On Defining and Sampling a Medial Axis Structure in State Space

Jory Denny and Jeremy LeCrone and Vadim Kudlay and Andre Shannon

***Abstract*— Many physical robot systems operate in safety critical settings, for example self-driving vehicles, unmanned aerial vehicles, or industrial arms. Motion planning methods often only consider a minimal safety requirement, that is staying a set distance away from risky areas, i.e., maintain a minimum 'geometric' clearance from obstacles. However, this approach is not suitably adaptable to spaces with narrow passages of variable widths. For holonomic systems, extensive work has been considered in maximizing safety by planning on, or near, the medial axis of configuration space — the set of configurations equidistant from two or more obstacles.**

In this work, we provide a novel look at the medial axis of a free state space for non-holonomic systems. Specifically, we observe that the classical definition using equidistant states might be disconnected for any robot that is not globally small-time locally controllable. Instead, we consider using the singularity set of the system dependent clearance function, and conjecture that this is always a connected structure in state space. Further, we provide algorithms for approximating clearance in non-holonomic systems and approximately sampling our new medial axis structure. We provide experimental validation of improving clearance compared with randomized sampling of state space.

I. INTRODUCTION

Many applications of robotics are safety critical [1], for example, self-driving vehicles, unmanned aerial vehicles, or any armed robot operating near humans. There are differing philosophies on safety ranging from maintaining above a minimum distance from obstacles [2], maximizing distance from obstacles [3], or producing smooth and predictable behaviors of robot systems [4] depending on the sub-area of robotics. In motion planning, the distance from obstacles, i.e., *geometric clearance*, is very commonly considered. There are distinct trade-offs between maintaining above a minimum clearance vs maximizing clearance, e.g., computational cost vs reasonable guarantees of safety.

Maximal clearance is highly correlated to the topological structure and geometric composition of the space. The *medial axis* of a space is the set of points with maximal clearance properties, see Figure 1(a), and has been extensively explored in geometry and motion planning [5], [6]. There are a few equivalent definitions of medial axis for a geometric space, for example: (1) the set of points equidistant to two or more obstacles or (2) the clearance function's singularity (or 'ridge') set, i.e., the points on which the gradient is not defined. Geometric algorithms computing explicit representations and sampling-based motion planners alike often exploit these definitions in their core algorithmic approaches. Despite its widespread appeal, computing explicit representations is limited to low dimensional spaces (≤ 3), and even sampling-based planning methods exploiting its structure are limited to configuration spaces of holonomic systems.

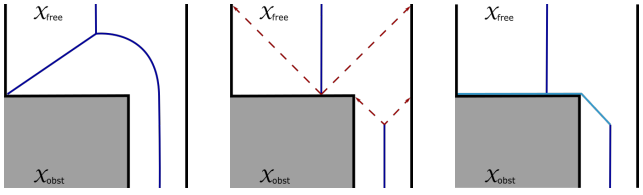

(a) Geometric MA (b) Equidistant MA (c) Singularity MA

Fig. 1. A point robot that can set its angle of motion in the range $[\pi/4, 3\pi/4]$ while maintaining constant upward velocity. Various medial axis definitions shown in blue: (a) the geometric medial axis with respect to position only, (b) the equidistant medial axis with respect to valid motions (specific equidistant motions shown in dashed red), and (c) our contribution, the singular set medial axis with respect to valid motions, (jump discontinuities highlighted in light blue).

In this work, we provide a novel perspective for defining the medial axis of a free state space for non-holonomic systems. In this setting, we observe that classical approaches to defining medial axes produce non-equivalent structures — contrasting their equivalence observed in configuration spaces for holonomic systems. Rather than geometric distance, we consider distance between states determined by feasible motions of a control system and provide examples showing the inequivalence of (1) the set of states equidistant to two or more obstacles and (2) the singularity set of the clearance function, i.e., the states at which the gradient of the clearance function is not defined, as shown in Figure 1. Further, our analysis shows that this inequivalence only occurs when the non-holonomic system is not small-time locally controllable everywhere and that this causes the structure produced by definition (1) to be disconnected. To provide utility to these structures, we produce algorithms for approximating the clearance of a state and approximately sampling the medial axis. Our specific contributions include:

- an analysis of two novel medial axis definitions in state space – characterized by either equidistance to the boundary or the singularities of clearance – showing their inequivalence and connectivity properties;
- approximation algorithms for computing the clearance of a state and sampling the singularity set of the clearance function; and
- experimental validation of the approach demonstrating improved clearance compared to randomized sampling of state space.

This work was not supported by any organization

Jory Denny, Jeremy LeCrone, Vadim Kudlay, and Andre Shannon are with the Department of Mathematics and Computer Science, University of Richmond, VA 23173, USA {jdenny, jlecrone, vadim.kudlay, andre.shannon}@richmond.edu

Jory Denny is the Director of SpiRoL: Spider Robotics Lab.

II. PRELIMINARIES

To fully understand our novel approach to the medial axis for non-holonomic systems, we review holonomic systems, the traditional definitions of medial axis and non-holonomic systems. We also review related work regarding motion planning on or near the medial axis.

A. Holonomic Systems and the Medial Axis

Degrees of freedom (DOFs) parameterize a unique placement of a robot (e.g., center of mass position, orientation, joint angles, etc.) in its two- or three-dimensional world, or *workspace*. A configuration $q = \langle q_1, q_2, \ldots, q_d \rangle$ is a specification of the values for the d DOFs, where q_i is the ith DOF. Systems that can be defined through position-only constraints are deemed *holonomic*. The set of all possible configurations is called the *configuration space* ($\mathcal{C}_{space}$) [7], and it can be partitioned into two main subsets: *free space* ($\mathcal{C}_{free}$) and *obstacle space* ($\mathcal{C}_{obst}$). These two subsets represent all possible valid, e.g., collision-free, and invalid configurations respectively.

In general, it is infeasible to explicitly compute any complete representation of $\mathcal{C}_{free}$ [8]. To overcome this inherent difficulty, sampling-based motion planners [9] avoid an explicit representation of $\mathcal{C}_{obst}$ by noting that a configuration q can be classified into $\mathcal{C}_{free}$ or $\mathcal{C}_{obst}$ efficiently by performing a workspace collision detection test between the robot placed at q and the environment. The boundary of $\mathcal{C}_{obst}$ is denoted $\partial\mathcal{C}_{obst}$ and is commonly referred to as the contact space.

For holonomic systems, the *clearance* $\text{CLR}(q)$ of a configuration q is the minimum distance from q to $\partial\mathcal{C}_{obst}$. The nearest configuration to q in $\partial\mathcal{C}_{obst}$ is called the *witness* to the clearance value. The *medial axis* of $\mathcal{C}_{free}$, denoted $\mathcal{M}_{\mathcal{C}_{free}}$, is the set of configurations with maximal clearance. The set $\mathcal{M}_{\mathcal{C}_{free}}$ is precisely characterized by any of the following equivalent definitions:

- the closure of the set of all configurations equidistant to two or more points on the contact surface, i.e. configurations with two or more unique witnesses;
- all configurations corresponding to centers of maximally inscribed hyperspheres in $\mathcal{C}_{free}$;
- all configurations not contained within an optimal path from another point to $\partial\mathcal{C}_{obst}$;
- all 'quenching' points for a 'grassfire' wavefront moving inwards from $\partial\mathcal{C}_{obst}$;
- all singular 'ridge' points of $clr(\cdot)$, i.e., the configurations at which the gradient of $clr(\cdot)$ is not defined.

We highlight the equivalence of the first and last definitions in particular, as these two statements are most relevant to our contributions. We will refer to $\mathcal{M}_{\mathcal{C}_{free}}$ as the *geometric* medial axis.

B. Non-holonomic Systems

A *non-holonomic constraint* is one that cannot be expressed as position-only, e.g., a wheeled vehicle where velocity is constrained to always be oriented in the same direction as its wheels. A *non-holonomic system* is one with non-holonomic constraints. We express the system's *state* by considering the configuration q, its velocity $\dot{q}$, and any other relevant derivatives of q. The set of all states is denoted $\mathcal{X}_{space}$ and can be split into $\mathcal{X}_{free}$, $\mathcal{X}_{obst}$, and $\partial\mathcal{X}_{obst}$, similar to a configuration space. Typically, for non-holonomic systems a set of allowable controls, $\mathcal{U}_{space}$, define admissible forces that act upon the robot driving it to a new state. Thus, a control system $\dot{x} = f(x, u)$ defines the evolution from one state to another. Additional constraints can be imposed on motion, e.g., disallowing collisions with obstacles, which is collectively termed a *kinodynamic* system. Thus, the motion planning problem for a non-holonomic system is seeking to find a time-varying series of allowable controls to drive a system from a start state to a goal region while remaining entirely within $\mathcal{X}_{free}$.

An important feature influencing motion planning for any kinodynamic system is whether or not it is *small-time locally controllable* (STLC) in all of $\mathcal{X}_{free}$. To clarify STLC, first define the *reachable* set from x in time t, $R(t, x)$, to be all states attainable from x in less than t time, applying any admissible controls from $\mathcal{U}_{space}$. Then, a system is STLC at a state x if x is contained in the interior of $R(t, x)$ for all (small) $t > 0$.

C. Related Work

A common solution to the motion planning problem is the utilization of sampling-based methods, e.g., Probabilistic RoadMaps (PRMs) [9]. Most often, sampling-based motion planners randomly select and progressively expand and connect robot configurations to form an approximate graph representation of $\mathcal{C}_{free}$, called a *roadmap*. A roadmap encodes valid states as its nodes and transitions between them as its edges. In order to find a path, a starting state and goal region are connected to the roadmap, and a feasible path is extracted from it. While many approaches attempt to improve sampling by generating configurations near $\partial\mathcal{C}_{obst}$, e.g., [10]–[13], they do not properly maintain any safety guarantees, e.g., maximal distance from the obstacles.

Many purely geometric approaches have been proposed to construct or plan directly on the medial axis of the workspace [14], [15], combine it with potential fields [16], or extend it to higher order graph structures, e.g., the Generalized Voronoi Graph [17], [18]. Other work approximates the 'skeleton' of the workspace medial axis and can be useful for sampling [19]–[21]. For example, these methods randomly sample points on the surfaces of a set of partially overlapping maximal spheres, and then use this to guide a sampling-based planner. However, collectively these methods are expensive and can only be used in 2D or 3D workspaces.

Optimal and near-optimal sampling-based planners [22] can be configured to optimize for clearance in conjunction with path length. However, this adds a prohibitive cost to optimal planning as is demonstrated in [23]. Another category of planners, Transition-based planning [24], can be configured to optimize for clearance. This method however does not directly sample or plan along the medial axis.

A category of sampler, Medial Axis Probabilistic RoadMap (MAPRM) [6], retracts configurations to the me-

dial axis of $\mathcal{C}_{free}$. This is guaranteed to sample inside of narrow passages more often than uniform random sampling. These methods have been applied to high dimensions [25], uniformly sampling the medial axis [26], and RRTs [23].

Recent work uses Support Vector Machine (SVM) classification to quickly approximate the medial axis [27]. In this method, obstacles are given unique labels for classification and samples are assigned labels according to their nearest obstacle. The medial axis can then be approximated by applying a max-margin push on a sample until its label changes.

Other approaches utilize and reason about clearance during planning, but do not directly alter a node's clearance. For example, volume-based sampling with RRT* [28] utilizes a node's clearance to define a hypersphere in $\mathcal{C}_{free}$ which is entirely visible to the node. In this way, an efficient extension and biasing technique to improve the efficiency of RRT was created.

Further, clearance can be used to analyze and deform already planned paths [3], [29]. In many cases, these methods rely on high clearance corridors.

All of these related approaches operate with a purely geometric medial axis framework in a workspace or configuration space only. To the best of the authors' knowledge, the medial axis for non-holonomic systems has not been deeply explored.

III. DEFINING A MEDIAL AXIS IN STATE SPACE

Upon first observation, a geometric medial axis does not necessarily align with the 'safest' states for even simple dynamical systems. Consider Dubin's car — a car that can only move forward and change its steering angle — operating in a space with two infinitely long parallel walls many car lengths apart. The geometric medial axis is the middle line parallel and equidistant to both walls. While this sometimes coincides with 'safe' states of a Dubin's car, it often does not. In fact, states in which the car's rear is against one of the walls are perfectly safe, as they are the farthest from any wall given the available controls of the system. In this way, any synonymous definition of the medial axis, encoding the 'safest' states of the system must consider the system dynamics and admissible controls.

First, we begin by defining an appropriate notion of clearance for a kinodynamic system (Definition 1) and its corresponding witness sets (Definition 2). A trajectory π is a parameterized (bounded) curve in $\mathcal{X}_{space}$ satisfying the control system $\dot{x} = f(x, u)$ for some admissible choice of time varying controls from $\mathcal{U}_{space}$. A trajectory moves an initial state $x \in \mathcal{X}_{free}$ to a terminal state $x' \in \mathcal{X}_{space}$ and c_π denotes the cost of π, e.g., arclength, time, energy used, etc. Let $\Pi_{\partial\mathcal{X}_{obst}}(x)$ denote the set of all *collision trajectories* π taking x to some $x' \in \partial\mathcal{X}_{obst}$, while remaining entirely in $\mathcal{X}_{free}$ except at the *terminus* x'. Then,

Definition 1: the **clearance**, $\text{CLR}(x)$, of a state $x \in \mathcal{X}_{free} \cup \partial\mathcal{X}_{obst}$ is the minimum cost of all trajectories in $\Pi_{\partial\mathcal{X}_{obst}}(x)$, that is to say:

$$\text{CLR}(x) = \min_{\pi \in \Pi_{\partial\mathcal{X}_{obst}}(x)} c_\pi$$

(we define $\text{CLR}(x) = 0$ for all $x \in \partial\mathcal{X}_{obst}$) and

Definition 2: the **witness set** of $x \in \mathcal{X}_{free}$, $\text{WIT}(x)$, is the set of all termini $x' \in \partial\mathcal{X}_{obst}$ realizing the clearance of x, i.e., all states $x' \in \partial\mathcal{X}_{obst}$ for which there exists $\pi \in \Pi_{\partial\mathcal{X}_{obst}}(x)$ terminating at x' such that $c_\pi = \text{CLR}(x)$.

Next, we introduce definitions for a medial axis structure in $\mathcal{X}_{free}$ akin to the definitions introduced for $\mathcal{M}_{\mathcal{C}_{free}}$. Herein, we are similarly motivated by the goal of identifying a medial axis structure $\mathcal{M}_{\mathcal{X}_{free}}$ in $\mathcal{X}_{free}$ through which trajectories should pass to maximize 'safety' from collisions.

Definition 3: The **multiple witness medial axis**, $\mathcal{M}_{pts}$, is the closure of the set of states $x \in \mathcal{X}_{free}$ with two or more witness points, that is $|\text{WIT}(x)| \geq 2$.

Definition 4: The **singular medial axis**, $\mathcal{M}_{sing}$, is the set of states $x \in \mathcal{X}_{free} \cup \partial\mathcal{X}_{obst}$ at which the gradient of $\text{CLR}(\cdot)$ is undefined.

At first glance, the reader may assume that these two sets are equal, recalling that the relevant statements in the geometric setting were equivalent, both producing the structure $\mathcal{M}_{\mathcal{C}_{free}}$, see Section II. As our first major contribution, we will show that these two definitions are in fact nonequivalent for general kinodynamic systems. To accomplish this, we introduce a simple example and observe that $\mathcal{M}_{sing}$ not only includes 'ridge' points x (coinciding with states in $\mathcal{M}_{pts}$ at which CLR is continuous but fails to be differentiable), but also includes states x at which CLR is discontinuous. We note that the presence of discontinuities in CLR for kinodynamic systems appears to be the core reason these definitions differ on $\mathcal{X}_{free}$, whereas the definitions coincide in $\mathcal{C}_{free}$ where discontinuities of CLR will never arise.

Theorem 1: $\mathcal{M}_{pts} \not\equiv \mathcal{M}_{sing}$

Proof: We introduce a simple example with constant upward 'drift' we refer to as the *Galaga* system[1]. In this system, we have a point moving in a subset of two dimensional space subject to constant upward velocity (in the x_2-direction) and limited controls on left / right movements (in the x_1-direction).

Specifically, consider a simple robot with state described by its location $x = (x_1, x_2)$ in $\mathcal{X}_{space} = \mathbb{R}^2$, subject to the control system

$$(\dot{x}_1, \dot{x}_2) = (u, 1) \qquad \text{for} \quad u \in \mathcal{U}_{space} = [-1, 1].$$

For any $x \in \mathcal{X}_{space}$ and $t > 0$, note that $R(t, x)$ is a linear cone with a flat horizontal front, opening upward from x with edges forming angles $\pi/4$ and $3\pi/4$ relative to the positive x_1-axis. We further introduce obstacles $\mathcal{X}_{obst}$ so that its complement

$$\mathcal{X}_{free} = (-1, 1) \times (-\infty, 0] \cup (-5, 1) \times (0, \infty),$$

i.e, free space is a straight narrow vertical alley whose left wall abruptly opens into a wider alley above $x_2 = 0$.

[1]The Galaga system was inspired by the 1981 classic arcade game of the same name and its subsequent incarnations.

Applying Definition 3, we find the multiple witness medial axis is composed of two vertical lines, located at the center of respective portions of the hallway $\mathcal{X}_{free}$, i.e.

$$\mathcal{M}_{pts} = \{(0, x_2) : x_2 \leq -1\} \cup \{(-2, x_2) : x_2 \geq 0\}.$$

Note that the lower portion of this structure terminates at $x_2 = -1$; the precise location above which $R(\cdot, x)$ no longer intersects the 'closer' obstacle at $x_1 = -1$. This is shown in Figure 1(b).

Applying Definition 4, we find the singular medial axis includes all of $\mathcal{M}_{pts}$ (all points on $\mathcal{M}_{pts}$ are 'ridge' points at which CLR is continuous but not differentiable), while $\mathcal{M}_{sing}$ includes two additional line segments in $\mathcal{X}_{space}$ at which CLR has discontinuities, shown in Figure 1(c). Namely, note that points below the line segment connecting $(0, -1)$ to the corner $(-1, 0)$ are closest to the wall $x_1 = -1$, while those above this line segment are closest to the wall at $x_1 = 1$. Thus, CLR is discontinuous at all points on this straight line segment. Meanwhile, CLR is also discontinuous at all points on the horizontal 'shelf' portion of $\partial\mathcal{X}_{obst}$, at $x_2 = 0$, where CLR is trivial, while all nearby points $x \in \mathcal{X}_{free}$ have nontrivial clearance. Summarizing, we have

$$\mathcal{M}_{sing} = \mathcal{M}_{pts} \cup \{(x_1, -x_1 - 1) : -1 \leq x_1 \leq 0\} \cup \{(x_1, 0) : -5 \leq x_1 \leq 0\},$$

which proves the theorem. ■

We also explored a variety of extensions to the Galaga system with more physically relevant applications:

• We consider a drift-free model with states again given by position $x = (x_1, x_2)$ in $\mathcal{X}_{space} = \mathbb{R}^2$, subject to controls

$$(\dot{x}_1, \dot{x}_2) = (v \cos\theta, v \sin\theta).$$

With constraints on velocity and angle of motion – say $v \in [v_{min}, 1]$ and $\theta \in [\pi/4, 3\pi/4]$ – the reachable sets $R(t, x)$ take on nearly the same structure as before, though the front of $R(t, x)$ is a circular arc here, rather than a flat horizontal edge. If we introduce the same obstacles $\mathcal{X}_{obst}$ as before and assume $v_{min} \geq 0$, then $\mathcal{M}_{pts}$ and $\mathcal{M}_{sing}$ are identical to those detailed in the Galaga model. On the other hand, if we allow for any amount of 'downward' motion, i.e., letting $v_{min} < 0$, the clearance function CLR experiences no discontinuities in this case and the structures $\mathcal{M}_{pts}$ and $\mathcal{M}_{sing}$ coincide. Note that this system is a simplification of the common Unicycle model.

• Next, we consider a model with second-order controls on x_2, where states are given by $x = (x_1, x_2, \dot{x}_2)$ in $\mathcal{X}_{space} = \mathbb{R}^3$. The system is governed by controls

$$(\dot{x}_1, \ddot{x}_2) = (u_1, u_2) \qquad \text{for} \quad (u_1, u_2) \in \mathcal{U}_{space} = [-1, 1]^2,$$

restricted to $\mathcal{X}_{free}$. This is a three-dimensional analog of the suddenly widening alley considered previously.. In particular, set $\mathcal{X}_{free}$ so that $x_1 \in (-1, 1)$ when $x_2 \leq 0$, $x_1 \in (-5, 1)$ when $x_2 > 0$, and further restrict $\dot{x}_2 \geq v_{min} > 0$ for all positions (x_1, x_2). Regarding the bound $v_{min} > 0$, consider an airplane flying at constant elevation through a long, narrow passage, then v_{min} may indicate the minimum airspeed the aircraft can maintain without stalling.

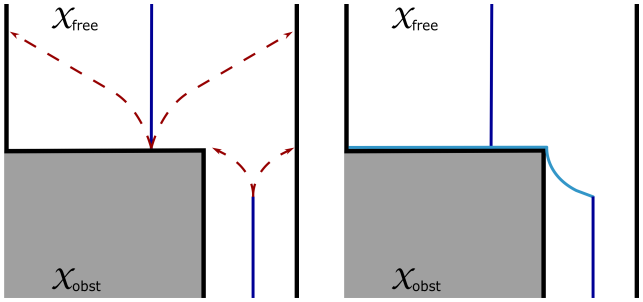

(a) Equidistant MA (b) Singular MA

Fig. 2. A point robot that can control x_1 velocity and x_2 acceleration, with limitation $\dot{x}_2 \geq v_{min}$. Optimal paths minimize arclength. Cross sections of medial axis structures at a slice $\dot{x}_2 > v_{min}$: (a) the equidistant medial axis with optimal motions shown in dashed red (concave parabolas continued by straight lines with slope $v_{min}/\dot{x}_1$), and (b) the singular set medial axis, with discontinuities exhibiting the complimentary convex parabolic structure.

For this model, we take c_π to be the arclength of the path π in $\mathcal{X}_{space}$. Considering the projection of paths to the x_1x_2-plane, optimal controls always decrease $\dot{x}_2$ (applying $u_2 = -1$), so optimal paths follow convex parabolic paths until either the terminus is reached, or until $\dot{x}_2 = v_{min}$ at which point the path continues as a straight line with $\dot{x}_2$ constant and $\dot{x}_1$ optimally chosen. This setting again produces a disconnected multiple witness medial axis, $\mathcal{M}_{pts}$ containing $(x_1, x_2, \dot{x}_2)$ if $x_1 = -2$ and $x_2 \geq 0$, or else $x_1 = 0$ and $x_2 \leq \rho(\dot{x}_2)$, where $\rho(\dot{x}_2) = \dot{x}_2 - \frac{1}{2}$ when $\dot{x}_2 \geq v_{min} + 1$ and $\rho(\dot{x}_2) = \frac{1}{2}(\dot{x}_2 - v_{min})^2 + v_{min}$ when $v_{min} \leq \dot{x}_2 < v_{min} + 1$. While the singular medial axis contains additional structures connecting the ends of $\mathcal{M}_{pts}$ to the corner $(-1, 0, \dot{x}_2)$ that take the form of optimal paths reflected about the origin. In particular, for $\dot{x}_2 \geq v_{min} + 1$ fixed, $\mathcal{M}_{sing}$ contains the segment of the parabolic arc $2x_2 = \dot{x}_2^2 - (x_1 + (1 - \dot{x}_2))^2$ between $x_1 = -1$ and $x_1 = 0$. Moreover, when $v_{min} \leq \dot{x}_2 < v_{min} + 1$, $\mathcal{M}_{sing}$ contains points on the parabolic arc $x_2 = v_{min}(\dot{x}_2 - v_{min}) + \frac{1}{2}(\dot{x}_2 - v_{min})^2$ from $x_1 = -1$ to $x_1 = \dot{x}_2 - v_{min} - 1$, continued by the line $x_2 = \dot{x}_2(1 + x_1) + \frac{1}{2}(\dot{x}_2 - v_{min})^2$ from $x_1 = \dot{x}_2 - v_{min} - 1$ to $x_1 = 0$. A visual comparison of $\mathcal{M}_{pts}$ and $\mathcal{M}_{sing}$ for this model are shown Figure 2.

Upon further reflection, we recognized that the systems in which $\mathcal{M}_{sing} = \mathcal{M}_{pts}$ are STLC everywhere, while the systems in which the definitions differ fail to be STLC everywhere. Intuitively, this can be thought of in a few ways. First, the notion of distance in non-STLC systems is asymmetric, i.e., going from state x to x' costs different from a path taking x' to x, if such a reverse path even exists. Second, we can consider a propagating wave-front formulation of medial axis similar to the grassfire construction in geometric settings, i.e., the set of quenching points of a wave-front propagating uniformly from $\partial\mathcal{C}_{obst}$. Applying similar reasoning to the nonholonomic setting, we see that in an STLC everywhere system, wave-fronts still propagate in all directions from $\partial\mathcal{X}_{obst}$, see Figure 3(a). Whereas, in non-STLC systems, wave-fronts propagate in limited directions corresponding to reverse integration along

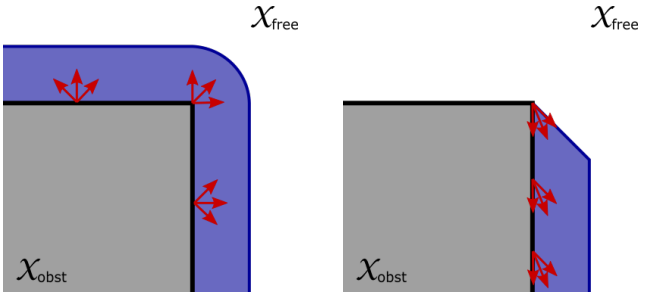

(a) STLC System (b) Non-STLC System

Fig. 3. Wavefront propagations of $\partial\mathcal{X}_{obst}$: (a) an STLC everywhere system and (b) a non-STLC everywhere system.

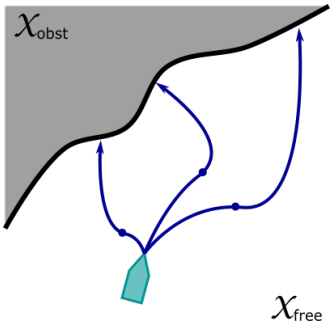

Fig. 4. Pose of a state (teal) for a car-like system and three random bang-bang trajectories (blue). Dots show the "time" at which the control switched along the trajectory. The left-most trajectory would be used to approximate clearance in this case, as it collides with $\partial\mathcal{X}_{obst}$ first in time.

limited local controls available in the system, see Figure 3(b). Thus, every portion of the boundary of a wave in an STLC system is a proper propagating wave-front, while non-STLC systems (with proper configurations for $\mathcal{X}_{obst}$) produce wave boundary portions that are stationary.

Because of its lengthy nature, we only provide a proof sketch for this second main theoretical contribution.

Theorem 2: Given a system that is not STLC everywhere. Then, there exists an $\mathcal{X}_{free} \subset \mathcal{X}_{space}$ such that $\mathcal{M}_{pts} \not\equiv \mathcal{M}_{sing}$.

Proof: [Sketch] Let $x \in \mathcal{X}_{space}$ be a state at which the system is not STLC. By definition, x is contained in the boundary of the reachable set $R(t, x)$ for all t sufficiently small. Construct $\mathcal{X}_{obst}$ so that $\partial\mathcal{X}_{obst}$ coincides locally with the boundary of $R(t, x)$, moving from x. This will create a discontinuity of CLR at x. Continue constructing $\mathcal{X}_{obst}$ by extending $\partial\mathcal{X}_{obst}$ until you enclose $R(t, x)$ in the bounded complement $\mathcal{X}_{free}$. Since this space contains a discontinuity of the clearance function, clearly, $\mathcal{M}_{pts} \not\equiv \mathcal{M}_{sing}$. ■

Given the properties discussed about $\mathcal{M}_{sing}$, as it contains $\mathcal{M}_{pts}$ as a subset and appears to always be a connected structure, we propose this as the ideal definition of the medial axis of $\mathcal{X}_{free}$. For the rest of the paper, we refer to $\mathcal{M}_{sing}$ as $\mathcal{M}_{\mathcal{X}_{free}}$.

Notice that Theorem 2 only shows existence of a specific construction of $\mathcal{X}_{obst}$, without exploring essential properties of $\mathcal{X}_{obst}$ that lead to discontinuities of CLR. We leave it to future work to more deeply explore this and other theoretical properties of our novel definitions. Specifically, like the geometric setting, we are interested in exploring the connectivity properties of $\mathcal{M}_{\mathcal{X}_{free}}$. We conjecture that $\mathcal{M}_{\mathcal{X}_{free}}$ is always connected and that a retract exists from $\mathcal{X}_{free}$ onto $\mathcal{M}_{\mathcal{X}_{free}}$.

IV. APPROXIMATING CLEARANCE

The core operation of medial-axis approaches is determining the clearance of a configuration, see Section II-C, as a side effect of the computation, a witness configuration is easily obtained. An exact algorithm for our definition of clearance in $\mathcal{X}_{free}$, Definition 1, would necessitate modeling and solving instances of high-dimensional boundary value problems. However, even modeling $\partial\mathcal{X}_{obst}$ is computationally difficult. Instead, we propose an approximate algorithm for computing clearance and a witness of a state similar in spirit to prior work [25].

Algorithm 1 Approximate Clearance of a State CLR$(\cdot)$

Input: State x, number of trajectories n, probability p, time-step Δt
Output: Approximate clearance and witness

Trajectories sequence $\Pi = (\pi_1, \ldots, \pi_n) \leftarrow (x, \ldots, x)$
Controls $U = (u_1, \ldots, u_n)$
for $i \leftarrow 1$ **to** n **do** ◁ Initialze each trajectory
  $u_i \leftarrow$ RANDOM_CONTROL()
Boolean $col \leftarrow$ FALSE
Step $t \leftarrow 0$
repeat
  $t \leftarrow t + 1$
  for $i \leftarrow 1$ **to** n **do** ◁ Step along each trajectory
    $\pi_i^t \leftarrow f(\pi_i^{t-1}, u_i, \Delta t)$
    $col \leftarrow col \vee$ IN_COLLISION(π_i^t)
    if RAND() $< p$ **then**
      $u_i \leftarrow random\_control()$
until col
Witness trajectory $\pi^\dagger =$ FIRST_COLLISION(Π)
return $\left(\text{LENGTH}(\pi^\dagger), \pi^{\dagger t}\right)$

To approximate the clearance of a state, n random bang-bang trajectories are chosen and analyzed for the one that collides with $\partial\mathcal{X}_{obst}$ first in time. The algorithm is shown in Algorithm 1. The algorithm's complexity lies in an incremental evaluation of the candidate trajectory set. To initialize each trajectory π_i a control u_i is chosen based on a policy, either fully random or optimal. Then, at each iteration t, the next state in time is computed through numerical integration of a time step Δt, e.g., using Fourth Order Runge-Kutta, and checked for validity. Additionally, each trajectory has a small probability p to switch controls during the iteration, based upon the same policy as initialization. p defines a binomial distribution allowing variable time for switching controls. Once a collision is found, a binary search, FIRST_COLLISION$(\cdot)$, is performed between the last two states of a trajectory for a more precise computation of the clearance. If more than one trajectory collides in the same

iteration, the binary search reveals the shorter trajectory. An example execution of this algorithm is shown in Figure 4.

Overall, the complexity of this approach requires $O(nt)$ collision checks per execution, where t is the number of steps before the first trajectory collision. This can be improved by using a binary stepping approach to $O(n \log t)$, but requires more extensive book-keeping to ensure that no obstacles are jumped over. In this paper, we implemented the more expensive approach. Further, we note that this is a Monte-Carlo algorithm. This implies that a greater number of input trajectories will increase the accuracy of the approximation — under mild assumptions on the smoothness of $\partial \mathcal{X}_{obst}$.

V. APPROXIMATING SAMPLING OF $\mathcal{M}_{\mathcal{X}_{free}}$

One of the primary goals of this work is generating samples on $\mathcal{M}_{\mathcal{X}_{free}}$. Our approach is based on the Uniform Medial Axis Probabilistic RoadMap (UMAPRM) technique [26]. However, we had to extend this algorithm in a significant way because UMAPRM was never built for approximate medial-axis detection. We decided on this approach because of its properties — namely it uniformly samples a medial-axis structure. The main idea of this approach is to sample and discretize a random line segment in state space, and then look at properties of adjacent samples to detect a medial axis crossing.

The algorithm is shown in Algorithm 2. It begins the same as UMAPRM by sampling and discretizing a random line segment in state space into a sequence of states, X. After, the clearance of each state $x_i \in X$ is approximated using Algorithm 1. Then, passing over these clearance values, the algorithm searches for one of two things: (1) a 'big' jump in clearance value, representing a possible detection of a jump discontinuity, and (2) a 'plateau' in the clearance value, representing a local maximum of clearance. After, the adjacent pairs of states are further analyzed through a binary search to yield samples very close to the medial axis. In case (2), this will always be successful. However, in case (1), a false detection might have been made and false positives are further filtered out at this point of the algorithm. The final set of samples is returned. An illustration of the algorithm is shown in Figure 5.

The overall time complexity of this operation will be $O(n)$ approximate clearance computations, where $n = \frac{l}{r}$, i.e., the number of discretized states along each line segment. In combination with the analysis of clearance, this yields an approximately cubic time algorithm for generating samples. Clearly, this is inefficient and there is significant room for improvement. Our contribution in this work is theoretical in nature, and we pose these algorithms as a first step to sampling on or near $\mathcal{M}_{\mathcal{X}_{free}}$. In the future, we will invest more heavily in developing efficient approximation schemes of both clearance and sampling.

VI. EXPERIMENTAL ANALYSIS

Because our method is currently of theoretical interest instead of direct application, we provide a simple experimental validation of our theory. Specifically, our theoretical analysis

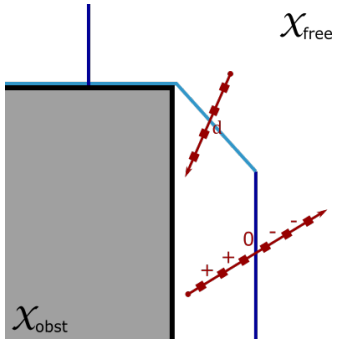

Fig. 5. Galaga system with an example $\mathcal{M}_{\mathcal{X}_{free}}$ structure shown in blue. Jump discontinuities are highlighted in light blue. Line segments are randomly sampled and discretized (red). Clearance deltas (positive, negative, zero, or above a threshold d) are used to find the medial axis crossings.

Algorithm 2 Sample $\mathcal{M}_{\mathcal{X}_{free}}$

Input: Sample length l, step length r, plateau size k, discontinuity threshold d
Output: Set of samples M on, or near, $\mathcal{M}_{\mathcal{X}_{free}}$

State $x \leftarrow$ RANDOM_STATE()
Direction $\vec{v} \leftarrow$ RANDOM_DIRECTION()
Sequence size $n \leftarrow \lceil \frac{l}{r} \rceil$
States $X = (x_1, \ldots x_n) \leftarrow$ SAMPLE_LINE$(x, x + l\vec{v})$
Clearances $C = (c_1, \ldots, c_n) \leftarrow (\text{CLR}(x_1), \ldots, \text{CLR}(n))$
Candidate pairs set $M \leftarrow \emptyset$
for $i \leftarrow 1$ **to** $n - 1$ **do** ◁ Discontinuities
  if $|c_i - c_{i+1}| > d$ **then**
    $M \leftarrow M \cup \{(x_i, x_{i+1})\}$
for $i \leftarrow k + 1$ **to** $n - k$ **do** ◁ Plateaus
  if IS_PLATEAU(X, C, i, k) **then**
    $M \leftarrow M \cup \{(x_{i-1}, x_{i+1})\}$
return BINARY_SEARCH_FOR_XMA(M)

indicates that our sampling mechanism for generating states on $\mathcal{M}_{\mathcal{X}_{free}}$ will yield higher clearance samples than uniform random sampling of $\mathcal{X}_{free}$.

A. Setup

We implemented our medial axis sampling routine, denoted with `MA`, and uniform random sampling of state space, denoted with `Uniform` in C++ using the GNU gcc compiler version 7.4.0. All experiments were run on Ubuntu 18.04 with an Intel® Core™ i5-8500 CPU at 2.8 GHz and 16 GB of RAM.

For our simple study, we validated our theoretical insights using our `Galaga` motion model and the second-order system similar to an aircraft, denoted `Airplane`, see Section III for details of their kinematic models. Both robots are represented as small cubes mimicking point sizes. We deployed both systems in two simulated environments varying complexity, see Figure 6. Their details are as follows:

1) `Simple` (Figure 6(a)) – a single box on the left of an infinite corridor creating an opening. Our theoretical analysis demonstrated that $\mathcal{M}_{pts}$ is disconnected and inequivalent to $\mathcal{M}_{\mathcal{X}_{free}}$.

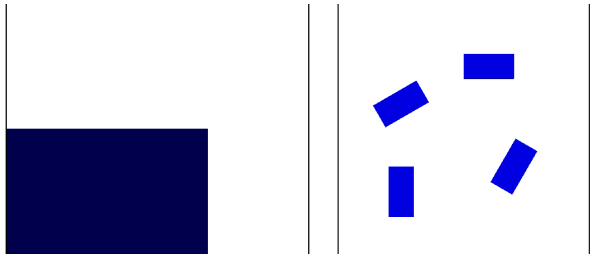

(a) `Simple` (b) `Complex`

Fig. 6. Environments used to evaluate clearance algorithms. (a) A simple block and (b) A scattered arrangement of rectangular obstacles in infinite corridors (left and right bound shown in black. The sampling area is the visible area, top to bottom.

2) `Complex` (Figure 6(b)) – four rectangular obstacles in an infinite corridor providing a much more complicated $\mathcal{M}_{\mathcal{X}_{free}}$ structure for the systems.

Sampling areas for the infinite corridors are shown in the figures.

We used a weighted Euclidean distance metric to measure distance favoring positional components of a state to velocity components. To measure clearance for our sampling mechanism (Algorithm 1), we used 25 random trajectories and a probability $p = 0.02$ using either optimal controls (denoted `MAOptimal`) or random controls (denoted `MARandom`). The time step of all numerical integration is $\Delta t = 0.05$ seconds. Our sampler uses a length $l = 4$ and a step length r the size of an environmentally dependent resolution. Its other parameters of plateau size k and discontinuity threshold d are tuned to the environment-system combination.

For all experiment-system combinations, we invoked the sampler routine 2500 times. Afterward, we measure the successful samples, execution time, and the average clearance of the successful samples. To measure the clearance, we used 50 random trajectories and a probability $p = 0.05$ of using optimal controls — we chose these values to acquire a more accurate measure of the sampler's clearance. All results are averaged over 10 runs. Avg clearance is shown in Figure 7.

B. Discussion

Analysis of our primary metric in this study, clearance, shown in Figure 7, clearly demonstrates that our medial sampler achieves between 1.25× and 1.75× improvement compared to `Uniform` across the environment-system combinations. Further, there is minimal differences between `MAOptimal` and `MARandom`. A deeper glance at which detection component generates samples in the MA methods, i.e., either a discontinuity or plateau, conveys that the `MARandom` approach has a bit more noise and detects more discontinuities. However, in environments like `Complex`, `MAOptimal` can sometimes randomly miss its closest witness state and produce false positives. In the future, we would like to perform a deeper study into the false-positive rate and false-negative rate of our detection heuristics. This should provide further insight into designing a better clearance and sampling algorithms in non-holonomic systems.

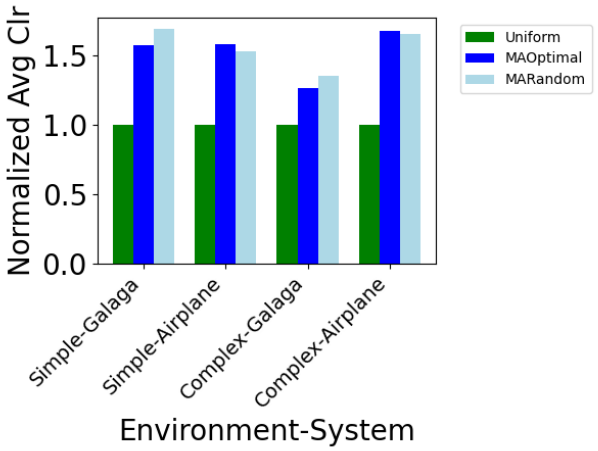

Fig. 7. Average sample clearance fpr `MAOptimal` (blue) and `MARandom` (light blue) normalized to `Uniform` (green) for each of the four environment-system combinations.

We omit reporting exact time figures, as this is not the focus of the paper. However, for reference, the method is approximately five orders of magnitude slower than uniform random sampling. We remind the reader that this is not surprising as a uniform random sampling requires a single CD invocation per attempt, while our method requires a Monte-Carlo approximation of clearance at many states along a line segment in $\mathcal{X}_{free}$. We plan to tackle the efficiency and applicability of our sampler in the future.

In terms of success rates of the samplers, while we omit exact numbers, uniform sampling is approximately 2.5 times as successful as our medial axis sampler. This is mainly due to the probability of the underlying approach to our sampler — sampling line segments of $\mathcal{X}_{free}$. Our success rate is coupled to the probability of sampling a random line segment that intersects $\mathcal{M}_{\mathcal{X}_{free}}$. This will almost always be smaller than the probability of uniform random sampling — the ratio of the hypervolume of $\mathcal{X}_{free}$ to the hypervolume of $\mathcal{X}_{obst}$.

Our experiment successfully validates our theoretical contribution and supports our hypothesis. Specifically, we showed that sampling $\mathcal{M}_{\mathcal{X}_{free}}$ will yield high clearance states, which supports our conjecture that this is the ideal structure for representing the collection of 'safest' states.

VII. CONCLUSION

In this work, we have proposed and explored definitions for a medial axis of free space for nonholonomic robots. Specifically, we showed the inequivalence of classical definitions of a medial axis, and we proved that this holds for non-STLC systems. Further, we proposed algorithms for approximating the clearance of a state in state space, and approximately sampling our proposed medial axis $\mathcal{M}_{\mathcal{X}_{free}}$. We validated our algorithms in a variety of settings to show that the states sampled truly improve upon the clearance.

In the future, we will explore the connectivity properties of our $\mathcal{M}_{\mathcal{X}_{free}}$ definition. We will also look for more efficient

algorithmic approaches to sampling and exploiting $\mathcal{M}_{\mathcal{X}_{free}}$ for use in sampling-based motion planners.

ACKNOWLEDGEMENT

We would like to acknowledge the members, past and present, of the Spider Robotics Lab (SpiRoL) at the University of Richmond for their assistance on this work.

REFERENCES

[1] M. Y. Jung and P. Kazanzides, "An architectural approach to safety of component-based robotic systems," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 3360–3366.

[2] J. Chen, Y. Zhou, Q. Lv, K. K. Deveerasetty, and H. Ugochi Dike, "A review of autonomous obstacle avoidance technology for multi-rotor uavs," in *2018 IEEE International Conference on Information and Automation (ICIA)*, 2018, pp. 244–249.

[3] R. Geraerts, "Planning short paths with clearance using explicit corridors," in *2010 IEEE International Conference on Robotics and Automation*, 2010, pp. 1997–2004.

[4] P. A. Lasota, T. Song, and J. A. Shah, *A Survey of Methods for Safe Human-Robot Interaction*, 2017.

[5] H. Blum, "A transformation for extracting new descriptors of shape," in *Models for the Perception of Speech and Visual Form*, W. W. Dunn, Ed. MIT Press, 1967, pp. 362–381.

[6] S. A. Wilmarth, N. M. Amato, and P. F. Stiller, "Maprm: a probabilistic roadmap planner with sampling on the medial axis of the free space," in *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C)*, vol. 2, May 1999, pp. 1024–1031 vol.2.

[7] T. Lozano-Pérez and M. A. Wesley, "An algorithm for planning collision-free paths among polyhedral obstacles," *Communications of the ACM*, vol. 22, no. 10, pp. 560–570, October 1979.

[8] J. H. Reif, "Complexity of the mover's problem and generalizations," in *Proc. IEEE Symp. Foundations of Computer Science (FOCS)*, San Juan, Puerto Rico, October 1979, pp. 421–427.

[9] L. E. Kavraki, P. Švestka, J. C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Trans. Robot. Automat.*, vol. 12, no. 4, pp. 566–580, August 1996.

[10] N. M. Amato, O. B. Bayazit, L. K. Dale, C. Jones, and D. Vallejo, "OBPRM: an obstacle-based PRM for 3D workspaces," in *Proc. Int. Wksp. Alg. Found. Robot. (WAFR)*. Natick, MA, USA: A. K. Peters, Ltd., 1998, pp. 155–168.

[11] V. Boor, M. H. Overmars, and A. F. van der Stappen, "The gaussian sampling strategy for probabilistic roadmap planners," in *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C)*, vol. 2, May 1999, pp. 1018–1023 vol.2.

[12] J. Denny and N. M. Amato, "Toggle PRM: A coordinated mapping of C-free and C-obstacle in arbitrary dimension," in *Alg. Found. Robot. X*. Springer, 2013, pp. 297–312, (WAFR '12).

[13] H. Yeh, S. Thomas, D. Eppstein, and N. M. Amato, "UOBPRM: A uniformly distributed obstacle-based PRM," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct 2012, pp. 2655–2662.

[14] F. Aurenhammer, "Voronoi diagrams: A survey of a fundamental geometric data structure," *ACM Comput. Surv.*, vol. 23, no. 3, pp. 345–405, Sep. 1991.

[15] J.-C. Latombe, *Robot Motion Planning*. Boston, MA: Kluwer Academic Publishers, 1991.

[16] C. J. Ong, S. S. Keerthi, E. Huang, and E. G. Gilbert, "Equidistance diagram — a new roadmap for path plannning," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 1999, pp. 682–687.

[17] H. Choset and J. Burdick, "Sensor-based exploration: The hierarchial generalized voronoi graph," *Int. J. Robot. Res.*, vol. 19, no. 2, pp. 96–125, 2000.

[18] M. Foskey, M. Garber, M. C. Lin, and D. Manocha, "A voronoi-based hybrid motion planner for rigid bodies," in *Proc. IEEE Int. Conf. Intel. Rob. Syst. (IROS)*, 2001, pp. 55–60.

[19] L. Guibas, C. Holleman, and L. Kavraki, "A probabilistic roadmap planner for flexible objects with a workspace medial-axis-based sampling approach," in *Proc. IEEE Int. Conf. Intel. Rob. Syst. (IROS)*, vol. 1, 1999, pp. 254–259.

[20] C. Holleman and L. E. Kavraki, "A framework for using the workspace medial axis in prm planners," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, vol. 2, San Franasisco, CA, 2000, pp. 1408–1413.

[21] Y. Yang and O. Brock, "Adapting the sampling distribution in prm planners based on an approximated medial axis," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, vol. 5, 2004, pp. 4405–4410.

[22] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.

[23] J. Denny, E. Greco, S. Thomas, and N. M. Amato, "MARRT: Medial axis biased rapidly-exploring random trees," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Hong Kong, China, May 2014, pp. 90–97.

[24] L. Jaillet, J. Cortés, and T. Siméon, "Sampling-based path planning on configuration-space costmaps," *Trans. Rob.*, vol. 26, no. 4, pp. 635–646, Aug. 2010.

[25] Jyh-Ming Lien, S. L. Thomas, and N. M. Amato, "A general framework for sampling on the medial axis of the free space," in *2003 IEEE International Conference on Robotics and Automation (Cat. No.03CH37422)*, vol. 3, Sep. 2003, pp. 4439–4444 vol.3.

[26] H. C. Yeh, J. Denny, A. Lindsey, S. Thomas, and N. M. Amato, "Umaprm: Uniformly sampling the medial axis," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 5798–5803.

[27] Guilin Liu and J. Lien, "Fast medial-axis approximation via max-margin pushing," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2015, pp. 3262–3267.

[28] A. C. Shkolnik and R. Tedrake, "Sample-based planning with volumes in configuration space," *ArXiv*, vol. abs/1109.3145, 2011.

[29] R. Geraerts and M. H. Overmars, "Creating high-quality paths for motion planning," *Int. J. Robot. Res.*, vol. 26, no. 8, pp. 845–863, 2007.