

Probabilistic Temporal Logic Falsification of Cyber-Physical Systems*

Houssam Abbas[†], Georgios Fainekos[‡], Sriram Sankaranarayanan[§],
Franjo Ivančić[¶], and Aarti Gupta^{||}

December 22, 2011

Abstract

We present a Monte-Carlo optimization technique for finding system behaviors that falsify a Metric Temporal Logic (MTL) property. Our approach performs a random walk over the space of system inputs guided by a robustness metric defined by the MTL property. Robustness is guiding the search for a falsifying behavior by exploring trajectories with smaller robustness values. The resulting testing framework can be applied to a wide class of Cyber-Physical Systems (CPS). We show through experiments on complex system models that using our framework can help automatically falsify properties with more consistency as compared to other means such as uniform sampling.

1 Introduction

Model Based Design (MBD) and automatic code generation are becoming the development methodologies of choice for safety critical applications. Most prominently, such design methodologies have been adopted by the automotive, medical and aerospace industries [45, 23] where correctness of the end product is of paramount importance. The types of systems in these industrial domains are particularly challenging because software is controlling the safe operation of a physical system. Such systems are also known as Cyber-Physical Systems (CPS). One of the pressing challenges in the MBD of CPS is how to verify the correctness of the developed model of the system as early as possible in the design cycle.

In answering such a problem, one must first specify what is an appropriate mathematical model that captures the behavior of the system and, second, what is an appropriate specification framework that has a nice mathematical structure that can help in analyzing the mathematical model of the system. One such popular mathematical framework for CPS modeling is hybrid automata [33]. Unfortunately, in general, the verification problem for hybrid automata is undecidable even for simple safety requirements [34], i.e., there is no terminating algorithm that can answer whether a CPS ever enters a set of bad states. Thus, a lot of research has focused on discovering the classes of hybrid automata where the safety verification problem is decidable [5] and on reachability analysis and testing based techniques [61].

However, in many cases, the system requirements extend well beyond simple safety properties. For example, we might be interested in conditional requirements such that “if the temperature increases above 10 degrees and remains above 10 degrees for 1 min, then it should be drop below 10 degrees within 2 min and remain below 10 degrees for 30min.” Such specifications can be captured using Metric Temporal Logic (MTL) [38].

*This work was partially supported by NSF grants CNS-1017074, CNS-1116136 and CNS-1016994. ©ACM, (2011). This is the authors version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version will be published in ACM Transactions on Embedded Computing Systems (TECS).

[†]School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ

[‡]School of Computing, Informatics, and Decision Systems Engineering, Arizona State University, Tempe, AZ

[§]Department of Computer Science, University of Colorado, Boulder, CO

[¶]NEC Laboratories America, Princeton, NJ

^{||}NEC Laboratories America, Princeton, NJ

In this paper, we propose a testing-based technique for finding counterexamples to MTL properties for CPS through global minimization of a *robustness metric*. The global optimization is carried out using a Monte-Carlo technique that performs a random walk over the space of inputs consisting of initial states and control inputs. The robustness metric defines the satisfaction of an MTL property over a given trajectory as a real number, as opposed to the Boolean notion used in Logic. The sign of the metric for a given trajectory \mathbf{y} and formula φ indicates whether \mathbf{y} satisfies φ (written as $\mathbf{y} \models \varphi$). Furthermore, “nearby” trajectories, defined using a metric over trajectories, whose distances from \mathbf{y} are smaller than its robustness also have the same outcome for the property φ as \mathbf{y} .

Given a robustness metric, finding a counterexample to a given property φ reduces to finding a trajectory \mathbf{y} that minimizes the robustness score w.r.t φ . This can be viewed as an optimization problem over the space of inputs of the system. However, in practice, this optimization problem is not necessarily guaranteed to be tractable. In almost all cases, the optimization problem (objective function and constraints) cannot be written down in a closed functional form. Nevertheless, such optimization problems can often be solved satisfactorily using Monte-Carlo techniques that perform a random walk in order to sample from a probability distribution defined implicitly by the robustness metric [54]. Over the long run, the random walk converges to a stationary distribution over the input space such that inputs with smaller values of robustness are sampled more frequently than inputs with larger values. Furthermore, Monte-Carlo techniques do not require the distribution itself to be known in a closed form. Instead, these techniques simply require the ability to compare the values (ratio) of the probability density function at two given points in the search space. In practice, this reduces to simulating the system using the sampled inputs.

The contributions of this work can be summarized as follows:

1. We show that metrics used for robust testing naturally define objective functions that enable us to cast the problem of falsifying MTL properties into a global optimization problem.
2. We demonstrate the use of hit-and-run Monte-Carlo samplers to carry out this optimization in the presence of (possibly non-convex) constraints over the inputs.
3. We extend our notions to CPS using quasi-metrics to provide a notion of robustness for hybrid trajectories w.r.t properties that can involve discrete as well as continuous state variables.

Our approach is applicable even if the property has been proven using a verification technique. In such cases, our technique obtains system trajectories that have low robustness values w.r.t the requirements. In practice, finding non-robust trajectories may imply designs with smaller safety margins. Traditional testing or verification techniques do not consider such trajectories using Boolean notions of temporal satisfaction. Our approach is readily applicable to *Simulink/Stateflow*TM (S/S) models, since simulating the system is the only primitive needed. We have implemented our approach in the Matlab (TM) toolbox S-TALIRO [8] and use it to discover counterexamples to MTL properties. We establish that random walks guided by robustness metrics can often falsify MTL properties that cannot be easily falsified using blind (uniform random) search.

Preliminary results of this work have appeared in [47], while the architecture of our toolbox S-TALIRO has appeared in [8]. In this paper, we reformulate the problem and its solution into a more general framework, we present the proofs that were omitted from [47], we provide new hybrid metrics in Section 4 and we perform more thorough experimental analysis using our toolbox S-TALIRO.

2 Preliminaries

In this section, we provide a formal and concise definition of the problem that this work addresses. Then, we introduce metrics and we utilize them to provide continuous semantics for Metric Temporal Logic (MTL) specifications over continuous time trajectories. We will be using the following notation: \mathbb{R} is the set of real numbers; $\overline{\mathbb{R}}$ is the closure of the reals, i.e., $[-\infty, +\infty]$; \mathbb{R}^+ is the set of positive real numbers and $\overline{\mathbb{R}}_+$ its closure, i.e., $\overline{\mathbb{R}}_+ = [0, +\infty]$; \mathbb{N} is the set of natural numbers (including 0) and $\mathbb{N}_\infty = \mathbb{N} \cup \{+\infty\}$; \mathbb{Z} is the set of integers and $\mathbb{Z}_\infty = \mathbb{Z} \cup \{\pm\infty\}$. Given sets A and B , B^A defines the set of all functions from A to B and $\mathcal{P}(A)$ denotes the powerset of A .

2.1 Problem Definition

In this work, we take a very general approach in modeling real-time embedded systems that interact with physical systems that have non-trivial dynamics. Such systems are also referred to as hybrid systems or Cyber-Physical Systems (CPS). In the following, we will be using the term hybrid systems since it is more concise. However, we would like to caution the reader against associating hybrid systems with hybrid automata [3] since the scope of our work is more general.

We view a system Σ as a mapping from a compact set of initial conditions X_0 and input signals $\mathbf{U} \subseteq U^R$ to output signals Y^R . Here, R is a bounded time domain equipped with a metric d_R , U is a compact set of possible input values at each point in time (input space) and Y is the set of output values (output space). This view of a system is standard in signals and systems [40]. We impose four assumptions / restrictions on the systems that we consider:

1. The input signals (if any) must be parameterizable using a finite number of parameters. That is, there exists a function \mathfrak{U} such that for any $u \in \mathbf{U}$, there exist two parameter vectors $\lambda = [\lambda_1 \dots \lambda_m]^T \in \Lambda$, where Λ is a compact set, and $\tau = [\tau_1 \dots \tau_m]^T \in R^m$ such that for all $t \in R$, $u(t) = \mathfrak{U}(\lambda, \tau)(t)$.
2. The output space Y must be equipped with a generalized metric \mathbf{d} which contains a subspace Z equipped with a metric d .
3. For a specific initial condition x_0 and input signal u , there must exist a unique output signal \mathbf{y} defined over the time domain R . That is, the system Σ is deterministic and we implicitly assume that the system does not exhibit Zeno behaviors [44].
4. For considering the convergence of our sampling scheme, we assume that the space of inputs is bounded and discretized to a large but finite set. In practice, any representation of the input through a vector of floating point numbers inside the computer must be finite and, therefore, implicitly discretizes the space of inputs. Thus, this assumption does not pose a restriction.

Under Assumption 3, a system Σ can be viewed as a function $\Delta_\Sigma : X_0 \times \mathbf{U} \rightarrow Y^R$ which takes as an input an initial condition $x_0 \in X_0$ and an input signal $u \in \mathbf{U}$ and it produces as output a signal $\mathbf{y} : R \rightarrow Y$ (also referred to as *trajectory*). When the output signals are only a function of the initial condition, i.e., $\Delta_\Sigma : X_0 \rightarrow Y^R$, then the system Σ is called *autonomous*. In either case, the set of all output signals of Σ will be denoted by $\mathcal{L}(\Sigma)$. That is, $\mathcal{L}(\Sigma) = \{\mathbf{y} \mid \exists x_0 \in X_0 . \exists u \in \mathbf{U} . \mathbf{y} = \Delta_\Sigma(x_0, u)\}$ or in case of autonomous systems $\mathcal{L}(\Sigma) = \{\mathbf{y} \mid \exists x_0 \in X_0 . \mathbf{y} = \Delta_\Sigma(x_0)\}$.

Assumption (1) is necessary in order to define a feasible search problem over a potentially infinite function space. Assumption (2) is not a restrictive condition since any reasonable hybrid system must have some output space that can be equipped with a non-trivial metric. However, Assumption (3) is a fundamental restriction in our current approach. In Section 3.1, we provide some brief remarks on how this assumption may be lifted.

Our high level goal is to infer the correctness of the system Σ by observing its response (output signals) to particular input signals and initial conditions. In particular, we are interested in finding witnesses, i.e., output signals, which prove that a requirement or specification is not satisfied by the system. The process of discovering such witnesses is usually referred to as *falsification*.

Example 2.1 *As a motivating example, we will consider the Automatic Transmission example which was also considered in [64]. This is a slightly modified version of the Automatic Transmission model provided by Mathworks as a Simulink demo¹. It is a model of an automatic transmission controller (see Fig. 1) with the following modifications. The only input to the system is the throttle schedule, while the break schedule is set simply to 0 for the duration of the simulation which is 30 sec, i.e., $R = [0, 30]$. Finally, the system has two outputs the speed of the engine ω (RPM) and the speed of the vehicle v , i.e., $Y = \mathbb{R}^2$ and $\mathbf{y}(t) = [\omega(t) v(t)]^T$ for all $t \in [0, 30]$.*

Internally, the system has two 2 continuous-time state variables: the vehicle speed v and engine speed ω . That is, for this example, the output of the system is the same as the continuous state of the system. Initially, the vehicle is at rest at time 0, i.e., $X_0 = \{[0 \ 0]^T\}$ and $x_0 = \mathbf{y}(0) = [0 \ 0]^T$. Therefore, the output trajectories depend only on

¹Available at: <http://www.mathworks.com/products/simulink/demos.html>

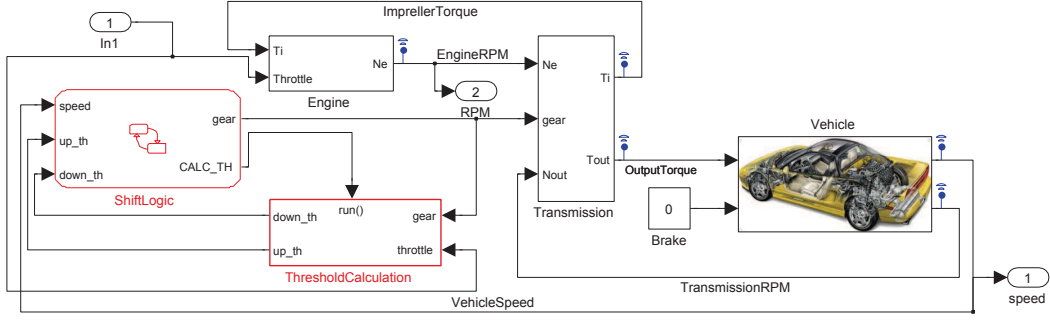


Figure 1: The modified version of the Simulink (TM) Automatic Transmission Demo `sldemo_autotrans`.

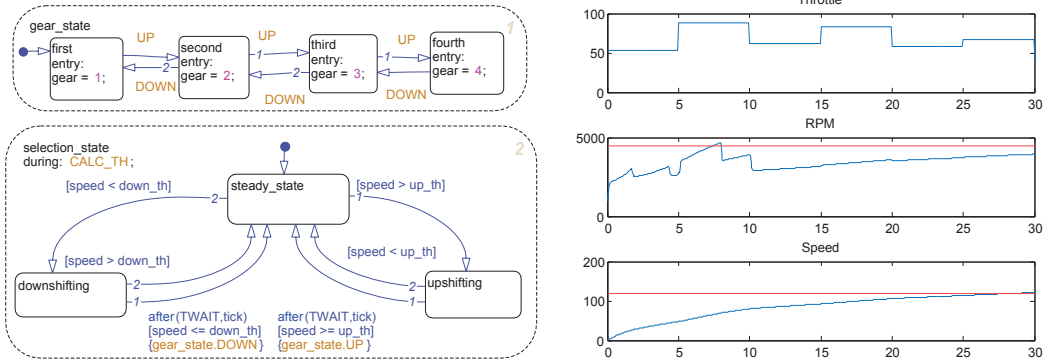


Figure 2: Example 2.1. *Left*: The switching logic for the automatic drivetrain; *Right*: A input signal and the corresponding output signals that falsify the specification.

the input signal u which models the throttle, i.e., $y = \Delta_{\Sigma}(u)$. The throttle at each point in time can take any value between 0 (fully closed) to 100 (fully open). Namely, $u(t) \in U = [0, 100]$ for each $t \in [0, 30]$. We remark that the system is deterministic, i.e., under the same input u , we will always observe the same output y .

We will assume that a system specification requires that the vehicle speed v is always under 120km/h or that the engine speed ω is always below 4500RPM. Our goal is to falsify the above specification. In other words, we would like to generate tests such that the vehicle speed v and the engine speed ω exceed the values 120km/h and 4500RPM, respectively. Such a falsifying system trajectory appears in Fig. 2.

The model contains 69 blocks out of which there are 2 integrators (i.e., 2 continuous state variables), 3 look-up tables, 3 look-up 2D tables and a Stateflow chart. The Stateflow chart (see Fig. 2 for a schematic) contains two concurrently executing Finite State Machines (FSM) with 4 and 3 states, respectively. Even though this is a small size model and the specification is a simple bounded time reachability requirement, it already exhibits all the complexities, i.e., look-up tables, switching conditions which depend on inputs, that prevent formal modeling and analysis using the state of the art tools, e.g., SpaceEx [29]. \diamond

Reachability requirements as described in Example 2.1 do not suffice to specify all system behaviors in practice. This is especially true for real-time embedded systems wherein richer properties such as timing requirements, sequencing of events, conditional requirements, stability and so on are equally important. Metric Temporal Logic (MTL) introduced by Koymans [38] is a popular formalism that can express such properties. Our objective in this work is to provide efficient tools for the falsification of bounded time MTL properties for CPS.

Problem 2.1 (MTL Falsification) For an MTL specification φ , the MTL falsification problem consists of finding an




Figure 3: Overview of the solution to the MTL Falsification of CPS.

output signal y of the system Σ starting from some valid initial state $x_0 \in X_0$ under an input signal $u \in \mathbf{U}$ such that y does not satisfy specification φ .

An overview of our proposed solution to Problem 2.1 appears in Fig. 3. The sampler produces a point x_0 from the set of initial conditions and a vector of parameters λ that characterize the control input signal u . These are passed to the system simulator which returns an execution trace (output trajectory). The trace is then analyzed by the MTL robustness analyzer which returns a robustness value. In turn, the robustness score computed is used by the stochastic sampler to decide on a next input to analyze. If in this process, a falsifying trace is found, it is returned to the user, who can then proceed to examine it inside the system modeling environment.

In this paper, not only we provide an efficient solution to Problem 2.1, but we are also able to provide a measure of how robustly the system satisfies or not an MTL property. That is, our falsification framework does not have to return the first falsifying trajectory it detects, but it can continue searching for the least possible robust system behavior. Similarly, even if the system is not falsifiable, our tool returns the least robust correct behavior that was detected. Such information can be valuable to the system designer.

2.2 Metrics and Distances

When given a collection of objects, it is frequently necessary to reason about how “close” these objects are to each other. In other words, we need a way to measure or compute the distance between any two objects in the collection. In mathematics, the distance between two objects that belong to a set Y can be quantified by a metric \mathbf{d} . The pair (Y, \mathbf{d}) is called a metric space.

Metrics arise very naturally in control and analysis of physical systems [60]. Interesting metrics can also be defined in computation theory with a number of diverse applications [57]. In either case, the interest in defining metrics is usually to show that a function is contractive (and, thus, to prove some notion of stability [60] or utilize a fixed-point computation [57]) or that we can define an interesting topology [37]. Here, our interest in metrics is different. We are interested in quantifying set membership questions, i.e., how deep is the object within the set it belongs or how far away is from the set it should belong. At a high level, quantification of set membership questions is the subject of study in fuzzy mathematics [9]. The fundamental difference between fuzzy set theory and our work is that fuzzy set theory abstracts away any topological information regarding the degree of membership. Such topological information is vital in our case as we will demonstrate in Section 2.3. Next, we briefly review the notion of generalized metrics and we refer the reader to [57] and the references therein for a more detailed exposition.

Def. 2.1 (Positively Ordered Commutative Monoid) • A semigroup $(V, +)$ is a set V together with a binary operation $+$ such that (i) the set is closed under $+$ and (ii) $+$ is associative.

- A monoid is a semigroup which has an identity element $\mathbf{0}$, i.e., for any $v \in V$, $v + \mathbf{0} = \mathbf{0} + v = v$.
- A commutative monoid is a monoid whose binary operation is commutative.

- An ordered monoid $(V, +, \preceq)$ is a monoid with an (partial) order relation \preceq which is compatible with $+$, i.e., $v_1 \preceq v_2$ implies $v_1 + v_3 \preceq v_2 + v_3$ and $v_3 + v_1 \preceq v_3 + v_2$ for all $v_1, v_2, v_3 \in V$.
- A positively ordered monoid is an ordered monoid such that for all $v \in V$, $\mathbf{0} \preceq v$.

Def. 2.2 (Generalized Metric) Let $(V, +, \preceq)$ be a positively ordered commutative monoid and Y be an arbitrary set. A generalized metric \mathbf{d} is a function $\mathbf{d} : Y \times Y \mapsto V$ which satisfies the following properties for $y_1, y_2, y_3 \in Y$:

Identity $\mathbf{d}(y_1, y_2) = \mathbf{0}$ iff $y_1 = y_2$,

Symmetry $\mathbf{d}(y_1, y_2) = \mathbf{d}(y_2, y_1)$, and

Triangle Inequality $\mathbf{d}(y_1, y_3) \preceq \mathbf{d}(y_1, y_2) + \mathbf{d}(y_2, y_3)$.

If V also has an absorbing element ∞ , i.e., for any $v \in V$, $v + \infty = \infty + v = \infty$, then \mathbf{d} is called an extended generalized metric. If the Symmetry condition is dropped from the definition, then \mathbf{d} is termed a generalized quasi-metric. If $(V, +, \preceq)$ is $(\mathbb{R}_+, +, \leq)$ with the usual addition $+$ and total order \leq , then we drop the term ‘‘generalized’’ from the terminology and denote the metric by d .

Using a generalized metric \mathbf{d} , we can define the distance of a point $y \in Y$ from a set $S \subseteq Y$. Intuitively, this distance is the shortest distance from y to all the points in S . In a similar way, the depth of a point y in a set S is defined to be the shortest distance of y from the boundary of S .

Def. 2.3 (Distance, Depth, Signed Distance [11] §8) Let $y \in Y$ be a point, $S \subseteq Y$ be a set and \mathbf{d} be a generalized metric on Y . Then, we define the

- Distance from y to S to be $\mathbf{dist}_{\mathbf{d}}(y, S) := \inf\{\mathbf{d}(y, y') \mid y' \in S\}$
- Signed Distance from y to S to be

$$\mathbf{Dist}_{\mathbf{d}}(y, S) := \begin{cases} -\mathbf{dist}_{\mathbf{d}}(y, S) & \text{if } y \notin S \\ \mathbf{dist}_{\mathbf{d}}(y, Y \setminus S) & \text{if } y \in S \end{cases}$$

We should point out that we use the extended definition of supremum and infimum. In other words, the supremum of the empty set is defined to be bottom element of the domain, while the infimum of the empty set is defined to be the top element of the domain. For example, when we reason over $\overline{\mathbb{R}}$, then $\sup \emptyset := -\infty$ and $\inf \emptyset := +\infty$.

Also of importance is the notion of an open ball of radius ε centered at a point $y \in Y$. Given a generalized metric \mathbf{d} , a radius $\varepsilon \in V$ and a point $y \in Y$, the open ε -ball (or neighborhood) centered at y is defined as $\mathcal{B}_{\mathbf{d}}(y, \varepsilon) = \{y' \in Y \mid \mathbf{d}(y, y') \prec \varepsilon\}$. The previous definition of a neighborhood includes all points y' which have distance from y less than ε . Since in this work we also use quasi-metrics, we also need the notion of *neighborhood-to*. The neighborhood-to includes all points y' which have distance to y less than ε . Similar to $\mathcal{B}_{\mathbf{d}}$, we define $\mathcal{N}_{\mathbf{d}}(y, \varepsilon) = \{y' \in Y \mid \mathbf{d}(y', y) \prec \varepsilon\}$.

Finally, in order to reason in time about the system behavior, we need to define metrics over signal spaces. If \mathbf{y} and \mathbf{y}' are two system output signals $\mathbf{y}, \mathbf{y}' : R \rightarrow Y$ that take values in a generalized metric space (Y, \mathbf{d}) , we will use $\rho_{\mathbf{d}}$ to denote the metric $\rho_{\mathbf{d}}(\mathbf{y}, \mathbf{y}') = \sup_{t \in R} \{\mathbf{d}(\mathbf{y}(t), \mathbf{y}'(t))\}$.

2.3 Robustness of Trajectories

With the help of metrics we can now provide a robust interpretation (semantics) to MTL formulas. Details are available in our previous work [26]. In this section, we refer to output signals simply as *signals*.

Def. 2.4 (MTL Syntax) Let AP be the set of atomic propositions and \mathcal{I} be any non-empty interval of $\overline{\mathbb{R}}_+$. The set MTL of all well-formed MTL formulas is inductively defined as $\varphi ::= \top \mid p \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \mathcal{U}_{\mathcal{I}} \varphi$, where $p \in AP$ and \top is true.

For (real-time) hybrid systems, the atomic propositions label subsets of the output space Y . An *observation map* $\mathcal{O} : AP \rightarrow \mathcal{P}(Y)$ maps each proposition $p \in AP$ to a set $\mathcal{O}(p)$ to a subset of Y . We require that for all $p \in AP$, $\emptyset \subset \mathcal{O}(p) \subset Y$. We emphasize here that the results in [26] require that the output space Y is equipped with an extended metric d . In Section 4, we relax this requirement and we demonstrate how these results are extended to output spaces which are equipped with a generalized quasi-metric.

We provide semantics that maps an MTL formula φ and a signal $\mathbf{y}(t)$ to a value drawn from the linearly ordered set $\overline{\mathbb{R}}$. The semantics for the atomic propositions evaluated for $\mathbf{y}(t)$ consists of the distance between $\mathbf{y}(t)$ and the set $\mathcal{O}(p)$ labeling atomic proposition p . Intuitively, this distance represents how robustly the point $\mathbf{y}(t)$ lies within (or is outside) the set $\mathcal{O}(p)$. If this distance is zero, then the smallest perturbation of the point y can affect the outcome of $y \in \mathcal{O}(p)$. We denote the robust valuation of the formula φ over the signal \mathbf{y} at time t by $\llbracket \varphi, \mathcal{O} \rrbracket_d(\mathbf{y}, t)$. Formally, $\llbracket \cdot, \cdot \rrbracket_d : (MTL \times \mathcal{P}(Y)^{AP}) \rightarrow (Y^R \times R \rightarrow \overline{\mathbb{R}})$.

Def. 2.5 (Robust Semantics) Consider a metric space (Y, d) , where d is an extended metric. Let $\mathbf{y} \in Y^R$, $c \in \overline{\mathbb{R}}$ and $\mathcal{O} \in \mathcal{P}(Y)^{AP}$, then the robust semantics of any formula $\varphi \in MTL$ with respect to \mathbf{y} is recursively defined as follows

$$\begin{aligned} \llbracket \top, \mathcal{O} \rrbracket_d(\mathbf{y}, t) &:= +\infty \\ \llbracket p, \mathcal{O} \rrbracket_d(\mathbf{y}, t) &:= \mathbf{Dist}_d(\mathbf{y}(t), \mathcal{O}(p)) \\ \llbracket \neg\varphi_1, \mathcal{O} \rrbracket_d(\mathbf{y}, t) &:= -\llbracket \varphi_1, \mathcal{O} \rrbracket_d(\mathbf{y}, t) \\ \llbracket \varphi_1 \vee \varphi_2, \mathcal{O} \rrbracket_d(\mathbf{y}, t) &:= \max(\llbracket \varphi_1, \mathcal{O} \rrbracket_d(\mathbf{y}, t), \llbracket \varphi_2, \mathcal{O} \rrbracket_d(\mathbf{y}, t)) \\ \llbracket \varphi_1 \mathcal{U}_{\mathcal{I}} \varphi_2, \mathcal{O} \rrbracket_d(\mathbf{y}, t) &:= \sup_{t' \in (t+R\mathcal{I})} \min(\llbracket \varphi_2, \mathcal{O} \rrbracket_d(\mathbf{y}, t'), \inf_{t < t'' < t'} \llbracket \varphi_1, \mathcal{O} \rrbracket_d(\mathbf{y}, t'')) \end{aligned}$$

where $t \in R$ and $t+R\mathcal{I} = \{\tau \mid \exists \tau' \in \mathcal{I}. \tau = t + \tau'\} \cap R$.

Example 2.2 The requirement expressed in natural language in Example 2.1 can be formally written as: $\phi_0^{AT} = \Box p_1^{AT} \vee \Box p_2^{AT}$, where each atomic proposition p_i^0 is mapped to: $\mathcal{O}(p_1^{AT}) = [120, +\infty) \times \mathbb{R}$ and $\mathcal{O}(p_2^{AT}) = \mathbb{R} \times [4500, +\infty)$, respectively. From the designer perspective, it might be easier to conceptualize the falsification problem as a test generation problem and, therefore, pose the formal requirement as the negation of the behavior that she/he would like to observe, i.e., $\phi_1^{AT} = \neg(\Diamond p_1^{AT} \wedge \Diamond p_2^{AT})$. Under the semantics of Def. 2.5, the two formulas are equivalent. \diamond

For the purposes of the following discussion, let $(\mathbf{y}, t, \mathcal{O}) \models \varphi$ denote the standard Boolean MTL satisfiability. For clarity in the presentation, we define the satisfiability relation for the base case, i.e., for atomic propositions: $p \in AP$, $(\mathbf{y}, t, \mathcal{O}) \models \varphi$ if $\mathbf{y}(t) \in \mathcal{O}(p)$. Note that Boolean MTL satisfiability reduces to an application of Def. 2.5 wherein the negation is defined to be the Boolean negation and the metric d is the discrete metric: for $y_1, y_2 \in Y$, $d(y_1, y_2) = 0$ if $y_1 = y_2$ and $d(y_1, y_2) = 1$ if $y_1 \neq y_2$. It is easy to show that if the signal satisfies the property, then its robustness is non-negative and, similarly, if the signal does not satisfy the property, then its robustness is non-positive. The following result holds [26].

Theorem 2.1 Given an output space (Y, d) , where d is an extended metric, a formula $\varphi \in MTL$, an observation map $\mathcal{O} \in \mathcal{P}(Y)^{AP}$ and an output signal $\mathbf{y} \in Y^R$, the following hold:

1. If $(\mathbf{y}, t, \mathcal{O}) \models \varphi$, then $\llbracket \varphi, \mathcal{O} \rrbracket_d(\mathbf{y}, t) \geq 0$. In other words, if \mathbf{y} satisfies the formula φ at time instant $t \geq 0$, then its distance valuation is non-negative. Conversely, if $\llbracket \varphi, \mathcal{O} \rrbracket_d(\mathbf{y}, t) > 0$, then $(\mathbf{y}, t, \mathcal{O}) \models \varphi$.
2. If $(\mathbf{y}, t, \mathcal{O}) \not\models \varphi$, then $\llbracket \varphi, \mathcal{O} \rrbracket_d(\mathbf{y}, t) \leq 0$. In other words, if \mathbf{y} does not satisfy the formula φ at time instant $t \geq 0$, then its distance valuation is non-positive. Conversely, if $\llbracket \varphi, \mathcal{O} \rrbracket_d(\mathbf{y}, t) < 0$, then $(\mathbf{y}, t, \mathcal{O}) \not\models \varphi$.
3. If for some $t \in \mathbb{R}^+$, $\varepsilon = \llbracket \varphi, \mathcal{O} \rrbracket_d(\mathbf{y}, t) \neq 0$, then for all $\mathbf{y}' \in B_{\rho_d}(\mathbf{y}, |\varepsilon|)$, we have $(\mathbf{y}, t, \mathcal{O}) \models \varphi$ if and only if $(\mathbf{y}', t, \mathcal{O}) \models \varphi$. I.e., ε defines a robustness tube around the trajectory such that other “nearby” trajectories lying inside this tube also satisfy φ .

Theorem 2.1 establishes the robust semantics of MTL as a natural measure of signal robustness. Namely, a signal is ε robust with respect to an MTL specification φ , if it can tolerate perturbations up to size ε and still maintain its current Boolean truth value. Alternatively, a signal with the opposite outcome for φ , if it exists, has a distance of at least ε away.

This is the main differentiating property from other works that also consider quantitative semantics for temporal logics such as [18, 39]. Namely, our semantics maintain the topological information which can be used to define neighborhoods for signals, while in quantitative or fuzzy semantics such information is lost. A more thorough comparison with other quantitative logics is provided in [26].

3 Falsifying Systems with Metric Output Spaces

In this section, we provide the basic formulation of MTL falsification as a global minimization of the robustness metric defined in Section 2 when the output space (Y, d) is a metric space, i.e., when $(Y, \mathbf{d}) = (Z, d)$, and describe a Monte-Carlo technique to solve this global optimization problem.

Let Σ be a system as defined in Section 2.1. Let φ be a given MTL property that we wish to falsify. Given a signal \mathbf{y} , we have defined a robustness metric $\llbracket \varphi, \mathcal{O} \rrbracket_d(\mathbf{y}, t)$ that denotes how robustly \mathbf{y} satisfies (or falsifies) φ from time t onwards. For the following discussion, we assume a fixed label map \mathcal{O} and always interpret the truth (and robustness) of MTL formulas evaluated at the starting time $t = 0$. Let $\mathcal{D}_\varphi(\mathbf{y}) = \llbracket \varphi, \mathcal{O} \rrbracket_d(\mathbf{y}, 0)$ denote the robustness metric for \mathbf{y} under these assumptions.

The robustness metric \mathcal{D}_φ maps each output signal \mathbf{y} to a real number r . The sign of r indicates whether $\mathbf{y} \models \varphi$ and its magnitude $|r|$ measures its robustness. Ideally, for the MTL verification problem, we would like to prove that $\inf_{\mathbf{y} \in \mathcal{L}(\Sigma)} \mathcal{D}_\varphi(\mathbf{y}) > \varepsilon > 0$ where ε is a desired robustness threshold. For the MTL falsification problem (Problem 2.1), we attempt to solve the problem:

$$\text{Find } \mathbf{y} \in \mathcal{L}(\Sigma) \text{ s.t. } \mathcal{D}_\varphi(\mathbf{y}) < 0 \quad (1)$$

More generally, given a robustness threshold $\varepsilon \geq 0$, we would like to solve the problem:

$$\text{Find } \mathbf{y} \in \mathcal{L}(\Sigma) \text{ s.t. } \mathcal{D}_\varphi(\mathbf{y}) < \varepsilon \quad (2)$$

In this work, we provide a solution to either problem through the optimization problem:

$$\mathbf{y}^* = \arg \min_{\mathbf{y} \in \mathcal{L}(\Sigma)} \mathcal{D}_\varphi(\mathbf{y}) \quad (3)$$

If $\mathcal{D}_\varphi(\mathbf{y}^*) < \varepsilon$, then we have produced a counterexample that can be used for debugging.

In the following, we provide parameterizations of the search space and a Monte-Carlo sampling method that will help us solve (3).

3.1 Autonomous Systems

The space of output signals is not the true search space for the falsification problem. For instance, it is hard to explore the space of trajectories directly while guaranteeing that each trajectory considered is valid. Fortunately, for deterministic systems, we may associate each input $x_0 \in X_0$ with a unique trajectory \mathbf{y} and vice-versa. Let $\mathcal{F}_\varphi(x_0) = \mathcal{D}_\varphi(\Delta_\Sigma(x_0))$ denote the robustness of the trajectory obtained corresponding to the initial state $x_0 \in X_0$. Therefore, the optimization can be expressed over the space of inputs as follows:

$$\min_{x_0 \in X_0} \mathcal{F}_\varphi(x_0) \quad (4)$$

The components of the vector x_0 are the search variables of the problem and the optimization is carried out subject to the constraints in X_0 .

Continuous trajectories are hard to compute precisely, even when the analytical form of the solution of the system is known. Thus, trajectories have to be *approximated* numerically. An approximate *simulation function* $\hat{\Delta}_\Sigma$ that

supports *robust evaluation* of the given property φ should guarantee that for some finite sampling \tilde{R} of the bounded time domain R , for $\tilde{\mathbf{y}} = \tilde{\Delta}_\Sigma(x_0)$ and for $\mathbf{y} = \Delta_\Sigma(x_0)$, $|\llbracket \phi, \mathcal{O} \rrbracket_d(\mathbf{y}, t) - \llbracket \phi, \mathcal{O} \rrbracket_d(\tilde{\mathbf{y}}, t)| \leq \epsilon$, for all $t \in \tilde{R}$, for a sufficiently small positive ϵ . Such a robust simulation function suffices, in practice, to resolve properties that may be of interest to the system designers. An appropriate simulation function can be obtained for a large class of ODEs using numerical simulation techniques of an appropriate order such as *Runge-Kutta* or Taylor-series methods with adaptive step sizes [50]. Numerical integration schemes can also be adapted to provide reliable bounds ϵ on the distance between the actual and the numerical solution. Thus, the robustness value $\mathcal{D}_\varphi(\mathbf{y})$ can be approximated by a value $\tilde{\mathcal{D}}_\varphi(\tilde{\mathbf{y}})$ using the set of sample points $\tilde{\mathbf{y}}$ obtained by a numerical integrator. Details on how $\tilde{\mathcal{D}}_\varphi(\tilde{\mathbf{y}})$ can be computed can be found in [26].

Unfortunately, for a trajectory $\tilde{\mathbf{y}}$ obtained as the output of a numerical integrator with known error bounds, the trace distance function may no longer satisfy $\tilde{\mathcal{D}}_\varphi(\tilde{\mathbf{y}}) \geq 0$ whenever $\mathbf{y} \models \varphi$. Instead, we may conclude the existence of some interval $[-\epsilon_2, \epsilon_1]$ for some $\epsilon_1, \epsilon_2 \geq 0$, such that if $\tilde{\mathcal{D}}_\varphi(\tilde{\mathbf{y}}) \leq -\epsilon_2$, then $\mathbf{y} \not\models \varphi$ and if $\tilde{\mathcal{D}}_\varphi(\tilde{\mathbf{y}}) \geq \epsilon_1$ then $\mathbf{y} \models \varphi$. In general, we may not draw any conclusions if $-\epsilon_1 \leq \tilde{\mathcal{D}}_\varphi(\tilde{\mathbf{y}}) \leq \epsilon_2$. Furthermore, the bounds ϵ_1, ϵ_2 are often unknown for a given system. Nevertheless, the presence of such a bound implies that it still makes sense to perform the optimization using a numerically simulated trajectory $\tilde{\mathbf{y}}$. Thus, our optimization problem becomes:

$$\min_{x_0 \in X_0} \mathcal{F}_\phi(x_0) = \min_{x_0 \in X_0} \tilde{\mathcal{D}}_\varphi(\tilde{\Delta}_\Sigma(x_0)). \quad (5)$$

In practice, even minimally “robust” simulated trajectories will often be of great interest to system designers even if mathematically speaking they do not violate the property under consideration.

Remark 3.1 *If the user is willing to tolerate additional computational cost, then it is possible to bound the inaccuracies of the numerical simulation even under the presence of floating-point errors [27]. Then, these bounds can be used to provide bounds on the robustness of the actual continuous-time trajectory [26].*

The resulting optimization problem (5) can be quite complex, unlikely to be convex for all but the simplest of cases. Furthermore, the objective function \mathcal{F} though computable for any given input through simulation, is not expressible in a closed form. Directly obtaining gradients, Hessians and so on is infeasible for all but the simplest of cases. We now present Monte-Carlo techniques that can solve such global optimization problems through a randomized technique that mimics gradient descent in many cases.

Remark 3.2 (Non-deterministic Systems) *For non-deterministic and stochastic systems, a single input can be associated with multiple (possibly infinitely many) behaviors. For stochastic systems, we may evaluate $\mathcal{F}_\varphi(x)$ as an expectation obtained by sampling a large but finite set of trajectories. Non-deterministic systems can often be determinized by adding new input variables to represent the non-deterministic choice. For the most part, we consider deterministic systems in this paper. For instance, systems modeled in formalisms such as Simulink/Stateflow diagrams (TM) are deterministic, at least in theory.*

3.2 Monte-Carlo Sampling

The Monte-Carlo techniques presented here are based on *acceptance-rejection* sampling [13, 7]. These techniques were first introduced in statistical physics, wherein, they were employed to simulate the behavior of particles in various potentials [30]. Variations of Monte-Carlo techniques are also widely used for solving global optimization problems [54]. In this paper, we focus on a class of Monte-Carlo sampling techniques known as Markov-Chain Monte-Carlo (MCMC) techniques. These techniques are based on random walks over a Markov chain that is defined over the space of inputs.

ALGORITHM 1: Monte-Carlo sampling algorithm.

Input: X_0 : Input Space, $f(\cdot)$: Robustness Function, ε : Robustness threshold, $\text{PS}(\cdot)$: Proposal Scheme

Output: $x \in X_0$

```
begin
1   Choose some initial input  $x \in X_0$ ;
2   while ( $f(x) \geq \varepsilon$ ) do
      /* Select  $x'$  using Prop. Scheme */
3      $x' \leftarrow \text{PS}(x)$ ;
4      $\alpha \leftarrow \exp(-\beta(f(x') - f(x)))$ ;
5      $r \leftarrow \text{UniformRandomReal}(0, 1)$ ;
6     if ( $r \leq \alpha$ ) then /* Accept proposal? */
7        $x \leftarrow x'$ ;
    end
end
```

We first present the basic sampling algorithm for drawing samples from a probability distribution and then the technique of *hit-and-run* sampling that respects the (convex) constraints on the input space due to X_0 . Let $f(x) = \mathcal{F}_\varphi(x)$ be a computable robustness function, given a property φ . We seek to minimize f over the inputs in the set X_0 . We wish to sample X_0 such that any two points $x, x' \in X_0$ with robustness values $f(x)$ and $f(x')$ are sampled with probability proportional to $\frac{e^{-\beta f_\varphi(x)}}{e^{-\beta f_\varphi(x')}}$, where $\beta > 0$ is a “temperature” parameter explained in the following.

Algorithm 1 shows the schematic implementation of the algorithm. Each iteration of the sampler generates a new *proposal* $x' \in X_0$ from the current sample x using some *proposal scheme* defined by the user (Line 3). The objective $f(x')$ is computed for this proposal. Subsequently, we compute the ratio $\alpha = e^{-\beta(f(x') - f(x))}$ (Line 4) and accept the proposal randomly, with probability α (Line 5). Note that if $\alpha \geq 1$ (i.e, $f(x') \leq f(x)$), then the proposal is accepted with certainty. Even if $f(x') > f(x)$ the proposal may still be accepted with some non-zero probability. If the proposal is accepted then x' becomes a new sample. Failing this, x remains the current sample. In general, MCMC techniques require the design of a *proposal scheme* for choosing a proposal x' given the current sample x . The convergence of the sampling to the underlying distribution defined by f , depends critically on the choice of this proposal distribution.

Proposal Scheme: A proposal scheme is generally defined by a probability distribution $P(x'|x)$ that specifies the probability of proposing a new sample input x' given the current sample x . In general, there are three requirements that a proposal scheme needs to satisfy so that its use in Algorithm 1 converges to the distribution defined by $f(x)$.

Detailed Balance $f(x')P(x'|x) = f(x)P(x|x')$ (see [13]).

Irreducibility Given any two inputs $x, x' \in X_0$, it should be possible with nonzero probability to generate a series of proposals x, x_1, x_2, \dots, x' that takes us from input x to x' . This is necessary in order to guarantee that the entire input state space can be covered.

Aperiodicity The greatest common divisor of the lengths of all non-zero probability cycles starting from a state x of the chain must be 1.

Convergence: Convergence of the sampling scheme guarantees that eventually after drawing a large but finite number of samples, the distribution of the samples approaches the distribution defined by the robustness function f . We will discuss convergence under the simplifying but practically relevant assumption of discreteness.

We assume that the space of inputs X_0 is *bounded and discrete*, consisting of a large but finite number of points. This assumption is always relevant in practice, since the inputs in X_0 that we consider are finitely represented floating point numbers inside a computer. As a result, the proposal scheme P defines a discrete Markov chain on the space of inputs. Convergence of MCMC sampling follows directly from the convergence of random walks on ergodic Markov Chains [51, 13, 54].

The robustness function $f(x)$ over X_0 induces a probability distribution:

$$p(x) = \frac{1}{M} e^{-\beta f(x)},$$

ht




Figure 4: Hit-and-run proposal scheme.

where M is an unknown normalizing factor added to ensure that the probabilities add up to one. Suppose Algorithm 1 were run to generate a large number of samples N . Let γ denote the frequency function mapping subsets of the input space to the number of times a sample was drawn from the set. Let $P(S) = \sum_{x \in S} p(x)$ denote the volume of the probability function for a set $S \subseteq X_0$.

Theorem 3.1 *In the limit, the acceptance rejection sampling technique (almost surely) generates samples according to the distribution p , $P(S) = \lim_{N \rightarrow \infty} \frac{\gamma(S)}{N}$*

As a direct consequence, one may conclude, for instance, that an input x_1 with $f(x_1) = -100$ is *more likely* to be sampled as compared to some other input x_2 with $f(x_2) = 100$ in the *long run*.

It is possible, in theory, to prove assertions about the number N of samples required for the sampled distribution to converge within some distance to the desired distribution governed by $e^{-\beta f_\varphi(x)}$. This rate of convergence is governed by the *mixing time* of the *Markov chain* on the inputs defined by the proposal scheme. This time is invariably large (polynomial in the number of input points), and depends on the proposal scheme used [51].

Importance of β : The overall algorithm itself can be seen as a *randomized* gradient descent, wherein at each step a new point x' in the search space is compared against the current sample. The probability of moving the search to the new point follows an exponential distribution on the difference in their robustness values: $p \sim e^{-\beta(f(x')-f(x))}$. In particular, if $f(x') \leq f(x)$, the new sample is accepted with certainty. Otherwise, it is accepted with probability $e^{-\beta(f(x')-f(x))}$. Informally, larger values of β ensure that only reductions to $f(x)$ are accepted whereas smaller values correspondingly increase the probability of accepting an increase in $f(x)$. As a result, points with lower values of f are sampled with an exponentially higher probability as compared to points with a higher value of the function f .

Adapting β . One of the main drawbacks of Algorithm 1 is that, based on nature of the distribution, the sampling may get “trapped” in *local minima*. This typically results in numerous proposals getting rejected and few being accepted. Even though we are guaranteed eventual convergence, the presence of local minima slows down this process, in practice. We therefore periodically adjust the values of β (and also the proposal scheme) to ensure that the ratio of accepted samples vs. rejected samples remains close to a fixed value (1 in our experiments). This is achieved by monitoring the acceptance ratio during the sampling process and adjusting β based on the acceptance ratio. A high acceptance ratio indicates that β needs to be increased, while a low acceptance rate indicates that β needs to be reduced.

Proposal Schemes. It is relatively simple to arrive at viable schemes for generating new proposals. However, designing a scheme that works well for the underlying problem requires a process of experimentation. For instance, it suffices to simply choose an input x' uniformly at random from the inputs, regardless of the current sample. However, such a scheme does not provide many advantages over uniform random sampling. In principle, given a current sample x , the choice of the next sample x' must depend on x .

A typical proposal scheme samples from a normal distribution centered at x with a suitably adjusted standard deviation (using some covariance matrix H). The covariance can be adjusted periodically based, once again, on the observed samples as well as the acceptance ratio. A smaller standard deviation around x yields samples whose




Figure 5: (a) Time trajectory violating the property $\square_{[0,2]} \neg a$, where $\mathcal{O}(a) = [-1.6, -1.4] \times [-.9, -1.1]$ along with the scatter plot of sampled inputs and (b) robustness value as a function of the simulation step number.

robustness values differ very little from $f(x)$, thus increasing the acceptance ratio. However, it is hard to respect the constraint $x' \in X_0$ using such a proposal scheme.

Hit-and-run proposal scheme. Hit-and-run schemes are useful in the presence of input domains such as $X_0 \subseteq \mathbb{R}^n$. For simplicity, we assume that X_0 is convex. Therefore, any line segment in some direction v starting from x has a maximum offset δ_M such that the entire segment between x and $x + \delta v$ lies inside X_0 . At each step, we propose a new sample x' based on the current sample x . This is done in three steps:

1. Choose a random unit vector v uniformly (or using a Gaussian distribution) (Cf. Fig. 4). In practice, one may choose a random vector h and generate a unit vector using $v = \frac{h}{|h|_2}$.
2. Discover the interval $[\delta_m, \delta_M]$, such that

$$\forall \delta \in [\delta_m, \delta_M], x + \delta v \in X_0.$$

In other words, v yields a line segment containing the point x along the directions $\pm v$ and $[\delta_m, \delta_M]$ represent the minimum and maximum offsets possible along the direction v starting from x . If X_0 is a polyhedron, bounds $[\delta_m, \delta_M]$ may be obtained efficiently by using a variant of the *minimum ratio test*. For a more complex convex set X_0 , value of δ_m (resp. δ_M) may be obtained by solving the one dimensional optimization problem $\min(\max) \delta$ s.t. $x + \delta v \in X_0$, by using a *bisection* procedure given an initial guess on $[\delta_m, \delta_M]$.

3. Finally, we choose a value $\delta \in [\delta_m, \delta_M]$ based on some probability distribution with a mean around 0. The variance of this distribution is an important parameter that can be used to control the acceptance ratio (along with β) to accelerate convergence.

Hit-and-run samplers can also be used for non-convex input domains such as unions of polytopes and so on. A detailed description of the theory behind such sampling techniques is available elsewhere [59, 54].

However, care must be taken to ensure that the input space X_0 is not *skewed* along some direction v' . In the worst case, we may imagine X_0 as a straight line segment. In such cases, the hit-and-run proposal scheme fails to generate new samples. This is remedied by adjusting the scheme for selecting unit directions to take the skew of X_0 , embedding of X_0 inside a subspace spanned by the independent variables and, finally, applying a suitable transformation to X_0 that aids in sampling.

In practice, hit and run samplers can work over non-convex, disconnected domains. Theoretical results on these samplers are very promising. Smith [58] proves the asymptotic convergence of hit and run sampling over arbitrary open subsets of \mathbb{R}^n . Lovasz [42, 43] has further demonstrated convergence in time $O(n^3)$ for hit and run sampling of uniform distribution over a convex body in n dimensions. Algorithms for global optimization such as *hide-and-peek* [53] and *improving hit-and-run* [63] have combined hit-and-run sampling with Monte-Carlo to generate global optimization techniques.

Example 3.1 Let $\mathbf{y}(t) = [\mathbf{y}_1(t) \ \mathbf{y}_2(t)]^T$. Consider the time varying system

$$\frac{d\mathbf{y}(t)}{dt} = \begin{bmatrix} \frac{d\mathbf{y}_1(t)}{dt} \\ \frac{d\mathbf{y}_2(t)}{dt} \end{bmatrix} = \begin{bmatrix} \mathbf{y}_1(t) - \mathbf{y}_2(t) + 0.1t \\ \mathbf{y}_2(t) \cos(2\pi\mathbf{y}_2(t)) - \mathbf{y}_1(t) \sin(2\pi\mathbf{y}_1(t)) + 0.1t \end{bmatrix}$$

with initial condition $\mathbf{y}(0) = x_0 \in X_0 = [-1, 1] \times [-1, 1]$. In this case, $Y = \mathbb{R}^2$ and, thus, we choose to use the Euclidean metric. We wish to falsify the property $\square_{[0,2]} \neg a$, wherein $\mathcal{O}(a) = [-1.6, -1.4] \times [-.9, -1.1]$. Our simulation uses a numerical ODE solver with a fixed time step over the time interval $t \in R = [0, 2]$. Figure 5(a) shows the trajectory the falsifies our safety property using the hit-and-run sampler and the scatter plot consisting of the samples generated by the Monte-Carlo sampler. Figure 5(b) plots the robustness of the trajectory at each simulation step. We observe that the sampling is concentrated in the more promising regions in the set of initial conditions. \diamond

3.3 Non-autonomous Systems

We now consider extensions to non-autonomous CPS. Again, for pragmatic reasons, we focus on the approximation $\tilde{\mathbf{y}} = \tilde{\Delta}_\Sigma(x_0, \tilde{u})$ of the actual trajectory $\mathbf{y} = \Delta_\Sigma(x_0, u)$. Here, the input signal \tilde{u} is a discrete-time approximation of the actual continuous-time input signal u . Therefore, in a naive search for a falsifying input signal, we may consider each sampling instance as a search variable. However, such an approach is infeasible for long simulation times with fast sampling rates.

Our goal is to recast the search for control input signals \tilde{u} in terms of a search in the set of parameters $\lambda \in \Lambda$ and $\tau \in R^m$, where $m \ll |\tilde{R}|$, i.e., m is substantially smaller than the number of samples from R . Since we have assumed that the input signal space can be parameterized on λ and τ , we can produce a discrete-time approximation $\tilde{u} = \tilde{\mathcal{U}}(\lambda, \tau)$ to $u = \mathcal{U}(\lambda, \tau)$ and, thus, we are able to represent realistic input signals. Now, our optimization problem becomes:

$$\min_{\langle x_0, \lambda, \tau \rangle \in X_0 \times \Lambda \times R^m} f(x_0, \lambda, \tau) = \min_{\langle x_0, \lambda, \tau \rangle \in X_0 \times \Lambda \times R^m} \tilde{D}_\varphi(\tilde{\Delta}_\Sigma(x_0, \tilde{\mathcal{U}}(\lambda, \tau))). \quad (6)$$

In practical terms, there exist numerous ways to parameterize the space of control inputs. We discuss a few such parameterizations below:

Piece-wise Constant Input: We partition the overall time interval $R = [0, T]$ into a set of intervals $\bigcup_{i=1}^m [\tau_{i-1}, \tau_i)$, wherein $\tau_0 = 0$ and $\tau_m = T$. For each interval $[\tau_{i-1}, \tau_i)$, $i \geq 1$, the control $u(t)$ is restricted to be a constant value λ_{i-1} .

Piece-wise Linear Input: Piece-wise constant control may be extended to piecewise linear controls. Once again, we partition $R = [0, T]$ into m disjoint intervals. For each interval $[\tau_{i-1}, \tau_i]$, we restrict the form of each control input to be piece-wise linear, i.e., for $t \in [t_{i-1}, t_i)$, we have $u(t) = (1 - \alpha(t))\lambda_{i-1} + \alpha(t)\lambda_i$ where $\alpha(t) = (-\tau_{i-1}) / (\tau_i - \tau_{i-1})$.

Spline Functions. We can choose a family of spline functions $\mathcal{U}_S(\lambda, \tau)$. Details on utilizing splines to represent control input signals can be found in [20].

Example 3.2 In order to parameterize the input signal space of Example 2.1, we used a piece-wise constant signal with 7 control points uniformly distributed over the time domain $[0, 30]$. That is, our search for a minima is performed over a bounded 7 dimensional space. Furthermore, since the output space Y is \mathbb{R}^2 , we are using the Euclidean metric for the distance computations in the formula defined in Example 2.2. The outcome of S-TALIRO appears in Fig. 2. As evident from the figure, the vehicle speed and the engine rotation indeed reach the specified thresholds. The Simulink model was simulated 41 times for this particular test. \diamond

4 Falsifying Systems with Generalized Quasi-Metric Output Spaces

In the previous sections, we demonstrated that MTL falsification of systems is possible as long as we can define a non-trivial metric on the output space. However, specifications on CPS usually have requirements on both the

discrete output space of the system and the continuous output space. Unfortunately, it is not straightforward to define metrics over such hybrid (discrete & continuous) output spaces. Therefore, in order to formulate and analyze such specifications, we need to relax our constraint on the system having metric output spaces.

Example 4.1 *Let us revisit Example 2.1. We are looking to generate tests such that the system visits each state in the state chart `selection_state` (see Fig. 2), i.e., `steady_state`, `upshifting` and `downshifting` when the vehicle speed exceeds 79. In this case, the output trajectory \mathbf{y} of the system model must not only contain information about the physical system quantities, i.e., engine rotation and vehicle speed, but also about the current state in the Stateflow chart. Therefore, the temporal logic analysis must be performed over the output space $Y = \{\text{steady_state}, \text{upshifting}, \text{downshifting}\} \times \mathbb{R}^2$. \diamond*

In this section, we first generalize Theorem 2.5 to signals over generalized quasi-metric output spaces. Then, we introduce a modeling formalism for hybrid systems and two interesting generalized quasi-metrics on output trajectories of such hybrid systems.

4.1 Robustness of Signals over Generalized Quasi-Metrics Spaces

The only requirement in the definition of the robust semantics of MTL formulas (Section 2.3) is that both the trajectory under study and the specifications take values from the same space. We can prove (see Appendix) by induction on the structure of formula φ that Theorem 2.1 also holds in the case where the metric d is replaced by a generalized quasi-metric \mathbf{d} .

Theorem 4.1 *Given an output space (Y, \mathbf{d}) , where \mathbf{d} is an extended generalized quasi-metric, a formula $\varphi \in MTL$, an observation map $\mathcal{O} \in \mathcal{P}(Y)^{AP}$ and an output signal $\mathbf{y} \in Y^R$, then:*

1. *If $(\mathbf{y}, t, \mathcal{O}) \models \varphi$, then $\llbracket \varphi, \mathcal{O} \rrbracket_{\mathbf{d}}(\mathbf{y}, t) \succeq \mathbf{0}$. Conversely, if $\llbracket \varphi, \mathcal{O} \rrbracket_{\mathbf{d}}(\mathbf{y}, t) \succ \mathbf{0}$, then $(\mathbf{y}, t, \mathcal{O}) \models \varphi$.*
2. *If $(\mathbf{y}, t, \mathcal{O}) \not\models \varphi$, then $\llbracket \varphi, \mathcal{O} \rrbracket_{\mathbf{d}}(\mathbf{y}, t) \preceq \mathbf{0}$. Conversely, if $\llbracket \varphi, \mathcal{O} \rrbracket_{\mathbf{d}}(\mathbf{y}, t) \prec \mathbf{0}$, then $(\mathbf{y}, t, \mathcal{O}) \not\models \varphi$.*
3. *If for some time $t \in R$, $\varepsilon = \llbracket \varphi, \mathcal{O} \rrbracket_{\mathbf{d}}(\mathbf{y}, t) \neq \mathbf{0}$, then for all $\mathbf{y}' \in \mathcal{B}_{\rho_{\mathbf{d}}}(\mathbf{y}, |\varepsilon|)$, we have $(\mathbf{y}, t, \mathcal{O}) \models \varphi$ if and only if $(\mathbf{y}', t, \mathcal{O}) \models \varphi$.*

Note that now the definition of the robustness valuation function for a formula φ over a signal \mathbf{y} at time t is a function $\llbracket \cdot, \cdot \rrbracket_{\mathbf{d}} : (MTL \times \mathcal{P}(Y)^{AP}) \rightarrow (Y^R \times R \rightarrow \mathbb{V})$. The set \mathbb{V} must include the set V of the positively ordered monoid $(V, +, \preceq)$ in the definition of the generalized quasi-metric \mathbf{d} and, also, it must be ordered under the same ordering relation \preceq . Furthermore, appropriate definitions of negation and absolute value are required as well as careful treatment of the absorbing elements (if any). Essentially, we need $(\mathbb{V}, +, \preceq)$ to be an Abelian group with two absorbing elements $\pm\infty$.

4.2 Generalized Quasi-Metrics for Hybrid Signals

In order to define quasi-metrics for hybrid signals, we need to take into account some information about the structure of the system that generates the output signals. Here, we will be using a generalization of hybrid automata [3] as a basic modeling language for CPS. We remark that our formalism resembles hierarchical hybrid systems [4].

Def. 4.1 (Hybrid System) *A hybrid system \mathcal{H} consists of components $\langle H, H_0, Y, \mathbf{U}, \mathbf{O}, \mathbf{G}, \mathbf{R}, \mathbf{D}, \rightarrow \rangle$, wherein,*

- *$H = L \times X$ is the state space of the system and L is a finite set of locations (modes or control locations),*
- *$H_0 \subseteq H$ represents the set of initial conditions,*
- *$Y = L \times Z$ is the output space, where (Z, d) is a metric space*
- *$\mathbf{U} \subseteq U^R$ is the set of possible input signals,*

- $\mathbf{O} : X \rightarrow Z$ is an output map,
- $\mathbf{G} : L \times L \times X \times U \rightarrow \mathcal{P}(Z)$ is the guard set for the transitions between control locations,
- $\mathbf{R} : (L \times L) \rightarrow (X \rightarrow X)$ is the reset function for the transitions between control locations,
- $\mathbf{D} : L \rightarrow (X \times R \times U \rightarrow X^R)$ is a mapping of each control location $\ell \in L$ to a deterministic subsystem, which given an initial condition x_0 , an initial time t_0 and an input signal u , returns the unique state trajectory of the subsystem $\mathbf{x}_\ell = \mathbf{D}_\ell(x_0, t_0, u)$, and, finally,
- $\rightarrow \subseteq L \times L$ is a set of (discrete) transitions such that for each $\langle \ell_1, \ell_2 \rangle \in \rightarrow$, i.e., $\ell_1 \rightarrow \ell_2$, the system moves from $\ell_1 \in L$ to $\ell_2 \in L$ if the output state $z = \mathbf{O}(x)$ of the system before the transition is in the set $\mathbf{G}(\ell_1, \ell_2, x, u)$ and after the transition is at the point $z' = \mathbf{O}(x')$ where $x' = \mathbf{R}_{\langle \ell_1, \ell_2 \rangle}(x)$.

We remark that our definition of a hybrid system allows each control location to be any arbitrary subsystem as long as it is deterministic and its state can be fully described by the function \mathbf{D}_ℓ . For example, each control location can be a hybrid system as well. The reason behind utilizing such a general model is that we are not necessarily interested in the whole structure of the hybrid system, but only on its part that is directly related to the functional specification that we are trying to falsify.

Example 4.2 *The Simulink/Stateflow model in Example 2.1 has state-space*

$$\{\text{first, second, third, fourth}\} \times \{\text{steady_state, upshifting, downshifting}\} \times \mathbb{R}^2.$$

In Example 4.1, the specification requirements focus only on the state chart `selection_state`. Therefore, our hybrid system will have the following components of interest:

- $L = \{\text{steady_state, upshifting, downshifting}\}$ and \rightarrow as defined in Fig. 2.
- $X = \{\text{first, second, third, fourth}\} \times \mathbb{R}^2$, $Z = \mathbb{R}^2$ and \mathbf{O} is the projection of X on \mathbb{R}^2 .

The reset function \mathbf{R} changes the state of the state chart `gear_state` and the guard \mathbf{G} is computed by the `ThresholdCalculation` block in the Simulink model in Fig. 1. Note that \mathbf{G} depends on x and u . However, we are not interested in the components \mathbf{R} and \mathbf{G} in this example. \diamond

A *timed trace* of a hybrid system is a *finite*² sequence of states $\langle t, \ell, x \rangle \in R \times L \times X$ of the form $\langle t_0, \ell_0, x_0 \rangle, \langle t_1, \ell_1, x_1 \rangle, \langle t_2, \ell_2, x_2 \rangle, \dots$, such that *initially*, at time t_0 , we have $\langle \ell_0, x_0 \rangle \in H_0$, and for each consecutive state pair $\langle t_i, \ell_i, x_i \rangle$, we

- either make discrete transition from ℓ_i to ℓ_{i+1} and set $x_{i+1} = \mathbf{R}_{\langle \ell_i, \ell_{i+1} \rangle}(x_i)$
- or we evolve under the subsystem \mathbf{D}_{ℓ_i} from x_i to x_{i+1} , i.e., $x_{i+1} = \mathbf{D}_{\ell_i}(x_i, t_i, u)(t_{i+1})$.

A hybrid system \mathcal{H} is *deterministic* iff starting from some initial state $\langle t_0, \ell_0, x_0 \rangle$ there exists a unique timed trace. Given a timed trace, we can construct a hybrid system trajectory $\mathbf{y} : R \rightarrow Y$ by setting $\mathbf{y}(t) = \langle \mathbf{l}(t), \mathbf{z}(t) \rangle$ for $t \in [t_i, t_{i+1})$, where $\mathbf{l}(t) = \ell_i$ and $\mathbf{z}(t) = \mathbf{O}(\mathbf{x}(t))$ with $\mathbf{x}(t) = \mathbf{D}_{\ell_i}(x_i, t_i, u)(t)$. Therefore, again, we may view a hybrid system as a function $\Delta_{\mathcal{H}}$ from the set of initial conditions H_0 and the input signals U to output signals Y^R .

Let $\tilde{\Delta}_{\mathcal{H}}(h_0, \tilde{u})$ represent the approximate simulation function for a deterministic hybrid system \mathcal{H} . We assume that $\tilde{\Delta}_{\mathcal{H}}(h_0, \tilde{u})$ approximates the time trajectories with some given tolerance bound ϵ by adjusting the integration method. In practice, this may be harder to achieve for hybrid systems than for purely continuous systems due to the problem of robust event detection [22]. However, assuming that such a simulator is available (see [55] for conditions), we may translate the trace fitness function defined for continuous simulations to hybrid simulations with discrete transitions.

Specifications for hybrid automata involve a sequence of locations of the discrete subsystem. The simplest such property being the (un)reachability of a given “error” location. As a result, continuous state distance based on a norm

²Again, we implicitly assume that the system does not exhibit Zeno behaviors [44].

(or a metric distance) does not, in general, provide a true notion of distance between the specification and the trace. This is especially true in the presence of discrete transitions with reset maps. For the case of hybrid systems with reset maps, the robustness metrics used in Section 3 cannot be used to compare the hybrid states (ℓ, z) and (ℓ', z') in terms of some norm distance between z and z' . Therefore, structural considerations based on the graph that connects the different modes of the hybrid automata have to be considered while designing fitness functions. We now consider (generalized quasi-) metrics for hybrid automata.

First, we have to define what is the distance between two modes of the hybrid automaton. We claim that a reasonable metric is the *shortest path distance* between two locations. A similar metric was used for guiding the exploration in a model checker for hybrid systems in [6]. Intuitively, the shortest path distance provides us with a measure of how close we are to a desirable or undesirable operating mode of the automaton. Such information is especially useful in the class of falsification algorithms that we consider in this paper.

In the following, given hybrid automaton \mathcal{H} , we let $\Gamma(\mathcal{H}) = (L, \rightarrow)$ represent the directed graph formed by its discrete modes and transitions. The shortest path distance from node ℓ to node ℓ' in the graph $\Gamma(\mathcal{H})$ will be denoted by $\pi(\ell, \ell')$. Note that $\pi(\ell, \ell') = \infty$ iff there is no path from ℓ to ℓ' in the graph $\Gamma(\mathcal{H})$. It is well known (and it is easy to verify) that the shortest path distance satisfies all the criteria for a quasi-metric.

The shortest path metric can be computed on-the-fly by running a Breadth First Search (BFS) [15] algorithm on the graph. It is well known that BFS runs in linear time on the size of the input graph. However, it is preferable to use an all-pairs shortest path algorithm [15] to precompute the distances between all pairs of control locations of the hybrid automaton. In our implementation, we are using the Floyd-Warshall algorithm which has running time $\Theta(|L|^3)$.

In order to reason over output trajectories \mathbf{y} in the hybrid state space Y , we need to introduce a generalized distance function [57]. In the following, we will denote the hybrid space $L \times Z$ by \mathbb{H} to indicate that a metric is defined over a particular space. Let $\mathbf{d}_h : \mathbb{H} \times \mathbb{H} \rightarrow \mathbb{V}_+^\infty$, where $\mathbb{V}_+^\infty = \mathbb{V}_+ \cup \{+\infty, +\infty\}$ and $\mathbb{V}_+ = \mathbb{N} \times \mathbb{R}_+$, with definition for $h = \langle \ell, z \rangle \in \mathbb{H}$ and $h' = \langle \ell', z' \rangle \in \mathbb{H}$,

$$\mathbf{d}_h(h, h') = \begin{cases} \langle 0, d(z, z') \rangle & \text{if } \ell = \ell' \\ \langle \pi(\ell, \ell'), \min_{\ell'' \in \partial \mathcal{N}_\pi(\ell, \ell')} \mathbf{dist}_d(z, \mathbf{G}^t(\ell, \ell'')) \rangle & \text{otherwise} \end{cases}$$

where π is the shortest path metric, d is a metric on Z and $\partial \mathcal{N}_\pi(\ell, \ell') = \mathbf{Nxt}(\ell) \cap \mathcal{N}_\pi(\ell', \pi(\ell, \ell'))$. Here, $\mathbf{Nxt}(\ell) = \{\ell' \in L \mid \ell \rightarrow \ell'\}$ and \mathbf{G}^t denotes that the guard set may be changing with respect to time. Informally, $\partial \mathcal{N}_\pi(\ell, \ell')$ is the “boundary” of all locations which are closer to ℓ' than ℓ and may be visited from ℓ within one transition. Therefore, when the two points h, h' are in the same control location, then the distance computation reduces to the distance computation between the points in the continuous state space. When the two points h, h' are in different control locations, then the distance is the path distance between the two control locations “weighted” by the distance to the closest guard that will enable the transition to the next control location that reduces the path distance. Essentially, the last condition is a heuristic that gives preference to shortest paths.

Next, we need to define an appropriate addition $+$ and a partial order \preceq such that the triplet $(\mathbb{V}_+, +, \preceq)$ is a positively ordered commutative monoid. First, the addition is defined component-wise, that is, for $\langle k, r \rangle, \langle k', r' \rangle \in \mathbb{V}_+$, we define

$$\langle k, r \rangle + \langle k', r' \rangle = \langle k + k', r + r' \rangle$$

The commutativity property is immediately satisfied. Second, we order the set using the dictionary order. Given $\langle k, r \rangle, \langle k', r' \rangle \in \mathbb{Z}_\infty \times \overline{\mathbb{R}}$, we define the order relation \prec as

$$\langle k, r \rangle \prec \langle k', r' \rangle \text{ iff } \begin{cases} k < k' & \text{if } k \neq k' \\ r < r' & \text{if } k = k' \end{cases}$$

It is easy to verify that the dictionary order is compatible with the addition as defined for \mathbb{V}_+ . Hence, \mathbb{V}_+ has a smallest element, namely $\mathbf{0} = \langle 0, 0 \rangle$, and \mathbb{V}_+^∞ has an absorbing element, namely $+\infty = \langle +\infty, +\infty \rangle$, which is also the least upper bound. Finlay, Proposition A.1 in the Appendix demonstrates that the generalized distance \mathbf{d}_h satisfies the identity and triangle inequality properties. In other words, \mathbf{d}_h is a generalized quasi-metric on \mathbb{H} .

The generalized distance function \mathbf{d}_h requires computations of a point to each guard set in a control location. This may potentially increase the computational load or it could be the case that the computation of the distance to the

ALGORITHM 2: Parallel Monte-Carlo sampling algorithm.

Input: $H_0 \times \Lambda \times R^m$: Input Space, $f(\cdot)$: Robustness Function, ε : Robustness Threshold,

$\text{PS}(\cdot)$: Proposal Scheme

Output: $\langle h, \lambda, \tau \rangle \in H_0 \times \Lambda \times R^m$

begin

```
1   Choose some initial input  $\langle h, \lambda, \tau \rangle \in H_0 \times \Lambda \times R^m$ ;  
2   while ( $f(h, \lambda, \tau) \geq \varepsilon$ ) do  
    /* Select  $\langle h', \lambda', \tau' \rangle$  using the Proposal Scheme */  
3    $\langle h', \lambda', \tau' \rangle \leftarrow \text{PS}(\langle h, \lambda, \tau \rangle)$ ;  
4    $\alpha_1 \leftarrow \exp(-\beta_1(f_1(h', \lambda', \tau') - f_1(h, \lambda, \tau)))$ ;  
5    $\alpha_2 \leftarrow \exp(-\beta_2(f_2(h', \lambda', \tau') - f_2(h, \lambda, \tau)))$ ;  
6    $r \leftarrow \text{UniformRandomReal}(0, 1)$ ;  
7   if ( $((f_1(h', \lambda', \tau') = f_1(h, \lambda, \tau)) \wedge (r \leq \alpha_2)) \vee ((f_1(h', \lambda', \tau') \neq f_1(h, \lambda, \tau)) \wedge (r \leq \alpha_1)))$  then  
8      $\langle h, \lambda, \tau \rangle \leftarrow \langle h', \lambda', \tau' \rangle$ ;  
   end  
end
```

guard might not be possible (for example, in certain Simulink/Stateflow models). Therefore, we also introduce the generalized distance function $\mathbf{d}_h^0 : \mathbb{H} \times \mathbb{H} \rightarrow \mathbb{V}_+^\infty$ with definition

$$\mathbf{d}_h^0(h, h') = \begin{cases} \langle 0, d(z, z') \rangle & \text{if } \ell = \ell' \\ \langle \pi(\ell, \ell'), 0 \rangle & \text{if } \ell \neq \ell' \text{ and } \pi(\ell, \ell') < +\infty \\ \langle +\infty, +\infty \rangle & \text{otherwise} \end{cases}$$

In this case, the distance function ignores the guard sets and simply checks whether the 2 points are in the same control location or not. The distance function \mathbf{d}_h^0 is a generalized quasi-metric as well.

Therefore, we are in position to reason about hybrid system trajectories by utilizing the MTL robustness Definition 2.5 and Theorem 4.1. Now the atomic propositions can map to subsets of \mathbb{H} placing, thus, requirements not only on the continuous state-space, but also on the mode of the hybrid system. Informally, a robustness value of $\langle k, r \rangle = \llbracket \varphi, \mathcal{O} \rrbracket_{\mathbf{d}_h}(y, t)$ will mean the following:

- If $k = 0$ and $r \neq 0$, then we can place a tube of radius $|r|$ around the continuous part of the trajectory which will guarantee equivalence under the MTL formula. Moreover, it is required that at each point in time t , the locations are the same for all such trajectories.
- If $k > 0$, then the specification is satisfied and, moreover, the trajectory is at least k discrete transitions away from being falsified.
- If $k < 0$, then the specification is falsified and, moreover, the trajectory is at least k discrete transitions away from being satisfied.

Remark 4.1 Note that both functions $\mathbf{Dist}_{\mathbf{d}_h}$ and $\mathbf{Dist}_{\mathbf{d}_h^0}$ never evaluate to some value of the form $\langle k, \pm\infty \rangle$ with $k \in \mathbb{Z}$ (see Proposition A.3). This is important because the temporal logic robustness value is now going to be a member of the set $\mathbb{V}^\infty = \mathbb{V} \cup \{\pm\infty\}$ where $\mathbb{V} = \mathbb{Z} \times \mathbb{R}$. In order for the triplet $(\mathbb{V}, +, \preceq)$ to be an ordered Abelian group and, thus, the robust MTL semantics to have a proper definition of negation, each member of \mathbb{V} must have an inverse. The negation for the MTL robust semantics induced by the aforementioned metrics is simply the pairwise negation. In Proposition A.3, we also demonstrate how the distance functions $\mathbf{Dist}_{\mathbf{d}_h}$ and $\mathbf{Dist}_{\mathbf{d}_h^0}$ can be computed based on the well known understood distance functions \mathbf{Dist}_π and \mathbf{Dist}_d .

4.3 Monte Carlo Sampling

One of the issues that arise when giving generalized (or “hybrid”) robust semantics to MTL formulas is how to sample over the space $H_0 \times \Lambda \times R^m$. Recall that $\Lambda \times R^m$ is the space of parameters that parameterize the input signals. In

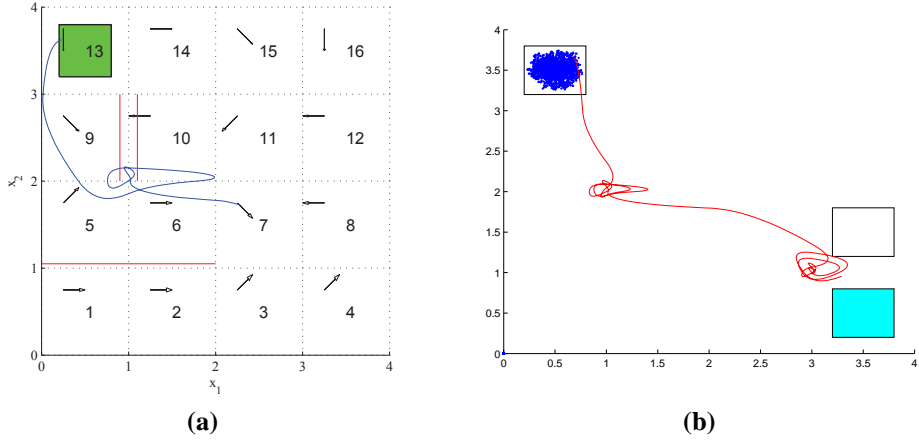


Figure 6: Example 4.3: (a) The environment of the vehicle benchmark. The arrows indicate the direction of the vector field in each location and the numbers the id of each location. The green box indicates the set of initial conditions projected on the position plane. The plotted trajectory falsifies the specification $\varphi_2 = \square(p_{21} \rightarrow (p_{22}\mathcal{R}\neg p_{23}))$. The red lines indicate the predicates in φ_2 . (b) The scatter plot of the sampled initial positions projected on the position plane, along with the least robust trajectory of the vehicle benchmark example for specification $\varphi_1 = (\neg p_{11}) \mathcal{U}_{[0,25.0]} p_{12}$.

other words, what is the probability distribution induced by the robustness function f ? In general, this issue can only be addressed in a case-by-case scenario depending on the generalized metric \mathbf{d} that is utilized.

In this work, for the generalized quasi-metric \mathbf{d}_h , we propose to use a Parallel Metropolis coupled Markov chain Monte Carlo algorithm (see Algorithm 2). For a point $\langle h_0, \lambda, \tau \rangle \in H_0 \times \Lambda \times R^m$, the robustness function is now $f(h_0, \lambda, \tau) = \mathcal{D}_\varphi(\hat{\Delta}_{\mathcal{H}}(h_0, \tilde{\mathbf{u}}(\lambda, \tau)))$. If $f(h_0, \lambda, \tau) = \langle k, r \rangle \in \mathbb{V}^\infty$, then we define $f_1(h_0, \lambda, \tau) = k \in \mathbb{Z}_\infty$ and $f_2(h_0, \lambda, \tau) = r \in \mathbb{R}$. In brief, in Algorithm 2, an input $\langle h_1, \lambda_1, \tau_1 \rangle$ will be more likely sampled over an input $\langle h_2, \lambda_2, \tau_2 \rangle$, if $f_1(h_1, \lambda_1, \tau_1) = f_1(h_2, \lambda_2, \tau_2)$ and $f_2(h_1, \lambda_1, \tau_1) << f_2(h_2, \lambda_2, \tau_2)$, or, if $f_1(h_1, \lambda_1, \tau_1) \neq f_1(h_2, \lambda_2, \tau_2)$ and $f_1(h_1, \lambda_1, \tau_1) << f_1(h_2, \lambda_2, \tau_2)$. The discussion in Section 3.2 on the importance of β and the proposal schemes still applies. Similarly, we can define a sampling algorithm for the metric \mathbf{d}_h^0 .

4.4 Examples

The new hybrid notion of MTL robustness is useful in the context of testing for hybrid systems. Namely, our original definition of MTL robustness places requirements only on the observable continuous-time trajectories of the system while ignoring the underlying discrete dynamics. The new robustness notion can structurally distinguish system trajectories that might have similar robustness values otherwise. Thus, it can be used to guide our search algorithms towards less robust system modes. Moreover, we can now impose different requirements at different operating modes of the system. This was not possible before.

Example 4.3 Consider a complex instance of the vehicle benchmark [28] shown in Fig. 6. The benchmark studies a hybrid automaton \mathcal{H} with 4×4 discrete locations and 4 continuous variables x_1, x_2, x_3, x_4 that form the state vector $x = [x_1 \ x_2 \ x_3 \ x_4]^T$. We refer to the vectors $[x_1 \ x_2]^T$ and $[x_3 \ x_4]^T$ as the position and the velocity of the system, respectively. The structure of the hybrid automaton can be better visualized in Fig. 6. The invariant set of every (i, j) location is an 1×1 box that constraints the position of the system, while the velocity can flow unconstrained. The guards in each location are the edges and the vertices that are common among the neighboring locations.

Each location has affine constant dynamics with drift. In detail, in each location (i, j) of the hybrid automaton,

the system evolves under the differential equation $\dot{x} = Ax - Bu(i, j)$ where

$$u(i, j) = [\sin(\pi C(i, j)/4) \quad \cos(\pi C(i, j)/4)]^T \text{ and}$$

$$A = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -1.2 & 0.1 \\ 0 & 0 & 0.1 & -1.2 \end{bmatrix} \quad B = \begin{bmatrix} 0 & 0 \\ -1.2 & 0.1 \\ 0.1 & -1.2 \end{bmatrix} \quad C = \begin{bmatrix} 4 & 2 & 3 & 4 \\ 3 & 6 & 5 & 6 \\ 1 & 2 & 3 & 6 \\ 2 & 2 & 1 & 1 \end{bmatrix}$$

Our first goal is to find an initial state in the set $H_0 = \{13\} \times [0.2, 0.8] \times [3.2, 3.8] \times [-0.4, 0.4] \times [-0.4, 0.4]$ that will falsify the formula $\varphi_1 = (\neg p_{11}) \mathcal{U}_{[0, 25.0]} p_{12}$, wherein atomic proposition p_{11} refers to the shaded rectangle (bottom right) in Fig. 6(b) and p_{12} to the unshaded rectangle above in Fig. 6(b). In detail, $\mathcal{O}(p_{11}) = \{4\} \times [3.2, 3.8] \times [0.2, 0.8] \times \mathbb{R}^2$ and $\mathcal{O}(p_{12}) = \{8\} \times [3.2, 3.8] \times [1.2, 1.8] \times \mathbb{R}^2$. Informally, φ_1 says that the system should reach p_{12} within 25 time units without passing through p_{11} . The least robust (≈ 0) trajectory found by our algorithm is shown in Fig. 6(b) along with the scatter plot for the samples. Note that it could be the case that the system is correct with the respect to the specification, but it is definitely not robustly correct.

The second specification that we consider for falsification is $\varphi_2 = \square(p_{21} \rightarrow (p_{22} \mathcal{R} \neg p_{23}))$ where \mathcal{R} is the temporal operator release and $\mathcal{O}(p_{21}) = \{10\} \times \{x \in \mathbb{R}^4 \mid x_1 \geq 1.1\}$, $\mathcal{O}(p_{22}) = \{5, 6\} \times \{x \in \mathbb{R}^4 \mid x_2 \leq 1.05\}$ and $\mathcal{O}(p_{23}) = \{9\} \times \{x \in \mathbb{R}^4 \mid x_1 \leq 0.9\}$. Essentially, we are looking for a trajectory that satisfies $\neg \varphi_2 = \diamond(p_{21} \wedge ((\neg p_{22}) \mathcal{U} p_{23}))$, i.e., a trajectory that exceeds value 1.1 in x_1 only in location 10 and, then it goes below value 0.9 in x_1 only in location 9, but it never drops below value 1.05 in x_2 in locations 5 and 6 (x_2 may take any value in any other location besides 5 and 6 or x_2 in any location must stay above 1.05). \diamond

5 Experiments

We have implemented our techniques and, in particular, the new metrics inside our Matlab toolbox S-TALIRO [8]. Our toolbox is general enough to interact with various means for modeling CPS including Simulink/Stateflow models. We currently support full time bounded MTL for continuous as well as hybrid time trajectories. We remark that all the benchmark problems are distributed with S-TALIRO at <https://sites.google.com/a/asu.edu/s-taliro/> which also includes all the MTL specifications used in this section.

We performed a comparison of our implementation (MC) against a simple uniform random (UR) exploration of the state-space. Both MC and UR are each run for a maximum number of 1000 tests, terminating early if a falsifying trajectory is found. Since these techniques are randomized, each experiment was repeated 100 times (runs) under different seeds in order to obtain statistically significant results. Uniform random exploration provides an ideal measure of the difficulty of falsifying a property over a given input. Its rate of success empirically quantifies the difficulty of falsifying a given property. Finally, we have already argued about the importance of obtaining the least robust trajectory where falsification cannot be achieved. To this end, we compare the set of minima found using MC as well as that using UR and the corresponding running times.

Table 1 reports on the results of our comparison on two benchmark problems using different MTL properties or problem instances. The first benchmark problem is the Automatic Transmission (AT) model considered in Example 2.1. We consider a number of MTL specifications of increasing difficulty to falsify. The formulas ϕ_2^{AT} , ϕ_4^{AT} and ϕ_5^{AT} are of the form $\phi_i^{AT} = \neg \diamond(p_1^{AT} \wedge \diamond_{\mathcal{I}_i} p_2^{AT})$, where $\mathcal{I}_2 = [0, 10]$, $\mathcal{I}_4 = [0, 7.5]$ and $\mathcal{I}_5 = [0, 5]$. Formula ϕ_1^{AT} is described in Example 2.2 while ϕ_3^{AT} is $\neg \diamond(p_1^{AT} \wedge \diamond_{[0, 10]} p_3^{AT})$, where $\mathcal{O}(p_3^{AT}) = [125, +\infty) \times \mathbb{R}$. The second benchmark is a Simulink model of a 3rd order $\Delta - \Sigma$ modulator whose description can be found in [17]. The 3rd order $\Delta - \Sigma$ modulator has unknown initial conditions in the set $[-0.1, 0.1]^3$ and a one dimensional input signal that takes values in a set $[u_m, u_M]$. The problem instances in Table 1 indicate the bounds on the input signal $[u_m, u_M]$. The specification for the $\Delta - \Sigma$ modulator is that the state of the system should always remain in the set $[-1, 1]^3$.

We find that the performance varies depending on the ease with which the property can be violated by means of uniformly sampling the input space. If the property can be easily falsified, then it is advantageous to utilize uniform random search. MC for easy problem instances seems to converge and get trapped at local minima. In practice, we may periodically reset the MC simulation using random restarts. However, such restarts were not used in our experimental comparison. The use of MC is clearly advantageous when the problem is challenging. In hard problem instances, MC can falsify the specification when UR fails to falsify. Moreover, even when falsification fails, MC still computes lower

Table 1: Experimental Comparison of Monte-Carlo (MC) vs. Uniform Random (UR) falsification on benchmark problems with Euclidean output spaces. Each instance was run for 100 times and each run was executed for a maximum of 1000 tests. Legend: **#Fals.:** the number of runs falsified, **Robustness:** $\langle \text{min, average, variance} \rangle$ of the runs that were not falsified, **Time:** $\langle \text{min, average, max} \rangle$ time in seconds per run.

Problem	ψ	#Fals.		Robustness		Time (sec)	
		MC	UR	MC	UR	MC	UR
AT	ϕ_1^{AT}	97	100	$\langle 2.54, 7, 48.5 \rangle$	-	$\langle 0.2, 11, 92 \rangle$	$\langle 0.2, 3, 16 \rangle$
AT	ϕ_2^{AT}	96	100	$\langle 3.03, 137, 6.6 \cdot 10^4 \rangle$	-	$\langle 0.2, 16, 94 \rangle$	$\langle 0.2, 10, 48 \rangle$
AT	ϕ_3^{AT}	51	0	$\langle 8 \cdot 10^{-4}, 0.42, 1.2 \rangle$	$\langle 0.04, 0.96, 0.35 \rangle$	$\langle 7, 61, 94 \rangle$	$\langle 93, 94, 99 \rangle$
AT	ϕ_4^{AT}	0	0	$\langle 5.86, 5.95, 0.02 \rangle$	$\langle 5.91, 6.06, 0.01 \rangle$	$\langle 92, 93, 93 \rangle$	$\langle 92, 92, 93 \rangle$
AT	ϕ_5^{AT}	0	0	$\langle 0.15, 0.41, 2.55 \rangle$	$\langle 0.25, 0.57, 0.06 \rangle$	$\langle 93, 93, 94 \rangle$	$\langle 92, 93, 94 \rangle$
$P_{[-0.45, 0.45]}^{\Delta-\Sigma}$	$\phi_{\Delta-\Sigma}$	84	81	$\langle 0.00, 0.04, 4.6 \cdot 10^{-4} \rangle$	$\langle 0.00, 0.01, 1.2 \cdot 10^{-4} \rangle$	$\langle 0.2, 19, 41 \rangle$	$\langle 0.2, 19, 43 \rangle$
$P_{[-0.4, 0.4]}^{\Delta-\Sigma}$	$\phi_{\Delta-\Sigma}$	58	40	$\langle 0.00, 0.06, 7.9 \cdot 10^{-4} \rangle$	$\langle 0.00, 0.03, 2.2 \cdot 10^{-4} \rangle$	$\langle 0.7, 26, 39 \rangle$	$\langle 0.3, 30, 38 \rangle$
$P_{[-0.35, 0.35]}^{\Delta-\Sigma}$	$\phi_{\Delta-\Sigma}$	21	1	$\langle 0.00, 0.07, 2.1 \cdot 10^{-3} \rangle$	$\langle 0.01, 0.06, 7.9 \cdot 10^{-4} \rangle$	$\langle 4.1, 35, 49 \rangle$	$\langle 5.4, 37, 44 \rangle$

minimum and average robustness values with the same computational cost. Further experimental results on a simpler benchmark problem can be found in [47]. The results in [47] also confirm the aforementioned conclusions.

Table 2 compares the performance of the falsification algorithm on benchmark problems with hybrid output space. We compared UR with MC on two benchmark problems on various temporal logic formulas of increasing difficulty to falsify. The first benchmark problem was AT. As opposed to the previous experiments, the specifications now not only place conditions on the continuous state of the system, but also on the discrete locations. In detail, we are looking to generate tests such that the system goes through all the states in the state chart `selection_state` (see Fig. 2), i.e., `steady_state`, `upshifting` and `downshifting`, while the vehicle speed exceeds a threshold v . In detail, for $i = 6, 7, 8$, we consider the formulas $\phi_i^{AT} = \neg(\diamond p_{1,i}^{AT} \wedge \diamond p_{2,i}^{AT} \wedge \diamond p_{3,i}^{AT})$, where $p_{1,i}^{AT} = \{\text{steady_state}\} \times [v_i, +\infty) \times \mathbb{R}$, $p_{2,i}^{AT} = \{\text{downshifting}\} \times [v_i, +\infty) \times \mathbb{R}$ and $p_{3,i}^{AT} = \{\text{upshifting}\} \times [v_i, +\infty) \times \mathbb{R}$ with $v_6 = 79$, $v_7 = 79.5$ and $v_8 = 80$. Since S-TALiRO does not support yet automatic extraction of guard conditions, we compared only UR with MC using the metric d_h^0 for the distance computations.

The second example that we consider is the Navigation (NV) benchmark problem from [28]. This is a hybrid automaton benchmark problem and both the control locations and the guards of the transitions are available to us. Thus, we compared the performance of the Monte Carlo sampling algorithm under the metrics d_h and d_h^0 with the performance of Uniform Random sampling under the d_h metric. The problem instance that is used in our experiments is presented in [47]. We performed a number of experiments using the following formulas:

1. $\phi_1^{NV} = (\neg p_{11}^{NV}) \mathcal{U}_{[0.25, 0]} p_{12}^{NV}$, where $\mathcal{O}(p_{11}^{NV}) = \{4\} \times [3.2, 3.8] \times [0.2, 0.8] \times \mathbb{R}^2$ and $\mathcal{O}(p_{12}^{NV}) = \{8\} \times [3.2, 3.8] \times [1.2, 1.8] \times \mathbb{R}^2$;
2. $\phi_2^{NV} = \square(p_{21}^{NV} \rightarrow (p_{22}^{NV} \mathcal{R} \neg p_{23}^{NV}))$, where \mathcal{R} is the temporal operator *release* and $\mathcal{O}(p_{21}^{NV}) = \{10\} \times \{x \in \mathbb{R}^4 \mid x_1 \geq 1.1\}$, $\mathcal{O}(p_{22}^{NV}) = \{5, 6\} \times \{x \in \mathbb{R}^4 \mid x_2 \leq 1.05\}$ and $\mathcal{O}(p_{23}^{NV}) = \{9\} \times \{x \in \mathbb{R}^4 \mid x_1 \leq 0.9\}$;
3. $\phi_3^{NV} = \square(p_{31}^{NV} \rightarrow \square \neg p_{32}^{NV})$, where $\mathcal{O}(p_{31}^{NV}) = \{10\} \times \{x \in \mathbb{R}^4 \mid x_1 \geq 1.05 \wedge x_2 \geq 2\}$, $\mathcal{O}(p_{32}^{NV}) = \{5\} \times \{x \in \mathbb{R}^4 \mid x_1 \leq 1 \wedge x_2 \leq 1.95\}$;

Table 2: Experimental Comparison of Monte-Carlo (MC) vs. Uniform Random (UR) falsification on benchmark problems with hybrid output spaces. Each instance was run for 100 times and each run was executed for a maximum of 1000 tests. Legend: **#Fals.**: the number of runs falsified, **Time**: $\langle \text{min, average, max} \rangle$ time in seconds per run, **MC-H**: MC with metric d_h , **MC-H0**: MC with metric d_h^0 .

Problem	ψ	#Fals.			Time		
		MC-H	MC-H0	UR	MC-H	MC-H0	UR
AT	ϕ_6^{AT}	-	93	86	-	$\langle 0.4, 24, 138 \rangle$	$\langle 0.4, 56, 139 \rangle$
AT	ϕ_7^{AT}	-	94	55	-	$\langle 0.1, 25, 128 \rangle$	$\langle 0.6, 81, 127 \rangle$
AT	ϕ_8^{AT}	-	0	0	-	$\langle 110, 115, 139 \rangle$	$\langle 109, 111, 115 \rangle$
NV _[0,25]	ϕ_1^{NV}	63	68	34	$\langle 4.2, 542, 831 \rangle$	$\langle 34, 545, 865 \rangle$	$\langle 44, 623, 817 \rangle$
NV _[0,12]	ϕ_2^{NV}	100	100	100	$\langle 1.1, 24, 140 \rangle$	$\langle 1.7, 25, 168 \rangle$	$\langle 0.9, 22, 108 \rangle$
NV _[0,12]	ϕ_3^{NV}	100	100	100	$\langle 0.8, 8.7, 62 \rangle$	$\langle 0.8, 17, 503 \rangle$	$\langle 0.7, 4.0, 22 \rangle$
NV _[0,12]	ϕ_4^{NV}	100	100	100	$\langle 1.2, 18, 85 \rangle$	$\langle 1.4, 26, 66 \rangle$	$\langle 0.8, 35, 427 \rangle$
NV _[0,12]	ϕ_5^{NV}	38	47	5	$\langle 21.0, 419, 595 \rangle$	$\langle 15, 390, 584 \rangle$	$\langle 9.4, 404, 437 \rangle$

4. $\phi_4^{NV} = \square(p_{41}^{NV} \rightarrow \square \neg p_{32}^{NV})$, where $\mathcal{O}(p_{41}^{NV}) = \{10\} \times \{x \in \mathbb{R}^4 \mid x_1 \geq 1.2 \wedge x_2 \geq 2\}$; and
5. $\phi_5^{NV} = \square(p_{41}^{NV} \rightarrow \square \neg p_{42})$, where $\mathcal{O}(p_{42}^{NV}) = \{5\} \times \{x \in \mathbb{R}^4 \mid x_1 \leq 1 \wedge x_2 \leq 1.9\}$.

First, we observe that on easy problem instances, i.e., ϕ_2^{NV} - ϕ_4^{NV} , the performance of all algorithms is comparable in terms of computation time. On hard problem instances, both MC-H and MC-H0 outperform UR in terms of numbers of falsifications.

The experimental results indicate that the best way to approach hybrid system falsification / verification is with a layered approach. Assuming that at the initial design stages the errors are abundant, then it is preferable to run random sampling for the falsification process. As the system design becomes more mature, then Monte Carlo sampling with the new metrics introduced in this paper can be utilized for the falsification. When the level of confidence in the system design has increased and potentially the system design is robust enough, then the designer may use a reachability analysis algorithm (for example SpaceEx [29]). However, we remark that currently reachability analysis tools cannot handle arbitrary MTL specifications. A more detailed discussion on system verification that compares the advantages/disadvantages of falsification and reachability methods can be found in [1, 2].

6 Related work

Due to the known undecidability results in the analysis of hybrid systems [3] and the state explosion problem of the reachability computation algorithms (see [35] for some related references), a lot of recent research activity has concentrated on testing approaches to the verification of continuous and hybrid systems [36, 64].

The use of Monte Carlo techniques for model checking has been considered previously by Grosu and Smolka [32]. Whereas Grosu and Smolka consider random walks over the automaton defined by the system itself, our technique defines random walks over the input state space. These are, in general, distinct approaches to the problem. In practice, our approach does not have the limitation of being restricted by the topology of the system's state transition graph. Depending on this topology, the probability of visiting states deeper in the graph can sometimes be quite small in pathological cases. On the other hand, Grosu et al.'s technique can be extended readily to the case of systems with control inputs without requiring a finite parameterization of the control. We are currently investigating the possibility of combining both types of random walks in a single framework. Previous work by some of the authors in this work considered Monte-Carlo techniques for finding bugs in programs [56]. However, our previous efforts did not have the systematic definition of robustness that we employ here.

There exist two main approaches to the testing problem of hybrid systems. The first approach is focused on choosing inputs and/or parameters in a systematic fashion so as to cover the state-space of the system [21, 10, 12, 46, 49]. These approaches are mainly based on the theory of rapidly exploring random trees (RRTs). The other approach

is based on the notion of robust simulation trajectory [19, 31, 35, 41]. In robust testing, a simulation trajectory can represent a neighborhood of trajectories achieving, thus, better coverage guarantees. Recently, the authors in [16] have made the first steps in bridging these two aforementioned approaches.

On the research front of falsification/verification of temporal logic properties through testing, the results are limited [48, 52, 24]. The work that is the closest to ours appears in [52]. The authors of that work develop a different notion of robustness for temporal logic specifications, which is also used as a fitness function for optimization problems. Besides the differences in the application domain, i.e., [52] focuses on parameter estimation for biological systems, whereas our paper deals with the falsification of hybrid systems, the two works have also several differences at the theoretical and computational levels. At the theoretical level, we have introduced a new metric for hybrid spaces which enables reasoning over hybrid trajectories, while at the computational level our approach avoids set operations, e.g., union, complementation etc, which, in general, increase the computational load.

Younes and Simmons, and more recently, Clarke et al. have proposed the technique of *Statistical Model Checking* (SMC). SMC targets stochastic system models such as continuous-time Markov chains [62] or Stochastic Hybrid Automata (SHA) [14]. For example, in order to model imperfect sensors in Example 2.1, we may add Gaussian noise to the sensor that reads the engine speed. Then, the resulting system would be a SHA. The goal of SMC is to assess the probability that a system satisfies a given *probabilistic* temporal logic property φ . This probability can be safely approximated using Wald’s probabilistic ratio test. SMC, like our technique, requires a simulator to be available for the system, but not a transition relation representation. In contrast to SMC, our approach is guided by a robustness metric towards less robust trajectories. On the other hand, the complex nature of the system and the robustness metrics imply that we cannot yet provide guarantees on whether our algorithm has converged to the global minimum of the temporal logic robustness function. However, this is an on-going endeavor.

Remark 6.1 *Our method does not try to assess the probability of failure, but to detect a failure. That is, our goal is to provide the engineer with tools in order to detect design problems in the system rather than perform a failure analysis. In our framework, if a failure is detected, then the designer has a counterexample to work with in order to “debug” the system. Moreover, if a failure is not detected, then the designer is still provided with the least robust behavior found. The fact that the system might be correct with probability one does not imply that the system is robustly correct. Therefore, we view SMC and our approach as complementary. In an MBD cycle, the model should be first assessed for its robustly correct behavior, and, then, a failure analysis should be performed under various failure models and requirements.*

7 Conclusions

Embedded systems require the verification of elaborate specifications such as those that can be expressed in MTL. The undecidability of the MTL verification problem over such complex continuous systems mandates the use of lightweight formal methods that usually involve testing. In this paper, we have presented a testing framework for the Metric Temporal Logic (MTL) falsification of hybrid systems using Monte-Carlo optimization techniques. The use of hit-and-run Monte-Carlo optimization is required in order to overcome the difficulties in handling the complex system dynamics as well as the nonlinearities in the objective function. Moreover, in order to enable more efficient search in hybrid state-spaces, a generalized distance function was introduced.

Experimental results indicate the superiority of our testing framework over random search on the hard benchmark examples. The advantages of our approach are not limited only to the fact that we can falsify arbitrary systems, but also that we can provide robustness guarantees even to systems that have been proven correct. The techniques and the methods that were introduced in this paper have been implemented in our Matlab toolbox S-TALIRO [8].

Acknowledgments

The authors would like to thank Truong Nghiem and Professor George Pappas for the useful discussions. The authors would also like to thank the reviewers for the very careful reading of the manuscript and their numerous comments which have improved the paper.

References

- [1] Houssam Abbas and Georgios Fainekos. Linear hybrid system falsification through descent. Technical Report arXiv:1105.1733, 2011.
- [2] Houssam Abbas and Georgios Fainekos. Linear hybrid system falsification through local search. In *Automated Technology for Verification and Analysis*, volume 6996 of *LNCS*, pages 503–510. Springer, 2011.
- [3] R. Alur, C. Courcoubetis, N. Halbwachs, T. A. Henzinger, P.-H. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine. The algorithmic analysis of hybrid systems. *Theoretical Computer Science*, 138(1):3–34, 1995.
- [4] R. Alur, T. Dang, J.M. Esposito, Y. Hur, F. Ivancic, V. Kumar, I. Lee, P. Mishra, G. J. Pappas, and O. Sokolsky. Hierarchical modeling and analysis of embedded systems. *Proceedings of the IEEE*, 91(1):11–28, 2003.
- [5] R. Alur, T. A. Henzinger, G. Lafferriere, and George J. Pappas. Discrete abstractions of hybrid systems. *Proceedings of the IEEE*, 88(2):971–984, 2000.
- [6] Rajeev Alur, Thao Dang, and Franjo Ivančić. Progress on reachability analysis of hybrid systems using predicate abstraction. In *Hybrid Systems: Computation and Control*, volume 2623 of *LNCS*, pages 4–19. Springer, 2003.
- [7] Christophe Andrieu, Nando De Freitas, Arnaud Doucet, and Michael I. Jordan. An introduction to MCMC for machine learning. *Machine Learning*, 50:5–43, 2003.
- [8] Yashwanth Singh Rahul Annapureddy, Che Liu, Georgios E. Fainekos, and Sriram Sankaranarayanan. S-taliro: A tool for temporal logic falsification for hybrid systems. In *Tools and algorithms for the construction and analysis of systems*, volume 6605 of *LNCS*, pages 254–257. Springer, 2011.
- [9] H. Bandemer and S. Gottwald. *Fuzzy sets, fuzzy logic, fuzzy methods, with applications*. Wiley, 1995.
- [10] A. Bhatia and E. Frazzoli. Incremental search methods for reachability analysis of continuous and hybrid systems. In *HSCC*, volume 2993 of *LNCS*, pages 142–156. Springer, 2004.
- [11] Stephen Boyd and Stephen Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. cf. <http://www.stanford.edu/~boyd/cvxbook.html>.
- [12] M. Branicky, M. Curtiss, J. Levine, and S. Morgan. Sampling-based planning, control and verification of hybrid systems. *IEE Proc.-Control Theory Appl.*, 153(5):575–590, 2006.
- [13] Siddhartha Chib and Edward Greenberg. Understanding the Metropolis-Hastings algorithm. *The American Statistician*, 49(4):327–335, Nov 1995.
- [14] Edmund Clarke, Alexandre Donze, and Axel Legay. Statistical model checking of analog mixed-signal circuits with an application to a third order $\delta - \sigma$ modulator. In *Hardware and Software: Verification and Testing*, volume 5394/2009 of *Lecture Notes in Computer Science*, pages 149–163, 2009.
- [15] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Cliff Stein. *Introduction to Algorithms*. MIT Press/McGraw-Hill, second edition, September 2001.
- [16] T. Dang, A. Donze, O. Maler, and N. Shalev. Sensitive state-space exploration. In *Proc. of the 47th IEEE CDC*, pages 4049–4054, Dec. 2008.
- [17] Thao Dang, Alexandre Donzé, and Oded Maler. Verification of analog and mixed-signal circuits using hybrid system techniques. In *5th International Conference on Formal Methods in Computer-Aided Design*, volume 3312 of *LNCS*, pages 21–36. Springer, 2004.
- [18] Luca de Alfaro, Marco Faella, and Mariëlle Stoelinga. Linear and branching metrics for quantitative transition systems. In *Proceedings of the 31st ICALP*, volume 3142 of *LNCS*, pages 97–109. Springer, 2004.

- [19] Alexandre Donzé and Oded Maler. Systematic simulation using sensitivity analysis. In *HSCC*, volume 4416 of *LNCS*, pages 174–189. Springer, 2007.
- [20] Magnus Egerstedt and Clyde Martin. *Control Theoretic Splines: Optimal Control, Statistics, and Path Planning*. Princeton University Press, 2009.
- [21] Joel M. Esposito, Jongwoo Kim, and Vijay Kumar. Adaptive RRTs for validating hybrid robotic control systems. In *Proceedings of the International Workshop on the Algorithmic Foundations of Robotics*, 2004.
- [22] Joel M. Esposito and Vijay Kumar. An asynchronous integration and event detection algorithm for simulating multi-agent hybrid systems. *ACM Trans. Model. Comput. Simul.*, 14(4):363–388, 2004.
- [23] Esterel Technologies. Scade success stories. Online, April 2011. <http://www.esterel-technologies.com/technology/success-stories/>.
- [24] Georgios E. Fainekos, Antoine Girard, and George J. Pappas. Temporal logic verification using simulation. In *FORMATS*, volume 4202 of *LNCS*, pages 171–186. Springer, 2006.
- [25] Georgios E. Fainekos and George J. Pappas. Robustness of temporal logic specifications for finite state sequences in metric spaces. Technical Report MS-CIS-06-05, Dept. of CIS, Univ. of Pennsylvania, May 2006.
- [26] Georgios E. Fainekos and George J. Pappas. Robustness of temporal logic specifications for continuous-time signals. *Theoretical Computer Science*, 410(42):4262–4291, 2009.
- [27] Georgios E. Fainekos, Sriram Sankaranarayanan, Franjo Ivančić, and Aarti Gupta. Robustness of model-based simulations. In *IEEE Real-Time Systems Symposium*, pages 345–354. IEEE Press, 2009.
- [28] Ansgar Fehnker and Franjo Ivančić. Benchmarks for hybrid systems verification. In *HSCC*, volume 2993 of *LNCS*, pages 326–341. Springer, 2004.
- [29] Goran Frehse, Colas Le Guernic, Alexandre Donz, Scott Cotton, Rajarshi Ray, Olivier Lebeltel, Rodolfo Ripado, Antoine Girard, Thao Dang, and Oded Maler. Spaceex: Scalable verification of hybrid systems. In *Proceedings of the 23d CAV*, 2011.
- [30] Daan Frenkel and Berend Smit. *Understanding Molecular Simulation: From Algorithms to Applications*. Academic Press, 1996.
- [31] Antoine Girard and George J. Pappas. Verification using simulation. In *HSCC*, volume 3927 of *LNCS*, pages 272 – 286. Springer, 2006.
- [32] Radu Grosu and Scott Smolka. Monte carlo model checking. In *TACAS*, volume 3440 of *LNCS*, pages 271–286. Springer, 2005.
- [33] Thomas A. Henzinger. The theory of hybrid automata. In *Logic In Computer Science (LICS 1996)*, pages 278–292. IEEE Computer Society Press, 1996.
- [34] Thomas A. Henzinger, Peter W. Kopke, Anuj Puri, and Pravin Varaiya. What’s decidable about hybrid automata? *J. Comput. Syst. Sci.*, 57(1):94–124, 1998.
- [35] A. Agung Julius, Georgios E. Fainekos, Madhukar Anand, Insup Lee, and George J. Pappas. Robust test generation and coverage for hybrid systems. In *HSCC*, number 4416 in *LNCS*, pages 329–342. Springer, 2007.
- [36] Jim Kapinski, Bruce H. Krogh, Oded Maler, and Olaf Stursberg. On systematic simulation of open continuous systems. In *HSCC*, volume 2623 of *LNCS*, pages 283–297. Springer, 2003.
- [37] R. Kopperman. All topologies come from generalized metrics. *Amer. Math. Monthly*, 95:89–97, 1988.

- [38] Ron Koymans. Specifying real-time properties with metric temporal logic. *Real-Time Systems*, 2(4):255–299, 1990.
- [39] K. B. Lamine and F. Kabanza. Using fuzzy temporal logic for monitoring behavior-based mobile robots. In M.H. Hamza, editor, *Proceedings of the IASTED International Conference Robotics and Applications*, page 116122. IASTED/ACTA Press, 2000.
- [40] Edward A. Lee and Pravin Varaiya. *Structure and Interpretation of Signals and Systems*. Addison Wesley, 2003.
- [41] Flavio Lerda, James Kapinski, Edmund M. Clarke, and Bruce H. Krogh. Verification of supervisory control software using state proximity and merging. In *HSCC*, volume 4981 of *LNCS*, pages 344–357. Springer, 2008.
- [42] Laszlo Lovasz. Hit-and-run is fast and run. *Mathematical Programming*, 86:443–461, 1999.
- [43] Laszlo Lovasz and Santosh S. Vempala. Hit-and-run from a corner. *SIAM Journal on Computing*, 35(4):985–1005, 2006.
- [44] J. Lygeros, K. H. Johansson, S. N. Simic, J. Zhang, and S. Sastry. Dynamical properties of hybrid automata. *IEEE Transactions on Automatic Control*, 48:2–17, 2003.
- [45] Mathworks. Simulink user stories. Online, April 2011. <http://www.mathworks.com/products/simulink/userstories.html>.
- [46] Tarik Nahhal and Thao Dang. Test coverage for continuous and hybrid systems. In *CAV*, volume 4590 of *LNCS*, pages 449–462. Springer, 2007.
- [47] Truong Nghiem, Sriram Sankaranarayanan, Georgios E. Fainekos, Franjo Ivancic, Aarti Gupta, and George J. Pappas. Monte-carlo techniques for falsification of temporal properties of non-linear hybrid systems. In *Proceedings of the 13th ACM International Conference on Hybrid Systems: Computation and Control*, pages 211–220. ACM Press, 2010.
- [48] E. Plaku, Lydia E. Kavraki, and Moshe Y. Vardi. Falsification of ltl safety properties in hybrid systems. In *TACAS*, volume 5505 of *LNCS*, pages 368 – 382. Springer, 2009.
- [49] Erion Plaku, Lydia E. Kavraki, and Moshe Y. Vardi. Hybrid systems: From verification to falsification. In *CAV*, volume 4590 of *LNCS*, pages 463–476. Springer, 2007.
- [50] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, Cambridge (UK) and New York, 2nd edition, 1992.
- [51] Dana Randall. Rapidly mixing markov chains with applications in computer science and physics. *Computing in Science and Engineering*, 8(2), March 2006.
- [52] A. Rizk, G. Batt, F. Fages, and S. Soliman. On a continuous degree of satisfaction of temporal logic formulae with applications to systems biology. In *6th International Conference on Computational Methods in Systems Biology*, number 5307 in *LNCS*, pages 251–268. Springer, 2008.
- [53] H.E. Romeign and R.L. Smith. Simulated annealing for constrained global optimization. *Journal of Global Optimization*, 5:101–126, 1994.
- [54] Reuven Y. Rubinstein and Dirk P. Kroese. *Simulation and the Monte Carlo Method*. Wiley Series in Probability and Mathematical Statistics, 2008.
- [55] Ricardo G. Sanfelice and Andrew R. Teel. Dynamical properties of hybrid systems simulators. *Automatica*, 46(2):239–248, 2010.
- [56] Sriram Sankaranarayanan, Richard M. Chang, Guofei Jiang, and Franjo Ivancic. State space exploration using feedback constraint generation and monte-carlo sampling. In *ESEC/SIGSOFT FSE*, pages 321–330. ACM, 2007.

- [57] Anthony Karel Seda and Pascal Hitzler. Generalized distance functions in the theory of computation. *The Computer Journal*, 53(4):bxm108443–464, 2008.
- [58] R.L. Smith. Monte Carlo procedures for generating points uniformly distributed over bounded regions. *Operations Research*, 38(3):1296–1308, 1984.
- [59] Robert L. Smith. The hit-and-run sampler: a globally reaching markov chain sampler for generating arbitrary multivariate distributions. In *Proceedings of the 28th conference on Winter simulation*, pages 260–264. IEEE Computer Society, 1996.
- [60] Eduardo D. Sontag. *Mathematical Control Theory: Deterministic Finite Dimensional Systems*. Springer, 2nd edition, 1998.
- [61] S. Tripakis and T. Dang. *Model-Based Design for Embedded Systems*, chapter Modeling, Verification and Testing using Timed and Hybrid Automata, pages 383–436. CRC Press, 2009.
- [62] Håkan L. S. Younes and Reid G. Simmons. Statistical probabilistic model checking with a focus on time-bounded properties. *Information & Computation*, 204(9):1368–1409, 2006.
- [63] A.B. Zabinsky, R.L. Smith, J.F. MacDonald, H.E. Romeijn, and D.E. Kaufman. Improving hit-and-run for global optimization. *Journal of Global Optimization*, 3:171–192, 1993.
- [64] Qianchuan Zhao, Bruce H. Krogh, and Paul Hubbard. Generating test inputs for embedded control systems. *IEEE Control Systems Magazine*, August:49–57, 2003.

APPENDIX

Proof: [of Theorem 4.1] The proof is by induction on the structure of the formula.

1. We will present only the base cases, since the other cases are identical with those in the proofs in [26] and [25].
 - If $\llbracket p, \mathcal{O} \rrbracket_{\mathbf{d}}(\mathbf{y}, t) \succ \mathbf{0}$, then by definition $\mathbf{Dist}_{\mathbf{d}}(\mathbf{y}(t), \mathcal{O}(p)) \succ \mathbf{0}$, which implies that $\mathbf{y}(t) \in \mathcal{O}(p)$ and, thus, that $(\mathbf{y}, t, \mathcal{O}) \models p$.
 - If $(\mathbf{y}, t, \mathcal{O}) \models p$, then by definition $\mathbf{y}(t) \in \mathcal{O}(p)$, which implies that $\mathbf{Dist}_{\mathbf{d}}(\mathbf{y}(t), \mathcal{O}(p)) \succeq \mathbf{0}$, and, thus, that $\llbracket \phi, \mathcal{O} \rrbracket_{\mathbf{d}}(\mathbf{y}, t) \succeq \mathbf{0}$.

Note that the equality in the first case fails when the signal value $\mathbf{y}(t)$ is right on the boundary of the set $\mathcal{O}(p)$, i.e., $\mathbf{y}(t) \in \partial\mathcal{O}(p)$. If $\llbracket p, \mathcal{O} \rrbracket_{\mathbf{d}}(\mathbf{y}, t) = \mathbf{0}$, then we cannot distinguish whether $(\mathbf{y}, t, \mathcal{O}) \models p$ or $(\mathbf{y}, t, \mathcal{O}) \not\models p$.
2. Similar to the previous proof.
3. We will present the base case and the negation (the other cases are based on the definition of supremum and infimum over the partial order \preceq of \mathbf{d} and are similar to the negation).
 - **Base case:**
 - If $\llbracket p, \mathcal{O} \rrbracket_{\mathbf{d}}(\mathbf{y}, t) = \varepsilon \succ \mathbf{0}$, then $(\mathbf{y}, t, \mathcal{O}) \models p$ and by definition $\mathbf{dist}_{\mathbf{d}}(\mathbf{y}(t), Y \setminus \mathcal{O}(p)) = \varepsilon \succ \mathbf{0}$, which implies that $\mathcal{B}_{\mathbf{d}}(\mathbf{y}(t), \varepsilon) \subseteq \mathcal{O}(p)$. Since $\mathbf{y}' \in \mathcal{B}_{\rho_{\mathbf{d}}}(\mathbf{y}, \varepsilon)$, we have $\rho_{\mathbf{d}}(\mathbf{y}, \mathbf{y}') = \sup_{t \in R} \mathbf{d}(\mathbf{y}(t), \mathbf{y}'(t)) \prec \varepsilon$. That is, $\mathbf{d}(\mathbf{y}(t), \mathbf{y}'(t)) \prec \varepsilon$ and, thus, $\mathbf{y}'(t) \in \mathcal{B}_{\mathbf{d}}(\mathbf{y}(t), \varepsilon) \subseteq \mathcal{O}(p)$. Hence, $(\mathbf{y}', t, \mathcal{O}) \models p$.
 - Similar to the previous case.
 - **Negation:**
 - Positive case: If $\llbracket \neg\phi, \mathcal{O} \rrbracket_{\mathbf{d}}(\mathbf{y}, t) = \varepsilon \succ \mathbf{0}$, then (i) $(\mathbf{y}, t, \mathcal{O}) \models \neg\phi$, i.e., $(\mathbf{y}, t, \mathcal{O}) \not\models \phi$, and (ii) $\llbracket \phi, \mathcal{O} \rrbracket_{\mathbf{d}}(\mathbf{y}, t) = -\varepsilon \prec \mathbf{0}$. Then, by (ii) and the induction hypothesis we have that for all $\mathbf{y}' \in \mathcal{B}_{\rho_{\mathbf{d}}}(\mathbf{y}, \varepsilon)$, $(\mathbf{y}', t, \mathcal{O}) \not\models \phi$.

– Negative case: Similar to the previous case.

We chose to present negation in order to demonstrate the properties that the negation must satisfy.

Proposition A.1 *The generalized distance function \mathbf{d}_h is a quasi-metric.*

Proof: We will need to demonstrate that the identity property and the triangle inequality hold. In the following, we let $h_i = (\ell_i, z_i) \in \mathbb{H}$ with $i = 1, 2, 3$.

Identity: Since π is a quasi-metric, we have $\pi(\ell_1, \ell_2) = 0$ iff $\ell_1 = \ell_2$. Since d is a metric, we have $d(z_1, z_2) = 0$ iff $z_1 = z_2$. Hence, $\mathbf{d}_h(h_1, h_2) = \langle 0, 0 \rangle$ iff $h_1 = h_2$.

Triangle Inequality: We need to show that for all $h_1, h_2, h_3 \in \mathbb{H}$, $\mathbf{d}_h(h_1, h_2) \leq \mathbf{d}_h(h_1, h_3) + \mathbf{d}_h(h_3, h_2)$. We proceed by case by case analysis:

1. **Case $\ell_1 = \ell_2 = \ell_3$:** Then,

$$\begin{aligned} \mathbf{d}_h(h_1, h_2) &= \langle 0, d(z_1, z_2) \rangle \leq \langle 0, d(z_1, z_3) + d(z_3, z_2) \rangle \\ &= \langle 0, d(z_1, z_3) \rangle + \langle 0, d(z_3, z_2) \rangle = \mathbf{d}_h(h_1, h_3) + \mathbf{d}_h(h_3, h_2) \end{aligned}$$

2. **Case $\ell_1 = \ell_2 \neq \ell_3$:** Then, $\pi(\ell_1, \ell_3) > 0$ and $\pi(\ell_3, \ell_2) > 0$ and

$$\begin{aligned} \mathbf{d}_h(h_1, h_2) &= \langle 0, d(z_1, z_2) \rangle \leq \langle \pi(\ell_1, \ell_3) + \pi(\ell_3, \ell_2), 0 \rangle = \langle \pi(\ell_1, \ell_3), 0 \rangle + \langle \pi(\ell_3, \ell_2), 0 \rangle \\ &\leq \left\langle \pi(\ell_1, \ell_3), \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_3)} \mathbf{dist}_d(z_1, \mathbf{G}^t(\ell_1, \ell)) \right\rangle + \left\langle \pi(\ell_3, \ell_2), \min_{\ell \in \partial \mathcal{N}_\pi(\ell_3, \ell_2)} \mathbf{dist}_d(z_3, \mathbf{G}^t(\ell_3, \ell)) \right\rangle \\ &= \mathbf{d}_h(h_1, h_3) + \mathbf{d}_h(h_3, h_2) \end{aligned}$$

3. **Case $\ell_1 \neq \ell_2$ and $\ell_1 = \ell_3$:** Then,

$$\mathbf{d}_h(h_1, h_2) = \left\langle \pi(\ell_1, \ell_2), \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_2)} \mathbf{dist}_d(z_1, \mathbf{G}^t(\ell_1, \ell)) \right\rangle$$

But, $\pi(\ell_1, \ell_2) = 0 + \pi(\ell_3, \ell_2) = \pi(\ell_1, \ell_3) + \pi(\ell_3, \ell_2)$, and, also,

$$\begin{aligned} \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_2)} \mathbf{dist}_d(z_1, \mathbf{G}^t(\ell_1, \ell)) &= \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_2)} \inf\{d(z_1, z) \mid z \in \mathbf{G}^t(\ell_1, \ell)\} \\ &\leq \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_2)} \inf\{d(z_1, z_3) + d(z_3, z) \mid z \in \mathbf{G}^t(\ell_1, \ell)\} \\ &= d(z_1, z_3) + \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_2)} \inf\{d(z_3, z) \mid z \in \mathbf{G}^t(\ell_1, \ell)\} \\ &\stackrel{(\ell_1 = \ell_3)}{=} d(z_1, z_3) + \min_{\ell \in \partial \mathcal{N}_\pi(\ell_3, \ell_2)} \mathbf{dist}_d(z_3, \mathbf{G}^t(\ell_3, \ell)) \end{aligned}$$

$$\begin{aligned} \text{Thus, } \mathbf{d}_h(h_1, h_2) &= \left\langle \pi(\ell_1, \ell_2), \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_2)} \mathbf{dist}_d(z_1, \mathbf{G}^t(\ell_1, \ell)) \right\rangle \\ &\leq \left\langle \pi(\ell_1, \ell_3) + \pi(\ell_3, \ell_2), d(z_1, z_3) + \min_{\ell \in \partial \mathcal{N}_\pi(\ell_3, \ell_2)} \mathbf{dist}_d(z_3, \mathbf{G}^t(\ell_3, \ell)) \right\rangle \\ &= \left\langle \pi(\ell_1, \ell_3), d(z_1, z_3) \right\rangle + \left\langle \pi(\ell_3, \ell_2), \min_{\ell \in \partial \mathcal{N}_\pi(\ell_3, \ell_2)} \mathbf{dist}_d(z_3, \mathbf{G}^t(\ell_3, \ell)) \right\rangle \end{aligned}$$

$$= \mathbf{d}_h(h_1, h_3) + \mathbf{d}_h(h_3, h_2)$$

4. **Case $\ell_1 \neq \ell_2$ and $\ell_2 = \ell_3$:** Then,

$$\mathbf{d}_h(h_1, h_2) = \left\langle \pi(\ell_1, \ell_2), \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_2)} \mathbf{dist}_d(z_1, \mathbf{G}^t(\ell_1, \ell)) \right\rangle$$

But, $\pi(\ell_1, \ell_2) = \pi(\ell_1, \ell_3) + 0 = \pi(\ell_1, \ell_3) + \pi(\ell_3, \ell_2)$, and, also,

$$\begin{aligned} \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_2)} \mathbf{dist}_d(z_1, \mathbf{G}^t(\ell_1, \ell)) &\stackrel{(\ell_2 = \ell_3)}{=} \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_3)} \mathbf{dist}_d(z_1, \mathbf{G}^t(\ell_1, \ell)) \\ &\leq \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_3)} \mathbf{dist}_d(z_1, \mathbf{G}^t(\ell_1, \ell)) + d(z_3, z_2) \end{aligned}$$

since $d(z_3, z_2) \geq 0$. Thus,

$$\begin{aligned} \mathbf{d}_h(h_1, h_2) &= \left\langle \pi(\ell_1, \ell_2), \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_2)} \mathbf{dist}_d(z_1, \mathbf{G}^t(\ell_1, \ell)) \right\rangle \\ &\leq \left\langle \pi(\ell_1, \ell_3) + \pi(\ell_3, \ell_2), \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_3)} \mathbf{dist}_d(z_1, \mathbf{G}^t(\ell_1, \ell)) + d(z_3, z_2) \right\rangle \\ &= \left\langle \pi(\ell_1, \ell_3), \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_3)} \mathbf{dist}_d(z_1, \mathbf{G}^t(\ell_1, \ell)) \right\rangle + \left\langle \pi(\ell_3, \ell_2), d(z_3, z_2) \right\rangle \\ &= \mathbf{d}_h(h_1, h_3) + \mathbf{d}_h(h_3, h_2) \end{aligned}$$

5. **Case $\ell_1 \neq \ell_2, \ell_1 \neq \ell_3$ and $\ell_2 \neq \ell_3$:** Then, $\pi(\ell_1, \ell_2) \leq \pi(\ell_1, \ell_3) + \pi(\ell_3, \ell_2)$, and

$$\begin{aligned} \mathbf{d}_h(h_1, h_2) &= \left\langle \pi(\ell_1, \ell_2), \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_2)} \mathbf{dist}_d(z_1, \mathbf{G}^t(\ell_1, \ell)) \right\rangle \\ &\leq \langle \pi(\ell_1, \ell_3) + \pi(\ell_3, \ell_2), 0 \rangle = \langle \pi(\ell_1, \ell_3), 0 \rangle + \langle \pi(\ell_3, \ell_2), 0 \rangle \\ &\leq \left\langle \pi(\ell_1, \ell_3), \min_{\ell \in \partial \mathcal{N}_\pi(\ell_1, \ell_3)} \mathbf{dist}_d(z_1, \mathbf{G}^t(\ell_1, \ell)) \right\rangle + \left\langle \pi(\ell_3, \ell_2), \min_{\ell \in \partial \mathcal{N}_\pi(\ell_3, \ell_2)} \mathbf{dist}_d(z_3, \mathbf{G}^t(\ell_3, \ell)) \right\rangle \\ &= \mathbf{d}_h(h_1, h_3) + \mathbf{d}_h(h_3, h_2) \end{aligned}$$

Proposition A.2 The generalized distance function \mathbf{d}_h^0 is a quasi-metric.

Proof: The proof is similar to the proof of Proposition A.1.

Proposition A.3 Let the current point be $h = \langle \ell, z \rangle$ and $\mathcal{O}(p) = L_p \times Z_p$, then $\mathbf{Dist}_{\mathbf{d}_h^0}(h, \mathcal{O}(p)) \neq \langle k, \pm\infty \rangle$ for any $k \in \mathbb{Z}$. Similarly for $\mathbf{Dist}_{\mathbf{d}_h}(h, \mathcal{O}(p))$.

Proof: Actually, we will show that $\mathbf{Dist}_{\mathbf{d}_h^0}(h, \mathcal{O}(p)) = \langle k, \pm\infty \rangle$ iff $k = \pm\infty$.

1. $h \notin \mathcal{O}(p)$ and $\ell \notin L_p$ and if L_p is not reachable from ℓ , then for any $\ell' \in L_p$, we have $\pi(\ell, \ell') = +\infty$. Thus, $\partial \mathcal{N}_\pi(\ell, \ell') = \emptyset$ and $\min_{\ell'' \in \partial \mathcal{N}_\pi(\ell, \ell')} \mathbf{dist}_d(z, \mathbf{G}^t(\ell, \ell'')) = +\infty$. Hence, $\mathbf{Dist}_{\mathbf{d}_h}(h, \mathcal{O}(p)) = -\mathbf{dist}_{\mathbf{d}_h}(h, \mathcal{O}(p)) = \langle -\infty, -\infty \rangle$. Also, $\mathbf{Dist}_{\mathbf{d}_h^0}(h, \mathcal{O}(p)) = \langle -\infty, -\infty \rangle$ by definition.

2. If $h \notin O(p)$ and $\ell \notin L_p$ and if L_p is reachable from ℓ , then $\partial\mathcal{N}_\pi(\ell, \ell') \neq \emptyset$ since (i) at least one of the neighbors of ℓ will have distance to L_p less than $\mathbf{dist}_\pi(\ell, L_p)$ and (ii) we have assumed that $\mathbf{G}^t(\ell, \ell') \neq \emptyset$ for all $\ell' \in \mathbf{Nxt}(\ell)$. Then, $\mathbf{dist}_d(z, \mathbf{G}^t(\ell, \ell'')) < +\infty$ for all $\ell'' \in \partial\mathcal{N}_\pi(\ell, \ell')$. Let

$$\ell^* \in \arg \min \left\{ \left\langle \pi(\ell, \ell'), \min_{\ell'' \in \partial\mathcal{N}_\pi(\ell, \ell')} \mathbf{dist}_d(z, \mathbf{G}^t(\ell, \ell'')) \right\rangle \mid \ell' \in L_p \right\}$$

and set $\delta^* = \min_{\ell'' \in \partial\mathcal{N}_\pi(\ell, \ell^*)} \mathbf{dist}_d(z, \mathbf{G}^t(\ell, \ell'')) < +\infty$. Therefore, $\mathbf{Dist}_{\mathbf{d}_h}(h, \mathcal{O}(p)) = -\mathbf{dist}_{\mathbf{d}_h}(h, \mathcal{O}(p)) = \langle -\pi(\ell, \ell^*), -\delta^* \rangle$. Finally, by definition, we have $\mathbf{Dist}_{\mathbf{d}_h^0}(h, \mathcal{O}(p)) = -\mathbf{dist}_{\mathbf{d}_h}(h, \mathcal{O}(p)) = \langle -\mathbf{dist}_\pi(\ell, L_p), 0 \rangle$.

3. If $h \notin O(p)$, but $\ell \in L_p$, i.e., $z \notin Z_p$, then $\mathbf{Dist}_{\mathbf{d}_h}(h, \mathcal{O}(p)) = -\mathbf{dist}_{\mathbf{d}_h}(h, \mathcal{O}(p)) = -\mathbf{dist}_{\mathbf{d}_h}(h, (L_p \setminus \{\ell\} \times Z_p) \cup (\{\ell\} \times Z_p)) = -\min\{\mathbf{dist}_{\mathbf{d}_h}(h, L_p \setminus \{\ell\} \times Z_p), \mathbf{dist}_{\mathbf{d}_h}(h, \{\ell\} \times Z_p)\} = -\mathbf{dist}_{\mathbf{d}_h}(h, \{\ell\} \times Z_p) = \langle 0, -\mathbf{dist}_d(z, Z_p) \rangle$. However, $\mathbf{dist}_d(z, Z_p) < +\infty$ since $\emptyset \subset \mathcal{O}(p) \subset Y$ by assumption. Similarly for $\mathbf{Dist}_{\mathbf{d}_h^0}(h, \mathcal{O}(p))$.
4. If $h \in O(p)$ and $Z_p \subset Z$, then $\mathbf{Dist}_{\mathbf{d}_h}(h, \mathcal{O}(p)) = \mathbf{dist}_{\mathbf{d}_h}(h, Y \setminus \mathcal{O}(p)) = \mathbf{dist}_{\mathbf{d}_h}(h, ((L \setminus L_p) \times Z) \cup (L \times (Z \setminus Z_p))) = \min\{\mathbf{dist}_{\mathbf{d}_h}(h, (L \setminus L_p) \times Z), \mathbf{dist}_{\mathbf{d}_h}(h, L \times (Z \setminus Z_p))\} = \mathbf{dist}_{\mathbf{d}_h}(h, L \times (Z \setminus Z_p)) = \langle 0, \mathbf{dist}_d(z, Z_p) \rangle$ since $\ell \in L_p \subseteq L$. However, $\mathbf{dist}_d(z, Z_p) = 0 < +\infty$ since $\emptyset \subset \mathcal{O}(p) \subset Y$ by assumption. Similarly for $\mathbf{Dist}_{\mathbf{d}_h^0}(h, \mathcal{O}(p))$.
5. If $h \in O(p)$ and $Z_p = Z$, i.e., $L_p \subset L$, then $\mathbf{Dist}_{\mathbf{d}_h}(h, \mathcal{O}(p)) = \mathbf{dist}_{\mathbf{d}_h}(h, Y \setminus \mathcal{O}(p)) = \mathbf{dist}_{\mathbf{d}_h}(h, ((L \setminus L_p) \times Z) \cup (L \times (Z \setminus Z_p))) = \mathbf{dist}_{\mathbf{d}_h}(h, ((L \setminus L_p) \times Z) \cup (L \times \emptyset)) = \mathbf{dist}_{\mathbf{d}_h}(h, (L \setminus L_p) \times Z)$. Now, we have two cases:
- if $L \setminus L_p$ is reachable from ℓ , then as in case (2), we have $\mathbf{Dist}_{\mathbf{d}_h}(h, \mathcal{O}(p)) = \langle \pi(\ell, \ell^*), \delta^* \rangle$ with $\delta^* < +\infty$.
 - if $L \setminus L_p$ is not reachable from ℓ , then $\mathbf{dist}_\pi(\ell, L \setminus L_p) = +\infty$ and as in case (1), we also have $\partial\mathcal{N}_\pi(\ell, \ell') = \emptyset$ for all $\ell' \in L \setminus L_p$. Thus, $\mathbf{Dist}_{\mathbf{d}_h}(h, \mathcal{O}(p)) = \langle +\infty, +\infty \rangle$.

Similarly, we can derive the value of $\mathbf{Dist}_{\mathbf{d}_h^0}(h, \mathcal{O}(p))$.

This concludes the proof since we have considered all possible cases.