



Sampling polynomial trajectories for LTL verification

Daniel Selvaratnam^{a,*}, Michael Cantoni^a, J.M. Davoren^a, Iman Shames^b

^a Department of Electrical and Electronic Engineering, University of Melbourne, Parkville, VIC 3010, Australia

^b College of Engineering and Computer Science, Australian National University, Canberra, ACT 2600, Australia



ARTICLE INFO

Article history:

Received 5 March 2021

Received in revised form 28 September 2021

Accepted 26 October 2021

Available online 3 November 2021

Communicated by J.-F. Raskin

Keywords:

MITL

Runtime verification

Dense-time

Continuous

Trace

Root isolation

ABSTRACT

This paper concerns the verification of continuous-time polynomial spline trajectories against linear temporal logic specifications (LTL without ‘next’). Each atomic proposition is assumed to represent a state space region described by a multivariate polynomial inequality. The proposed approach samples a trajectory strategically, to capture every one of its region transitions. This yields a discrete word called a trace, which is amenable to established formal methods for path checking. The original continuous-time trajectory is shown to satisfy the specification if and only if its trace does. General topological conditions on the sample points are derived that ensure a trace is recorded for arbitrary continuous paths, given arbitrary region descriptions. Using techniques from computer algebra, a trace generation algorithm is developed to satisfy these conditions when the path and region boundaries are defined by polynomials. The proposed POLYTRACE algorithm has polynomial complexity in the number of atomic propositions, and is guaranteed to produce a trace of any polynomial path. Its performance is demonstrated via numerical examples and a case study from robotics.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

Path checking [1] is about testing whether a particular signal (or word) satisfies a given formal specification. By contrast, the aim in model checking [2] is to establish whether all outputs from a model satisfy the specification. Path checking is generally much easier than model checking [3], particularly for infinite state models, for which the model checking problem is usually undecidable. Hence runtime verification, a popular approach to formal system verification [4], abandons model checking in favour of path checking, at the expense of exhaustive coverage. Path checking algorithms for discrete words are well studied [1,3,5]. We focus here on path checking continuous-time signals. Signal Temporal Logic (STL) is introduced in [6] as part of a framework for path checking the outputs of continuous-time and hybrid systems. STL is defined there as a fragment of Metric Temporal Logic (MTL) [7] in which each atomic proposition is associated with a region of the state-space. A procedure is then developed in [6] for path checking continuous-time signals against STL specifications with bounded temporal operators. This procedure relies on sampling the continuous-time signal in a manner assumed dense enough to capture all signal transitions between the relevant state space regions [6]. Satisfying this assumption is a decidedly non-trivial problem. Here, we present a sampling strategy that guarantees it, assuming that the signal to be verified is a polynomial spline, and that the state space regions of interest are semi-algebraic sets.

* Corresponding author.

E-mail address: selvaratnamd@unimelb.edu.au (D. Selvaratnam).

A useful comparison between path checking and model checking is provided in [5, Section 1.1]. A recurring issue relates to word length: most temporal logics are interpreted over infinite words, but we are practically constrained to store and process only finite words. In discrete-time, one way to resolve this is by restricting attention to *lasso words*, which are infinite words of the form $\alpha \cdot \beta^\omega$, for finite α, β . Path checking algorithms for such words are developed in [1,3]. Other ways to tackle the finite length issue for discrete words are discussed in [8]. In continuous time, a different approach has been to only path check fragments of MTL with bounded temporal operators [6,9,10], the satisfaction of which are fully determined by signal prefixes over finite time horizons. In contrast to this, we focus on Linear Temporal Logic (LTL) specifications, which can be considered fragments of MTL restricted to unbounded temporal operators. This choice is motivated by the control systems literature [11–13], in which controllers for continuous-time systems are synthesised to meet closed-loop LTL specifications. We address the verification problem, not the synthesis problem. The proposed sampling strategy generates a *trace* of the continuous-time signal to be verified, in line with the definition of such from [13]. Intuitively, a trace is a discrete word that captures every region transition taken by the continuous-time signal. In Section 4.2 we formalise the relationship between traces and the *timed state sequences* (TSSs) of [14]. We also identify a class of continuous-time signals having lasso word traces, and argue these are the only continuous-time signals with unbounded domains that are practically amenable to LTL path checking.

The notion of a robustness margin has proved valuable for path checking continuous-time signals by sampling. Intuitively, a robustness margin quantifies how robustly a formal specification is satisfied or how severely it is violated. In order to conclude that a continuous-time signal satisfies the specification, its sampled version is typically required to satisfy a corresponding discrete-time specification, with a sufficiently large positive robustness margin. Such an approach is proposed in [9], in which MTL semantics that can be interpreted over both discrete and continuous-time signals are developed. The signals are assumed to take values in a metric space, and a robustness margin is defined based on the metric. Moreover, the continuous-time signals are assumed to satisfy a generalised global Lipschitz condition [9, Assumption 35], and an upper bound on the sampling time is also required. Then, continuous-time satisfaction can be inferred from sufficiently robust discrete-time satisfaction. The same approach is adopted in [15]. The drawback is that the required sampling frequency may not be known *a priori*. Thus, if the robustness margin is not large enough for a particular sampled signal, then violation of the specification for the continuous-time signal cannot be inferred. The test can only be repeated with a different sampling strategy, which may again yield an inconclusive result. It is indeed possible for a continuous-time signal to satisfy or violate a specification with zero robustness, in which case the test never yields a conclusive result. A distinctive feature of our approach is that it does not rely on robustness margins at all. Since our algorithm is guaranteed to produce a trace, the trace can be checked against any LTL specification to always yield a conclusive result, even when the robustness margin is zero. To generate the trace, our sampling strategy does not depend on the specification imposed, neither does it rely on an upper bound between sampling times, thereby avoiding unnecessary sampling over intervals in which no transitions occur. This comes at the cost of considering a narrower class of continuous signals and regions of interest, specifically, those that are fully characterised by polynomials. However, such polynomial signals need not satisfy [9, Assumption 35].

Building on [9], subsequent works have estimated the robustness margin of a continuous-time signal from its samples by computing the robustness margin of its reconstructed version. Linear interpolation between sample points is assumed in [10,16]. This approach is generalised in [17] to allow for reconstruction via polynomial spline interpolation. Error bounds between the robustness margin of the original signal and the reconstructed spline are also derived. Our approach relates to the computation of the spline robustness margins in [17], but with some key differences. Firstly, we focus only on evaluating the boolean satisfaction of a specification by a signal, without attempting to compute its robustness margin. This allows us to verify edge cases that have zero (or almost zero) robustness, which would otherwise yield an inconclusive result. Secondly, [17] relies on numerical root finding to find inflection and intersection points of polynomials, without specifying a particular algorithm. Numerical root finding can be expensive. It also introduces errors, because, in general, polynomial roots are irrational, and cannot be found exactly for polynomials of degree five and higher. The choice of root finding algorithm is important, because this determines the overall complexity of the algorithm, and some root finders are not guaranteed to find every root (e.g. Newton’s method [18, Chapter 9]). Polynomial roots are also central to our approach, however we adopt root isolation techniques from computer algebra, which rely on symbolic computation over the rationals, and come with theoretical guarantees. The roots do not actually need to be found, only isolated. As a result, no numerical errors are introduced during our computations. Finally, the procedure in [17] appears to be sound, but it is stated without proof. A key contribution of our work is to provide a rigorous proof of correctness.

The restriction of scope to polynomial signals and semi-algebraic regions of interest may at first appear severe, however this class of verification problems is of practical significance. The use of polynomial splines for signal reconstruction in [17] has already been mentioned. Robotic motion planning is also frequently solved by constructing a polynomial spline through a sequence of waypoints [19–22]. Our algorithm can be used to verify a proposed motion plan against mission requirements expressed in LTL. This is illustrated by a case study in Section 7.2. A special case of this problem, focused only on detecting collisions, is treated in [23]. Collision detection is also the focus of [24], which assumes rational rather than polynomial spline paths. Semi-algebraic regions include half-spaces, circles, ellipses, hyperbolas, and finite intersections and unions thereof. These suffice to approximate obstacles or target regions to any desired degree of fidelity. Polynomials are of theoretical interest as well, because the celebrated Weierstrass Approximation Theorem assures us that any continuous function can be approximated uniformly and arbitrarily closely by a polynomial. Polynomials, then, are worthy of special attention, and our work exploits their properties for verification, by using the appropriate tools from computer algebra.

Overview: At its core, our algorithm involves sampling a continuous-time trajectory on either side of (and sometimes at) every crossing of the boundary of a region of interest. Since boundaries are a topological property, Section 2.2 provides some basic definitions and results from point-set topology, which are used along the way. We also rely on the preliminaries from algebra reviewed in Section 2.3 to exploit the algebraic properties of polynomials for their verification. The general problem we consider is formulated precisely in Section 3, which reviews the continuous-time semantics of MITL, and then presents LTL without next as a fragment of MITL. Section 4 lays the theoretical foundations of the paper. A distinction is made between two types of continuous signals: paths and trajectories, the former having a finite trace, and the latter an infinite one. This enables us to verify the class of trajectories that can be decomposed into a finite number of paths, because such trajectories have lasso word traces. Sections 4.1 and 4.3 address the decomposition of trajectories into paths, and the related construction of trajectory traces from path traces. The first main result of the paper is Theorem 1 in Section 4.4, which presents topological conditions to ensure that samples are taken on either side of every boundary crossing, and when necessary, at the crossings themselves. These conditions are stated in terms of sampling functions, which are introduced in Section 4.2 to describe arbitrary sampling strategies. The theory developed in Section 4 is applied in Section 5 to polynomial paths travelling through semi-algebraic regions of interest. Section 5.2, in particular, identifies the algebraic conditions that require samples to be taken exactly on a boundary. This precedes Theorem 2, which expresses the topological conditions of Theorem 1 in terms of polynomial roots. Section 6 then constructs the POLYTRACE algorithm to satisfy the conditions of Theorem 2. An important result is Lemma 7 (the Isolation Lemma), which permits a trajectory to be sampled exactly on a boundary, with only approximate knowledge of the crossing point. POLYTRACE relies on root isolation algorithms and a root existence test from computer algebra, which are discussed in Sections 6.2 and 6.3. The final main result is Theorem 3, which establishes the correctness of POLYTRACE. Numerical examples demonstrate the capabilities of the algorithm in Section 7, and a case study applies it to the verification of a robot motion plan. In particular, the case study imposes an LTL specification that can only be satisfied with zero robustness. Such specifications simply cannot be verified using existing methods, but yield readily to POLYTRACE. Concluding remarks are then made in Section 8.

2. Preliminaries

2.1. Notation

Let \mathbb{R} denote the real numbers, \mathbb{Q} the rational numbers, \mathbb{Z} the integers, and $\mathbb{N} := \{0, 1, \dots\}$ the natural numbers. For any $M \subset \mathbb{N}$, let $M^\ominus := \{m \in M \mid m + 1 \in M\}$. The subset relation is denoted by \subset , and the strict subset relation by \subsetneq . A set is *countable* iff it is either finite or countably infinite. If $f : A \rightarrow B$, then let $f[X] \subset B$ denote the image of $X \subset A$, and $f^{-1}Y \subset A$ the preimage of $Y \subset B$. The composition of $g : B \rightarrow C$ with $f : A \rightarrow B$ is denoted by $g \circ f : A \rightarrow C$. The restriction of f to $D \subset A$ is denoted by $f|_D$. For two integers $i \leq j$, let $a_{i:j} := (a_i, a_{i+1}, \dots, a_{j-1}, a_j)$. For products taken over the empty set, define $\prod_{i \in \emptyset} (\cdot) := 1$.

2.2. Topology

Let X be a topological space. A set $U \subset X$ is a neighbourhood of $x \in X$ iff U is open and $x \in U$. A topological space is *connected* iff it is not the union of two of its disjoint open subsets [25, §23]. A point $x \in X$ is an *isolated point* of X iff $\{x\}$ is open [26, Section 1.3]. Let $\text{iso}(X)$ denote the set of isolated points of X . Moreover, let $\text{int}(A)$ and $\text{cl}(A)$ denote the interior and closure of $A \subset X$, respectively. Then $\text{bd}(A) := \text{cl}(A) \setminus \text{int}(A)$ denotes the boundary of A , and $\text{ext}(A) := \text{int}(X \setminus A)$ the exterior of A . A point $x \in \text{bd}(A)$ iff every neighbourhood of x intersects both A and $X \setminus A$ [26, Proposition 1.3.1]. The result below is a corollary of this.

Proposition 1. For any topological space X , $\text{bd}(X) = \emptyset$.

Any $S \subset X$ becomes a topological space when endowed with the subset topology. That is, $V \subset S$ is declared open in S iff $V = S \cap U$ for some $U \subset X$ open in X . Then S is referred to as a subspace of X . It follows from the definition of the subset topology that $s \in \text{iso}(S)$ iff there exists $U \subset X$, open in X , such that $U \cap S = \{s\}$. We refer to $V \subset S$ as a S -neighbourhood of $s \in S$ iff V is open in S and $s \in V$.

Proposition 2. Let $A \subset X$, where X is a topological space. If X contains no isolated points, then $\text{iso}(A) \subset \text{bd}(A)$.

Proof. Let $a \in \text{iso}(A) \subset A$. There then exists $U \subset X$ open in X such that

$$U \cap A = \{a\}. \quad (1)$$

Choose any neighbourhood V of a . Now $a \in U \cap V \neq \emptyset$. Furthermore, $U \cap V$ is open in X . If $U \cap V = \{a\}$, then $a \in \text{iso}(X)$. Assume X contains no isolated points. Then $\{a\} \subsetneq U \cap V$, and there exists $x \in U \cap V$ such that $x \neq a$. If $x \in A$, then $x \in U \cap A$, which contradicts (1). Thus $x \in X \setminus A$, and $x \in V$, so $V \cap (X \setminus A) \neq \emptyset$. Clearly also $a \in V \cap A \neq \emptyset$. Thus $a \in \text{bd}(A)$. \square

Proposition 3. Let $A, B \subset X$, where X is a topological space. Then $\text{bd}(A \cap B) \subset \text{bd}(A) \cup \text{bd}(B)$.

Proof. By [26, Theorem 1.3.2, (iv)], $\text{bd}(A \cap B) \subset (\text{bd}(A) \cap \text{cl}(B)) \cup (\text{bd}(B) \cap \text{cl}(A))$. \square

Lemma 1. Let $A, B \subset X$, where X is a topological space. Then $\text{iso}(A \cap B) \setminus \text{bd}(B) \subset \text{iso}(A)$.

Proof. Suppose

$$s \in \text{iso}(A \cap B) \subset (A \cap B). \tag{2}$$

Then $V \cap A \cap B = \{s\}$, for some neighbourhood V of s . Suppose in addition that $s \notin \text{bd}(B)$. Then (2) implies $s \in B \setminus \text{bd}(B) = \text{int}(B)$ [26, Theorem 1.3.2, (i)]. Thus,

$$U \subset B \tag{3}$$

for some neighbourhood U of s . Observe that $V \cap U$ is also a neighbourhood of s . In particular, $s \in V \cap U$, and (2) then implies $s \in V \cap U \cap A$. Furthermore, $(V \cap U \cap A) \subset (V \cap B \cap A) = \{s\}$ by (3), and therefore $V \cap U \cap A = \{s\}$, which implies $s \in \text{iso}(A)$. \square

Suppose $A \subset S \subset X$, and let $\text{int}(A|S)$, $\text{cl}(A|S)$, $\text{bd}(A|S)$, $\text{ext}(A|S)$ respectively denote the interior, closure, boundary and exterior of A relative to S (i.e., with respect to the subset topology on S).

Remark 1. Since S is a topological space when endowed with the subset topology, Propositions 2, 3 and Lemma 1 all hold when $\text{bd}(\cdot)$ is replaced with $\text{bd}(\cdot|S)$, provided that $A, B \subset S \subset X$.

We will assume the standard topology on \mathbb{R}^n , which in particular implies that \mathbb{R}^n (and hence \mathbb{R}) have no isolated points.

Lemma 2. Let $U \subset S \subset \mathbb{R}$, and suppose $I \cap U \neq \emptyset$ for some interval $I \subset S$. Then

$$I \cap \text{bd}(U|S) = \emptyset \implies I \subset U.$$

Proof. Firstly, $S = \text{int}(U|S) \cup \text{bd}(U|S) \cup \text{ext}(U|S)$, and therefore

$$I = (I \cap \text{int}(U|S)) \cup (I \cap \text{bd}(U|S)) \cup (I \cap \text{ext}(U|S)). \tag{4}$$

By definition of the subset topology, $I \cap \text{int}(U|S)$ and $I \cap \text{ext}(U|S)$ are open in I . Suppose $I \cap \text{bd}(U|S) = \emptyset$. Since $U \cap I \neq \emptyset$, this implies $I \cap \text{int}(U|S) \neq \emptyset$. If $I \cap \text{ext}(U|S) \neq \emptyset$, then (4) implies I is the union of two non-empty disjoint open subsets. However, I is an interval and therefore connected, which is a contradiction. Thus $I \cap \text{ext}(U|S) = \emptyset$, and (4) then implies $I \subset \text{int}(U|S) \subset U$. \square

2.3. Algebra

Let $R[x_1, \dots, x_n]$ denote the ring of polynomials in the variables x_1, \dots, x_n with coefficients in R , where R is a commutative ring with multiplicative identity [27, Chapter 9]. Since \mathbb{R} is an infinite integral domain, we do not distinguish between a polynomial $p \in \mathbb{R}[x_1, \dots, x_n]$ and the polynomial function $p : \mathbb{R}^n \rightarrow \mathbb{R}$ associated with it [28, Chapter 3, Section 5.1]. Let $\text{deg}(p)$ denote the degree of p . In what follows, we are particularly concerned with the Euclidean domain $\mathbb{R}[s]$ of univariate polynomials with real coefficients. Take note of the standard definitions below.

- A polynomial $u \in \mathbb{R}[s]$ divides $p \in \mathbb{R}[s]$, denoted by $u \mid p$, iff $p = uq$ for some $q \in \mathbb{R}[s]$, which can also be written as $q = p/u$. The expression $u \nmid p$ asserts there is no such q .
- A polynomial is *monic* if its leading coefficient is 1.
- Let $u = \text{gcd}(p, q)$ denote the unique monic greatest common divisor of p and q . If $p \neq 0$, then $\text{gcd}(p, 0) = p/\text{lc}(p)$, where $\text{lc}(p) \in \mathbb{R} \setminus \{0\}$ is the leading coefficient of p . Also, $\text{gcd}(0, 0)$ is undefined.
- Polynomials p, q are *coprime* iff $\text{gcd}(p, q) = 1$.
- A polynomial $p \in \mathbb{R}[s]$ is *reducible* if $p = uv$ for some $u, v \in \mathbb{R}[s] \setminus \mathbb{R}$. A polynomial is *irreducible* iff it is both non-constant and not reducible [18, Chapter 25].
- A polynomial $p \in \mathbb{R}[s]$ is *square-free* [29] over \mathbb{R} iff $\forall g \in \mathbb{R}[s], (g^2 \mid p \implies g \in \mathbb{R})$.
- Let p' denote the derivative of $p \in \mathbb{R}[s]$.

Proposition 4. Let $p \in \mathbb{R}[s]$ be non-constant, and let $v := \text{gcd}(p, p')$. Then $q := p/v \in \mathbb{R}[s]$ is square-free. Furthermore, $(u \mid p \implies u \mid q)$, for any irreducible $u \in \mathbb{R}[s]$.

Proof. Proposition 6.6 of [28] states that q is square-free. Now it follows from [29, Corollary 1] that $q = \prod_{i=1}^m f_i$, where $p = \prod_{i=1}^m f_i^{e_i}$ is the irreducible factorisation [28, Chapter 3.4] of p . If u is an irreducible factor of p , then $u|f_i$ for some i , which implies $u|q$. \square

Corollary 1. Let $p \in \mathbb{R}[s]$ be non-constant and $q := p/\gcd(p, p')$. Then $p(s) = 0$ iff $q(s) = 0$.

The polynomial $q = p/\gcd(p, p')$ is referred to as the *square-free part* of p . Proposition 4 demonstrates that q contains every irreducible factor of p , but without multiplicities.

Proposition 5. Let $p, u \in \mathbb{R}[s]$ be non-zero. Then $u^2 | p$ iff both $u | p$ and $u | p'$.

Proof. The result holds trivially if $u \in \mathbb{R}$. Assume u is non-constant. If $p = gu^2$ for some $g \in \mathbb{R}[s]$, then $p' = g'u^2 + 2gu$, and therefore $u | p'$. Suppose now that $p = bu$ for some $b \in \mathbb{R}[s]$. Then p is non-constant, and

$$p' = b'u + bu' \neq 0 \tag{5}$$

Recall that $\deg(u') < \deg(u)$, and therefore $u \nmid u'$. If in addition $u | p'$, then (5) implies $b = cu$ for some $c \in \mathbb{R}[s]$, which in turn yields $p = cu^2$. \square

For any polynomial $p(s) = \sum_{i=0}^d p_i s^i \in \mathbb{R}[s]$, the inequalities $\|p\|_\infty \leq \|p\|_2 \leq d\|p\|_\infty$ hold, where $\|p\|_\infty := \max\{|p_i| \mid 0 \leq i \leq d\}$ and $\|p\|_2 := \sqrt{\sum_{i=0}^d p_i^2}$ [30, Chapter 3].

3. Problem formulation

3.1. Trajectories and observation maps

Definition 1. A trajectory is a continuous function $x : [0, \infty) \rightarrow X$.

A trajectory represents the evolution of a system over continuous time. We refer to X as the *state space* of the trajectory, which at this stage need only be a topological space. Note the variable x is used to represent both state values and trajectories, but its type is always clear from the context.

Given a finite set G of atomic propositions, an *observation map* is a function $h : X \rightarrow 2^G$, which associates each state with a set of propositions. Specifically, $h(x) \subset G$ is viewed as the set of propositions satisfied by the state $x \in X$. Consequently, the set of states that satisfy the proposition $g \in G$ is given by $\llbracket g \rrbracket := \{x \in X \mid g \in h(x)\}$. Any state space region that can be generated by taking a finite number of set unions, intersections and complements over the family $\llbracket g \rrbracket, g \in G$ is called a *region of interest* generated by h .

Remark 2. Consider a map $h : X \rightarrow A$, with A a finite set. The state space X can be partitioned into the family of subsets $h^{-1}\{a\}, a \in A$. By definition,¹ h is constant over each $h^{-1}\{a\}$. Moreover, if h is an observation map, then $A = 2^G$ and $h^{-1}\{a\} = \left(\bigcap_{g \in a} \llbracket g \rrbracket\right) \setminus \left(\bigcup_{g \in G \setminus a} \llbracket g \rrbracket\right)$, which shows that every $h^{-1}\{a\}$ is a region of interest generated by h .

Clearly $h \circ x(t)$ is the set of atomic propositions satisfied by the trajectory $x : [0, \infty) \rightarrow X$ at time $t \geq 0$. Thus, the composition $h \circ x : [0, \infty) \rightarrow 2^G$ is the primary object to be verified. Since it is a map from the non-negative reals into a finite set, we can expect $h \circ x$ to be piecewise constant when the trajectory and observation map are well-behaved. This property is typically expressed as a finite variability condition, which requires that $h \circ x$ change value only a finite number of times over any bounded time interval [14].

Definition 2 (Interval sequence). An interval sequence $\mathcal{I} = \{I_k \subset [0, \infty) \mid k \in \mathbb{N}\}$ is a set of non-empty intervals satisfying

- 2.a) $\forall k \geq 0, \sup I_k = \inf I_{k+1} < \infty$
- 2.b) $0 \in I_0$ and $\sup I_k \rightarrow \infty$ as $k \rightarrow \infty$
- 2.c) $\forall k \geq 0, I_k \cap I_{k+1} = \emptyset$
- 2.d) $\forall k \geq 0, \sup I_k \in I_k \cup I_{k+1}$.

Definition 2 implies $\bigcup \mathcal{I} = [0, \infty)$, and it is equivalent to [14, Definition 2.2.1], but has been stated here in a manner that is easier for us to check. Note it is impossible for both I_k and I_{k+1} to be singletons, as this would contradict 2.a)

¹ Recall that $h^{-1}B$ is the preimage of $B \subset A$ under h , as per Section 2.1. Here, $B = \{a\}$.

together with 2.c). A function z with domain S is constant over $I \subset S$ iff $z(s_1) = z(s_2)$ for all $s_1, s_2 \in I$. By this definition, every function is constant over the empty set.

Definition 3 (*Finite variability*). A function $z : S \rightarrow A$, with $S \subset [0, \infty)$ an interval, is of *finite variability* iff there exists an interval sequence \mathcal{I} such that z is constant over $I \cap S$ for every $I \in \mathcal{I}$.

A trajectory x is said to be of finite variability under an observation map h [12, Definition 3] iff $h \circ x$ is of finite variability. Only pathological pairs of trajectories and observation maps fail to meet this criterion. We exclude such cases from consideration, because they are of no practical interest.

3.2. Temporal logic semantics for continuous-time signals

In order to reason precisely about $h \circ x$, a temporal logic that can be interpreted over continuous-time signals is required. We adopt MITL [14, Section 2], which can be interpreted over finite variability signals via *timed state sequence* (TSS) representations. Note the definition below is a restatement of [14, Definition 2.2.1].

Definition 4. A TSS over a finite set G is a pair (ζ, \mathcal{I}) such that \mathcal{I} is an interval sequence and $\zeta : \mathbb{N} \rightarrow 2^G$.

This differs from the definition adopted in [9], which replaces each interval in \mathcal{I} with a single point sampled from it. Given a TSS (ζ, \mathcal{I}) over G , and an MITL formula φ over G , the satisfaction relation $(\zeta, \mathcal{I}) \models \varphi$ is defined in [14, Definition 2.3.2.1]. A grammar that generates MITL formulae is presented in [14, Definition 2.3.1.1]. These precise definitions are not critical to the development of this paper so we do not repeat them here, but focus instead on presenting the pertinent relationships between our work and the wider literature.

Definition 5. A TSS (ζ, \mathcal{I}) , with $\mathcal{I} = \{I_k \subset [0, \infty) \mid k \in \mathbb{N}\}$, is a *TSS representation* of $z : [0, \infty) \rightarrow 2^G$ iff

$$\forall k \geq 0, \forall t \in I_k, z(t) := \zeta(k).$$

Clearly every TSS represents a unique continuous-time signal, which by Definition 3 is of finite variability. Furthermore, given $z : [0, \infty) \rightarrow 2^G$ of finite variability, it is clear how a TSS representation (ζ, \mathcal{I}) may be constructed from the interval sequence in Definition 3. Finally by [14, Remark 2.3.2.2], if (ζ_1, \mathcal{I}_1) and (ζ_2, \mathcal{I}_2) are both TSS representations of $z : [0, \infty) \rightarrow 2^G$, then for any MITL formula φ ,

$$(\zeta_2, \mathcal{I}_2) \models \varphi \iff (\zeta_1, \mathcal{I}_1) \models \varphi.$$

MITL satisfaction can therefore be extended to continuous-time finite variability signals as follows.

Definition 6 (*MITL satisfaction*). Let G be a finite set of atomic propositions, φ an MITL formula over G , and $z : [0, \infty) \rightarrow 2^G$ a function of finite variability. The relation $z \models \varphi$ holds iff there exists a TSS representation (ζ, \mathcal{I}) of z such that $(\zeta, \mathcal{I}) \models \varphi$. Otherwise, $z \not\models \varphi$.

Having defined the satisfaction relation for well-behaved continuous-time signals precisely, we are now in a position to state the general problem.

Problem 1 (*MITL verification*). Given an MITL formula φ over a finite set of atomic propositions G , and a trajectory $x : [0, \infty) \rightarrow X$ of finite variability under an observation map $h : X \rightarrow 2^G$, determine whether $h \circ x \models \varphi$.

The works [6,9] address this problem for various fragments of MITL. In this paper, we restrict attention to a fragment of MITL that has temporal operators only over $[0, \infty)$, which corresponds to LTL without the next operator.

3.3. LTL as a fragment of MITL

The MITL operator \mathcal{U}_I in [14, Definition 2.3.2.1] is a strict non-matching until operator [31], which only imposes requirements at future times. The standard semantics of LTL [32] are stated in terms of the non-strict non-matching until \mathcal{U} , which includes the current time within its scope. The relationship between the two is formalised below.

Definition 7 (*LTL without next*). An MITL formula φ , over a finite set of atomic propositions G , is an LTL_{\circlearrowleft} formula over G iff it is generated by the restricted grammar

$$\varphi ::= g \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \mathcal{U} \varphi,$$

where $g \in G$, and $\varphi_1 \mathcal{U} \varphi_2 := \varphi_2 \vee (\varphi_1 \wedge (\varphi_1 \mathcal{U}_{[0, \infty)} \varphi_2))$.

In the context of trajectory verification, nothing is lost by omitting the LTL next operator from the above grammar. That operator would only have meaning for a fixed sampling strategy, which we do not wish to assume. In its place, nested \mathcal{U} operators could be used to impose a desired sequence of transitions without specifying when they must occur. More generally, for LTL_{\odot} formulas, the precise timing information in $h \circ x$ is not pertinent to the verification. The MITL satisfaction question for a TSS then reduces to an LTL satisfaction question for a discrete word, as per the remark below.

Remark 3. If φ is an LTL_{\odot} formula, then $(\zeta, \mathcal{I}) \models \varphi$ holds for the TSS (ζ, \mathcal{I}) iff $\zeta \models \varphi$ under standard discrete-time LTL semantics.

Remark 4. Definition 7 adopts the non-strict non-matching until \mathcal{U} for consistency with the standard discrete-time semantics of LTL. However, any of the other variants in [31] could just as well be adopted, without affecting the results in this paper. The only requirement is that the same until variant be adopted for both LTL_{\odot} and its discrete-time counterpart.

Given LTL path checking has been well studied, and algorithmic solutions to it exist [1,3], the missing step in the LTL_{\odot} verification of $h \circ x$ is the construction of the first component ζ of its TSS representation, which we call an infinite trace.

Definition 8 (Infinite trace). A function $\zeta : \mathbb{N} \rightarrow 2^G$ is an *infinite trace* of $z : [0, \infty) \rightarrow 2^G$ iff there exists an interval sequence \mathcal{I} such that (ζ, \mathcal{I}) is a TSS representation of z .

Thus we have shown that for LTL_{\odot} requirements, Problem 1 reduces to Problem 2 below.

Problem 2 (Infinite trace construction). Given a trajectory $x : [0, \infty) \rightarrow X$ of finite variability under an observation map $h : X \rightarrow 2^G$, construct an infinite trace of $h \circ x$.

The remainder of this paper is devoted to Problem 2. A complete algorithmic solution is provided for particular classes of trajectories and observation maps.

4. Sampling trajectories

4.1. Trajectories and paths

A path is a standard object in topology, which can be used to describe the geometry and ordering of portions of a trajectory, while abstracting away the timing information.

Definition 9. A *path* is a continuous function $r : [0, 1] \rightarrow X$.

A path has a compact domain, which forces its image to also be compact. This need not be true of a trajectory, and so not every trajectory can be fully represented by a path. Consider however the path $r(s) = x(a(1 - s) + bs)$, which describes the restriction of trajectory x to $[a, b]$. The coordinate transformation $s \mapsto a(1 - s) + bs$ preserves the geometry and direction of x . All the information relevant to the LTL verification of $h \circ x$ within the time interval $[a, b]$ is retained in $h \circ r$.

The relationship between paths and trajectories in continuous time mirrors the relationship between finite and infinite words in discrete time. As shown below, the prefix of a trajectory is a path, and the suffix of a trajectory is a trajectory.

Definition 10 (Trajectory prefix and suffix). Let $x : [0, \infty) \rightarrow X$ be a trajectory. Then for a given $T \geq 0$, the path $x_{\leq T} : [0, 1] \rightarrow X$, $x_{\leq T}(s) := x(Ts)$ is the *prefix* of x , and the trajectory $x_{\geq T} : [0, \infty) \rightarrow X$, $x_{\geq T}(t) := x(t + T)$ is the *suffix* of x .

Definition 11 (Path prefix and suffix). Let $r : [0, 1] \rightarrow X$ be a path. Then for a given $T \in [0, 1]$, the path $r_{\leq T} : [0, 1] \rightarrow X$, $r_{\leq T}(s) = r(Ts)$ is the *prefix* of r , and the path $r_{\geq T} : [0, 1] \rightarrow X$, $r_{\geq T}(s) := r(T + (1 - T)s)$ is the *suffix* of r .

The concatenation operation defined below is called the path product in [25].

Definition 12 (Path concatenation). Let $r_1, r_2 : [0, 1] \rightarrow X$ be paths such that $r_1(1) = r_2(0)$. Then the concatenation $r_1 \cdot r_2 : [0, 1] \rightarrow X$ is defined by

$$r_1 \cdot r_2(s) := \begin{cases} r_1(2s), & \text{if } s \in [0, \frac{1}{2}], \\ r_2(2s - 1), & \text{if } s \in (\frac{1}{2}, 1]. \end{cases}$$

The concatenation of two paths is a path [25, Chapter 9]. Trajectories may evolve along the same path in different ways. These correspond to having $x = r \circ u$ for different mappings $u : [0, \infty) \rightarrow [0, 1]$, which we refer to as *motions*.

Definition 13. A continuous map $u : [0, \infty) \rightarrow [0, 1]$ is a *direct motion* iff it is non-decreasing, $u(0) = 0$, and $u(t) \rightarrow 1$ as $t \rightarrow \infty$.

If u is a direct motion, then $r \circ u$ is called a *direct trajectory*, which either arrives at the endpoint $r(1)$ in finite time, or approaches it asymptotically. Although many direct motions are possible, we can expect the resulting trajectories to retain similar LTL properties when defined along the same path.

Definition 14. A *loop* is a path r such that $r(1) = r(0)$.

The map $\text{mod} : [0, \infty) \rightarrow [0, 1)$, given by $\text{mod}(s) := s - \lfloor s \rfloor$, returns s modulo 1.

Definition 15. A mapping $u : [0, \infty) \rightarrow [0, 1]$ is a *cyclic motion* iff there exists a non-decreasing surjection $v : [0, \infty) \rightarrow [0, \infty)$ such that $u = \text{mod} \circ v$.

If r is a loop, and u is a cyclic motion, the trajectory $r \circ u$ moves perpetually around the loop in one direction. Such trajectories are referred to as *cyclic*. They need not be periodic, because they can move around the loop at different speeds each time. Thus, with the aid of path concatenation and direct/cyclic motions, a rich class of trajectories can be constructed from a finite number of paths.

Remark 5. This framework of paths, motions and trajectories is consistent with the decomposition of the robotic motion planning problem into a *path planning* stage followed by a *path tracking* stage, discussed in [33]. Path planning generates the path r , and path tracking prescribes a motion u along the path. Together, they yield the trajectory $x = r \circ u$, which constitutes the motion plan.

4.2. Sampling functions for traces

Let (ζ, \mathcal{I}) be a TSS representation of $h \circ x$. According to Definition 5, the infinite trace ζ is sampled from $h \circ x$ once within each interval of \mathcal{I} . A sampling function can be used to represent this, as done in [9].

Definition 16 (Infinite sampling function). A map $\sigma : \mathbb{N} \rightarrow [0, \infty)$ is an *infinite sampling function* iff σ is strictly increasing, $\sigma(0) = 0$ and $\sigma(k) \rightarrow \infty$ as $k \rightarrow \infty$.

Observe that if σ is an infinite sampling function, then $h \circ x \circ \sigma : \mathbb{N} \rightarrow 2^{\mathcal{G}}$ is a discrete word, but not necessarily an infinite trace (see Definition 8). In the same way, paths can be sampled to obtain finite words.

Definition 17 (Finite sampling function). A map $\sigma : \{0, \dots, K\} \rightarrow [0, 1]$ is a *finite sampling function* iff σ is strictly increasing, $\sigma(0) = 0$ and $\sigma(K) = 1$.

Definition 18 (Sampling function). A *sampling function* is a finite sampling function or an infinite sampling function.

We refer to the values $\sigma(k)$ taken by a sampling function as *checkpoints*. Beyond Definition 16, extra conditions on σ are required to guarantee that $h \circ x \circ \sigma$ is an infinite trace. These are made explicit in the definition below, which generalises Definition 8 to allow for finite traces.

Definition 19 (Trace). Let A be a non-empty finite set, $\mathcal{K} \subset \mathbb{N}$ an interval containing 0, and either $S = [0, 1]$ or $S = [0, \infty)$. Then $\zeta : \mathcal{K} \rightarrow A$ is a *trace* of $z : S \rightarrow A$ iff there exist both a sampling function $\sigma : \mathcal{K} \rightarrow S$ and set $S^* := \{s_k^* \in S \mid k \in \mathcal{K}^{\ominus}\}$ such that all of the following hold:

- 19.a) $\zeta = z \circ \sigma$
- 19.b) $\sup \sigma[\mathcal{K}] = \sup S \in \{1, \infty\}$
- 19.c) $\forall k \in \mathcal{K}^{\ominus}, \sigma(k) \leq s_k^* \leq \sigma(k+1)$
- 19.d) $\forall k \in \mathcal{K}^{\ominus}, \forall s \in [\sigma(k), s_k^*), z(s) = z \circ \sigma(k)$
- 19.e) $\forall k \in \mathcal{K}^{\ominus}, \forall s \in (s_k^*, \sigma(k+1)], z(s) = z \circ \sigma(k+1)$
- 19.f) $\forall k \in \mathcal{K}^{\ominus}, z(s_k^*) \in \{z \circ \sigma(k), z \circ \sigma(k+1)\}$.

The sampling function σ is then said to *record* a trace of z . The checkpoints $\sigma(k)$ are the sample points, which are always distinct by Definitions 16 – 18. The set S^* includes all times (or path parameter values) at which z changes value, but z is not required to change value at every $s_k^* \in S$. It is possible that $s_k^* = s_{k+1}^*$ for some k , and this case is discussed further in Remark 8.

Remark 6 (Finiteness of trace). Clause 19.b) requires that σ be a finite sampling function when $S = [0, 1]$, and an infinite sampling function when $S = [0, \infty)$. Thus a trace of $h \circ y$ must have a finite domain if y is a path, and domain \mathbb{N} if y is a trajectory.

Remark 7 (Origins of trace definition). Definition 19 is based on [13, Definition 2], which in turn derives from [11,12]. Note some differences that are introduced here for convenience. Firstly, Definition 19 defines a trace in terms of a sampling function, and generalises [13, Definition 2] to permit traces with finite domains. Secondly, [13, Definition 2] is stated for x directly, whereas Definition 19 applies to $h \circ x$. Specifically, $\zeta : \mathbb{N} \rightarrow 2^G$ is a trace of $h \circ x$ under Definition 19 iff it is a trace of x under [13, Definition 2].

The following result establishes that Definitions 8 and 19 are consistent when $S = [0, \infty)$. Its proof is more involved than the simplicity of the result might suggest, because of the number of clauses in both Definitions 19 and 2 that must be verified.

Proposition 6. A mapping $\zeta : \mathbb{N} \rightarrow 2^G$ is a trace of $z : [0, \infty) \rightarrow 2^G$ iff it is an infinite trace of z .

Proof. (\implies) Suppose ζ is a trace of z . Then by Remark 6, there exist an infinite sampling function $\sigma : \mathbb{N} \rightarrow [0, \infty)$ and set $S^* := \{s_k^* \geq 0 \mid k \in \mathbb{N}\}$ satisfying all clauses of Definition 19. By Clause 19.a), $\zeta = z \circ \sigma$. Since σ is strictly increasing, Clause 19.c) implies $0 \leq s_0^* \leq s_1^* \dots$, where each $s_k^* \in S^*$. An interval sequence \mathcal{I} is now constructed by choosing $s_{k-1}^*, s_k^* \in S^*$ as the endpoints of every $I_k \in \mathcal{I}$. To ensure mutual disjointness, the right endpoint s_k^* should be included in I_k iff

$$\psi(k) : z(s_k^*) = \zeta(k) \text{ and } s_k^* < s_{k+1}^*, \tag{6}$$

where the predicate ψ , defined over all $k \in \mathbb{N}$, has been introduced to clarify the subsequent presentation. Let

$$I_0 := \begin{cases} [0, s_0^*], & \text{if } \psi(0) \\ [0, s_0^*), & \text{if } \neg\psi(0) \end{cases}, \tag{7}$$

and for each $k > 0$,

$$I_k := \begin{cases} [s_{k-1}^*, s_k^*], & \text{if } \neg\psi(k-1) \wedge \psi(k) \\ (s_{k-1}^*, s_k^*], & \text{if } \psi(k-1) \wedge \psi(k) \\ [s_{k-1}^*, s_k^*), & \text{if } \neg\psi(k-1) \wedge \neg\psi(k) \\ (s_{k-1}^*, s_k^*), & \text{if } \psi(k-1) \wedge \neg\psi(k) \end{cases}. \tag{8}$$

Clearly the above case definition covers all possibilities. It must be shown that $\mathcal{I} := \{I_k \mid k \in \mathbb{N}\}$ is an interval sequence. Each I_k is an interval by construction, so it must now be shown that they are non-empty. If $s_0^* > 0$, then $I_0 \neq \emptyset$. Suppose on the other hand that $s_0^* = 0$. Then because σ is strictly increasing,

$$0 = s_0 = \sigma(0) < \sigma(1) \leq s_1^*$$

by Clause 19.c), which implies $\psi(0)$ is true. Equation (7) in turn implies $I_0 = [0, 0] = \{0\}$. Consider now $k > 0$. If $s_{k-1}^* < s_k^*$, then $I_k \neq \emptyset$. Suppose instead that $s_{k-1}^* = s_k^*$. Then

$$\sigma(k-1) < s_{k-1}^* = \sigma(k) = s_k^* < \sigma(k+1) \leq s_{k+1}^*,$$

because σ is strictly increasing. This implies $\psi(k)$ is true and $\psi(k-1)$ is false, and (8) in turn implies $I_k = [s_{k-1}^*, s_k^*] = \{s_k^*\}$. This establishes that $I_k \neq \emptyset$ for all $k \in \mathbb{N}$, as required by Definition 2. Given this, Clause 2.a) clearly follows from (7) – (8). To establish Clause 2.b), first note that $0 \in I_0$ by (7). Now since $\sigma(k) \rightarrow \infty$ by Definition 16, Clause 19.c) then implies $s_k^* = \sup I_k \rightarrow \infty$ as $k \rightarrow \infty$. Turning to the next clause, Equations (7) – (8) imply $I_k \cap I_{k+1} \subset \{s_k^*\}$ for all $k \geq 0$. If $s_k^* \in I_k$, then (7) – (8) imply $\psi(k)$ is true, which further implies $s_k^* \notin I_{k+1}$. This establishes Clause 2.c). Conversely, if $s_k^* \notin I_k$, then $\psi(k)$ is false, which implies $s_k^* \in I_{k+1}$ by (8), establishing Clause 2.d). Thus \mathcal{I} is an interval sequence.

It remains to be shown that (ζ, \mathcal{I}) is a TSS representation of z . Choose any $s \in I_0 \subset [0, s_0^*]$. If $s \in [0, s_0^*)$, then Clause 19.d) implies $z(s) = z \circ \sigma(0) = \zeta(0)$. Suppose instead that $s = s_0^*$. Since $s \in I_0$, (7) then implies $\psi(0)$ is true, by which $z(s) = z(s_0^*) = \zeta(0)$. Thus $z(s) = \zeta(0)$ for all $s \in I_0$. Now choose any $k > 0$ and $s \in I_k \subset [s_{k-1}^*, s_k^*]$. If $s \in (s_{k-1}^*, s_k^*) = (s_{k-1}^*, \sigma(k)) \cup [\sigma(k), s_k^*)$, Clauses 19.d) – 19.e) imply $z(s) = z \circ \sigma(k) = \zeta(k)$. Alternatively, if $s = s_k^*$, then $\psi(k)$ is true by (8), implying $z(s) = z(s_k^*) = \zeta(k)$. The only remaining possibility for $s \in I_k$ is that $s = s_{k-1}^*$. In that case, $\psi(k-1)$ is false, which implies that $z(s_{k-1}^*) \neq \zeta(k-1)$ or $s_{k-1}^* = s_k^*$. If the latter, then $s = s_k^*$ and it has already been shown that $z(s) = \zeta(k)$. If the former, then Clause 19.f) implies $z(s) = z(s_{k-1}^*) = \zeta(k)$. It has been shown that, for any $k \geq 0$, for any $s \in I_k$, $z(s) = \zeta(k)$, satisfying Definition 5. Therefore ζ is an infinite trace of z by Definition 8.

(\impliedby) Suppose (ζ, \mathcal{I}) is a TSS representation of z . A sampling function σ and a set S^* must now be constructed to satisfy all clauses of Definition 19. Definition 2 implies each $I_k \in \mathcal{I}$ is non-empty. Define $S^* := \{s_k^* \geq 0 \mid k \in \mathbb{N}\}$ where

$s_k^* := \sup I_k$ for all $I_k \in \mathcal{I}$. Let $\sigma : \mathbb{N} \rightarrow [0, \infty)$ be a function such that $\sigma(0) = 0$ and $\sigma(k) \in I_k$ for all $k > 0$ (e.g., choose $\sigma(k)$ to be the midpoint of I_k). It then follows from Definition 5 that $\zeta = z \circ \sigma$, satisfying Clause 19.a). Note also that $0 \in I_0$ by Clause 2.b). Clause 2.a) now implies that σ is non-decreasing, and Clause 2.c) then that it is strictly increasing. Clauses 2.a) and 2.b) also imply that $\inf I_{k+1} = \sup I_k \rightarrow \infty$, and therefore $\sigma(k+1) \rightarrow \infty$. Thus σ is an infinite sampling function by Definition 16, which in particular, satisfies Clause 19.b). To establish Clause 19.c), $\sigma(k) \leq \sup I_k = s_k^* = \inf I_{k+1} \leq \sigma(k+1)$ for all $k \geq 0$. To establish Clause 19.d), choose any $k \geq 0$. If $\sigma(k) = s_k^*$, then Clause 19.d) holds trivially for that k . Assume instead that $\sigma(k) < s_k^*$ and choose any $s \in [\sigma(k), s_k^*]$. Since $\sigma(k) \in I_k$ and $s_k^* = \sup I_k$, it follows that $s \in I_k$. Since also $\sigma(k) \in I_k$, Definition 5 then implies $z(s) = z \circ \sigma(k) = \zeta(k)$. To establish Clause 19.e), choose any $k \geq 0$. If $s_k^* = \sigma(k+1)$, then Clause 19.e) holds trivially for that k . Assume instead that $s_k^* < \sigma(k+1)$ and choose any $s \in (s_k^*, \sigma(k+1)]$. Since $\sigma(k+1) \in I_{k+1}$ and $s_k^* = \sup I_k = \inf I_{k+1}$, it follows that $s \in I_{k+1}$. Since also $\sigma(k+1) \in I_{k+1}$, Definition 5 then implies $z(s) = z \circ \sigma(k+1) = \zeta(k+1)$. Finally, Clause 19.f) is directly implied by Clause 2.d) together with Definition 5. \square

Remark 8. The proof of Proposition 6 reveals that S^* in Definition 19 consists of the endpoints of the intervals in the corresponding interval sequence in Definition 8. If some $I_m \in \mathcal{I}$ is a singleton, then $s_{m-1}^* = \sigma(m) = s_m^*$, which makes Clause 19.d) vacuous at $k = m$ and Clause 19.e) vacuous at $k = m - 1$. In this case, it is possible that all three of the following hold:

$$z(s) = z \circ \sigma(m - 1) \text{ for all } s \in [\sigma(m - 1), s_m^*), \tag{9}$$

$$z(s) = z \circ \sigma(m + 1) \text{ for all } s \in (s_m^*, \sigma(m + 1)], \tag{10}$$

$$z \circ \sigma(m - 1) \neq z(s_m^*) \neq z \circ \sigma(m + 1), \tag{11}$$

indicating that z only holds its value at s_m^* instantaneously. To record a trace, the sampling function must place checkpoint $\sigma(m)$ exactly at s_m^* . A uniform sampling strategy almost surely fails to achieve this. The coming sections develop a trace generation algorithm capable of computing $z(s_m^*) = h \circ x(s_m^*)$ exactly, for polynomial trajectories x , and observation maps h that only generate semi-algebraic regions of interest.

4.3. Manipulating traces

In this subsection it is shown how traces of paths may be combined, and traces of trajectories decomposed, into traces of paths. We first specify some relevant operations on discrete words.

Definition 20 (Word operations). The definitions below hold for any $\alpha : \mathcal{K} \rightarrow A$ and $\beta : \mathcal{L} \rightarrow A$, with $\mathcal{L} \subset \mathbb{N}$ an interval and $\mathcal{K} \subset \mathbb{N}$ a finite interval such that $0 \in \mathcal{K} \cap \mathcal{L}$. Define $\mathcal{K} \oplus \mathcal{L} := \{k + l \in \mathbb{N} \mid k \in \mathcal{K}, l \in \mathcal{L}\}$.

- 20.a) *Letter repetition:* $\alpha^\omega : \mathbb{N} \rightarrow \{a\}$, given $a \in A$.
- 20.b) *Word repetition:* $\alpha^\omega : \mathbb{N} \rightarrow A$, with $\alpha^\omega(k) := \alpha(k \bmod |\mathcal{K}|)$ for all $k \in \mathbb{N}$.
- 20.c) *Word prefix:* $\beta_{<k} := \beta \upharpoonright_{\{l \in \mathcal{L} \mid l < k\}}$ and $\beta_{\leq k} := \beta \upharpoonright_{\{l \in \mathcal{L} \mid l \leq k\}}$, given $k \in \mathbb{N}$.
- 20.d) *Word concatenation:* $\alpha \cdot \beta : \mathcal{K} \oplus \mathcal{L} \rightarrow A$, where

$$\alpha \cdot \beta(k) := \begin{cases} \alpha(k), & \text{if } k \in \mathcal{K} \\ \beta(k - |\mathcal{K}|), & \text{if } k \notin \mathcal{K} \end{cases}$$

20.e) *Letter concatenation:* $\alpha \cdot a := \alpha \cdot \vec{a}$, where $\vec{a} : \{0\} \rightarrow \{a\}$, given $a \in A$.

Remark 9. Every non-decreasing surjection $\gamma : [0, 1] \rightarrow [0, 1]$ is continuous, and can be viewed as a path reparameterisation.

Proposition 7 (Path reparameterisation). Let $r : [0, 1] \rightarrow X$ be a path, $\gamma : [0, 1] \rightarrow [0, 1]$ a non-decreasing surjection, and $h : X \rightarrow A$, with A finite. Every trace of $h \circ r$ is a trace of $h \circ r \circ \gamma$.

Proof. Let ζ be a trace of $h \circ r$. By Remark 6, there exist a finite sampling function $\sigma : \{0, \dots, K\} \rightarrow [0, 1]$ and set $S^* = \{s_k^* \in [0, 1] \mid 0 \leq k < K\}$ that satisfy all clauses of Definition 19. Clause 19.a) implies $\zeta = h \circ r \circ \sigma$. First observe that, for any $s \in [0, 1]$, $\gamma^{-1}\{s\} \subset [0, 1]$ is non-empty because γ is surjective, and closed because γ is continuous (see Remark 9). Define $\theta : \{0, \dots, K\} \rightarrow [0, 1]$ as

$$\theta(k) := \begin{cases} 0, & \text{if } k = 0 \\ \max \gamma^{-1}\{\sigma(k)\}, & \text{if } 0 < k \leq K \end{cases}, \tag{12}$$

and construct the set $U^* := \{u_k^* \in [0, 1] \mid 0 \leq k < K\}$ by choosing

$$u_k^* := \max \gamma^{-1}\{s_k^*\} \tag{13}$$

for all $0 \leq k < K$. Given that σ and S^* satisfy all clauses of Definition 19 for $z = h \circ r$, the task is now to show that θ is a sampling function, and that θ and U^* satisfy all clauses of Definition 19 for $z = h \circ r \circ \gamma$.

By Definition 17, $\sigma(0) = \gamma(0) = 0$, implying that $\gamma \circ \theta(k) = \sigma(k)$ for all $0 \leq k \leq K$. It follows that $\gamma \circ \theta = \sigma$, and therefore $\zeta = h \circ r \circ \gamma \circ \theta$. Thus θ satisfies Clause 19.a). Furthermore, σ is injective by Definition 17, which implies $\gamma^{-1}\{\sigma(j)\} \cap \gamma^{-1}\{\sigma(k)\} = \emptyset$ whenever $j \neq k$. This in turn implies that θ is injective, and hence strictly increasing because γ and σ are both non-decreasing. Recall that $\gamma(1) = \sigma(K) = 1$, which implies that $\theta(K) = 1$. This establishes that θ is a finite sampling function, which in particular satisfies Clause 19.b). By Clause 19.c), $\gamma \circ \theta(k) \leq \gamma(u_k^*) \leq \gamma \circ \theta(k+1)$ for all $0 \leq k < K$. For $k = 0$, $\theta(0) = 0 \leq u_0^*$ by definition. Consider now any $0 < k < K$. If $\gamma \circ \theta(k) < \gamma(u_k^*)$, then $\theta(k) < u_k^*$ because γ is non-decreasing. Alternatively if $\gamma \circ \theta(k) = \gamma(u_k^*)$, then it follows from (12) and (13) that $\theta(k) = u_k^*$. In the same way, it can be shown that $u_k^* \leq \theta(k+1)$ for all $0 \leq k < K$, establishing Clause 19.c). Turning to Clause 19.d), let $0 \leq k < K$ and suppose $\theta(k) \leq s < u_k^*$. Then $\sigma(k) = \gamma \circ \theta(k) \leq \gamma(s) \leq \gamma(u_k^*) = s_k^*$, which implies $h \circ r \circ \gamma(s) = h \circ r \circ \sigma(k) = h \circ r \circ \gamma \circ \theta(k)$. Clause 19.e) can be established in a similar manner, and Clause 19.f) holds because $\sigma = \gamma \circ \theta$ and $s_k^* = \gamma(u_k^*)$ for all $0 \leq k < K$. \square

Corollary 2. Let $r : [0, 1] \rightarrow X$ be a path, $u : [0, \infty) \rightarrow [0, 1]$ a direct motion, and $h : X \rightarrow A$, with A finite. Then for any $T \geq 0$, every trace of $h \circ r_{\leq u(T)}$ is a trace of $h \circ (r \circ u)_{\leq T}$.

Proof. If $u(T) = 0$, then $u(t) = 0$ for all $t \in [0, T]$, and therefore $r_{\leq u(T)}(s) = r(0) = (r \circ u)_{\leq T}(s)$ for all $s \in [0, 1]$. If instead $u(T) > 0$, then $(r \circ u)_{\leq T} = r_{\leq u(T)} \circ \gamma$, where $\gamma : [0, 1] \rightarrow [0, 1]$, $\gamma(s) := \frac{u(Ts)}{u(T)}$ is a non-decreasing surjection. \square

Propositions 8 – 12 below are consequences of the preceding developments, stated here without proof.

Proposition 8 (Snipping Lemma). Let $r : [0, 1] \rightarrow X$ be a path, and $h : X \rightarrow A$, with A finite. Suppose $\sigma : \mathcal{K} \rightarrow [0, 1]$ is a sampling function that records the trace ζ of $h \circ r$, and $T \in [\sigma(k), \sigma(k+1)]$ for some $k \in \mathcal{K}^\ominus$. If $h \circ r(T) = \zeta(k+1)$, then $\zeta_{\leq k+1}$ is a trace of $h \circ r_{\leq T}$. Otherwise, $\zeta_{\leq k} \cdot \zeta(k)$ is a trace of $h \circ r_{\leq T}$.

The concatenation of $\zeta_{\leq k}$ with $\zeta(k)$ in the Snipping Lemma is necessary to handle the case where $k = 0$, because Definition 19 requires the domain of a trace to contain at least two elements, but the domain of $\zeta_{\leq 0}$ is a singleton.

Proposition 9 (Path concatenation). Let $r_1, r_2 : [0, 1] \rightarrow X$ be paths such that $r_1(1) = r_2(0)$, and $h : X \rightarrow A$, with A finite. If ζ_1 is a trace of $h \circ r_1$, and ζ_2 a trace of $h \circ r_2$, then $\zeta_1 \cdot \zeta_2$ is a trace of $h \circ (r_1 \cdot r_2)$.

We refer to a trajectory x as *invariant* under h iff $h \circ x$ is constant.

Proposition 10 (Invariant trajectories). Let $x : [0, \infty) \rightarrow X$ be a trajectory, and $h : X \rightarrow A$, with A finite. If $h \circ x$ is constant, then $(h \circ x(0))^\omega$ is a trace of $h \circ x$.

Proposition 11 (Trajectory suffixes and prefixes). Let $x : [0, \infty) \rightarrow X$ be a trajectory, and $h : X \rightarrow A$, with A finite. For any $T \geq 0$, if ζ_1 is a trace of $h \circ x_{\leq T}$ and ζ_2 a trace of $h \circ x_{\geq T}$, then $\zeta_1 \cdot \zeta_2$ is a trace of $h \circ x$.

Corollary 3 (Invariant suffixes). Let $x : [0, \infty) \rightarrow X$ be a trajectory, and $h : X \rightarrow A$, with A finite. Suppose there exists $T \geq 0$ such that $h \circ x_{\geq T}$ is constant. If $\zeta : \{0, \dots, K\} \rightarrow A$ is a trace of $h \circ x_{\leq T}$, then $\zeta \cdot \zeta(K)^\omega$ is a trace of $h \circ x$.

Proof. The result follows directly from Propositions 10 and 11. \square

Proposition 12 (Cyclic trajectories). Let $r : [0, 1] \rightarrow X$ be a loop, $u : [0, \infty) \rightarrow [0, 1]$ a cyclic motion, and $h : X \rightarrow A$, with A finite. If ζ is a trace of $h \circ r$, then ζ^ω is a trace of $h \circ r \circ u$.

Corollary 4 (Cyclic suffixes). Let $x : [0, \infty) \rightarrow X$ be a trajectory, and $h : X \rightarrow A$, with A finite. Suppose there exist a loop $r : [0, 1] \rightarrow X$, cyclic motion $u : [0, \infty) \rightarrow [0, 1]$, and $T \geq 0$ such that $x_{\geq T} = r \circ u$. If ζ_1 is a trace of $h \circ x_{\leq T}$ and ζ_2 a trace of $h \circ r$, then $\zeta_1 \cdot \zeta_2^\omega$ is a trace of $h \circ x$.

Proof. The result follows directly from Propositions 11 and 12. \square

Surprisingly, direct trajectories require a more verbose result. The complication arises when one asymptotically approaches a point on the boundary of a region of interest, because it never actually gets there. Notwithstanding this, all direct trajectories have invariant suffixes, as established below.

Proposition 13 (Direct trajectories). Let $r : [0, 1] \rightarrow X$ be a path, $u : [0, \infty) \rightarrow [0, 1]$ a direct motion, and $h : X \rightarrow A$, with A finite. Suppose $\zeta : \{0, \dots, K\} \rightarrow A$ is a trace of $h \circ r$. Then, for $x := r \circ u$, there exists $T \geq 0$ such that $h \circ x_{\geq T}$ is constant. Moreover, if at least one of the following hold:

$$\exists t \geq 0, u(t) = 1, \quad (14)$$

$$\exists \varepsilon \in [0, 1), \forall s \in (\varepsilon, 1], h \circ r(s) = \zeta(K), \quad (15)$$

then $\zeta \cdot \zeta(K)^\omega$ is a trace of $h \circ x$. Otherwise, $\zeta_{<K} \cdot \zeta(K-1)^\omega$ is a trace of $h \circ x$.

Proof. Since ζ is a trace of $h \circ r$, by Remark 6, there exist a finite sampling function $\sigma : \{0, \dots, K\} \rightarrow [0, 1]$ and set $S^* = \{s_k^* \in [0, 1] \mid 0 \leq k < K\}$ that satisfy all clauses of Definition 19. Clause 19.a) implies $\zeta = h \circ r \circ \sigma$. There are now three cases to consider, which cover all possibilities.

- Suppose (14) holds, and choose any $T \geq 0$ such that $u(T) = 1$. Since u is non-decreasing, it follows that $u(t) = 1$ for all $t \geq T$. Noting that $\sigma(K) = 1$, it follows that $h \circ r \circ u(t) = \zeta(K)$ for all $t \geq T$, implying that $h \circ (r \circ u)_{\geq T}$ is constant. Now Corollary 2 implies that ζ is a trace of $h \circ (r \circ u)_{\leq T}$ because it is a trace of $h \circ r$, and applying Corollary 3 then yields the result.
- Suppose (15) holds. Choose $\varepsilon \in [0, 1)$ as in (15). Since u is non-decreasing and $u(t) \rightarrow 1$,

$$\exists T \geq 0, \forall t \geq T, \max\{\varepsilon, \sigma(K-1)\} < u(t) \leq 1 = \sigma(K), \quad (16)$$

by which (15) implies $h \circ r \circ u(t) = \zeta(K)$ for all $t \geq T$. This shows that $h \circ (r \circ u)_{\geq T}$ is constant, and in particular that $h \circ r \circ u(T) = \zeta(K)$. Noting also that $u(T) \in (\sigma(K-1), \sigma(K)]$ by (16), the Snipping Lemma then implies that ζ is a trace of $h \circ r_{\leq u(T)}$. Corollary 2 in turn implies that ζ is a trace of $h \circ (r \circ u)_{\leq T}$, and applying Corollary 3 then yields the result.

- Suppose (14) and (15) are both false. If $s_{K-1}^* < 1$, then Clause 19.e) implies (15), which is a contradiction. Thus $s_{K-1}^* = 1$. Since $u(t) \rightarrow 1$, it follows that

$$\exists T \geq 0, \forall t \geq T, \sigma(K-1) < u(t) < s_{K-1}^* = \sigma(K) = 1. \quad (17)$$

Thus, $h \circ r \circ u(t) = \zeta(K-1)$ for all $t \geq T$. This shows that $h \circ (r \circ u)_{\geq T}$ is constant, and in particular that $h \circ r \circ u(T) = \zeta(K-1)$. Noting also that $u(T) \in (\sigma(K-1), \sigma(K))$ by (17), the Snipping Lemma then implies that $\zeta_{<K} \cdot \zeta(K-1)$ is a trace of $h \circ r_{\leq u(T)}$. Corollary 2 in turn implies that $\zeta_{<K} \cdot \zeta(K-1)$ is also a trace of $h \circ (r \circ u)_{\leq T}$, and applying Corollary 3 yields the result. \square

Remark 10. Only trajectories with cyclic or invariant suffixes are practically amenable to path checking, because in the absence of a model for the underlying dynamics, these are the only types of trajectories that offer information infinitely far into the future. Corollaries 3 and 4 establish that both have lasso word traces constructed from the traces of at most two paths. LTL verification of these traces can be performed using the discrete path-checking algorithms of [1,3].

Remark 10 implies that all practical solutions to Problem 2 can be constructed from solutions to the problem below, which is conducive to an algorithmic solution because it always has a finite output. Just as for trajectories, we say that a path r is of finite variability under an observation map h iff $h \circ r$ is of finite variability.

Problem 3 (Finite trace construction). Given a path $r : [0, 1] \rightarrow X$ of finite variability under an observation map $h : X \rightarrow 2^G$, construct a trace of $h \circ r$.

Observe that this places no restrictions on the LTL formula that can be satisfied. For any LTL formula φ , there exists a non-deterministic Buchi automaton (NBA) that accepts all and only the words satisfying φ [34, Corollary 23]. Furthermore, if the language accepted by an NBA is nonempty, it contains a lasso word [35, Proposition 3.3.1]. The remark below is a corollary of this.

Remark 11. An LTL formula is satisfiable iff there exists a lasso word that satisfies it.

Remark 12 (Implications for robotic motion planning). In Corollary 4 and Propositions 12 – 13, the trace of $h \circ r \circ u$ depends primarily on the path r , and is invariant to the motion u under only mild conditions. In the context of robotic motion planning (see Remark 5), this means that verification of the entire motion plan depends only on the path r , provided that the path tracking stage is restricted to produce a suitable class of motions.

4.4. Topological conditions on sampling functions

The result below provides a topological condition that is sufficient for $h \circ r$ or $h \circ x$ to be constant over an interval.

Corollary 5. *Let $z : S \rightarrow A$, with A finite and $S \subset \mathbb{R}$ an interval, and let $I \subset S$ also be an interval. If*

$$\forall a \in A, I \cap \text{bd}(z^{-1}\{a\} | S) = \emptyset, \tag{18}$$

then z is constant over I .

Proof. The result is trivial for $I = \emptyset$, so assume $I \neq \emptyset$. For all $a \in A$, $z^{-1}\{a\} \subset S$. In particular, choose $a = z(s^*)$ for some $s^* \in I$. Clearly $s^* \in I \cap z^{-1}\{a\}$, and if (18) holds, then Lemma 2 implies $I \subset z^{-1}\{a\}$. Thus, $z[I] \subset \{a\}$. \square

A necessary topological condition for the situation described in Remark 8 is given in the next result. Specifically, that situation may only occur at isolated points.

Lemma 3. *Let $z : S \rightarrow A$, with A finite and $S \subset \mathbb{R}$ an interval. Let $s_1, s^*, s_2 \in S$, where $s_1 \leq s^* \leq s_2$, and suppose z is constant on $[s_1, s^*)$ and on $(s^*, s_2]$. If $z(s_1) \neq a \neq z(s_2)$, where $a := z(s^*)$, then $s^* \in \text{iso}(z^{-1}\{a\})$.*

Proof. Suppose $z(s_1) \neq a = z(s^*) \neq z(s_2)$. Then clearly $s_1 < s^* < s_2$. Choose $c_1 \in (s_1, s^*)$ and $c_2 \in (s^*, s_2)$. By assumption, $z(s) = z(s_1) \neq a$ for all $s_1 \leq s < s^*$, and $z(s) = z(s_2) \neq a$ for all $s^* < s \leq s_2$. Thus $(c_1, c_2) \cap z^{-1}\{a\} = \{s^*\}$, which implies $s^* \in \text{iso}(z^{-1}\{a\})$. \square

We now combine these results to derive sufficient conditions for a sampling functions to record a trace. Though abstract, the result is crucial to the construction of our algorithm in Section 6. Although the focus of subsequent sections is on paths, notice that the result is also applicable to trajectories.

Theorem 1. *Suppose that either $S = [0, 1]$ and $\mathcal{K} \subset \mathbb{N}$ is finite, or $S = [0, \infty)$ and $\mathcal{K} = \mathbb{N}$. Let $r : S \rightarrow X$ be continuous, and $h : X \rightarrow A$, with A finite. Define $R_a := r^{-1}h^{-1}\{a\}$ for $a \in A$, and*

$$\Gamma := \bigcup_{a \in A} \text{bd}(R_a | S), \quad \Theta := \bigcup_{a \in A} \text{iso}(R_a). \tag{19}$$

If $\sigma : \mathcal{K} \rightarrow S$ is a sampling function such that

- i) $\Theta \subset \sigma[\mathcal{K}]$,
- ii) $\forall k \in \mathcal{K}^\ominus, |\Gamma \cap [\sigma(k), \sigma(k+1)]| \leq 1$,

then $h \circ r \circ \sigma$ is a trace of $h \circ r$.

Before proceeding with the proof, we pause to interpret this theorem. Each $H_a := h^{-1}\{a\} \subset X$ is a state space region over which h is constant, and the family $H_a, a \in A$ partitions X (see Remark 2). Given the path (or trajectory) $r, R_a \subset S$ is the set of parameter values s at which $r(s) \in H_a$. The set $\Gamma \subset S$ contains only the parameter values at which r enters or leaves some H_a , and $\Theta \subset S$ contains the parameter values at which the path locally touches some H_a at only a single point. Condition i) forces a checkpoint at such isolated points, and Condition ii) ensures there can be no more than one boundary crossing between consecutive checkpoints. Since S is an interval and $R_a \subset S$, it follows that $\Theta \subset \Gamma$ by Proposition 2.

Proof. Suppose Conditions i) and ii) hold. Clauses 19.a) and 19.b) of Definition 19 follow directly from the theorem hypotheses. Below, we construct the set $S^* := \{s_k^* \in S \mid k \in \mathcal{K}^\ominus\}$, which is then shown to satisfy the remaining clauses of Definition 19. For all $k \in \mathcal{K}^\ominus$, define

$$s_k^* := \begin{cases} s, & \text{if } \Gamma \cap [\sigma(k), \sigma(k+1)] = \{s\} \\ \sigma(k), & \text{if } \Gamma \cap [\sigma(k), \sigma(k+1)] = \emptyset \end{cases}. \tag{20}$$

Condition ii) ensures the above case definition is exhaustive. Clearly $\sigma(k) \leq s_k^* \leq \sigma(k+1)$ for all $k \in \mathcal{K}^\ominus$, satisfying Clause 19.c).

Suppose, for a contradiction, that there exists $k \in \mathcal{K}^\ominus$ for which $[\sigma(k), s_k^*) \cap \Gamma \neq \emptyset$. Then there exists $s \in \Gamma$ such that

$$\sigma(k) \leq s < s_k^*. \tag{21}$$

Since $s \in [\sigma(k), s_k^*) \subset [\sigma(k), \sigma(k+1)]$ and $s \in \Gamma$, Condition ii) implies $\Gamma \cap [\sigma(k), \sigma(k+1)] = \{s\}$. Then $s = s_k^*$ by (20), which contradicts (21). Thus, $[\sigma(k), s_k^*) \cap \Gamma = \emptyset$ for all $k \in \mathcal{K}^\ominus$. Further, using (19),

$$\forall k \in \mathcal{K}^\ominus, \forall a \in A, [\sigma(k), s_k^*] \cap \text{bd}(R_a | S) = \emptyset.$$

Corollary 5 then implies $h \circ r$ is constant over $[\sigma(k), s_k^*]$, whereby

$$\forall k \in \mathcal{K}^\ominus, \forall s \in [\sigma(k), s_k^*], h \circ r(s) = h \circ r \circ \sigma(k), \tag{22}$$

which establishes Clause 19.d).

Suppose now there exists $k \in \mathcal{K}^\ominus$ for which $(s_k^*, \sigma(k+1)] \cap \Gamma \neq \emptyset$. Then there exists $s \in \Gamma$ such that

$$s_k^* < s \leq \sigma(k+1). \tag{23}$$

Since $s \in (s_k^*, \sigma(k+1)] \subset [\sigma(k), \sigma(k+1)]$ and $s \in \Gamma$, Condition ii) implies $\Gamma \cap [\sigma(k), \sigma(k+1)] = \{s\}$. Then $s = s_k^*$ by (20), which contradicts (23). Thus, $(s_k^*, \sigma(k+1)] \cap \Gamma = \emptyset$ for all $k \in \mathcal{K}^\ominus$. Moreover, using (19),

$$\forall k \in \mathcal{K}^\ominus, \forall a \in A, (s_k^*, \sigma(k+1)] \cap \text{bd}(R_a | S) = \emptyset.$$

Corollary 5 then implies $h \circ r$ is constant over $(s_k^*, \sigma(k+1)]$, therefore

$$\forall k \in \mathcal{K}^\ominus, \forall s \in (s_k^*, \sigma(k+1)], h \circ r(s) = h \circ r \circ \sigma(k+1), \tag{24}$$

which establishes Clause 19.e).

Now let $k \in \mathcal{K}^\ominus$, and suppose that both the following hold:

$$h \circ r \circ \sigma(k) \neq h \circ r(s_k^*), \tag{25}$$

$$h \circ r(s_k^*) \neq h \circ r \circ \sigma(k+1). \tag{26}$$

Together with (22) and (24), Lemma 3 then implies $s_k^* \in \text{iso}(R_a | S) \subset \Theta$, where $a = h \circ r(s_k^*)$. Condition i) further implies $s_k^* = \sigma(j)$ for some $j \in \mathcal{K}$, and (20) in turn implies $\sigma(k) \leq \sigma(j) \leq \sigma(k+1)$. However (25) and (26) imply $j \notin \{k, k+1\}$. Since there is no natural number strictly between k and $k+1$, a contradiction ensues. Therefore, for any $k \in \mathcal{K}^\ominus$, (25) and (26) cannot both be true, and Clause 19.f) has been established. \square

Remark 13. The final paragraph of the proof demonstrates that if (9) – (11) of Remark 8 hold for $z = h \circ r$, then $s_m^* \in \Theta$.

We conclude this section with a lemma that simplifies the application of Theorem 1 to an observation map, which has codomain $A = 2^{\mathcal{G}}$. The result introduces a computationally cheaper over-approximation $\hat{\Gamma}$ to the set of boundary crossing points Γ .

Lemma 4. Let $S \subset \mathbb{R}$ be an interval, $r : S \rightarrow X$ a continuous function, and $h : X \rightarrow 2^{\mathcal{G}}$, with \mathcal{G} finite. Let $\llbracket g \rrbracket := \{x \in X \mid g \in h(x)\}$ for $g \in \mathcal{G}$, and $R_a := r^{-1}h^{-1}\{a\}$ for $a \subset \mathcal{G}$. Then

$$\Gamma := \bigcup_{a \in 2^{\mathcal{G}}} \text{bd}(R_a | S) \subset \bigcup_{g \in \mathcal{G}} \text{bd}(r^{-1}\llbracket g \rrbracket | S) =: \hat{\Gamma}. \tag{27}$$

Proof. Let $\llbracket g \rrbracket^c := X \setminus \llbracket g \rrbracket$. For any $a \subset \mathcal{G}$,

$$h^{-1}\{a\} = \left(\bigcap_{g \in a} \llbracket g \rrbracket \right) \setminus \left(\bigcup_{g \in \mathcal{G} \setminus a} \llbracket g \rrbracket \right) = \left(\bigcap_{g \in a} \llbracket g \rrbracket \right) \cap \left(\bigcap_{g \in \mathcal{G} \setminus a} \llbracket g \rrbracket^c \right),$$

and therefore

$$R_a = \left(\bigcap_{g \in a} r^{-1}\llbracket g \rrbracket \right) \cap \left(\bigcap_{g \in \mathcal{G} \setminus a} (r^{-1}\llbracket g \rrbracket)^c \right), \tag{28}$$

where $r^{-1}(\cdot)^c := S \setminus r^{-1}(\cdot)$. Applying Proposition 3, and noting that the boundary of a set is equal to the boundary of its complement,

$$\begin{aligned} \text{bd}(R_a | S) &\subset \left(\bigcup_{g \in a} \text{bd}(r^{-1}\llbracket g \rrbracket | S) \right) \cup \left(\bigcup_{g \in \mathcal{G} \setminus a} \text{bd}((r^{-1}\llbracket g \rrbracket)^c | S) \right) \\ &= \bigcup_{g \in \mathcal{G}} \text{bd}(r^{-1}\llbracket g \rrbracket | S). \end{aligned}$$

Since the above hold for any $a \subset \mathcal{G}$, (27) then follows. \square

The set $\hat{\Gamma}$ in (27) is a union of $|\mathcal{G}|$ -many sets, a saving when compared with Γ , which is a union of $2^{|\mathcal{G}|}$ -many sets. The set inclusion in Lemma 4 permits us to replace Γ with $\hat{\Gamma}$ in Condition ii) of Theorem 1, without otherwise changing the result.

5. Sampling polynomial paths

5.1. Polynomial path and region assumptions

In this section, we focus on solving a special case of Problem 3 in which r is a polynomial and h only generates semi-algebraic regions of interest. Proposition 9 on path concatenation allows the results of this section to be directly extended to polynomial splines.

Assumption 1 (Polynomial paths and region boundaries). Let $M, n \geq 1$.

i) Given $\mathcal{G} := \{g_1, \dots, g_M\}$, let

$$h : \mathbb{R}^n \rightarrow 2^{\mathcal{G}}, h(x) := \{g_i \in \mathcal{G} \mid g_i(x) \leq 0\}, \tag{29}$$

where $g_i : \mathbb{R}^n \rightarrow \mathbb{R}$ is continuously differentiable for each $1 \leq i \leq M$.

ii) Suppose in addition that $g_i \in \mathbb{R}[x_1, \dots, x_n]$ for each $1 \leq i \leq M$.

iii) Let $r = \rho|_{[0,1]}$ for some $\rho \in \mathbb{R}[s]^n$.

Thus, $M \geq 1$ is the number of atomic proposition and $n \geq 1$ is the dimension of the state space. Under Clause ii), h can be made to generate any semi-algebraic region of interest by choosing M and the g_i appropriately. Note that $\rho : \mathbb{R} \rightarrow \mathbb{R}^n$ in Clause iii) is a vector function, each component of which is a univariate polynomial, and thus $r : [0, 1] \rightarrow \mathbb{R}^n$. The key property that makes an algorithmic solution possible is that every $g_i \circ \rho$ is also a univariate polynomial, having a finite number of roots if non-zero. Lemma 6 and Corollary 7 in the next subsection characterise the isolated points discussed in Remarks 8 and 13. They provide insight in a more general context that only requires the g_i to be continuously differentiable, which is the reason for separating Clauses i) and ii) in Assumption 1.

Take note of the relationship

$$\forall g_i \in \mathcal{G}, r^{-1}[[g_i]] = (g_i \circ r)^{-1}(-\infty, 0]. \tag{30}$$

The following lemma assures us that if $s \in [0, 1]$ lies on the relative boundary of some $r^{-1}[[g_i]]$, then for that i , we have $g_i \circ r(s) = 0$.

Lemma 5. Let $g : S \rightarrow \mathbb{R}$ be continuous, with $S \subset \mathbb{R}$ an interval. Then

$$\text{bd}(g^{-1}(-\infty, 0] \mid S) \subset g^{-1}\{0\}.$$

Proof. Let $G := g^{-1}(-\infty, 0]$ and $s^* \in S$. Suppose that $g(s^*) > 0$. Since g is continuous, there exists some U that is an S -neighbourhood of s^* , such that $g(s) > 0$ for all $s \in U$. Clearly $U \cap G = \emptyset$, and therefore s^* is not a boundary point of G relative to S . Similarly, if $g(s^*) < 0$, there exists some U that is an S -neighbourhood of s^* , such that $g(s) < 0$ for all $s \in U$. In this case $U \subset G$, and therefore s^* is not a boundary point of G relative to S . This proves that $g(s^*) \neq 0 \implies s^* \notin \text{bd}(G|S)$. \square

The next result implies that if $g_i \circ r$ is the zero polynomial, or has no roots in $[0, 1]$, then the relative boundary of $r^{-1}[[g_i]] \subset [0, 1]$ is empty.

Corollary 6. Let $g : S \rightarrow \mathbb{R}$ be continuous, with $S \subset \mathbb{R}$ an interval. If $\text{bd}(g^{-1}(-\infty, 0] \mid S) \neq \emptyset$, then $\{0\} \subsetneq g[S]$.

Proof. Suppose $\text{bd}(g^{-1}(-\infty, 0] \mid S) \neq \emptyset$. Then Lemma 5 implies $\{0\} \subset g[S]$. If $\{0\} = g[S]$, then $\text{bd}(g^{-1}(-\infty, 0] \mid S) = \text{bd}(S \mid S) = \emptyset$ by Proposition 1, which is a contradiction. \square

We now pause to consider polynomial trajectories, which (if non-zero) have unbounded images. This case violates Clause iii) of Assumption 1, however the result below reveals that such trajectories always have invariant suffixes when the remaining clauses hold. Therefore, according to Corollary 3, the trace of a polynomial trajectory can still be constructed from the trace of a suitable prefix, which does satisfy Assumption 1. Hence, traces of polynomial trajectories can also be computed within the current framework.




Fig. 1. Points $s_1, s_2, s_3 \in [0, 1]$ depict the three types of isolated points in (31) – (33), respectively, from each of which follow i) – iii) of Corollary 7, respectively.

Proposition 14 (Unbounded polynomial trajectories). Let Clauses i) and ii) of Assumption 1 hold, $\rho \in \mathbb{R}[t]^n$, $x = \rho \upharpoonright_{[0, \infty)}$, and $L \geq 0$. Define

$$\mathcal{C} := \{i \in \{1, \dots, M\} \mid \{0\} \subsetneq g_i \circ x[0, \infty)\},$$

$$P := \prod_{i \in \mathcal{C}} g_i \circ \rho.$$

If $P(s) \neq 0$ for all $s \geq L$, then $h \circ x_{\geq L}$ is constant.

Proof. Suppose $h \circ x_{\geq L}$ is non-constant. Then by Corollary 5,

$$\exists a \in 2^G, [L, \infty) \cap \text{bd}((h \circ x)^{-1}\{a\} \mid [0, \infty)) \neq \emptyset.$$

Noting that $(h \circ x)^{-1}\{a\} = x^{-1}h^{-1}\{a\} =: R_a$ for $a \subset G$, this implies

$$[L, \infty) \cap \bigcup_{a \in 2^G} \text{bd}(R_a \mid [0, \infty)) \neq \emptyset,$$

and Lemma 4 in turn implies

$$[L, \infty) \cap \bigcup_{g \in G} \text{bd}(x^{-1}[[g]] \mid [0, \infty)) \neq \emptyset.$$

Recalling that $x^{-1}[[g_i]] = (g_i \circ x)^{-1}(-\infty, 0]$ for every $g_i \in G$,

$$[L, \infty) \cap \bigcup_{i=1}^M \text{bd}((g_i \circ x)^{-1}(-\infty, 0] \mid [0, \infty)) \neq \emptyset.$$

Lemma 5 and Corollary 6 then imply

$$\exists s \geq L, \exists i \in \mathcal{C}, g_i \circ x(s) = 0.$$

Thus $P(s) = 0$ for some $s \geq L$, which establishes the contrapositive of the result. \square

Remark 14. To obtain an invariant $x_{\geq L}$ in Proposition 14, L can be any strict upper bound on the real roots of P . Various upper bounds for polynomial roots are provided in [28, Theorem 4.2].

5.2. Isolated point characterisation

The next result characterises the isolated points in Θ in terms of the relative boundaries of the $r^{-1}[[g_i]]$, and the derivatives $(g_i \circ r)'$. A more intuitive restatement is then provided in Corollary 7, along with an illustration in Fig. 1.

Lemma 6. Suppose Clause i) of Assumption 1 holds. Let $r : [0, 1] \rightarrow \mathbb{R}^n$ also be continuously differentiable, and $[[g_i]] := \{x \in \mathbb{R}^n \mid g_i(x) \leq 0\}$. If

$$s \in \Theta := \bigcup_{a \in 2^G} \text{iso}(R_a),$$

where $R_a := r^{-1}h^{-1}\{a\}$ for $a \subset G$, then at least one of the following hold:

$$\exists i \in \{1, \dots, M\}, s \in \text{bd}(r^{-1}[\![g_i]\!] \mid [0, 1]) \cap \{0, 1\} \tag{31}$$

$$\exists i, j \in \{1, \dots, M\}, j \neq i, s \in \text{bd}(r^{-1}[\![g_i]\!] \mid [0, 1]) \cap \text{bd}(r^{-1}[\![g_j]\!] \mid [0, 1]) \tag{32}$$

$$\exists i \in \{1, \dots, M\}, s \in \text{bd}(r^{-1}[\![g_i]\!] \mid [0, 1]) \wedge (g_i \circ r)'(s) = 0. \tag{33}$$

Proof. Note that $S := [0, 1]$ has no isolated points, and $R_a \subset S$ for every $a \in G$. Then recalling Remark 1, Proposition 2 implies $\text{iso}(R_a) \subset \text{bd}(R_a \mid S)$, and therefore $\Theta \subset \Gamma := \bigcup_{a \in G} \text{bd}(R_a \mid S)$.

Let $s \in \Theta \subset \Gamma$. Then Lemma 4 implies there exists $i \in \{1, \dots, M\}$ such that

$$s \in \text{bd}(r^{-1}[\![g_i]\!] \mid S) = \text{bd}((g_i \circ r)^{-1}(-\infty, 0] \mid S). \tag{34}$$

This in turn implies $g_i \circ r(s) = 0$ by Lemma 5, and therefore

$$s \in (g_i \circ r)^{-1}(-\infty, 0] = r^{-1}[\![g_i]\!]. \tag{35}$$

To obtain a contradiction, suppose (31) – (33) are all false. Since (34) is true and (31) false, $s \in (0, 1) = \text{int}(S)$. Furthermore, (33) is false, and therefore $(g_i \circ r)'(s) \neq 0$. This implies s is not a local minimum of $g_i \circ r$, so for any S -neighbourhood V of s , there exists $\hat{s} \in V$ such that $g_i \circ r(\hat{s}) \leq g_i \circ r(s) = 0$. Therefore

$$s \notin \text{iso}((g_i \circ r)^{-1}(-\infty, 0]) = \text{iso}(r^{-1}[\![g_i]\!]). \tag{36}$$

Since $s \in \Theta$, there exists $a \in G$ such that

$$s \in \text{iso}(R_a) \subset R_a. \tag{37}$$

Now recall (28), which is restated below:

$$R_a = \left(\bigcap_{g \in a} r^{-1}[\![g]\!] \right) \cap \left(\bigcap_{g \in G \setminus a} (r^{-1}[\![g]\!])^c \right).$$

From (35) it follows that $g_i \in a$, because if $g_i \notin a$, then $s \in (r^{-1}[\![g_i]\!])^c$. The above equation can therefore be re-written as

$$R_a = r^{-1}[\![g_i]\!] \cap G_a, \tag{38}$$

where

$$G_a := \left(\bigcap_{g \in a \setminus \{g_i\}} r^{-1}[\![g]\!] \right) \cap \left(\bigcap_{g \in G \setminus a} (r^{-1}[\![g]\!])^c \right).$$

Given (38), Lemma 1 implies

$$\text{iso}(R_a) \setminus \text{bd}(G_a \mid S) \subset \text{iso}(r^{-1}[\![g_i]\!]), \tag{39}$$

and applying Proposition 3 to the definition of G_a ,

$$\begin{aligned} \text{bd}(G_a \mid S) &\subset \left(\bigcup_{g \in a \setminus \{g_i\}} \text{bd}(r^{-1}[\![g]\!] \mid S) \right) \cup \left(\bigcup_{g \in G \setminus a} \text{bd}((r^{-1}[\![g]\!])^c \mid S) \right) \\ &= \bigcup_{g \in G \setminus \{g_i\}} \text{bd}(r^{-1}[\![g]\!] \mid S). \end{aligned}$$

Since (32) is also false, $s \notin \text{bd}(r^{-1}[\![g_j]\!] \mid S)$ for every $j \neq i$. Then clearly $s \notin \text{bd}(G_a \mid S)$, and together with (37), (39) implies $s \in \text{iso}(r^{-1}[\![g_i]\!])$, which contradicts (36). \square

Corollary 7 (Necessary conditions for isolated points). Under the hypotheses of Lemma 6, let

$$C := \{i \in \{1, \dots, M\} \mid \{0\} \subsetneq g_i \circ r[[0, 1]]\}.$$

If $s \in \Theta$ then at least one of the following hold:

- i) $s \in \{0, 1\}$ and there exists $i \in C$ such that $g_i \circ r(s) = 0$
- ii) there exist $i, j \in C$ with $i \neq j$ such that $g_i \circ r(s) = g_j \circ r(s) = 0$
- iii) there exists $i \in C$ such that $g_i \circ r(s) = (g_i \circ r)'(s) = 0$.

Proof. Applying (30) and Lemma 5 to (31) – (33) yields i) – iii), respectively, but for $1 \leq i, j \leq M$. Corollary 6 additionally confirms that $i, j \in \mathcal{C}$. \square

Remark 15. Consider the case in Remark 8 for $z = h \circ r$. Equations (9) – (11) imply $s_m^* \in (0, 1)$, and Remark 13 implies $s_m^* \in \Theta$. Corollary 7 now makes it clear that, for $s = s_m^*$, ii) and iii) are the only possible conditions under which (9) – (11) occur. Specifically, s_m^* is a root of two distinct non-zero $g_i \circ r$, or a shared root of some non-zero $g_i \circ r$ with its derivative. Thus, potential isolated points can be identified by testing for the presence of double roots.

In view of Fig. 1, we refer to points satisfying (32) and (33) as *double crossing* and *bounce* points, respectively.

5.3. Algebraic conditions on sampling functions

We are now ready to state the main result, which specialises Theorem 1 to the setting described by Assumption 1.

Theorem 2. Suppose Assumption 1 holds. Define

$$\mathcal{C} := \{i \in \{1, \dots, M\} \mid \{0\} \subsetneq g_i \circ r[[0, 1]]\}, \tag{40}$$

let $P \in \mathbb{R}[s]$ be a non-zero polynomial such that

$$\prod_{i \in \mathcal{C}} g_i \circ \rho \mid P, \tag{41}$$

and $V := \gcd(P, P')$. If $\sigma : \mathcal{K} \rightarrow [0, 1]$ is a finite sampling function such that

$$V^{-1}\{0\} \cap [0, 1] \subset \sigma[\mathcal{K}], \tag{42}$$

$$\forall k \in \mathcal{K}^\ominus, \left| P^{-1}\{0\} \cap [\sigma(k), \sigma(k+1)] \right| \leq 1, \tag{43}$$

then $h \circ r \circ \sigma$ is a trace of $h \circ r$.

Once again, we pause for interpretation before proceeding with the proof. Recall that r is the restriction of the vector polynomial ρ to $[0, 1]$. Equation (40) can therefore be expanded to

$$\mathcal{C} := \{i \in \{1, \dots, M\} \mid g_i \circ \rho \neq 0 \wedge \exists s \in [0, 1], g_i \circ \rho(s) = 0\}. \tag{44}$$

Thus \mathcal{C} selects all non-zero $g_i \circ \rho$ with roots in $[0, 1]$ to be factors of P in (41). Referring to Theorem 1, all boundary points in Γ are roots of P . Similarly, any isolated points in Θ are repeated roots of P , and therefore roots of V . Conditions (42) – (43) then ensure the conditions of Theorem 1 are satisfied. The next section considers how to construct a sampling function that satisfies (42) – (43).

Proof. Let $S := [0, 1]$, $R_a := r^{-1}h^{-1}\{a\}$ for $a \in \mathbb{G}$, and define Γ and Θ as in (19). First note that each $g_i \circ \rho \in \mathbb{R}[s]$, because $\rho \in \mathbb{R}[s]^n$ and $g_i \in \mathbb{R}[x_1, \dots, x_n]$ by Assumption 1. The following chain of equalities and inclusions is now established:

$$\Gamma \subset \bigcup_{g \in \mathbb{G}} \text{bd}(r^{-1}[[g]] \mid S) = \bigcup_{i=1}^M \text{bd}(g_i \circ r^{-1}(-\infty, 0] \mid S) = \bigcup_{i \in \mathcal{C}} \text{bd}(g_i \circ r^{-1}(-\infty, 0] \mid S) \tag{45}$$

$$\subset \bigcup_{i \in \mathcal{C}} g_i \circ r^{-1}\{0\} \subset \bigcup_{i \in \mathcal{C}} (g_i \circ \rho)^{-1}\{0\} \subset P^{-1}\{0\}. \tag{46}$$

The first inclusion of (45) follows from Lemma 4, and (30) establishes the subsequent equality. For any $1 \leq i \leq M$, if $i \notin \mathcal{C}$, the contrapositive of Corollary 6 implies $\text{bd}(g_i \circ r^{-1}(-\infty, 0] \mid S) = \emptyset$, which in turn establishes the second equality of (45). The first inclusion of (46) follows from Lemma 5, and the second holds because r is a restriction of ρ . If $g_i \circ \rho(s) = 0$ for $i \in \mathcal{C}$ and $s \in \mathbb{R}$, then (41) implies $P(s) = 0$, which establishes the final inclusion. If (43) is satisfied, then (45) – (46) imply Condition ii) of Theorem 1 is also satisfied. Also since $g_i \circ \rho$ is non-zero for all $i \in \mathcal{C}$, P can indeed be non-zero, as assumed.

Now suppose $s^* \in \Theta \subset [0, 1]$. Then at least one of i) – iii) of Corollary 7 hold true. If i), then $s^* \in \{0, 1\} \subset \sigma[\mathcal{K}]$ by Definition 17. Consider now the latter two points. If ii) holds, then $g_i \circ \rho(s^*) = g_j \circ \rho(s^*) = 0$ for some distinct $i, j \in \mathcal{C}$, and it follows from (41) that $(s - s^*)^2 \mid P$. If iii), then $g_i \circ \rho(s^*) = (g_i \circ \rho)'(s^*) = 0$ for some $i \in \mathcal{C}$, which implies $(s - s^*)^2 \mid g_i \circ \rho$ by Proposition 5, which in turn implies $(s - s^*)^2 \mid P$. In either case then, $(s - s^*)^2 \mid P$. Now $Q := P/V$ is square-free by Proposition 4, and $P = QV$. If $(s - s^*) \nmid V$, then $(s - s^*)^2 \nmid Q$, which is a contradiction. Thus $V(s^*) = 0$, and $s^* \in V^{-1}\{0\} \cap [0, 1]$. If (42) holds, then $s^* \in \sigma(\mathcal{K})$. It has been shown that $\Theta \subset \sigma[\mathcal{K}]$, satisfying Condition i) of Theorem 1. \square

Remark 16. Condition (41) permits some flexibility in the choice of P , which is useful for the algorithm designer. This issue is discussed further in Section 6.5, but it is worth noting that there is no clear choice of P that yields optimal computational complexity for all problem instances.

Remark 17. Throughout this section, it has not been necessary to explicitly assume finite variability in any of the results. In fact, Assumption 1 guarantees that $h \circ x$ is of finite variability, because all univariate polynomials have a finite number of roots.

6. Algorithmic implementation

6.1. Overview of sampling strategy

According to Theorem 2, we can compute a trace of $h \circ r$ by constructing a sampling function that places checkpoints

- 1) on either side of every root of P in $[0, 1]$ to satisfy (43); and
- 2) exactly at every root of V in $[0, 1]$ to satisfy (42).

Since P and V are univariate polynomials, we can exploit some concepts and subroutines from Computer Algebra to do this algorithmically. In particular, we require a root existence test and a root isolation algorithm.

Before proceeding with a discussion of these subroutines, recall that any polynomial $p \in \mathbb{Q}[s]$ with rational coefficients can be scaled to form a polynomial $q \in \mathbb{Z}[s]$ with integer coefficients, without changing its roots. Specifically, let $\text{RATIONALTOINTEGER}(p) = q$ multiply p by the product of the denominators of its coefficients. Its arithmetic complexity is $O(\deg(p)^2)$.

6.2. Root existence test

The result below, based on Descartes’ Rule of Signs, provides a computationally robust test to determine whether a univariate polynomial has any real roots in an open interval.

Proposition 15. Let $p \in \mathbb{R}[s]$ be non-zero, $d := \deg(p)$, and

$$q(s) := (s + 1)^d p\left(\frac{as + b}{s + 1}\right), \tag{47}$$

where $a < b$. All the non-zero coefficients of q have the same sign iff p has no roots in the interval (a, b) .

Proof. As noted in [36, Section 3], the Möbius transformation $(s \mapsto \frac{as+b}{s+1})$ is a bijection from $(0, \infty)$ to (a, b) . Furthermore, $q \in \mathbb{R}[s]$. Thus for any $s > 0$, $(q(s) = 0 \iff p(\hat{s}) = 0)$, where $\hat{s} := \frac{as+b}{s+1} \in (a, b)$. Let $\text{var}(q)$ denote the number of sign changes in the sequence of coefficients of q [36]. If $\text{var}(q) = 0$, then q has no roots in $(0, \infty)$ by [37, Theorem 4], which implies p has no roots in (a, b) . Conversely, if p has no roots in (a, b) , then q has no roots in $(0, \infty)$, which implies $\text{var}(q) = 0$ by [37, Theorem 5]. \square

Suppose $p \in \mathbb{R}[s]$ is given by $p(s) = \sum_{i=0}^d p_i s^i \in \mathbb{R}[s]$, and consider the following operators of type $\mathbb{R}[s] \rightarrow \mathbb{R}[s]$ for $s, \lambda \in \mathbb{R}$:

$$(\mathbf{R}p)(s) := \sum_{i=0}^d p_{d-i} s^i = \begin{cases} s^d p(s^{-1}), & \text{if } s \neq 0 \\ p_d, & \text{if } s = 0 \end{cases},$$

$$(\mathbf{T}_\lambda p)(s) := \sum_{i=0}^d \left(\sum_{j=i}^d \binom{j}{i} p_j \lambda^{j-i} s^i \right) = p(s + \lambda),$$

$$(\mathbf{C}_\lambda p)(s) := \sum_{i=0}^d p_i \lambda^i s^i = p(\lambda s).$$

The transformed polynomial q in (47) can then be computed according to $q = \mathbf{T}_1 \mathbf{R} \mathbf{C}_{b-a} \mathbf{T}_a p$. Operator \mathbf{R} is the *reciprocal transformation* [38, Definition 4, Remark 6], which simply reverses the order of the coefficients. The *Taylor shift* [30, Chapter 4.1] operation \mathbf{T}_λ can be performed in $O(d^2)$ arithmetic operations [30, Theorem 4.3].

Remark 18. The existence of roots of $p \in \mathbb{Q}[s]$ in the interval (a, b) can be determined in $O(\deg(p)^2)$ arithmetic operations. This can obviously be extended to the closed interval by evaluating $p(a)$ and $p(b)$.

6.3. Root isolation subroutine

Given a polynomial $p \in \mathbb{R}[s]$ with root $s^* \in \mathbb{R}$, a bounded interval is an *isolating interval* for s^* iff it is either open or a singleton, contains s^* , and contains no other roots of p . An isolating interval for s^* is a *strict isolating interval* iff it is open and its closure contains no other roots of p .

We focus below on constructing isolating intervals for the roots of polynomials in $(0, 1)$. For computational reasons, attention is restricted to polynomials with rational coefficients and intervals with rational endpoints.

Definition 21 (*Root isolation algorithm*). Given $p \in \mathbb{Q}[s]$, a root isolation algorithm returns a set of intervals

$$\mathcal{J} := \{I_j \subset (0, 1) \mid 1 \leq j \leq J\}, \tag{48}$$

of cardinality $J \leq \deg(p)$, such that all of the following hold:

- 21.a) each I_j is either open or a singleton
- 21.b) each I_j has rational endpoints
- 21.c) $\sup I_j \leq \inf I_{j+1}$ for all $1 \leq j < J$
- 21.d) every root of p in $(0, 1)$ is contained in some I_j
- 21.e) every I_j contains exactly one root of p .

The output (48) is represented in pseudocode by $I_{1:J} = \text{ROOTISOLATION}(p)$. Clearly each $I_j \in \mathcal{J}$ is an isolating interval for some root of p . Letting $a_j := \inf I_j$ and $b_j := \sup I_j$ for $1 \leq j \leq J$, the interval endpoints satisfy

$$0 \leq a_1 \leq b_1 \leq a_2 \leq b_2 \leq \dots \leq a_J \leq b_J \leq 1.$$

The Vincent, Collins and Akritas (VCA) Algorithm [39], also known as the Modified Uspensky or Descartes method, is a well-studied root isolation algorithm that performs a bisection search using Descartes Rule of Signs. Its output complies with Definition 21, and the algorithm is presented concisely in [36, Algorithm 1].

Remark 19. If $p \in \mathbb{Z}[s]$ is square-free, then the bit complexity of the VCA algorithm is $O(d^5(L + \log d)^2)$, where $d = \deg(p)$ and $\|p\|_\infty < 2^L$ [40, Theorem 4.1].

Our strategy is to use a root isolation algorithm to place a checkpoint on either side of the roots of P . However, Definition 21 permits roots at the boundaries of the isolating intervals, which leads to problems. In order to guarantee (43), all the isolating intervals must be strict. We therefore adopt the more stringent definition below. Note that singleton intervals are not permitted.

Definition 22 (*Strict root isolation algorithm*). Given any $p \in \mathbb{Q}[s]$, a strict root isolation algorithm returns a set of intervals

$$\mathcal{J}^* := \{(a_j, b_j) \subset (0, 1) \mid 1 \leq j \leq J\},$$

of cardinality $J \leq \deg(p)$, such that all of the following hold:

- 22.a) $a_j, b_j \in \mathbb{Q}$ for all $1 \leq j \leq J$
- 22.b) $0 \leq a_1 < b_1 \leq a_2 < b_2 \leq \dots \leq a_J < b_J \leq 1$
- 22.c) every root of p in $(0, 1)$ is contained in some interval (a_j, b_j)
- 22.d) every interval (a_j, b_j) contains exactly one root of p
- 22.e) $p(a_1) \neq 0 \neq p(b_J)$.

Remark 20. It follows from Definition 22 that, for all $1 \leq j \leq J$, the closed interval $[a_j, b_j]$ contains exactly one root of p . Also observe there is at most one root of p in $[0, a_1]$, which can only occur if $p(0) = 0$. Similarly, there is at most one root of p in $[b_J, 1]$, which only occurs when $p(1) = 0$. The proof of Theorem 3 in Section 6.4 relies on these properties to establish (43).

Using lower bounds on polynomial root separation [41], the output of a root isolation algorithm is readily modified into the output of a strict root isolation algorithm. Algorithm 2, STRICTISOLATINGINTERVALS, offers a procedure for doing this. Since this procedure is somewhat distracting, and motivated by aspects of the main algorithm yet to be presented, the details are deferred to Algorithm 2 in Appendix A.

6.4. POLYTRACE algorithm

Recall the two-step procedure outlined in Section 6.1. Step 1) can be achieved by means of a root isolation algorithm. Step 2) is problematic, because in general the roots of V cannot be computed exactly. And even if they could, the roots may not be rational. Approximate roots will not do, because the roots of V extract isolated points, which by definition capture information absent from their neighbouring points. The following result demonstrates that, even if an isolated point $s^* \in [0, 1]$ is not known precisely, the set $h \circ r(s^*)$ needed to construct a trace can still be evaluated by means of a root existence test. Only a strict isolating interval for s^* is required.

Lemma 7 (Isolation Lemma). *Under the Assumptions of Theorem 2, suppose there exist $0 \leq a < s^* < b \leq 1$ such that*

$$\{s \in [a, b] \mid P(s) = 0\} = \{s^*\}. \tag{49}$$

Then $h \circ r(s^*) = h \circ r(a) \cup \{g_i \in G \mid \exists \hat{s} \in (a, b), g_i \circ \rho(\hat{s}) = 0\}$.

The Isolation Lemma tells us that, given a strict isolating interval (a, b) for the root $s^* \in [0, 1]$ of P , the set $h \circ r(s^*)$ can be computed by evaluating $h \circ r(a)$, and then adding to it all g_i for which $g_i \circ \rho$ has a root in (a, b) .

Proof. Define $E := \{g_i \in G \mid \exists \hat{s} \in (a, b), g_i \circ \rho(\hat{s}) = 0\}$. It is first shown that (i) $E \subset h \circ r(s^*)$, then that (ii) $h \circ r(a) \cup E \subset h \circ r(s^*)$, and finally that (iii) $h \circ r(s^*) \subset h \circ r(a) \cup E$.

Observe that

$$\forall s \in [0, 1], h \circ r(s) = h \circ \rho(s) = \{g_i \in G \mid g_i \circ \rho(s) \leq 0\}. \tag{50}$$

(i) Suppose $g_i \in E$. Then $g_i \circ \rho(\hat{s}) = 0$ for some $\hat{s} \in (a, b)$. Recall (44), restated below:

$$C := \{i \in \{1, \dots, M\} \mid g_i \circ \rho \neq 0 \wedge \exists s \in [0, 1], g_i \circ \rho(s) = 0\}. \tag{44}$$

If $i \notin C$, then (44) implies $g_i \circ \rho = 0$ and therefore $g_i \circ \rho(s^*) = 0$, which in turn implies $g_i \in h \circ r(s^*)$ by (50). Suppose instead that $i \in C$. Then $g_i \circ \rho \mid P$ by (41), and furthermore $P(\hat{s}) = 0$. The assumption (49) then implies $\hat{s} = s^*$, by which $g_i \circ \rho(s^*) = 0$ and therefore $g_i \in h \circ r(s^*)$. Thus $E \subset h \circ r(s^*)$.

(ii) Suppose now that $g_i \in h \circ r(a)$, which implies

$$g_i \circ \rho(a) \leq 0. \tag{51}$$

If $i \notin C$, then by (44) either $g_i \circ \rho = 0$, which implies $g_i \circ \rho(s^*) = 0$, or $g_i \circ \rho$ has no roots in $[0, 1]$, which implies $g_i \circ \rho(a) < 0$, and moreover, that $g_i \circ \rho(s^*) < 0$ due to continuity. In either case, $g_i \circ \rho(s^*) \leq 0$, hence $g_i \in h \circ r(s^*)$. Suppose instead that $i \in C$. Then $g_i \circ \rho \mid P$ by (41), and since $s^* \in (a, b)$, (49) implies $P(a) \neq 0$. Since $g_i \circ \rho \mid P$, it follows that $g_i \circ \rho(a) \neq 0$, and therefore $g_i \circ \rho(a) < 0$ by (51). If $g_i \circ \rho(s^*) \leq 0$, then $g_i \in h \circ r(s^*)$ by (50). If $g_i \circ \rho(s^*) > 0$, then by the Intermediate Value Theorem there exists $\hat{s} \in (a, s^*) \subset (a, b)$ such that $g_i \circ \rho(\hat{s}) = 0$, and therefore $g_i \in E \subset h \circ r(s^*)$. Thus $h \circ r(a) \subset h \circ r(s^*)$, which establishes that $h \circ r(a) \cup E \subset h \circ r(s^*)$.

(iii) Finally, suppose $g_i \in h \circ r(s^*)$, which implies $g_i \circ \rho(s^*) \leq 0$. If $g_i \circ \rho(s^*) = 0$, then $g_i \in E$. Suppose now $g_i \circ \rho(s^*) < 0$. If $g_i \circ \rho(a) \leq 0$, then $g_i \in h \circ r(a)$. Alternatively if $g_i \circ \rho(a) > 0$, then by the Intermediate Value Theorem there exists $\hat{s} \in (a, s^*) \subset (a, b)$ such that $g_i \circ \rho(\hat{s}) = 0$, and therefore $g_i \in E$. Thus $h \circ r(s^*) \subset h \circ r(a) \cup E$. \square

Lemma 7 is exploited in Lines 21 – 24 of Algorithm 1, the correctness of which is now established.

Theorem 3. *Suppose Assumption 1 is satisfied. In addition, let $G = \{1, \dots, M\}$, $\rho \in \mathbb{Q}[s]^n$, and $g_i \in \mathbb{Q}[x_1, \dots, x_n]$ for all $1 \leq i \leq M$. Then the output $\zeta = \text{POLYTRACE}(\rho, g_{1:M})$ of Algorithm 1 is a trace of $h \circ r$.*

Proof. First note that a checkpoint is placed at 0 and 1 by virtue of Lines 2 and 32, to satisfy Definition 17.

Lines 3 – 9 form

$$P := \prod_{i \in C} g_i \circ \rho, \tag{52}$$

where C is given by (44). This definition clearly satisfies (41), and V is then computed in Line 10 as prescribed by Theorem 2. The square-free part Q of P is computed in Line 11. Strict isolating intervals for the roots of Q in $(0, 1)$ are computed in Lines 12 – 16. By Corollary 1, these are also strict isolating intervals for the roots of P in $(0, 1)$, because P and Q share exactly the same roots. Checkpoints are placed at the start a_j of each isolating interval in Line 19, and at the end of the final isolating interval b_j in Line 27. Any remaining roots of P in $[0, 1]$ are included in the checkpoints at 0 and 1. As noted in Remark 20, there is exactly one root of P in $[a_j, a_{j+1}]$ for any $1 \leq j < J$. Furthermore, there is at most one root in $[0, a_1]$ and in $[b_J, 1]$. Thus, (43) is satisfied.

Algorithm 1 Computes a trace of $h \circ \rho \upharpoonright_{[0,1]}$, where $\rho \in \mathbb{Q}[s]^n$, $g_1 \dots g_M \in \mathbb{Q}[x_1, \dots, x_n]$, and $h(x) := \{i \in \{1, \dots, M\} \mid g_i(x) \leq 0\}$.

```

1: function POLYTRACE( $\rho, g_{1:M}$ )
2:    $\zeta(0) := h \circ \rho(0)$ 
3:    $P := 1$ 
4:   for  $1 \leq i \leq M$  do
5:      $p_i := \text{RATIONALTOINTEGER}(g_i \circ \rho)$ 
6:     if  $p_i$  has roots in  $[0,1]$  and  $\deg(p_i) > 0$  then
7:        $P \leftarrow P \times p_i$ 
8:     end if
9:   end for
10:   $V := \text{GCD}(P, P')$ 
11:   $Q := P/V$ 
12:   $I_{1:J} := \text{ROOTISOLATION}(Q)$ 
13:  for  $1 \leq j \leq J$  do
14:     $E_j := \{i \in \{1, \dots, M\} \mid p_i \text{ has roots in } I_j\}$ 
15:  end for
16:   $(a_{1:J}, b_{1:J}) := \text{STRICTISOLATINGINTERVALS}(p_{1:M}, I_{1:J}, E_{1:J})$ 
17:   $k := 1$ 
18:  for  $1 \leq j \leq J$  do
19:     $\zeta(k) := h \circ \rho(a_j)$ 
20:     $k \leftarrow k + 1$ 
21:    if  $\deg(V) > 0$  then
22:       $\zeta(k) := \zeta(k-1) \cup E_j$ 
23:       $k \leftarrow k + 1$ 
24:    end if
25:  end for
26:  if  $J > 0$  then
27:     $\zeta(k) := h \circ \rho(b_j)$ 
28:  else
29:     $\zeta(k) := h \circ \rho(\frac{1}{2})$ 
30:  end if
31:   $k \leftarrow k + 1$ 
32:   $\zeta(k) := h \circ \rho(1)$ 
33:  return  $\zeta$ 
34: end function

```

Line 21 tests for potential isolated points by checking the degree of V . If $V = 1$, then V has no roots in $[0, 1]$. Otherwise, $\deg(V) > 1$, and Lines 21 – 24 exploit Lemma 7 to evaluate $h \circ r$ at the unique root of P within every strict isolating interval. Since $P = QV$, this includes all roots of V in $(0, 1)$. Any additional roots of V in $[0, 1]$ are included in the checkpoints at 0 and 1. Thus, (42) is satisfied.

Finally, if no $g_i \circ \rho$ has roots in $(0, 1)$, the empty set is returned in Line 12, implying that $J = 0$. In this case, $\sigma(0) = 0$ and $\sigma(2) = 1$. It is still possible that P has roots at both endpoints $\{0, 1\}$. Therefore, a checkpoint $\sigma(1)$ is placed at their midpoint in Line 29, to ensure there is no more than one root in $[\sigma(k), \sigma(k+1)]$ for each $k \in \{0, 1\}$, as required by (43). □

6.5. Computational complexity of POLYTRACE

The complexity of Algorithm 1 depends heavily on the subroutines chosen for the individual steps. In this section, it is assumed that for every $1 \leq i \leq M$, g_i has no more than N non-zero coefficients and $\deg(p_i) \leq D$, where each $p_i \in \mathbb{Z}[s]$ is formed in Line 5 of Algorithm 1. To understand how N and D can be derived from upper bounds on $\deg(\rho)$, $\deg(g_i)$, and the state space dimension n , refer to Appendix B.

- *Formation of P* : forming each $g_i \circ \rho$ requires $O(ND^2)$ arithmetic operations (see Appendix B). Each root existence test in Line 6 requires $O(D^2)$ arithmetic operations. These are both done M times. Finally, $\deg(P) \leq MD$, and therefore forming P requires $O(M^2D^2)$ arithmetic operations [30, Lemma 3.15]. Overall, Lines 3 – 9 require $O(MD^2(M+N))$ arithmetic operations.
- *G.c.d computation*: [42, Algorithm 10.1], which has $O(M^2D^2)$ arithmetic complexity, performs Lines 10–11 and returns both Q and V with integer coefficients. Implementations with polynomial bit complexity exist [42, Remark 10.19]. See [18, Chapters 3, 6, 11] for a detailed discussion of the available algorithms for g.c.d computation, and their complexities.
- *Root isolation*: this step has bit complexity $O(M^5D^5(L + \log MD)^2)$ for $Q \in \mathbb{Z}[s]$ by Remark 19, where $L \in \mathbb{N}$ is any constant such that $\|Q\|_\infty \leq 2^L$.
- *Strict isolating intervals*: the arithmetic complexity of Algorithm 2 is $O(MD)$.
- *Computing E_1, \dots, E_J* : this requires at most MD root existence tests for polynomials of degree D or less, yielding an overall arithmetic complexity of $O(MD^3)$ for Lines 13 – 15.

While a more detailed analysis of the growth in polynomial coefficients sizes is required to obtain a precise estimate, the bit complexity of POLYTRACE is clearly polynomial. The root isolation step appears likely to dominate.

Now as noted in Remark 16, the choice P affects the complexity of the algorithm. Given the above complexity estimate of the root isolation step, omitting unnecessary factors from P in (41) potentially offers significant improvements. The definition (52) adopted by Algorithm 1 constructs the lowest degree polynomial that satisfies (41). However, this requires the root existence test in Line 6, which adds $O(MD^2)$ complexity. An alternative definition

$$P := \prod_{\deg(g_i \circ \rho) > 0} g_i \circ \rho$$

dispenses with the root existence test, but makes no attempt to reduce $\deg(P)$. Both approaches have the same worst-case asymptotic complexity, because it may well be that all $g_i \circ \rho$ have roots in $[0, 1]$, making the test redundant. If it is known *a priori* that only a few boundaries are likely to be crossed within a given segment, we expect that (52) leads to a net reduction in complexity.

7. Numerical examples

7.1. Randomly generated examples

Fig. 2 depicts a randomly generated instance of Problem 3 involving a two-dimensional polynomial spline path, and atomic propositions that are mapped to either ellipses or half-spaces. Algorithm 1 has been implemented in Wolfram Mathematica 12.1, taking advantage of the native `ROOTINTERVALS` root isolation function. The output of `POLYTRACE` for each segment is concatenated and included in Fig. 2, after a post-processing step that removes consecutive repetitions from the trace. Region 8 has been constructed to induce a bounce point between the first and second waypoints, and Regions 6 and 7 to induce a double crossing point between the second and third waypoints. Equipped with this knowledge, the reader is invited to visually confirm the trace produced.

The total computation time for the trace in Fig. 2 is 93.75 ms. Compare this with a computation time of 1.983 s for the scenario in Fig. 3, which retains the same path while increasing the number of atomic propositions sevenfold.

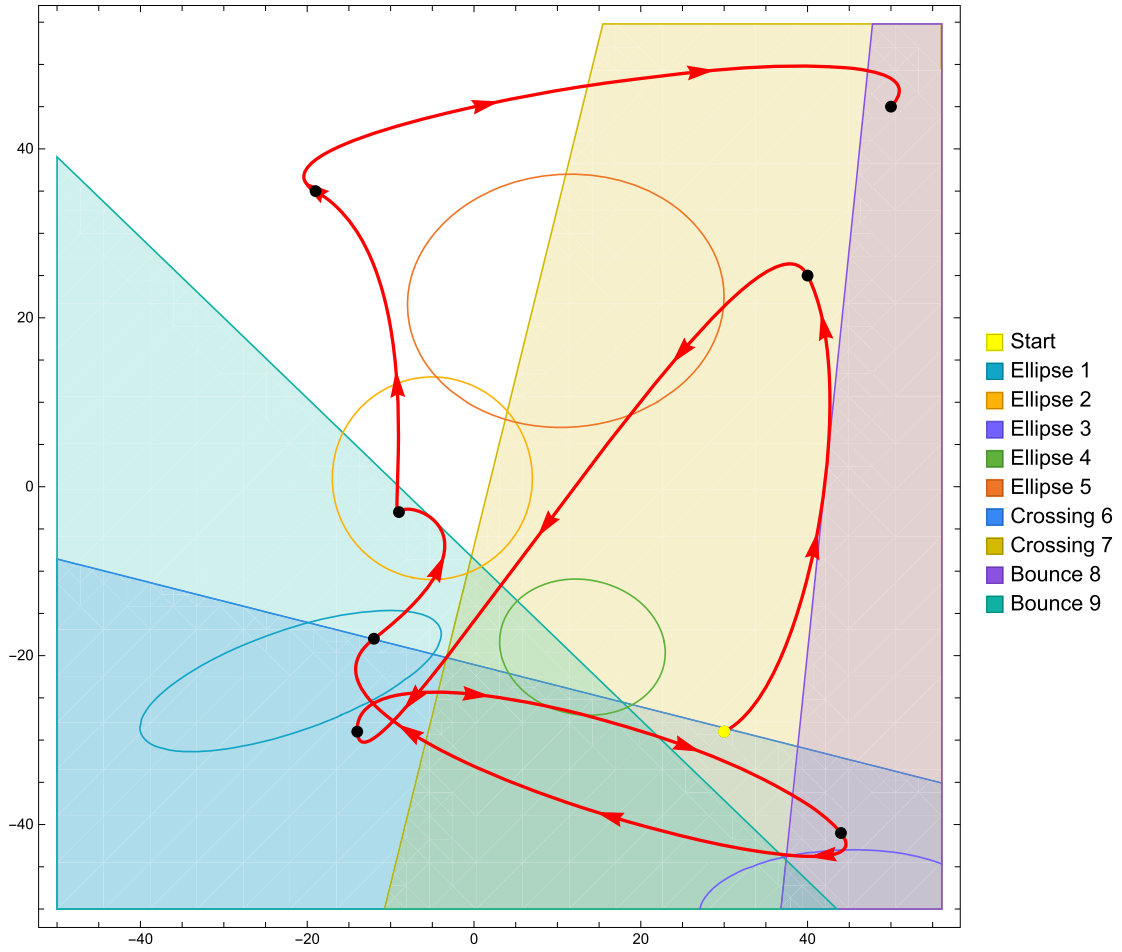
7.2. Case study: verification of a robot motion plan

Consider a mobile robot on a factory floor, tasked with fetching a hazardous chemical compound and returning to a target zone, without colliding with obstacles. The chemical is stored in a 0.5 m diameter drum, and the robot also has a circular perimeter, 1.5 m in diameter. The drum attaches to the side of the robot, via a latching system. Both robot and drum must remain within a predefined 20 m x 20 m safe zone, which is guaranteed to be clear of moving obstacles throughout the operation. The drum, and any static obstacles, are detected using a system of cameras, and are represented within the robot’s map by enclosing ellipses. The surface of the robot must make contact with the surface of the drum in order to trigger the latching system. The unstable nature of the chemical compound requires contact to be made gently, and this imposes precise requirements on the robot trajectory. Here, we encode these requirements in LTL_{\diamond} , and use `POLYTRACE` to certify a cubic spline proposed by a path planner for the robot. The subsequent path tracking stage (see Remark 5) is guaranteed to produce only direct motions. Since this is a safety critical mission, the resulting motion plan must be verified independently, prior to execution. Here, as explained in Remark 12, the verification of the motion plan depends only on the path proposed by the path planner.

Fig. 4 depicts the path superimposed on the robot’s map. Obstacles are shown in solid black, the drum in red, and the safe zone corresponds exactly to the plot axes limits. We model the robot as a disk. Since the path is to be tracked by the centre of the robot, collisions can be avoided by expanding the obstacles in the map by the radius of the robot. These expanded obstacles are represented by Ovals 4, 6 and 8. In general, the expanded obstacles are no longer ellipses, but are instead defined by the 8th-order polynomial in [43, Example 4]. After the robot has latched onto the drum, collision avoidance for their composite shape can be achieved by further expanding the obstacles by the diameter of the drum, yielding Ovals 5, 7 and 9. Similarly, Halfspaces 10, 12, 14 and 16 enclose the safe zone contracted by the robot radius, and Halfspaces 11, 13, 15 and 17 contract it again by the drum diameter. Oval 2 expands the drum by the robot radius, so that the robot touches the drum iff its centre touches Oval 2. Finally, the target zone is represented by Oval 1.

We now work with the set of atomic propositions $\mathcal{G} := \{g_1, \dots, g_{17}\}$. The observation map h associates each g_i with the i th region (either Oval i or Halfspace i) in Fig. 4. In order to specify the surface contact requirement, we have also defined the region ‘Exterior 3’ by the polynomial $g_3 := -g_2 \in \mathbb{Q}[x_1, x_2]$. Thus, Exterior 3 is the complement of the interior of Oval 2, and the two regions only intersect on their common boundary. The formula $\varphi_C := \square(g_2 \rightarrow g_3)$ then permits the robot to make contact with the surface of the drum, without penetrating its interior. The complete mission requirement can be expressed in LTL_{\diamond} as $\varphi := \varphi_C \wedge \varphi_G \wedge \square\varphi_I \wedge \varphi_D$, where

$$\begin{aligned} \varphi_G &:= \diamond(g_2 \wedge \diamond\square g_1), \\ \varphi_I &:= \neg(g_4 \vee g_6 \vee g_8 \vee g_{10} \vee g_{12} \vee g_{14} \vee g_{16}), \\ \varphi_D &:= \neg(g_5 \vee g_7 \vee g_9 \vee g_{11} \vee g_{13} \vee g_{15} \vee g_{17}), \end{aligned}$$



trace=({6, 7}, {7}, {7, 8}, {7}, {5, 7}, {7}, {7, 9}, {6, 7, 9}, {6, 9}, {6, 7, 9}, {4, 6, 7, 9}, {6, 7, 9}, {6, 7}, {6, 7, 8}, {3, 6, 7, 8}, {6, 7, 8}, {6, 7}, {6, 7, 9}, {6, 9}, {1, 6, 9}, {1, 9}, {9}, {2, 9}, {2}, ∅, {7}, {7, 8}).

Fig. 2. Polynomial spline path connecting 8 waypoints under 9 atomic propositions. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

$$\varphi_D := \square(g_2 \rightarrow \square\varphi_0).$$

The mission goal is expressed by φ_C , which requires the robot to eventually make contact with the drum, before entering the target zone and remaining there. Formula φ_I prohibits the robot from colliding with the inner boundaries of the obstacles and ‘unsafe’ zone, which have been expanded by the robot radius. Formula φ_O does the same for the outer boundaries, which have been expanded by the drum diameter as well. Formula φ_D only requires φ_O to hold from the point of contact with the drum.

The trace of the proposed path, computed by POLYTRACE, is given at the bottom of Fig. 4. This path $r : [0, 1] \rightarrow \mathbb{R}^2$ can then be composed with a direct motion $u : [0, \infty) \rightarrow [0, 1]$ to generate the direct trajectory $x = r \circ u$. Invoking Proposition 13, the resulting trace is

$$h \circ x \circ \sigma = \{g_1, g_3, g_{13}\}, \{g_1, g_3\}, \{g_3\}, \{g_3, g_7\}, \{g_3, g_5, g_7\}, \{g_3, g_7\}, \{g_3\}, \{g_2, g_3\}, \{g_3\}, \{g_1, g_3\}^\omega,$$

where σ is the corresponding sampling function. It is straightforward to verify that $h \circ x \circ \sigma \models \varphi$, which holds for any direct motion u produced by the path tracking stage.

Remark 21. The surface contact requirement φ_C is an example of an LTL_\square formula that can only be satisfied with zero robustness. To the best of the authors’ knowledge, there are no other approaches in the literature capable of verifying a continuous-time trajectory against such a specification, because the existing verification approaches rely on strictly positive robustness margins. For a real-world implementation, it is, of course, acceptable to relax φ_C by expanding the contact boundary into a contact area. In practice, a sufficiently small expansion does not jeopardise the mission, because the use of




Fig. 4. The proposed path for the robot is shown in red. Black filled regions represent obstacles, and the red disk at (3,-7) represents the drum. The boundaries of Oval 2 and Exterior 3 coincide. Units are in meters.

feedback control makes the latching system robust to small deviations from the ideal path. However, the existing approaches, which are based on [9], still demand high frequency sampling at a rate that is not known beforehand. The tighter the expansion, the harder it becomes to verify any path proposed. As discussed in the introduction, multiple attempts may be made without guarantee of an outcome. By contrast, even in the most extreme cases, POLYTRACE never fails to return a conclusive result, as this example serves to demonstrate.

8. Conclusion

This paper culminates in an algorithmic solution to a special case of Problem 3, in which the path is a polynomial, and the observation map generates semi-algebraic regions of interest. First, the more general problem of verifying a trajectory under an observation map against a temporal logic specification is formally stated as Problem 1. Its formulation relies on the MITL semantics of [14], which can be interpreted over continuous-time signals of finite variability via their TSS representations. The trace of such a signal is identified with the first component of its TSS representation. The specification language LTL_{∞} is then introduced as a fragment of MITL, and Remark 3 emphasises that for LTL_{∞} formulas, MITL satisfaction of a continuous-time signal can be inferred from LTL satisfaction of its trace. Narrowing the scope to LTL_{∞} formulas, the focus of the paper shifts to Problem 2: generating a trace of a trajectory. As argued in Remark 10, the only trajectories that are practically amenable to path checking have infinite traces that are lasso words composed of the finite traces of paths. Thus, we arrive at Problem 3: generating the trace of a path. This can be done by sampling the path judiciously, and Theorem 1 presents topological conditions for a sampling function to achieve this. Section 5 applies this result to sample polynomial paths under observation maps that only generate semi-algebraic regions of interest, deriving more concrete conditions in terms of the roots of univariate polynomials. The POLYTRACE algorithm proposed in Section 6 invokes root isolation and root existence subroutines from computer algebra to satisfy these conditions. One novel feature of this algorithm is its ability to sample at the isolated points characterised in Corollary 7. Correctness of the algorithm is proved, and its complexity shown

to be polynomial in the number of atomic propositions. Numerical examples and a case study are provided in Section 7, including an example of an LTL specification that cannot be verified using other existing methods.

An obvious extension to our work is to relax the restriction to LTL_{∞} specifications by permitting bounded temporal operators. In this case, satisfaction no longer depends exclusively on the trace, so new analysis is required. The precise locations of the polynomial roots are then of greater importance, because they correspond to crossing times. Furthermore, we expect the polynomial assumptions to simplify the calculation of robustness margins, and improve on the margin estimates in [9,10,16,44]. This paper has only considered the verification problem, but it is possible that the synthesis problem (i.e., generating polynomial trajectories guaranteed to satisfy formal specifications) may also benefit from the methods we have discussed.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This research has received funding from the Australian Government through the Defence Cooperative Research Centre for Trusted Autonomous Systems. Other project partners include BAE Systems Australia, DST Group, The University of Adelaide, and The University of Melbourne. The authors are grateful for their partnership, which has sparked many fruitful discussions.

Appendix A. Adjustment of non-strict isolating intervals

Here we describe a procedure for adjusting the output of a root isolation algorithm to comply with Definition 22. Let $a_j := \inf I_j$ and $b_j := \sup I_j$ for all $I_j \in \mathcal{J}$, where \mathcal{J} complies with Definition 21. There are two problems to be dealt with, both involving roots at the boundary of an isolating interval. In general, Definition 21 does not permit the endpoints to be roots, except for two special cases.

- *Singleton isolating intervals:* a singleton isolating interval $I_j = \{a_j\} = \{b_j\}$ contains only a root, and must be expanded into an open isolating interval without absorbing a second root. If $b_{j-1} < a_j$, then the new left endpoint of I_j can be easily chosen in the interval (b_{j-1}, a_j) . If however $b_{j-1} = a_j$, then both a_j and b_{j-1} must be shifted to the left by an amount less than the distance between the root a_j and the unique root in I_{j-1} . This is done by Lines 15 – 20 of Algorithm 2. The right endpoint b_j is treated in a similar manner by Lines 21 – 26.
- *Roots at 0 or 1:* consider the case where 0 is both a root and the left endpoint of I_1 . This is permitted by Definition 21, but violates Clause 22.e). Here I_1 cannot be a singleton, because otherwise $I_1 = \{0\} \not\subseteq (0, 1)$. In this case, a_1 must be shifted to the right by an amount less than the distance between 0 and the unique root in I_1 . This is done by Lines 7 – 11. If 1 is both a root and the right endpoint of I_j , then b_j is treated in a similar manner by Lines 29 – 33.

In the discussion above, the endpoints may be shifted by any amount less than the minimum root separation of the polynomial being processed. For any $p \in \mathbb{Z}[s]$ of degree $d \geq 1$, the MINROOTSEP subroutine in Algorithm 2 returns

$$\varepsilon := \frac{1}{\left\lceil d^{\frac{3d}{2}} \right\rceil \|p\|_{\infty}^{d-1}} \leq d^{-\frac{d+2}{2}} \|p\|_2^{1-d},$$

which is strictly less than the minimum root separation of p by [28, Theorem 4.6]. The ceiling operator $\lceil \cdot \rceil$ is invoked to ensure a rational output. Observe that ε^{-1} grows super-exponentially with $\deg(p)$, and the length of its binary representation is $O(d \log d + d\ell)$, where $\|p\|_{\infty} \leq 2^{\ell}$. In POLYTRACE, $p = P$ and $\deg(P) \leq MD$, where M is the number of atomic propositions. In practice, the binary representation of ε becomes extremely long with even a small number of atomic propositions, severely impacting the computing time. Fortunately our knowledge of p_1, \dots, p_M , each of degree bounded by D , can be used to our advantage.

Remark 22. Let (a, b) be an isolating interval for the root s_1 of $p \in \mathbb{R}[s]$, and let $s_2 \in \{a, b\}$ be a root of p on the boundary of that interval. Suppose that $p_1 p_2 \mid p$, where it is known that $p_1 \in \mathbb{R}[s]$ has a root in (a, b) , and that $p_2(s_2) = 0$. Since any root of p_1 is a root of p , it follows that s_1 is the unique root of p_1 in (a, b) . Clearly then $|s_1 - s_2| > \text{MINROOTSEP}(p_1 p_2)$. Furthermore, there can be no other roots of p between s_1 and s_2 .

Lines 10, 17, 23 and 32 of Algorithm 2 exploit Remark 22, which allows us to compute ε for a polynomial of degree $2D$, rather than MD . This involves choosing elements from the sets $E_j := \{i \in \{1, \dots, M\} \mid p_i \text{ has roots in } I_j\}$ for $1 \leq j \leq J$, which are computed by Line 14 of Algorithm 1.

Algorithm 2 Converts isolating intervals $\{I_j \subset (0, 1) \mid 1 \leq j \leq J\}$ for $p = \prod_{i=1}^M p_i$ into strict isolating intervals, assuming $p_1, \dots, p_M \in \mathbb{Z}[s]$.

```

1: function STRICTISOLATINGINTERVALS( $p_{1:M}, I_{1:J}, E_{1:J}$ )
2:    $b_0 := 0$ 
3:   for  $1 \leq j \leq J$  do
4:      $a_j := \inf I_j, b_j := \sup I_j$ 
5:   end for
6:    $a_{J+1} := 1$ 
7:   if  $p(a_1) = 0$  and  $a_1 = 0$  then
8:     Choose  $m \in E_1$ 
9:     Find  $1 \leq i \leq M$  such that  $p_i(0) = 0$ 
10:     $a_1 \leftarrow \text{MINROOTSEP}(p_i \times p_m)$ 
11:   end if
12:   for  $1 \leq j \leq J$  do
13:     if  $a_j = b_j$  then
14:       Choose  $m \in E_j$ 
15:       if  $a_j = b_{j-1}$  then ▷ Never entered when  $j = 1$ .
16:         Choose  $i \in E_{j-1}$ 
17:          $a_j, b_{j-1} \leftarrow a_j - \text{MINROOTSEP}(p_i \times p_m)$ 
18:       else
19:          $a_j \leftarrow \frac{a_j + b_{j-1}}{2}$ 
20:       end if
21:       if  $b_j = a_{j+1}$  then ▷ Never entered when  $j = J$ .
22:         Choose  $i \in E_{j+1}$ 
23:          $b_j, a_{j+1} \leftarrow b_j + \text{MINROOTSEP}(p_i \times p_m)$ 
24:       else
25:          $b_j \leftarrow \frac{b_j + a_{j+1}}{2}$ 
26:       end if
27:     end if
28:   end for
29:   if  $p(b_J) = 0$  and  $b_J = 1$  then
30:     Choose  $m \in E_J$ 
31:     Find  $1 \leq i \leq M$  such that  $p_i(1) = 0$ 
32:      $b_J \leftarrow 1 - \text{MINROOTSEP}(p_i \times p_m)$ 
33:   end if
34:   return  $(a_{1:J}, b_{1:J})$ 
35: end function

```

```

1: function MINROOTSEP( $q$ ) ▷ Requires  $q \in \mathbb{Z}[s]$ 
2:    $d := \text{deg}(q)$ 
3:   return  $1 / \left( \lceil d^{\frac{3d}{2}} \rceil \cdot \|q\|_{\infty}^{d-1} \right)$ 
4: end function

```

Appendix B. Polynomial composition complexity

If $\text{deg}(g) \leq b$, then the general form of $g \in \mathbb{R}[x_1, \dots, x_n]$ is

$$g(x) = \sum_{i_1 + \dots + i_n \leq b} a_{i_1, \dots, i_n} x_1^{i_1} \dots x_n^{i_n},$$

subject to the obvious restriction $i_1, \dots, i_n \geq 0$. For any $f = (f_1, \dots, f_n) \in \mathbb{R}[s]^n$,

$$g \circ f = \sum_{i_1 + \dots + i_n \leq b} a_{i_1, \dots, i_n} f_1^{i_1} \dots f_n^{i_n}. \tag{53}$$

Regarding arithmetic complexity, if $\text{deg}(f_i) \leq d$ for all i , then $\text{deg}(f_1^{i_1} \dots f_n^{i_n}) \leq bd$. There are $\binom{b+n}{n}$ summands in (53) [42, Lemma 8.5], and forming each term $f_1^{i_1} \dots f_n^{i_n}$ requires at most $b^2 d^2$ operations [30, Lemma 3.15]. Thus, $\text{deg}(g \circ f) \leq bd$, and $g \circ f$ can be computed in $O\left(b^2 d^2 \binom{b+n}{n}\right)$ arithmetic operations.

We can also use the polynomial norm inequalities in [30, Chapter 3] to place a bound on coefficient sizes. Suppose that $\|f_j\|_{\infty} \leq 2^\ell$ for all $1 \leq j \leq n$. Then

$$\left\| \prod_{j=1}^n f_j^{i_j} \right\|_{\infty} \leq \left\| \prod_{j=1}^n f_j^{i_j} \right\|_1 \leq \prod_{j=1}^n \|f_j\|_1^{i_j} \leq \prod_{j=1}^n d \|f_j\|_{\infty}^{i_j} \leq d^n \prod_{j=1}^n 2^{\ell i_j} = d^n 2^{\ell b}.$$

If in addition $|a_{i_j}| < 2^m$ for every coefficient of g in (53), then $\|g \circ f\|_{\infty} \leq d^n 2^{m\ell b}$.

References

- [1] D. Basin, B.N. Bhatt, D. Traytel, Optimal proofs for linear temporal logic on lasso words, in: S.K. Lahiri, C. Wang (Eds.), *Automated Technology for Verification and Analysis*, in: *Lecture Notes in Computer Science*, Springer International Publishing, Cham, 2018, pp. 37–55.
- [2] C. Baier, J.-P. Katoen, *Principles of Model Checking*, MIT Press, 2008.
- [3] N. Markey, P. Schnoebelen, Model checking a path, in: R. Amadio, D. Lugiez (Eds.), *CONCUR 2003 - Concurrency Theory*, in: *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg, 2003, pp. 251–265.
- [4] K. Havelund, G. Roşu, Runtime verification - 17 years later, in: C. Colombo, M. Leucker (Eds.), *Runtime Verification*, in: *Lecture Notes in Computer Science*, Springer International Publishing, Cham, 2018, pp. 3–17.
- [5] A. Bauer, M. Leucker, C. Schallhart, Runtime verification for LTL and TLTL, *ACM Trans. Softw. Eng. Methodol.* 20 (Sept. 2011) 14:1–14:64.
- [6] O. Maler, D. Nickovic, Monitoring temporal properties of continuous signals, in: Y. Lakhnech, S. Yovine (Eds.), *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems*, in: *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg, 2004, pp. 152–166.
- [7] R. Koymans, Specifying real-time properties with metric temporal logic, *Real-Time Syst.* 2 (Nov. 1990) 255–299.
- [8] C. Eisner, D. Fisman, J. Havlicek, Y. Lustig, A. Mclsaac, D. Van Campenhout, Reasoning with temporal logic on truncated paths, in: W.A. Hunt, F. Somenzi (Eds.), *Computer Aided Verification*, in: *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg, 2003, pp. 27–39.
- [9] G.E. Fainekos, G.J. Pappas, Robustness of temporal logic specifications for continuous-time signals, *Theor. Comput. Sci.* 410 (Sept. 2009) 4262–4291.
- [10] A. Donzé, O. Maler, Robust satisfaction of temporal logic over real-valued signals, in: K. Chatterjee, T.A. Henzinger (Eds.), *Formal Modeling and Analysis of Timed Systems*, in: *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg, 2010, pp. 92–106.
- [11] M. Kloetzer, C. Belta, A fully automated framework for control of linear systems from temporal logic specifications, *IEEE Trans. Autom. Control* 53 (Feb. 2008) 287–297.
- [12] J. Liu, N. Ozay, U. Topcu, R.M. Murray, Synthesis of reactive switching protocols from temporal logic specifications, *IEEE Trans. Autom. Control* 58 (July 2013) 1771–1785.
- [13] T. Wongpiromsarn, U. Topcu, A. Lamperski, Automata theory meets barrier certificates: temporal logic verification of nonlinear systems, *IEEE Trans. Autom. Control* 61 (Nov. 2016) 3344–3355.
- [14] R. Alur, T. Feder, T.A. Henzinger, The benefits of relaxing punctuality, *J. ACM* 43 (Jan. 1996) 116–146.
- [15] G. Babin, Y. Aït-Ameur, S. Nakajima, M. Pantel, Refinement and proof based development of systems characterized by continuous functions, in: X. Li, Z. Liu, W. Yi (Eds.), *Dependable Software Engineering: Theories, Tools, and Applications*, in: *Lecture Notes in Computer Science*, Springer International Publishing, Cham, 2015, pp. 55–70.
- [16] A. Donzé, T. Ferrère, O. Maler, Efficient robust monitoring for STL, in: N. Sharygina, H. Veith (Eds.), *Computer Aided Verification*, in: *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg, 2013, pp. 264–279.
- [17] H. Abbas, Y.V. Pant, R. Mangharam, Temporal logic robustness for general signal classes, in: *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*, Montreal, Quebec, Canada, ACM, Apr. 2019, pp. 45–56.
- [18] J. von zur Gathen, J. Gerhard, *Modern Computer Algebra*, Cambridge University Press, New York, United States, 2013.
- [19] D. Burke, A. Chapman, I. Shames, Fast spline trajectory planning: minimum snap and beyond, arXiv:2105.01788 [cs, eess], May 2021.
- [20] T. Mercy, R.V. Parys, G. Pipeleers, Spline-based motion planning for autonomous guided vehicles in a dynamic environment, *IEEE Trans. Control Syst. Technol.* 26 (Nov. 2018) 2182–2189.
- [21] B. Lau, C. Sprunk, W. Burgard, Kinodynamic motion planning for mobile robots using splines, in: *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, St. Louis, MO, USA, Oct. 2009, pp. 2427–2433.
- [22] Z. Shiller, Y. Gwo, Dynamic motion planning of autonomous vehicles, *IEEE Trans. Robot. Autom.* 7 (Apr. 1991) 241–249.
- [23] S. Yoon, D. Lee, J. Jung, D.H. Shim, Spline-based RRT* using piecewise continuous collision-checking algorithm for car-like vehicles, *J. Intell. Robot. Syst.* 90 (June 2018) 537–549.
- [24] R. Wein, O. Ilushin, G. Elber, D. Halperin, Continuous path verification in multi-axis NC-machining, in: *Proceedings of the Twentieth Annual Symposium on Computational Geometry*, SCG '04, Association for Computing Machinery, New York, NY, USA, June 2004, pp. 86–95.
- [25] J.R. Munkres, *Topology*, 2nd ed., Prentice Hall, Inc., 2000.
- [26] R. Engelking, *General Topology*, Heldermann, 1989.
- [27] D.S. Dummit, R.M. Foote, *Abstract Algebra*, Wiley, 2004.
- [28] M. Mignotte, *Mathematics for Computer Algebra*, Springer, New York, 1992.
- [29] P. Gianni, B. Trager, Square-free algorithms in positive characteristic, *Appl. Algebra Eng. Commun. Comput.* 7 (Jan. 1996) 1–14.
- [30] J. Gerhard, *Modular Algorithms in Symbolic Summation and Symbolic Integration*, *Lecture Notes in Computer Science*, Springer-Verlag, Berlin, Heidelberg, 2005.
- [31] C.A. Furiá, M. Rossi, On the expressiveness of MTL variants over dense time, in: J.-F. Raskin, P.S. Thiagarajan (Eds.), *Formal Modeling and Analysis of Timed Systems*, in: *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg, 2007, pp. 163–178.
- [32] P. Wolper, Constructing automata from temporal logic formulas: a tutorial*, in: E. Brinksma, H. Hermans, J.-P. Katoen (Eds.), *Lectures on Formal Methods and Performance Analysis: First EEF/Euro Summer School on Trends in Computer Science*, Bergen Dal, The Netherlands, July 3–7, 2000 Revised Lectures, in: *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg, 2001, pp. 261–277.
- [33] D. Verscheure, B. Demeulenaere, J. Swevers, J. De Schutter, M. Diehl, Time-optimal path tracking for robots: a convex optimization approach, *IEEE Trans. Autom. Control* 54 (Oct. 2009) 2318–2327.
- [34] M.Y. Vardi, An automata-theoretic approach to linear temporal logic, in: F. Moller, G. Birtwistle (Eds.), *Logics for Concurrency: Structure Versus Automata*, in: *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg, 1996, pp. 238–266.
- [35] B. Khoussainov, A. Nerode, *Automata Theory and Its Applications*, Birkhäuser, Boston, MA, 2001.
- [36] K. Mehlhorn, M. Sagraloff, A deterministic algorithm for isolating real roots of a real polynomial, *J. Symb. Comput.* 46 (Jan. 2011) 70–90.
- [37] G.E. Collins, R. Loos, Real zeros of polynomials, in: B. Buchberger, G.E. Collins, R. Loos (Eds.), *Computer Algebra: Symbolic and Algebraic Computation*, in: *Computing Supplementum*, Springer, Vienna, 1982, pp. 83–94.
- [38] W. Krandick, K. Mehlhorn, New bounds for the Descartes method, *J. Symb. Comput.* 41 (Jan. 2006) 49–66.
- [39] G.E. Collins, A.G. Akritas, Polynomial real root isolation using Descartes's rule of signs, in: *Proceedings of the Third ACM Symposium on Symbolic and Algebraic Computation*, SYMSAC '76, Association for Computing Machinery, New York, NY, USA, Aug. 1976, pp. 272–275.
- [40] A. Eigenwillig, V. Sharma, C.K. Yap, Almost tight recursion tree bounds for the Descartes method, in: *Proceedings of the 2006 International Symposium on Symbolic and Algebraic Computation*, ISSAC '06, Association for Computing Machinery, New York, NY, USA, July 2006, pp. 71–78.
- [41] G.E. Collins, Polynomial minimum root separation, *J. Symb. Comput.* 32 (Nov. 2001) 467–473.
- [42] S. Basu, R. Pollack, M.-F. Roy, *Algorithms in Real Algebraic Geometry*, *Algorithms and Computation in Mathematics*, second ed., Springer-Verlag, Berlin, Heidelberg, 2006.
- [43] J. Sendra, Brief atlas of offset curves, in: L. Lambán, A. Romero, J. Rubio (Eds.), *Contribuciones Científicas en Honor de Mirian Andrés Gómez*, Universidad de la Rioja, 2010, pp. 483–505.
- [44] J.V. Deshmukh, A. Donzé, S. Ghosh, X. Jin, G. Juniwal, S.A. Seshia, Robust online monitoring of signal temporal logic, *Form. Methods Syst. Des.* 51 (Aug. 2017) 5–30.