.4GHz processor with 8GB of memory, running a 64-bit Windows 7 OS.
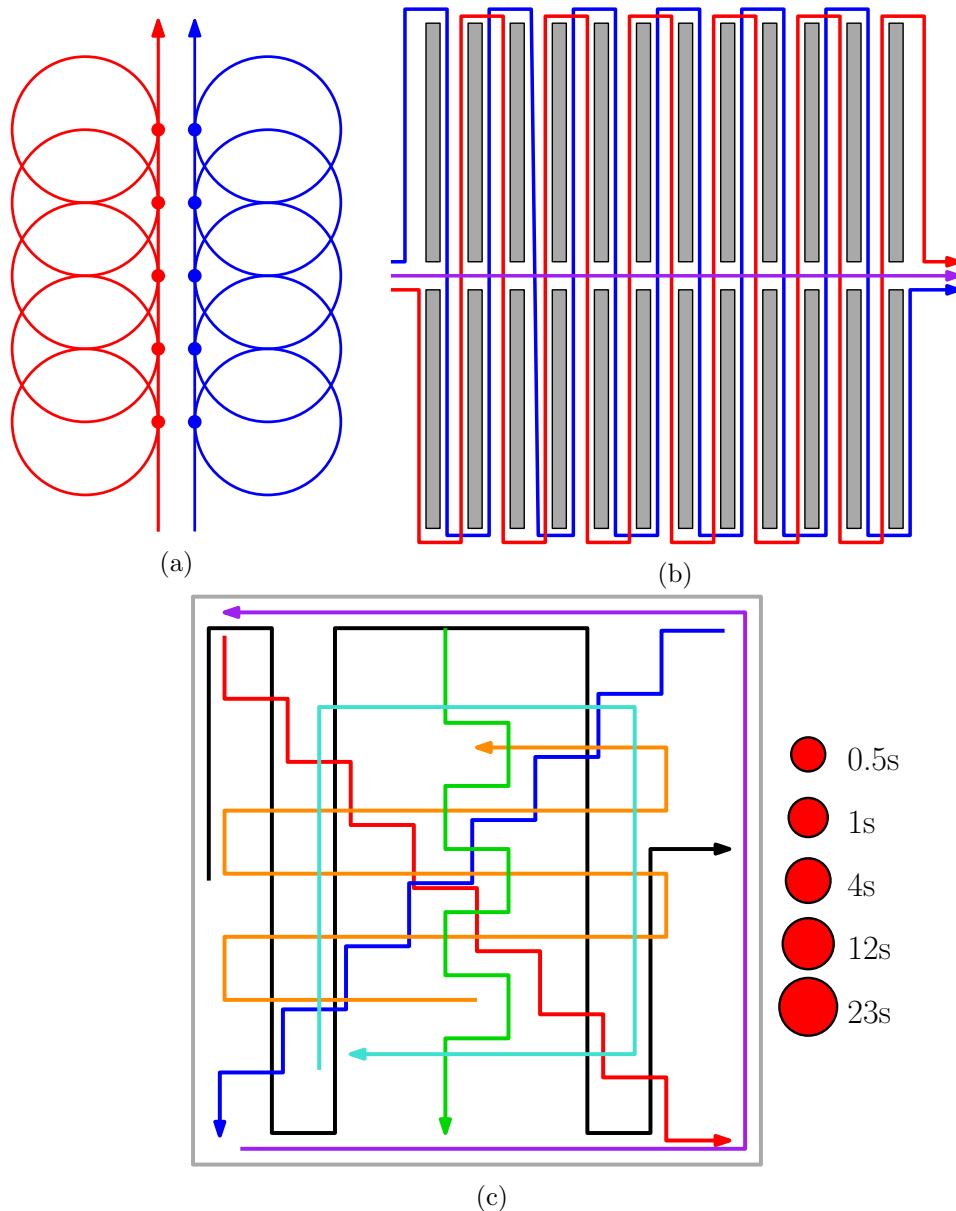
(a)

(b)

(c)

Figure 2: (a) Scenario for **(P1)**, for $d = 2$: the right (blue) curve consists of five circular loops of radius 0.15, where the entrance and exit point to each circle is indicated by a bullet. The left red curve is similarly defined, and the two curves are separated by a vertical distance of 0.04. The optimal matching of cost 0.34 is obtained in the following manner: when a given circle of the red curve is traversed, the position along the blue curve is fixed to the entrance point of the circle directly to the right of the traversed circle, and vice versa. (b) Scenario for **(P2)**: two agents moving along the red and blue curves, respectively, must minimize the distance to the leader on the horizontal purple curve while maintaining visibility with the leader. (c) Scenario for **(P3)**: finding the safest coordination between $d = 7$ moving agents whose routes are depicted as red, blue, turquoise, orange, green, black, and purple curves. The red circles on the right represent the amount of separation obtained by BTT for the specified running time: given a plan returned by BTT, it induces a maximal radius of a disc, such that when d copies of this disc are centered on the moving agents, they are guaranteed to remain collision-free.

6.1 Scenarios

We test BTT on several challenging problems, which are described below. We mention that the underlying implementation of BTT and T-RRT* does not change between the different scenarios and we only change the subroutine responsible for the computation of cost-map values.

(P1) Fréchet matching: The goal is to find a traversal which *minimizes* the pairwise distance between the traversal points along $d \in \{2, 3, 4\}$ curves. Recall that given $d \geq 2$ curves $\sigma_1, \ldots, \sigma_d : [0, 1] \rightarrow \mathbb{R}^2$, the cost map $\mathcal{M} : [0, 1]^d \rightarrow \mathbb{R}_+$ induced by the problem of Fréchet matching [3] is defined to be $\mathcal{M}(\tau_1, \ldots, \tau_d) = \max_{1 \leq i < j \leq d} \|\sigma_i(\tau_i) - \sigma_j(\tau_j)\|$ for $\tau_1, \ldots, \tau_d \in [0, 1]$. The scenario for $d = 2$ is depicted in Fig. 2a. For $d = 3$ we add an identical blue curve, and for $d = 4$ another red curve. Note that the optimum is the same for all values of d here.

(P2) Leader following: The problem consists of three curves, where the purple one (Fig. 2b) represents the motion of a "leader", whereas the blue and the red curves represent the "followers". The goal is to plan the motion of the three agents such that at least one of the followers sees the (current position of the) leader at any given time, where the view of each agent can be obstructed by the gray walls. In addition, we must minimize the distance between the leader to its closest follower at a given time. The optimal solution is attained in the following manner: while the leader passes between two vertically-opposite walls the blue follower keeps an eye on it while the red follower goes around a similar set of walls. Then, when the leader becomes visible to the red agent, the blue agent catches up with the other two agents in an analogous fashion, and this process is repeated.

(P3) Safest coordination: This problem consists of finding the *safest coordination* among $d = 7$ agents moving along predefined routes, which represents a complex traffic intersection. The goal is to find the traversal which *maximizes* the pairwise distance between the d traversal points. This can be viewed as "anti"-Fréchet coordination. Each route is drawn in a different color, where the direction is indicated by an arrow (see Fig. 2c).

6.2 Results

We report on the running times and the obtained cost values of BTT and T-RRT*, after averaging over ten runs in each setting.

We start with **(P1)**. For $d \in \{2, 3\}$, BTT rapidly converges to the optimum (see Fig. 3). For $d = 4$ the convergence is slower, although a value relatively close to the optimum is reached (observe that the trivial solution in which the curves proceed simultaneously induces a cost of 0.64). Note that BTT's performance for $d = 3$ is better than that of T-RRT* for $d = 2$. A similar phenomenon occurs for BTT with $d = 4$ and T-RRT* for $d = 3$. Lastly, for $d = 4$ T-RRT* manages to find a solution that is only slightly better than the trivial one after 60s.

In **(P2)** BTT attains the (near-)optimal solution described earlier after 13s on average. Here T-RRT* is unable to obtain any solution within 60s. In **(P3)**, again, T-RRT* is unable to obtain a solution within 60s. In contrast, BTT obtains a solution in less than a second. Observe that the solution obtained by it for 22 seconds maintains a rather large safety distance between the agents in such a complex setting (see caption of Fig. 2c). Both in **(P2)** and **(P3)** T-RRT*'s transition test leads to an extremely slow exploration of the parameter space. We mention that in these cases RRT* is able to find an initial solution quite rapidly, albeit one of poor quality. Moreover, the convergence rate of RRT* is very slow and hardly manages to improve over the solution obtained initially.

Now we analyze in greater detail the running time of BTT as a function of the number of samples n for the specific scenario in **(P3)**. The running time is dominated by three components: (i) construction of the RTG NN-search data structure; (ii) NN-search calls; (iii) computation

Figure 3: Results of BTT and T-RRT* on **(P1)** for increasing number of curves $2 \leq d \leq 4$, indicated in parenthesis. The black dashed line represents the optimal cost 0.34. In order to display the quality of the solution of BTT as a function of the running time, rather than the number of samples, we ran BTT on an increasing number of samples, with ten runs for each number. For each n we plot the average cost and the average running time. We note that the standard deviation of the running time is very low, and so the average accurately captures the typical running time. For instance, for $d = 3$ the standard deviation ranges between 0.002 and 0.01.

of the cost map, which corresponds to collision detection in standard motion planning. BTT's proportion of running time spent between these components changes with n (see Fig. 4). In particular, as n increases so does the proportion of NN calls. This trend can be observed in other planners as well [26].

7   Discussion and future work

We have introduced BTT for sampling-based bottleneck pathfinding over cost maps. We showed that it manages to cope with complex scenarios of the problem and outperforms T-RRT* in all tests. In addition to its efficiency, BTT requires the tuning of only a single parameter: while the number of samples n cannot be determined a priori, the connection radius r_n specified in Theorem 1 proves to be sufficient for all tested cases. We also provide theoretical justification for this phenomenon in the same theorem. Note that this connection radius is widely used in non-monotone settings of motion planning (see OMPL [45]). Thus, we believe that r_n can be fixed to the aforementioned value. In contrast, T-RRT* requires the tuning of two more parameters that are unique to that planner, as well as those that are inherited from RRT and RRT*.

We do note that BTT might benefit from biasing some of its search towards the target as is the case in RRT-type planners. We intend to pursue this direction in future work. We also mention that since BTT is based on Dijkstra's shortest-path algorithm, it may be applicable to other cost functions such as integral cost and path length, by modifying the edge cost function (lines 12,13) and the underlying discrete path planner (line 8) in Algorithm 1.

Currently the main obstacles to using BTT in more complex settings of bottleneck pathfinding, e.g., higher dimensions, is the high memory consumption of the RTG data structure. In particular, for $n = 10^6$ and $d = 7$ the program may exceed the 4GB limit provided by the OS. Interestingly, in all the tested cases BTT examines only a small portion of the actual vertex set
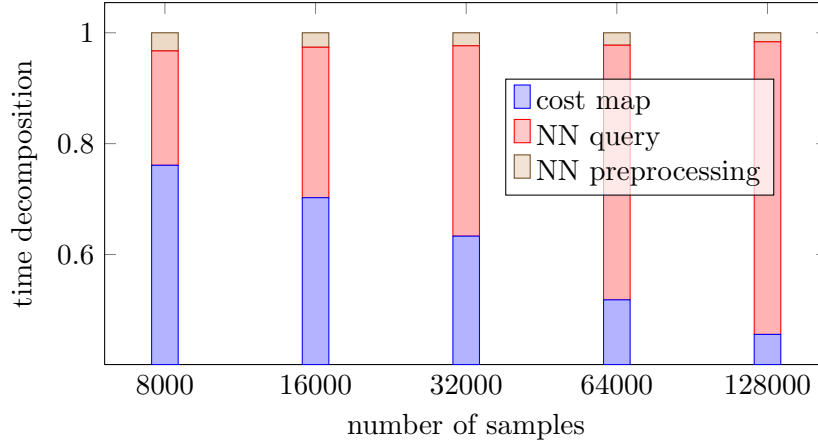
Figure 4: Time percentage for each of the main components of BTT for **(P3)**.

of $\mathcal{G}_n$ (at most 5%). Moreover, the examined vertices are usually grouped together in contiguous regions of the parameter space. This calls for incremental versions of BTT and RTG which generate samples and preprocess them in an online fashion.

On the theoretical side, we aim to refine the statement made in Theorem 1. Currently we can only ensure that a solution that is "almost monotone" exists, but does a "fully-monotone" solution over $\mathcal{G}_n$ exists as well? We conjecture that this statement is true.

References

[1] A. Adler, M. de Berg, D. Halperin, and K. Solovey. Efficient multi-robot motion planning for unlabeled discs in simple polygons. *T-ASE*, 12(4):1309–1317, 2015.

[2] P. K. Agarwal, K. Fox, and O. Salzman. An efficient algorithm for computing high-quality paths amid polygonal obstacles. In *SODA*, pages 1179–1192, 2016.

[3] H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves. *Int. J. Comput. Geometry Appl.*, 5:75–91, 1995.

[4] F. Altché, X. Qian, and A. de La Fortelle. Time-optimal coordination of mobile robots along specified paths. *CoRR*, abs/1603.04610, 2016.

[5] O. Arslan and P. Tsiotras. Use of relaxation methods in sampling-based algorithms for optimal motion planning. In *ICRA*, pages 2421–2428, 2013.

[6] K. Buchin, M. Buchin, and C. Wenk. Computing the Fréchet distance between simple polygons. *Comput. Geom.*, 41(1-2):2–20, 2008.

[7] K. Buchin, M. Buchin, and A. Schulz. Fréchet distance of surfaces: Some simple hard cases. In *ESA*, pages 63–74, 2010.

[8] K. Buchin, M. Buchin, W. Meulemans, and W. Mulzer. Four soviets walk the dog - with an application to Alt's conjecture. In *SODA*, pages 1399–1413, 2014.

[9] K. Buchin, M. Buchin, M. Konzack, W. Mulzer, and A. Schulz. Fine-grained analysis of problems on curves. In *EuroCG*, 2016.

[10] S. Cafieri and N. Durand. Aircraft deconfliction with speed regulation: new models from mixed-integer optimization. *Journal of Global Optimization*, 58(4):613–629, 2014.

[11] E. W. Chambers, É. C. de Verdière, J. Erickson, S. Lazard, F. Lazarus, and S. Thite. Homotopic Fréchet distance between curves or, walking your dog in the woods in polynomial time. *Comput. Geom.*, 43(3):295–311, 2010.

[12] A. F. Cook and C. Wenk. Geodesic Fréchet distance inside a simple polygon. *ACM Transactions on Algorithms*, 7(1):9, 2010.

[13] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms (3. ed.)*. MIT Press, 2009.

[14] P. Dasler and D. M. Mount. On the complexity of an unregulated traffic crossing. In *WADS*, pages 224–235, 2015.

[15] M. de Berg and M. J. van Kreveld. Trekking in the alps without freezing or getting tired. *Algorithmica*, 18(3):306–323, 1997.

[16] D. Devaurs, T. Siméon, and J. Cortés. Optimal path planning in complex cost spaces with sampling-based algorithms. *T-ASE*, 13(2):415–424, 2016.

[17] R. Geraerts and M. Overmars. Creating high-quality paths for motion planning. *I. J. Robotics Res.*, 26(8):845–863, 2007.

[18] S. Har-Peled and B. Raichel. The Fréchet distance revisited and extended. *ACM Transactions on Algorithms*, 10(1):3, 2014.

[19] R. Holladay and S. Srinivasa. Distance metrics and algorithms for task space path optimization. In *IROS*, 2016. to appear.

[20] J. E. Hopcroft, J. T. Schwartz, and M. Sharir. On the complexity of motion planning for multiple independent objects; PSPACE-hardness of the "Warehouseman's problem". *I. J. Robotic Res.*, 3(4):76–88, 1984.

[21] L. Jaillet, J. Cortés, and T. Siméon. Sampling-based path planning on configuration-space costmaps. *Trans. Robotics*, 26(4):635–646, 2010.

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

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

[24] L. E. Kavraki, M. N. Kolountzakis, and J. Latombe. Analysis of probabilistic roadmaps for path planning. In *ICRA*, pages 3020–3025, 1996.

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

[26] M. Kleinbort, O. Salzman, and D. Halperin. Efficient high-quality motion planning by fast all-pairs r-nearest-neighbors. In *ICRA*, pages 2985–2990, 2015.

[27] J. J. Kuffner and S. M. LaValle. RRT-Connect: An efficient approach to single-query path planning. In *ICRA*, pages 995–1001, 2000.

[28] J. S. B. Mitchell. Shortest paths and networks. In *Handbook of Discrete and Computational Geometry, Second Edition.*, pages 607–641. 2004.

[29] M. Muja and D. G. Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. In *VISSAPP*, pages 331–340. INSTICC Press, 2009.

[30] O. Nechushtan, B. Raveh, and D. Halperin. Sampling-diagram automata: A tool for analyzing path quality in tree planners. In *WAFR*, pages 285–301, 2010.

[31] C. Ó'Dúnlaing and C. Yap. A "retraction" method for planning the motion of a disc. *J. Algorithms*, 6(1):104–111, 1985.

[32] F. Pokorny, K. Goldberg, and D. Kragic. Topological trajectory clustering with relative persistent homology. In *ICRA*, pages 16–23, 2016.

[33] B. Raveh, A. Enosh, and D. Halperin. A little more, a lot better: Improving path quality by a path-merging algorithm. *Trans. Robotics*, 27(2):365–371, 2011.

[34] J. H. Reif. Complexity of the movers problem and generalizations: Extended abstract. In *FOCS*, pages 421–427, 1979.

[35] O. Salzman and D. Halperin. Asymptotically near-optimal RRT for fast, high-quality motion planning. *Trans. Robotics*, 32(3):473–483, 2016.

[36] M. Sharir. Algorithmic motion planning. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry, Second Edition.*, pages 1037–1064. Chapman and Hall/CRC, 2004.

[37] T. Simeon, S. Leroy, and J. P. Laumond. Path coordination for multiple mobile robots: a resolution-complete algorithm. *Trans. Robotics*, 18(1):42–49, 2002.

[38] K. Solovey and D. Halperin. On the hardness of unlabeled multi-robot motion planning. In *RSS*, 2015.

[39] K. Solovey and D. Halperin. Sampling-based bottleneck pathfinding with applications to Fréchet matching. In *ESA*, pages 76:1–76:16, 2016.

[40] K. Solovey, J. Yu, O. Zamir, and D. Halperin. Motion planning for unlabeled discs with optimality guarantees. In *RSS*, 2015.

[41] K. Solovey, O. Salzman, and D. Halperin. New perspective on sampling-based motion planning via random geometric graphs. In *RSS*, 2016.

[42] K. Solovey, O. Salzman, and D. Halperin. Finding a needle in an exponential haystack: Discrete RRT for exploration of implicit roadmaps in multi-robot motion planning. *I. J. Robotic Res.*, 35(5):501–513, 2016.

[43] D. Spensieri, R. Bohlin, and J. S. Carlson. Coordination of robot paths for cycle time minimization. In *CASE*, pages 522–527, 2013.

[44] P. G. Spirakis and C.-K. Yap. Strong NP-hardness of moving many discs. *Information Processing Letters*, 19(1):55–59, 1984.

[45] I. A. Şucan, M. Moll, and L. E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation*, 19(4):72–82, 2012. `http://ompl.kavrakilab.org`.

[46] The CGAL Project. *CGAL user and reference manual.* CGAL editorial board, 4.8 edition, 2016. URL `http://www.cgal.org/`.

[47] M. Turpin, N. Michael, and V. Kumar. Concurrent assignment and planning of trajectories for large teams of interchangeable robots. In *ICRA*, pages 842–848, 2013.

[48] J. P. van den Berg and M. H. Overmars. Prioritized motion planning for multiple robots. In *IROS*, pages 430–435, 2005.

[49] C. Voss, M. Moll, and L. E. Kavraki. A heuristic approach to finding diverse short paths. In *ICRA*, pages 4173–4179, 2015.

[50] R. Wein, J. P. van den Berg, and D. Halperin. The visibility-voronoi complex and its applications. *Comput. Geom.*, 36(1):66–87, 2007.

[51] R. Wein, J. P. van den Berg, and D. Halperin. Planning high-quality paths and corridors amidst obstacles. *I. J. Robotic Res.*, 27(11-12):1213–1231, 2008.