

Agile Temporal Discretization for Symbolic Optimal Control

Adrien Janssens

Adrien Banse

Julien Calbert

Raphaël M. Jungers

Abstract—As control systems grow in complexity, abstraction-based methods have become essential for designing controllers with formal guarantees. However, a key limitation of these methods is their reliance on discrete-time models, typically obtained by discretizing continuous-time systems with a fixed timestep. This discretization leads to two major problems: when the timestep is small, the abstraction includes numerous stuttering and spurious trajectories, making controller synthesis suboptimal or even infeasible; conversely, a large time step may also render control design infeasible due to a lack of flexibility. In this work, drawing inspiration from Reinforcement Learning concepts, we introduce temporal abstractions, which allow for a flexible timestep. We provide a method for constructing such abstractions and formally establish their correctness in controller design. Furthermore we show how to apply these to optimal control under reachability specifications. Finally we showcase our methods on two numerical examples, highlighting that our approach leads to controllers that achieve a lower worst-case control cost.

I. INTRODUCTION

The growing complexity of engineered dynamical systems has made controller design increasingly arduous. In the literature, this paradigm shift has been coined by the academic and industrial communities as the *cyber-physical revolution* [12], [3], [14]. On the other hand, it is crucial to provide safety and/or performance guarantees for such systems [13]. In particular, we focus on *reachability analysis* [2]: designing a controller that drives the system from an initial set to a target set while avoiding obstacles. Being able to solve such problems for complex dynamical systems is crucial in many applications (see [1], [8] for applications to autonomous driving and robotic aerial vehicles).

In this context, *abstraction-based control*, also known as symbolic control, is a promising tool to provide formal reachability guarantees [19]. Every abstraction-based control method follows the same systematic procedure:

- 1) *Abstraction* – The concrete system is approximated by a finite-state mathematical object (the *abstract* system);
- 2) *Abstract controller synthesis* – A discrete controller is designed for the abstract system;

R. M. Jungers is a FNRS honorary Research Associate. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme under grant agreement No 864017 - L2C, from the Horizon Europe programme under grant agreement No101177842 - Unimaas, and from the ARC (French Community of Belgium)- project name: SIDDARTA. Adrien Banse is supported by the French Community of Belgium in the framework of a FNRS/FRIA grant.

A. Janssens, A. Banse, J. Calbert and R. Jungers are with the ICTEAM, UCLouvain (Louvain-la-Neuve, Belgium). E-mail addresses: adrien.janssens@student.uclouvain.be, {adrien.banse, julien.calbert, raphael.jungers}@uclouvain.be

- 3) *Concretization* – From the discrete controller, a concrete controller is refined for the original (*concrete*) dynamical system.

The validity of this approach and the procedure to concretize the discrete controller are provided by a mathematical *simulation relation* [19], [6] between the abstract and concrete systems, such as the *feedback refinement relation* (FRR) [17].

Continuous-time deterministic dynamical systems are abstracted by partitioning the state space into a finite set of *cells*, selecting a finite set of inputs, considering discrete-time dynamics with timestep τ . The abstract states correspond to the cells, and a transition between s_1 and s_2 with input u is added if the set of reachable states from the cell s_1 intersects with the cell s_2 . In this work, we focus on the time discretization. Indeed, in the context of reachability analysis, fixing a too small constant timestep τ leads to conservatism for at least two reasons:

- 1) *Stuttering transitions* – If the timestep is small, then the reachable set from an initial cell may intersect with the cell itself, leading to a self-loop, or *stuttering transition* in the automaton. When performing reachability analysis on the abstract system, a possible behaviour is that the state stays infinitely in the initial cell, whereas this may be impossible in reality. This kind of transition is illustrated in yellow on Figure 1.
- 2) *Spurious trajectories* – At each step of a trajectory in the abstract system, the information of where *exactly* the state is within the cell is lost. As a consequence, an error is propagated with the number of steps, and more and more trajectories are valid in the abstraction but cannot happen in reality: such trajectories are called *spurious trajectories*. This phenomenon is also illustrated in red (two small timesteps) and in blue (one long timestep) in Figure 1.

This conservatism leads either to infeasible problems, or to sub-optimal solutions in the context of optimal control. On the other hand, considering small timesteps allows for more flexibility, and is sometimes needed in order to ensure feasibility of the problem.

Contributions: In this work, we propose to take into account this trade-off by introducing *temporal abstractions*, in which multiple timesteps are allowed. We prove the correctness of such abstractions, thereby providing formal guarantees. We then provide the necessary theoretical framework to use these abstractions in the context of optimal control under reachability specifications. While our approach does not directly tackle the curse of dimensionality inherent

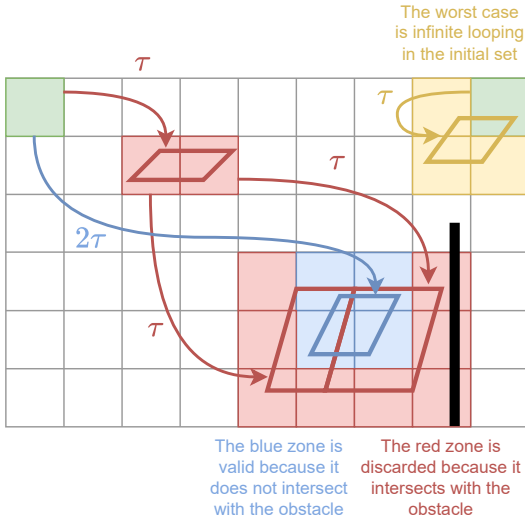


Fig. 1: Illustration of stuttering transitions (yellow) and spurious trajectories (red and blue). **Yellow:** the timestep is too small, and therefore the reachable set intersects with the starting cell. **Red:** from the initial green cell, two cells are reachable with a jump of length τ . At the next step, the two reachable sets from both cells are computed. Applying the same input for 2τ is discarded because 3 cells intersect with the black obstacle. **Blue:** we now directly consider a timestep corresponding to a larger timestep, 2τ . The reachable set intersects with 4 cells that do not intersect with the obstacle, making the input valid.

to abstraction-based control, numerical experiments show that using multiple timesteps yields lower worst-case cost controllers than those obtained with a finer input space discretization.

Related works: Our work takes inspiration from the notion of *temporal abstraction* in the Reinforcement Learning community (see e.g. [18] for a survey), but aims at combining this powerful framework and the formal guarantees that abstraction-based methods can provide. The works [10], [11] have already investigated flexible timesteps in the context of abstraction-based techniques but for safety analysis, and with a different abstraction construction.

Outline: The rest of this paper is organized as follows. In Section II, we introduce all the necessary formalism to ensure correctness of the approach. In Section III, we introduce the concept of *temporal abstractions* and how to apply them to an optimal control problem. Finally, in Section IV we demonstrate the quality of our method on two examples.

Notations: Given two sets $\mathcal{X}_1, \mathcal{X}_2$, we define a *single-valued map* as $f : \mathcal{X}_1 \rightarrow \mathcal{X}_2$, while a *set-valued map* is defined as $F : \mathcal{X}_1 \rightarrow 2^{\mathcal{X}_2}$, where $2^{\mathcal{X}_2}$ is the power set of \mathcal{X}_2 , i.e., the set of all subsets of \mathcal{X}_2 . The image of a subset $\Omega \subseteq \mathcal{X}_1$ under $F : \mathcal{X}_1 \rightarrow 2^{\mathcal{X}_2}$ is denoted $F(\Omega)$. Given two sets \mathcal{X}_1 and \mathcal{X}_2 , we identify a binary *relation* $R \subseteq \mathcal{X}_1 \times \mathcal{X}_2$ with set-valued maps, i.e., $R(x_1) = \{x_2 \mid (x_1, x_2) \in R\}$ and

$R^{-1}(x_2) = \{x_1 \mid (x_1, x_2) \in R\}$. A relation $R \subseteq \mathcal{X}_1 \times \mathcal{X}_2$ is said to be *strict* if $R(x_1) \neq \emptyset$ for all $x_1 \in \mathcal{X}_1$. When \mathcal{X}_2 is finite, the mapping $R : \mathcal{X}_1 \rightarrow 2^{\mathcal{X}_2}$ is referred to as a *quantizer*. The set of all signals that take their values in \mathcal{B} are defined on intervals of the form $[0; N] = [0; N - 1]$ is denoted \mathcal{B}^∞ , $\mathcal{B}^\infty = \bigcup_{N \in \mathbb{N} \cup \{\infty\}} \mathcal{B}^{[0; N]}$. Given a set-valued map $f : \mathcal{X} \rightarrow 2^{\mathcal{Y}}$ and $\mathbf{x} \in \mathcal{X}^{[0; N]}$, we denote by $f(\mathbf{x}) = \{\mathbf{y} \in \mathcal{Y}^{[0; N]} \mid \forall k \in [0; N) : y(k) \in f(x(k))\}$.

II. PRELIMINARIES

In this section, we introduce the control formalism and the symbolic control framework required to ensure the correctness of the proposed approach.

A. Control framework

Definition 1: A *transition system* is a triple $\mathcal{S} = (\mathcal{X}, \mathcal{U}, F)$, where \mathcal{X} and \mathcal{U} are respectively the sets of *states* and *inputs*, and $F : \mathcal{X} \times \mathcal{U} \rightarrow 2^{\mathcal{X}}$ is the *transition map* such that

$$x(k+1) \in F(x(k), u(k)). \quad (1)$$

A tuple $(\mathbf{x}, \mathbf{u}) \in \mathcal{X}^{[0; N]} \times \mathcal{U}^{[0; N]}$ is a *trajectory* of length N of \mathcal{S} starting at $x(0)$ if $N \in \mathbb{N} \cup \{\infty\}$ and (1) holds for all $k \in [0; N - 1]$. \triangle

Given a transition system $\mathcal{S} = (\mathcal{X}, \mathcal{U}, F)$, we introduce the set-valued operator of *available inputs*, defined as $\mathcal{U}_{\mathcal{S}}(x) = \{u \in \mathcal{U} \mid F(x, u) \neq \emptyset\}$, which gives the set of inputs u available at a given state x . For a given $u \in \mathcal{U}$, the *reachable set* of $\Omega \subseteq \mathcal{X}$ is defined as $F(\Omega, u) := \bigcup_{x \in \Omega} F(x, u)$.

We use a set-valued function to model perturbations and non-determinism in a unified formalism. We say that the transition system is deterministic if for every state $x \in \mathcal{X}$ and input $u \in \mathcal{U}$, $F(x, u)$ is either empty or a singleton

A *finite* transition system refers to a system characterized by finitely many states and inputs. We define the *behaviour* of a system as its set of *maximal trajectories*.

Definition 2: Given the system \mathcal{S} in Definition 1, the set

$$\mathcal{B}(\mathcal{S}) = \{(\mathbf{x}, \mathbf{u}) \mid (\mathbf{x}, \mathbf{u}) \text{ is a trajectory of } \mathcal{S} \text{ on } [0; N], \text{ and if } N < \infty \text{ then } F(x(N-1), u(N-1)) = \emptyset\}, \quad (2)$$

is called the *behaviour* of \mathcal{S} . \triangle

For clarity, we focus on *static* controllers, though most results extend naturally to dynamical controllers.

Definition 3: A *static controller* for a transition system $\mathcal{S} = (\mathcal{X}, \mathcal{U}, F)$ is a set-valued map $\mathcal{C} : \mathcal{X} \rightarrow 2^{\mathcal{U}}$ satisfying $\forall x \in \mathcal{X}, \mathcal{C}(x) \subseteq \mathcal{U}(x)$. The *domain* of \mathcal{C} is defined as $\text{dom}(\mathcal{C}) = \{x \in \mathcal{X} \mid \mathcal{C}(x) \neq \emptyset\}$. The *controlled system* is the transition system $(\mathcal{X}, \mathcal{U}, F_{\mathcal{C}})$ where

$$x^+ \in F_{\mathcal{C}}(x, u) \Leftrightarrow u \in \mathcal{C}(x) \wedge x^+ \in F(x, u),$$

and is denoted $\mathcal{C} \times \mathcal{S}$. \triangle

This controller is *static* because the allowed inputs at each state depend only on the current state, not on the past trajectory. We now define the control problem.

Definition 4: Given a transition system $\mathcal{S} = (\mathcal{X}, \mathcal{U}, F)$, a *specification* Σ is any subset $\Sigma \subseteq (\mathcal{X} \times \mathcal{U})^\infty$. The system \mathcal{S}

satisfies Σ if $\mathcal{B}(\mathcal{S}) \subseteq \Sigma$. A *control problem* is a pair (\mathcal{S}, Σ) , and a controller \mathcal{C} solves it if $\mathcal{B}(\mathcal{C} \times \mathcal{S}) \subseteq \Sigma$. \triangle

B. Abstraction-based control

Given the concrete system $\mathcal{S}_1 = (\mathcal{X}_1, \mathcal{U}_1, F_1)$ and the abstract system $\mathcal{S}_2 = (\mathcal{X}_2, \mathcal{U}_2, F_2)$, we aim to construct a concrete controller $\mathcal{C}_1 : \mathcal{X}_1 \rightarrow 2^{\mathcal{U}_1}$ from an abstract controller $\mathcal{C}_2 : \mathcal{X}_2 \rightarrow 2^{\mathcal{U}_2}$, as part of the third step of the abstraction approach described in the introduction. To do so, the relation $R \subseteq \mathcal{X}_1 \times \mathcal{X}_2$ must impose conditions on local dynamics, ensuring that input choices lead to consistent state transitions across associated states. One such relation is the *feedback refinement relation*.

Definition 5 ([17, Definition V.2]): Given two transition systems $\mathcal{S}_1 = (\mathcal{X}_1, \mathcal{U}_1, F_1)$ and $\mathcal{S}_2 = (\mathcal{X}_2, \mathcal{U}_2, F_2)$, a relation $R \subseteq \mathcal{X}_1 \times \mathcal{X}_2$ is a *feedback refinement relation* from \mathcal{S}_1 to \mathcal{S}_2 , denoted $\mathcal{S}_1 \preceq_R^{\text{FRR}} \mathcal{S}_2$, if $\forall (x_1, x_2) \in R$, the conditions

$$\mathcal{U}_{\mathcal{S}_2}(x_2) \subseteq \mathcal{U}_{\mathcal{S}_1}(x_1) \quad (3)$$

$$\forall u \in \mathcal{U}_{\mathcal{S}_2}(x_2) : R(F_1(x_1, u)) \subseteq F_2(x_2, u) \quad (4)$$

hold. \triangle

The feedback refinement relation benefits from a simple concretization scheme as established by the following proposition.

Proposition 1 ([17, Theorem V.4]): Let $\mathcal{S}_1 = (\mathcal{X}_1, \mathcal{U}_1, F_1)$ and $\mathcal{S}_2 = (\mathcal{X}_2, \mathcal{U}_2, F_2)$ be transition systems, and let $R \subseteq \mathcal{X}_1 \times \mathcal{X}_2$. If $\mathcal{S}_1 \preceq_R^{\text{FRR}} \mathcal{S}_2$, then for any controller \mathcal{C}_2 , the controller $\mathcal{C}_1 = \mathcal{C}_2 \circ R$ ensures $\mathcal{B}(\mathcal{C}_1 \times \mathcal{S}_1) \subseteq R^{-1}(\mathcal{B}(\mathcal{C}_2 \times \mathcal{S}_2))$. Consequently, if $\Sigma_1 \subseteq R^{-1}(\Sigma_2)$ and \mathcal{C}_2 solves $(\mathcal{S}_2, \Sigma_2)$, then \mathcal{C}_1 solves $(\mathcal{S}_1, \Sigma_1)$.

Although \mathcal{S}_2 need not be finite, practical use of graph-theoretic tools for controller synthesis requires a finite abstraction (see Section IV).

C. Continuous-time systems

In this work, we focus on the control of continuous-time systems.

Definition 6: A *continuous-time system* is a triple $\mathcal{D} = (\mathcal{X}, \mathcal{U}, f)$, where $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{U} \subseteq \mathbb{R}^m$ are respectively the sets of *states* and *inputs*, and $f : \mathbb{R}^n \times \mathcal{U} \rightarrow \mathbb{R}^n$ is the vector field such that

$$\dot{x}(t) = f(x(t), u(t)). \quad (5)$$

We assume that $f(\cdot, u)$ is locally Lipschitz for all $u \in \mathcal{U}$. \triangle

In the sequel, ϕ denotes the general solution of the system associated with (5) for constant inputs. That is, if $x_0 \in \mathbb{R}^n$, $u \in \mathcal{U}$, and $f(\cdot, u)$ is locally Lipschitz, then $\phi(\cdot, x_0, u)$ is the unique solution of the initial value problem $\dot{x} = f(x, u)$, $x(0) = x_0$ [9].

III. MAIN RESULTS

In this section, we introduce *temporal abstractions*. We start by providing formal definitions of both the concrete systems \mathcal{S}_1 and the abstract systems \mathcal{S}_2 . We then provide conditions on the abstract system \mathcal{S}_2 for $\mathcal{S}_1 \preceq_R^{\text{FRR}} \mathcal{S}_2$ to hold (see Theorem 1). We then show how to use the latter for

optimal control of a system under reach-avoid specifications. We show that the developed abstraction-based techniques lead to optimality guarantees (see Theorem 2).

A. Temporal abstractions

We begin by introducing a general framework to frame a continuous-time system as a transition system without loss of generality, by extending the input space:

Definition 7: Consider a continuous-time system $\mathcal{D} = (\mathcal{X}, \mathcal{U}, f)$. This system can be framed as a transition system $\mathcal{S}_1 = (\mathcal{X}_1, \mathcal{U}_1, F_1)$ as follows:

- The state space is unchanged $\mathcal{X}_1 = \mathcal{X}$;
- The input space \mathcal{U}_1 is defined as $\mathcal{U}_1 := \mathcal{U} \times [0; \infty)$;
- The transition function $F_1 : \mathcal{X}_1 \times \mathcal{U}_1 \rightarrow 2^{\mathcal{X}_1}$ is given by

$$F_1(x, (u, t)) := \{\phi(t, x, u)\}. \quad (6)$$

We call this system the *multi-sampled system*. \triangle

Now, in this work, we apply abstraction-based control to the continuous-time system $\mathcal{D} = (\mathcal{X}, \mathcal{U}, f)$ (see Definition 6). Consequently, this approach involves constructing both a transition system $\mathcal{S}_2 = (\mathcal{X}_2, \mathcal{U}_2, F_2)$ and a relation $R \subseteq \mathcal{X}_1 \times \mathcal{X}_2$ such that $\mathcal{S}_1 \preceq_R^{\text{FRR}} \mathcal{S}_2$, where $\mathcal{S}_1 = (\mathcal{X}_1, \mathcal{U}_1, F_1)$ is the transition system associated with \mathcal{D} as in Definition 7. As proven below, the following conditions on the abstract system \mathcal{S}_2 and R are sufficient for $\mathcal{S}_1 \preceq_R^{\text{FRR}} \mathcal{S}_2$ to hold:

- The input space \mathcal{U}_2 is given by

$$\mathcal{U}_2 = \mathcal{U}'_2 \times \mathcal{T} \subseteq \mathcal{U}_1 = \mathcal{U} \times [0, \infty), \quad (7)$$

where $\mathcal{U}'_2 \subseteq \mathcal{U}$ represents a subset of the concrete system's inputs, and $\mathcal{T} \subseteq [0, \infty)$ is a subset of times;

- For all $u_2 = (u, \tau) \in \mathcal{U}_2$, the transition map satisfies

$$\begin{aligned} F_2(x_2, u_2) \\ \supseteq \{x_2^+ \mid F_1(R^{-1}(x_2), (u, \tau)) \cap R^{-1}(x_2^+) \neq \emptyset\}. \end{aligned} \quad (8)$$

The finiteness of \mathcal{U}_2 implies that only a finite number of inputs and sampling times are selected from the original system. As a result, the abstract system \mathcal{S}_2 can be interpreted as a *sampled* version of \mathcal{S}_1 . When \mathcal{T} is a singleton, i.e., $\mathcal{T} = \{\tau\}$ for some $\tau \in [0, \infty)$, it corresponds to the classical discretization setting with a fixed timestep τ .

Remark 1: In this section, we remain general. However, in Section IV, we will provide an explicit procedure for constructing such abstract system \mathcal{S}_2 and relation R . \triangle

Now, we prove the correctness of our method by showing that, for any $\mathcal{S}_1, \mathcal{S}_2$ and R such as defined above, $\mathcal{S}_1 \preceq_R^{\text{FRR}} \mathcal{S}_2$. For the remainder of this paper, we assume that R is strict (see *Notations* paragraph in Section I for a definition).

Theorem 1: Let $\mathcal{D} = (\mathcal{X}, \mathcal{U}, f)$ be a continuous-time system with its associated transition system $\mathcal{S}_1 = (\mathcal{X}_1, \mathcal{U}_1, F_1)$ as in (6). Consider a transition system $\mathcal{S}_2 = (\mathcal{X}_2, \mathcal{U}_2, F_2)$ and a strict relation $R \subseteq \mathcal{X}_1 \times \mathcal{X}_2$, with \mathcal{U}_2 as in (7), and R and F_2 as in (8). Then it holds that $\mathcal{S}_1 \preceq_R^{\text{FRR}} \mathcal{S}_2$. \triangle

Proof: For all $(x_1, x_2) \in R$ and $(u, \tau) \in \mathcal{U}_2(x_2)$, we have $F_2(x_2, u_2) \neq \emptyset$, which by construction ensures $(x, u) \in \mathcal{U}_1(x_1)$. Given $\{x_1^+\} = F(x_1, (u, \tau))$, the strictness of R implies $R(x_1^+) \neq \emptyset$ and $R(x_1^+) \subseteq F_2(x_2, (u, \tau))$. \blacksquare

B. Application to optimal control

We begin by defining the considered reach-avoid specification (see Definition 4).

Definition 8: Given a transition system $\mathcal{S} = (\mathcal{X}, \mathcal{U}, F)$, a reach-avoid specification associated with sets $\mathcal{I}, \mathcal{F}, \mathcal{O} \subseteq \mathcal{X}$ such that $\mathcal{F} \cap \mathcal{O} = \emptyset$ is defined as

$$\Sigma^{\text{Reach}} = \{(\mathbf{x}, \mathbf{u}) \in (\mathcal{X} \times \mathcal{U})^\infty \mid x(0) \in \mathcal{I} \Rightarrow \exists N \in \mathbb{Z}_{\geq 0} : (x(N) \in \mathcal{F} \wedge \forall k \in [0; N[: x(k) \notin \mathcal{O}])\}, \quad (9)$$

which enforces that all states in the initial set \mathcal{I} will reach the target \mathcal{F} in finite time while avoiding obstacles in \mathcal{O} . We use the abbreviated notation $\Sigma^{\text{Reach}} = [\mathcal{I}, \mathcal{F}, \mathcal{O}]$ to denote the specification (9). \triangle

In addition to specifying the state behaviour of the closed-loop system, transition costs can be incorporated to address optimal control problems aiming to minimize the *worst-case* cost.

Definition 9 (Optimal control problem): Given a transition system $\mathcal{S} = (\mathcal{X}, \mathcal{U}, F)$, a reach-avoid specification $\Sigma^{\text{Reach}} = [\mathcal{I}, \mathcal{F}, \mathcal{O}]$, and a transition cost function $c : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}_{>0} \cup \{\infty\}$, the *optimal control problem* involves designing an *optimal controller* $\mathcal{C}^* : \mathcal{X} \rightarrow 2^{\mathcal{U}}$ that minimizes the *cost function*

$$\mathcal{L}(\mathcal{C}) = \begin{cases} \sup_{x_0 \in \mathcal{I}} l_{\mathcal{C}}(x_0), & \text{if } \mathcal{C} \text{ solves } (\mathcal{S}, \Sigma^{\text{Reach}}), \\ \infty, & \text{otherwise.} \end{cases} \quad (10)$$

Here, $l_{\mathcal{C}}(x) : \mathcal{I} \rightarrow \mathbb{R}_{\geq 0}$ is defined as

$$l_{\mathcal{C}}(x_0) = \sup_{(\mathbf{x}, \mathbf{u}) \in \mathcal{B}(\mathcal{C} \times \mathcal{S}) \mid x(0) = x_0} \sum_{i=0}^{N(\mathbf{x})-1} c(x(i), u(i)), \quad (11)$$

where $N(\mathbf{x}) = \min_{k \in \mathbb{Z}_{\geq 0} \mid x(k) \in \mathcal{F}} k$, i.e., $x(N)$ is the first occurrence in the sequence \mathbf{x} that belongs to \mathcal{F} . \triangle

Given a transition system $\mathcal{S} = (\mathcal{X}, \mathcal{U}, F)$, a *value function* is a function $v : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$. Given a *transition cost function* $c : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}_{>0} \cup \{\infty\}$ and a value function v , we denote by B_c the *Bellman operator* defined as

$$[B_c(v)](x) = \min_{u \in \mathcal{U}_{\mathcal{S}}(x)} \left(c(x, u) + \max_{x^+ \in F(x, u)} v(x^+) \right). \quad (12)$$

The *controller associated* with a value function v is given by $\mathcal{C}(x) = \operatorname{argmin}_{u \in \mathcal{U}_{\mathcal{S}}(x)} (c(x, u) + \max_{x^+ \in F(x, u)} v(x^+))$.

A value function v that satisfies the Bellman equation [4]

$$v(x) = [B_c(v)](x), \quad (13)$$

is commonly referred to as a *Bellman value function*. Its computation is generally intractable for infinite state systems. Instead, we consider *suboptimal value functions* [15, Definition 7].

Definition 10 (Suboptimal value function): Given a transition system $\mathcal{S} = (\mathcal{X}, \mathcal{U}, F)$ and a set $\mathcal{X}_v \subseteq \mathcal{X}$, a value function $v : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ is a *suboptimal value function* with transition cost function $c : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}_{>0} \cup \{\infty\}$, if for all $x \in \mathcal{X} \setminus \mathcal{X}_v$, $v(x) = \infty$, and for all $x \in \mathcal{X}_v$, $v(x)$ is finite and $v(x) \geq [B_c(v)](x)$. \triangle

The following proposition demonstrates that the controller

associated with a suboptimal value function provides an upper bound on the optimal worst-case cost (see Definition 9). This proposition is inspired by works such as [16]. However, the variant we propose for robust control of set-valued dynamics is, to our knowledge, new to the literature.

Proposition 2: Consider a transition system $\mathcal{S} = (\mathcal{X}, \mathcal{U}, F)$, a reach-avoid specification $\Sigma^{\text{Reach}} = [\mathcal{I}, \mathcal{F}, \mathcal{O}]$, and a transition cost function $c : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}_{>0} \cup \{\infty\}$, such that there exists $b > 0$ with $c(x, u) \geq b$ for all $x \in \mathcal{X}$ and all $u \in \mathcal{U}_{\mathcal{S}}(x)$. Let $v : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a suboptimal value function such that $v(x) = 0$ if and only if $x \in \mathcal{F}$, $v(x) = \infty$ if $x \in \mathcal{O}$, and $v(x) > 0$ otherwise.

Then, the controller \mathcal{C} associated with v solves the control problem $(\mathcal{S}, [\operatorname{dom}(\mathcal{C}), \mathcal{F}, \mathcal{O}])$. If $\mathcal{I} \subseteq \operatorname{dom}(\mathcal{C})$, then \mathcal{C} also solves $(\mathcal{S}, \Sigma^{\text{Reach}})$. Moreover, the value function $v(x)$ provides an upper bound on the worst-case cost under \mathcal{C} , i.e., $v(x) \geq l_{\mathcal{C}}(x)$, and the overall performance satisfies $\mathcal{L}(\mathcal{C}) \geq \mathcal{L}(\mathcal{C}^*)$, with equality if v is a Bellman value function. \triangle

Proof: Since v is a suboptimal value function, for all $x \in \mathcal{X}_v$, where $\mathcal{X}_v = \{x \in \mathcal{X} \mid v(x) < \infty\}$, we have $v(x) \geq [B_c(v)](x)$. For $u \in \mathcal{C}(x)$, it holds that $\max_{x^+ \in F(x, u)} v(x^+) \leq v(x) - c(x, u) \leq v(x) - b < v(x)$ since $c(x, u) \geq b$. As a result, the function v decreases by at least b along trajectories of $\mathcal{C} \times \mathcal{S}$ for all $x \in \mathcal{X}_v$ towards the minimum of v , which corresponds to the target set \mathcal{F} while avoiding the obstacle \mathcal{O} . \blacksquare

The following theorem formalizes the transfer of surrogate value functions from the abstract system to the concrete system.

Theorem 2 ([7, Theorem 1]): Consider two transition systems $\mathcal{S}_1 = (\mathcal{X}_1, \mathcal{U}_1, F_1)$ and $\mathcal{S}_2 = (\mathcal{X}_2, \mathcal{U}_2, F_2)$, and a relation $R \subseteq \mathcal{X}_1 \times \mathcal{X}_2$ such that $\mathcal{S}_1 \stackrel{\text{FRR}}{\succeq}_R \mathcal{S}_2$. Given a transition cost c_1 for \mathcal{S}_1 , we consider a transition cost c_2 for \mathcal{S}_2 such that

$$c_2(x_2, u) \geq \max_{x_1 \in R^{-1}(x_2)} c_1(x_1, u). \quad (14)$$

If v_2 is a suboptimal value function for \mathcal{S}_2 with transition cost c_2 , then

$$v_1(x_1) = \min_{x_2 \in R(x_1)} v_2(x_2) \quad (15)$$

is a suboptimal value function with transition cost c_1 for \mathcal{S}_1 .

In the special case where R is deterministic, the value function defined in (15) reduces to $v_1 = v_2 \circ R$, as previously established in [15, Theorem 8]. This is the case in the classical abstraction approach [17], [19], where the abstract state space corresponds to a partition of (a subset of) the concrete state space. In practice, while the finiteness of \mathcal{S}_2 enables efficient computation of a Bellman value function v_2 using classical graph-theoretic tools such as dynamic programming [5, Chapter 1], Theorem 2 allows it to be translated into a suboptimal value function v_1 for the concrete system.

IV. NUMERICAL EXPERIMENTS

In this section, we illustrate the benefits of our approach on two different optimal control problems.

Before stating the numerical results, we outline the abstraction construction. We construct our abstractions \mathcal{S}_2 in the same fashion as in [17], but allow multiple timesteps. To construct the abstract system $\mathcal{S}_2 = (\mathcal{X}_2, \mathcal{U}_2, F_2)$, we consider a uniform partition of the concrete state space \mathcal{X}_1 . Formally, the abstract state space is defined as $\mathcal{X}_2 := \{q_1, \dots, q_{N_q}\}$, where N_q is the number of partition cells. The relation $R \subseteq \mathcal{X}_1 \times \mathcal{X}_2$ is defined such that $(x_1, q) \in R$ if and only if $x_1 \in \mathcal{X}_1$ lies in the cell represented by $q \in \mathcal{X}_2$. The transition function F_2 is computed via a growth bound [17, Definition VIII.2] to over-approximate the concrete *reachable sets* $F_1(R^{-1}(q), (u, \tau))$ for all $q \in \mathcal{X}_2$ and $(u, \tau) \in \mathcal{U}_2$. By Theorem 1, this ensures that $\mathcal{S}_1 \preceq_R^{\text{FRR}} \mathcal{S}_2$. The abstract reach-avoid specification Σ_2^{Reach} is constructed analogously to [17], ensuring that $\Sigma_1^{\text{Reach}} \subseteq R^{-1}(\Sigma_2^{\text{Reach}})$. Then, by Proposition 1, any controller \mathcal{C}_2 solving $(\mathcal{S}_2, \Sigma_2^{\text{Reach}})$ can be concretized into a controller $\mathcal{C}_1 = \mathcal{C}_2 \circ R$ that solves $(\mathcal{S}_1, \Sigma_1^{\text{Reach}})$. Finally, to ensure a fair comparison, we keep abstractions of the same size by fixing the number of input-timesteps pairs.

A. Effort-optimal control of a simple pendulum

To illustrate the theoretical framework of the previous section, we consider a frictionless simple pendulum. This continuous time system is given by $\mathcal{D} = (\mathcal{X}, \mathcal{U}, f)$, with $\mathcal{X} = [-\pi, 2\pi]$, $\mathcal{U} = [-6, 6]$ and dynamics defined by $f((\theta, \omega), u) = \begin{pmatrix} \omega \\ -\frac{g}{l} \sin(\theta) + u \end{pmatrix}$, where $g = 9.81[m/s^2]$ is gravity, l is the length of the rod and u is the input torque.

Given the transition system $\mathcal{S}_1 = (\mathcal{X}_1, \mathcal{U}_1, F_1)$ associated with \mathcal{D} as in Definition 7, we define the concrete specification $\Sigma_1^{\text{Reach}} = [\mathcal{I}_1, \mathcal{F}_1, \mathcal{O}_1]$ with $\mathcal{I}_1 = [-\frac{5\pi}{180}, \frac{5\pi}{180}] \times [-0.5, 0.5]$, $\mathcal{F}_1 = [\pi - \frac{17.5}{180}, \pi + \frac{17.5\pi}{180}] \times [-1, 1]$ and $\mathcal{O}_1 = \emptyset$.

We consider a cost function $c_1(x_1, (u, \tau)) := u^2\tau + \epsilon^1$. One can interpret this cost as a measure of the effort that is needed to reach the specification. The abstract transition cost is then given by $c_2(q, (u, \tau)) := u^2\tau + \epsilon$, which satisfies condition (14), since c_1 is independent of the concrete state x_1 . As a result, by Theorem 2, given a Bellman value function v_2 and its associated controller \mathcal{C}_2 for \mathcal{S}_2 , one can construct a suboptimal value function $v_1 = v_2 \circ R$ for \mathcal{S}_1 , which provides an upper bound on the worst-case cost under the concretized controller $\mathcal{C}_1 = \mathcal{C}_2 \circ R$.

Results are summarized in Table I. We can observe that with our method we get a feasible abstract problem, which outperforms controllers synthesized with a single timestep. An example trajectory as well as a representation of the suboptimal value function is given in Figure 2.

¹The term ϵ is added to meet the assumptions of Proposition 2, as without this, $u = 0$ leads to a zero cost. In practice, this ϵ can be chosen arbitrarily close to zero.




Fig. 2: Controlled trajectory of the simple pendulum. The red region denotes the target set \mathcal{F}_1 . The heatmap shows the computed suboptimal value function v_1 for \mathcal{S}_1 . Arrow colors indicate the timestep length used along the closed-loop trajectory—the lighter, the shorter.

Timestep(s) [s]	Worst-case effort guarantee ($\mathcal{L}(\mathcal{C}_1)$)
0.1	Infeasible
0.15	81.47
0.2	Infeasible
0.25	Infeasible
0.3	Infeasible
0.1, 0.15, 0.2, 25, 0.3	65.70

TABLE I: Worst-case effort guarantees for the simple pendulum problem for different individual sampling steps. The last row shows the results of our method combining multiple timesteps, with a coarser input space discretization.

B. Time-optimal path-planning

We consider a path-planning problem in which a vehicle must navigate through a two-dimensional maze. The system has three state variables: the position (x, y) in the plane and the orientation θ . It is controlled via two inputs: wheel speed u_1 and steering angle u_2 . This model has previously been studied in the context of abstraction-based control [17].

The continuous-time system is given by $\mathcal{D} = (\mathcal{X}, \mathcal{U}, f)$, with state space $\mathcal{X} = [-0.1, 10.1] \times [-0.1, 10.1] \times [-\pi - 0.4, \pi + 0.4]$, input space $\mathcal{U} \subseteq [-1, 1] \times [-1, 1]$, and dynamics defined by $f((x, y, \theta), (u_1, u_2)) = \begin{pmatrix} u_1 \cos(\alpha + \theta) \cos(\alpha)^{-1} \\ u_1 \sin(\alpha + \theta) \cos(\alpha)^{-1} \\ u_1 \tan(u_2) \end{pmatrix}$, where $\alpha = \arctan\left(\frac{\tan(u_2)}{2}\right)$.

Given the transition system $\mathcal{S}_1 = (\mathcal{X}_1, \mathcal{U}_1, F_1)$ associated with \mathcal{D} , we define the concrete specification Σ_1^{Reach} , which constrains only the first two state variables, as illustrated in Figure 3 and its caption.

We consider the concrete transition cost $c_1(x_1, (u, \tau)) := \tau$, which corresponds to minimizing the total time required to satisfy the reach-avoid specification. The abstract transition cost is $c_2(q, (u, \tau)) := \tau$, which satisfies condition (14), since c_1 is independent of the concrete state x_1 .

Results are summarized in Table II. Again, a combination of timesteps is used and the worst-case cost is improved. See on Figure 3 an example trajectory obtained with our approach.

Timestep(s) [s]	Worst-case time guarantee ($\mathcal{L}(\mathcal{C}_1)$) [s]
0.2	Infeasible
0.5	Infeasible
0.8	144.0
0.2, 0.5, 0.8	114.60
0.3, 0.7, 1.1, 1.5, 1.9	94.1

TABLE II: Worst-case time guarantees for the path planning problem, for different individual sampling steps. The last 2 rows shows the result of our method combining multiple timesteps, with a coarser input space discretization. The last row is obtained with the data-driven approach (Subsection IV-C).




Fig. 3: An example trajectory, projected onto the first two components of the system, for the path-planning problem. The initial set $\mathcal{I}_1 = \{x_0\}$, target set \mathcal{F}_1 , and obstacle set \mathcal{O}_1 are shown in green, red, and black, respectively. The color of the arrow indicates the used timestep.

C. Data-driven approach

Growth-bound functions provide over-approximations of the reachable sets $F_1(R^{-1}(q), (u, \tau))$, necessary so that Theorem 1 holds. However, they become more and more conservative as the timestep increases [17]. This can lead to infeasibility of the abstract problem by introducing non-determinism, and prevents us from experimenting with longer timesteps, as the over-approximation may intersect obstacles and be discarded.

To mitigate this, we propose a data-driven abstraction approach. For each abstract state-input pair $(q, (u, \tau))$ of the abstract system, we sample N points $x_1 \in R^{-1}(q)$. We observe $x_1^+ \in F_1(x_1, (u, \tau))$, and if there exists $q^+ \in \mathcal{X}_2$ such that $(x_1^+, q^+) \in R$, we add a transition $q^+ \in F_2(q, (u, \tau))$.

This has the effect of *under-approximating* the reachable set, as we may overlook a trajectory. As a result, Theorem 1 no longer applies. Using this data-driven approach, we are able to consider larger timesteps, as shown in Table II, which leads to improved expected worst-case cost. In Figure 4, we show a sample trajectory obtained with this method. As with the growth-bound approach, the controller uses a combination of different timesteps.

V. CONCLUSIONS AND FUTURE WORKS

In this work, we introduced a temporal abstraction framework for abstraction-based control of continuous-time sys-




Fig. 4: An example trajectory for the path-planning problem for the **data driven approach**. The color of the arrows indicate the duration of the timestep.

tems. By allowing multiple timesteps within a unified abstraction, we addressed the trade-off between conservatism and feasibility that arises in fixed-step discretizations. We proved that these abstractions preserve the feedback refinement relation, enabling formal guarantees to be carried over from the abstract to the concrete system. We further demonstrated the effectiveness of our approach through numerical experiments, showing improved worst-case performance without increasing the size of the abstractions.

Future work includes optimizing time and space discretization, reducing conservatism in reachability computation, and exploring probabilistic guarantees in data-driven settings.

REFERENCES

- [1] M. Althoff and J. M. Dolan. Online verification of automated road vehicles using reachability analysis. *IEEE Transactions on Robotics*, 30(4):903–918, 2014.
- [2] M. Althoff, G. Frehse, and A. Girard. Set propagation techniques for reachability analysis. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(1):369–395, 2021.
- [3] R. Alur. *Principles of cyber-physical systems*. MIT press, 2015.
- [4] R. Bellman. *Dynamic Programming*. Dover Publications, 1957.
- [5] D. Bertsekas. *Dynamic programming and optimal control: Volume 1*, volume 1. Athena scientific, 2005.
- [6] J. Calbert, A. Girard, and R. M. Jungers. Classification of simulation relations for symbolic control. *arXiv preprint arXiv:2410.06083*, 2024.
- [7] L. N. Egidio, T. A. Lima, and R. M. Jungers. State-feedback abstractions for optimal control of piecewise-affine systems. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 7455–7460. IEEE, 2022.
- [8] J. H. Gillula, G. M. Hoffmann, H. Huang, M. P. Vitus, and C. J. Tomlin. Applications of hybrid reachability analysis to robotic aerial vehicles. *The International Journal of Robotics Research*, 2011.
- [9] P. Hartman. *Ordinary differential equations*. SIAM, 2002.
- [10] E. Ivanova and A. Girard. Lazy safety controller synthesis with multi-scale adaptive-sampling abstractions of nonlinear systems. *IFAC-PapersOnLine*, 53(2):1837–1843, 2020.
- [11] Z. Kader, A. Saoud, and A. Girard. Safety controller design for incrementally stable switched systems using event-based symbolic models. In *18th European Control Conference (ECC)*, pages 1269–1274. IEEE, 2019.
- [12] K.-D. Kim and P. R. Kumar. Cyber-physical systems: A perspective at the centennial. *Proceedings of the IEEE*, 100(Special Centennial Issue):1287–1308, 2012.
- [13] J. C. Knight. Safety critical systems: challenges and directions. In *Proceedings of the 24th international conference on software engineering*, pages 547–550, 2002.
- [14] E. A. Lee. Introducing embedded systems: a cyber-physical approach. In *Proceedings of the 2009 Workshop on Embedded Systems Education*, pages 1–2, 2009.

- [15] B. Legat, J. Bouchat, and R. M. Jungers. Abstraction-based branch and bound approach to q-learning for hybrid optimal control. In *Proceedings of the 3rd Annual Learning for Dynamics & Control Conference*, 2021.
- [16] S. Meyn. *Control Systems and Reinforcement Learning*. 2022.
- [17] G. Reissig, A. Weber, and M. Rungger. Feedback refinement relations for the synthesis of symbolic controllers. *IEEE Transactions on Automatic Control*, 62(4):1781–1796, 2016.
- [18] R. S. Sutton, D. Precup, and S. Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211, 1999.
- [19] P. Tabuada. *Verification and control of hybrid systems: a symbolic approach*. Springer Science & Business Media, 2009.