

Optimal Sampling-Based Motion Planning under Differential Constraints: the Driftless Case

Edward Schmerling, Lucas Janson, Marco Pavone

Abstract—Motion planning under differential constraints is a classic problem in robotics. To date, the state of the art is represented by sampling-based techniques, with the Rapidly-exploring Random Tree algorithm as a leading example. Yet, the problem is still open in many aspects, including guarantees on the quality of the obtained solution. In this paper we provide a thorough theoretical framework to assess optimality guarantees of sampling-based algorithms for planning under differential constraints. We exploit this framework to design and analyze two novel sampling-based algorithms that are guaranteed to converge, as the number of samples increases, to an optimal solution (namely, the Differential Probabilistic RoadMap algorithm and the Differential Fast Marching Tree algorithm). Our focus is on driftless control-affine dynamical models, which accurately model a large class of robotic systems. In this paper we use the notion of convergence in probability (as opposed to convergence almost surely): the extra mathematical flexibility of this approach yields convergence rate bounds — a first in the field of optimal sampling-based motion planning under differential constraints. Numerical experiments corroborating our theoretical results are presented and discussed.

I. INTRODUCTION

Motion planning is a fundamental problem in robotics. It involves the computation of a sequence of actions that drives a robot from an initial condition to a terminal condition while avoiding obstacles, respecting kinematic/dynamical constraints, and possibly optimizing an objective function [1]. The basic problem, where a robot does not have any constraints on its motion and only an obstacle-free solution is required, is well-understood and solved for a large number of practical scenarios [2]. On the other hand, robots do usually have stringent kinematic/dynamical (in short, differential) constraints on their motion, which in most settings need to be properly taken into account. There are two main approaches [2]: (i) a decoupling approach, in which the problem is decomposed in steps of computing a geometric collision-free path (neglecting the differential constraints), smoothing the path to satisfy the motion constraints, and finally reparameterizing the trajectory so that the robot can execute it, or (ii) a direct approach, in which the differentially-constrained motion planning problem (henceforth referred to as DMP problem) is solved in one shot. The first approach, while fairly common in practice, has several drawbacks, including

the computation of very inefficient trajectories, failure in finding a trajectory due to the decoupling scheme itself, and inflated information requirements [2]. This motivates a quest for efficient algorithms that *directly* solve the DMP problem.

However, directly finding a feasible, let alone optimal, solution to the DMP problem is difficult (note that the basic version without differential constraints is already PSPACE-hard [3, 1], which implies NP-hard). Early work on this topic dates back to more than two decades ago [4], but the problem, especially when optimality is taken into account, is still open in many aspects [5, 2], including algorithms with practical convergence rates, guarantees on the quality of the obtained solution, and class of dynamical systems that can be addressed. To date, the state of the art is represented by sampling-based techniques, where an *explicit* construction of the configuration space is avoided and the configuration space is probabilistically “probed” with a sampling scheme.

Arguably, the most successful algorithm for DMP is the sampling-based rapidly-exploring tree algorithm (RRT) [6], which incrementally builds a tree of trajectories by randomly sampling points in the configuration space. Lately, several variations of the RRT algorithm, referred to as RRT*, stemming from [7] and its kinodynamic extension [8], have been considered to ensure that the cost of the computed trajectory converges to the optimal cost as the number of sampled points goes to infinity [8, 9, 10, 11, 12]. These works, while providing strong experimental validation, only provide proof sketches that do not fully address many of the complications that arise in extending asymptotic optimality arguments from the geometric case to differentially constrained paths. For example, rewiring the RRT tree within a local volume containing (in expectation) a log fraction of previous samples is not sufficient in itself to claim optimality, as in [10, 11]. Additional assumptions on trajectory approximability must be stated and verified for the differential constraints in question. Such requirements are discussed in [8], but it is not clear how assuming the existence of forward-reachable trajectory approximations is sufficient for a “ball-to-ball” proof technique that requires backward approximations as well. A different approach to asymptotically optimal planning has recently been proposed by STABLE SPARSE RRT which achieves optimality through random control propagation instead of connecting existing samples using a steering subroutine [13]. This paper, like the RRT* variations, follows a steering approach, although it may be considered less general, as it is our view that leveraging as much knowledge as possible of the differential constraints while planning is necessary in order to have a hope of planning in real-time.

Edward Schmerling is with the Institute for Computational & Mathematical Engineering, Stanford University, Stanford, CA 94305, schmrlng@stanford.edu.

Lucas Janson is with the Department of Statistics, Stanford University, Stanford, CA 94305, ljanson@stanford.edu.

Marco Pavone is with the Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305, pavone@stanford.edu.

This work was supported by an Early Career Faculty grant from NASA’s Space Technology Research Grants Program (Grant NNX12AQ43G).

Statement of Contributions: The objective of this paper is to provide a theoretical framework to study optimality guarantees of sampling-based algorithms for the DMP problem, and to exploit this framework to design efficient sampling-based algorithms that are guaranteed to asymptotically converge to an optimal solution. The focus of this paper is on driftless control-affine (DCA) dynamical systems of the form

$$\dot{x}(t) = \sum_{i=1}^m g_i(x(t))u_i(t),$$

where the available motions of trajectories $x(t)$ are linear combinations given by input control functions $u_i(t)$ and their corresponding actions at each point in space $g_i(x)$. Our work is written from the perspective of optimizing trajectory arc length, but applies also to cost metrics satisfying similar properties, which we discuss. This model is often the result of nonholonomic constraints that the kinematic variables of the system must satisfy [14]. A large class of robotic systems can be modeled as DCA, including mobile robots with wheels that roll without slipping [1], multi-fingered robotic hands [15], rigid bodies with zero angular momentum undergoing re-orientation [16], and systems with nonholonomic actuators [14]. The DCA model, however, rules out the presence of dynamics with drift, which in many important cases (e.g., spacecraft control) can not be neglected. Specifically, the contribution of this paper is threefold. First, we show that *any* trajectory of a DCA system may be “traced” arbitrarily well by connecting randomly distributed points from a sufficiently large sample set covering the configuration space. We will refer to this property as *probabilistic exhaustivity*, as opposed to probabilistic completeness [1], where the requirement is that *at least* one trajectory is traced with a sufficiently large sample set. Probabilistic exhaustivity is a key tool in proving asymptotic optimality, and is of independent interest. Second, we introduce two novel sampling-based algorithms for the solution to the DMP with DCA systems, namely the Differential Probabilistic RoadMap algorithm (DPRM*) and the Differential Fast Marching Trees algorithm (DFMT*). Third, by leveraging the property of probabilistic exhaustivity for DCA systems, we rigorously show that both DPRM* and DFMT* are asymptotically optimal. We note that in this paper we use the notion of convergence in probability (as opposed to convergence almost surely, as in, e.g., [11]): the extra mathematical flexibility of this approach yields convergence rate bounds — a first in the field of optimal sampling-based motion planning under differential constraints. Our approach is inspired by [11] and [17].

Organization: This paper is structured as follows. In Section II we provide a review of some key results in differential geometry that will be extensively used in the paper. In Section III we formally define the problem we wish to solve, while in Section IV we prove the aforementioned probabilistic exhaustivity property for DCA systems. In Section V we present the DPRM* and DFMT* algorithms, and in Section VI we prove their asymptotic optimality (together with a convergence rate characterization). In Section VII we

provide implementation details for the proposed algorithms, and we study them via numerical experiments. Finally, in Section VIII, we draw some conclusions and we discuss directions for future work.

II. BACKGROUND MATERIAL

In this section we provide some definitions and a brief review of key results in differential geometry, on which we will rely extensively later in the paper. Let $\mathcal{M} \subset \mathbb{R}^n$ be the manifold defining a configuration space. Within this space let us consider driftless control-affine (DCA) dynamical systems of the form

$$\dot{x}(t) = \sum_{i=1}^m g_i(x(t))u_i(t), \quad x \in \mathcal{M}, u \in U, \quad (1)$$

where the available motions of trajectories $x(t)$ are linear combinations given by input control functions $u_i(t)$ and their corresponding actions at each point in space $g_i(x)$. We shall assume in this paper that g_1, \dots, g_m are smooth vector fields on \mathcal{M} , and that the control set $U \subset \mathbb{R}^m$ is closed and bounded. We also assume U is symmetric about the origin so that the system is time-reversible and 0 is in the interior of the convex hull of U . This last condition ensures that the local possibilities for motion at each point appear as a linear space spanned by the g_i , a fact essential to the forthcoming controllability discussion [18]. We denote the driftless control-affine system of equation (1) as $\Sigma = (\mathcal{M}, g, U)$. A function $x : [0, T] \rightarrow \mathcal{M}$ is called a *dynamically feasible* trajectory, alternatively path, if there exists a corresponding control function $u(t)$ with which it satisfies Σ . All trajectories discussed in this paper are dynamically feasible unless otherwise noted.

A. Arc Length and Sub-Riemannian Distance

The *arc length* of a path $x(t)$ is defined as

$$\ell(x) := \int_0^T \|\dot{x}(t)\| dt,$$

where $\|\dot{x}(t)\| = \|\dot{x}(t)\|_2 = \sqrt{\langle \dot{x}(t), \dot{x}(t) \rangle}$ is computed using the standard Euclidean inner product on the tangent spaces of \mathcal{M} . The arc length function induces a *sub-Riemannian distance* d on \mathcal{M} , defined for $x_1, x_2 \in \mathcal{M}$ as $d(x_1, x_2) := \inf_x \ell(x)$, where the infimum is taken over dynamically feasible trajectories $x(t)$ connecting x_1 and x_2 . Note that for driftless control-affine systems, time-reversibility implies that d is symmetric and indeed a metric. The *sub-Riemannian ball* may be defined in analogy to the standard Euclidean ball (i.e., $B^e(x, \varepsilon) = \{y \in \mathcal{M} : \|x - y\| \leq \varepsilon\}$) according to $B(x, \varepsilon) := \{y \in \mathcal{M} : d(x, y) \leq \varepsilon\}$. Note that by definition $B(x, \varepsilon) \subset B^e(x, \varepsilon)$.

B. Controllability and Reachable Sets

We now present a series of results regarding system controllability following the discussion in [19] and [20]. As noted above, the vector fields g_1, \dots, g_m characterizing the system Σ represent a set of possible motions for a trajectory within \mathcal{M} . More precisely at each point $p \in \mathcal{M}$ the vectors

$g_1(p), \dots, g_m(p)$ span a linear subspace of local directions within the tangent space $T_p\mathcal{M}$. For a vector field Y on \mathcal{M} let $\Phi_{Y,t} : \mathcal{M} \rightarrow \mathcal{M}$ denote its local flow, the function that maps an initial state to the state obtained by following the vector field Y for time t . That is, $\Phi_{Y,t}(p) := y(t)$ where $y : [0, t] \rightarrow \mathcal{M}$ is a solution to the initial value problem $\dot{y}(\tau) = Y(y(\tau))$, $y(0) = p$. Commutators of flows, akin to control switching in (1), allow for local motions transverse to the g_i to be achieved while satisfying the differential constraints. Given two vector fields Y, Z on \mathcal{M} and a starting point $p \in \mathcal{M}$, we have the approximation for small t :

$$p + t^2[Y, Z](p) = \Phi_{Y,t} \circ \Phi_{Z,t} \circ \Phi_{Y,-t} \circ \Phi_{Z,-t}(p) + O(t^3),$$

where the Lie bracket $[Y, Z]$ is a third vector field which may be computed with respect to a coordinate system as $[Y, Z] = J_Z X - J_X Z$, where J_Y and J_Z are the Jacobian matrices of Y and Z respectively. Computing Lie brackets allows us to characterize all directions that are possible from each p , in addition to those given by $g_1(p), \dots, g_m(p)$.

Let $\mathcal{L} = \mathcal{L}(g_1, \dots, g_m)$ be the distribution, equivalently the set of local vector subspaces, generated by the vector fields g_1, \dots, g_m . We define recursively $\mathcal{L}^1 = \mathcal{L}$,

$$]\mathcal{L}^{k+1} = \mathcal{L}^k + [\mathcal{L}^1, \mathcal{L}^k]$$

where $[\mathcal{L}^1, \mathcal{L}^k] = \text{Span}\{[Y, Z] : Y \in \mathcal{L}^1, Z \in \mathcal{L}^k\}$. Then \mathcal{L}^k is the distribution generated by the iterated Lie brackets of the g_i with k terms or fewer. The Lie hull of \mathcal{L} is $\text{Lie}(\mathcal{L}) := \cup_{k \geq 1} \mathcal{L}^k$. Let $L^k(x)$ denote the vector space corresponding to $x \in \mathcal{M}$ in \mathcal{L}^k .

The vector fields g_1, \dots, g_m are said to be *bracket generating* if $\text{Lie}(\mathcal{L})(x) = T_x\mathcal{M}$ for all $x \in \mathcal{M}$. This requirement is also referred to as Chow's condition, or the linear algebra rank condition, and means that arbitrary local motion may be achieved by composing motions along the control directions g_i . In fact, provided that the g_i are bracket generating, any two points $x_1, x_2 \in \mathcal{M}$ may be connected by a trajectory satisfying Σ , that is $d(x_1, x_2) < \infty$. The remainder of this section develops tighter bounds on $d(x_1, x_2)$ that will be used in our asymptotic optimality proofs of planning algorithms.

Chow's condition implies that for all $x \in \mathcal{M}$, there exists a smallest integer $s = s(x)$ such that $L^{s(x)} = T_x\mathcal{M}$. Indeed, we have $L^1(x) \subset L^2(x) \subset \dots \subset L^{s(x)}(x) = T_x\mathcal{M}$. Set $n_k(x) = \dim L^k(x)$. The integer list $(n_1(x), \dots, n_{s(x)}(x))$ is called the growth vector of \mathcal{L} at x . A point x is called a *regular point* if there exists an open neighborhood of \mathcal{M} around x such that the growth vector is constant; otherwise x is said to be a *singular point*.

We now further assume that every $x \in \mathcal{M}$ is a regular point, so that the growth vector $(n_1, \dots, n_s = n)$ is constant over the whole configuration manifold. Fix a base point $x_0 \in \mathcal{M}$. Using the bracket-generating assumption we select a local orthonormal frame for $T_{x_0}\mathcal{M}$ of vector fields Y_1, \dots, Y_n as follows: the set $\{Y_1 = g_1, \dots, Y_{n_1} = g_m\}$ spans \mathcal{L} near x_0 ; $\{Y_1, \dots, Y_{n_2}\}$ spans \mathcal{L}^2 near x_0 ; $\{Y_1, \dots, Y_{n_3}\}$ spans \mathcal{L}^3 near x_0 ; and so on. Define the weights $w_i = k$ if $Y_i(x_0) \in \mathcal{L}^k(x_0)$ and $Y_i(x_0) \notin \mathcal{L}^{k+1}(x_0)$. Applying a procedure developed in [19], the coordinate system y_i

corresponding to this local frame may be transformed into a *privileged* coordinate system z_i by a polynomial change of coordinates of the form

$$\begin{aligned} z_1 &= y_1, \\ z_2 &= y_2 + \text{poly}_2(y_1), \\ z_3 &= y_3 + \text{poly}_3(y_1, y_2), \\ &\vdots \\ z_n &= y_n + \text{poly}_n(y_1, \dots, y_{n-1}), \end{aligned} \tag{2}$$

where $\text{poly}_i(\cdot)$, $i = 2, \dots, n$, denotes a polynomial function that includes only terms of degree ≥ 2 and $< w_i$. From the triangular structure of (2), it is clear that the inverse transformation $y = z + \text{poly}'(z)$ is of the same form. Given privileged coordinates z_i , define the *pseudonorm* at x_0 as

$$\|z\|_{x_0} := \max\{|z_1|^{1/w_1}, \dots, |z_n|^{1/w_n}\}.$$

Using this pseudonorm we define the w -*weighted box* of size ε at x_0 as the point set $\text{Box}^w(\varepsilon) := \{z \in \mathbb{R}^n : \|z\|_{x_0} \leq \varepsilon\}$. We use the notation $\text{Box}^w(x_0, \varepsilon)$ for the corresponding locus of points in \mathcal{M} given by the coordinates z_i .

Theorem II.1 (Ball-box theorem [19]). *Fix a point $x_0 \in \mathcal{M}$ and a system of privileged coordinates z_1, \dots, z_n at x_0 . Then there exist positive constants $a(x_0), A(x_0) > 0$, and $\sigma(x_0) > 0$ such that for all x with $d(x_0, x) < \sigma(x_0)$,*

$$a(x_0)\|z(x)\|_{x_0} \leq d(x_0, x) \leq A(x_0)\|z(x)\|_{x_0}. \tag{3}$$

Constructing the coordinate system z thus gives structure to how sub-Riemannian distance behaves locally; this structure may be used for steering, e.g. [21]. It can be shown that there exists a continuously varying system of privileged coordinates on \mathcal{M} so that the inequality (3) holds at all x_0 for continuous positive functions $a(\cdot), A(\cdot)$, and $\sigma(\cdot)$ on \mathcal{M} . Let us assume that the system Σ is sufficiently regular such that there exist bounds $0 < a_{\min} \leq a(x) \leq A(x) \leq A_{\max} < \infty$ and $\sigma(x) \geq \sigma_{\min} > 0$ for all $x \in \mathcal{M}$. We state a pair of lemmas (whose proofs are provided in the Appendix) concerning how privileged coordinates relate to the Euclidean notions of volume and distance.

Lemma II.2 (Box volume). *Fix $x \in \mathcal{M}$. The volume of $\text{Box}^w(x, r)$ is given by $\mu(\text{Box}^w(x, r)) = r^D$ where $D = \sum_{i=1}^n w_i$.*

Lemma II.3 (Distance comparison). *Fix a point $x_0 \in \mathcal{M}$ and a system of privileged coordinates z_1, \dots, z_n at x_0 . Then there exists a positive constant $\theta(x_0) > 0$ such that for all x with $\|x_0 - x\| \leq \theta(x_0)$,*

$$\|x_0 - x\| \leq d(x_0, x) \leq 2A_{\max}\|x_0 - x\|^{1/s}.$$

III. PROBLEM FORMULATION

Let $\mathcal{M} \subset \mathbb{R}^n$ be the manifold defining a configuration space of a robotic system. Let $\mathcal{M}_{\text{obs}} \subset \mathcal{M}$ be the obstacle region, such that $\mathcal{M} \setminus \mathcal{M}_{\text{obs}}$ is an open set, and denote the obstacle-free space as $\mathcal{M}_{\text{free}} = \text{cl}(\mathcal{M} \setminus \mathcal{M}_{\text{obs}})$. The starting configuration x_{init} is an element of $\mathcal{M}_{\text{free}}$, and the goal region

$\mathcal{M}_{\text{goal}}$ is an open subset of $\mathcal{M}_{\text{free}}$. The trajectory planning problem is denoted by the tuple $(\Sigma, \mathcal{M}_{\text{free}}, x_{\text{init}}, \mathcal{M}_{\text{goal}})$, where Σ , as discussed in Section II, denotes a driftless control-affine system. A dynamically feasible trajectory x is *collision-free* if $x(t) \in \mathcal{M}_{\text{free}}$ for all $t \in [0, T]$. A trajectory x is said to be *feasible* for the trajectory planning problem $(\Sigma, \mathcal{M}_{\text{free}}, x_{\text{init}}, \mathcal{M}_{\text{goal}})$ if it is dynamically feasible, collision-free, $x(0) = x_{\text{init}}$, and $x(T) \in \text{cl}(\mathcal{M}_{\text{goal}})$.

Let \mathcal{X} be the set of all feasible paths. A *cost function* for $(\Sigma, \mathcal{M}_{\text{free}}, x_{\text{init}}, \mathcal{M}_{\text{goal}})$ is a function $c : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ from the set of paths to the nonnegative real numbers; in this paper we consider the cost function $c(x) = \ell(x)$ defined as the arc length of x with respect to the Euclidean metric in \mathcal{M} . The objective is to find the feasible path with minimum associated cost; this minimum is achieved as long as U is closed and bounded [1]. The optimal trajectory planning problem is then defined as follows:

Optimal motion planning for driftless systems: Given a trajectory planning problem $(\Sigma, \mathcal{M}_{\text{free}}, x_{\text{init}}, \mathcal{M}_{\text{goal}})$ and an arc length function $c : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$, find a feasible path x^* such that $c(x^*) = \min\{c(x) : x \text{ is feasible}\}$. If no such path exists, report failure.

Our analysis will rely on two key sets of assumptions, relating, respectively, to the underlying system Σ and the problem-specific parameters $\mathcal{M}_{\text{free}}, x_{\text{init}}, \mathcal{M}_{\text{goal}}$.

1) *Assumptions on system:* As in Section II, we require from Σ that a) the vector fields g_1, \dots, g_m are bracket generating, b) the configuration space \mathcal{M} contains only regular points, c) there exist constants $a_{\min}, A_{\max}, \sigma_{\min}$ such that Theorem II.1 holds with these values at all $x_0 \in \mathcal{M}$, and d) there exists a bounding constant θ_{\min} such that $0 < \theta_{\min} < \theta(x_0)$ for all $x_0 \in \mathcal{M}$, where $\theta(x_0)$ is as defined in Lemma II.3. Assumption (a) is a basic requirement for the system to be controllable, let alone optimally controlled, while (b), (c), and (d) ensure that there are no extreme regions of the configuration space which would require an unbounded sample density to capture their geometry. These assumptions will be collectively referred to as A_{Σ} .

2) *Assumptions on problem parameters:* We require that the goal region $\mathcal{M}_{\text{goal}}$ has *regular boundary*, that is there exists $\xi > 0$ such that $\forall y \in \partial\mathcal{M}_{\text{goal}}$, there exists $z \in \mathcal{M}_{\text{goal}}$ with $B^e(z, \xi) \subseteq \mathcal{M}_{\text{goal}}$ and $y \in \partial B^e(z, \xi)$. This requirement that the boundary of the goal region has bounded curvature ensures that a point sampling procedure may expect to select points in the goal region near any point on its boundary.

We also make requirements on the *clearance* of a trajectory, i.e., its “distance” from \mathcal{M}_{obs} , standard for sampling-based methods [22]. For a given $\delta > 0$, the δ -interior of $\mathcal{M}_{\text{free}}$ is defined as the set of all states that are at least a Euclidean distance δ away from any point in \mathcal{M}_{obs} . A collision-free path x is said to have strong δ -clearance if it lies entirely inside the δ -interior of $\mathcal{M}_{\text{free}}$. A collision-free path x is said to have weak δ -clearance if there exists a path x' that has strong δ -clearance and there exists a homotopy ψ , with $\psi(0) = x$ and $\psi(1) = x'$ that satisfies the following three properties: (a) $\psi(\alpha)$ is a dynamically feasible path

for all $\alpha \in (0, 1]$, (b) $\lim_{\alpha \rightarrow 0} c(\psi(\alpha)) = c(x)$, and (c) for all $\alpha \in (0, 1]$, $\psi(\alpha)$ has strong δ_{α} -clearance for some $\delta_{\alpha} > 0$. Properties (a) and (b) are required since pathological obstacle sets may be constructed that squeeze all optimum-approximating homotopies into undesirable motion. In practice, however, as long as $\mathcal{M}_{\text{free}}$ does not contain any passages of infinitesimal width, the fact that Σ is bracket-generating will allow every trajectory to be weak δ -clear.

IV. PROBABILISTIC EXHAUSTIVITY OF SAMPLING SCHEMES UNDER DRIFTLESS CONSTRAINTS

In this section we prove a key result characterizing random sampling schemes for motion planning under driftless differential constraints: any feasible trajectory through the configuration space \mathcal{M} is “traced” arbitrarily well by connecting randomly distributed points from a sufficiently large sample set covering the configuration space. We will refer to this property as *probabilistic exhaustivity*. Note that this notion is much stronger than the usual notion of probabilistic completeness in motion planning, where the requirement is that *at least one* feasible trajectory is traced. The notion of probabilistic exhaustivity, besides being a result of independent interest, is a strong tool in proving asymptotic optimality of sampling-based motion planning algorithms, as will be shown in Section V (specifically, we focus on differential variants of PRM* [22] and FMT* [17]).

Let $x : [0, T] \rightarrow \mathcal{M}$ satisfy the system (1). Given a set of waypoints $\{y_m\}_{m=1}^M \subset \mathcal{M}$, we associate a dynamically feasible trajectory $y^* : [0, S] \rightarrow \mathcal{M}$ that connects the nodes y_1, \dots, y_M in order so that each connection is locally optimal, i.e. each path segment connecting y_m to y_{m+1} has length $d(y_m, y_{m+1})$. We consider the waypoints $\{y_m\}$ to (ε, r) -trace the trajectory x if: a) $d(y_m, y_{m+1}) \leq r$ for all m , b) the cost of y^* is bounded as $c(y^*) \leq (1 + \varepsilon)c(x)$, and c) the distance from any point of y^* to x is no more than r , i.e. $\min_{t \in [0, T]} d(y(s), x(t)) \leq r$ for all $s \in [0, S]$. In the context of sampling-based motion planning, we may expect to find closely tracing $\{y_m\}$ as a subset of the sampled points, provided the sample size is large. We formalize this notion in Theorem IV.5, the proof of which requires three technical lemmas (the proofs of these lemmas are provided in the Appendix).

Lemma IV.1. *Let x be a dynamically feasible trajectory and consider a partition of the time interval $0 = \tau_1 < \tau_2 < \dots < \tau_M = T$. Suppose that $\{y_m\} \subset \mathcal{M}$ satisfy a) $y_m \in B(x(\tau_m), \rho)$ for all $m \in \{1, \dots, M\}$, and b) more than a $(1 - \alpha)$ fraction of the y_m satisfy $y_m \in B(x(\tau_m), \beta\rho)$ for a parameter $\beta \in (0, 1)$. Then the cost $c(y^*)$ of the trajectory y^* sequentially connecting the nodes y_1, \dots, y_M is upper bounded as*

$$c(y^*) \leq c(x) + 2M\rho(\beta + \alpha - \alpha\beta).$$

Remark IV.2. *If we further assume that $y_1 = x(0)$, then the bound $c(y^*) \leq c(x) + 2(M - 1)\rho(\beta + \alpha - \alpha\beta)$ holds.*

Let $\text{SampleFree}(n)$ denote a set of n points sampled independently and identically from the uniform distribution

on $\mathcal{M}_{\text{free}}$.

Lemma IV.3. Fix $n \in \mathbb{N}$, $\alpha \in (0, 1)$, and let S_1, \dots, S_M be disjoint subsets of $\mathcal{M}_{\text{free}}$ with

$$\mu(S_m) = \mu(S_1) \geq \left(\frac{2 + \log(1/\alpha)}{n} \right) e^2 \mu(\mathcal{M}_{\text{free}}),$$

for each m . Let $V = \text{SampleFree}(n)$ and define

$$K_n := \#\{m \in \{1, \dots, M\} : S_m \cap V = \emptyset\}.$$

Then $\mathbb{P}(K_n \geq \alpha M) \leq \frac{e^{-\alpha M}}{1 - e^{-\alpha}}$.

Lemma IV.4. Fix $n \in \mathbb{N}$ and let T_1, \dots, T_M be subsets of $\mathcal{M}_{\text{free}}$, possibly overlapping, with

$$\mu(T_m) = \mu(T_1) \geq \kappa (\log n / n) \mu(\mathcal{M}_{\text{free}})$$

for each m and some constant $\kappa > 0$. Let $V = \text{SampleFree}(n)$ and denote by E_m the event that $T_m \cap V = \emptyset$ for each m . Then

$$\mathbb{P}\left(\bigvee_{m=1}^M E_m\right) \leq Mn^{-\kappa}.$$

Before stating the theorem and proof in full, we sketch our approach for proving probabilistic exhaustivity. Given a path to be traced with waypoints from a sample set, we tile the span of the path with two sequences of concentric sub-Riemannian balls – a sequence of “small” balls and a sequence of “large” balls. With high probability, all but a tiny α fraction of the small balls will contain a point from the sample set (Lemma IV.3), and for any small balls that don’t contain such a point we ensure that the concentric large ball does (Lemma IV.4). We take these points as a sequence of waypoints which tightly follows the reference path with few exceptions, and never has a gap over any section of the reference path when it deviates. We then use the metric inequality for d to bound the total cost of the waypoint trajectory (Lemma IV.1).

Theorem IV.5 (Probabilistic exhaustivity). Let Σ be a DCA system satisfying the assumptions A_Σ and suppose $x : [0, T] \rightarrow \mathcal{M}_{\text{free}}$ is a dynamically feasible trajectory with strong δ -clearance, $\delta > 0$. Let $V = \{x_{\text{init}}\} \cup \text{SampleFree}(n)$, $\varepsilon > 0$, and for fixed n consider the event A_n that there exist $\{y_m\}_{m=1}^M \subset V$ which (ε, r_n) -trace x , where

$$r_n = 4A_{\max}(1 + \eta)^{1/D} \left(\frac{\mu(\mathcal{M}_{\text{free}})}{D} \right)^{1/D} \left(\frac{\log n}{n} \right)^{1/D}$$

for a parameter $\eta \geq 0$. Then, as $n \rightarrow \infty$, the probability that A_n does not occur is asymptotically bounded as $\mathbb{P}(A_n^c) = O(n^{-\eta} \log^{-1/D} n)$.

Proof. Note that in the case $c(x) = 0$ we may pick $y_1 = x(0)$ to be the only waypoint and the result is trivial. Therefore assume $c(x) > 0$. Fix n sufficiently large so that $r_n/2 \leq \sigma_{\min}$. Take $x(\tau_{n,m})$ to be points spaced along x at sub-Riemannian distances $r_n/2$; more precisely let $\tau_1 = 0$, and for $m = 2, 3, \dots$ consider

$$\tau_{n,m} = \min(\{\tau \in (\tau_{n,m-1}, 1) : d(x(\tau), x(\tau_{n,m-1})) \geq r_n/2\}).$$

Let M_n be the first m for which the set is empty; take $\tau_{n,M_n} = T$. Note that since the distance d is an infimum over all feasible trajectories, one of which is the segment of x between $\tau_{n,m}$ and $\tau_{n,m+1}$, we have the bound $M_n \leq \lceil 2c(x)/r_n \rceil$.

We now make the identification $\alpha = \beta = \varepsilon/2$ anticipating the application of Lemma IV.1. Take $\rho_n = r_n/4$ and define a sequence of sub-Riemannian balls $B_{n,1}, \dots, B_{n,M_n}$ centered along the trajectory x by $B_{n,m} = B(x_m, \rho_n)$, where $x_m = x(\tau_{n,m})$ for $m \in \{1, \dots, M_n\}$. Within these balls define a concentric sequence of smaller, disjoint balls $B_{n,m}^\beta = B(x_m, \beta\rho_n)$ for each $m \in \{1, \dots, M_n\}$. Note that since $\text{Box}^w(x_m, \rho_n/A_{\max}) \subset B_{n,m}$, we have the volume lower bound $\mu(B_{n,m}) \geq (\rho_n/A_{\max})^D$ and similarly $\mu(B_{n,m}^\beta) \geq (\beta\rho_n/A_{\max})^D$. Denote $K_n^\beta := \#\{m \in \{1, \dots, M_n\} : B_{n,m}^\beta \cap V = \emptyset\}$. We consider the event \tilde{A}_n that every large ball $B_{n,m}$, as well as at least a $(1 - \alpha)$ fraction of the small balls $B_{n,m}^\beta$, contains at least one point of V :

$$\tilde{A}_n = \{K_n^\beta < \alpha M_n\} \wedge \bigwedge_{m=1}^{M_n} \{B_{n,m} \cap V \neq \emptyset\}.$$

We claim that \tilde{A}_n implies the event A_n that there exist (ε, r_n) -tracing $\{y_m\} \subset V$. If \tilde{A}_n holds, then we select waypoints $\{y_m\}_{m=1}^{M_n} \subset V$ such that $y_m \in B_{n,m}$ for every point, as well as $y_m \in B_{n,m}^\beta$ for at least a $(1 - \alpha)$ fraction of the points. In particular let us select $y_1 = x(0)$.

First note that $y_m \in B_{n,m}$ for each m implies

$$\begin{aligned} d(y_m, y_{m+1}) &\leq d(y_m, x_m) + d(x_m, x_{m+1}) \\ &\quad + d(x_{m+1}, y_{m+1}) \\ &\leq r_n/4 + r_n/2 + r_n/4 = r_n. \end{aligned}$$

Next, applying Remark IV.2 we have

$$\begin{aligned} c(y^*) &\leq c(x) + 2(M_n - 1)\rho_n(\beta + \alpha - \alpha\beta) \\ &\leq c(x) + 2(2c(x)/r_n)(r_n/4)\varepsilon \\ &\leq (1 + \varepsilon)c(x). \end{aligned}$$

Finally we check that any point $y^*(t)$ is within distance r_n from a point on x . To see this suppose that $y^*(t)$ lies on the shortest path connecting y_m to y_{m+1} . Note that

$$d(y_m, y^*(t)) + d(y^*(t), y_{m+1}) = d(y_m, y_{m+1}) \leq r_n.$$

Then we may write

$$\begin{aligned} d(x_m, y^*(t)) + d(x_{m+1}, y^*(t)) &\leq d(x_m, y_m) \\ &\quad + d(x_{m+1}, y_{m+1}) + r_n \\ &\leq (3/2)r_n, \end{aligned}$$

so that $\min\{d(x_m, y^*(t)), d(x_{m+1}, y^*(t))\} \leq (3/4)r_n < r_n$. Thus the waypoints $\{y_m\}$ (ε, r_n) -trace x .

Now making the identifications $S_m = B_{n,m}^\beta$ and $T_m = B_{n,m}$ in Lemmas IV.3 and IV.4 respectively, we see that by considering $n \geq N_1$ sufficiently large so that $\log N_1 \geq$

$D\beta^{-D}(2 + \log(1/\alpha))e^2$ (satisfying the volume assumption of Lemma IV.3), we compute the union bound

$$\begin{aligned} \mathbb{P}(A_n^c) &\leq \mathbb{P}(\tilde{A}_n^c) \\ &\leq \mathbb{P}(K_n^\beta \geq \alpha M_n) + \mathbb{P}\left(\bigvee_{m=1}^{M_n} \{B_{n,m} \cap V = \emptyset\}\right) \\ &\leq \frac{e^{-\alpha M_n}}{1 - e^{-n}} + M_n n^{-(1+\eta)/D}. \end{aligned}$$

Now, $c(x) > 0$ and $r_n = \Theta((\log n/n)^{1/D})$ together imply that $M_n = \Theta((n/\log n)^{1/D})$. The second term in the bound dominates as $n \rightarrow \infty$, and $\mathbb{P}(A_n^c) = O(n^{-\eta} \log^{-1/D} n)$. \square

V. OPTIMAL SAMPLING-BASED ALGORITHMS FOR DRIFTLESS CONTROL-AFFINE SYSTEMS

In this section we present two algorithms for the motion planning problem with driftless control-affine systems. The first algorithm, named the Differential Probabilistic RoadMap algorithm (DPRM*), is a derivation of the PRM* algorithm presented in [22], while the second algorithm, named the Differential Fast Marching Tree algorithm (DFMT*), is a derivation of the FMT* algorithm presented in [17]. This section provides a description of both algorithms, while the next section focuses on their theoretical characterization (chiefly, their asymptotic optimality property).

As in Section IV, let $\text{SampleFree}(k)$ be a function that returns a set of $k \in \mathbb{N}$ states sampled independently and identically from the uniform distribution on $\mathcal{M}_{\text{free}}$. These sampled states are connected as vertices in a graph from which a solution trajectory will be computed. Given two vertices x_1 and x_2 , we denote with the edge (x_1, x_2) an optimal cost trajectory from x_1 to x_2 neglecting obstacle constraints. Let $\text{CollisionFree}(x_1, x_2)$ denote the boolean function which returns true if and only if the edge (x_1, x_2) does not intersect an obstacle. Given a set of vertices V , a state $x \in \mathcal{M}$, and a threshold $r > 0$, let $\text{Near}(V, x, r)$ be a function that returns the set of states $\{v \in V : \|v\|_x < r/a_{\min}\}$. Given a graph $G = (V, E)$, where V is the vertex set and E is the edge set, and a vertex $x \in V$, let $\text{Cost}(x, G)$ be the function that returns the cost of the shortest path in the graph G between the vertices x_{init} and x . Let $\text{Path}(x, G)$ be the function that returns the path achieving that cost.

The DPRM* algorithm is given in Algorithm 1, while the DFMT* algorithm is given in Algorithm 2. DPRM* works by sampling a set of points within $\mathcal{M}_{\text{free}}$ and connecting each state to every state in a local neighborhood around it, provided that the connection is collision free. The resulting graph spans $\mathcal{M}_{\text{free}}$, with the local connections combining to yield a global “roadmap” for traveling between any two states, not just x_{init} and the goal. The least cost path in the graph between x_{init} and $\mathcal{M}_{\text{goal}}$, computed, e.g., using the Dijkstra’s algorithm, is output by DPRM*.

The DFMT* algorithm essentially implements a streamlined version of DPRM* by performing a “lazy” dynamic programming recursion, this time during the state connection phase instead of as a last step, to grow a *tree* of trajectories

Algorithm 1 Differential Probabilistic RoadMap (DPRM*)

```

1  $V \leftarrow \{x_{\text{init}}\} \cup \text{SampleFree}(n)$ ;  $E \leftarrow \emptyset$ 
2 for each  $v \in V$  do
3   for each  $u \in \text{Near}(V \setminus \{v\}, v, r_n)$  do
4     if  $\text{CollisionFree}(u, v)$  then
5        $E \leftarrow E \cup \{(u, v)\}$ 
6     end if
7   end for
8 end for
9  $V^* \leftarrow \text{argmin}_{v \in V \cap \mathcal{M}_{\text{goal}}} \text{Cost}(v, G = (V, E))$ 
10 if  $V^* \neq \emptyset$  then
11    $v^* \leftarrow \text{random vertex in } V^*$ 
12   return  $\text{Path}(v^*, G = (V, E))$ 
13 else
14   return Failure
15 end if

```

which moves steadily outward in cost-to-come space. In the Algorithm 2 outline, the set W consists of all of the nodes that have not yet been added into the tree, while H is comprised only of nodes that are in the tree. In particular, while H keeps track of nodes which have already been added to the tree, nodes are removed from H if they are not near enough to the edge of the expanding tree to actually have any new connections made with W (see [17] for further details). At each iteration, DFMT* examines the neighborhood of a state in H and *only* considers locally-optimal (assuming no obstacles) connections for potential inclusion in the tree (see line (10)). By only checking for collision on the locally-optimal (assuming no obstacles) connection, as opposed to every possible connection (essentially what is done in DPRM*), DFMT* saves a large (indeed unbounded as the number of vertices increases) number of collision-check computations. A more detailed explanation of the meaning of the dual sets H and W and of the philosophy behind DFMT* can be found in [17].

The main differences of these algorithms with respect to their geometric counterparts (i.e., PRM* and FMT*) are that (i) near vertices lie within the privileged coordinate box $\text{Box}^w(x, r/a_{\min})$ rather than within the Euclidean ball, and (ii) the edges connecting vertices are optimal *trajectories* rather than straight lines. Indeed, attempting connections within the sub-Riemannian ball $B(x, r)$ would be the best analogy to the geometric planning case, and is preferable for efficiency if possible. The structure of this set is difficult to compute exactly, however, which motivates the choice $\text{Box}^w(x, r/a_{\min})$ as a tractable approximation (as also suggested in [11]). Checking connections within this superset of $B(x, r)$ induces an extra run time factor of at most $(A_{\max}/a_{\min})^D$, a bound on the volume ratio between the two sets.

VI. ASYMPTOTIC OPTIMALITY OF DFMT* AND DPRM*

In this section, we prove the asymptotic optimality of DFMT* and DPRM* (obtained as a simple corollary). We conclude the section by providing a discussion of the results.

Algorithm 2 Differential Fast Marching Tree (DFMT*)

```
1  $V \leftarrow \{x_{\text{init}}\} \cup \text{SampleFree}(n)$ ;  $E \leftarrow \emptyset$ 
2  $W \leftarrow V \setminus \{x_{\text{init}}\}$ ;  $H \leftarrow \{x_{\text{init}}\}$ 
3  $z \leftarrow x_{\text{init}}$ 
4 while  $z \notin \mathcal{M}_{\text{goal}}$  do
5    $H_{\text{new}} \leftarrow \emptyset$ 
6    $X_{\text{near}} = \text{Near}(V \setminus \{z\}, z, r_n) \cap W$ 
7   for  $x \in X_{\text{near}}$  do
8      $Y_{\text{near}} \leftarrow \text{Near}(V \setminus \{x\}, x, r_n) \cap H$ 
9      $y_{\text{min}} \leftarrow \arg \min_{y \in Y_{\text{near}}} \{\text{Cost}(y, T = (V, E)) + d(y, x)\}$ 
10    if  $\text{CollisionFree}(y_{\text{min}}, x)$  then
11       $E \leftarrow E \cup \{y_{\text{min}}, x\}$ 
12       $H_{\text{new}} \leftarrow H_{\text{new}} \cup \{x\}$ 
13       $W \leftarrow W \setminus \{x\}$ 
14    end if
15  end for
16   $H \leftarrow (H \cup H_{\text{new}}) \setminus \{z\}$ 
17  if  $H = \emptyset$  then
18    return Failure
19  end if
20   $z \leftarrow \arg \min_{y \in H} \{\text{Cost}(y, T = (V, E))\}$ 
21 end while
22 return  $\text{Path}(z, T = (V, E))$ 
```

A. Asymptotic Optimality of DFMT* and DPRM*

The following theorem presents a result comparing the output cost of DFMT* to the cost of any feasible trajectory.

Theorem VI.1 (DFMT* cost comparison). *Let $(\Sigma, \mathcal{M}_{\text{free}}, x_{\text{init}}, \mathcal{M}_{\text{goal}})$ be a trajectory planning problem satisfying the assumptions A_Σ and suppose $x : [0, T] \rightarrow \mathcal{M}$ is a feasible path with strong δ -clearance, $\delta > 0$. Assume further that x extends into the interior of $\mathcal{M}_{\text{goal}}$, i.e. there exists $\gamma > 0$ such that $B(x(T), \gamma) \subset \mathcal{M}_{\text{goal}}$. Let c_n denote the cost of the path returned by DFMT* with n vertices using a radius*

$$r_n = 4A_{\text{max}}(1 + \eta)^{1/D} \left(\frac{\mu(\mathcal{M}_{\text{free}})}{D} \right)^{1/D} \left(\frac{\log n}{n} \right)^{1/D}$$

for a parameter $\eta \geq 0$. Then for fixed $\varepsilon > 0$

$$\mathbb{P}(c_n > (1 + \varepsilon)c(x)) = O(n^{-\eta} \log^{-1/D} n).$$

Proof. The proof of this theorem is conceptually similar to the one of Theorem 1 in [17]. The details are provided in the Appendix. \square

We are now in a position to state the optimality result for DFMT* (note that the result also provides a convergence rate bound). The optimality of DPRM* will follow as a corollary.

Theorem VI.2 (DFMT* asymptotic optimality). *Let $(\Sigma, \mathcal{M}_{\text{free}}, x_{\text{init}}, \mathcal{M}_{\text{goal}})$ be a trajectory planning problem, satisfying the assumptions A_Σ and with $\mathcal{M}_{\text{goal}}$ ξ -regular, such that there exists an optimal path $x^* : [0, T^*] \rightarrow \mathcal{M}$ with weak δ -clearance for some $\delta > 0$. Let c^* denote the arc*

length of x^* , and let c_n denote the cost of the path returned by DFMT* with n vertices using the radius

$$r_n = 4A_{\text{max}}(1 + \eta)^{1/D} \left(\frac{\mu(\mathcal{M}_{\text{free}})}{D} \right)^{1/D} \left(\frac{\log n}{n} \right)^{1/D}$$

for a parameter $\eta \geq 0$. Then for fixed $\varepsilon > 0$

$$\mathbb{P}(c_n > (1 + \varepsilon)c^*) = O(n^{-\eta} \log^{-1/D} n).$$

Proof. Note that $c^* = 0$ implies $x_{\text{init}} \in \text{cl}(\mathcal{M}_{\text{goal}})$, and DFMT* will terminate immediately and optimally. In this case the result is trivial, therefore assume $c^* > 0$. Let ψ denote the homotopy given by the weak δ -clearance definition for x^* . Using the fact that $\lim_{\alpha \rightarrow 0} c(\psi(\alpha)) = c^*$, we pick an $\alpha_0 > 0$ such that $c(\psi(\alpha_0)) \leq (1 + \varepsilon/4)c^*$. Denote $x := \psi(\alpha_0)$. We now extend x into the interior of $\mathcal{M}_{\text{goal}}$ so that we may apply Theorem VI.1. Since x^* is optimal, $x(T) = x^*(T^*) \in \partial\mathcal{M}_{\text{goal}}$ and the ξ -regularity of $\mathcal{M}_{\text{goal}}$ means there exists $z \in \mathcal{M}_{\text{goal}}$ with $B^e(z, \xi) \subseteq \mathcal{M}_{\text{goal}}$ and $y \in \partial B^e(x(T), \xi)$. Pick $\rho = \min \left\{ \xi, \frac{\delta_{\alpha_0}}{2}, \frac{1}{2A_{\text{max}}} \left(\frac{\varepsilon c^*}{4} \right)^s, \theta_{\text{min}} \right\}$, and take z' on the straight line segment between $x(T)$ and z with $\|x(T) - z'\| \leq \rho$. Then $B(z', \rho) \subset B^e(z', \rho) \subset B^e(z, \xi) \subset \mathcal{M}_{\text{goal}}$. The last two terms in the minimum above ensure, from Lemma II.3, that $d(x(T), z') \leq (\varepsilon/4)c^*$. Consider the extension x' of x constructed by concatenating x with the shortest sub-Riemannian path between $x(T)$ and z' . The cost of x' is bounded as

$$c(x') \leq c(x) + (\varepsilon/4)c^* \leq (1 + \varepsilon/2)c^*.$$

The strong $(\delta_{\alpha_0}/2)$ -clearance of x' is established by noting for any p along the path between $x(T)$ and z' :

$$\begin{aligned} \inf_{a \in \mathcal{M}_{\text{obs}}} \|p - a\| &\geq \inf_{a \in \mathcal{M}_{\text{obs}}} \|x(T) - a\| - \|p - x(T)\| \\ &\geq \delta_{\alpha_0} - \delta_{\alpha_0}/2 = \delta_{\alpha_0}/2. \end{aligned}$$

Then by Theorem VI.1, with the approximation factor $\varepsilon/4$, we have for $\varepsilon \in (0, 1)$ that

$$\begin{aligned} \mathbb{P}(c_n > (1 + \varepsilon)c^*) &\geq \mathbb{P}(c_n > (1 + \varepsilon/4)(1 + \varepsilon/2)c^*) \\ &= O(n^{-\eta} \log^{-1/D} n). \end{aligned}$$

For $\varepsilon \geq 1$ the desired result follows from the monotonicity in ε of the above probability. \square

Remark VI.3. *There is an implicit dependence on the problem parameters and approximation factor ε in the convergence rate bound given above in Theorem VI.2; these fixed parameters influence the threshold $n > N$ after which the stated asymptotic bound holds.*

We conclude this section with the asymptotic optimality result for DPRM*.

Corollary VI.4 (DPRM* asymptotic optimality). *Let $(\Sigma, \mathcal{M}_{\text{free}}, x_{\text{init}}, \mathcal{M}_{\text{goal}})$ be a trajectory planning problem, satisfying the assumptions A_Σ and with $\mathcal{M}_{\text{goal}}$ ξ -regular, such that there exists an optimal path $x^* : [0, T^*] \rightarrow \mathcal{M}$ with weak δ -clearance for some $\delta > 0$. Let c^* denote the arc*

length of x^* , and let c_n denote the cost of the path returned by DPRM* with n vertices using the radius

$$r_n = 4A_{\max}(1 + \eta)^{1/D} \left(\frac{\mu(\mathcal{M}_{\text{free}})}{D} \right)^{1/D} \left(\frac{\log n}{n} \right)^{1/D}$$

for a parameter $\eta \geq 0$. Then for fixed $\varepsilon > 0$

$$\mathbb{P}(c_n > (1 + \varepsilon)c^*) = O(n^{-\eta} \log^{-1/D} n).$$

Proof. Given the same sample set V , the trajectory tree constructed by DFMT* is a subgraph of the roadmap graph constructed by DPRM*. Hence the cost of the path returned by DPRM* is upper bounded by that of DFMT*. The desired result then follows immediately from Theorem VI.2. \square

B. Discussion

We note that there is room for improvement in the choice of constants for the theorems and lemmas above (particularly in the connection radius r_n), which for the sake of clarity were not pushed as tight as possible. As a consequence, the asymptotic rate bound $O(n^{-\eta} \log^{-1/D} n)$ may be achieved with a connection radius smaller by a constant factor. The convergence rate bounds for DFMT* and DPRM*, stated in terms of the sample size, are also tunable to specific implementations and planning problems. A greater local connection radius achieves a faster asymptotic convergence rate, but such a choice may ultimately result in greater algorithm execution time as more connections are checked.

The essential property of the system Σ underpinning the asymptotic optimality results above is the ball-box Theorem II.1. Indeed, although the results we present nominally apply only to arc length, the proofs hold for any system with a cost metric d on \mathcal{M} satisfying the ball-box inequality (3). For example, if the ball-box inequality holds for arc length on a subset of dimensions of \mathcal{M} , DFMT* and DPRM* asymptotic optimality hold with respect to that subspace metric, a fact we make use of in our simulations. The constant a_{\min} ensures that the sub-Riemannian ball $B(x, r)$ is covered by the connection neighborhood of x (i.e. the algorithms search far enough in each direction), while the constant A_{\max} provides a lower bound on $\mu(B(x, r))$ that ensures the sample set will intersect the neighborhood with high probability (i.e. the algorithms examine enough volume along an optimal path). If checking state membership in balls $B(x, r)$ is computationally efficient, then algorithmic performance gains may be achieved by using $\text{Near}(V, x, r) = V \cap B(x, r)$, while maintaining the guarantees of Theorems VI.1 and VI.2. Another consequence of this ball-box analysis foundation is the possibility of accommodating *approximate* steering techniques (as in [21]) for computing local, obstacle-free state connections (one of the bottlenecks for DPRM* and DFMT*). As long as DFMT* and DPRM* have access to a local planner that can connect two states exactly, but possibly with some approximation factor to the optimal distance, then the global optimality theorems hold up to that same factor.

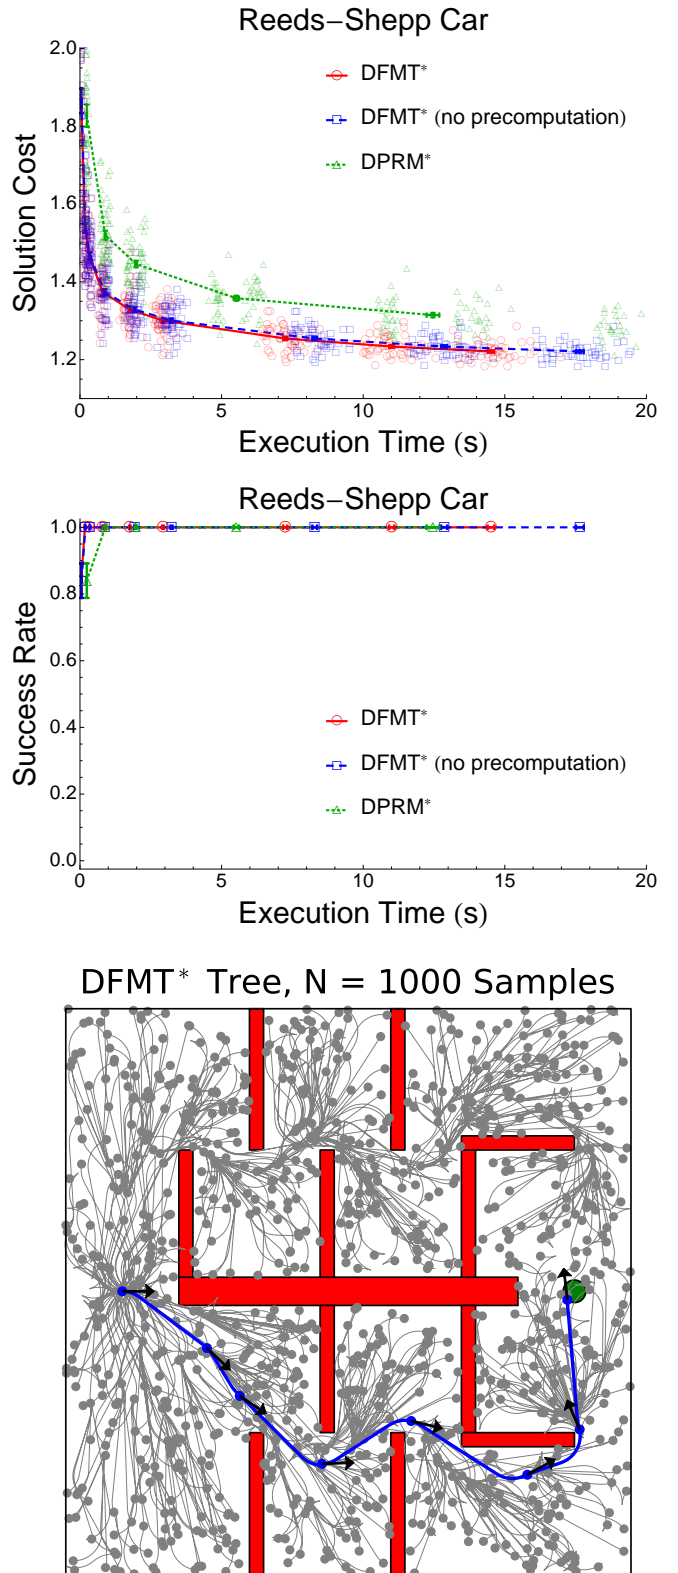


Fig. 1. Top: Simulation results for the Reeds-Shepp car system with a maze obstacle set. The error bars in each axis represent plus and minus one standard error of the mean for a fixed sample size n . Bottom: Example DFMT* tree for $n = 1000$. The feasible path returned is highlighted in blue with car direction at each sample state denoted by an arrow.

VII. IMPLEMENTATION AND NUMERICAL EXPERIMENTS

The DFMT* and DPRM* algorithms were implemented in Julia and run using a Unix operating system with a 2.0 GHz processor and 8 GB of RAM. Near neighbor sets were precomputed and cached at algorithm initialization after the sample set was selected. Note that for batch-processing (as opposed to “anytime”) algorithms such as DFMT* and DPRM*, one can precompute both near neighbor sets and sub-Riemannian distance ahead of time, as they do not depend on the obstacle configuration — the price to pay is a moderate increase in memory requirements.

We tested DFMT* and DPRM* on the Reeds-Shepp car system [23], a well-known example of a regular driftless control-affine system. Within the configuration manifold $\mathcal{M} = \mathbb{R}^2 \times S^1$, trajectories $x(t) = (x^1(t), x^2(t), \theta(t))$ are subject to equation (1) with $g_1(x) = (\cos \theta, \sin \theta, 0)$ and $g_2(x) = (0, 0, 1)$. With the car constrained to move at unit speed with turning radius R , the control set is $U = \{-1, 1\} \times [-1/R, 1/R]$. The privileged coordinate system at each point x is defined by the vectors $g_1(x)$, $g_2(x)$, and $[g_1(x), g_2(x)] = (\sin \theta, -\cos \theta, 0)$ with associated weight vector $(1, 1, 2)$. We aim to optimize planar arc length in the first two dimensions of \mathcal{M} , in which case the constants for the ball-box Theorem II.1 are computed as $a_{\min} = \sqrt{2R}$ and $A_{\max} = 2\sqrt{2R}$ for $\sigma_{\min} = R$.

An analytic solution to the Reeds-Shepp optimal steering problem in the absence of obstacles is known [23], and consists of checking forty-six candidate curves for the one of minimum length. We precomputed and cached the optimal curve type between a state at the origin and any general state (using spatial symmetry this gives the optimal steering curve type between any two states in Reeds-Shepp space); the time for this operation was not included in our accounting of algorithm execution time. We used this data to implement the variants of DFMT* and DPRM* discussed in Section VI-B where the exact sub-Riemannian ball is used for the *Near* function instead of the privileged coordinate box. Collisions with obstacles were detected by approximating trajectories by piecewise linear interpolations, for which polytope intersection is simple to compute.

The simulation results are summarized in Figure 1. A maze was used for \mathcal{M}_{obs} , and our DFMT* and DPRM* implementations were run 50 times each on sample sizes up to $n = 4000$ in the case of DFMT* and $n = 1000$ in the case of DPRM*. We plot results for DFMT* run both with and without the precomputed neighbor sets and distance cache; the solutions returned are identical and differ only in their timing data. For the Reeds-Shepp system, local connection computation and collision checking takes the majority of the computation time, and the speed increase from the cache is modest. For a greater speedup at the cost of more memory usage, all local *collision-free* connections could be cached ahead of time as well, leaving only collision checking with \mathcal{M}_{obs} (the only detail specific to the problem instance) to be computed online.

VIII. CONCLUSIONS

In this paper we presented a thorough theoretical framework to assess optimality guarantees of sampling-based algorithms for planning with driftless control-affine systems. We exploited this framework to design and analyze two novel sampling-based algorithms that are guaranteed to converge, as the number of samples increases, to an optimal solution. Our theoretical framework, which relies on the key notion of probabilistic exhaustivity, separates the probabilistic analysis of random sampling from the discrete analysis of combinatorial optimization algorithms, the two main elements in determining the performance of batch-processing motion planners (such as DFMT* and DPRM*). We hope that this analysis framework can prove useful to analyze other types of algorithms for this class of problems. We note that our analysis approach yields convergence rate guarantees, which shed additional light on the behavior of sampling-based algorithms for differentially-constrained motion planning problems.

This paper leaves numerous important extensions open for further research. One extension that appears quite tractable is generalizing the cost function from arc length to functions that may depend on state-space position and/or control effort. For example, the techniques established in this paper may be used, with a few modifications, to prove DFMT* and DPRM* variants for positionally weighted arc length costs, in which certain spatial regions incur higher or lower costs for traversal. A second extension of great interest is the design and analysis of algorithms that address systems with *drift* (of critical importance, for example, for spacecraft motion planning). The specific analysis techniques used in this paper rely heavily on the fact that driftless control-affine systems can be viewed as metric spaces, but asymmetric adaptations may be imagined within the same proof framework. Third, we plan to study bidirectional versions of DFMT* and DPRM*. Finally, determining how algorithm tuning parameters affect asymptotic convergence rates and, ultimately, run time merits more extensive theoretical and experimental analysis, especially for DCA systems beyond the basic proof-of-concept Reeds-Shepp example presented in this paper.

REFERENCES

- [1] S. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.
- [2] S. M. LaValle, “Motion planning,” *IEEE Robotics Automation Magazine*, vol. 18, no. 2, pp. 108–118, 2011.
- [3] J. H. Reif, “Complexity of the mover’s problem and generalizations,” in *20th Annual Symposium on Foundations of Computer Science*, 1979, pp. 421–427.
- [4] B. Donald, P. Xavier, J. Canny, and J. Reif, “Kinodynamic motion planning,” *Journal of the Association for Computing Machinery*, vol. 40, no. 5, pp. 1048–1066, 1993.
- [5] K. I. Tsianos, I. A. Sucas, and L. E. Kavraki, “Sampling-based robot motion planning: Towards realistic applications,” *Computer Science Review*, vol. 1, no. 1, pp. 2–11, 2007.
- [6] S. M. LaValle and J. J. Kuffner, “Randomized kinodynamic planning,” *International Journal of Robotics Research*, vol. 20, no. 5, pp. 378–400, 2001.
- [7] 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.

- [8] —, “Optimal kinodynamic motion planning using incremental sampling-based methods,” in *Proc. IEEE Conf. on Decision and Control*, 2010, pp. 7681–7687.
- [9] A. Perez, R. Platt, G. Konidaris, L. Kaelbling, and T. Lozano-Perez, “LQR-RRT*: Optimal sampling-based motion planning with automatically derived extension heuristics,” in *Proc. IEEE Conf. on Robotics and Automation*, 2012, pp. 2537–2542.
- [10] G. Goretkin, A. Perez, r. Platt Jr., and G. Konidaris, “Optimal sampling-based planning for linear-quadratic kinodynamic systems,” in *Proc. IEEE Conf. on Robotics and Automation*, 2013.
- [11] S. Karaman and E. Frazzoli, “Sampling-based optimal motion planning for non-holonomic dynamical systems,” in *Proc. IEEE Conf. on Robotics and Automation*, 2013.
- [12] D. J. Webb and J. van den Berg, “Kinodynamic RRT*: Optimal motion planning for systems with linear differential constraints,” in *Proc. IEEE Conf. on Robotics and Automation*, 2013.
- [13] Y. Li, Z. Littlefield, and K. Bekris, “Sparse methods for efficient asymptotically optimal kinodynamic planning,” *Workshop on Algorithmic Foundations of Robotics*, 2014, available at <http://arxiv.org/pdf/1407.2896v2.pdf>.
- [14] R. T. M’Closkey III, “Exponential stabilization of driftless nonlinear control systems,” Ph.D. dissertation, 1995.
- [15] R. M. Murray and S. Sastry, *A mathematical introduction to robotic manipulation*. CRC press, 1994.
- [16] N. E. Leonard and P. Krishnaprasad, “Averaging for attitude control and motion planning,” in *Proc. IEEE Conf. on Decision and Control*, 1993, pp. 3098–3104.
- [17] 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.
- [18] F. Bullo and A. Lewis, *Geometric control of mechanical systems*, ser. Texts in Applied Mathematics. Springer, 2005, vol. 49.
- [19] A. Bellaïche, “The tangent space in sub-riemannian geometry,” in *Sub-Riemannian Geometry*, A. Bellaïche and J. Risler, Eds. Birkhauser, 1996, pp. 1–78.
- [20] R. Montgomery, *A tour of subriemannian geometries, their geodesics, and applications*, ser. Mathematical Surveys and Monographs. American Mathematical Society, 2002, vol. 91.
- [21] F. Jean, G. Oriolo, and M. Vendittelli, “A globally convergent steering algorithm for regular nonholonomic systems,” in *Proc. IEEE Conf. on Decision and Control*, 2005, pp. 7514–7519.
- [22] 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.
- [23] J. Reeds and R. Shepp, “Optimal paths for a car that goes both forwards and backwards,” *Pacific Journal of Mathematics*, vol. 145, no. 2, 1990.
- [24] M. Penrose, *Random Geometric Graphs*. Oxford University Press, 2003.

APPENDIX

Lemma II.2 (Box volume). *Fix $x \in \mathcal{M}$. The volume of $\text{Box}^w(x, r)$ is given by $\mu(\text{Box}^w(x, r)) = r^D$ where $D = \sum_{i=1}^n w_i$.*

Proof. From the formula (2) we see that the Jacobian matrix J_Z for the change of coordinates between the local frame y_1, \dots, y_n and privileged coordinates z_1, \dots, z_n is lower-triangular with all diagonal elements equal to 1; thus the Jacobian determinant $|J_Z|$ is identically 1. Since $Y_1(x), \dots, Y_n(x)$ are orthonormal, the change of coordinates from the Euclidean standard basis to the coordinates y_1, \dots, y_n , with Jacobian J_Y , satisfies $|J_Y| = 1$. Then $\mu(\text{Box}^w(x, r)) = \int_{\|z\|_x \leq r} |J_Z^{-1}| |J_Y^{-1}| dz = \int_{\|z\|_x \leq r} 1 dz = r^{\sum_{i=1}^n w_i}$. \square

Lemma II.3 (Distance comparison). *Fix a point $x_0 \in \mathcal{M}$ and a system of privileged coordinates z_1, \dots, z_n at x_0 . Then*

there exists a positive constant $\theta(x_0) > 0$ such that for all x with $\|x_0 - x\| \leq \theta(x_0)$,

$$\|x_0 - x\| \leq d(x_0, x) \leq 2A_{\max} \|x_0 - x\|^{1/s}.$$

Proof. The left-hand inequality is a consequence of the definition of d . For the right-hand inequality, we note for each coordinate that $|z_i(x)^{1/w_i}| = |y_i(x)|^{1/w_i} + O(\|y(x)\|^{2/w_i}) \leq \|x_0 - x\|^{1/w_i} + O(\|x_0 - x\|^{2/w_i})$, so then $d(x_0, x) \leq A_{\max} \|z(x)\|_{x_0} \leq 2A_{\max} \|x_0 - x\|^{1/s}$ for small $\|x_0 - x\|$. \square

Lemma IV.1. *Let x be a dynamically feasible trajectory and consider a partition of the time interval $0 = \tau_1 < \tau_2 < \dots < \tau_M = T$. Suppose that $\{y_m\} \subset \mathcal{M}$ satisfy a) $y_m \in B(x(\tau_m), \rho)$ for all $m \in \{1, \dots, M\}$, and b) more than a $(1 - \alpha)$ fraction of the y_m satisfy $y_m \in B(x(\tau_m), \beta\rho)$ for a parameter $\beta \in (0, 1)$. Then the cost $c(y^*)$ of the trajectory y^* sequentially connecting the nodes y_1, \dots, y_M is upper bounded as*

$$c(y^*) \leq c(x) + 2M\rho(\beta + \alpha - \alpha\beta).$$

Proof. Using the triangle inequality for the metric d we compare:

$$\begin{aligned} c(y^*) &= \sum_{m=1}^{M-1} d(y_m, y_{m+1}) \leq \sum_{m=1}^{M-1} [d(y_m, x(\tau_m)) \\ &\quad + d(x(\tau_m), x(\tau_{m+1})) + d(x(\tau_{m+1}), y_{m+1})] \\ &= \sum_{m=1}^{M-1} d(x(\tau_m), x(\tau_{m+1})) + 2 \sum_{m=1}^M d(x(\tau_m), y_m) \\ &\quad - d(x_1, y_1) - d(x_M, y_M) \\ &\leq c(x) + 2(M\beta\rho + \alpha M(1 - \beta)\rho) \leq c(x) + 2M\rho(\beta + \alpha - \alpha\beta). \end{aligned}$$

\square

Lemma IV.3. *Fix $n \in \mathbb{N}$, $\alpha \in (0, 1)$, and let S_1, \dots, S_M be disjoint subsets of $\mathcal{M}_{\text{free}}$ with*

$$\mu(S_m) = \mu(S_1) \geq \left(\frac{2 + \log(1/\alpha)}{n} \right) e^2 \mu(\mathcal{M}_{\text{free}}),$$

for each m . Let $V = \text{SampleFree}(n)$ and define

$$K_n := \#\{m \in \{1, \dots, M\} : S_m \cap V = \emptyset\}.$$

Then $\mathbb{P}(K_n \geq \alpha M) \leq \frac{e^{-\alpha M}}{1 - e^{-n}}$.

Proof. The proof relies on a Poissonization argument. Let \tilde{n} be a random variable drawn from a Poisson distribution with parameter n/e^2 (denoted as $\text{Poisson}(n/e^2)$). Consider the set of nodes $\tilde{V} := \text{SampleFree}(\tilde{n})$. Let \tilde{K}_n be the Poissonized analogue of K_n , namely

$$\tilde{K}_n := \#\{m \in \{1, \dots, M\} : S_m \cap \tilde{V} = \emptyset\}.$$

From the definition of \tilde{V} , we can see that $\mathbb{P}(K_n \geq \alpha M) =$

$\mathbb{P}(\tilde{K}_n \geq \alpha M \mid \tilde{n} = n)$. Thus, we have □

$$\begin{aligned}
\mathbb{P}(\tilde{K}_n \geq \alpha M) &= \sum_{j=0}^{\infty} \mathbb{P}(\tilde{K}_n \geq \alpha M \mid \tilde{n} = j) \mathbb{P}(\tilde{n} = j) \\
&\geq \sum_{j=0}^n \mathbb{P}(\tilde{K}_n \geq \alpha M \mid \tilde{n} = j) \mathbb{P}(\tilde{n} = j) \\
&\geq \sum_{j=0}^n \mathbb{P}(\tilde{K}_n \geq \alpha M \mid \tilde{n} = n) \mathbb{P}(\tilde{n} = j) \quad (4) \\
&= \mathbb{P}(\tilde{K}_n \geq \alpha M \mid \tilde{n} = n) \mathbb{P}(\tilde{n} \leq n) \\
&= \mathbb{P}(K_n \geq \alpha M) \mathbb{P}(\tilde{n} \leq n) \\
&\geq (1 - e^{-n}) \mathbb{P}(K_n \geq \alpha M),
\end{aligned}$$

The third line follows from the fact that $\mathbb{P}(\tilde{K}_n \geq \alpha M \mid \tilde{n} = j)$ is nonincreasing in j , and the last line follows from a tail approximation of the Poisson distribution [24, p. 17] and the fact that $\mathbb{E}[\tilde{n}] = n/e^2$.

The locations of the nodes in \tilde{V} are distributed as a spatial Poisson process with intensity $n/(e^2\mu(\mathcal{M}_{\text{free}}))$. This means that for a Lebesgue-measurable region $R \subseteq \mathcal{M}_{\text{free}}$, the number of nodes in R is distributed as a Poisson random variable with distribution $\text{Poisson}((n/e^2)\mu(R)/\mu(\mathcal{M}_{\text{free}}))$, independent of the number of nodes in any region disjoint with R [22, Lemma 11]. Since by assumption the S_m are all disjoint, we get that for $m \in \{1, \dots, M\}$,

$$\mathbb{P}(S_m \cap \tilde{V} = \emptyset) = e^{-(n/e^2)\mu(S_m)/\mu(\mathcal{M}_{\text{free}})}.$$

Therefore, \tilde{K}_n is distributed according to a binomial distribution, in particular according to a distribution $\text{Binomial}(M, e^{-(n/e^2)\mu(S_1)/\mu(\mathcal{M}_{\text{free}})})$. Noting that $e^{-(n/e^2)\mu(S_1)/\mu(\mathcal{M}_{\text{free}})} \leq \alpha/e^2$ by assumption, we have that $\mathbb{E}[\tilde{K}_n] \leq \alpha M/e^2$. So from a tail approximation to the Binomial distribution [24, p. 16], $\mathbb{P}(\tilde{K}_n \geq \alpha M) \leq e^{-\alpha M}$ and thus plugging into inequality (4)

$$\mathbb{P}(K_n \geq \alpha M) \leq \frac{e^{-\alpha M}}{1 - e^{-n}}.$$

□

Lemma IV.4. Fix $n \in \mathbb{N}$ and let T_1, \dots, T_M be subsets of $\mathcal{M}_{\text{free}}$, possibly overlapping, with

$$\mu(T_m) = \mu(T_1) \geq \kappa (\log n/n) \mu(\mathcal{M}_{\text{free}})$$

for each m and some constant $\kappa > 0$. Let $V = \text{SampleFree}(n)$ and denote by E_m the event that $T_m \cap V = \emptyset$ for each m . Then

$$\mathbb{P}\left(\bigvee_{m=1}^M E_m\right) \leq Mn^{-\kappa}.$$

Proof. We union bound the probability of any E_m occurring:

$$\begin{aligned}
\mathbb{P}\left(\bigvee_{m=1}^M E_m\right) &\leq \sum_{m=1}^M \mathbb{P}(T_m \cap V = \emptyset) \\
&= \sum_{m=1}^M \left(1 - \frac{\mu(T_m)}{\mu(\mathcal{M}_{\text{free}})}\right)^n \\
&\leq Me^{-\left(\frac{\mu(T_m)}{\mu(\mathcal{M}_{\text{free}})}\right)n} \leq Mn^{-\kappa}.
\end{aligned}$$

Theorem VI.1 (DFMT* cost comparison). Let $(\Sigma, \mathcal{M}_{\text{free}}, x_{\text{init}}, \mathcal{M}_{\text{goal}})$ be a trajectory planning problem satisfying the assumptions A_Σ and suppose $x: [0, T] \rightarrow \mathcal{M}$ is a feasible path with strong δ -clearance, $\delta > 0$. Assume further that x extends into the interior of $\mathcal{M}_{\text{goal}}$, i.e. there exists $\gamma > 0$ such that $B(x(T), \gamma) \subset \mathcal{M}_{\text{goal}}$. Let c_n denote the cost of the path returned by DFMT* with n vertices using a radius

$$r_n = 4A_{\max}(1 + \eta)^{1/D} \left(\frac{\mu(\mathcal{M}_{\text{free}})}{D}\right)^{1/D} \left(\frac{\log n}{n}\right)^{1/D}$$

for a parameter $\eta \geq 0$. Then for fixed $\varepsilon > 0$

$$\mathbb{P}(c_n > (1 + \varepsilon)c(x)) = O(n^{-\eta} \log^{-1/D} n).$$

Proof. Consider n so that $r_n \leq \min\{\gamma, \delta a_{\min}/2A_{\max}\}$, and apply Theorem IV.5 to produce, with probability at least $1 - O(n^{-\eta} \log^{-1/D} n)$, a sequence of waypoints $\{y_m\}_{m=1}^M \subset V$ which (ε, r_n) -trace x . We claim that in the event that such $\{y_m\}$ exist, the DFMT* algorithm will return a path with cost upper bounded as $c_n \leq c(y^*) \leq (1 + \varepsilon)c(x)$. It is clear that the desired result follows from this claim.

Assume the existence of (ε, r_n) -tracing $\{y_m\}$. Note that our upper bound on r_n implies that $\text{Box}^w(y_m, r_n/a_{\min})$ intersects no obstacles. This follows from the inclusion $\text{Box}^w(y_m, r_n/a_{\min}) \subset B(y_m, A_{\max}r_n/a_{\min}) \subset B^e(y_m, A_{\max}r_n/a_{\min})$, and the Euclidean distance bound

$$\begin{aligned}
\inf_{a \in \mathcal{M}_{\text{obs}}} \|y_m - a\| &\geq \inf_{a \in \mathcal{M}_{\text{obs}}} \|x_m - a\| - \|y_m - x_m\| \\
&\geq 2 \left(\frac{A_{\max}}{a_{\min}}\right) r_n - r_n \geq \left(\frac{A_{\max}}{a_{\min}}\right) r_n.
\end{aligned}$$

Consider running DFMT* to completion and for each y_m , let $c(y_m)$ denote the cost-to-come of y_m in the generated tree. If y_m is not contained in any edge of the tree, we set $c(y_m) = \infty$. We show by induction that

$$\min(c(y_m), c_n) \leq \sum_{k=1}^{m-1} d(y_k, y_{k+1}), \quad (5)$$

for all $m \in \{2, \dots, M\}$.

The base case $m = 2$ is trivial, since the first step in the DFMT* algorithm is to make every collision-free connection between $x_{\text{init}} = y_1$ and the nodes contained in $\text{Box}^w(x, r_n/a_{\min})$, which will include y_2 and, thus, $c(y_2) = d(y_1, y_2)$. Now suppose (5) holds for $m - 1$; that means that one of the following four statements must hold.

1. $c_n \leq \sum_{k=1}^{m-2} d(y_k, y_{k+1})$,
2. $c(y_{m-1}) \leq \sum_{k=1}^{m-2} d(y_k, y_{k+1})$ and DFMT* ends before considering y_m ,
3. $c(y_{m-1}) \leq \sum_{k=1}^{m-2} d(y_k, y_{k+1})$ and $y_{m-1} \in H$ when y_m is first considered,
4. $c(y_{m-1}) \leq \sum_{k=1}^{m-2} d(y_k, y_{k+1})$ and $y_{m-1} \notin H$ when y_m is first considered.

Case 1: $c_n \leq \sum_{k=1}^{m-2} d(y_k, y_{k+1}) \leq \sum_{k=1}^{m-1} d(y_k, y_{k+1})$.

Case 2: $c(y_{m-1}) < \infty$ implies that y_{m-1} enters H at some point during DFMT*. The case that y_m goes unconsidered means that the algorithm terminates before x_{m-1} is ever minimum-cost element of H . Since the end-node of the solution returned must have been the minimum-cost element of H , $c_n \leq c(y_{m-1}) \leq \sum_{k=1}^{m-2} d(y_k, y_{k+1}) \leq \sum_{k=1}^{m-1} d(y_k, y_{k+1})$.

Case 3: Since $\text{Box}^w(y_m, r_n/a_{\min})$ intersects no obstacles, y_m must be connected to some parent when it is first considered. y_{m-1} is a candidate, so $c(y_m) \leq c(y_{m-1}) + d(y_{m-1}, y_m) \leq \sum_{k=1}^{m-1} d(y_k, y_{k+1})$.

Case 4: When y_m is first considered, it is because there exists $z \in \text{Box}^w(y_m, r_n/a_{\min})$ such that z is the minimum-cost element of H . Again since $\text{Box}^w(y_m, r_n/a_{\min})$ intersects no obstacles and contains at least one node in H , y_m must be connected to some parent when it is first considered. Since $c(y_{m-1}) < \infty$, there is a well-defined path $\mathcal{P} = \{v_1 = x_{\text{init}}, v_2, \dots, v_q = y_{m-1}\}$ giving its cost-to-come in the DFMT* tree at termination. Let $w = v_j$, where $j = \max_{i \in \{1, \dots, q\}} \{i : v_i \in H \text{ when } y_m \text{ is first considered}\}$. Then there are two subcases, either $w \in B(y_m, r_n)$ or $w \notin B(y_m, r_n)$. If $w \in B(y_m, r_n)$, then,

$$\begin{aligned} c(y_m) &\leq c(w) + d(w, y_m) \\ &\leq c(w) + d(w, y_{m-1}) + d(y_{m-1}, y_m) \\ &\leq c(y_{m-1}) + d(y_{m-1}, y_m) \leq \sum_{k=1}^{m-1} d(y_k, y_{k+1}). \end{aligned}$$

If $w \notin B(y_m, r_n)$, then $r_n \leq d(w, y_m) \leq d(w, y_{m-1}) + d(y_{m-1}, y_m)$, so we have

$$\begin{aligned} c(y_m) &\leq c(z) + d(z, y_m) \leq c(w) + r_n \\ &\leq c(w) + d(w, y_{m-1}) + d(y_{m-1}, y_m) \\ &\leq c(y_{m-1}) + d(y_{m-1}, y_m) \leq \sum_{k=1}^{m-1} d(y_k, y_{k+1}). \end{aligned}$$

Thus the inductive step holds in all cases and therefore (5) holds for all m . In particular taking $m = M$, and noting that $r_n \leq \gamma$ implies that $y_M \in \mathcal{M}_{\text{goal}}$, this means that $c_n \leq c(y_M) \leq \sum_{k=1}^{M-1} d(y_k, y_{k+1}) = c(y^*)$ as desired. \square