2016 IEEE 55th Conference on Decision and Control (CDC)
ARIA Resort & Casino
December 12-14, 2016, Las Vegas, USA

Discretized Boundary Methods for Computing Smallest Forward Invariant Sets*

Paul Varnell[1], Shayok Mukhopadhyay[2], and Fumin Zhang[3]

***Abstract*— This paper presents a multi-dimensional method to compute forward invariant sets (FIS) that are tight approximations of the smallest FIS of nonlinear perturbed systems modeled by differential inclusions. We formulate the problem as a discretized optimal boundary search, using methods from computational topology and non-smooth analysis to ensure invariance constraints for piecewise linear boundaries of FIS. We solve this optimal boundary search problem using a greedy search method with backtracking to find the optimal boundary, which defines the smallest FIS that can be represented in the discretized search space.**

I. INTRODUCTION

There are many important examples of control systems with strict safety and performance requirements that must be met in spite of the presence of uncertainties. For example, teleoperated surgery, industrial robots cooperatively working with humans, and aircraft flight controls are all systems where the violation of certain constraints could result in destruction of expensive equipment or even loss of human life. These examples also require exceptional control performance in order to accomplish their intended purpose. Achieving both safety and performance is made more difficult by uncertainties in the environment resulting from human operators and other complex unmodeled dynamics.

Controllers for these kinds of systems need to guarantee that state constraints are not violated, without sacrificing potential performance due to overly conservative estimates. Forward invariant sets (FIS) provide a general way to analyze state constraints for perturbed systems [1]. An FIS of a system is a subset of the state space that trajectories will never leave. FIS are most commonly used in Lyapunov analysis, where the sublevel sets of a Lyapunov function are FIS containing an equilibrium point of an unperturbed system. Input-to-State Stability (ISS) techniques [2], [3] are used to find FIS for perturbed systems. However, given bounds on the perturbation, ISS methods typically will not find tight approximations of the smallest FIS. Tighter estimates on the smallest FIS allow us to use higher performance controllers while still ensuring that state constraints are satisfied under given assumptions about the perturbations.

This paper proposes a computational technique to find FIS that are tight approximations of the smallest FIS of perturbed systems. This is the first method capable of this for perturbed nonlinear continuous-time systems with any finite number of state dimensions. We formulate this problem as a discretized optimal boundary search, and use methods from computational topology and non-smooth analysis to enforce invariance constraints for piecewise linear boundaries of FIS. We solve this optimal boundary search problem using a greedy search method with backtracking to find the optimal boundary, which defines the smallest FIS that can be represented in the discretized search space.

Invariant sets have a long been used to design and analyze perturbed systems [1]. There are existing algorithms for computing FIS for discrete-time [4] and continuous-time [5] linear systems. There are also a large number of algorithms for computing finite-time and infinite-time reachable sets [6]–[8], including optimization based methods [9]–[11], level set and boundary propagation methods [12], [13], and trajectory tube approaches [14]–[16]. Additionally, there are a few related algorithms for computing regions of attraction [17].

For autonomous systems, smallest FIS are essentially the same as infinite-time reachable sets. However, reachable set algorithms are not ideal for computing FIS. Infinite-time reachable set algorithms produce approximations of the smallest FIS, but these approximations are not usually guaranteed to be FIS. Also, existing methods to compute infinite-time reachable sets are limited to linear or polynomial systems. Finally, most reachable set methods involve iteratively integrating trajectories to find a set at each point in time. This is unnecessary and inefficient for computing FIS, which can be evaluated using only boundary conditions.

The method in this paper is inspired by the 2D version [18], [19] and they share some important features. Both approaches discretize the boundary search space into piecewise linear segments, which are searched to find an optimal boundary within the discretized space. The most important difference is this new method works for systems with more than two state variables, as it uses multi-dimensional simplexes instead of line segments to discretize the boundary space. Instead of using path-planning techniques to perform the optimal boundary search as in [18], [19], we use a greedy search with backtracking that can handle the increased complexity of the multi-dimensional discretization.

II. BACKGROUND

A. Differential Inclusions

Differential inclusions can model systems that evolve over time in multiple possible ways, such as systems with perturbations or control inputs. Defining the state $x(t) \in \mathcal{X} \subseteq \mathbb{R}^n$,

*The research work is supported by ONR grants N00014-10-10712 (YIP) and N00014-14-1-0635; and NSF grants CE-1032285, IIS-1319874, and CMMI-1436284.

[1]Georgia Institute of Technology paul.varnell@gatech.edu
[2]American University of Sharjah smukhopadhyay@aus.edu
[3]Georgia Institute of Technology fumin@gatech.edu

a differential inclusion

$$\dot{x}(t) \in F(x(t)) \ , \tag{1}$$

constrains the time derivative of the state to a set of possible values defined by the set-valued map $F(x) \subset \mathbb{R}^n$.

For our purposes, differential inclusions provide a notationally convenient way to study perturbed systems. For example, a nonlinear system $\dot{x}(t) = f(x(t), \delta(t))$ with unknown perturbation $\delta(t) \in \mathcal{P} \subseteq \mathbb{R}^m$ is equivalent to a differential inclusion (1) with $F(x) = \{f(x,\delta) : \delta \in \mathcal{P}\}$. In the next subsection, we will show how the FIS of perturbed systems can be characterized succinctly in terms of F.

In this paper, we assume that $F(x)$ is bounded and non-empty for each value of $x \in \mathcal{X}$ and that F is Lipschitz continuous with respect to the Hausdorff metric, meaning that there exists a Lipschitz constant $\ell \geq 0$ such that

$$H\left(F(x_1), F(x_2)\right) \leq \ell \left\|x_1 - x_2\right\| \tag{2}$$

for all $x_1, x_2 \in \mathcal{X}$, where the Hausdorff metric H [20] is

$$H(S_1, S_2) = \max\left\{ h(S_1, S_2),\, h(S_2, S_1) \right\} \quad , \tag{3}$$

$$h(A, B) = \sup_{a \in A} \inf_{b \in B} \|a - b\| \quad . \tag{4}$$

For a differential inclusion representing a perturbed system $F(x) = \{f(x,\delta) : \delta \in \mathcal{P}\}$, certain properties of f and $\mathcal{P}$ can ensure that our assumptions on F hold. Specifically, if for each $x \in \mathcal{X}$, $f(x, \delta)$ is uniformly bounded over all $\delta \in \mathcal{P}$, then $F(x)$ is bounded for each $x \in \mathcal{X}$. If f is uniformly Lipschitz continuous in x (i.e. $\|f(x_1,\delta) - f(x_2,\delta)\| \leq \ell\|x_1 - x_2\|$ for all $x_1, x_2 \in \mathcal{X}$, $\delta \in \mathcal{P}$), then F is also Lipschitz continuous with the same Lipschitz constant.

When F is Lipschitz continuous, absolutely continuous solutions to (1) are guaranteed to exist for every initial state $x_0 \in \mathcal{X}$ [20], and we use the term Φ to denote the set of these solutions.

B. Forward Invariant Sets (FIS)

Forward invariance is a property that can be used to study the robustness and performance of a system.

Definition 1: Given a system with solutions Φ, a subset of the state space $S \subseteq \mathcal{X}$ is *forward invariant* or *strongly forward invariant* if for all solutions $\phi \in \Phi$, if $\phi(t_0) \in S$ for some time t_0, then $\phi(t) \in S$ for all times $t \geq t_0$.

The robustness of an equilibrium point or a nominal trajectory can be measured by a forward invariant set (FIS) that contains them. The smallest FIS is of particular interest, as it is a tight measure of robustness. Given a seed set $\mathcal{E} \subset \mathcal{X}$, such as $\mathcal{E} = \{x_e\}$ for an equilibrium point x_e with $f(x_e, 0) = 0$ or $\mathcal{E} = \{x^*(t) : t \geq 0\}$ for a nominal trajectory $x^*(t)$, we denote the smallest FIS containing $\mathcal{E}$ as $\mathcal{S}_{F,\mathcal{E}}$. Defining the set of all FIS containing $\mathcal{E}$ as $\mathscr{R}_{F,\mathcal{E}}$, the smallest FIS $\mathcal{S}_{F,\mathcal{E}}$ is the least element in $\mathscr{R}_{F,\mathcal{E}}$, meaning that $\mathcal{S}_{F,\mathcal{E}} \in \mathscr{R}_{F,\mathcal{E}}$ satisfies $\mathcal{S}_{F,\mathcal{E}} \subseteq S$ for all $S \in \mathscr{R}_{F,\mathcal{E}}$.

In most cases, it is impractical to directly verify that a set is forward invariant or to construct forward invariant sets by examining all of the trajectories of a system. Instead, if a system has continuous solutions, we can evaluate if a set is forward invariant by checking that trajectories never leave the set through its boundary. We do this by comparing the angle between the dynamics F and the normal to the set's boundary. In general, the boundary of an arbitrary FIS may not be a smooth manifold, which is traditionally required to define normal vectors. Our method will use sets with piecewise linear boundaries that are non-smooth at the intersections between the linear segments. However, for any arbitrary set S the *outward-pointing normal cone* [20] is

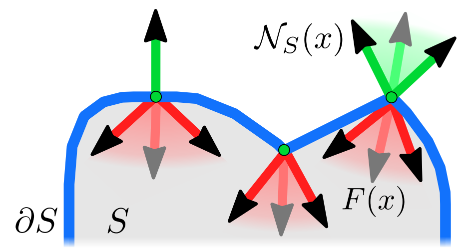

Fig. 1: The invariance conditions (6) are satisfied at the three boundary points shown labeled with the directions of the dynamics F and the normal directions $\mathcal{N}_S$.

$$\mathcal{N}_S(x) = \{\, \zeta \in \mathbb{R}^n : \exists \epsilon > 0 \text{ with } \inf_{s \in S} \|x + \epsilon\zeta - s\| = \epsilon\|\zeta\| \,\} \ . \tag{5}$$

Along the boundary of S, which we denote as ∂S, the normal cone $\mathcal{N}_S$ is the set of outward-pointing normal directions. Fig. 1 illustrates three boundary points with the corresponding normal cone at those points. On the left of Fig. 1 is a point where ∂S is locally smooth, and the normal cone is equivalent to the traditional normal vector, on the right is a non-smooth point where S is locally convex and the normal cone is non-trivial, and in the center is a non-smooth point where S is locally concave and the normal cone is equal to the zero vector $\mathcal{N}_S(x) = \{0\}$.

Theorem 1: For a system as in (1) with F Lipschitz continuous, a closed set $S \subseteq \mathcal{X}$ is an FIS if and only if

$$\langle F(x), \mathcal{N}_S(x) \rangle \leq 0 \text{ for all } x \in \partial S \quad . \tag{6}$$

Proof: This theorem is a minor variation [21, Theorem 3.8] (although with different notation), where we are using the additional fact that $\mathcal{N}_S(x) = \{0\}$ for $x \in \text{int}(S)$. ■

The invariance conditions in (6) relate the angle between the outward-pointing normal cone $\mathcal{N}_S$ and the direction of the dynamics F, so that the set S is invariant if the elements of F point "into" the set, as illustrated in Fig. 1. In the rest of this paper, we will use the invariance conditions (6) to construct sets that are the boundaries of FIS.

C. Using Oriented Boundaries to Define Sets

Our method will construct sets that are the boundaries of FIS, which is possible because a) the invariance conditions (6) for a set are only relevant at its boundary and b) we can entirely specify a set by its boundary and normal cones. Specifically, given a closed set $\Gamma \subset \mathcal{X}$ with empty interior $\Gamma = \partial\Gamma$ and given a set-valued function $N(x) \subseteq R^n$ defined for $x \in \Gamma$, we will call the pair (Γ, N) an *oriented boundary* if it defines a unique closed set $\mathcal{C}_{\Gamma,N}$ of which Γ is the boundary, so that $\partial\mathcal{C}_{\Gamma,N} = \Gamma$ and $\mathcal{N}_{\mathcal{C}_{\Gamma,N}} \equiv N$.

In general, it may be difficult to determine if an arbitrary pair (Γ, N) is an oriented boundary, but in Section IV-B we will show how this is relatively easy when Γ is piecewise linear. Our approach to finding FIS will involve incrementally constructing piecewise linear oriented boundaries.

III. Computing Smallest Forward Invariant Sets

Given dynamics F and a seed set $\mathcal{E}$, our goal is to find the smallest FIS $\mathcal{S}_{F,\mathcal{E}}$. We formulate this as a constrained optimization problem to find the oriented boundary of $\mathcal{S}_{F,\mathcal{E}}$.

Problem 1: Find (Γ, N) that satisfies the constraints

1) (boundary) (Γ, N) defines a unique set $\mathcal{C}_{\Gamma,N}$ with $\partial\mathcal{C}_{\Gamma,N} = \Gamma$ and $\mathcal{N}_{\mathcal{C}_{\Gamma,N}} \equiv N$
2) (seed) $\mathcal{E} \subseteq \mathcal{C}_{\Gamma,N}$
3) (invariance) $\langle F(x), N(x)\rangle \leq 0$ for all $x \in \Gamma$

and (optimality conditions) if some other (Γ', N') satisfies these constraints, then $\mathcal{C}_{\Gamma,N} \subseteq \mathcal{C}_{\Gamma',N'}$.

In general, it is not tractable to find closed-form solutions to this problem. Instead, we want to find iterative numerical approximations that will converge to the solution. Furthermore, in many applications, we want the approximate solution to strictly satisfy the constraints of Problem 1.

This paper focuses on finding an approximate solution to Problem 1 that strictly satisfies the constraints. This solution will be an FIS that is close the smallest containing the seed set $\mathcal{E}$. The fact that the solutions we find will strictly meet the constraints is an important feature of our approach that distinguishes it from most other related numerical methods.

Before we can proceed to finding approximate FIS solutions to Problem 1, we must first consider how to discretize the problem in a way that makes it computationally tractable.

IV. Discretizing Boundary Search Space

We will use a finite collection of oriented piecewise linear segments to represent possible oriented boundaries and we will show in the following sections how this representation is amenable for solving Problem 1.

A. Background: Oriented Simplexes

A (geometric) *simplex* is a generalization of triangles and tetrahedrons to higher dimensions. In $\mathbb{R}^n$, a *k-simplex* Δ with $0 \leq k \leq n$, also called a *k-face*, is defined by the convex hull of a set of (k+1) vertex points $\{v_0, ..., v_k\}$ with $v_i \in \mathbb{R}^n$

$$\Delta = \left\{ \sum_{i=0}^{k} c_i v_i \;:\; c_i \geq 0,\; \sum_{i=0}^{k} c_i = 1 \right\} \quad , \tag{7}$$

where we assume that the matrix

$$B = \left[(v_1 - v_0), ..., (v_k - v_0)\right] \in \mathbb{R}^{n\times k} \tag{8}$$

has linearly independent columns. A k-simplex contains (k+1) number of (k-1)-simplexes, each defined by removing one of the original vertices $\{v_0, ..., v_{i-1}, v_{i+1}, ..., v_k\}$. This can be repeated to obtain lower dimensional simplexes. We will refer to these (k-l)-simplexes as *(k-l)-sub-faces*. Fig. 2 shows examples of k-simplexes for $k \in \{0, 1, 2, 3\}$.

When the vertices are ordered $(v_0, ..., v_k)$, the columns of B form an ordered basis that *orients* the simplex, as shown

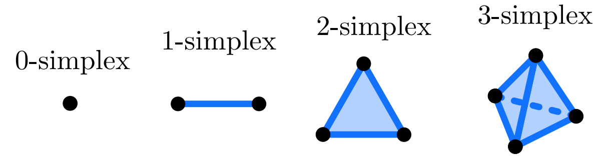

Fig. 2: Simplexes of different dimensions. A 1-simplex has two 0-simplexes, a 2-simplex has three 1-simplexes, etc.

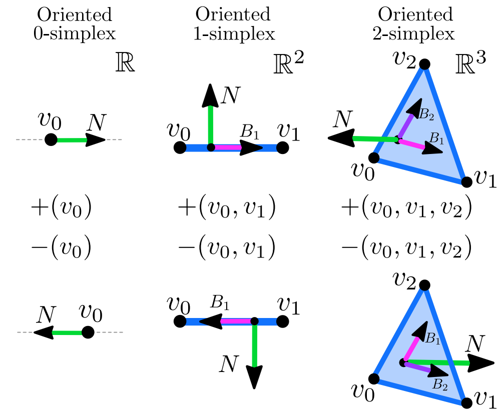

Fig. 3: Oriented (n-1)-simplexes shown in $\mathbb{R}^n$ with normals (green arrows) and ordered basis vectors (purple arrows, not to scale). Note that a 0-simplex has only 1 vertex, so it can be oriented by a sign (+ or −) but not by a vertex ordering.

in Fig. 3. There are only two possible unique orientations of a given simplex [22], and the orientation is reversed whenever pairs of vertices are exchanged in the ordering. For example, (v_0, v_1) has opposite orientation as (v_1, v_0). In typical notation, the vertices are kept in some default ordering and the simplex is labeled with a minus sign $(-)$ if it has opposite orientation to the default.

For an (n-1)-simplex, an orientation also defines a oriented unit normal vector $N = m(B)$ using the right-hand convention, as shown in Fig. 3, where

$$m(B) = \overrightarrow{\widetilde{m}(B)} \in \mathbb{R}^n \qquad \widetilde{m}_i(B) = (-1)^{n+i} \det B_{\setminus i} \tag{9}$$

and where $B_{\setminus i}$ is formed by removing the ith row from B.

B. Oriented Simplicial Boundaries

In this section, consider a collection of k-simplexes $I = \{\Delta_1, ..., \Delta_r\}$ where $\Delta_i = s_i\,(a_i, v_0, ..., v_{k-1})$, with signs $s_i \in \{+, -\}$ defined accordingly, that are intersecting at the shared (k-1)-sub-face $\overline{\Delta} = (v_0, ..., v_{k-1})$.

When two oriented k-simplexes intersect at a shared (k-1)-sub-face, the intersection is *consistent* if their orientations are equivalent at the intersection. When oriented simplexes intersect consistently, they can be treated as a single oriented object. Using the notation in the previous paragraph, two oriented k-simplexes Δ_1 and Δ_2 intersect consistently if $s_1 = -s_2$; see [22] for a definition using the induced orientations of the sub-face. If two oriented simplexes do not intersect consistently, then the intersection is *inconsistent*. Fig. 4 shows some examples of consistent and inconsistent intersections for $k \in \{1, 2\}$.

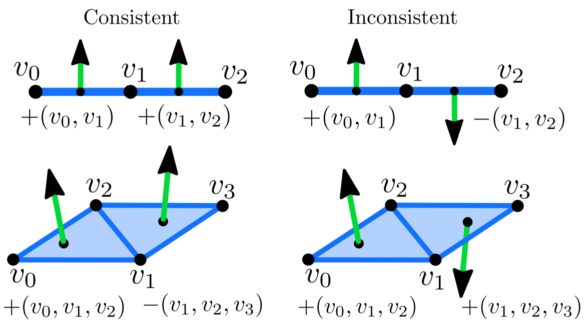

Fig. 4: Consistent and inconsistent intersections of 1- and 2-simplexes, shown with normal vectors (green arrows).

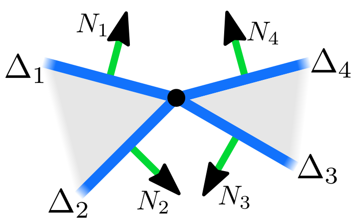

Fig. 5: The intersection of $\{\Delta_1, \Delta_2, \Delta_3, \Delta_4\}$ is consistent, but the intersection of any three of the faces is inconsistent.

We define an intersection of multiple oriented (n-1)-simplexes at a shared (n-2)-sub-face to be *consistent* if each pair of adjacent (n-1)-simplexes intersect consistently. Consistently intersecting (n-1)-simplexes locally define disjoint interior and exterior sets, as seen in Fig. 5. We use this property to define oriented boundaries made up of simplexes. To make this definition more rigorous, we first define the angle between two (n-1)-simplexes as

$$\theta(\Delta_1, \Delta_2) = \operatorname{atan2}(s_1 s_2 \langle R_1, N_2 \rangle, s_1 s_2 \langle N_1, N_2 \rangle) \quad (10)$$

where $R_1 = m([\overline{B}, N_1])$ and where $\overline{B}$ is the matrix (8) associated with the shared sub-face $\overline{\Delta}$. The positive angle $\theta(\Delta_1, \Delta_2) \in [0, 2\pi)$ is measured by rotating Δ_1 in the direction opposite of N_1 until reaching Δ_2. For example, in Fig. 5, the angle $\theta(\Delta_1, \Delta_2) \approx \pi/3$.

We can use the angle between faces to rotationally order them around a shared sub-face. We can represent an intersection by the signed, cyclic ordering $I \equiv ((\delta_1, \Delta_{\sigma(1)}), ..., (\delta_r, \Delta_{\sigma(r)}))$ where σ is a permutation with $\delta_1\, \theta(\Delta_{\sigma(1)}, \Delta_{\sigma(i)})) \leq \delta_1\, \theta(\Delta_{\sigma(1)}, \Delta_{\sigma(i+1)}))$ and the signs $\delta_i \in \{+, -\}$ are chosen so $\delta_i\, s_{\sigma(i)}$ is the same for all i. For example, the intersection in Fig. 5 can be represented as $((+, \Delta_1), (-, \Delta_2), (+, \Delta_3), (-, \Delta_4))$ or equivalently as $((-, \Delta_1), (+, \Delta_4), (-, \Delta_3), (+, \Delta_2))$ or $((+, \Delta_2), (-, \Delta_1), (+, \Delta_4), (-, \Delta_3))$.

With this representation, an intersection I is consistent if the signs alternate $\delta_i = -\delta_{i+1}$ for all i. The caption of Fig. 5 indicates which subsets of the faces intersect consistently or inconsistently, as can be verified using this definition.

Given a set $\mathcal{D}$ of (n-1)-simplexes, we define $\mathcal{I}_\mathcal{D} \subseteq 2^\mathcal{D}$ so that each element $I \in \mathcal{I}_\mathcal{D}$ is the set of all faces involved in the intersection at a particular shared (n-2)-sub-face.

Proposition 1: Given a finite set $\mathcal{D}$ of oriented (n-1)-simplexes that only intersect at shared sub-faces, $\mathcal{D}$ is an oriented boundary if and only if each $I \in \mathcal{I}_\mathcal{D}$ is consistent.

Proof: ($\Rightarrow$) If there exists an $I \in \mathcal{I}_\mathcal{D}$ that is not consistent, then there will be a local neighborhood around the (n-2)-simplex associated with the intersection I where the oriented normal vectors will not agree on which region is the interior or exterior of the set, so in this case, $\mathcal{D}$ cannot be an oriented boundary.

($\Leftarrow$) If each intersection $I \in \mathcal{I}_\mathcal{D}$ is consistent, we can topologically "cut" the boundary at each intersection I into locally consistent pairs of faces. The topological space formed in this way is equivalent to one or more combinatorial (n-1)-pseudomanifolds [23]. As shown in [23] and the references therein, when a combinatorial (n-1)-pseudomanifold is bijectively mapped onto a set $\Gamma \subset \mathbb{R}^n$, the set Γ divides $\mathbb{R}^n$ into two connected components S_int and S_ext. These correspond to the interior and exterior of the set defined by the boundary Γ, with $S_\text{int} \cup \Gamma \cup S_\text{ext} = \mathbb{R}^n$, $S_\text{int} \cap S_\text{ext} = \emptyset$, and $\partial S_\text{int} = \partial S_\text{ext} = \Gamma$. In our case, the mappings from the pseudomanifolds to $\Gamma_i \subset \mathbb{R}^n$ may be singular at each $I \in \mathcal{I}_\mathcal{D}$, but since the boundaries Γ_i do not cross each other, we preserve the separation between the interior sets $S_{\text{int},i}$ defined by each pseudomanifold mapping to Γ_i. The combined boundary $\Gamma = \bigcup_i \Gamma_i$ separates the space into the interior $S_\text{int} = (\bigcup_i S_{\text{int},i})$ and exterior $S_\text{ext} = (\bigcup_i S_{\text{ext},i}) \setminus S_\text{int}$ satisfying the properties to make Γ an oriented boundary. ■

This theorem allows us to determine if a set $\mathcal{D}$ of oriented (n-1)-simplexes is an oriented boundary by checking the consistency of each intersection $I \in \mathcal{I}_\mathcal{D}$. We will call $\mathcal{D}$ that meet these conditions *oriented simplicial boundaries* and define $\mathcal{C}_\mathcal{D}$ as the unique closed set of which $\mathcal{D}$ is the boundary.

Given a search space of (n-1)-simplexes, our algorithm will iteratively construct an oriented simplicial boundary that is the smallest possible FIS containing some given seed set that can be represented in the discretized space.

V. TESTING FORWARD INVARIANCE

In order to ensure the oriented simplicial boundary that our algorithm constructs is an FIS, we can choose the orientation of each simplex to meet the invariance constraints (6) if possible, or remove the simplex if it cannot meet the invariance constraints for either normal orientation.

A. Normal Cones of Oriented Simplicial Boundaries

Since oriented simplicial boundaries are non-smooth at the intersections between faces, we must compute its normal cone at these points to test the invariance conditions (6). As we will show, this is not a problem when F is Lipschitz, as we can essentially ignore the non-smooth boundary points.

It is possible to derive an analytical expression for the normal cone $\mathcal{N}_{\mathcal{C}_\mathcal{D}}(x)$ of an oriented simplicial boundary $\mathcal{D}$, but the expression is extremely lengthy. Instead, we use the much more compact estimate for the normal cone

$$\mathcal{N}_{\mathcal{C}_\mathcal{D}}(x) \subseteq \operatorname{cone}(\{N_i : \Delta_i \in \mathcal{D} \text{ with } x \in \Delta_i\}) \quad (11)$$

where for some finite set A, the operator $\operatorname{cone}(A) = \{\sum_{a \in A} \lambda_a a : \lambda_a \geq 0\}$. When all intersections between faces are convex, the relationship (11) holds with equality.

This estimate for the normal cone (11) is sufficient to test the invariance condition on an oriented simplicial boundary.

Proposition 2: Given an oriented simplicial boundary $\mathcal{D}$, if $\langle F(x), N_i \rangle \leq 0$ for all $\Delta_i \in \mathcal{D}$ and all $x \in \Delta_i$, then $\langle F(x), \mathcal{N}_{\mathcal{C}_\mathcal{D}}(x) \rangle \leq 0$ for all $x \in \bigcup \mathcal{D}$.

Proof: We know that for all $x \in \bigcup \mathcal{D}$

$$\langle F(x), \mathcal{N}_{\mathcal{C}_\mathcal{D}}(x) \rangle \subseteq \langle F(x), \text{cone}\{N_i : \Delta_i \in \mathcal{D} \text{ s.t. } x \in \Delta_i\} \rangle$$
$$= \left\{ \left\langle v, \sum_{\substack{\Delta_i \in \mathcal{D}: \\ \text{s.t } x \in \Delta_i}} \lambda_i N_i \right\rangle : v \in F(x), \lambda_i \geq 0 \right\}$$
$$= \left\{ \sum_{\substack{\Delta_i \in \mathcal{D}: \\ \text{s.t } x \in \Delta_i}} \lambda_i \langle v, N_i \rangle : v \in F(x), \lambda_i \geq 0 \right\}$$

Therefore, if $\langle v, N_i \rangle \leq 0$ for each $v \in F(x)$ and $x \in \Delta_i$, we know that $\langle F(x), \mathcal{N}_{\mathcal{C}_\mathcal{D}}(x) \rangle \leq 0$. ■

Proposition 2 provides sufficient conditions to test if an oriented simplicial boundary satisfies the invariance conditions (6). It allows us to show that (6) holds at the non-smooth intersections, without having to directly check the normal cones at those points.

B. Computational Verification

We want to computationally verify that an oriented simplicial boundary satisfies the invariance conditions (6) without testing all of the infinite number of boundary points. When the dynamics F are Lipschitz continuous, if a given normal direction non-strictly satisfies the invariance conditions (6) at a boundary point, then the normal direction will also satisfy the invariance conditions at nearby boundary points.

Proposition 3: Assuming F is ℓ-Lipschitz continuous and given a point $x_0 \in \mathcal{X}$, a unit normal vector $N \in \mathbb{R}^n$ with $\|N\| = 1$, and a constant $\epsilon > 0$, if $\langle F(x_0), N \rangle \leq -\epsilon$, then $\langle F(x), N \rangle \leq 0$ for all $x \in \mathcal{X}$ with $\|x - x_0\| \leq \epsilon/\ell$.

Proof: For all $x \in \mathcal{X}$

$$\langle F(x), N \rangle \leq \sup_{v \in F(x)} \langle v, N \rangle$$
$$= \sup_{v \in F(x)} \inf_{v_0 \in F(x_0)} \langle v - v_0, N \rangle + \sup_{v_0 \in F(x_0)} \langle v_0, N \rangle$$

If $\langle F(x_0), N \rangle \leq -\epsilon$, then $\sup_{v_0 \in F(x_0)} \langle v_0, N \rangle \leq -\epsilon$, by definition of the supremum. Using the Cauchy-Schwarz inequality $\langle v - v_0, N \rangle \leq \|v - v_0\| \, \|N\| = \|v - v_0\|$, so

$$\langle F(x), N \rangle \leq \sup_{v \in F(x)} \inf_{v_0 \in F(x)} \|v - v_0\| - \epsilon$$
$$= h(F(x), F(x_0)) - \epsilon \leq H(F(x), F(x_0)) - \epsilon$$
$$\leq \ell \|x - x_0\| - \epsilon$$

where the last step uses the Lipschitz property of F. Therefore, if $\|x - x_0\| \leq \epsilon/\ell$, then $\langle F(x), N \rangle \leq 0$ ■

Using Propositions 2 and 3, we can guarantee that an oriented simplicial boundary satisfies the invariance conditions (6) in some cases. We do this by showing that for each face $\Delta_i \in \mathcal{D}$ there exists a collection of test points $x_j \in \Delta_i$ where $\langle F(x_j), N_i \rangle \leq -\epsilon_i$ with $\epsilon_i > 0$ and the test points are sufficiently close together so that $\Delta_i \subseteq \bigcup_j \{x \in \mathcal{X} : \|x - x_j\| \leq \epsilon_j/\ell\}$. If we can find such points for each face, then the entire boundary satisfies the invariance conditions.

As a part of our algorithm, we will only include an oriented face in the search space if we can guarantee that it satisfies the invariance conditions (6). This way, any oriented simplicial boundary we find is guaranteed to be an FIS. For any face, it will only ever be possible to prove that one of the two orientations satisfies the invariance conditions. This is because the only case where both orientations satisfy the invariance conditions is when $\langle F(x), N_i \rangle = 0$ for all $x \in \Delta_i$, which can only be verified using an infinite number of test points. When we cannot guarantee that either orientation satisfies the invariance conditions, either because a) $\langle F(x), N_i \rangle > 0$ at some points on the face and $\langle F(x), N_i \rangle < 0$ at other points or b) we have only a limited number of test points and could not cover the face with these points, then we will eliminate that face from the search space.

VI. DISCRETIZED FIS SEARCH

A. Selection of Discretized Search Space

The first step of our search for the boundary of the smallest FIS is to discretize the boundary space into a finite set of (n-1)-simplexes that only intersect at shared sub-faces. Once an orientation has been selected for these faces, they will serve as the search space for our discretized algorithm. The faces should be chosen to cover the region of interest of the state space, i.e. the region surrounding the seed set $\mathcal{E}$. As in [18], [19], it is advantageous to start with some larger FIS obtained from another method, which can reduce the size of the region that needs to be searched. The search set of faces can be generated in a very simple way, such as a tiling of the region of interest, or by more advanced adaptive sampling methods that will be explored in future research.

Each time we first encounter a face during the search, we either assign it an orientation or remove it from the search space, as described in the previous section. We will also remove any faces that intersect with the interior of the seed set $\text{int}(\mathcal{E})$, as these faces could not be part of a boundary that fully contains $\mathcal{E}$. Since we will either assign a unique orientation to a face or remove the face from the search the first time it is visited, we will act as if each face in the search always has an assigned orientation. With this perspective, the discretized search space is a set of oriented (n-1)-simplexes that only intersect at shared sub-faces.

B. Discretized Problem Statement

Problem 2: Given a finite set $\mathcal{D}$ of oriented (n-1)-simplexes that only intersect at shared sub-faces, find the subset $S^* \subseteq \mathcal{D}$ that satisfies the constraints

1) (boundary) Each intersection in $I \in \mathcal{I}_{S^*}$ is consistent (as defined in Section IV-B)
2) (seed) $\mathcal{E} \subseteq \mathcal{C}_{S^*}$

and (optimality conditions) if some other $S \subseteq \mathcal{D}$ satisfies these constraints, then $\mathcal{C}_{S^*} \subseteq \mathcal{C}_S$.

This is a discretized version of Problem 1. Instead of finding the smallest possible FIS boundary, this problem seeks the smallest possible boundary in the discretized space $\mathcal{D}$ of oriented simplexes, which we denote as $\mathcal{S}_{\mathcal{D}} = S^*$. The invariance constraints in Problem 1 are enforced implicitly in Problem 2 by only allowing oriented simplexes to be a part of the discretized space $\mathcal{D}$ if they satisfy the invariance constraints. If any set S exists that satisfies constraints 1) and 2) of Problem 2, then $\mathcal{S}_{\mathcal{D}}$ exists and is unique. Instead of using a brute force search to find $\mathcal{S}_{\mathcal{D}}$, we use a greedy local search with backtracking.

C. Starting Faces

We use a method to similar [18], [19] to pick the starting faces to begin our search for the smallest discretized FIS $\mathcal{S}_{\mathcal{D}}$. For now, we will assume that the seed set $\mathcal{E}$ is path connected, but later we will show how this assumption can be removed. We first choose a smooth non-self-intersecting path $\rho : [0, 1] \to \mathcal{X}$ with the start point $\rho(0)$ in the seed set $\mathcal{E}$ and end point $\rho(1)$ outside of the region of interest. A simple choice is to make this path linear, which is similar to [18], [19]. However, we must also choose the path so that it only intersects faces transversally and it does not intersect with any sub-face in the discretized boundary space. In other words, the path can not intersect faces tangentially or at their "edges." If we inadvertently choose a path that violates these conditions, then we can locally deform the path to restore the conditions. With these assumptions, the path only intersects one face at a time, and we can order the faces that the path intersects $(\Delta_1, ..., \Delta_k)$ by the order in which they were intersected, i.e. the intersection times $s_i \in [0, 1]$ with $s_i < s_{i+1}$ have $\rho(s_i) \in \Delta_i$. We can also identify which faces were intersected on the outward-pointing or inward-pointing side, as specified by the normal orientation of the faces. Specifically, if $\langle \dot{\rho}(s_i), N_i \rangle > 0$, then face Δ_i was intersected on its inward-pointing side, if $\langle \dot{\rho}(s_i), N_i \rangle < 0$, then face Δ_i was intersected on its outward-pointing side, and $\langle \dot{\rho}(s_i), N_i \rangle \neq 0$ because of the transversality assumption.

We will search these starting faces in the order they are intersected by the path, ignoring faces that are intersected on their outward-pointing side. We try to find an oriented simplicial boundary containing each face in the order, and only move on to the next face if we find that no oriented simplicial boundary can contain the previous face.

Proposition 4: Assuming the seed set $\mathcal{E}$ is path connected, given a finite set $\mathcal{D}$ of oriented (n-1)-simplexes that only intersect at shared sub-faces and given a collection of faces $(\Delta_1, ..., \Delta_k)$ ordered by their intersection of a path ρ as defined above, the smallest discretized FIS $\mathcal{S}_{\mathcal{D}}$ contains the face with the lowest index in the order Δ_i that is intersected on its inward-pointing side by the path ρ and such that there exists an oriented simplicial boundary containing Δ_i.

Proof: Suppose Δ_i is the face with the lowest index in the order such that there exists an oriented simplicial boundary containing Δ_i. Since none of the faces before Δ_i are in any oriented simplicial boundary, no oriented simplicial boundaries will cross between $\mathcal{E}$ and Δ_i along ρ. Therefore, if ρ intersects Δ_i on its outward-pointing side, then any oriented simplicial boundary containing Δ_i will be path connected from its exterior to $\mathcal{E}$, so $\mathcal{E}$ must be in the exterior of the set enclosed by the boundary. Similarly, if Δ_i is intersected on its inward-pointing side, then any oriented simplicial boundary that contains Δ_i will enclose the seed set $\mathcal{E}$ in its interior. Furthermore, any oriented simplicial boundary that encloses $\mathcal{E}$ in its interior must also either enclose Δ_i in its interior or contain Δ_i. If an oriented simplicial boundary encloses Δ_i in its interior, then it is not the smallest discretized FIS, since there exists another oriented simplicial boundary containing Δ_i, and thus the intersection of the sets enclosed by these two FIS would be smaller. Therefore $\mathcal{S}_{\mathcal{D}}$ must contain Δ_i. ■

Proposition 4 tells us that we should look for the first face that is intersected by the path on its inward-pointing side and is contained by some discretized FIS, as this face will be part of the smallest discretized FIS $\mathcal{S}_{\mathcal{D}}$. Starting from a face like this, we will use similar reasoning to construct the rest of $\mathcal{S}_{\mathcal{D}}$ by choosing the next faces locally optimally.

We can also extend these results to the case when the seed set $\mathcal{E}$ has multiple path connected components. The simplest way to do this is to apply these results for each path connected component of $\mathcal{E}$ independently, and then join the resulting oriented simplicial boundaries.

D. Locally Optimal Faces

Now that we have a way to pick a starting face for the search, we can build upon this face by adding faces that must be in the smallest discretized FIS, if a discretized FIS exists that contains these faces. We will iteratively construct sets of included and removed faces by examining each intersection in the existing included set, removing faces that are inconsistent, and including additional faces that are locally optimal, as will be defined in this subsection.

At an intersection $I \in \mathcal{I}_{\mathcal{D}}$, for a face $\Delta \in I$, given a set of removed faces $R \in \mathcal{D}$, if it exists, we define $\mathcal{B}_{\Delta,I,R}$ the *locally optimal* face to Δ not in R as the face not in R that is consistent to Δ with minimum interior angle at intersection I. If the modified intersection is represented as $I \backslash R \equiv (..., (+, \Delta), (-, \Delta'), ...)$, then $\mathcal{B}_{\Delta,I,R} = \Delta'$.

In the following proposition, we say that an oriented simplicial boundary that contains S and does not contain R is an oriented simplicial boundary for (S, R).

Proposition 5: Given a finite set $\mathcal{D}$ of oriented (n-1)-simplexes that only intersect at shared sub-faces, given sets $S \subset \mathcal{D}$ and $R \subset \mathcal{D} \backslash S$, given an intersection $M \in \mathcal{I}_S$ with associated intersection $I \in \mathcal{I}_{\mathcal{D}}$ with $M \subseteq I$, and given a face $\Delta \in M$, if there exists an oriented simplicial boundary for $((S \cup \mathcal{B}_{\Delta,I,R}), R)$, then the smallest oriented simplicial boundary for (S, R) will also contain $\mathcal{B}_{\Delta,I,R}$.

Proof: Any oriented simplicial boundary S' for (S, R) must also include some element consistent to Δ at I. If S' contained such a consistent element $\Delta' \in I$ with $\Delta' \neq \mathcal{B}_{\Delta,I,R}$, then Δ' would be locally exterior to $\mathcal{B}_{\Delta,I,R}$ at I. Because of this, if there exists an oriented simplicial boundary S'' for $((S \cup \mathcal{B}_{\Delta,I,R}), R)$, then S' would not be the smallest oriented simplicial boundary for (S, R), since the intersection of the sets enclosed by S' and S'' would be smaller. Therefore, the smallest oriented simplicial boundary for (S, R) must also contain $\mathcal{B}_{\Delta,I,R}$. ■

E. Conditionally Included and Removed Faces

As a part of the search for $\mathcal{S}_{\mathcal{D}}$, we will iteratively include locally optimal faces, as described in the previous subsection. We will also remove faces that, if added, would make the currently included faces inconsistent, even if additional faces are included. Given sets $S \subset \mathcal{D}$ and $R \subset \mathcal{D} \backslash S$ and given an intersection $M \in \mathcal{I}_S$ with associated intersection I, we will remove any face $\Delta \in I$ such that M' is not consistent for any M' with $(M \cup \Delta) \subset M' \subset (I \backslash R)$.

We can represent these conditionally included and removed faces by a graph $G = (S \cup R, E)$, where the nodes of the graph are faces that are included S or removed R and where the directed edges represent the conditions under which each face was either added or removed. For an included face $\Delta \in S$, if we add another face Δ' that is locally optimal to Δ, then we create an edge $(\Delta, \Delta') \in E$. Furthermore, if there was another removed edge $\Delta'' \in R$ that would have been locally optimal to Δ instead of Δ' if Δ'' had not been removed, then we create another edge $(\Delta'', \Delta') \in E$. If we conditionally remove a face as described in the previous paragraph, then we create edges pointing to it from any face that is currently in S or R at the intersection. Using this representation, we can sometimes prove that $S = \mathcal{S}_{\mathcal{D}}$.

Theorem 2: Assuming the seed set $\mathcal{E}$ is path connected, given a finite set $\mathcal{D}$ of oriented (n-1)-simplexes that only intersect at shared sub-faces and given a graph $G = (S \cup R, E)$ as described above, if G is an ordered graph that is rooted at a start face Δ_i, as in Proposition 4, and if the locally optimal face to any terminal node of G in S is also in S, then $S = \mathcal{S}_{\mathcal{D}}$.

Proof: By construction, the set S is consistent at every intersection in $\mathcal{I}_S$. Since G is ordered, we can inductively prove, beginning at the terminal nodes in S and finishing at the root Δ_i, that a node in S is in the smallest discretized FIS for all of the higher level nodes in G. Since the graph is rooted at the start face Δ_i, as in Proposition 4, this is the smallest discretized FIS that encloses the seed set $\mathcal{E}$. ■

F. Greedy Search with Backtracking

Our goal is to now construct a graph of faces $G = (S \cup R, E)$ as in Theorem 2. We can do this by iteratively adding and removing faces as described in the previous subsection. The last difficulty is that during this process, we may be forced to remove some faces that were already added to S. This happens whenever we discover an intersection $M \in \mathcal{I}_S$ that is inconsistent regardless of any faces that we may additionally include, so M' is inconsistent for any M' with $M \subset M' \subset (I \backslash R)$. We can resolve this by removing one or more elements in M from S, along with its corresponding edges in the graph G. Then, in order to prevent the same situation from happening again, we remove any inconsistent faces at the intersection. Regardless of the details, this kind of backtracking procedure should not effect the optimally of the result, since as long as we meet the conditions of Theorem 2, we will have the optimal solution.

References

[1] F. Blanchini, "Survey paper: Set invariance in control," *Automatica (Journal of IFAC)*, vol. 35, no. 11, pp. 1747–1767, 1999.

[2] E. D. Sontag, "Input to state stability: Basic concepts and results," in *Nonlinear and optimal control theory*. Springer, 2008, pp. 163–220.

[3] M. Malisoff and F. Mazenc, *Constructions of Strict Lyapunov Functions*. Springer Science & Business Media, 2009.

[4] S. Raković, E. C. Kerrigan, K. I. Kouramas, and D. Q. Mayne, "Invariant approximations of the minimal robust positively invariant set," *Automatic Control, IEEE Transactions on*, vol. 50, no. 3, pp. 406–410, 2005.

[5] S. V. Raković and K. I. Kouramas, "Invariant approximations of the minimal robust positively invariant set via finite time aumann integrals," in *Decision and Control, 2007 46th IEEE Conference on*. IEEE, 2007, pp. 194–199.

[6] J. E. Gayek, "A survey of techniques for approximating reachable and controllable sets," in *Decision and Control, 1991., Proceedings of the 30th IEEE Conference on*. IEEE, 1991, pp. 1724–1729.

[7] O. Maler, "Computing reachable sets: An introduction," Tech. Rep., 2008.

[8] E. C. Kerrigan and J. M. Maciejowski, "Invariant sets for constrained nonlinear discrete-time systems with application to feasibility in model predictive control," in *Decision and Control, 2000. Proceedings of the 39th IEEE Conference on*, vol. 5. IEEE, 2000, pp. 4951–4956.

[9] P. Varaiya, "Reach set computation using optimal control," in *Verification of Digital and Hybrid Systems*. Springer, 2000, pp. 323–331.

[10] A. B. Kurzhanski and P. Varaiya, "Dynamic optimization for reachability problems," *Journal of Optimization Theory and Applications*, vol. 108, no. 2, pp. 227–251, 2001.

[11] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin, "A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games," *Automatic Control, IEEE Transactions on*, vol. 50, no. 7, pp. 947–957, 2005.

[12] S. Osher and R. Fedkiw, *Level set methods and dynamic implicit surfaces*. Springer Science & Business Media, 2006, vol. 153.

[13] J. Rieger, "Robust boundary tracking for reachable sets of nonlinear differential inclusions," *Foundations of Computational Mathematics*, vol. 15, no. 5, pp. 1129–1150, 2015.

[14] D. Limon, I. Alvarado, T. Alamo, and E. Camacho, "Robust tube-based mpc for tracking of constrained linear systems with additive disturbances," *Journal of Process Control*, vol. 20, no. 3, pp. 248–260, 2010.

[15] D. Q. Mayne, E. C. Kerrigan, E. Van Wyk, and P. Falugi, "Tube-based robust nonlinear model predictive control," *International Journal of Robust and Nonlinear Control*, vol. 21, no. 11, pp. 1341–1353, 2011.

[16] M. Farina and R. Scattolini, "Tube-based robust sampled-data mpc for linear continuous-time systems," *Automatica*, vol. 48, no. 7, pp. 1473–1476, 2012.

[17] F. Camilli, L. Grüne, and F. Wirth, "A generalization of zubov's method to perturbed systems," *SIAM Journal on Control and Optimization*, vol. 40, no. 2, pp. 496–515, 2001.

[18] S. Mukhopadhyay and F. Zhang, "A path planning approach to compute the smallest robust forward invariant sets," in *Proc. of 2014 American Control Conference (ACC2014)*, Portland, OR, 2014, pp. 1851–1856.

[19] S. Mukhopadhyay, "Robust forward invariant sets for nonlinear systems," Ph.D. dissertation, Georgia Institute of Technology, August 2014. [Online]. Available: http://hdl.handle.net/1853/52311

[20] J.-P. Aubin and A. Cellina, *Differential inclusions: set-valued maps and viability theory*. Springer Science & Business Media, 2012, vol. 264.

[21] F. H. Clarke, Y. S. Ledyaev, R. J. Stern, and P. R. Wolenski, *Nonsmooth Analysis and Control Theory*, ser. Graduate Texts in Mathematics. Springer-Verlag New York, 1998, vol. 178.

[22] A. Hatcher, *Algebraic Topology*. Cambridge University Press, Dec. 2001.

[23] M. Khachan, P. Chenin, and H. Deddi, "Digital pseudomanifolds, digital weakmanifolds and jordan–brouwer separation theorem," *Discrete applied mathematics*, vol. 125, no. 1, pp. 45–57, 2003.