Closing the Gap between Discrete Abstractions and Continuous Control: Completeness via Robustness and Controllability*

Jun Liu[0000−0001−8762−2299]

University of Waterloo, Waterloo Ontario N2L 3G1, Canada
j.liu@uwaterloo.ca

Abstract. A central theoretical question surrounding abstraction-based control of continuous nonlinear systems is whether one can decide through algorithmic procedures the existence of a controller to render the system to satisfy a given specification (e.g., safety, reachability, or more generally a temporal logic formula). Known algorithms are mostly sound but not complete in the sense that they return a correct controller upon termination, but do not offer guarantees of finding a controller if one exists. Completeness of abstraction-based nonlinear control in the general setting, therefore, remains an open question. This paper investigates this theoretical question and presents two sets of main results. First, we prove that sampled-data control of nonlinear systems with temporal logic specifications is robustly decidable in the sense that, given a continuous-time nonlinear control system and a temporal logic formula, one can algorithmically decide whether there exists a robust sampled-data control strategy to realize this specification when the right-hand side of the system is slightly perturbed by a small disturbance. Second, we show that under the assumption of local nonlinear controllability of the nominal system around an arbitrary trajectory that realizes a given specification, we can always construct a (robust) sampled-data control strategy via a sufficiently fine discrete abstraction. In a sense, this shows that temporal logic control for controllable nonlinear systems is decidable.

Keywords: Nonlinear systems · Temporal logic · Control synthesis · Completeness· Robustness · Controllability

1 Introduction

The control of dynamical systems to satisfy formal specifications (e.g., temporal logics) has received considerable attention in the past decade [4, 30]. This is partially motivated by the increasing demand of autonomous decision making by physical systems (e.g., mobile robots) in uncertain environments to achieve more complex tasks [6, 10, 12]. Many system relations have been proposed as

* Supported by the Natural Sciences and Engineering Research Council of Canada, the Canada Research Chairs Program, and the Ontario Early Researcher Award Program. The paper has an extended version with Appendix available at [19].

abstractions of nonlinear systems [22, 24, 27, 28, 32]. Such abstractions are desirable for several reasons. First, they are sound in the sense that they can be used to design provably correct controllers with respect to a given formal specification. Second, they are often finite (e.g., finite transition systems) and the original control design problem over an infinite state space can be effectively solved as a search problem over a finite structure. Third, the construction of these abstractions can be automated with the aid of a computer.

One of the main drawbacks of abstraction-based approaches is their computational cost, which is often incurred when a finer and finer abstraction is used in the hope of finding a controller when a coarser abstraction fails to yield one. However, without theoretical guarantees on completeness, i.e., if a control strategy exists, then it can be found by an abstraction-based approach, such computational efforts can be futile. This motivates the research in this paper.

Related work: We review several results in the literature that are most relevant to the result presented in this paper. In [31], it is shown that bisimilar (equivalent) symbolic models exist for controllable discrete-time linear systems and, as a result, temporal logic control for discrete-time controllable linear systems is decidable. For nonlinear systems, the authors of [27] showed that approximately bisimilar models can be constructed for incrementally stable systems [1]. The assumption of incremental stability essentially allows one to construct a deterministic transition system that can approximate a sampled-data representation of the original nonlinear system to any degree of precision. For nonlinear systems without the incremental stability assumption, the authors of [32] showed that symbolic models that approximately alternatingly simulate the sampled-data representation of a general nonlinear control system can be constructed. Because a sampled-data representation is used in [27,32], inter-sample behaviours are not considered in such approximations. The authors of [25] (see also [20, 24]) considered partition-based transition systems as over-approximations of nonlinear systems for synthesizing controllers for temporal logic specifications. Because no time-discretization is used, correctness guarantee is proved for the continuous-time trajectories. In [21, 22], the authors proposed a notion of robust abstractions of continuous-time nonlinear systems using grid-based approximations. A salient feature of such abstractions is that they under-approximate the control space so that all controls used by the abstractions can be implemented by the original system. At the same time, they over-approximate the reachable sets of the original system under a control so that correctness can be guaranteed (behaviours of the original system are included by the behaviours of the abstract system). In addition, the work in [21, 22] also tackled the problem of synthesizing robust controllers by introducing robustness margins in the abstractions and reasoned inter-sampling behaviours so that correctness is proved in continuous-time semantics of linear temporal logic. In [28], the authors proposed feedback refinement relations that can be used for control design for systems modelled by difference inclusions. This system relation has the same feature of under-approximating the control space, while over-approximating the reachable sets of the original system.

Nonetheless, all the above mentioned abstraction techniques are sound but not complete, with the exception of [27, 31], where additional assumptions on system dynamics are needed (controllable linear and incrementally stable, respectively). In [18], a notion of completeness for abstractions of discrete-time nonlinear systems is proved using a robustness argument (termed as robust completeness). It is shown that with sufficient computational sources, one can construct a finite transition system that robustly abstracts a discrete-time nonlinear system and, at the same time, is robustly abstracted by a slightly perturbed version of the same system. We also note that in [13–17] robust completeness is achieved for invariance and reachability type specifications and beyond using interval analysis for direct control synthesis on the continuous state space without first constructing abstractions. All these results focus on discrete-time control systems. The decidability via robustness for *continuous-time* control systems remains an *open* question.

Main contributions: In this paper, motivated by the above open question, we establish two sets of theoretical results on abstraction-based control of continuous-time nonlinear systems. First, we prove that robustly complete abstractions of continuous-time control systems exist under a mild assumption (i.e., local Lipschitz continuity) on system dynamics and use this to show decidability of robust realization of temporal logic formulas for continuous-time nonlinear systems by using a sampled-data control strategy. Furthermore, we show that under a suitable assumption on local controllability of the nominal system around a satisfying trajectory, it is guaranteed that a sufficiently fine discrete abstraction will return a (robust) controller. This, in a sense, shows that temporal logic control for controllable nonlinear systems is decidable.

2 Problem Formulation

2.1 Continuous-time control system

Consider a *continuous-time nonlinear control system* of the form:

$$x' = f(x, u), \tag{1}$$

where $x \in X \subseteq \mathbb{R}^n$ is the system state and $u \in U \subseteq \mathbb{R}^m$ is the control input. We assume that $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ satisfies the basic regularity assumptions (e.g., local Lipschitz continuity) such that, given any sufficiently regular control input signal and any initial condition, there exists a unique local solution to (1).

A *trajectory* of (1) is a pair $(\mathbf{x}, \mathbf{u})$, where $\mathbf{x} : \mathbb{R}^+ \to X$ is a state trajectory, $\mathbf{u} : \mathbb{R}^+ \to U$ is an input trajectory, and $(\mathbf{x}, \mathbf{u})$ satisfies (1) in the sense that $\mathbf{x}'(t) = f(\mathbf{x}(t), \mathbf{u}(t))$ for all $t \geq 0$.

A *(sampled-data) control strategy* with sampling period $\tau > 0$ for (1) is a partial function of the form:

$$\sigma(x_0, \cdots, x_i) = u_i \in U, \ \forall i = 0, 1, 2, \cdots, \tag{2}$$

where $x_0, \cdots, x_i$ is a finite sequence of sampled states taken at sampling times $t_0 = 0, \cdots, t_i$ and u_i is a constant control input. The sampling times t_0, t_1, t_2, $\cdots$ satisfy $t_{i+1} - t_i = \tau$ for all $i \geq 0$, where $\tau > 0$ is the sampling period that represents the duration for which the constant u_i is applied to the system.

A *σ-controlled trajectory* is a trajectory $(\mathbf{x}, \mathbf{u})$ resulting from executing the control strategy σ, where $\mathbf{u}$ is defined by $\mathbf{u}(t) = u_i$ for $t \in [t_i, t_{i+1})$, where $t_i = i\tau$ and u_i is determined by (2).

Given a positive integer N, a control strategy σ is said to have *dwell time* N, if each control input u_i is used for a multiple of N times, that is, if $i = mN$ for some integer m, then

$$u_i = u_{i+1} = \cdots = u_{i+N-1}. \tag{3}$$

This can be easily encoded by a control strategy with a simple counter. This seemingly peculiar definition plays a role later on in proving completeness for any fixed, but not necessarily small, sampling period.

2.2 δ-perturbed control system

Given a scalar $\delta \geq 0$, a *δ-perturbation* of the continuous-time nonlinear control system (1) is the differential inclusion

$$x' \in f(x, u) + \delta\mathbb{B}, \tag{4}$$

where $f(x, u) + \delta\mathbb{B}$ denotes the unit closed ball (in infinity norm) centered at $f(x, u)$. A *trajectory* of (4) is a pair $(\mathbf{x}, \mathbf{u})$, where $\mathbf{x} : \mathbb{R}^+ \to X$ is a state trajectory, $\mathbf{u} : \mathbb{R}^+ \to U$ is an input trajectory, and $(\mathbf{x}, \mathbf{u})$ satisfies (4) in the sense that $\mathbf{x}'(t) \in f(\mathbf{x}(t), \mathbf{u}(t)) + \delta\mathbb{B}$ for all $t \geq 0$.

We call system (1) the nominal system and denote it by $\mathcal{S}$. The δ-perturbation of $\mathcal{S}$ defined by (4) is denoted by $\mathcal{S}_\delta$. Apparently, $\mathcal{S}_0$ is exactly $\mathcal{S}$.

2.3 Linear temporal logic and labelling function

We consider linear-time properties described by linear temporal logic (LTL) [26]. In particular, we consider LTL without the next operator ($\text{LTL}_{\setminus\bigcirc}$) and we refer the readers to [3] (or the Appendix of [19]) for its syntax and semantics. We also assume that the formulas are written in a positive form, where negations of atomic propositions are replaced with new atomic propositions (for details, see the Appendix of the extended version [19]).

In the following, we need to reason about satisfaction of LTL formulas by continuous-time trajectories and by discrete-time sequences, and in particular, the implication between the two. For this purpose, we need to introduce the notion of an ε-strengthening of a labelling function [18]. For $\varepsilon > 0$, a labelling function $\mathcal{L}_\varepsilon : \mathbb{R}^n \to 2^\Pi$ is said to be the *ε-strengthening* of another labelling function $\mathcal{L} : \mathbb{R}^n \to 2^\Pi$, if $\pi \in \mathcal{L}_\varepsilon(x)$ *if and only if* $\pi \in \mathcal{L}(y)$ for all $y \in x + \varepsilon\mathbb{B}$. The Appendix includes an illustration of strengthening labelling functions.

The following proposition (proved in the Appendix) relates different strengthening of labelling functions, which is used later in the proof of the main theorems.

Proposition 1. *Given $\varepsilon_2 \geq \varepsilon_1 \geq 0$, let $\mathcal{L}_{\varepsilon_1}$ be the ε_1-strengthening of a labelling function $\mathcal{L} : \mathbb{R}^n \to 2^{\Pi}$, $(\mathcal{L}_{\varepsilon_1})_{\varepsilon_2}$ be the ε_2-strengthening of $\mathcal{L}_{\varepsilon_1}$, and $\mathcal{L}_{\varepsilon_1+\varepsilon_2}$ be the $(\varepsilon_1+\varepsilon_2)$-strengthening of $\mathcal{L}$. Then $\mathcal{L}_{\varepsilon_1+\varepsilon_2}(x) \subseteq (\mathcal{L}_{\varepsilon_1})_{\varepsilon_2}(x)$ for all $x \in \mathbb{R}^n$.*

2.4 Robust decidability of sampled-data control

Given a temporal logic formula φ together with a labelling function $\mathcal{L}$, we would like to design a sampled-data control strategy such that the resulting continuous-time state trajectories of $\mathcal{S}_\delta$ satisfy $(\varphi, \mathcal{L})$. If such a control strategy exists, we say $(\varphi, \mathcal{L})$ is *realizable* for $\mathcal{S}_\delta$ (by a sampled-data control strategy).

We formulate the robust decidability problem for system (1) as follows.

Problem 1 (Robust decidability). Given a temporal logic formula φ, a labelling function $\mathcal{L}$, a sampling period $T > 0$, numbers $\delta_2 > \delta_1 \geq 0$ and $\varepsilon > 0$, decide which one of the following is true:

- There exists (and one can algorithmically construct) a sampled-data control strategy with sampling period T for $\mathcal{S}_{\delta_1}$ to realize the specification $(\varphi, \mathcal{L})$;
- There does not exist a sampled-data control strategy with sampling period T for $\mathcal{S}_{\delta_2}$ to realize the specification $(\varphi, \mathcal{L}_\varepsilon)$.

We shall answer this question under the following assumption.

Assumption 1 *The sets X and U are compact and f is locally Lipschitz in both x and u. Let L be the Lipschitz constant (in infinity norm) of f w.r.t. both x and u on $X \times U$.*

3 Transition Systems and Finite Abstractions

In this section, we define finite abstractions of $\mathcal{S}_\delta$ that can be used to synthesize sampled-data control strategies for $\mathcal{S}_\delta$. Due to space limit, the proofs of the preliminary results (Propositions 2–4) are relegated to the Appendix in [19].

3.1 Transition systems

Definition 1. A *transition system* is a tuple $\mathcal{T} = (Q, A, R)$, where

- Q is the set of states and A is the set of actions;
- $R \subseteq Q \times A \times Q$ is the transition relation.

For each action $a \in A$ and $q \in Q$, we define the a-successor of q by

$$\mathrm{Post}_{\mathcal{T}}(q, a) = \{q' : q' \in Q \text{ s.t. } (q, a, q') \in R\}.$$

To simplify the presentation, we assume in this paper that, for the transition systems under consideration, every action is admissible for every state in the sense that $\mathrm{Post}_{\mathcal{T}}(q, a) \neq \emptyset$ for all $q \in Q$ and all $a \in A$.

An *execution* of $\mathcal{T}$ is an infinite alternating sequence of states and actions $\rho = q_0, a_0, q_1, a_1, q_2, a_2, \cdots$, where q_0 is some initial state and $(q_i, a_i, q_{i+1}) \in R$ for all $i \geq 0$. The *path* resulting from the execution ρ above is the sequence $\mathrm{Path}(\rho) = q_0, q_1, q_2 \cdots$. A *control strategy* κ for a transition system $\mathcal{T}$ is a partial function $\kappa : (q_0, q_1, \cdots, q_i) \mapsto a_i$ that maps the state history to the next action. An *κ-controlled execution* of a transition system $\mathcal{T}$ is an execution of $\mathcal{T}$, where for each $i \geq 0$, the action a_i is chosen according to the control strategy κ; κ-controlled paths are defined in a similar fashion. A dwell-time control strategy is defined in the same way as that for $\mathcal{S}_\delta$ in (3).

3.2 Transition systems for sampled-data control systems

With a fixed sampling period $\tau > 0$, we define the transition system representation of $\mathcal{S}_\delta$ as follows.

Definition 2. *The system $\mathcal{S}_\delta$ with a sampling period $\tau > 0$ can be interpreted as a transition system $\mathcal{T}_{\delta,\tau} = (Q, A, R)$, by defining*

- *$Q = X$ and $A = U$;*
- *$(x_0, u, x_1) \in R$ if and only if there exists a trajectory $\mathbf{x} : [0, \tau] \to X$ such that $x(0) = x_0$, $x_1 = x(\tau)$, and $\mathbf{x}'(s) \in f(\mathbf{x}(s), u) + \delta\mathbb{B}$ for all $s \in [0, \tau]$.*

If $\delta = 0$, we simply write $\mathcal{T}_{\delta,\tau}$ as $\mathcal{T}_\tau$.

We say that an execution ρ of $\mathcal{T}_{\delta,\tau}$ satisfies an $\mathrm{LTL}_{\setminus\bigcirc}$ formula φ with a labelling function $\mathcal{L}$, written as $\rho \vDash (\varphi, \mathcal{L})$, if and only if $\mathrm{Path}(\rho) \vDash (\varphi, \mathcal{L})$. For a control strategy κ for $\mathcal{T}_{\delta,\tau}$, if all κ-controlled executions of $\mathcal{T}_{\delta,\tau}$ satisfy φ with respect to L, we write $(\mathcal{T}_{\delta,\tau}, \kappa) \vDash (\varphi, \mathcal{L})$. If such a control strategy κ exists, we say that $(\varphi, \mathcal{L})$ is *realizable* for $\mathcal{T}_{\delta,\tau}$.

The following proposition relates realizability of a temporal logic formula φ on a continuous-time control system with sampled-data control strategies of different sampling periods.

Proposition 2. *Let φ be a temporal logic formula over Π and $\mathcal{L} : X \to 2^\Pi$ be a labelling function. Suppose that $T = N\tau$, where N is a positive integer.*

1. *If $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}_\delta$ with a sampled-data control strategy with sampling period T, then $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}_\delta$ with a sampled-data control strategy with sampling period τ and dwell time N.*
2. *Conversely, if $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}_\delta$ with a sampled-data control strategy with sampling period τ and dwell time N, then $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}_\delta$ with a sampled-data control strategy with sampling period T.*

By the assumption that X and U are compact sets, we can define $M = \max_{x \in X,\, u \in U} |f(x, u)|$, where $|\cdot|$ is the infinity norm (throughout the paper). The following proposition relates realizability of a temporal logic formula φ on a sampled-data transition system $(\mathcal{T}_{\delta,\tau})$ and a continuous-time system $(\mathcal{S}_\delta)$. The main technical part is to show how discrete-time and continuous-time semantics of temporal logic formulas imply each other.

Proposition 3 (Inter-sample correctness). *Let φ be a temporal logic formula over Π. Let $\mathcal{L} : X \to 2^{\Pi}$ be a labelling function and $\mathcal{L}_\varepsilon$ be an ε-strengthening of $\mathcal{L}$. Suppose that $\varepsilon \geq (M + \delta)\tau/2$.*

1. *If $(\varphi, \mathcal{L}_\varepsilon)$ is realizable for $\mathcal{T}_{\delta,\tau}$ with a dwell-time N control strategy, then $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}_\delta$ with a sampled-data control strategy with sampling period τ and dwell-time N.*
2. *Conversely, if $(\varphi, \mathcal{L}_\varepsilon)$ is realizable for $\mathcal{S}_\delta$ with a sampled-data control strategy with sampling period τ and dwell-time N, then $(\varphi, \mathcal{L})$ is realizable for $\mathcal{T}_{\delta,\tau}$ with a dwell-time N control strategy.*

3.3 Abstraction

We define control abstraction of transition system that preserves realizability of temporal logic specifications.

Definition 3. Given two transition systems $\mathcal{T}_i = (Q_i, A_i, R_i)$, $i = 1, 2$, a relation $\alpha \subseteq Q_1 \times Q_2$ is said to be an *abstraction* from $\mathcal{T}_1$ to $\mathcal{T}_2$, if the following conditions are satisfied:

(i) for all $q_1 \in Q_1$, there exists $q_2 \in Q_2$ such that $(q_1, q_2) \in \alpha$ (i.e., $\alpha(q_1) \neq \emptyset$);
(ii) for $a_2 \in A_2$, there exists $a_1 \in A_1$ such that, for all $q_2 \in Q_2$ and $q_1 \in \alpha^{-1}(q_2)$,

$$\alpha(\mathrm{Post}_{\mathcal{T}_1}(q_1, a_1)) \subseteq \mathrm{Post}_{\mathcal{T}_2}(q_2, a_2). \tag{5}$$

If such a relation α exists, we say that $\mathcal{T}_2$ *abstracts* $\mathcal{T}_1$ and write $\mathcal{T}_1 \preceq_\alpha \mathcal{T}_2$ or simply $\mathcal{T}_1 \preceq \mathcal{T}_2$. When both Q_1 and Q_2 are subsets of $\mathbb{R}^n$, we say that α is of granularity $\eta > 0$, if for every $q_2 \in Q_2$, $\alpha^{-1}(q_2) \subseteq q_2 + \eta\mathbb{B}$.

The following proposition shows that the abstraction relation defined above is sound in the sense of preserving realization of temporal logic specifications.

Proposition 4 (Soundness). *Consider transition systems $\mathcal{T}_1 = (Q_1, A_1, R_1)$ and $\mathcal{T}_2 = (Q_2, A_2, R_2)$ such that $\mathcal{T}_1 \preceq_\alpha \mathcal{T}_2$. Suppose that Q_1 and Q_2 are subsets of $X \subseteq \mathbb{R}^n$. Let $\mathcal{L} : X \to 2^{\Pi}$ be a labelling function. Let N be a positive integer.*

- *Suppose that α is proposition preserving with respect to L, defined as $L(q_2) \subseteq L(q_1)$ for all $(q_1, q_2) \in \alpha$. Then $(\varphi, \mathcal{L})$ is realizable for $\mathcal{T}_2$ implies that $(\varphi, \mathcal{L})$ is realizable for $\mathcal{T}_1$.*
- *Suppose that α is of granularity $\eta > 0$ and let $\mathcal{L}_\eta$ denote an η-strengthening of $\mathcal{L}$. Then $(\varphi, \mathcal{L}_\eta)$ is realizable for $\mathcal{T}_2$ implies that $(\varphi, \mathcal{L})$ is realizable for $\mathcal{T}_1$.*

Moreover, a dwell-time N strategy for $\mathcal{T}_2$ can be implemented by a dwell-time N strategy for $\mathcal{T}_1$.

4 Robustly Complete Abstraction and Robust Decidability

In this section, we present the main results on robustly complete abstractions and robust decidability of continuous-time control via discrete abstractions.

4.1 Robustly complete abstraction

The key technical result for proving robustly decidability of sampled-data control for nonlinear system is the following result on the possibility of constructing an arbitrarily accurate abstraction of the nonlinear system in the sense that for any $\delta_2 > \delta_1 \geq 0$, one can find a finite transition system $\mathcal{T}$ such that $\mathcal{T}$ abstracts $\mathcal{S}_{\delta_1}$ while $\mathcal{S}_{\delta_2}$ abstracts $\mathcal{T}$. Hence, realizability of a specification by $\mathcal{S}_{\delta_2}$ would imply realizability of the same specification by $\mathcal{S}_{\delta_1}$.

Theorem 1 (Robust completeness). *Given any $\delta_2 > \delta_1 \geq 0$, we can choose $\tau > 0$ and compute a finite transition system $\mathcal{T}$ such that $\mathcal{T}_{\delta_1,\tau} \preceq \mathcal{T} \preceq \mathcal{T}_{\delta_2,\tau}$.*

Proof. We construct $\mathcal{T} = (Q, A, R)$ as follows. Let $\eta > 0$ and $\mu > 0$ be parameters to be chosen. Let Q consist of the centres of the grid cells in $[\mathbb{R}^n]_\eta$ that have a non-empty intersection with X. Let A consist of the centres of the grid cells in $[\mathbb{R}^m]_\mu$ that have a non-empty intersection with U. Because U and X are compact sets, Q and A are both finite. We define a relation $\alpha \subseteq X \times Q$ by $(x, q) \in \alpha$ if and only if $|x - q| \leq \frac{\eta}{2}$. Clearly, α^{-1} is a relation on $Q \times X$. Define $R \subseteq (Q, A, Q)$ by $(q, a, q_1) \in R$ if and only if

$$|q_1 - (q + \tau f(q, a))| \leq \frac{\eta}{2} + \frac{\eta}{2} e^{L\tau} + (\frac{\delta_1}{L} + \frac{\mu}{2})(e^{L\tau} - 1) + \frac{M(e^{L\tau} - L\tau - 1)}{L}. \tag{6}$$

We show that, if η, μ, and τ are chosen sufficiently small, we have $\mathcal{T}_{\delta_1,\tau} \preceq_\alpha \mathcal{T} \preceq_{\alpha^{-1}} \mathcal{T}_{\delta_2,\tau}$. Condition (i) in Definition 3 is clearly satisfied by both α and α^{-1}.

We verify that condition (ii) holds for $\mathcal{T}_{\delta_1,\tau} \preceq_\alpha \mathcal{T}$, that is, for $q \in Q$ and $a \in A$, there exists $u \in U$ such that

$$\alpha(\mathrm{Post}_{\mathcal{T}_{\delta_1,\tau}}(x, u)) \subseteq \mathrm{Post}_{\mathcal{T}}(q, a); \tag{7}$$

for all $x \in \alpha^{-1}(q)$. Pick $u \in U$ with $|u - a| \leq \frac{\mu}{2}$. Given $x_1 \in \mathrm{Post}_{\mathcal{T}_{\delta_1,\tau}}(x, u)$, there exists a trajectory $\mathbf{x} : [0, \tau] \to X$ such that $\mathbf{x}(0) = x$, $\mathbf{x}(\tau) = x_1$, and $\mathbf{x}'(s) \in f(\mathbf{x}(s), u) + \delta_1 \mathbb{B}$ for all $s \in [0, \tau]$. Define $\mathbf{x}_\tau(t) = q + tf(q, a)$ for $t \in [0, \tau]$. We have

$$\begin{aligned}
&|\mathbf{x}'(t) - \mathbf{x}'_\tau(t)| \leq |f(\mathbf{x}(t), u) - f(q, a)| + \delta_1 \\
&\leq |f(\mathbf{x}(t), u) - f(\mathbf{x}_\tau(t), u)| + |f(\mathbf{x}_\tau(t), u) - f(q, u)| + |f(q, u) - f(q, a)| + \delta_1 \\
&\leq L\,|\mathbf{x}(t) - \mathbf{x}_\tau(t)| + L\,|\mathbf{x}_\tau(t) - q| + L\,|u - a| + \delta_1 \\
&\leq L\,|\mathbf{x}(t) - \mathbf{x}_\tau(t)| + LMt + \frac{L\mu}{2} + \delta_1, \quad t \in [0, \tau].
\end{aligned} \tag{8}$$

By Gronwall's inequality (see, e.g., [2]), we have

$$\begin{aligned}
&|x_1 - (q + \tau f(q, u))| = |\mathbf{x}(\tau) - \mathbf{x}_\tau(\tau)| \\
&\leq |x - q|\, e^{L\tau} + \int_0^\tau (LMs + \frac{L\mu}{2} + \delta_1) e^{L(\tau - s)} ds \\
&\leq \frac{\eta}{2} e^{L\tau} + (\frac{\delta_1}{L} + \frac{\mu}{2})(e^{L\tau} - 1) + \frac{M(e^{L\tau} - L\tau - 1)}{L}.
\end{aligned}$$

By (6), this shows $\alpha(x_1) \subseteq \mathrm{Post}_{\mathcal{T}}(q,a)$. Hence (7) holds.

We next verify that condition (ii) holds for $\mathcal{T} \preceq_{\alpha^{-1}} \mathcal{T}_{\delta_2,\tau}$, that is, for $x \in X$ and $u \in U$, there exists $a \in A$ such that

$$\alpha^{-1}(\mathrm{Post}_{\mathcal{T}}(q,a)) \subseteq \mathrm{Post}_{\mathcal{T}_{\delta_2,\tau}}(x,u); \tag{9}$$

for all $q \in \alpha(x)$. Pick a be the center of the grid cell in $[\mathbb{R}^m]_\mu$ that contains u. Given $y_1 \in \alpha^{-1}(\mathrm{Post}_{\mathcal{T}}(q,a))$, there exists $q_1 \in \mathrm{Post}_{\mathcal{T}}(q,a)$ such that $|y_1 - q_1| \leq \frac{\eta}{2}$. By the definition of $\mathrm{Post}_{\mathcal{T}}(q,a)$, we have

$$|q_1 - (q + \tau f(q,a))| \leq \frac{\eta}{2} + \frac{\eta}{2}e^{L\tau} + (\frac{\delta_1}{L} + \frac{\mu}{2})(e^{L\tau} - 1) + \frac{M(e^{L\tau} - L\tau - 1)}{L}.$$

Consider the trajectory $\mathbf{x} : [0,\tau] \to X$ such that $\mathbf{x}(0) = x$, $\mathbf{x}(\tau) = x_1$, and $\mathbf{x}'(s) \in f(\mathbf{x}(s),u)$. By a similar argument as in (8), we can show

$$|x_1 - (q + \tau f(q,a))| \leq \frac{\eta}{2}e^{L\tau} + \frac{\mu}{2}(e^{L\tau} - 1) + \frac{M(e^{L\tau} - L\tau - 1)}{L}.$$

Hence, by the triangle inequality,

$$|y_1 - x_1| \leq \eta + \eta e^{L\tau} + (\frac{\delta_1}{L} + \mu)(e^{L\tau} - 1) + \frac{2M(e^{L\tau} - L\tau - 1)}{L} \tag{10}$$

Define

$$\mathbf{z}(\theta) = \mathbf{x}(\theta) + \frac{\theta}{\tau}(y_1 - x_1), \quad \theta \in [0,\tau].$$

Then $\mathbf{z}(0) = \mathbf{x}(0) = x$ and $\mathbf{z}(\tau) = y_1$, and

$$\mathbf{z}'(\theta) \in f(\mathbf{x}(\theta),u) + \frac{1}{\tau}(y_1 - x_1). \tag{11}$$

Note that

$$|\mathbf{z}(\theta) - \mathbf{x}(\theta)| = \left|\frac{\theta}{\tau}[y_1 - x_1]\right| \leq |y_1 - x_1|, \quad \theta \in [0,\tau]. \tag{12}$$

Since $0 \leq \delta_1 < \delta_2$, we can choose τ, μ, η sufficiently small such that

$$[\eta + \eta e^{L\tau} + (\frac{\delta_1}{L} + \mu)(e^{L\tau} - 1) + \frac{2M(e^{L\tau} - L\tau - 1)}{L}][L + \frac{1}{\tau}] < \delta_2. \tag{13}$$

To see this is possible, choose, e.g., $\eta = \tau^2$ and $\mu = \tau$, and note that the limit of the left-hand side as $\tau \to 0$ is given by $\lim_{\tau\to 0} \delta_1 \frac{e^{L\tau}-1}{L\tau} = \delta_1$. It follows from (10)–(13) and Lipschitz continuity of f that $\mathbf{z}'(\theta) \in f(\mathbf{z}(\theta),u) + \delta_2 B$. Hence $y_1 \in \mathrm{Post}_{\mathcal{T}_{\delta_2,\tau}}(x,u)$ and (9) holds. □

Remark 1. In the proof, we choose the simplest possible validated bounds on a one-step reachable set, i.e., a forward Euler scheme with an error bound. This suffices to prove the required convergence to show approximate completeness. With the template provided by the proof of Theorem 1, one can in fact use any accurate over-approximation of the one-step reachable set for $\mathcal{S}_{\delta_1}$ to replace (6) for defining the transitions in $\mathcal{T}$ and then show that this over-approximation is contained in the actual one-step reachable set of $\mathcal{S}_{\delta_2}$.

Remark 2. Theorem 1 (as well as the problem formulation in the paper) only considers sampled-data control strategies. Later in Section 4.3, we also discuss how to approximate arbitrary measurable control signals under the L^1 norm, which plays a role later in proving a notion of completeness via controllability.

4.2 Robust decidability

The following theorem is an immediate consequence of Theorem 1 and states that sampled-data control for nonlinear system is robustly decidable.

Theorem 2 (Robust decidability). *Given a temporal logic specification φ, a sampling period $T > 0$, any $\delta_2 > \delta_1 \geq 0$, and any $\varepsilon > 0$. Let $\mathcal{L} : X \to 2^{\Pi}$ be a labelling function and $\mathcal{L}_\varepsilon$ be an ε-strengthening of $\mathcal{L}$. Then there exists a decision procedure that determines which one of the following holds:*

- *there exists (and one can algorithmically construct) a sampled-data control strategy with sampling period T such that $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}_{\delta_1}$; or*
- *$(\varphi, \mathcal{L}_\varepsilon)$ is not realizable for $\mathcal{S}_{\delta_2}$ with a sampled-data control strategy with sampling period T.*

Proof. Suppose that $(\varphi, \mathcal{L}_\varepsilon)$ is realizable for $\mathcal{S}_{\delta_2}$ with a sampled-data control strategy with sampling period T. Let N be a positive integer and $\tau = \frac{T}{N}$. Let $\varepsilon_1 = \frac{(M+\delta_1)\tau}{2}$ and $\varepsilon_2 = \frac{(M+\delta_2)\tau}{2}$. Choose τ sufficiently small such that

$$\frac{(2M + \delta_1 + \delta_2)\tau}{2} = \varepsilon_1 + \varepsilon_2 \leq \varepsilon. \tag{14}$$

Let $\mathcal{L}_{\varepsilon_1}$ be the ε_1-strengthening of $\mathcal{L}$ and $(\mathcal{L}_{\varepsilon_1})_{\varepsilon_2}$ denote the ε_2-strengthening of $\mathcal{L}_{\varepsilon_1}$. Let $\mathcal{L}_{\varepsilon_1+\varepsilon_2}$ be the $(\varepsilon_1+\varepsilon_2)$-strengthening of $\mathcal{L}$. By the definition of strengthening a labeling function and Proposition 1, we have $\mathcal{L}_\varepsilon(x) \subseteq \mathcal{L}_{\varepsilon_1+\varepsilon_2}(x) \subseteq (\mathcal{L}_{\varepsilon_1})_{\varepsilon_2}(x)$ for all $x \in X$. Hence, by the semantics of $\mathrm{LTL}_{\setminus\bigcirc}$, $(\varphi, (\mathcal{L}_{\varepsilon_1})_{\varepsilon_2})$ is realizable for $\mathcal{S}_{\delta_2}$ with a sampled-data control strategy with sampling period T.

By Proposition 2, $(\varphi, (\mathcal{L}_{\varepsilon_1})_{\varepsilon_2})$ is realizable for $\mathcal{S}_{\delta_2}$ with a sampled-data control strategy with sampling period τ and dwell-time N. By Proposition 3, $(\varphi, \mathcal{L}_{\varepsilon_1})$ is realizable for $\mathcal{T}_{\delta_2,\tau}$ with a dwell-time N control strategy, because $\varepsilon_2 \geq \frac{(M+\delta_2)\tau}{2}$ (indeed equal). Construct $\mathcal{T}$ by Theorem 1 so that $\mathcal{T}_{\delta_1,\tau} \preceq \mathcal{T} \preceq \mathcal{T}_{\delta_2,\tau}$. By Proposition 4, $(\varphi, \mathcal{L}_{\varepsilon_1})$ is realizable for $\mathcal{T}$ and hence also for $\mathcal{T}_{\delta_1,\tau}$ with a dwell-time N control strategy. By Proposition 3, $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}_{\delta_1}$ with a sampled-data control strategy with sampling period τ and dwell-time N, because $\varepsilon_1 \geq \frac{(M+\delta_1)\tau}{2}$. Finally, by Proposition 2 again, $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}_{\delta_1}$ with a sampled-data control strategy with sampling period T. One can algorithmically construct such a control strategy by synthesizing a dwell-time N controller strategy for the finite transition system $\mathcal{T}$. For the case there is not necessarily a proposition preserving partition, we can choose $\varepsilon_1 = \frac{(M+\delta_1)\tau+\eta}{2}$ and $\varepsilon_2 = \frac{(M+\delta_2)\tau+\eta}{2}$ to account for mismatch by an abstraction with granularity $\frac{\eta}{2}$. In this case, we can choose τ and η sufficiently small such that

$$\eta + \frac{(2M + \delta_1 + \delta_2)\tau}{2} = \varepsilon_1 + \varepsilon_2 \leq \varepsilon. \tag{15}$$

On the other hand, by this construction of $\mathcal{T}$, if $(\varphi, \mathcal{L}_{\varepsilon_1})$ is not realizable for $\mathcal{T}$, then we can conclude that $(\varphi, \mathcal{L}_{\varepsilon})$ is not realizable for $\mathcal{S}_{\delta_2}$ with a sampled-data control strategy with sampling period T. □

A decision diagram summarizing the argument in the proof of Theorem 2 can be found in the Appendix of [19]. When there is no *a priori* fixed sampling period for the decision process, we can formulate the robust decidability theorem as follows, where it is proved that the problem can be solved for all sufficiently small sampling periods. The proof follows exactly from the proof of Theorem 2 with $N = 1$.

Theorem 3 (Robust decidability II). *Given a temporal logic specification φ, any $\delta_2 > \delta_1 \geq 0$, and any $\varepsilon > 0$. Let $\mathcal{L} : X \to 2^{\Pi}$ be a labelling function and $\mathcal{L}_{\varepsilon}$ be an ε-strengthening of $\mathcal{L}$. Then there exists some $\tau^* > 0$ (and one can explicitly compute it) such that, for each $\tau \in (0, \tau^*]$, there exists a decision procedure that determines which one of the following holds:*

- *there exists (and one can algorithmically construct) a sampled-data control strategy with sampling period τ such that $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}_{\delta_1}$; or*
- *$(\varphi, \mathcal{L}_{\varepsilon})$ is not realizable for $\mathcal{S}_{\delta_2}$ with a sampled-data control strategy with sampling period τ.*

Another version of robust decidability can be formulated as follows, which says that with one procedure, one can decide robust realizability by a sampled-data control strategy with *any* sampling period greater than a threshold value (e.g., a lower bound limited by the physical sampling frequency).

Theorem 4 (Robust decidability III). *Given a temporal logic specification φ, any $\delta_2 > \delta_1 \geq 0$, $\varepsilon > 0$, and $\tau^* > 0$. Let $\mathcal{L} : X \to 2^{\Pi}$ be a labelling function and $\mathcal{L}_{\varepsilon}$ be an ε-strengthening of $\mathcal{L}$. Then there exists some $\tau > 0$ (and one can explicitly compute it) and a decision procedure that determines which one of the following holds:*

- *there exists (and one can algorithmically construct) a sampled-data control strategy with sampling period τ such that $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}_{\delta_1}$; or*
- *$(\varphi, \mathcal{L}_{\varepsilon})$ is not realizable for $\mathcal{S}_{\delta_2}$ with a sampled-data control strategy with a sampling period $T \geq \tau^*$.*

To prove Theorem 4, we need the following lemma, which shows that, if $\delta_2 > \delta_1$, then system $\mathcal{T}_{\delta_1,\tau}$ can be abstracted by $\mathcal{T}_{\delta_2,\tau'}$ despite a slight mismatch between the sampling periods τ and τ'.

Lemma 1. *Given any $\tau^* > 0$ and $\delta_2 > \delta_1 \geq 0$, there exists $r^* > 0$ such that $\mathcal{T}_{\delta_1,T} \preceq_{id_X} \mathcal{T}_{\delta_2,T+r}$ for all $T \geq \tau^*$ and all $|r| \leq r^*$, where $id_X \subseteq X \times X$ is the identity relation.*

Proof. Choose any $x \in X$ and $u \in U$. Let $x_1 \in \mathrm{Post}_{\mathcal{T}_{\delta_1,T}}(x, u)$. We show that $x_1 \in \mathrm{Post}_{\mathcal{T}_{\delta_2,T+r}}(x, u)$. By definition, there exists a trajectory $\mathbf{x}$ such that $\mathbf{x}(0) =$

x, $\mathbf{x}(T) = x_1$, and $\mathbf{x}'(s) \in f(\mathbf{x}(s), u) + \delta_1 \mathbb{B}$ for all $s \in [0, T]$. Let $\mathbf{z}(s) = \mathbf{x}(\frac{T}{T+r}s)$ for $s \in [0, T]$. Then $\mathbf{z}(0) = x$, $\mathbf{z}(T + r) = x_1$ and

$$\begin{aligned}\mathbf{z}'(s) &= \frac{T}{T+r}\mathbf{x}'(\frac{T}{T+r}s) \in \frac{T}{T+r}f(\mathbf{z}(s), u) + \frac{T}{T+r}\delta_1\mathbb{B} \\ &\subseteq f(\mathbf{z}(s), u) - \frac{r}{T+r}f(\mathbf{z}(s), u) + \delta_1\mathbb{B} \\ &\subseteq f(\mathbf{z}(s), u) + (\frac{|r|\, M}{\tau^* - |r|} + \delta_1)\mathbb{B},\end{aligned}$$

where we assumed $|r|$ is sufficiently small so that $|r| \leq \tau^*$. Clearly, since $\delta_1 < \delta_2$, we can choose $r^* > 0$ so that $\frac{|r|M}{\tau^* - |r|} + \delta_1 < \delta_2$ for all $|r| \leq r^*$. Hence, $\mathbf{z}'(s) \in f(\mathbf{z}(s), u) + \delta_2\mathbb{B}$ and $x_1 = \mathbf{z}(T + r) \in \mathrm{Post}_{\mathcal{T}_{\delta_2, T+r}}(x, u)$. □

Now we can present the proof of Theorem 4.

Proof (Proof of Theorem 4).

Let ε_1 and ε_2 be as defined in the proof for Theorem 2. Choose δ_3 such that $\delta_2 > \delta_3 > \delta_1$. Let τ^*, η^*, and μ^* be chosen so that (13) and (14) (or (13) and (15) if a proposition preserving partition is not used), with δ_3 replacing δ_2 in (13), hold for all $\tau \leq \tau^*$, $\eta \leq \eta^*$, and $\mu \leq \mu^*$.

Suppose that $(\varphi, \mathcal{L}_\varepsilon)$ is realizable for $\mathcal{S}_{\delta_2}$ with a sampled-data control strategy with sampling period T. Without loss of generality, assume $\frac{\tau^*}{2} < T \leq \tau^*$. Otherwise, one can divide T by a positive integer number N so that $\frac{T}{N} \in (\tau^*/2, \tau^*]$ and $(\varphi, \mathcal{L}_\varepsilon)$ is realizable for $\mathcal{S}_{\delta_2}$ with a sampled-data control strategy with sampling period T/N (with dwell-time N).

Construct, by Theorem 1, $\mathcal{T}$ so that

$$\mathcal{T}_{\delta_3, T} \preceq \mathcal{T} \preceq \mathcal{T}_{\delta_2, T}. \tag{16}$$

Let $\tau \leq \frac{\tau^*}{2}$ be chosen (guaranteed by Lemma 1) so that

$$\mathcal{T}_{\delta_1, T+r} \preceq_{\mathrm{id}_X} \mathcal{T}_{\delta_3, T}. \tag{17}$$

for all $|r| \leq \tau$.

Let $\mathcal{L}_{\varepsilon_1}$, $(\mathcal{L}_{\varepsilon_1})_{\varepsilon_2}$ and $\mathcal{L}_{\varepsilon_1+\varepsilon_2}$ be as defined in the proof for Theorem 2. By Proposition 3, $(\varphi, \mathcal{L}_{\varepsilon_1})$ is realizable for $\mathcal{T}_{\delta_2, T}$, because $\varepsilon_2 \geq \frac{(M+\delta_2)T}{2}$. By Proposition 4 and (16), $(\varphi, \mathcal{L}_{\varepsilon_1})$ is realizable for $\mathcal{T}$ and hence also for $\mathcal{T}_{\delta_3, T}$. Let m be the largest integer such that $m\tau \leq T$. Then $|m\tau - T| \leq \tau$. By (17), we obtain $\mathcal{T}_{\delta_1, m\tau} = \mathcal{T}_{\delta_1, T+(m\tau - T)} \preceq_{\mathrm{id}_X} \mathcal{T}_{\delta_3, T}$. By Proposition 4 again, $(\varphi, \mathcal{L}_{\varepsilon_1})$ is realizable for $\mathcal{T}_{\delta_3, m\tau}$. By Proposition 3, $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}_{\delta_1}$ with a sampled-data control strategy with sampling period $m\tau$, because $\varepsilon_1 \geq \frac{(M+\delta_1)m\tau}{2}$. Finally, by Proposition 2, $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}_{\delta_1}$ with a sampled-data control strategy with sampling period τ. One can algorithmically construct such a control strategy by synthesizing a control strategy for $\mathcal{T}$ to realize $(\varphi, \mathcal{L}_{\varepsilon_1})$. □

Remark 3. Theorem 4 is sharp in the sense that, if one does not impose a lower bound on the sampling period, it is possible to show that a one-dimensional

continuous-time control system can robustly simulate any Turing machine. Since the undecidable Halting problem can be encoded as a reachability problem, one cannot expect to have a general decision procedure in this setting. Details can be found in the Appendix of [19].

4.3 Dealing with arbitrary control signals

As mentioned in Remark 2, we formulated the problem only for sampled-data control strategies. In this section, we prove a lemma that shows this is assumption is without loss of generality. An input $\mathbf{u} : \mathbb{R}^+ \to U$ is called a sampled-data input with period $\tau > 0$ if there exists $t_0 = 0, \cdots, t_i, \cdots$ such that it is constant on each $[t_i, t_{i+1})$ and $t_{i+1} - t_i = \tau$.

Lemma 2. *Let $\mathbf{u} : [0,T] \to U$ be a measurable signal. Then for $\mu > 0$, there exists some $\tau > 0$, a finite subset of $[U]_\mu$ of U, and a sampled-data input $\mathbf{v} : [0,T] \to [U]_\mu$ with period greater than τ such that $\int_0^T |\mathbf{u}(s) - \mathbf{v}(s)|\, ds \leq \mu$.*

Proof. Since U is compact, $\mathbf{u} \in L^1([0,T], U)$. The conclusion follows from the fact that step functions defined on rational partitions and taking rational values are dense in $L^1([0,T], U)$ (see, e.g., [29, p. 152, Chapter 7]). □

4.4 Completeness via controllability

In this subsection, we show that under a suitable notion of local controllability around a trajectory that realizes a given temporal logic specification on the nominal control system (1), we can always construct a robust abstraction to find a sampled-data control strategy to robustly realize the same specification (subject to an ε-strengthening of the labelling function for an arbitrarily small ε). In other words, local controllability suffices to remove the robustness relaxation in our notion of completeness. We refer the readers to [9,23] for sufficient conditions on local controllability around a closed orbit [23] or reference trajectory [9].

Definition 4. *Let $\mathcal{O} \subseteq X$ be an open connected set. Let $T \geq 0$. A point X_T is $\mathcal{O}$-reachable from $x_0 \in X$ at time T if there exists a trajectory $(\mathbf{x}, \mathbf{u})$ such that $\mathbf{x}(0) = x_0$, $\mathbf{x}(T) = x_T$, and $\mathbf{x}(t) \in \mathcal{O}$ for all $t \in [0,T]$ and $\mathbf{u} : [0,T] \to U$ is a measurable signal. The set of all points X_T that are $\mathcal{O}$-reachable from x_0 at time T is denoted by $\mathcal{R}(x_0, T, \mathcal{O})$. Let $\mathcal{R}(x_0, \mathcal{O}) = \cup_{T \geq 0} \mathcal{R}(x_0, T, \mathcal{O})$. We say the control system (1) is controllable on $\mathcal{O}$ if $\mathcal{R}(x_0, \mathcal{O}) = \mathcal{O}$.*

Theorem 5 (Robust realizability via controllability). *Let φ be a temporal logic specification and $\mathcal{L} : X \to 2^{\Pi}$ be a labelling function. For any $\varepsilon > 0$, if there exists a trajectory $(\mathbf{x}, \mathbf{u})$ for (1) such that $\mathbf{x}$ satisfies $(\varphi, \mathcal{L}_\varepsilon)$ and (1) is controllable on an open set $\mathcal{O}$ containing $\mathbf{x}$, then there exists some $\delta > 0$ and a sampled-data control strategy for $\mathcal{S}_\delta$ to realize $(\varphi, \mathcal{L})$.*

Proof. For any $\varepsilon > 0$, if $\mathbf{x}$ satisfies $(\varphi, \mathcal{L}_\varepsilon)$, we can easily show that state trajectories that are ε-close to $\mathbf{x}$ satisfy $(\varphi, \mathcal{L})$. Hence, we only need to show that there exists a sampled-data control strategy for trajectories of $\mathcal{S}_\delta$ to stay in an ε-neighborhood of $\mathbf{x}$.

By the compactness of X, assume, without loss of generality, that the ε-neighborhood of the image of $\mathbf{x}$ is contained in $\mathcal{O}$. We construct a finite resolution sampled-data control strategy for $\mathcal{S}_\delta$ as follows. For any $x \in X$, let q be a grid point in $[X]_\eta$ such that $|x - q| \leq \frac{\eta}{2}$. Suppose that q is in an $\frac{\varepsilon}{2}$-ball around some point p on $\mathbf{x}$. Let $\mathbf{u}_1$ be a control signal that steers q to some q' on $\mathbf{x}$ at some time $T > 0$ according to dynamics of (1) without leaving an ε-ball of γ. This is always possible under the local controllability assumption. By Lemma 2, there exists a sampled-data signal $\mathbf{u}_2$ with period greater than some $\tau > 0$ such that $\int_0^T |\mathbf{u}_1(s) - \mathbf{u}_2(s)|\, ds \leq \mu$. We estimate the state trajectories of $\mathcal{S}_\delta$ under control of $\mathbf{v}$ as follows. Let $\mathbf{x}_1$ be the trajectory of the nominal system $\mathcal{S}$ under $\mathbf{u}_1$ starting from q and ending at q'. Let $\mathbf{x}_2$ be a trajectory of $\mathcal{S}_\delta$ starting from x. For $t \in [0, T]$, we have

$$\begin{aligned} |\mathbf{x}_1'(t) - \mathbf{x}_2'(t)| &\leq |f(\mathbf{x}_1(t), \mathbf{u}_1(t)) - f(\mathbf{x}_2(t), \mathbf{u}_2(t))| + \delta \\ &\leq |f(\mathbf{x}_1(t), \mathbf{u}_1(t)) - f(\mathbf{x}_2(t), \mathbf{u}_1(t))| \\ &\qquad + |f(\mathbf{x}_2(t), \mathbf{u}_1(t)) - f(\mathbf{x}_2, \mathbf{u}_2(t))| + \delta \\ &\leq L\,|\mathbf{x}_1(t) - \mathbf{x}_2(t)| + L\,|\mathbf{u}_1(t) - \mathbf{u}_2(t)| + \delta. \end{aligned} \tag{18}$$

By Gronwall's inequality (see, e.g., [2, p. 120]), we have

$$\begin{aligned} |\mathbf{x}_1(t) - \mathbf{x}_2(t)| &\leq |x - q|\, e^{Lt} + \int_0^t (L\,|\mathbf{u}_1(s) - \mathbf{u}_2(s)| + \delta) e^{L(t-s)} ds \\ &\leq \frac{\eta}{2} e^{Lt} + \frac{\delta}{L}(e^{Lt} - 1) + L\mu e^{Lt}, \quad \forall t \in [0, T]. \end{aligned}$$

From this estimate, we can see that for any fixed $T > 0$, we can choose η, $\delta > 0$, and $\mu > 0$ sufficient small such that the right-hand side of the above inequality is less than $\frac{\varepsilon}{2}$. It follows that the state trajectory $\mathbf{x}_2$ of $\mathcal{S}_\delta$ remains in a ε-neighborhood of $\mathbf{x}$ and, at $t = T$, is within a $\frac{\varepsilon}{2}$-neighborhood of q'. Since there is only a finite number of grid points of granularity $\frac{\eta}{2}$ in a compact domain, the choice of sampling period τ for the signals satisfying Lemma 2 can be fixed. Therefore, there exists a sampled-data control strategy for state trajectories of $\mathcal{S}_\delta$ to stay in an ε-neighborhood of $\mathbf{x}$, as long as they start in an $\frac{\varepsilon}{2}$-neighborhood of $x_0 = \mathbf{x}(0)$. Since $\mathbf{x}$ satisfies $(\varphi, \mathcal{L}_\varepsilon)$, all trajectories of $\mathcal{S}_\delta$ starting from the $\frac{\varepsilon}{2}$-neighborhood of x_0 satisfy $(\varphi, \mathcal{L})$. □

Remark 4. It is clear from the proof of Theorem 5 that the robust realizability result does not depend on the reference trajectory $\mathbf{x}$ satisfying (1), but on the path induced by the trajectory. Therefore, if system (1) has a control neighborhood that contains a path parameterizable by any trajectory satisfying the specification, the conclusion still holds. As a result, if a system is controllable in a neighborhood $\mathcal{O}$ and a specification is satisfiable by a path in $\mathcal{O}$, then the specification is robustly realizable by a sampled-data control strategy.

Combining the robust realizability result Theorem 5 and the result on robust completeness proved in previous sections, we can show that controllability implies completeness of discrete abstractions in the following sense.

Corollary 1 (Completeness via controllability). *Given a temporal logic specification φ and any $\varepsilon > 0$, let $\mathcal{L} : X \to 2^{\Pi}$ be a labelling function and $\mathcal{L}_\varepsilon$ be an ε-strengthening of $\mathcal{L}$. Suppose that $(\varphi, \mathcal{L}_\varepsilon)$ is satisfiable by a trajectory of (1) contained in a controllable neighborhood. Then, for any $\varepsilon' \in (0, \varepsilon)$, there exists $\delta > 0$, $\tau > 0$, and a finite transition system $\mathcal{T}$ such that $\mathcal{T}_\tau \preceq \mathcal{T} \preceq \mathcal{T}_{\delta,\tau}$ and $(\varphi, \mathcal{L}_{\varepsilon'})$ is realizable for $\mathcal{T}$. Moreover, a control strategy synthesized using $\mathcal{T}$ for $(\varphi, \mathcal{L}_{\varepsilon'})$ can be refined as a sampled-data control strategy for $\mathcal{S}$ that realizes $(\varphi, \mathcal{L})$.*

Proof. Choose any ε_1 and ε_2 such that $\varepsilon' < \varepsilon_1 < \varepsilon_2 < \varepsilon$. By Theorem 5, there exists $\delta > 0$ and $\tau > 0$ such that $(\varphi, \mathcal{L}_{\varepsilon_2})$ is realizable for $\mathcal{S}_\delta$. Suppose that τ is also chosen sufficiently small according to conditions of Proposition 3. Then $(\varphi, \mathcal{L}_{\varepsilon_1})$ is realizable for $\mathcal{T}_{\delta,\tau}$. The existence of $\mathcal{T}$ follows from Theorem 1. By Propositions 4, $(\varphi_{\varepsilon'}, \mathcal{L})$ is realizable for $\mathcal{T}$, provided that the time discretization τ and space discretization η (in the case of non-proposition preserving abstraction) are chosen sufficiently small. Continuing this reasoning for $\mathcal{T}_\tau$ and $\mathcal{S}$, we obtain that $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}$ with a sampled-data control strategy synthesized from the abstraction $\mathcal{T}$.

Remark 5. The completeness result Corollary 1 is stated with respect to the nominal system (1). It is clear from the reasoning in the proof that under the same assumption we can synthesize a robust sampled-data control strategy for $\mathcal{S}$ from the abstraction $\mathcal{T}$.

5 Conclusions

In this paper, we proved two sets of theoretical results on completeness of abstraction-based nonlinear control. First, we show that control synthesis for sampled-data nonlinear systems with temporal logic specifications is robustly decidable in the sense that if a robust control strategy exists, then a robust control strategy can be found using a sufficiently fine discretization. Second, we show that, under the assumption of nonlinear controllability around a trajectory that realizes a given specification, it is always possible to construct a sampled-data control strategy via a sufficiently fine discrete abstraction.

We see the main theoretical contributions of this work as showing the existence of robustly complete abstractions for continuous-time nonlinear control systems and that nonlinear controllability implies completeness. It is hoped that this work will motivate further research on computing tight abstractions of nonlinear control systems. In this regard, Theorem 1 on robust completeness can be viewed as a potential metric on closeness of abstractions and Corollary 1 can help provide assurance that control synthesis via discrete abstractions always works for controllable systems.

References

1. Angeli, D.: A lyapunov approach to incremental stability properties. IEEE Transactions on Automatic Control **47**(3), 410–421 (2002)
2. Aubin, J.P., Cellina, A.: Differential Inclusions: Set-valued Maps and Viability Theory. Springer (2012)
3. Baier, C., Katoen, J.P.: Principles of Model Checking. MIT press (2008)
4. Belta, C., Yordanov, B., Gol, E.A.: Formal Methods For Discrete-time Dynamical Systems. Springer (2017)
5. Clarke, E.M., Grumberg, O., Peled, D.: Model Checking. MIT Press (1999)
6. Fainekos, G.E., Girard, A., Kress-Gazit, H., Pappas, G.J.: Temporal logic motion planning for dynamic robots. Automatica **45**(2), 343–352 (2009)
7. Fainekos, G.E., Pappas, G.J.: Robustness of temporal logic specifications for continuous-time signals. Theoretical Computer Science **410**(42), 4262–4291 (2009)
8. Gradel, E., Thomas, W.: Automata, logics, and infinite games: a guide to current research, vol. 2500. Springer Science & Business Media (2002)
9. Hermes, H.: On local and global controllability. SIAM Journal on Control **12**(2), 252–261 (1974)
10. Kloetzer, M., Belta, C.: Temporal logic planning and control of robotic swarms by hierarchical abstractions. IEEE Transactions on Robotics **23**(2), 320–330 (2007)
11. Koiran, P., Cosnard, M., Garzon, M.: Computability with low-dimensional dynamical systems. Theoretical Computer Science **132**(1-2), 113–128 (1994)
12. Kress-Gazit, H., Wongpiromsarn, T., Topcu, U.: Correct, reactive, high-level robot control. IEEE Robotics & Automation Magazine **18**(3), 65–74 (2011)
13. Li, Y., Liu, J.: Invariance control synthesis for switched nonlinear systems: An interval analysis approach. IEEE Transactions on Automatic Control **63**(7), 2206–2211 (2018)
14. Li, Y., Liu, J.: Robustly complete reach-and-stay control synthesis for switched systems via interval analysis. In: Proc. of ACC (2018)
15. Li, Y., Liu, J.: Rocs: A robustly complete control synthesis tool for nonlinear dynamical systems. In: Proc. of HSCC. pp. 130–135 (2018)
16. Li, Y., Liu, J.: Robustly complete synthesis of memoryless controllers for nonlinear systems with reach-and-stay specifications. IEEE Transactions on Automatic Control **66**(3), 1199–1206 (2021)
17. Li, Y., Sun, Z., Liu, J.: A specification-guided framework for temporal logic control of nonlinear systems. arXiv preprint arXiv:2104.01385 (2021)
18. Liu, J.: Robust abstractions for control synthesis: completeness via robustness for linear-time properties. In: Proc. of HSCC. pp. 101–110. ACM (2017)
19. Liu, J.: Closing the gap between discrete abstractions and continuous control: Completeness via robustness and controllability. In: Proc. of FORMATS (2021), Extended version with Appendix available at: `https://www.math.uwaterloo.ca/~j49liu/papers/2021/liu2021closing.pdf`
20. Liu, J., Ozay, N., Topcu, U., Murray, R.: Synthesis of reactive switching protocols from temporal logic specifications. IEEE Trans. on Automatic Control **58**(7), 1771–1785 (2013)
21. Liu, J., Ozay, N.: Abstraction, discretization, and robustness in temporal logic control of dynamical systems. In: Proc. of HSCC. pp. 293–302 (2014)
22. Liu, J., Ozay, N.: Finite abstractions with robustness margins for temporal logic-based control synthesis. Nonlinear Analysis: Hybrid Systems **22**, 1–15 (2016)

23. Nam, K., Arapostathis, A.: A sufficient condition for local controllability of nonlinear systems along closed orbits. IEEE Transactions on Automatic Control **37**(3), 378–380 (1992)
24. Nilsson, P., Ozay, N., Liu, J.: Augmented finite transition systems as abstractions for control synthesis. Discrete Event Dynamic Systems **27**(2), 301–340 (2017)
25. Ozay, N., Liu, J., Prabhakar, P., Murray, R.M.: Computing augmented finite transition systems to synthesize switching protocols for polynomial switched systems. In: Proc. of ACC. pp. 6237–6244 (2013)
26. Pnueli, A.: The temporal logic of programs. In: Proc. of FOCS. pp. 46–57. IEEE (1977)
27. Pola, G., Girard, A., Tabuada, P.: Approximately bisimilar symbolic models for nonlinear control systems. Automatica **44**(10), 2508–2516 (2008)
28. Reissig, G., Weber, A., Rungger, M.: Feedback refinement relations for the synthesis of symbolic controllers. IEEE Transactions on Automatic Control **62**(4), 1781–1796 (2017)
29. Royden, H., Fitzpatrick, P.: Real Analysis. Printice-Hall (2010)
30. Tabuada, P.: Verification and Control of Hybrid Systems: A Symbolic Approach. Springer (2009)
31. Tabuada, P., Pappas, G.J.: Linear time logic control of discrete-time linear systems. IEEE Transactions on Automatic Control **51**(12), 1862–1877 (2006)
32. Zamani, M., Pola, G., Mazo, M., Tabuada, P.: Symbolic models for nonlinear control systems without stability assumptions. IEEE Transactions on Automatic Control **57**(7), 1804–1809 (2012)

A Linear Temporal Logic

We provide the syntax and semantics of linear temporal logic with the next operator ($\text{LTL}_{\setminus\bigcirc}$). The readers can refer to [3, 5] for standard definitions.

Syntax We can define the syntax of $\text{LTL}_{\setminus\bigcirc}$ over a set of atomic propositions Π inductively as follows:

- **true** and **false** are $\text{LTL}_{\setminus\bigcirc}$ formulas;
- an atomic proposition $\pi \in \Pi$ is an $\text{LTL}_{\setminus\bigcirc}$ formula;
- if φ and ψ are $\text{LTL}_{\setminus\bigcirc}$ formulas, then $\neg\varphi$, $\varphi \vee \varphi$, and $\varphi \mathcal{U} \varphi$ are $\text{LTL}_{\setminus\bigcirc}$ formulas.

Negation Normal Form (NNF): All $\text{LTL}_{\setminus\bigcirc}$ formulas can be transformed into negation normal form [5, p. 132], where

- all negations appear only in front of the atomic propositions[1];
- only the logical operators **true**, **false**, $\wedge$, and $\vee$ can appear; and
- only the temporal operators $\mathcal{U}$ and $\mathcal{R}$ can appear, where $\mathcal{R}$ is defined by $\varphi_1 \mathcal{R} \varphi_2 \equiv \neg(\neg\varphi_1 \mathcal{U} \neg\varphi_2)$, called the *dual until* operator.

For syntactic convenience, we can define additional temporal operators $\Box$ and $\Diamond$ by $\Box\varphi \equiv \texttt{false}\mathcal{R}\varphi$ and $\Diamond\varphi \equiv \texttt{true}\mathcal{U}\varphi$.

Semantics We consider two types of semantics for $\text{LTL}_{\setminus\bigcirc}$ formulas, namely, continuous-time and discrete-time semantics. To define semantics, an atomic proposition is interpreted as a subset of the state space on which the atomic proposition holds true. This is achieved by defining a *labelling function* $\mathcal{L} : \mathbb{R}^n \to 2^{\Pi}$ that maps a state to a set of propositions that hold true for this state.

Continuous-time semantics of $\text{LTL}_{\setminus\bigcirc}$: Given a continuous-time function $\xi : [0, \infty) \to \mathbb{R}^n$, we define $\xi, t \vDash (\varphi, \mathcal{L})$ with respect to an $\text{LTL}_{\setminus\bigcirc}$ formula φ and a labelling function $\mathcal{L}$ at time t inductively as follows:

- $\xi, t \vDash (\pi, \mathcal{L})$ if and only if $\pi \in \mathcal{L}(\xi(t))$;
- $\xi, t \vDash (\varphi_1 \vee \varphi_2, \mathcal{L})$ if and only if $\xi, t \vDash (\varphi_1, \mathcal{L})$ or $\xi, t \vDash (\varphi_2, \mathcal{L})$;
- $\xi, t \vDash (\varphi_1 \wedge \varphi_2, \mathcal{L})$ if and only if $\xi, t \vDash (\varphi_1, \mathcal{L})$ and $\xi, t \vDash (\varphi_2, \mathcal{L})$;
- $\xi, t \vDash (\varphi_1 \mathcal{U} \varphi_2, \mathcal{L})$ if and only if there exists $t' \geq 0$ such that $\xi, t + t' \vDash (\varphi_2, \mathcal{L})$ and for all $t'' \in [0, t')$, $\xi, t + t'' \vDash (\varphi_1, \mathcal{L})$;
- $\xi, t \vDash (\varphi_1 \mathcal{R} \varphi_2, \mathcal{L})$ if and only if, for all $t' \geq 0$, at least one of the following holds: $\xi, t + t' \vDash (\varphi_2, \mathcal{L})$ or there exists $t'' \in [0, t')$ such that $\xi, t + t'' \vDash (\varphi_1, \mathcal{L})$.

We write $\xi \vDash (\varphi, \mathcal{L})$ if $\xi, 0 \vDash (\varphi, \mathcal{L})$. If the labelling function is clear from the context, we simply write $\xi \vDash \varphi$.

Discrete-time semantics of $\text{LTL}_{\setminus\bigcirc}$: Given a sequence $\rho = \{x_i\}_{i=0}^{\infty}$ in $\mathbb{R}^n$, we define $\rho, i \vDash \varphi$ with respect to an $\text{LTL}_{\setminus\bigcirc}$ formula φ and a labelling function $\mathcal{L}$ inductively as follows:

[1] We assume that all negations can be effectively removed by introducing new atomic propositions corresponding to the negations of current ones.

- $\rho, i \vDash (\pi, \mathcal{L})$ if and only if $\pi \in \mathcal{L}(x_i)$;
- $\rho, i \vDash (\varphi_1 \vee \varphi_2, \mathcal{L})$ if and only if $\rho, i \vDash (\varphi_1, \mathcal{L})$ or $\rho, i \vDash (\varphi_2, \mathcal{L})$;
- $\rho, i \vDash (\varphi_1 \wedge \varphi_2, \mathcal{L})$ if and only if $\rho, i \vDash (\varphi_1, \mathcal{L})$ and $\rho, i \vDash (\varphi_2, \mathcal{L})$;
- $\rho, i \vDash (\varphi_1 \mathcal{U} \varphi_2, \mathcal{L})$ if and only if there exists $j \geq i$ such that $\rho, j \vDash (\varphi_2, \mathcal{L})$ and $\rho, k \vDash (\varphi_1, \mathcal{L})$ for all $k \in [i, j)$;
- $\rho, i \vDash (\varphi_1 \mathcal{R} \varphi_2, \mathcal{L})$ if and only if, for all $j \geq i$, at least one of the following holds: $\rho, j \vDash (\varphi_2, \mathcal{L})$ or there exists $k \in [i, j)$ such that $\rho, k \vDash (\varphi_1, \mathcal{L})$.

Similarly, we write $\rho \vDash (\varphi, \mathcal{L})$ if $\rho, 0 \vDash (\varphi, \mathcal{L})$. If the labelling function is clear from the context, we simply write $\rho \vDash \varphi$.

Remark 6. The use of linear temporal logic as specifications is a matter of choice. It should be straightforward to extend the results in this paper to ω-regular properties (by showing the soundness result Proposition 4 holds for properties expressed as a deterministic ω-automaton such as Rabin [8]).

B Illustration of Strengthening Labelling Functions

In the following, we illustrate the strengthening of labelling functions.

Remark 7. We use a simple example as shown in Figure 1 to illustrate the strengthening of labelling functions. Consider an atomic proposition p and the specification $\varphi = \diamond p$ (i.e., eventually reach p). Figure 1(a) shows that, while a continuous-time trajectory ξ satisfies φ, the sequence of sampled states from $\mathbf{x}$ (shown in red dots and denoted by ρ) does not satisfy φ. Similarly, consider the specification $\psi = \square \neg p$ (i.e., always not p). Then the sampled state sequence ρ satisfies ψ, whereas the continuous-time trajectory ξ violates it. To cope with this inherent mismatch caused by discretization, we use robust semantics of temporal logic (see [7]). In this paper, we introduce robust semantics by strengthening the labelling function. Figure 1(b) shows the effect of an ε-strengthening for labelling p. With a properly strengthened labelling function $\mathcal{L}_\varepsilon$ (cf. Proposition 3), we can show that $\xi \vDash (\varphi, \mathcal{L}_\varepsilon)$ implies $\rho \vDash (\varphi, \mathcal{L})$. Similarly, we can define robust semantics for $\psi = \square \neg p$ by a strengthened labelling function. To do so, however, requires that we introduce a new atomic proposition $q = \neg p$ to replace the negation (as we noted in the footnote of the previous page). Figure 1(c) illustrates this case. We can see that $\rho \vDash (\psi, \mathcal{L}_\varepsilon)$ implies $\xi \vDash (\psi, \mathcal{L})$, where now $\psi = \square q$.

Remark 8. We emphasize that the robust semantics introduced by a strengthened labelling function are by definition conservative; they are not meant to be equivalent to the original semantics. Consider a simple example with atomic propositions p and q. Suppose a labelling function $\mathcal{L}$ labels $p = [0, 1]$ and $q = [-1, 0]$ and the specification $\varphi = q \mathcal{U} p$. It is clear that, for any ε-strengthening of the labelling function with $\varepsilon > 0$, $(\varphi, \mathcal{L}_\varepsilon)$ (trivially) cannot be satisfied by any continuous trajectory, while $(\varphi, \mathcal{L})$ is satisfiable. The results of this paper are built upon such conservative robust semantics and aim to decide either a temporal logic formula is not satisfiable in strengthened semantics (by a perturbed

system) or it can be satisfiable in the original semantics (by the nominal system). A formal problem statement is given in the next subsection.

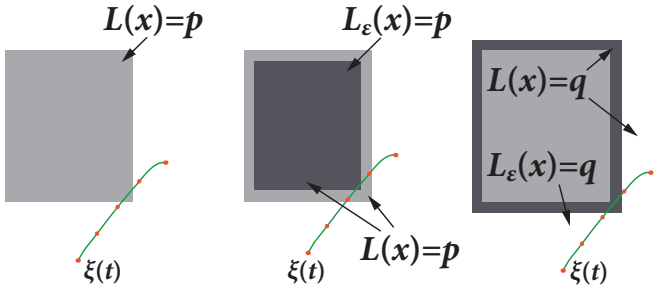

Fig. 1: An illustration of the strengthening of labelling functions (see Remark 7).

C A Decision Diagram for Roust Realizability

Figure 2 provides a decision diagram for roust realizability, which can be seen as an overview of the main results developed in the paper.

D Proof of Proposition 1

Proof (Proposition 1). Pick $\pi \in \mathcal{L}_{\varepsilon_1+\varepsilon_2}(x)$, then $\pi \in \mathcal{L}(y)$ for all $y \in x + (\varepsilon_1 + \varepsilon_2)\mathbb{B}$. To prove $\pi \in (\mathcal{L}_{\varepsilon_1})_{\varepsilon_2}(x)$, we have to show that $\pi \in \mathcal{L}_{\varepsilon_1}(z)$ for all $z \in x+\varepsilon_2\mathbb{B}$. Fix any such z, we verify $\pi \in \mathcal{L}_{\varepsilon_1}(z)$ by showing that $\pi \in \mathcal{L}(w)$ for all $w \in z+\varepsilon_1\mathbb{B}$. This is true by the triangle inequality $|w - x| \leq |w - z| + |z - x| \leq \varepsilon_1 + \varepsilon_2$.

E Proof of Proposition 2

Proof (Proposition 2). The proof of the above proposition is straightforward. A dwell-time N control strategy with sampling period τ corresponds exactly to a sampled-data control strategy with sampling period $T = N\tau$. Note that a control strategy can have memory and can easily encode consecutive use of the same control input for a finite number of times.

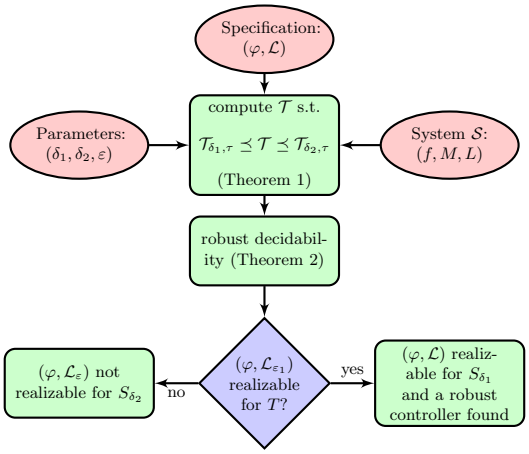

Fig. 2: A decision diagram for checking robust realizability: given a system S, a temporal logic specification φ with a labelling function $\mathcal{L}$, and parameters $\delta_2 > \delta_1 \geq 0$ and $\varepsilon > 0$, we can decide either $(\varphi, \mathcal{L})$ is realizable for $\mathcal{S}_{\delta_1}$ with a robust controller, or $(\varphi, \mathcal{L}_\varepsilon)$ is not realizable for $\mathcal{S}_{\delta_2}$. This is done by checking realizability of $(\varphi, \mathcal{L}_{\varepsilon_1})$ on $\mathcal{T}$, where $\mathcal{T}$ is a sufficiently precise abstraction of $\mathcal{S}_{\delta_1}$ constructed as in the proof of Theorem 1 by choosing the discretization parameters η, μ, τ sufficiently small according to (13), and ε_1 is chosen according to (14) in the proof of Theorem 2.

F Proof of Proposition 3

Proof (Proposition 3). The implementation of control strategy and preservation of dwell-time are straightforward. The main part is to show correctness of temporal logic formula. Suppose that a trajectory for $\mathcal{S}_\delta$ resulting from a control strategy is $\mathbf{x}(t)$. The corresponding path of an execution of $\mathcal{T}_{\delta,\tau}$ is given by $\rho = \mathbf{x}(0), \mathbf{x}(\tau), \mathbf{x}(2\tau), \cdots$.

We need to show that (1) $\rho \vDash (\varphi, \mathcal{L}_\varepsilon)$ implies $x \vDash (\varphi, \mathcal{L})$, and (2) $x \vDash (\varphi, \mathcal{L}_\varepsilon)$ implies $\rho \vDash (\varphi, \mathcal{L})$. The following proof, modelled after that for Theorem 4.1 in [22], is an inductive argument based on the structure of $\mathrm{LTL}_{\setminus\bigcirc}$ formulas. In fact, the proof for (1) is very similar to of for Theorem 4.1 in [22]. In the following, we prove case (2), that is, $x \vDash (\varphi, \mathcal{L}_\varepsilon)$ implies $\rho \vDash (\varphi, \mathcal{L})$. We do so by proving a stronger statement: for every $i \geq 0$, $x, t \vDash (\varphi, \mathcal{L}_\varepsilon)$ for some $t \in J_i = [i\tau - \frac{\tau}{2}, i\tau + \frac{\tau}{2}]$ implies $\rho, i \vDash (\varphi, \mathcal{L}_\varepsilon)$.

Case $\varphi = \pi$: Suppose that $x, t \vDash (\pi, \mathcal{L}_\varepsilon)$ for some $t \in J_i$, we have to show that $\pi \in \mathcal{L}(x(i\tau))$. This follows from $\pi \in \mathcal{L}_\varepsilon(x(t))$, $\varepsilon \geq (M + \delta)\tau/2$ and

$$|x(t) - x(i\tau)| \leq |x(t) - x(\tau_i)| \leq (M + \delta)\tau/2. \tag{19}$$

Case $\varphi = \varphi_1 \mathcal{R} \varphi_2$: Suppose that $x(t) \vDash (\varphi, \mathcal{L}_\varepsilon)$ for some $t \in J_i$. We need to show that $\rho, i \vDash (\varphi, \mathcal{L})$, that is, for all $j \geq i$, either $\rho, j \vDash (\varphi_2, \mathcal{L})$ holds or there exists some $k \in [i, j)$ such that $\rho, k \vDash (\varphi_1, \mathcal{L})$ holds. Since $x(t) \vDash (\varphi, \mathcal{L}_\varepsilon)$ for some $t \in J_i$, we know that for every $t' \geq t$, either $x(t') \vDash (\varphi, \mathcal{L}_\varepsilon)$ holds or there exists $s \in [t, t')$ such that $x(s) \vDash (\varphi, \mathcal{L}_\varepsilon)$ holds. Let $t' = j\tau - \frac{\tau}{2}$. If the former holds, we have $x(t') \vDash (\varphi, \mathcal{L}_\varepsilon)$ for $t' = j\tau - \frac{\tau}{2} \in J_j$. By the inductive assumption, this implies $\rho, j \vDash (\varphi_2, \mathcal{L})$. If the latter holds, there exists some interval J_k such that $s \in J_k$, $k \in [i, j)$, and $x(s) \vDash (\varphi, \mathcal{L}_\varepsilon)$. It follows by the inductive assumption that $\rho, k \vDash (\varphi_2, \mathcal{L})$.

The other cases are straightforward.

G Proof of Proposition 4

Proof (Proposition 4). The proof is similar to the proof of Theorem 1 in [18] and it is straightforward to handle the separate cases of proposition preserving and finite-granularity abstractions. Additional consideration has to be given to the dwell-time requirement. We provide a proof for soundness of dwell-time strategies as follows. Suppose that $\mathcal{T}_2$ has a dwell time N strategy σ_2. An implementation of σ_2 can be constructed as follows. At each time instant, we observe a concrete state $q_1 \in Q_1$ for $\mathcal{T}_1$ and map it to an abstract state $q_2 \in Q_2$ (after applying α). Suppose that, at state $q_2 \in Q_2$, σ_2 chooses to apply the same action a_2 for N consecutive times. We can let $\mathcal{T}_1$ use a dwell time N strategy σ_1 to follow the action a_1, corresponding to a_2 by condition (ii) of Definition 3, N times. By condition (ii) of Definition 3, it follows by induction that the reachable set of $\mathcal{T}_1$ from q_1 under action a_1 in N steps is contained (after applying α) in the reachable set of q_2 under action a_2 in N steps. Hence, after N steps, we can repeat the same procedure. This results in an implementation of a dwell-time N control strategy σ_1 for $\mathcal{T}_1$ from a dwell-time N control strategy σ_2 for $\mathcal{T}_2$.

H Robust Simulations of Turing Machines by a One-Dimensional Control System

We outline a procedure to robustly simulate any Turing machine by a one-dimensional control system. This gives a negative result when the lower bound on sampling period in Theorem 4 is removed.

Let TM be a Turing machine. One can encode its configuration in a rational number of the form $x_{q,h,r,s} = \frac{1}{2^p 3^h 5^r 7^s}$, where p, q, r, s are nonnegative integers [11, Proof of Theorem 5.6]. These rational numbers form a sequence in $(0, 1]$, with an accumulation point at 0. Consider a control system of the form $x' = u$, where $u \in [-1, 1]$ (for simplicity). We design a sampled-data control strategy to simulate TM (in multiple steps). Let $\delta \in (0, 1)$. Choose u_1 such that $u_1 + \delta < 0$ and u_2 such that $u_2 - \delta > 0$. For each configuration of TM encoded by y_i, there exists a control strategy for $\mathcal{S}_\delta$ to get from y_i to the encoding of the next configuration y_{i+1} as follows. If $y_{i+1} > y_i$, choose u_2 and a sampling time τ such

that $(u_2+\delta)\tau < \gamma$, where $\gamma > 0$ is such that $(y_{i+1}-\gamma, y_{i+1}-\gamma)$ does not include any encoding of the configurations of TM. The sampled-data control strategy is to apply u_2 until it reaches $(y_{i+1}-\gamma/2, y_{i+1}-\gamma/2)$. The case for $y_{i+1} < y_i$ is similar.