

Mesh Manifold Based Riemannian Motion Planning for Omnidirectional Micro Aerial Vehicles

Michael Pantic , Lionel Ott , Cesar Cadena , *Member, IEEE*, Roland Siegwart , *Fellow, IEEE*, and Juan Nieto , *Member, IEEE*

Abstract—This letter presents a novel on-line path planning method that enables aerial robots to interact with surfaces. We present a solution to the problem of finding trajectories that drive a robot towards a surface and move along it. Triangular meshes are used as a surface map representation that is free of fixed discretization and allows for very large workspaces. We propose to leverage planar parametrization methods to obtain a lower-dimensional topologically equivalent representation of the original surface. Furthermore, we interpret the original surface and its lower-dimensional representation as manifold approximations that allow the use of Riemannian Motion Policies (RMPs), resulting in an efficient, versatile, and elegant motion generation framework. We compare against several Rapidly-exploring Random Tree (RRT) planners, a customized CHOMP variant, and the discrete geodesic algorithm. Using extensive simulations on real-world data we show that the proposed planner can reliably plan high-quality near-optimal trajectories at minimal computational cost. The accompanying multimedia attachment demonstrates feasibility on a real OMAV. The obtained paths show less than 10% deviation from the theoretical optimum while facilitating reactive re-planning at kHz refresh rates, enabling flying robots to perform motion planning for interaction with complex surfaces.

Index Terms—Aerial systems, motion and path planning.

I. INTRODUCTION

RECENT developments in *micro aerial vehicle* (MAV) that can interact with and exert forces on the environment make a variety of new use cases such as contact inspection, spraying, and painting, to name a few, possible. The advent of *omnidirectional micro aerial vehicle* (OMAV) allows these interactions to take place in arbitrary orientations [1], which was previously impossible with traditional MAV due to their underactuated nature. Current planning methods [2] often use a discretized map such as an octree or a voxel grid to store occupancy information and plan collision-free trajectories using




Fig. 1. Left: Illustration of a typical planning run with all planners on the hilo scenario. Red is the proposed, green the DGEO, and brown the RRT*-Con planner. Blue shades correspond to the RRT*-Sam, yellow shades to the RRT*-Pro, and pink to CHOMP. Right: Example of an OMAV inspecting a surface using the proposed planner.

sampling-based or optimization-based algorithms. There is a large range of literature about planning through unknown spaces in continuously updating discretized maps while avoiding obstacles. In these use cases, the exact spatial location of the trajectory is of less importance as long as it remains collision-free. Using these existing methods for aerial interaction on surfaces with an OMAV can be cumbersome and ill-posed as the requirements are completely different. One problem is the scalability of map representations currently used for MAV motion planning. Three dimensional fixed-discretization map representations do not scale well in terms of resolution or map size. OMAVs however are able to perform centimeter-level manipulation tasks on very large a priori known workspaces that can be hundreds of meters in size, giving rise to the need for a very efficient high-resolution surface representation and an associated planning method. In most industrial, inspection, and outdoor use-cases, maps can be obtained up-front, which makes on-line mapping less important. Another important challenge is the trajectory generation itself, as aerial interaction takes place in relation to a surface rather than free space. Planning a high-precision path along a surface should exploit connectedness of the surface, which implicit representations such as *Truncated Signed Distance Fields* (TSDFs) are unable to do. Furthermore, the determination and planning of the relative orientation of an OMAV with respect to a surface is challenging in the presence of gimbal lock and singularities. These problems can be mitigated by using a surface map that contains a notion of orientation within the surface that is always valid.

A. Related Work

Many planning approaches designed specifically for MAV use a smooth spline representation of the trajectory and exploit

Manuscript received October 15, 2020; accepted February 6, 2021. Date of publication February 24, 2021; date of current version April 14, 2021. This letter was recommended for publication by Associate Editor N. Amato and Editor S. Thomas upon evaluation of the reviewers' comments. This work was supported in part by the National Center of Competence in Research (NCCR) on Digital Fabrication, and in part by NCCR Robotics through the Swiss National Science Foundation. The work of Juan Nieto was conducted when the author was with ETH Zürich. (*Corresponding author: Michael Pantic.*)

Michael Pantic, Lionel Ott, Cesar Cadena, and Roland Siegwart are with the Autonomous Systems Lab, ETH Zürich, 8092 Zürich, Switzerland (e-mail: mpantic@ethz.ch; lionel.ott@mavt.ethz.ch; cesarcadena.lerma@gmail.com; rsiegwart@ethz.ch).

Juan Nieto is with the Mixed Reality & AI Lab, Microsoft, 8001 Zürich, Switzerland, work conducted when the author was with ETH Zürich (e-mail: jnieto@ethz.ch).

Digital Object Identifier 10.1109/LRA.2021.3061869

their differential flatness. Typical examples are spline motion primitives [3] and optimization-based spline planning [4]. These methods generally assume that the MAV moves in free-space. When obstacles are to be avoided, optimization-based techniques that use collision gradients such as CHOMP [5] are popular. Similarly, collision gradients are used in [2] to optimize MAV trajectory splines in order to obtain safe free-space trajectories. Another option are sampling-based planners, such as rapidly-exploring random trees (RRT) [6] and its variants (e.g. RRT* [7], RRT-Connect [8]), to find obstacle-free paths. However, trajectory planning for aerial inspection and manipulation should not just avoid obstacles but actively follow surfaces and objects. Sampling-based planners can be used for this by constraining the sampling space. [9] uses an RRT planner that samples in a task space constrained by all admissible end-effector states. Another option is to sample unconstrained but project onto the closest admissible space [10]. While obstacle avoidance can allow inaccuracies as long as they are safe, surface following or interaction quality is more dependent on the surface representation quality and resolution. Implicit surface representations store occupancy information that allows the recovery of the surface. Common examples are octrees [11] and hash-based voxel grids [12]. While fixed-discretized representations are very efficient for retrieval and lookup of data, their practical scalability in terms of resolution and workspace is constrained by memory usage. Explicit surface representations directly store the boundary between occupied and free space. For 3D applications, common techniques are manifold splines [13], *Non-uniform rational B-Splines* (NURBs) [14] and triangular meshes. A major advantage of these representations is their geometric nature which does not rely on a fixed discretization or a fixed resolution and allows efficient exploitation of geometric neighborhood connectedness. Surfels [15] and point clouds are often used as a surface representation that is easy to obtain from sensor data. However, they do not encode surface connectivity, which could be exploited by planning algorithms. [16] uses triangular meshes for inspection planning but does not exploit connectedness, and instead samples states based on triangle normals and centers. Similarly, [17] uses meshes to represent free space without relying on fixed discretization and to extract traversability information, without explicitly using the connected nature of the mesh. The interpretation of a surface mesh as an approximation of a manifold yields a natural mathematical parametrization of the connectedness. A typical example is the calculation of discrete geodesics [18], which are defined as the shortest path between two points on a mesh manifold. The work in [19] introduces trajectory optimization on Riemannian Manifolds for obstacle-avoidance and field-of-view aware planning for MAVs but uses analytically represented manifolds. Another class of planners that exploit the manifold structure of task and configuration space are *Riemannian motion policies* (RMPs) [20]. By optimally combining multiple motion policies across manifolds, RMPs enable the formulation of planning and control problems in arbitrary Riemannian manifolds as long as a Jacobian that relates them locally is obtainable.

B. Contributions

In this letter, we propose a solution to the problem of efficiently planning trajectories that follow and/or approach a-priori known surfaces. We use triangular meshes as a surface representation that is not limited by discretization or resolution. To make use of this representation for interaction planning, we propose

a motion generation algorithm that is based on differential geometry principles and Riemannian motion policies [20]. Our contributions are the following:

- The formulation of an efficient $2D - 3D$ mesh manifold parametrization for planning.
- A highly efficient path planning framework based on mesh manifolds and Riemannian motion policies.
- Extensive experimental verification in simulation and comparison to a variety of sampling-based and optimization-based planning algorithms for surface following.

While in this letter we concentrate on the surface following use-case for OMAVs, the proposed approach of using meshes as an approximation to a mathematical manifold for RMP-based planning can be generalized to a variety of other applications, such as explicitly encoding nullspace manifolds.

C. High-Level System Overview

The concept of the proposed planner is easiest understood by a well-known analogy from cartography. In our daily lives we generally do not care about the spherical shape of our world. Instead, we use flattened, cartesian representations to plan a trajectory to a desired goal. Analogous, we flatten the 3D mesh surface map to a lower-dimensional representation using a suitable parametrization function. We then exploit the topological equivalence (*homeomorphism*) between the original surface and the flattened representation to induce acceleration fields as Riemannian motion policies that generate the desired trajectories along and towards the surface in 3D. Thereby we use the explicit surface encoding of a triangular mesh efficiently. In the remainder of this letter we detail the manifold approximation using meshes (Section II) and present an RMP planner using this approximation (Section III). We show extensive evaluations (Section IV) and discuss the results (Section V).

II. MESH MANIFOLDS

The defining property of our framework is the use of conventional triangular meshes as a computationally tractable approximation to a smooth Riemannian manifold that represents a surface in \mathbb{R}^3 . By using the mesh representation which is inherently geometric and has no fixed discretization, our system is not constrained to a specific resolution or extent of the map. In this section, we lay out the theoretical background of making the surface connectedness of a mesh easily accessible to the planner. We do so by interpreting the mesh as an approximation to a manifold embedded in a higher-dimensional space and obtaining a flattened, axis-aligned homeomorphic representation where the two axes of “*going along the surface*” are perpendicular to the remaining axis of moving “*towards or away*” of the surface. Such a mapping also provides explicit orientation information on the surface, as the angle with respect to the axes along the surface is always defined. To ensure the correctness of the proposed planner, we show that the approximations made are valid and additionally give strategies for computationally efficient implementation.

A. Prerequisites and Definitions

In the following, \mathcal{M} denotes a surface embedded in \mathbb{R}^d . Here, we refer to the mathematical definition of a surface - a $(d - 1)$ -dimensional manifold. We assume \mathcal{M} to be smooth, free of holes, have a defined boundary, and be homeomorphic to a disc (no self-intersections). We require \mathcal{M} to be Riemannian,

i.e. to have a smoothly varying positive-definite inner product on the tangent space at every point. While the smoothness criterion is needed for a sound theoretical base, the proposed algorithm works relatively well on not perfectly smooth geometries as demonstrated in the results section. Without loss of generality, we assume \mathcal{M} to be embedded in \mathbb{R}^3 in the following. The homeomorphic 2D representation of \mathcal{M} is subsequently denoted by \mathbf{m} . Intuitively, we use the flattened representation \mathbf{m} of the 3D surface \mathcal{M} , together with a one-to-one mapping between them, as a 2D coordinate representation that is embedded in the surface. We use triangular meshes as a surface representation that approximates the properties of an ideal manifold. To distinguish the ideal, mathematical manifold from its approximate triangular mesh implementation, we subsequently refer to the mesh as \mathcal{M}_Δ respectively \mathbf{m}_Δ .

B. Notation

We denote a point i that is part of the 3D mesh \mathcal{M}_Δ as $P_i = (x_i, y_i, z_i) \in \mathbb{R}^3$. Similarly, a point j on the 2D mesh \mathbf{m}_Δ is denoted by $p_j = (u_j, v_j) \in \mathbb{R}^2$. It is important to note that P_i , respectively p_j , refer to an arbitrary point that is part of the mesh surface. Vertices are a subset of these points and in the following referred to as \bar{P}_i , respectively \bar{p}_j . The mesh structure is formalized as the set of all vertices $V(\mathcal{M}_\Delta) = \{\bar{P}_0, \dots, \bar{P}_N\}$, respectively $V(\mathbf{m}_\Delta) = \{\bar{p}_0, \dots, \bar{p}_n\}$. A triangle formed by vertices and edges is formalized as an ordered triplet of vertices $T_{ijk} = \{\bar{P}_i, \bar{P}_j, \bar{P}_k\}$ on the 3D mesh, respectively $t_{ijk} = \{\bar{p}_i, \bar{p}_j, \bar{p}_k\}$ on the 2D mesh. Any arbitrary point on the mesh is part of at least one triangle.

C. Coordinate Mapping

We obtain the 2D representation \mathbf{m}_Δ of \mathcal{M}_Δ by applying a parametrization \mathcal{H} that maps each 3D vertex to a 2D vertex while retaining topology:

$$\begin{aligned} \bar{p}_k &= \mathcal{H}(\bar{P}_k) \\ \bar{p}_k &\in V(\mathbf{m}_\Delta) \subset \mathbb{R}^2, \bar{P}_k \in V(\mathcal{M}_\Delta) \subset \mathbb{R}^3, \end{aligned} \quad (1)$$

We assume \mathcal{H} to be a one-to-one mapping and to generate a valid triangulation \mathbf{m}_Δ that is topologically equivalent to \mathcal{M}_Δ . To obtain a mapping of an arbitrary point $p \in \mathbf{m}_\Delta$ to their corresponding point $P \in \mathcal{M}_\Delta$ and vice versa, we use standard barycentric coordinates [21]. Assuming $p \in \mathbf{m}_\Delta$ lies in triangle t_{ijk} we obtain the corresponding point $P \in \mathcal{M}_\Delta$ by first calculating the barycentric coordinates $\mathcal{B}(p, t_{ijk}) = \{\beta_1, \beta_2, \beta_3\}$, $\beta_{1..3} \in \mathbb{R}$ of p relative to 2D triangle t_{ijk} and then applying these barycentric coordinates to the corresponding 3D triangle T_{ijk} :

$$P \in \mathcal{M}_\Delta = \Gamma(p, t_{ijk}, T_{ijk}) = \mathcal{B}^{-1}(\mathcal{B}(p, t_{ijk}), T_{ijk}) \quad (2)$$

Note: Γ is a short-hand notation for the full mapping defined by \mathcal{B} and \mathcal{B}^{-1} . Figure 2 visualizes the mapping process. Additionally to the 2D-3D coordinate mapping we augment the 2D coordinates of \mathbf{m}_Δ with a third dimension h that is normal to the surface. For any triangle on \mathcal{M}_Δ , this dimension coincides with the normal of that triangle.




Fig. 2. Closeup of a pair of triangles on both mesh manifolds that shows the used notation. The left triangles are three dimensional while the right triangles are two dimensional. The two points P_i and p_i have the same barycentric position w.r.t. the orange triangle. The vectors V_i and v_i are related to each other by the Jacobian $J_{\mathbf{m}_\Delta \leftarrow \mathcal{M}_\Delta}$.

D. Selection of Mapping Function

One important property needed for this planner is a surface flattening method that provides a bijective one-to-one mapping \mathcal{H} . Algorithms such as *Tutte Barycentric Mapping* [22], *Discrete Authalic Parametrization* [23] and *Floater Mean Value Coordinates* [24] provide such a mapping. All three parametrizations can generate 2D homeomorphic discs with either square or circular boundaries. Furthermore, as in regular navigation a conformal (angle-preserving) mapping is strongly preferred as otherwise mappings of velocities and directions between 2D and 3D can be inaccurate. Only the *Floater Mean Value Coordinates* explicitly optimize for low angular distortion and are therefore chosen as our default implementation of \mathcal{H} . We use circular border parametrizations, as planners operating on square borders showed divergent behavior for points on the boundaries. A more detailed overview of the different methods can be found in [25].

E. Implementation

To facilitate efficient and fast translation of arbitrary points between \mathcal{M}_Δ and \mathbf{m}_Δ , the result of \mathcal{H} is computed for every vertex \bar{P} on startup and cached in forward and reverse hashmaps. This allows translation of vertices with a computational complexity of $\mathcal{O}(1)$. In order to obtain the closest triangle to an arbitrary point, we use accelerated queries in axis-aligned bounding-box (AABB) [26] trees for both \mathcal{M}_Δ and \mathbf{m}_Δ .

III. PLANNING FRAMEWORK

In this section, we explain how the previously described mesh manifolds are used to plan paths relative to a surface. Our framework is based on Riemannian Motion Policies (RMPs) [20]. The proposed planner is efficient and well defined due to two properties that follow from the mesh manifold approximations. Firstly, by exploiting the 2D nature of a surface embedded in 3D space we effectively reduce the dimensionality of the problem and thus mitigate scaling effects in terms of computational complexity and memory usage. Secondly, having a gravity-independent and singularity-free orientation on the surface allows the construction of a valid orientation of the OMAV at any point on the surface.

A. Riemannian Motion Policies

Riemannian Motion Policies (RMPs) [20] provide a framework to formulate and combine multiple motion policies on

different manifolds. In the following, a brief summary is given. Without loss of generality, we assume to have only two different manifolds, configuration space \mathcal{Q} and task space \mathcal{X} . We denote positions and their derivatives in configuration space as q, \dot{q}, \ddot{q} and similarly for the task space x, \dot{x}, \ddot{x} . We convert between the two spaces using a task-map $\phi(q) = x$. Additionally, a position-dependent analytic Jacobian $J_{\mathcal{X} \leftarrow \mathcal{Q}} = \frac{\partial \phi}{\partial q}$ that maps velocities and accelerations *locally* between the two spaces is needed. A Riemannian Motion Policy $\mathcal{P}_{\mathcal{X}}$ in the task space is defined as the tuple $(f, A)_{\mathcal{X}}$, where f is an acceleration function $f(x, \dot{x})$ and $A(x, \dot{x})$ the smoothly varying, positive semidefinite Riemannian metric associated with the policy. As shown in [20], two policies can be summed as a metric-weighted average

$$\mathcal{P}_{\mathcal{X}}^a + \mathcal{P}_{\mathcal{X}}^b = ((A^a + A^b)^+ (A^a f^a + A^b f^b), A^a + A^b)_{\mathcal{X}} \quad (3)$$

to provide an optimal solution for the combined system. This property is very powerful, as it allows the decomposition of complex problems into many simple policies that can be combined optimally.

Furthermore, to transform a policy from one space to another, the pullback operator is used:

$$\text{pull}_{\mathcal{Q}}((f, A)_{\mathcal{X}}) = ((J^T A J)^+ J^T A f, J^T A J)_{\mathcal{Q}}. \quad (4)$$

The result is an acceleration in \mathcal{Q} that executes a policy (or combination thereof) defined in taskspace \mathcal{X} .

B. Mesh Manifold as Task Space

We use $\mathcal{M}_{\Delta} \in \mathbb{R}^3$ as the configuration space \mathcal{Q} and \mathbf{m}_{Δ} as the task space \mathcal{X} . While in many applications, such as dexterous manipulation, the configuration space has a higher dimensionality than the task space, here we exploit the space transformation capabilities of RMPs to simplify the 3D planning problem and exploit surface connectedness. Due to the induced coordinate mapping on the surface, it is straightforward to generate spatiotemporal trajectories that follow the induced acceleration field on the surface and map them into \mathbb{R}^3 . We use the coordinate mapping defined in the previous section as task map and obtain the needed Jacobian $J_{\mathbf{m}_{\Delta} \leftarrow \mathcal{M}_{\Delta}} \in \mathbb{R}^{3 \times 3}$ analytically by exploiting the bijective triangle mapping between \mathcal{M}_{Δ} and \mathbf{m}_{Δ} and the mapping of barycentric coordinates for arbitrary points on the two meshes. By taking the partial derivatives of eq. (2) for all dimensions, we obtain the first 2 columns of $J_{\mathbf{m}_{\Delta} \leftarrow \mathcal{M}_{\Delta}}$. The last column follows from the definition of the h axis and corresponds to the normalized normal of the 3D triangle. For a point $P \in \mathcal{M}_{\Delta}$ that is an element of triangle T with the corresponding triangle $t \in \mathbf{m}_{\Delta}$, we obtain the Jacobian as follows:

$$J_{\mathbf{m}_{\Delta} \leftarrow \mathcal{M}_{\Delta}}(P) = \left[\frac{\partial}{\partial x} \Gamma(P, t, T) \quad \frac{\partial}{\partial y} \Gamma(P, t, T) \quad \mathcal{N}(T) \right] \quad (5)$$

where $\mathcal{N}(\cdot)$ returns the normalized normal vector for a triangle T . The analytical derivation of $J_{\mathbf{m}_{\Delta} \leftarrow \mathcal{M}_{\Delta}}$ is constant for all points P on a specific triangle T . This allows us to formulate policies that follow and approach \mathbf{m}_{Δ} , but execute them in 3D on an OMAV that operates on \mathcal{M}_{Δ} embedded in \mathbb{R}^3 .

C. Surface Attractor Policy

We decouple the surface following problem into two independent policies formulated on \mathbf{m}_{Δ} . The first policy, subsequently called $\mathcal{P}_{\mathbf{m}_{\Delta}}^{\perp}$, drives the trajectory onto the surface of \mathbf{m}_{Δ} . This policy corresponds to the generic attractor policy in [20] and is

defined as

$$\mathcal{P}_{\mathbf{m}_{\Delta}}^{\perp} = (f^{\perp}, A^{\perp}) \quad (6)$$

$$f^{\perp} = \alpha^{\perp} \cdot \mathbb{S}(\mathbf{0} - p_0) - \beta^{\perp} \dot{p}_0 = \ddot{p}_0 \quad (7)$$

$$A^{\perp} = \text{diag}(0, 0, 1) \in \mathbb{R}^{3 \times 3}, \quad (8)$$

where α^{\perp} and β^{\perp} are tuning parameters, and \mathbb{S} is the soft-normalization function

$$\mathbb{S}(z) = \frac{z}{|z| + \gamma \log(1 + \exp(\gamma|z|))} \quad (9)$$

with tuning parameter γ defined in [20], $p_0 \in \mathbf{m}_{\Delta}$ is the current position w.r.t. \mathbf{m}_{Δ} and \dot{p}_0 the current velocity. $\mathcal{P}_{\mathbf{m}_{\Delta}}^{\perp}$ induces an acceleration field in \mathbf{m}_{Δ} that points towards the surface. By pulling the policy from task space to configuration space using metric A^{\perp} , we constrain the policy to be only acting on the dimension perpendicular to the surface. The resulting acceleration field in \mathcal{M}_{Δ} respectively \mathbb{R}^3 smoothly drives the trajectory to a desired surface distance (usually 0).

D. Surface Following Policy

The second policy, $\mathcal{P}_{\mathbf{m}_{\Delta}}^{\rightarrow}$ drives the trajectory to a desired position $p_{des} = [u_{des}, v_{des}, h_{des}]$ on the surface. Due to the metric only u, v affect the execution of the policy.

$$\mathcal{P}_{\mathbf{m}_{\Delta}}^{\rightarrow} = (f^{\rightarrow}, A^{\rightarrow}) \quad (10)$$

$$f^{\rightarrow} = \alpha^{\rightarrow} \cdot \mathbb{S}(p_{des} - p_0) - \beta^{\rightarrow} \dot{p}_0 = \ddot{p}_0 \quad (11)$$

$$A^{\rightarrow} = \text{diag}(1, 1, 0) \in \mathbb{R}^{3 \times 3}, \quad (12)$$

The policy $\mathcal{P}_{\mathbf{m}_{\Delta}}^{\rightarrow}$, when pulled to the configuration space \mathcal{M}_{Δ} with metric A^{\rightarrow} , induces an acceleration field that follows the surface and drives the trajectory to the specified goal position.

E. Parameter Tuning

For each policy a parameter set $\mathcal{T} = \{\alpha, \beta, \gamma\}$ is needed. By varying $\alpha^{\perp}, \beta^{\perp}$ against α^{\rightarrow} and β^{\rightarrow} , the relative strength and aggressiveness of the two policies can be chosen. This has an especially large influence on the trajectory planning from free-space towards a goal on the surface. Depending on the relative strength, the contact with the surface is sought as soon as possible, gradually or towards the end of the trajectory. For all experiments we used the following values: $\mathcal{T}^{\rightarrow} = \{0.7, 13.6, 0.4\}$, $\mathcal{T}^{\perp} = \{20.0, 30.0, 0.01\}$. The tuning process is relatively intuitive, stable, and fast, as the results of a re-tuning can be visualized in real-time.

F. Orientation and Offset

At any location on \mathbf{m}_{Δ} we can trivially determine the u and v axis direction, as in any regular Cartesian coordinate system. This is a side benefit of using a planar parametrization of an explicit surface representation. By mapping these directions onto \mathcal{M}_{Δ} we obtain an on-surface orientation that is always defined, regardless of the actual 3D surface orientation. By calculating the normalized inverse of the Jacobian $J_{\mathbf{m}_{\Delta} \leftarrow \mathcal{M}_{\Delta}}$ at a specific point P , we obtain the on-surface orientation as column vectors. As an example, we derive the orientation R for aligning an OMAV body x -axis with direction tangent to the surface along u , and the body z -axis with the direction perpendicular to the surface. The resulting rotation matrix $R \in \mathbb{R}^3$ is therefore




Fig. 3. Rendering of all three scenarios, not to scale. *hilo* originates from a laser scan of a thin-shelled concrete roof and *rhone* is generated from high-res digital elevation maps of a glacier. *curve* is artificially generated and serves to demonstrate sampling bottlenecks and planning difficulties in highly variable geometry under ideal conditions.

TABLE I
METRIC SIZES AND MESHING DETAILS OF EACH SCENARIO

	Faces	Extent [m]	Surface [m ²]	Setup [s]
rhone	126.4k	625.9 × 463.1 × 88.6	160802.1	2.31
curve	20.0k	37.5 × 35.2 × 25.0	4234.1	0.18
hilo	16.1k	21.2 × 9.4 × 5.2	209.6	0.13

constructed as

$$R = \begin{bmatrix} -J_{\mathbf{m}_\Delta \leftarrow \mathcal{M}_\Delta}^{-1}[:, 1]^T \\ J_{\mathbf{m}_\Delta \leftarrow \mathcal{M}_\Delta}^{-1}[:, 1] \times J_{\mathbf{m}_\Delta \leftarrow \mathcal{M}_\Delta}^{-1}[:, 3]^T \\ J_{\mathbf{m}_\Delta \leftarrow \mathcal{M}_\Delta}^{-1}[:, 3]^T \end{bmatrix}^T, \quad (13)$$

where $[:, i]$ selects the i -th 1-based column. Due to the induced distortion by mapping \mathcal{H} , the 3D representation of the u and v axis are not necessarily orthogonal, therefore we only use the u axis and the normal and obtain the third direction by cross-product (second row in eq. 13). The mesh-based planning method applies identically for meshes constructed based on the original meshes. One use-case would be the inflation of the mesh with methods such as the one presented in [27] to guarantee a certain distance from the surface without self-intersections.

IV. EXPERIMENTS

We evaluate the proposed planning method in three different scenarios, shown in Fig. 3 and Table I. For each scenario, 100 planning tasks are generated by randomly sampling a start and end location on the surface and solved by each planner variant. The planner performances are evaluated in simulation, and feasibility demonstrated on a real OMAV.¹

A. Comparison Planners

We compare against a broad range of geometric, sampling-based, and optimization-based planners. All compared methods solve the same problem: to find the shortest path between two points while being constrained to the surface. As a geometric planner, we use the widely used theoretically-optimal (shortest path) discrete geodesic algorithm (subsequently abbreviated as DGEO) proposed in [18] and implemented in CGAL [28].

¹In the written paper we concentrate on the simulation experiments as the bare execution of the planned trajectory mostly measures controller performance and not planning quality. See accompanying multimedia attachment for videos of the OMAV experiments.

Three variants of RRT* implemented using OMPL [29] represent sampling-based planners, while CHOMP [5] is used as an optimization-based planner.² The variants of the RRT* algorithm differ in how they randomly generate new planning states (here, positions in 3D space). The *sampling* variant (called RRT*-Sam), leverages the explicit surface representation of the mesh to uniformly sample positions directly on the mesh surface. Transitions between states are considered valid if paths between positions are within 1 cm of the surface, verified at 5% steps along the path. This is a rather loose verification but needed to clear sampling bottlenecks (discussed in the next section). RRT*(*) is often used with task-space constrained samplers [9]. Here we use it as an example to show the behavior of constrained sampling in difficult geometries, as any extension needs to clear the same sampling bottlenecks. The *projecting* variant [10], (called RRT*-Pro) mimics planning using an implicit surface representation. Therefore, any direct addressing of the surface is impossible and sampling cannot be performed on the surface. Instead, states are randomly sampled in the encompassing 3D volume and are then projected onto the surface using a Jacobian. For RRT*-Pro all paths between states are considered valid, even if they leave the surface. However, the resulting path is smoothed and again projected back onto the surface using the Jacobian. Note that RRT*-Pro is evaluated on analytic geometries in [10], which makes geometric lookups and Jacobian calculation more efficient, whereas here it suffers from more expensive operations on a mesh map. Both, RRT*-Sam and RRT*-Pro use a fixed time budget and return the shortest path found within the allocated time or a failure state. For both variants we evaluate a time budget of 1 s and 0.25 s, which are indicated by “1” respectively “1/4”. In contrast, the *connecting* RRT (called RRT*-Con) starts sampling from start and goal independently and terminates as soon as a valid connection is found or the allocated time (1 s) has passed. The CHOMP planner uses an adjusted potential function, as its published version is used to avoid obstacles. We observed almost impossible convergence by using a potential function without slack around the ideal state (= path on the surface). Therefore, we define zero potential if within 0.1 m of the surface, quadratic within 1.1 m and linear otherwise. The parameters and configuration for all used planners are constant across all scenarios. Our evaluation additionally serves to provide insight into the underlying nature of the surface following problem and to show that the ability to exploit an explicit surface representation *and* its connectedness is highly beneficial. Fig. 1 shows an example trajectory for each planner.

B. Success Rate

The success rates for all 8 tested planners are plotted in Fig. 4. For the RMP planner, the success criterion is fulfilled if the trajectory is within 0.5 cm of the desired target and at rest. The discrete geodesic planner is guaranteed to converge exactly and the sampling-based planners are considered successful if an exact connected solution is returned. The proposed planner successfully converged for all trajectories. The sampling-based planners were relatively successful on the moderate geometry of scenario *hilo*. The variants that are able to exploit the explicit

²To facilitate fair comparison, we use state-of-the-art C++ implementations with as little changes as possible, usually limited to adjustments in the cost function implementation to enable all planners to use the same mesh mapping back-end e.g. for distance lookups.

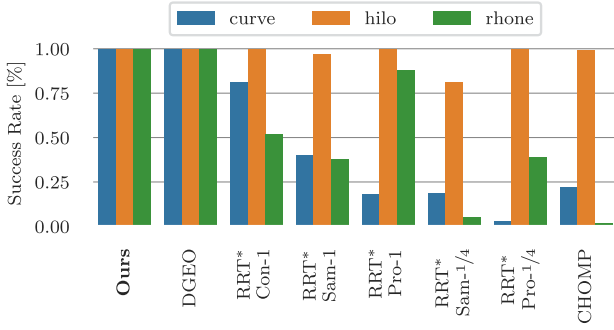


Fig. 4. Success rates as percentages for each planner. The color indicating the scenario is the same for all subsequent plots.




Fig. 5. Average duration for planning a full trajectory from start to goal. Note the logarithmic scale.

surface connectedness (*Con/Pro*) consistently outperformed the projecting planner on the difficult geometry of scenario *curve*. This can be attributed to the non-uniform sampling w.r.t. the surface in the *RRT*-Sam* variants. CHOMP worked well on the mid-sized, easy scenario *hilo* but got stuck in local minima often (*curve*) or would need to be re-tuned to account for the largest map (*rhone*). On all subsequent plots, only successful trajectories are shown. Note that this can introduce biases in the data, as e.g. all successfully planned paths for some of the RRT variants on scenario *curve* are strictly within the planar parts of the surface and do not go across one of the bends.

C. Planning Duration

Another important metric is the duration to reach a successful planning state. For the fixed-time variants of RRT* the duration is constant within timer resolution. For all RRT-based planners as well as the discrete geodesic algorithm, only the actual solving time (including path interpolation) is counted and setup times are excluded. The execution time of CHOMP highly depends on the scenario and task to solve. The results are displayed in Fig. 5. The RMP planner by itself only generates acceleration fields. In order to facilitate a fair comparison, we measure the time it takes to evaluate this acceleration field at a temporal resolution of 100 Hz and integrate the acceleration and velocity using a trapezoidal integrator to obtain the full trajectory. For execution on a real robot, this would not be necessary as we can simply evaluate the policies at the current position and velocity to follow the trajectory at any time. Additionally, the time to obtain the mapping \mathcal{H} once on startup is indicated in Table I




Fig. 6. Trajectory smoothness for all planner-scenario combinations. The plot is truncated at a smoothness of 0.85 - there are values below this. The discrete geodesic algorithm provides a lower bound of smoothness of the mesh, as it follows the mesh exactly at all times.

for each scenario. Over all results presented here, the RMP planner needed on the order of just $10 \mu s$ per such iteration, which corresponds to seamless replanning at rates of > 10 kHz. This enables smooth replanning for semi-manual steering on the surface, reactive planning with onboard sensors, or reacting to controller deviations.

D. Smoothness

For execution on the real robot, path smoothness is a desired property as changes in direction should be followed gradually and not in a jerky motion. Here we evaluate the three dimensional angular similarity between subsequent segments of a trajectory. Formally, the angular similarity ρ between two vectors A, B is defined as

$$\rho = 1 - \frac{1}{\pi} \cos^{-1} \left(\frac{A \cdot B}{|A||B|} \right), \quad (14)$$

and the smoothness as the average ρ over a full trajectory. Values very close to 1.0 indicate very little angular changes and short average segments, whereas values below 0.95 become visibly jagged. 0.5 corresponds to an average change of angle of 90 degrees. Fig. 6 visualizes the trajectory smoothness over all evaluated planners. The RRT-based planners produce on average a lot less smooth trajectories than the proposed planner. Especially RRT-Connect results in quite jagged trajectories, as it terminates on the first found connecting path. The very smooth trajectories produced by the proposed planners can largely be attributed to their physical nature and high sampling rate. As the planner produces an acceleration field that is integrated, the trajectory must always change gradually (at least on a very local scale). Paths obtained by RRT variants could be smoothed and post-processed further e.g. by fitting splines. However, this is not part of the planner per se and could add arbitrary increases of runtime and impairment of accuracy. CHOMP is forced to generate smooth paths due to the smoothness cost used in the optimization.

E. Surface Distance

We evaluate the surface following quality by measuring the distance between mesh surface and the obtained trajectories at 1 cm intervals along the trajectory. As shown in Fig. 7, the proposed planner shows below 1 mm deviation on average from the true surface for paths shorter than 10 m and the *hilo*

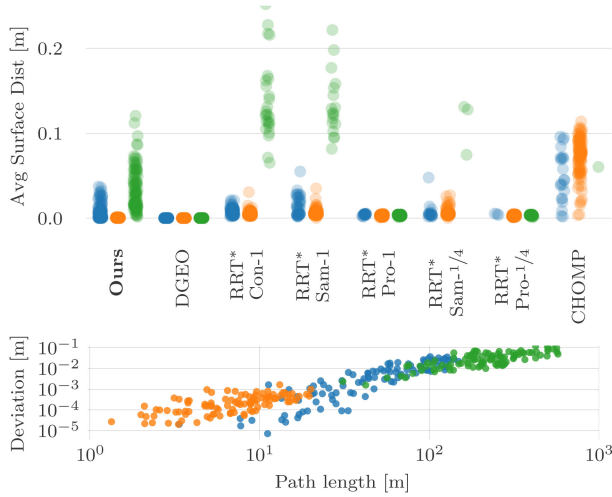


Fig. 7. Top: Average surface distance for each planned trajectory. Bottom: Scatter plot for all trajectories generated by the proposed planner that shows a correlation of path length and average surface distance deviation.

scenario. Deviations on the large *rhone* mesh can be explained by the physical nature of the planner. Trajectories of multiple 100 meters lead to larger velocities with the current tuning. Yet, there might be sharp changes in slope on the meshes, which the planner smooths to a certain extent as the trajectory is only affected by the resulting acceleration field. Depending on the use case, it can be advisable to adjust the planner tuning for very large maps. The projection-based planners show less deviation, as they explicitly project the obtained trajectory onto the nearest surface. However, depending on the sampling quality and geometry this can lead to invalid paths as the projection might not be uniquely defined. The observed deviations of the CHOMP planner are in large parts due to the needed slack in the potential function. Note that the plot only accounts for successfully planned trajectories and the success rate is relatively low on some scenarios for some planners (see Fig. 4).

F. Path Length

To show the optimality of the obtained paths as well as the effect of the induced distortion of \mathcal{H} for the RMP planner, we compare trajectory lengths to the theoretical optimum obtained by the discrete geodesic algorithm. The length ratio used here is defined as the trajectory length divided by the trajectory length for the same problem as obtained by the discrete geodesic algorithm. The closer to 1.0 the ratio, the closer to the theoretical optimum without any smoothing of abrupt edges or corners. Fig. 8 shows the results. Except for RRT*-Con, most of the trajectories obtained by the RRT* variants are within reasonable bounds with a few outliers. However, as only successful plans are considered, there is a bias towards straight, simple paths. CHOMP tends to generate paths that are too short, i.e. intersect geometry. The trajectories obtained with the proposed RMP-based planner are close to the optimum without outliers, effectively showing that the distortion induced by the mapping \mathcal{H} seems of relatively small impact for the practical problems presented in this letter.




Fig. 8. Length ratios for all planners. The discrete geodesic algorithm is used as the benchmark and therefore its length ratios are by definition 1.0.




Fig. 9. Left: Variations of α^\perp , where the darkest shade of green corresponds to a value of 0.1 and the brightest to 25.6. Right: Variations of α^\parallel between 0.2 (darkest) and 25.6. The green ball marks the start position, and the red ball the goal position. The blue trajectory corresponds to the tuning evaluated in the rest of the experiments.

G. Weighting of Policies

The proposed planner is able to plan trajectories from free-space towards a goal on a surface. The behavior of this transition can be tuned by changing the α^\perp and α^\parallel relative to each other. Intuitively, this can be interpreted as balancing the strength of the two policies relative to each other. Fig. 9 shows the impact of the two parameters for a given planning problem. As α^\perp , respectively the strength of the surface attraction policy, approaches 0 we obtain trajectories that stay at an approximately constant distance to the surface in curved parts, and exact constant distance in flat parts. Conversely, if the surface following policy is weaker relatively, the trajectory attaches to the surface as direct as possible. As the planner generates an acceleration field that can be evaluated at each time step, the tuning can be adjusted mid-trajectory at any time.

V. CONCLUSION

A novel path planning framework that combines Riemannian motion planning with mesh manifolds has been presented in this letter. The proposed framework solves the problem of approaching a surface and staying on a surface in 3D space in a mathematically elegant and real-world applicable way. The guarantees provided by the Riemannian motion policies combined with the proposed mesh-manifold rule out local-minima problems. Numerical inaccuracies could theoretically lead to non-optimal or non-terminating policies in rare cases, however we never observed such behavior in tests. We showed that our approach outperforms others in terms of performance, robustness, and execution time and works well on dissimilar scenarios with the same tuning. Our proposed algorithm does not suffer from local-minima or tuning problems like optimization-based approaches

and is not prone to sampling bottlenecks or geometrical ambiguity such as sampling-based planners. The proposed planner possesses a range of very powerful properties that motivate many future directions of work. The ability to follow a global near-optimum by the next best local direction at a very high rate allows the seamless integration of live sensor data to e.g. facilitate on-surface reactive and dynamic obstacle avoidance. Other directions include the addition of orientation policies and combination with learned policies based on surface properties.

REFERENCES

- [1] K. Bodie, M. Brunner, M. Pantic, S. Walser, P. Pfändler, U. Angst, R. Siegwart and J. Nieto, "Active interaction force control for contact-based inspection with a fully actuated aerial vehicle," *IEEE Trans. Robot.*, to be published, doi: [10.1109/TRO.2020.3036623](https://doi.org/10.1109/TRO.2020.3036623).
- [2] H. Oleynikova, C. Lanegger, Z. Taylor, M. Pantic, A. Millane, R. Siegwart and J. Nieto, "An open-source system for vision-based micro-aerial vehicle mapping, planning, and flight in cluttered environments," *J. Field Robot.*, vol. 37, pp. 642–666, 2020, doi: [10.1002/rob.21950](https://doi.org/10.1002/rob.21950).
- [3] M. W. Mueller, M. Hehn, and R. D'Andrea, "A computationally efficient motion primitive for quadcopter trajectory generation," *IEEE Trans. Robot.*, vol. 31, no. 6, pp. 1294–1310, Dec. 2015.
- [4] C. Richter, A. Bry, and N. Roy, "Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments," *Robotics Research. Springer Tracts in Advanced Robotics*. M. Inaba and P. Corke Eds. vol. 114, 2016. [Online]. Available: https://doi.org/10.1007/978-3-319-28872-7_37
- [5] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "CHOMP: Gradient optimization techniques for efficient motion planning," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2009, pp. 489–494.
- [6] S. M. LaValle and J. J. Kuffner Jr, "Randomized kinodynamic planning," *Int. J. Robot. Res.*, vol. 20, no. 5, May 2001, pp. 378–400.
- [7] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *Int. J. Robot. Res.*, vol. 30, no. 7, Jun. 2011, pp. 846–894.
- [8] J. J. Kuffner and S. M. LaValle, "Rrt-connect: An efficient approach to single-query path planning," in *Proc. Int. Conf. Robot. Automat. Millennium Conf. IEEE Int. Conf. Robot. Automat. Symp.*, 2000, pp. 995–1001.
- [9] M. Tognon, E. Cataldi, H. A. T. Chavez, G. Antonelli, J. Cortés, and A. Franchi, "Control-aware motion planning for task-constrained aerial manipulation," *IEEE Robot. Automat. Lett.*, vol. 3, no. 3, pp. 2478–2484, Jul. 2018.
- [10] Z. Kingston, M. Moll, and L. E. Kavraki, "Sampling-based methods for motion planning with constraints," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 1, no. 1, pp. 159–185, 2018.
- [11] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: An efficient probabilistic 3D mapping framework based on octrees," *Auton. Robot.*, vol. 34, pp. 189–206, 2013, doi: [10.1007/s10514-012-9321-0](https://doi.org/10.1007/s10514-012-9321-0).
- [12] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3D euclidean signed distance fields for on-board Mav planning," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, 2017, pp. 1366–1373.
- [13] X. Gu, Y. He, and H. Qin, "Manifold splines," *Graphical Models*, vol. 68, no. 3, pp. 237–254, May 2006.
- [14] L. Piegl and W. Tiller, *The NURBS Book*, Ser. Monographs in Visual Communication. Berlin, Heidelberg: Springer, 1997.
- [15] H. Pfister, M. Zwicker, J. Van Baar, and M. Gross, "Surfels: Surface elements as rendering primitives," in *Proc. 27th Annu. Conf. Comput. Graph. Interactive Technol.*, 2000, pp. 335–342.
- [16] A. Bircher *et al.*, "Structural inspection path planning via iterative view-point resampling with application to aerial robotics," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2015, pp. 6423–6430.
- [17] F. Ruetz *et al.*, "OVPC mesh: 3D free-space representation for local ground vehicle navigation," in *Proc. Int. Conf. Robot. Automat.*, 2019, *arXiv:1811.10266*.
- [18] S.-Q. Xin and G.-J. Wang, "Improving Chen and Han's algorithm on the discrete geodesic problem," *ACM Trans. Graph.*, vol. 28, no. 4, pp. 1–8, Aug. 2009.
- [19] M. Watterson, S. Liu, K. Sun, T. Smith, and V. Kumar, "Trajectory optimization on manifolds with applications to quadrotor systems," *Int. J. Robot. Res.*, vol. 39, no. 2-3, pp. 303–320, 2020.
- [20] N. D. Ratliff, J. Issac, D. Kappler, S. Birchfield, and D. Fox, "Riemannian motion policies," Jul. 2018, *arXiv:1801.02854 [cs]*.
- [21] A. F. Möbius, *Der Barycentrische Calcul: Ein Neues Hülfsmittel Zur Analytischen Behandlung Der Geometrie*. Munich, Germany: Barth, 1827.
- [22] W. T. Tutte, "How to draw a graph," in *Proc. London Math. Soc.*, pp. 743–767, 1963.
- [23] M. Desbrun, M. Meyer, and P. Alliez, "Intrinsic parameterizations of surface meshes," *Comput. Graph. Forum*, vol. 21, pp. 209–218, 2002, doi: [10.1111/1467-8659.00580](https://doi.org/10.1111/1467-8659.00580).
- [24] M. S. Floater, "Mean value coordinates," *Comput. Aided Geometric Desi.*, vol. 20, no. 1, pp. 19–27, 2003.
- [25] L. Saboret, P. Alliez, B. Lévy, M. Rouxel-Labbé, and A. Fabri, "Triangulated surface mesh parameterization," in *CGAL User and Reference Manual*, 5.1 ed. CGAL Editorial Board, 2020.
- [26] P. Alliez, S. Tayeb, and C. Wormser, "3D fast intersection and distance computation," in *CGAL User and Reference Manual*, 5.1 ed. CGAL Editorial Board, 2020.
- [27] S. Liu and C. C. L. Wang, "Fast intersection-free offset surface generation from freeform models with triangular meshes," *IEEE Trans. Automat. Sci. Eng.*, vol. 8, no. 2, pp. 347–360, Apr. 2011.
- [28] S. Kiazky, S. Lorient, and É. C. de Verdière, "Triangulated surface mesh shortest paths," in *CGAL User and Reference Manual*, 5.1 ed. CGAL Editorial Board, 2020.
- [29] I. A. Şucan, M. Moll, and L. E. Kavraki, "The open motion planning library," *IEEE Robot. Automat. Mag.*, vol. 19, no. 4, pp. 72–82, Dec. 2012.