The Complete Proof Theory of Hybrid Systems

André Platzer

November 17, 2011
CMU-CS-11-144

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA

This material is based upon work supported by the National Science Foundation under NSF CAREER Award CNS-1054246. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution or government.

Keywords: proof theory; hybrid dynamical systems; differential dynamic logic; axiomatization; completeness

Abstract

Hybrid systems are a fusion of continuous dynamical systems and discrete dynamical systems. They freely combine dynamical features from both worlds. For that reason, it has often been claimed that hybrid systems are more challenging than continuous dynamical systems and than discrete systems. We now show that, proof-theoretically, this is not the case. We present a complete proof-theoretical alignment that interreduces the discrete dynamics and continuous dynamics of hybrid systems. We give a sound and complete axiomatization of hybrid systems relative to continuous dynamical systems and a sound and complete axiomatization of hybrid systems relative to discrete dynamical systems. Thanks to our axiomatization, proving properties of hybrid systems is exactly the same as proving properties of continuous dynamical systems and again, exactly the same as proving properties of discrete dynamical systems. This fundamental cornerstone sheds light on the nature of hybridness and enables flexible and provably perfect combinations of discrete reasoning with continuous reasoning that lift to all aspects of hybrid systems and their fragments.

1 Introduction

Hybrid systems are dynamical systems that combine discrete dynamics and continuous dynamics. They play an important role in modeling systems that use computers to control physical systems. Hybrid systems feature (iterated) difference equations for discrete dynamics and differential equations for continuous dynamics. They, further, combine conditional switching, nondeterminism, and repetition. The theory of hybrid systems concluded that very limited classes of systems are undecidable [AM98, Hen96, CL00]. Most hybrid systems research since focused on practical approaches for efficient approximate reachability analysis for classes of hybrid systems [ADG03, Col07, PC07, GG09]. Undecidability also did not stop researchers in program verification from making impressive progress. This progress, however, concerned both the practice and the theory, where logic was the key to studying the theory beyond undecidability [Pra76, Coo78, HMP77, HKT00, Lei06].

We take a logical perspective, with which we study the logical foundations of hybrid systems and obtain interesting proof-theoretical relationships in spite of undecidability. We have developed a logic and proof calculus for hybrid systems [Pla08, Pla10b] in which it becomes meaningful to investigate concepts like "what is true for a hybrid system" and "what can be proved about a hybrid system" and investigate how they are related. Our proof calculus is *sound*, i.e., all it can prove is true. Soundness should be *sine qua non* for formal verification, but is so complex for hybrid systems [PC07] that it is often inadvertently forsaken. In logic, we can simply ensure soundness locally per proof rule.

More intriguingly, however, our logical setting also enables us to ask the converse: is the proof calculus *complete*, i.e., can it prove all that is true? A corollary to Gödel's incompleteness theorem shows, however, that hybrid systems do not have a sound and complete calculus that is effective, because both their discrete fragment and their continuous fragment alone are nonaxiomatizable since each can define integer arithmetic [Pla08, Theorem 2]. But logic can do better. The suitability of an axiomatization can still be established by showing completeness relative to a fragment [Coo78, HMP77]. This *relative completeness*, in which we assume we were able to prove valid formulas in a fragment and prove that we can then prove all others, also tells us how subproblems are related computationally. It tells us whether one subproblem dominates the others. Standard relative completeness [Coo78, HMP77], however, which works relative to the data logic, is inadequate for hybrid systems, whose complexity comes from the dynamics, not the data logic, first-order real arithmetic, which is decidable [Tar51].

In this paper, we answer an open problem about hybrid systems proof theory [Pla08]. We prove that differential dynamic logic ($\mathsf{d}\mathcal{L}$), which is a logic of hybrid systems, has a sound and complete axiomatization relative to its discrete fragment. This is the first discrete relative completeness result for hybrid systems.

Together with our previous result of a sound and complete axiomatization of hybrid systems relative to the continuous fragment, we obtain a complete alignment of the proof theories of hybrid systems, of continuous dynamical systems, and of discrete dynamical systems. Even though these classes of dynamical systems seem to have quite different intuitive expressiveness, their proof theories actually align perfectly and make them (provably) interreducable. Our $\mathsf{d}\mathcal{L}$ calculus can prove properties of hybrid systems exactly as good as properties of continuous systems can be

proved, which, in turn, our calculus can do exactly as good as discrete systems can be proved. Exactly as good as any one of those subquestions can be solved, dL can solve all others. Relative to the fragment for either system class, our dL calculus can prove all valid properties for the others. It lifts any approximation for the fragment perfectly to all hybrid systems. This also defines a relative decision procedure for dL sentences, because our completeness proofs are constructive.

On top of its theoretical value and the full provability alignment that our new result shows, our discrete completeness result is significant in that—in computer science and verification—programs are closer to being understood than differential equations. Well-established and (partially) automated machinery exists for classical program verification, which, according to our result, has unexpected direct applications in hybrid systems. Completeness relative to discrete systems increases the confidence that discrete computers can solve hybrid systems questions at all. Conversely, control theory provides valuable tools for understanding continuous systems. Previously, it had been just as hard to generalize discrete computer science techniques to continuous questions as it has been to generalize continuous control approaches to discrete phenomena, let alone to the mixed case of hybrid systems.

Overall, our results provide a perfect link between both worlds and allow—in a sound and complete, and constructive way—to combine the best of both worlds. dL allows discrete reasoning as well as continuous reasoning within one single logic and proof system. The dL calculus links and transfers one side of reasoning in a provably perfect (that is sound and complete) way to the other side. For whatever question about a hybrid system (or its fragments) a discrete approach is more natural or promising, dL lifts this reasoning in a perfect way to continuous systems, and to hybrid systems, and vice versa for any part where a continuous approach is more useful.

This complete alignment of the proof theories is a fundamental cornerstone for understanding hybridness and relations between discrete and continuous dynamics. In a nutshell, we can proof-theoretically equate:

"*hybrid* = *continuous* = *discrete*"

2 Differential Dynamic Logic

2.1 Regular Hybrid Programs

We use (regular) *hybrid programs* (HP) [Pla08] as hybrid system models. HPs form a Kleene algebra with tests [Koz97]. Atomic HPs are instantaneous discrete jump *assignments* $x := \theta$, *tests* $?\chi$ of a first-order formula[1] χ, and *differential equation (systems)* $x' = \theta \,\&\, \chi$ for a continuous evolution restricted to the domain of evolution χ. Compound HPs are generated from atomic HPs by nondeterministic choice ($\cup$), sequential composition (;), and Kleene's nondeterministic repetition (*). We use polynomials with rational coefficients as terms. HPs are defined by the following grammar (α, β are HPs, x a variable, θ a term possibly containing x, and χ a formula of first-order logic of real arithmetic):

$$\alpha, \beta \ ::= \ x := \theta \mid ?\chi \mid x' = \theta \,\&\, \chi \mid \alpha \cup \beta \mid \alpha; \beta \mid \alpha^*$$

[1] The test $?\chi$ means "if χ then *skip* else *abort*".

These operations can define all hybrid systems [Pla10b]. We, e.g., write $x' = \theta$ for the unrestricted differential equation $x' = \theta \,\&\, true$. We allow differential equation systems and use vectorial notation. Vectorial assignments are definable from scalar assignments (and ;).

A *state* ν is a mapping from variables to $\mathbb{R}$. We denote the value of term θ in ν by $[\![\theta]\!]_\nu$. Each HP α is interpreted semantically as a binary reachability relation $\rho(\alpha)$ over states, defined inductively by

- $\rho(x := \theta) = \{(\nu, \omega) : \omega = \nu \text{ except } \omega(x) = [\![\theta]\!]_\nu\}$
- $\rho(?\chi) = \{(\nu, \nu) : \nu \models \chi\}$
- $\rho(x' = \theta \,\&\, \chi) = \{(\varphi(0), \varphi(r)) \;:\; \varphi(t) \models x' = \theta \text{ and } \varphi(t) \models \chi \text{ for all } 0 \leq t \leq r \text{ for a solution } \varphi \text{ of any duration } r\}$; i.e., with $\varphi(t)(x') \overset{\text{def}}{=} \frac{\mathsf{d}\varphi(\zeta)(x)}{\mathsf{d}\zeta}(t)$, φ solves the differential equation and satisfies χ at all times [Pla08]
- $\rho(\alpha \cup \beta) = \rho(\alpha) \cup \rho(\beta)$
- $\rho(\alpha; \beta) = \rho(\beta) \circ \rho(\alpha)$
- $\rho(\alpha^*) = \bigcup_{n \in \mathbb{N}} \rho(\alpha^n)$ with $\alpha^{n+1} \equiv \alpha^n; \alpha$ and $\alpha^0 \equiv ?true$.

We refer to our book [Pla10b] for a comprehensive background. We also refer to [Pla10b] for an elaboration how the case $r = 0$ (in which the only condition is $\varphi(0) \models \chi$) is captured by the above definition. To avoid technicalities, we consider only polynomial differential equations, which are all smooth.

2.2 dL Formulas

The *formulas of differential dynamic logic* (dL) are defined by the grammar (where ϕ, ψ are dL formulas, θ_1, θ_2 terms, $\sim \in \{=, \geq, >\}$, x a variable, α a HP):

$$\phi, \psi \;::=\; \theta_1 \sim \theta_2 \mid \neg\phi \mid \phi \wedge \psi \mid \forall x\, \phi \mid [\alpha]\phi$$

The satisfaction relation $\nu \models \phi$ is as usual in first-order logic (of real arithmetic) with the addition that $\nu \models [\alpha]\phi$ iff $\omega \models \phi$ for all ω with $(\nu, \omega) \in \rho(\alpha)$. The operator $\langle\alpha\rangle$ dual to $[\alpha]$ is defined by $\langle\alpha\rangle\phi \equiv \neg[\alpha]\neg\phi$. Consequently, $\nu \models \langle\alpha\rangle\phi$ iff $\omega \models \phi$ for some ω with $(\nu, \omega) \in \rho(\alpha)$. Operators $\leq, <, \vee, \rightarrow, \leftrightarrow, \exists x$ can be defined as usual. A dL formula ϕ is *valid*, written $\vDash \phi$, iff $\nu \models \phi$ for all states ν.

2.3 Axiomatization

Our axiomatization of dL is shown in Figure 1. To highlight the logical essentials, we present a significantly simplified axiomatization in comparison to our earlier work [Pla08], which was tuned for automation. The axiomatization we use here is closer to that of Pratt's dynamic logic for

$$[:=]\quad [x:=\theta]\phi(x) \leftrightarrow \phi(\theta)$$

$$[?]\quad [?\chi]\phi \leftrightarrow (\chi \rightarrow \phi)$$

$$[']\quad [x'=\theta]\phi \leftrightarrow \forall t{\geq}0\, [x:=y(t)]\phi \qquad (y'(t)=\theta)$$

$$[\&]\quad \begin{array}{l}[x'=\theta\,\&\,\chi]\phi \\ \leftrightarrow \forall t_0{=}x_0\, [x'=\theta]\big([x'=-\theta](x_0 \geq t_0 \rightarrow \chi) \rightarrow \phi\big)\end{array}$$

$$[\cup]\quad [\alpha\cup\beta]\phi \leftrightarrow [\alpha]\phi \wedge [\beta]\phi$$

$$[;]\quad [\alpha;\beta]\phi \leftrightarrow [\alpha][\beta]\phi$$

$$[^{*n}]\quad [\alpha^*]\phi \leftrightarrow \phi \wedge [\alpha][\alpha^*]\phi$$

$$\mathrm{K}\quad [\alpha](\phi\rightarrow\psi)\rightarrow([\alpha]\phi\rightarrow[\alpha]\psi)$$

$$\mathrm{I}\quad [\alpha^*](\phi\rightarrow[\alpha]\phi)\rightarrow(\phi\rightarrow[\alpha^*]\phi)$$

$$\mathrm{C}\quad \begin{array}{l}[\alpha^*]\forall v{>}0\,(\varphi(v)\rightarrow\langle\alpha\rangle\varphi(v-1)) \\ \rightarrow \forall v\,(\varphi(v)\rightarrow\langle\alpha^*\rangle\exists v{\leq}0\,\varphi(v))\end{array} \qquad (v\notin\alpha)$$

$$\mathrm{B}\quad \forall x\,[\alpha]\phi\rightarrow[\alpha]\forall x\,\phi \qquad (x\notin\alpha)$$

$$\mathrm{V}\quad \phi\rightarrow[\alpha]\phi \qquad (FV(\phi)\cap BV(\alpha)=\emptyset)$$

$$\mathrm{G}\quad \frac{\phi}{[\alpha]\phi}$$

Figure 1: Differential dynamic logic axiomatization

conventional discrete programs [Pra76, HMP77]. We use the first-order Hilbert calculus (modus ponens and $\forall$-generalization) as a basis and allow all instances of valid formulas of first-order real arithmetic as axioms. The first-order theory of real-closed fields is decidable [Tar51]. We write $\vdash \phi$ iff $\mathsf{d}\mathcal{L}$ formula ϕ can be *proved* with $\mathsf{d}\mathcal{L}$ rules from $\mathsf{d}\mathcal{L}$ axioms (including first-order rules and axioms).

Axiom $[:=]$ is Hoare's assignment rule. Formula $\phi(\theta)$ is obtained from $\phi(x)$ by *substituting* θ for x, provided x does not occur in the scope of a quantifier or modality binding x or a variable of θ. A modality $[\alpha]$ containing $z:=$ or z' *binds* z (written $z\in BV(\alpha)$). In axiom $[']$, $y(\cdot)$ is the (unique [Wal98, Theorem 10.VI]) solution of the symbolic initial value problem $y'(t)=\theta, y(0)=x$. It goes without saying that variables like t are fresh in Figure 1. Axiom $[^{*n}]$ is the iteration axiom. Axiom K is the modal modus ponens from modal logic [HC96]. Axiom I is an induction schema for loops. Axiom C, in which v does not occur in α (written $v\notin\alpha$), is a variation of Harel's convergence rule, suitably adapted to hybrid systems over $\mathbb{R}$. Axiom B is the Barcan formula of first-order modal logic, characterizing anti-monotonic domains [HC96]. In order for it to be sound

for dL, x must not occur in α. The converse of B is provable[2] [HC96, BFC p. 245] and we also call it B. Axiom V is for vacuous modalities and requires that no free variable of ϕ (written $FV(\phi)$) is bound by α. The converse holds, but we do not need it. Rule G is Gödel's necessitation rule for modal logic [HC96]. Note that, unlike rule G, axiom V crucially requires the variable condition that ensures that the value of ϕ is not affected by running α.

We add the new modular dL axiom [&] that reduces differential equations with evolution domain constraints to differential equations without them by checking the evolution domain constraint backwards along the reverse flow. It checks χ backwards from the end up to the initial time t_0, using that $x' = -\theta$ follows the same flow as $x' = \theta$, but backwards. See Figure 2 for an

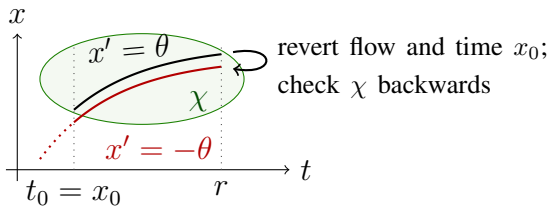

Figure 2: "There and back again" axiom [&] checks evolution domain along backwards flow over time

illustration. To simplify notation, we assume that the (vector) differential equation $x' = \theta$ in [&] already includes a clock $x_0' = 1$ for tracking time.

The following loop invariant rule *ind* derives from G and I. Convergence rule *con* derives from $\forall$-generalization, G, and C (like in C, v does not occur in α):

$$(ind)\ \frac{\phi \to [\alpha]\phi}{\phi \to [\alpha^*]\phi} \qquad (con)\ \frac{\varphi(v) \wedge v > 0 \to \langle\alpha\rangle\varphi(v-1)}{\varphi(v) \to \langle\alpha^*\rangle\exists v{\leq}0\, \varphi(v)}$$

While this is not the focus of this paper, we note that we have successfully used a refined sequent calculus variant of the Hilbert calculus in Figure 1 for automatic verification of hybrid systems, including trains, cars, and aircraft; see [Pla08, Pla10b].

[2] From $\forall x\, \phi \to \phi$, derive $[\alpha](\forall x\, \phi \to \phi)$ by G, from which K and propositional logic derive $[\alpha]\forall x\, \phi \to [\alpha]\phi$. Then, first-order logic derives $[\alpha]\forall x\, \phi \to \forall x\, [\alpha]\phi$, as x is not free in the antecedent.

$$(\overleftarrow{\Delta})\ [x' = f(x)]F \leftarrow \exists h_0{>}0\, \forall 0{<}h{<}h_0\, [(x := x + hf(x))^*]F \qquad (\textit{closed } F)$$

$$(\overrightarrow{\Delta})\ [x' = f(x)]F \to \forall t{\geq}0\, \exists h_0{>}0\, \forall 0{<}h{<}h_0\, [(x := x + hf(x))^*](t \geq 0 \to F) \qquad (\textit{open } F)$$

$$(\overleftrightarrow{\Delta})\ [x' = f(x)]F \leftrightarrow \forall t{\geq}0\, \exists \varepsilon_0{>}0\, \forall 0{<}\varepsilon{<}\varepsilon_0\, \exists h_0{>}0\, \forall 0{<}h{<}h_0\, [(x := x + hf(x))^*]\big(t \geq 0 \to \neg\mathcal{U}_\varepsilon(\neg F)\big) \qquad (\textit{open } F)$$

Figure 3: Discrete Euler approximation axioms (for $f \in C^2$, fresh variables, $\overrightarrow{\Delta}$ and $\overleftrightarrow{\Delta}$ assume $t' = -1$)

3 Continuous Completeness

We have shown that our previous $\mathsf{d}\mathcal{L}$ calculus [Pla08] is a sound and complete axiomatization of $\mathsf{d}\mathcal{L}$ relative to the continuous fragment (FOD). FOD is the *first-order logic of differential equations*, i.e., first-order real arithmetic augmented with formulas expressing properties of differential equations, that is, $\mathsf{d}\mathcal{L}$ formulas of the form $[x' = \theta]F$ with a first-order formula F. We prove that our simplified $\mathsf{d}\mathcal{L}$ axiomatization in Figure 1 is sound and complete relative to FOD (see Appendix):

Theorem 1 (Continuous relative completeness of $\mathsf{d}\mathcal{L}$)**.** *The $\mathsf{d}\mathcal{L}$ calculus is a* sound and complete axiomatization *of hybrid systems relative to FOD, i.e., every valid $\mathsf{d}\mathcal{L}$ formula can be derived from FOD tautologies.*

Axioms V and B are not needed for the proof of Theorem 1; see Appendix. They are included for subsequent results.

4 Discrete Completeness

We study completeness of $\mathsf{d}\mathcal{L}$ relative to the discrete fragment. We denote the *discrete fragment* of $\mathsf{d}\mathcal{L}$ by DL, i.e., the fragment without differential equations (for our purposes we can restrict DL to the operators $:=, ^*$ and allow either ; or vector assignments). The axiomatization in Figure 1 is *not* complete relative to the discrete fragment, since not all differential equations even have closed-form solutions, let alone polynomial solutions. We develop an extension of the $\mathsf{d}\mathcal{L}$ calculus that is complete relative to the discrete fragment by adding an axiom for differential equations.

4.1 Open Discrete Completeness

Axioms like $[']$ that require solutions for differential equations cannot be complete, because most differential equations do not have closed-form solutions. We can understand properties of differential equations from a discrete perspective using discretizations of the dynamics. The question is why that should be complete or even sound. All discretization schemes have errors. Could errors for difficult cases become so large that we cannot obtain conclusive evidence? Or could errors be so unmanageable that they may mislead us into concluding incorrect properties from approximations? Our first step for an answer is for open postconditions.

Theorem 2 (Soundness of approximation)**.** *The approximation axioms in Figure 3 are sound. To simplify notation, we assume that the (vector) differential equation $x' = f(x)$ in $\overrightarrow{\Delta}$ and $\overleftrightarrow{\Delta}$ already includes an extra clock $t' = -1$.*

Before we prove Theorem 2, we develop a number of auxiliary results and consider examples that show why the conditions for the axioms in Figure 3 are necessary. For a set $S \subseteq \mathbb{R}^n$ and $\varepsilon > 0$ we denote the open set $\{x : \|x - y\| < \varepsilon \text{ for a } y \in S\}$ around S by $\mathcal{U}_\varepsilon(S)$. $\overline{\mathcal{U}}_\varepsilon(S)$ is $\{x : \|x - y\| \leq \varepsilon \text{ for a } y \in S\}$. For a logical formula F with the free variable (vector) x and a term ε we define the formula representing the *ε-neighborhood* around F as

$$\mathcal{U}_\varepsilon(F) \overset{\text{def}}{\equiv} \exists y\, (\|x - y\| < \varepsilon \wedge F(y))$$

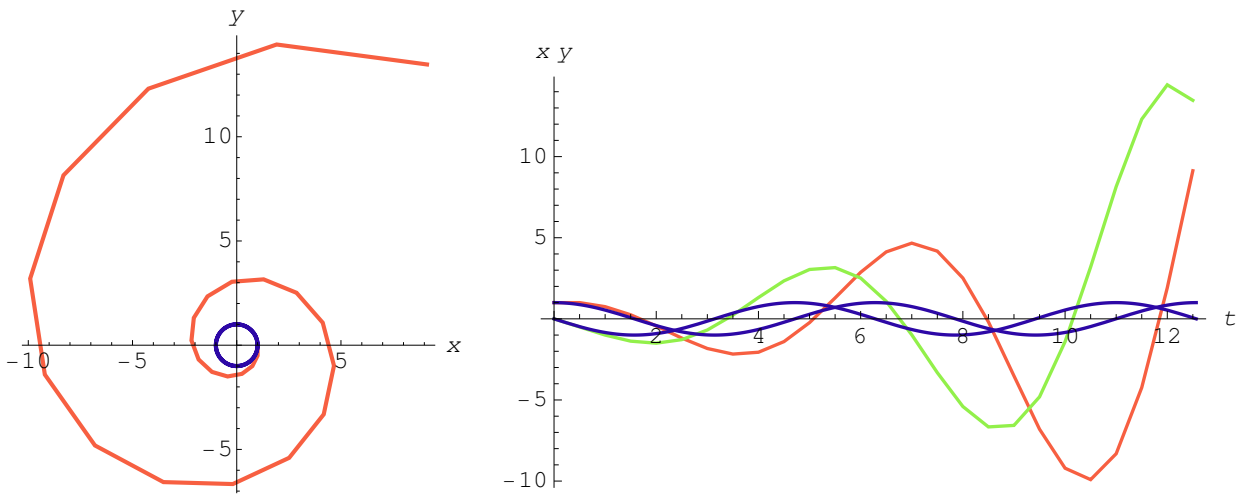

Figure 4: **(top)** Dark circle shows true solution, light line segments show Euler approximation for $h = \frac{1}{4}$ **(bottom)** Dark true bounded trigonometric solution and Euler approximation in lighter colors with increasing errors over time t

The logical formula $\mathcal{U}_\varepsilon(F)$ is indeed true for exactly those values of x that are within distance $<\varepsilon$ from a y satisfying F. Note the nontrivial similarities when comparing axiom $\overleftrightarrow{\Delta}$ with $[']$. The difference is that $[']$ requires a closed-form solution $y(t)$, whereas $\overleftrightarrow{\Delta}$ uses a repeated assignment of the right-hand side $f(x)$ of the differential equation. The latter is appropriate thanks to the extra quantifiers for the approximations. The conditions of the axioms in Figure 3 about F being open/closed are decidable over real-closed fields [Tar51].

Axiom $\overleftarrow{\Delta}$ is incomplete, since the following valid closed property is not provable by $\overleftarrow{\Delta}$, as no approximation, however small h is, works for all time horizons t (see Figure 4 for an illustration):

$$x^2 + y^2 \leq 1.1 \rightarrow [x' = y, y' = -x]x^2 + y^2 \leq 1.1$$

For completeness of approximation schemes, axiom $\overrightarrow{\Delta}$, thus, only states the existence of a h_0 for each time bound t. Axiom $\overleftarrow{\Delta}$ is also insufficient for another reason, because it would be unsound for open F, since the following formula is invalid (Figure 4):

$$x = 1 \wedge y = 0 \rightarrow [x' = y, y' = -x](x \leq 0 \rightarrow x^2 + y^2 > 1)$$

All Euler approximations stay in $x^2 + y^2 > 1$, e.g., when $x \leq 0$, but the dynamics only remains inside its closure $x^2 + y^2 \geq 1$. For the same reason, the converse of $\overrightarrow{\Delta}$ would be unsound for open F, and, thus, is insufficient. For closed F, instead, the converse of $\overrightarrow{\Delta}$ is sound and can be derived from $\overleftarrow{\Delta}$ and simple extra arguments. Unlike its converse, axiom $\overrightarrow{\Delta}$ itself, however, would not be sound for closed F, because no approximation for the following valid formula stays in $x^2 + y^2 = 1$ for any positive duration:

$$x^2 + y^2 = 1 \rightarrow [x' = y, y' = -x]x^2 + y^2 = 1$$

This property *only* holds in the limit case that defines the solution of the differential equation and does not hold for *any* approximation with piecewise polynomial functions. Soundness of axiom $\overrightarrow{\Delta}$ implies, however, that the converse of $\overrightarrow{\Delta}$ can completely prove by approximation that a system does not leave the closure $\overline{F}$ of a postcondition, provided the true dynamics never even leaves its interior $\mathring{F}$. The above examples show, however, that this pair of axioms is incomplete, because they do not align and only prove a weaker closed property and need a stronger open assumption.

To handle properties of differential equations by approximation schemes more completely, we use axiom $\overleftrightarrow{\Delta}$, instead, which, for each time bound t, in addition, quantifies universally over sufficiently small tolerances ε that the discrete approximation tolerates around the reachable states without violating F (as reflected in $\neg\mathcal{U}_\varepsilon(\neg F)$). It is this nesting of quantifiers where $\overleftarrow{\Delta}$ and $\overrightarrow{\Delta}$ "meet" in the sense that both directions of the implication hold. The equivalence axiom $\overleftrightarrow{\Delta}$ completely handles open F. But there are valid properties with closed postconditions F that are still not provable just by $\overleftrightarrow{\Delta}$. The following formula is valid (e.g., provable by a differential invariant [Pla10a]):

$$x^2+y^2\leq 1 \rightarrow [x'=y,y'=-x]x^2+y^2\leq 1 \tag{1}$$

Unfortunately, no Euler approximation for the dynamics, however small h is, satisfies $x^2+y^2\leq 1$ for any duration $t>0$; see Figure 4 for an illustration. The otherwise (i.e., using $\overleftrightarrow{\Delta}$) provable open property

$$x^2+y^2< 1.1 \rightarrow [x'=y,y'=-x]x^2+y^2< 1.1$$

illustrates that $\overleftrightarrow{\Delta}$ would be incomplete if we inverted the order of the quantifiers in $\overleftrightarrow{\Delta}$ to be $\exists\varepsilon{>}0\,\forall t{\geq}0$. Time-uniform approximations are rare. Our approach, instead, uses "proof-uniform" approximations, i.e., one proof for all t, not one value ε for all t. We will answer the question to what extent our approach can always work.

To justify $\overleftrightarrow{\Delta}$, we use an estimate of the global error of Euler approximations in a neighborhood of the solution [SB02, Theorem 7.2.2.3]. For the sake of a self-contained presentation, a proof of Theorem 3 is in the Appendix.

Theorem 3 (Global error)**.** *Let $f\in C^2$, $\hat{x}^0\in\mathbb{R}^n$, and x a solution on $[0,t]$ of $x'=f(x),x(0)=\hat{x}^0$. Let f be Lipschitz-continuous with Lipschitz-constant L on $\mathcal{U}_E(x([0,t]))$ for some $E>0$. Then there is an $h_0>0$ such that for all h with $0<h\leq h_0$ and all $n\in\mathbb{N}$ with $nh\leq t$, the sequence $\hat{x}^{n+1}=\hat{x}^n+hf(\hat{x}^n)$ satisfies:*

$$\|x(nh)-\hat{x}^n\|\leq\frac{h}{2}\max_{\zeta\in[0,t]}\|x''(\zeta)\|\frac{e^{Lt}-1}{L}$$

The following lemmas are proved in the Appendix.

Lemma 4 (Continuous distance)**.** *For a set $S\subseteq\mathbb{R}^n$ the distance $d(\cdot,S):\mathbb{R}^n\to\mathbb{R};x\mapsto\inf_{y\in S}\|x-y\|$ is a continuous map.*

Lemma 5. *Let $K\subseteq F$ a* compact *subset of an open set F. Then $\inf_{x\in K}d(x,F^{\complement})>0$ for complement $F^{\complement}$.*

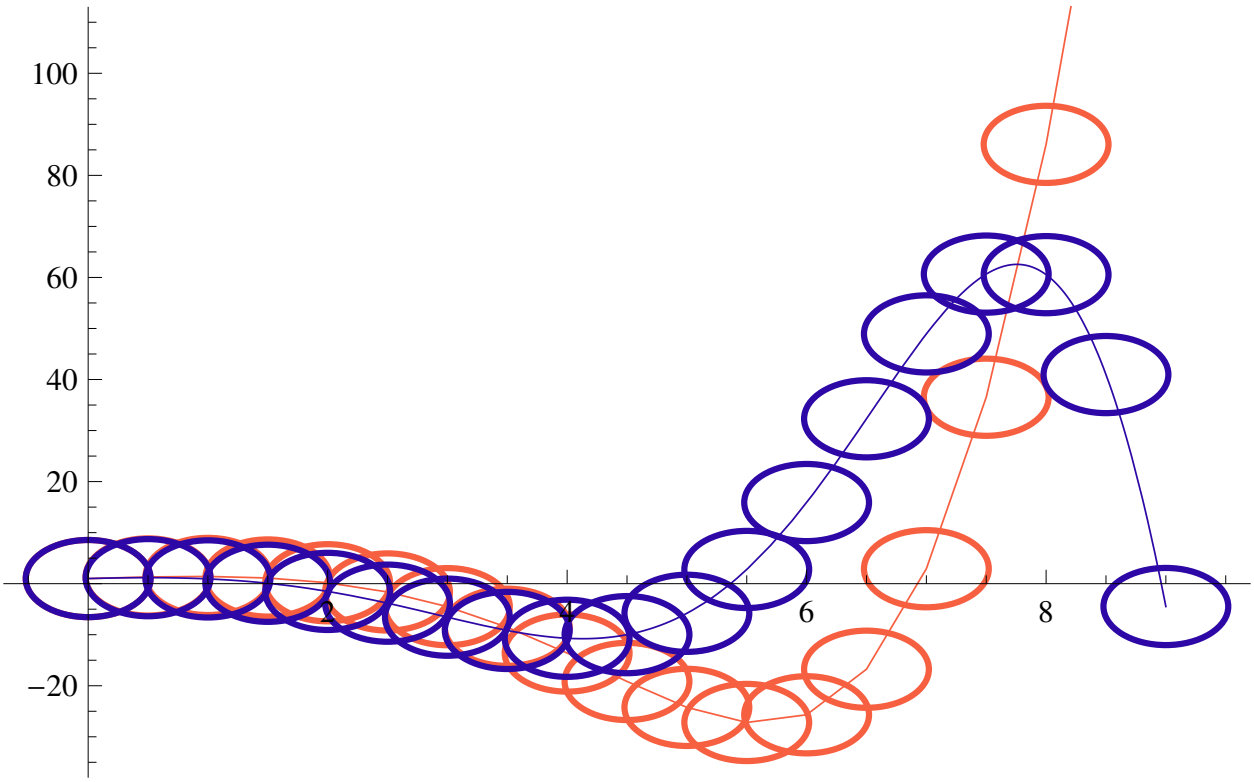

Figure 5: Dark partial covering for dark solution and light partial covering for light approximation

Equipped with this prelude of lemmas and cautionary examples we proceed to prove Theorem 2.

Proof of Theorem 2. $\overleftarrow{\Delta}$: Assume the antecedent is true in a state ν. In order to show that the succedent is true in ν, consider any solution $x(\cdot)$ of $x' = f(x)$ with initial value according to ν. Let $t \geq 0$ be the duration of $x(\cdot)$. We need to show that $x(t) \models F$. Since f is C^1, it is locally Lipschitz continuous and, thus, Lipschitz continuous on every compact subset (these conditions are equivalent for locally compact spaces). Fix an arbitrary $E > 0$. As a continuous image of the compact $x([0,t]) \times \overline{\mathcal{U}}_E(0)$ under addition, $U \overset{\text{def}}{=} \overline{\mathcal{U}}_E(x([0,t])) = \bigcup_{q \in x([0,t])} \overline{\mathcal{U}}_E(q)$ is compact. See Figure 5 for a partial illustration. Let L a Lipschitz constant for f on U. Consider any small h ($0 < h < h_0$ according to the antecedent). Let $\hat{x}^n$ be the value of variable x after n iterations of the discrete program in the antecedent of $\overleftarrow{\Delta}$. By Theorem 3, for sufficiently small h with $nh \leq t$:

$$\|x(nh) - \hat{x}^n\| \leq \underbrace{h \max_{\zeta \in [0,t]} \|x''(\zeta)\| \frac{e^{Lt} - 1}{2L}}_{C(t)} \overset{!}{<} \varepsilon \tag{2}$$

The last inequality holds on $[0,t]$ for all sufficiently small $h > 0$ for the following reason. Since f is C^1, the solution[3] $x(\cdot)$ is C^2. Given the initial state ν, the remaining factor $C(t)$ is a constant

[3] x solves $x' = f(x)$, hence $x \in C$. So the composition $x' = f(x)$ is continuous, hence, $x \in C^1$. Yet then again

depending on t, because the continuous function $x''(\zeta)$ is bounded on the compact set $[0,t]$. Here we need that L for $x'=f(x)$ is determined by ν and t and the choice of E. In short, for any $0<\varepsilon<E$ inequality (2) holds for all sufficiently small $h>0$ (also satisfying $h<h_0$) and all n with $nh\leq t$. Consider $n \overset{\text{def}}{=} \lfloor \frac{t}{h} \rfloor$, which satisfies $nh\leq t$ but $(n+1)h>t$. By mean-value theorem, there is a $\xi\in(nh,t)$ such that

$$\|x(t)-x(nh)\| = \|x'(\xi)\|(t-nh) = \|f(x(\xi))\|(t-nh) \leq \underbrace{\max_{\xi\in[0,t]} \|f(x(\xi))\|}_{=:D(t)}(t-nh) \overset{!}{<} \varepsilon \quad (3)$$

The last inequality holds for all sufficiently small $h>0$ (with $h<h_0$), because $nh\to t$ as $h\to 0$ with $n \overset{\text{def}}{=} \lfloor \frac{t}{h} \rfloor$ and $D(t)$ is a constant. Constant $D(t)$ is determined by t and the initial state for $x'=f(x)$ corresponding to ν, because the continuous function $f(x(\xi))$ is bounded on the compact set $[0,t]$. Combining (2) with (3) we obtain that for any $0<\varepsilon<E$ and all sufficiently small $h>0$ (still $h<h_0$) and $n \overset{\text{def}}{=} \lfloor \frac{t}{h} \rfloor$:

$$\|x(t)-\hat{x}^n\| \leq \|x(t)-x(nh)\| + \|x(nh)-\hat{x}^n\| < 2\varepsilon \quad (4)$$

By antecedent, $\hat{x}^n \models F$ for all these h and n. By (4), there, thus, is a sequence of $\hat{x}^n$ in F that converges to $x(t)$ as $h\to 0$. Thus, $x(t)\models F$, because F is ***closed***.

$\overrightarrow{\Delta}$: Assume $[x'=f(x)]F$ is true in a state ν, which fixes the initial state of the differential equation. According to Picard-Lindelöf [Wal98, Theorem 10.VI], let $x(\cdot)$ be the unique solution (of maximal duration) of $x'=f(x)$ starting with the initial value corresponding to ν. Consider any duration $t\geq 0$ for which $x(\cdot)$ is defined. By assumption, the compact set $x([0,t])$ lies in the region where F is true, which is ***open***. Thus Lemma 5 implies that there is a $\varepsilon_1 \overset{\text{def}}{=} \inf_{q\in x([0,t])} d(q,F^{\complement}) > 0$ so that the open ε_1 ball around each point of $x([0,t])$ is still in F. Here, $F^{\complement}$ is the region of states q with $q \not\models F$. Fix any $0<E<\varepsilon_1$. Then $U \overset{\text{def}}{=} \overline{\mathcal{U}}_E(x([0,t]))$ is in F by construction and, again, compact. Part of this construction is illustrated in Figure 5 Let L be a Lipschitz constant for f on U. Now (2), which follows from Theorem 3, implies for sufficiently small h with $nh\leq t$, that $\|x(nh)-\hat{x}^n\|<E$. Thus, $\hat{x}^n\models F$ for sufficiently small h with $nh\leq t$. Thus, $\exists h_0{>}0\,\forall 0{<}h{<}h_0\,[(x:=x+hf(x))^*](t\geq 0\to F)$ is true in ν where the initial time horizon t was arbitrary. Recall that the decreasing clock $t'=-1$ was assumed to be part of the differential equation $x'=f(x)$ for simplicity. Thus, $nh\leq t$ iff $t\geq 0$ holds after the loop. Note that h_0 depends on t. Relation (4) relates different points in time and bounds the maximum difference of solution $x(\cdot)$ and its discrete approximation $\hat{x}^n$ when they exist for different durations by choosing sufficiently small h.

$\overleftrightarrow{\Delta}$: Like in the proof for $\overrightarrow{\Delta}$, we assume that $\nu\models[x'=f(x)]F$ and, using that F is open, conclude that $\overline{\mathcal{U}}_E(x([0,t]))$ is in F for an $E>0$ that depends on ν and t. Thus (recall that t is a decreasing clock):

$$\nu\models[x'=f(x)]\big(t\geq 0\to\forall z\,(\|z-x\|<E\to F(z))\big) \quad (5)$$

the composition $x'=f(x)$ is C^1, because $f\in C^1$. Henceforth, $x\in C^2$.

By (2) we conclude for all $0 < \varepsilon < \frac{E}{2}$ and sufficiently small h with $nh \leq t$ that $\|x(nh) - \hat{x}^n\| < \varepsilon$. Thus,

$$\|x(nh) - z\| \leq \|x(nh) - \hat{x}^n\| + \|\hat{x}^n - z\| < 2\varepsilon \leq E$$

for all z with $\|\hat{x}^n - z\| < \varepsilon$. Hence, $F(z)$ holds by (5). Let ν_n the state reached after n iterations of the loop in $\overleftrightarrow{\Delta}$, then $\nu_n \models t \geq 0 \rightarrow \forall z\,(\|z - \hat{x}\| < \varepsilon \rightarrow F(z))$, as $\nu_n \models t \geq 0$ iff $\nu \models nh \leq t$, since t is a decreasing clock. Soundness of the "$\rightarrow$" direction of $\overleftrightarrow{\Delta}$ follows with the respective choice $\varepsilon_0 \overset{\text{def}}{=} \frac{E}{2}$ for each t and ν.

The converse "$\leftarrow$" direction of $\overleftrightarrow{\Delta}$ follows from the soundness of $\overleftarrow{\Delta}$ using that $\neg\mathcal{U}_\varepsilon(\neg F)$, which is equivalent to $\forall z\,(\|z - x\| < \varepsilon \rightarrow F(z))$, is closed since the union $\mathcal{U}_\varepsilon(S)$ is open for any S. The proof follows by observing that, for each time bound $t > 0$, the region $t{\geq}0 \rightarrow \neg\mathcal{U}_\varepsilon(\neg F)$ is closed for the purpose of $\overleftrightarrow{\Delta}$, because the solution $x(\cdot)$ cannot leave a closed region on a compact time interval $[0, t]$ unless it already leaves it on $[0, t)$. It is also easy to derive this direction formally from $\overleftarrow{\Delta}$ with corresponding arithmetic. □

To prove Theorem 2, one could simply try a finite covering of the open balls for U, which exists by compactness of $x([0, t])$. The ε neighborhoods of all points of an arbitrary finite covering, however, are not guaranteed to remain within F, see Figure 5 at $t \approx 6$.

4.2 Closed Discrete Completeness

Axiom $\overleftrightarrow{\Delta}$ handles open postconditions of differential equations but not closed postconditions. Even though the property in (1) is a closed region and not provable using $\overleftrightarrow{\Delta}$ alone, this property and other closed F are still provable indirectly using $\mathsf{d}\mathcal{L}$ axioms together with $\overleftrightarrow{\Delta}$. We need the following formula that we derive[4] when no free variable of ϕ is bound in α

$$(\mathbf{V}\vee)\quad \phi \vee [\alpha]\psi \leftrightarrow [\alpha](\phi \vee \psi)$$

Proposition 6. *For every (topologically) closed F, the following formula is provable in $\mathsf{d}\mathcal{L}$*

$$(\mathring{U})\quad [x' = f(x)]F \leftrightarrow \forall \check{\varepsilon}{>}0\,[x' = f(x)]\mathcal{U}_{\check{\varepsilon}}(F)$$

Proof. For a set $S \subseteq \mathbb{R}^n$ we denote its (topological) closure by $\overline{S}$. Since $\mathbb{R}^n$ has a regular topology:

$$\begin{aligned} x \in \overline{S} &\iff \forall \check{\varepsilon}{>}0\ \exists y \in S\ \|x - y\| < \check{\varepsilon} \\ &\iff \forall \check{\varepsilon}{>}0\ \mathcal{U}_{\check{\varepsilon}}(x) \cap S \neq \emptyset \\ &\iff \forall \check{\varepsilon}{>}0\ x \in \mathcal{U}_{\check{\varepsilon}}(S) \\ &\iff x \in \bigcap_{\check{\varepsilon}>0} \mathcal{U}_{\check{\varepsilon}}(S) \end{aligned}$$

[4] "$\rightarrow$": Trivially, $(\phi \vee [\alpha]\psi) \rightarrow (\phi \vee [\alpha]\psi)$, from which V derives $(\phi \vee [\alpha]\psi) \rightarrow ([\alpha]\phi \vee [\alpha]\psi)$. Thus, $(\phi \vee [\alpha]\psi) \rightarrow [\alpha](\phi \vee \psi)$ derives by a consequence [HC96, K4 p. 31] of G.
"$\leftarrow$": Conversely, K derives $[\alpha](\neg\phi \rightarrow \psi) \rightarrow ([\alpha]\neg\phi \rightarrow [\alpha]\psi)$, from which V derives $[\alpha](\neg\phi \rightarrow \psi) \rightarrow (\neg\phi \rightarrow [\alpha]\psi)$.

Set S is closed iff $S = \overline{S}$, i.e., iff $S = \bigcap_{\check{\varepsilon}>0} \mathcal{U}_{\check{\varepsilon}}(S)$. Since F is closed, the following equivalence is valid, hence, provable in real arithmetic

$$F \leftrightarrow \forall \check{\varepsilon}{>}0\, \mathcal{U}_{\check{\varepsilon}}(F) \quad \text{i.e.,} \quad F \leftrightarrow \forall \check{\varepsilon}\, (\neg(\check{\varepsilon}{>}0) \vee \mathcal{U}_{\check{\varepsilon}}(F))$$

Since $\check{\varepsilon}$ does not occur in the dynamics, both sides of $\mathring{U}$ are, thus, equivalent using B and V$\vee$. $\square$

With an extra quantifier, $\mathring{U}$ transforms closed postconditions to open postconditions, which $\overleftrightarrow{\Delta}$ handles. Recall that $\overleftarrow{\Delta}$ also handles closed postconditions, but, unlike $\overleftrightarrow{\Delta}$ together with $\mathring{U}$, axiom $\overleftarrow{\Delta}$ cannot prove them all.

4.3 Discrete Completeness of $\mathsf{d}\mathcal{L}_\Delta = \mathsf{d}\mathcal{L} + \Delta$

Locally closed postconditions (conjunctions $O \wedge C$ of a closed region C and an open O) are handled in a sound and complete way when combining $\overleftrightarrow{\Delta}$,$\mathring{U}$, and the following formula derived from K [HC96, K3 p. 28]

$$([]\wedge)\ \ [\alpha](\phi \wedge \psi) \leftrightarrow [\alpha]\phi \wedge [\alpha]\psi$$

One missing case is where postcondition F is a union $O \vee C$ of an open O and a closed C. We generalize the idea behind Proposition 6 to this case.

Proposition 7. *For every (topologically) open O and (topologically) closed C, the following formula is provable in $\mathsf{d}\mathcal{L}$*

$$(\check{U})\ \ [x' = f(x)](O \vee C) \leftrightarrow \forall \check{\varepsilon}{>}0\, [x' = f(x)](O \vee \mathcal{U}_{\check{\varepsilon}}(C))$$

Proof. As in the proof of Proposition 6, C is closed and $C \leftrightarrow \forall \check{\varepsilon}{>}0 \mathcal{U}_{\check{\varepsilon}}(C)$ valid, and, thus, provable in real arithmetic. Since $\check{\varepsilon}$ is fresh, we, thus, derive equivalence of both sides of $\check{U}$ using V$\vee$ and B

$$\begin{aligned}
[x' = f(x)](O \vee C) &\equiv [x' = f(x)](O \vee \forall \check{\varepsilon}{>}0 \mathcal{U}_{\check{\varepsilon}}(C)) \\
&\equiv [x' = f(x)] \forall \check{\varepsilon}{>}0\, (O \vee \mathcal{U}_{\check{\varepsilon}}(C)) \\
&\equiv \forall \check{\varepsilon}{>}0\, [x' = f(x)](O \vee \mathcal{U}_{\check{\varepsilon}}(C))
\end{aligned}$$ $\square$

Like $\mathring{U}$, $\check{U}$ reduces non-open postconditions to (quantified) open postconditions, which we then want to prove by $\overleftrightarrow{\Delta}$. Can we prove all resulting formulas when they are valid? More generally, can we prove all valid $\mathsf{d}\mathcal{L}$ formulas from discrete DL this way?

The $\mathsf{d}\mathcal{L}$ calculus is complete relative to the continuous fragment (Theorem 1), but incomplete relative to the discrete fragment. We study the $\mathsf{d}\mathcal{L}$ calculus in Figure 1 enriched with the approximation axiom $\overleftrightarrow{\Delta}$ in Figure 3 and denote this calculus by $\mathsf{d}\mathcal{L}_\Delta$. The $\mathsf{d}\mathcal{L}_\Delta$ calculus inherits completeness relative to the continuous fragment from Theorem 1. We now prove that $\mathsf{d}\mathcal{L}_\Delta$ is a *sound and complete axiomatization* of $\mathsf{d}\mathcal{L}$ relative to discrete DL, i.e., every valid $\mathsf{d}\mathcal{L}$ formula can be proved in the $\mathsf{d}\mathcal{L}_\Delta$ calculus from valid DL formulas.

In particular, we need to prove that $\mathsf{d}\mathcal{L}$ can express all required invariants and variants, and the resulting formulas with all their nested repetitions, assignments, differential equations and so on

are provable in the $\mathsf{d}\mathcal{L}_\Delta$ calculus from valid DL facts. This would be a tricky proof. Instead, we prove completeness in an unusual way. We leverage the fact that we have already proved $\mathsf{d}\mathcal{L}$ to be complete relative to the continuous fragment FOD in Theorem 1. Thus, every valid $\mathsf{d}\mathcal{L}$ formula can be proved in the $\mathsf{d}\mathcal{L}$ calculus (and the $\mathsf{d}\mathcal{L}_\Delta$ calculus) from valid FOD formulas. FOD is, in a sense, farthest away from $\mathsf{d}\mathcal{L}_\Delta$, because it only involves differential equations, which is precisely what is missing in DL. But by basing our proof on Theorem 1, we can piggyback on its proof how proofs about repetitions and interactions of discrete and continuous dynamics reduce in a sound and complete way to FOD formulas. So we only need to prove the remaining step that $\mathsf{d}\mathcal{L}_\Delta$ can prove all valid FOD formulas from DL tautologies, which is significantly easier than having to worry about all formulas of $\mathsf{d}\mathcal{L}$.

Theorem 8 (Discrete relative completeness of $\mathsf{d}\mathcal{L}_\Delta$)**.** *The $\mathsf{d}\mathcal{L}_\Delta$ calculus is a* sound and complete axiomatization *of hybrid systems relative to its discrete fragment DL, i.e., every valid $\mathsf{d}\mathcal{L}$ formula can be derived from DL tautologies.*

Proof. Theorems 1 and 2 show that the $\mathsf{d}\mathcal{L}_\Delta$ calculus is sound. We need to show that the $\mathsf{d}\mathcal{L}_\Delta$ calculus can prove all valid $\mathsf{d}\mathcal{L}$ formulas from instances of DL tautologies. By Theorem 1, $\mathsf{d}\mathcal{L}$ is complete relative to its continuous fragment, i.e., elementary properties of differential equations in FOD. Consequently, all valid $\mathsf{d}\mathcal{L}$ formulas can be proved in the $\mathsf{d}\mathcal{L}$ (and $\mathsf{d}\mathcal{L}_\Delta$) calculus from instances of valid FOD formulas. All that remains to be shown is that we can then prove all those valid FOD formulas from valid formulas of discrete DL in the $\mathsf{d}\mathcal{L}_\Delta$ calculus. Consider any valid FOD formula ϕ. We proceed by induction on the structure of ϕ and show that $\mathsf{d}\mathcal{L}_\Delta$ can (provably) translate ϕ into an equivalent DL formula $\phi^\#$ (with the same free variables), which can be proved by assumption. Observe that the construction of $\phi^\#$ from ϕ is effective.

1. When ϕ is a (valid) formula of first-order real arithmetic, then $\phi^\# \overset{\text{def}}{\equiv} \phi$ is already in DL and provable by assumption. First-order real arithmetic is even decidable by quantifier elimination [Tar51].

2. When ϕ is of the form $[x' = f(x)]F$ with a first-order (or semialgebraic) formula F of real arithmetic[5], then, by a standard boolean argument for normal forms applied to semialgebraic sets obtained by quantifier elimination [Tar51], F is *provably* equivalent to a formula of the form
$$\bigwedge_{i=1}^{m} \left(\bigvee_j p_{i,j} > 0 \vee \bigvee_k q_{i,k} \geq 0 \right)$$
with polynomials $p_{i,j}$ and $q_{i,k}$. As a preimage of an open set, the set $\{x \in \mathbb{R}^n : p_{i,j}(x) > 0\}$ is an open set, since $p_{i,j}$ is a continuous function. Dually, the set where $q_{i,k} \geq 0$ is a closed set, because it is the complement of the open set where $-q_{i,k} > 0$. As a union of open sets, the set where $O_i \overset{\text{def}}{\equiv} \bigvee_j p_{i,j} > 0$ holds is open. As a *finite* union of closed sets, the set where

[5] We can assume F to be semialgebraic, because, by Theorem 1, FOD does not need nested modalities since it has quantifiers.

$C_i \overset{\text{def}}{\equiv} \bigvee_k q_{i,k} \geq 0$ holds is closed. This gives the following (provable) equivalence:

$$\vdash F \leftrightarrow \bigwedge_{i=1}^{m} (O_i \vee C_i)$$

Formula $[]\wedge$, which derives from K, thus, derives

$$\vdash \phi \leftrightarrow \bigwedge_{i=1}^{m} [x' = f(x)](O_i \vee C_i)$$

With m uses of $\check{U}$, we derive

$$\vdash \phi \leftrightarrow \bigwedge_{i=1}^{m} \forall \check{\varepsilon}{>}0\, [x' = f(x)](O_i \vee \mathcal{U}_{\check{\varepsilon}}(C_i))$$

Since, for $\check{\varepsilon} > 0$, each $O_i \vee \mathcal{U}_{\check{\varepsilon}}(C_i)$ is open for every i, we, therefore, derive with m uses of axiom $\overset{\leftrightarrow}{\Delta}$ that $\vdash \phi \leftrightarrow \phi^{\#}$ where

$$\phi^{\#} \overset{\text{def}}{\equiv} \bigwedge_{i=1}^{m} \forall \check{\varepsilon}{>}0\, \psi(O_i \vee \mathcal{U}_{\check{\varepsilon}}(C_i))$$

By $\psi(O_i \vee \mathcal{U}_{\check{\varepsilon}}(C_i))$ we denote the DL formula in the right-hand side of axiom $\overset{\leftrightarrow}{\Delta}$ with $O_i \vee \mathcal{U}_{\check{\varepsilon}}(C_i)$ in place of F. Thus, $\vdash \phi \leftrightarrow \phi^{\#}$ is provable in the $\mathsf{d}\mathcal{L}_\Delta$ calculus, $\phi^{\#}$ is in DL, and, thus, provable by assumption.

3. When ϕ is of the form $[x' = f(x)\,\&\,\chi]F$, then it is by axiom $[\&]$ provably equivalent to a formula without evolution domain restrictions, which is structurally simpler and, thus, provable from DL by induction hypothesis.

4. When ϕ is of the form $\neg\psi$, then, by induction hypothesis, the simpler formula ψ is provably equivalent to the DL formula $\psi^{\#}$. This equivalence $\psi \leftrightarrow \psi^{\#}$ is provable in $\mathsf{d}\mathcal{L}_\Delta$ by induction hypothesis. Consequently, ϕ is (in $\mathsf{d}\mathcal{L}_\Delta$) provably equivalent to $\phi^{\#} \overset{\text{def}}{\equiv} \neg(\psi^{\#})$, which is a DL formula and, thus, provable by assumption.

5. When ϕ is of the form $\phi_1 \wedge \phi_2$, then ϕ is provable from DL by induction hypothesis, because both ϕ_1 and ϕ_2 can be turned into DL formulas $\phi_1^{\#}$ and $\phi_2^{\#}$, respectively, with provable $\phi_i \leftrightarrow \phi_i^{\#}$. Thus, $\phi_1 \wedge \phi_2 \leftrightarrow \phi_1^{\#} \wedge \phi_2^{\#}$ is provable in $\mathsf{d}\mathcal{L}_\Delta$.

6. When ϕ is of the form $\forall x\, \psi$, then, by induction hypothesis, ψ is provably equivalent to a DL formula $\psi^{\#}$, i.e., $\psi \leftrightarrow \psi^{\#}$ is provable in $\mathsf{d}\mathcal{L}_\Delta$. Thus, $\forall x\, \psi$ is, by congruence, provably equivalent to $\phi^{\#} \overset{\text{def}}{\equiv} \forall x\, (\psi^{\#})$, which is a DL formula and, thus, provable by assumption. $\square$

The proof of Theorem 8 and its base Theorem 1 and the other proofs in this section are constructive. Hence, there is a constructive way of proving dL formulas by systematic reduction to discrete program properties. The resulting formulas may be unnecessarily complicated, because of the way our proof reduces the completeness of dL_Δ relative to DL to the completeness of dL relative to FOD, which may require turning dL into continuous FOD and then back into discrete DL. Still, the proof is constructive and shows an upper bound on how quantifier alternations increase in the reduction. A more efficient reduction may be sought in practice. Thanks to our result, we now know that this reduction is possible at all.

Note that recursive reductions would be flawed. The validity of dL formulas reduces to that of FOD, which reduces to DL, which again reduces to FOD etc. But we need an approximation to handle either fragment, for we cannot otherwise break this cycle of mutual reductions. This makes approximations of either fragment (or even combined fragments) interesting and ensures that they lift to dL perfectly.

5 Relative Decidability

Our relative completeness results entail relative decidability results for free. Since our relative completeness proofs are constructive and the rules automatable [Pla08], they even define a relative decision procedure. The proof of relative decidability rests on the coincidence lemma for dL, which shows that only the values of free variables of a formula affect its truth-value.

Lemma 9 (Coincidence lemma)**.** *If the states ν and ω agree on all free variables of formula ϕ, then $\nu \models \phi$ iff $\omega \models \phi$.*

Proof. The proof is by a simple structural induction using the definitions of $\nu \models \cdot$ and $\rho(\cdot)$. □

Theorem 10 (Relative decidability)**.** *Validity of dL sentences (i.e., formulas without free variables) is decidable relative to either an oracle for continuous FOD or an oracle for discrete DL.*

Proof. Let ϕ by a sentence in dL and ν a state. Then either $\nu \models \phi$ or $\nu \not\models \phi$. Thus, either $\nu \models \phi$ or $\nu \models \neg\phi$. By coincidence lemma 9, however, $\nu \models \phi$ iff $\omega \models \phi$ for arbitrary ω, because the truth-value of dL formula ϕ is determined entirely[6] by the value of its free variables, of which there are none. Consequently, either $\vDash \phi$ or $\vDash \neg\phi$. In either case, Theorems 8 and 1 imply that the respective valid formula is provable in dL_Δ from valid DL (or FOD) formulas. □

6 Related Work

A general overview of hybrid systems and logics can be found in [ADG03, GG09, DN00, Pla10b]. Hybrid systems are undecidable and do not have finite-state bisimulations [Hen96, AHLP00], so abstractions and approximations are often used. Euler approximations are standard. Discrete approximations have been considered many times before [LT05, Col07, PC07]. Discretizations

[6] The semantics of dL function and predicate symbols is fixed.

have been used for linear systems [GG09], to obtain abstractions of fragments of hybrid systems [AHLP00, ADI06, Tiw08], and to approximate nonlinear systems by hybrid systems [HHWT98] or by piecewise linear dynamics [ADG03] when assuming that error bounds or Lipschitz constants are given. See [Hen96, Col07, PC07] for a discussion of the limits and decidability frontier. These are interesting uses of approximation. But we use approximations for a different, proof-theoretical purpose: to obtain a sound and complete axiomatization relative to properties of discrete programs.

Related approaches do not take a logic and proofs perspective. That made it difficult to formulate appropriate completeness notions, which are natural in logic. Previous completeness-type arguments for hybrid systems were restricted to bounded model checking [ADI06], continuous systems [Tiw08], discrete linear systems on compact domains that are assumed to be so robustly save that simulation is enough [GP06], or assume the system could be changed without affecting the property [HHWT98]. We, instead, prove full relative completeness of an expressive logic relative to a small fragment. Our results identify a more fundamental, proof-theoretical connection between discrete, continuous, and hybrid dynamics. They are also not limited to safety properties but extend to the full expressivity of $\mathsf{d}\mathcal{L}$.

Our notion of relative completeness is inspired by relative completeness for conventional programs, which has been pioneered by Cook [Coo78] and, for dynamic logic of conventional discrete programs [Pra76], by Harel et al. [HMP77, HKT00]. They show that Hoare's and Pratt's program logics are complete relative to an oracle for the first-order logic of the program data. Relative completeness is the standard approach to showing adequacy of calculi for undecidable classical program logics. Those completeness notions are inadequate for hybrid systems, however, because the data logic of hybrid systems is real arithmetic, hence decidable [Tar51]. It is not the data, but the dynamics proper, that causes incompleteness. We, thus, prove completeness relative to fragments of the dynamics.

As an alternative to arithmetical relative completeness notions, Leivant [Lei06] considered completeness of discrete program logics by alignment with proof schemes in higher-order logic. It is not clear how that would generalize to a compelling completeness notion for hybrid systems, whose semantics intimately depends on arithmetical models that are rich enough to give differential equations a well-defined semantics.

Discrete Turing machines have been encoded into classes of hybrid [AM98, Hen96, CL00] or continuous systems [Bra95, GCB07]. Our proof works the other way around and handles full hybrid systems. We use discrete dynamics to understand hybrid dynamics. Our results are also about provability not encodability.

7 Conclusions

We have presented a significantly simplified axiomatization of differential dynamic logic ($\mathsf{d}\mathcal{L}$), our logic for hybrid systems. We have introduced a new axiom for discrete approximation of differential equations based on Euler discretizations. We prove the calculus to be a sound and complete axiomatization of $\mathsf{d}\mathcal{L}$ relative to the continuous fragment (differential equations) and also a sound and complete axiomatization relative to the discrete fragment. Our results show that the proof theory of hybrid systems aligns completely with that of continuous systems *and* with that of

discrete systems. Our axiomatization defines a perfect lifting. Because our proofs are constructive, our axiomatization even defines relative decision procedures for $\mathsf{d}\mathcal{L}$ sentences. Our construction shows how quantifier alternations increase when interreducing dynamics.

Our complete alignment shows that any reasoning technique in one domain has a counterpart in the other. (In)variants, which are the predominant proof technique for loops, have differential (in)variants [Pla10a] as a counterpart of induction for differential equations. Our results indicate a high potential for identifying other practical consequences of our theoretical alignment. They also revitalize and justify the hope that control and computer science techniques *can* work together to understand hybrid systems and can even work together to understand purely discrete or purely continuous systems.

References

[ADG03] Eugene Asarin, Thao Dang, and Antoine Girard. Reachability analysis of nonlinear systems using conservative approximation. In Oded Maler and Amir Pnueli, editors, *HSCC*, volume 2623 of *LNCS*, pages 20–35. Springer, 2003.

[ADI06] Rajeev Alur, Thao Dang, and Franjo Ivancic. Predicate abstraction for reachability analysis of hybrid systems. *ACM Trans. Embedded Comput. Syst.*, 5(1):152–199, 2006.

[AHLP00] Rajeev Alur, Thomas Henzinger, Gerardo Lafferriere, and George J. Pappas. Discrete abstractions of hybrid systems. *Proc. IEEE*, 88(7):971–984, 2000.

[AM98] Eugene Asarin and Oded Maler. Achilles and the tortoise climbing up the arithmetical hierarchy. *J. Comput. Syst. Sci.*, 57(3):389–398, 1998.

[Bra95] Michael S. Branicky. Universal computation and other capabilities of hybrid and continuous dynamical systems. *Theor. Comput. Sci.*, 138(1):67–100, 1995.

[CL00] Franck Cassez and Kim Guldstrand Larsen. The impressive power of stopwatches. In *CONCUR*, pages 138–152, 2000.

[Col07] Pieter Collins. Optimal semicomputable approximations to reachable and invariant sets. *Theory Comput. Syst.*, 41(1):33–48, 2007.

[Coo78] Stephen A. Cook. Soundness and completeness of an axiom system for program verification. *SIAM J. Comput.*, 7(1):70–90, 1978.

[DN00] Jennifer Mary Davoren and Anil Nerode. Logics for hybrid systems. *IEEE*, 88(7):985–1010, July 2000.

[GCB07] Daniel Silva Graça, Manuel L. Campagnolo, and Jorge Buescu. Computability with polynomial differential equations. *Advances in Applied Mathematics*, 2007.

[GG09] Colas Le Guernic and Antoine Girard. Reachability analysis of hybrid systems using support functions. In Ahmed Bouajjani and Oded Maler, editors, *CAV*, volume 5643 of *LNCS*, pages 540–554. Springer, 2009.

[Göd31] Kurt Gödel. Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Mon. hefte Math. Phys.*, 38:173–198, 1931.

[GP06] Antoine Girard and George J. Pappas. Verification using simulation. In João P. Hespanha and Ashish Tiwari, editors, *HSCC*, volume 3927 of *LNCS*, pages 272–286. Springer, 2006.

[HC96] G. E. Hughes and M. J. Cresswell. *A New Introduction to Modal Logic*. Routledge, 1996.

[Hen96] Thomas A. Henzinger. The theory of hybrid automata. In *LICS*, pages 278–292, Los Alamitos, 1996. IEEE Computer Society.

[HHWT98] Thomas A. Henzinger, Pei-Hsin Ho, and Howard Wong-Toi. Algorithmic analysis of nonlinear hybrid systems. *IEEE T. Automat. Contr.*, 43:540–554, 1998.

[HKT00] David Harel, Dexter Kozen, and Jerzy Tiuryn. *Dynamic logic*. MIT Press, Cambridge, 2000.

[HMP77] David Harel, Albert R. Meyer, and Vaughan R. Pratt. Computability and completeness in logics of programs (preliminary report). In *STOC*, pages 261–268. ACM, 1977.

[Koz97] Dexter Kozen. Kleene algebra with tests. *ACM Trans. Program. Lang. Syst.*, 19(3):427–443, 1997.

[Lei06] Daniel Leivant. Matching explicit and modal reasoning about programs: A proof theoretic delineation of dynamic logic. In *LICS*, pages 157–168. IEEE Computer Society, 2006.

[LT05] Ruggero Lanotte and Simone Tini. Taylor approximation for hybrid systems. In Manfred Morari and Lothar Thiele, editors, *HSCC*, volume 3414 of *LNCS*, pages 402–416. Springer, 2005.

[Mor87] Michał Morayne. On differentiability of Peano type functions. *Colloquium Mathematicum*, LIII:129–132, 1987.

[PC07] André Platzer and Edmund M. Clarke. The image computation problem in hybrid systems model checking. In Alberto Bemporad, Antonio Bicchi, and Giorgio Buttazzo, editors, *HSCC*, volume 4416 of *LNCS*, pages 473–486. Springer, 2007.

[Pla08] André Platzer. Differential dynamic logic for hybrid systems. *J. Autom. Reas.*, 41(2):143–189, 2008.

[Pla10a] André Platzer. Differential-algebraic dynamic logic for differential-algebraic programs. *J. Log. Comput.*, 20(1):309–352, 2010.

[Pla10b] André Platzer. *Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics*. Springer, Heidelberg, 2010.

[Pra76] Vaughan R. Pratt. Semantical considerations on Floyd-Hoare logic. In *FOCS*, pages 109–121. IEEE, 1976.

[SB02] Josef Stoer and R. Bulirsch. *Introduction to Numerical Analysis*. Springer, New York, 3rd edition, 2002.

[Tar51] Alfred Tarski. *A Decision Method for Elementary Algebra and Geometry*. University of California Press, Berkeley, 2nd edition, 1951.

[Tiw08] Ashish Tiwari. Abstractions for hybrid systems. *Form. Methods Syst. Des.*, 32(1):57–83, 2008.

[Wal98] Wolfgang Walter. *Ordinary Differential Equations*. Springer, 1998.

A Euler Approximation Proofs

In this section of the appendix, we prove the lemmas from Section 4. For the sake of a self-contained presentation we report an explicit yet standard proof of the error bound for Euler approximation shown in Theorem 3. A more general result can be found in [SB02, Theorem 7.2.2.3].

Proof of Theorem 3. By $f \in C^2$ and footnote 3 we have $x \in C^2$. Consider the variation defined as $\check{x}^{n+1} = \check{x}^n + h\check{\Phi}(hn, \check{x}^n)$ with $\check{x}^0 = \hat{x}^0 = x(0)$ and

$$\check{\Phi}(\zeta, y) \stackrel{\text{def}}{=} \begin{cases} f(y) & \text{if } \|y - x(\zeta)\| \leq E \\ f\left(x(\zeta) + E\frac{y-x(\zeta)}{\|y-x(\zeta)\|}\right) & \text{if } \|y - x(\zeta)\| \geq E \end{cases}$$

Like f, $\check{\Phi}$ is continuous and Lipschitz-continuous in y with Lipschitz-constant L, but, by construction, for all $y \in \mathbb{R}^n$, because $\|x(\zeta) + E\frac{y-x(\zeta)}{\|y-x(\zeta)\|} - x(\zeta)\| \leq E$ for all $\zeta \leq t$. Consider any $n \in \mathbb{N}$. By Taylor approximation for x at nh we know for some $\xi \in (nh, (n+1)h)$ that

$$\begin{aligned}
&\|x((n+1)h) - \check{x}^{n+1}\| \\
&= \|x(nh) + x'(nh)h + \frac{x''(\xi)}{2}h^2 - \check{x}^n - h\check{\Phi}(nh, \check{x}^n)\| \\
&\stackrel{\text{ODE}}{=} \|x(nh) - \check{x}^n + (f(x(nh)) - \check{\Phi}(nh, \check{x}^n))h + \frac{x''(\xi)}{2}h^2\| \\
&= \|x(nh) - \check{x}^n + (\check{\Phi}(nh, x(nh)) - \check{\Phi}(nh, \check{x}^n))h + \frac{x''(\xi)}{2}h^2\| \\
&\leq \|x(nh) - \check{x}^n\| + Lh\|x(nh) - \check{x}^n\| + \frac{h^2}{2}\|x''(\xi)\| \\
&\leq (1 + Lh)\|x(nh) - \check{x}^n\| + \frac{h^2}{2}\max_{\zeta\in[0,t]}\|x''(\zeta)\|
\end{aligned}$$

This error bound holds for any $n \in \mathbb{N}$ starting with error $\|x(0) - \check{x}^0\| = 0$. Thus, recursively, for any n:

$$\begin{aligned}
&\|x(nh) - \check{x}^n\| \\
&\leq (1 + Lh)\|x((n-1)h) - \check{x}^{n-1}\| + \frac{h^2}{2}\max_{\zeta\in[0,t]}\|x''(\zeta)\| \\
&\leq (1 + Lh)\big((1 + Lh)\|x((n-2)h) - \check{x}^{n-2}\| \\
&\qquad + \frac{h^2}{2}\max_{\zeta\in[0,t]}\|x''(\zeta)\|\big) + \frac{h^2}{2}\max_{\zeta\in[0,t]}\|x''(\zeta)\| \\
&\leq \dots \\
&\leq \sum_{k=0}^{n}(1 + Lh)^k\frac{h^2}{2}\max_{\zeta\in[0,t]}\|x''(\zeta)\| \\
&\leq \frac{h^2}{2}\max_{\zeta\in[0,t]}\|x''(\zeta)\|\sum_{k=0}^{n}(e^{Lh})^k
\end{aligned}$$

$$\leq \frac{h^2}{2} \max_{\zeta\in[0,t]} \|x''(\zeta)\| \int_0^n e^{Lht} dt$$
$$\leq \frac{h^2}{2} \max_{\zeta\in[0,t]} \|x''(\zeta)\| \frac{e^{Lhn}-1}{Lh}$$

because $1 + Lh \leq e^{Lh}$ for $Lh \geq 0$, which can be seen by its power series expansion. The next-to-last inequality follows, because the sum is a particlar lower Riemann sum of the integral, since $e^{Lhk} \geq 0$ is monotone in k. Since $0 \leq hn \leq t$ is bounded and $E > 0$, there is an $h_0 > 0$ such that $\|x(nh) - \check{x}^n\| < E$ for all $0 \leq h \leq h_0$ and all $n \in \mathbb{N}$ with $nh \leq t$. Therefore, $\hat{x}^n = \check{x}^n$ for these h, n and

$$\|x(nh) - \hat{x}^n\| \leq \frac{h}{2} \max_{\zeta\in[0,t]} \|x''(\zeta)\| \frac{e^{Lt}-1}{L}$$

□

Proof of Lemma 4. Write $d(x,y) \stackrel{\text{def}}{=} \|x-y\|$ for $x, y \in \mathbb{R}^n$. $d(\cdot, S)$ satisfies the triangle inequality $d(x,S) = \inf_{z\in S} d(x,z) \leq \inf_{z\in S}(d(x,y) + d(y,z)) = d(x,y) + d(y,S)$. For $\varepsilon > 0$ and x, y with $d(x,y) < \delta := \varepsilon$ we, thus, know $d(x,S) - d(y,S) \leq d(x,y) < \varepsilon$. Also, $d(y,S) - d(x,S) \leq d(y,x) = d(x,y) < \varepsilon$. □

Proof of Lemma 5. Suppose $\inf_{x\in K} d(x, F^{\complement}) = 0$. Then there is a sequence $(x_n)_{n\in\mathbb{N}} \subseteq K$ with $d(x_n, F^{\complement}) \to 0$ as $n \to \infty$. By compactness of K, we can pass to a subsequence x_{n_k} such that $x_{n_k} \to x$ converges to an $x \in K$ as $k \to \infty$. By Lemma 4,

$$d(\lim_{k\to\infty} x_{n_k}, F^{\complement}) = \lim_{k\to\infty} d(x_{n_k}, F^{\complement}) = \lim_{n\to\infty} d(x_n, F^{\complement}) = 0$$

Now $x \in K \subseteq F$ implies $x \notin F^{\complement}$. Since $d(x, F^{\complement}) = \inf_{y\in F^{\complement}} d(x,y) = 0$, there is a sequence in $F^{\complement} \setminus \{x\}$ converging to x. Yet, $F^{\complement}$ is closed, hence $x \in F^{\complement}$, contradicting $x \in F$. □

B Continuous Completeness Proof

We will first prove the soundness direction of Theorem 1. Then it remains to prove the completeness direction of Theorem 1. In this appendix, we present a fully constructive proof of Theorem 1, following our proof structure from [Pla08]. Thanks to our significantly simplified axiomatization, the soundness and relative completeness proofs are much simplified. The relative completeness proof shows that for every valid $\mathsf{d}\mathcal{L}$ formula, there is a finite set of valid FOD formulas from which it can be derived in the $\mathsf{d}\mathcal{L}$ calculus.

Proof Outline. The (constructive) proof, which, in full, is contained in the remainder of this appendix, adapts the techniques of Cook [Coo78] and Harel [HMP77, HKT00] to the hybrid case. The decisive step is to show that every valid property of a repetition α^* can be proven by axioms I or C, respectively, with a sufficiently strong invariant or variant that is expressible in $\mathsf{d}\mathcal{L}$. For this, we show that $\mathsf{d}\mathcal{L}$ formulas can be expressed equivalently in FOD, and that valid $\mathsf{d}\mathcal{L}$ formulas can be derived from corresponding FOD axioms in the $\mathsf{d}\mathcal{L}$ calculus. In turn, the crucial step is to construct a finite FOD formula that characterizes the effect of unboundedly many repetitive hybrid transitions and just uses finitely many real variables. □

Natural numbers are definable in FOD [Pla08, Theorem 2]. For the sake of a complete presentation, we recall our proof.

Theorem 11 (Incompleteness)**.** *Both the discrete fragment and the continuous fragment of* $\mathsf{d}\mathcal{L}$ *are* not effectively axiomatisable*, i.e., they have no sound and complete effective calculus, because natural numbers are definable in both fragments.*

Proof. We prove that natural numbers are definable among the real numbers of $\mathsf{d}\mathcal{L}$ interpretations in both fragments. Then these fragments extend first-order *integer* arithmetic such that the incompleteness theorem of Gödel [Göd31] applies. Gödel's incompleteness theorem shows that no logic extending first-order integer arithmetic can have a sound and complete effective calculus. Natural numbers are definable in the discrete fragment without continuous evolutions using repetitive additions:

$$\mathit{nat}(n) \leftrightarrow \langle x := 0; (x := x + 1)^* \rangle\, x = n.$$

In the continuous fragment, an isomorphic copy of the natural numbers is definable using linear differential equations:

$$\mathit{nat}(n) \leftrightarrow \exists s{=}0\, \exists c{=}1\, \exists \tau{=}0\, \langle s' = c, c' = -s, \tau' = 1\rangle (s = 0 \wedge \tau = n).$$

These differential equations characterise $\sin$ and $\cos$ as unique solutions for s and c, respectively. Their zeros, as detected by τ, correspond to an isomorphic copy of natural numbers, scaled by π, i.e., $\mathit{nat}(n)$ holds iff n is of the form $k\pi$ for a $k \in \mathbb{N}$; see Figure 6. The initial values for s and c prevent the trivial solution identical to 0. □

Let the FOD formula $\mathit{nat}(x)$ be true iff x is a natural number. In this section, we abbreviate quantifiers over natural numbers by $\forall x : \mathbb{N}\ \phi$ and $\exists x : \mathbb{N}\ \phi$ for $\forall x\, (\mathit{nat}(x) \rightarrow \phi)$ and $\exists x\, (\mathit{nat}(x) \wedge \phi)$. Likewise, we abbreviate quantifiers over integers, e.g., by $\forall x : \mathbb{Z}\ \phi$.

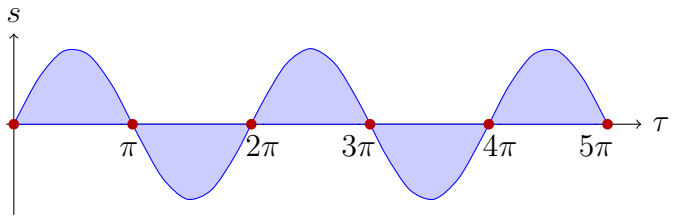

Figure 6: Characterisation of $\mathbb{N}$ as zeros of solutions of differential equations

B.1 Soundness of $\mathsf{d}\mathcal{L}$ Calculus

Before we turn to prove completeness, we first prove the soundness direction of Theorem 1. We state soundness as a separate theorem, because it is of independent interest:

Theorem 12 (Soundness of $\mathsf{d}\mathcal{L}$)**.** *The $\mathsf{d}\mathcal{L}$ calculus is sound, i.e., every provable formula is valid, i.e., true in all states.*

Proof. All axioms of the $\mathsf{d}\mathcal{L}$ calculus in Figure 1 are sound, i.e., all their instances valid.

[:=] Axiom [:=] is sound. For state ν, let ω be the unique state such that $(\nu,\omega) \in \rho(x := \theta)$. That is, $\omega = \nu$ except $\omega(x) = [\![\theta]\!]_\nu$. By the Substitution Lemma [Pla10b, Lemma 2.2] for admissible substitutions, $\omega \models \phi$ iff $\nu \models \phi^\theta_x$. Thus, $\nu \models [x := \theta]\phi$ iff $\nu \models \phi^\theta_x$.

[?] Axiom [?] is sound. Consider a state ν. If $\nu \models \chi$, then the only transition is $(\nu,\nu) \in \rho(?\chi)$, hence, $\nu \models [?\chi]\phi$ iff $\nu \models \phi$, which holds iff $\nu \models \chi \to \phi$. If, otherwise, $\nu \not\models \chi$, then $?\chi$ allows no transitions $(\nu,\omega) \in \rho(?\chi)$ hence $\nu \models [?\chi]\phi$ holds vacuously and $\nu \models \chi \to \phi$ holds vacuously, too.

[′] Axiom [′] is sound, because y is the solution (unique by Picard-Lindelöf [Wal98, Theorem 10.VI]) of the differential equation $y(t)' = \theta$ with symbolic initial values $y(0) = x$. Thus, $\nu \models [x' = \theta]\phi$ iff ϕ holds at all times $t \geq 0$ along $y(t)$. That is, $\nu \models [x' = \theta]\phi$ iff $\nu \models \forall t{\geq}0\, [x := y(t)]\phi$.

[&] Axiom [&] is sound, because the right-hand side checks χ along the reverse flow. Continuous evolution is reversible, i.e., the transitions of $x' = -\theta$ are inverse to those of $x' = \theta$. For this, consider $(\nu,\omega) \in \rho(x' = \theta)$, that is, let φ be the unique [Pla08, Lemma 1] solution of $x' = \theta$ of some duration r starting in state ν and ending in ω. Then ϱ defined as $\varrho(\zeta) = \varphi(r - \zeta)$, is of duration r, starts in ω and ends in ν. Furthermore, ϱ is a solution of $x' = -\theta$:

$$\begin{aligned}\frac{\mathsf{d}\varrho(t)(x)}{\mathsf{d}t}(\zeta) &= \frac{\mathsf{d}\varphi(r-t)(x)}{\mathsf{d}t}(\zeta) = \frac{\mathsf{d}\varphi(u)(x)}{\mathsf{d}u}\frac{\mathsf{d}(r-t)}{\mathsf{d}t}(\zeta)\\ &= -\frac{\mathsf{d}\varphi(u)(x)}{\mathsf{d}u}(\zeta) = -[\![\theta]\!]_{\varphi(\zeta)} = [\![-\theta]\!]_{\varphi(\zeta)}.\end{aligned}$$

Consequently, all evolutions of $[x' = -\theta]$ follow the same flow as $[x' = \theta]$, but backwards. The antecedent of the postcondition tests whether, along the reverse flow, χ has been true at all times until the starting time t_0; see Figure 2. The quantifier $\forall t_0 = x_0 \ldots$, which is

an abbreviation for $\forall t_0\,(t_0 = x_0 \rightarrow \dots)$, remembers the initial time x_0 in t_0. Recall that we assume x_0 to be a clock with the differential equation $x_0' = 1$ in the vectorial differential equation $x' = \theta$ to track time.

[$\cup$] Axiom [$\cup$] is sound. Since $\rho(\alpha \cup \beta) = \rho(\alpha) \cup \rho(\beta)$, we have that $(\nu, \omega) \in \rho(\alpha \cup \beta)$ iff $(\nu, \omega) \in \rho(\alpha)$ or $(\nu, \omega) \in \rho(\beta)$. Thus, $\nu \models [\alpha \cup \beta]\phi$ iff $\nu \models [\alpha]\phi$ and $\nu \models [\beta]\phi$.

[;] Axiom [;] is sound. Since $\rho(\alpha; \beta) = \rho(\beta) \circ \rho(\alpha)$, we have that $(\nu, \omega) \in \rho(\alpha; \beta)$ iff $(\nu, \mu) \in \rho(\alpha)$ and $(\mu, \omega) \in \rho(\beta)$ for some middle state μ. Hence, $\nu \models [\alpha; \beta]\phi$ iff $\mu \models [\beta]\phi$ for all μ with $(\nu, \mu) \in \rho(\alpha)$. That is $\nu \models [\alpha; \beta]\phi$ iff $\nu \models [\alpha][\beta]\phi$.

[^{*n}] Axiom [^{*n}] is sound. Since $\rho(\alpha^*) = \bigcup_{n \in \mathbb{N}} \rho(\alpha^n)$, there are two cases: α either repeats for 0 or for 1 or more iterations. Thus $\nu \models [\alpha^*]\phi$ iff $\nu \models [\alpha^0]\phi$ and $\nu \models [\alpha; \alpha^*]\phi$. Thus, by the soundness of [;], $\nu \models [\alpha^*]\phi$ iff $\nu \models \phi$ and $\nu \models [\alpha][\alpha^*]\phi$.

K Let $\nu \models [\alpha](\phi \rightarrow \psi)$ and $\nu \models [\alpha]\phi$. Consider any ω with $(\nu, \omega) \in \rho(\alpha)$. Then, $\omega \models \phi \rightarrow \psi$ and $\omega \models \phi$. Thus, $\omega \models \psi$, implying $\nu \models [\alpha]\psi$, since ω was arbitrary with $(\nu, \omega) \in \rho(\alpha)$.

I Let $\nu \models [\alpha^*](\phi \rightarrow [\alpha]\phi)$ and $\nu \models \phi$. Since $\rho(\alpha^*) = \bigcup_{n \in \mathbb{N}} \rho(\alpha^n)$, it is enough to show that $\nu \models [\alpha^n]\phi$ for all $n \in \mathbb{N}$. For $n = 0$, this follows from $\nu \models \phi$. Inductively, from $\nu \models [\alpha^n]\phi$, we show that $\nu \models [\alpha^{n+1}]\phi$. By soundness of [;], it is enough to show $\nu \models [\alpha^n][\alpha]\phi$. For any ω with $(\nu, \omega) \in \rho(\alpha^n)$, we know $\omega \models \phi$ and need to show $\omega \models [\alpha]\phi$. Yet, we also know $\omega \models \phi \rightarrow [\alpha]\phi$ by $\nu \models [\alpha^*](\phi \rightarrow [\alpha]\phi)$, because $\rho(\alpha^n) \subseteq \rho(\alpha^*)$.

C Let $\nu \models [\alpha^*]\forall v{>}0\,(\varphi(v) \rightarrow \langle\alpha\rangle\varphi(v-1))$ and $\nu \models \exists v\,\varphi(v)$. First note that v does not occur in α, hence its value does not change during α^* and does not affect the runs of α^*. We show $\nu \models \langle\alpha^*\rangle\exists v{\leq}0\,\varphi(v)$ by a well-founded induction along states ω with $(\nu, \omega) \in \rho(\alpha^*)$ satisfying $\omega \models \varphi(v)$ for some value of v. If $\omega \models \varphi(v)$ for a value of $v \leq 0$, we have $\omega \models \exists v{\leq}0\,\varphi(v)$, which implies $\nu \models \langle\alpha^*\rangle\exists v{\leq}0\,\varphi(v)$ by $(\nu, \omega) \in \rho(\alpha^*)$. Otherwise, if $\omega \models \varphi(v)$ for a value of $v > 0$, then by antecedent, we know $\omega \models v > 0 \wedge \varphi(v) \rightarrow \langle\alpha\rangle\varphi(v-1)$, because $(\nu, \omega) \in \rho(\alpha^*)$. Thus, $\omega \models \langle\alpha\rangle\varphi(v-1)$. Thus, there is a ω_1 with $(\omega, \omega_1) \in \rho(\alpha)$ such that $\omega_1 \models \varphi(v-1)$. The induction is, thus, well-founded, because the value of v decreases at least by 1, which it can only do finitely often down to the base case $v \leq 0$.

B Contrapositively, let $\nu \not\models [\alpha]\forall x\,\phi$. Thus, there is a state ω with $(\nu, \omega) \in \rho(\alpha)$ such that $\omega \not\models \forall x\,\phi$, because $\omega_x^d \not\models \phi$ where ω_x^d is like ω except for the value of x, which is $d \in \mathbb{R}$ in ω_x^d. Since B assumes x not to occur in α, its value does not change during α and does not affect runs of α. Thus, for the state ν_x^d that is like ν except for the value of x, which is d in ν_x^d, we have that $(\nu_x^d, \omega_x^d) \in \rho(\alpha)$. Hence, $\omega_x^d \not\models \phi$ implies $\nu_x^d \not\models [\alpha]\phi$, i.e., $\nu \not\models \forall x\,[\alpha]\phi$.

V Let ν with $\nu \models \phi$. Consider any ω with $(\nu, \omega) \in \rho(\alpha)$. Since V assumes α not to bind any variable that is free in ϕ, the free variables of ϕ cannot change their value when passing from ν to ω, hence $\nu \models \phi$ iff $\omega \models \phi$ by Coincidence Lemma 9.

G Rule G is (globally) sound, which we show by induction on the structure of the proof. The $\mathsf{d}\mathcal{L}$ axioms (and basic axioms of first-order logic and first-order real arithmetic) are sound,

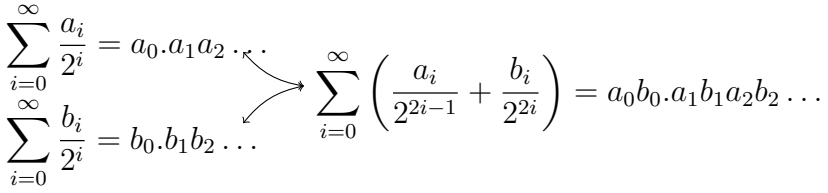

Figure 7: Fractional encoding principle of $\mathbb{R}$-Gödel encoding by bit interleaving

hence, the proof can only start from valid formulas. Let ϕ be provable, and let $[\alpha]\phi$ result from ϕ by application of G. The proof of ϕ has one step less than that of $[\alpha]\phi$, hence, by induction hypothesis, the proof of ϕ is sound, which means that ϕ is valid ($\vDash \phi$). That is, ϕ is true in all states ν, which implies that, in particular, ϕ is true ($\nu \models \phi$) in all states ω for which $(\nu, \omega) \in \rho(\alpha)$. Thus, $[\alpha]\phi$ is valid and its proof sound.

□

Soundness of the rules and axioms of the first-order Hilbert calculus are as usual. Modus ponens is obvious and $\forall$-generalization follows the pattern of G.

Next, we can turn to proving relative completeness.

B.2 Characterizing Real Gödel Encodings

As the central device for constructing a FOD formula that captures the effect of unboundedly many repetitive hybrid transitions and just uses finitely many real variables, we prove that a real version of Gödel encoding is definable in FOD. That is, we give a FOD formula that reversibly packs finite sequences of real values into a single real number. The standard prime power constructions for natural number pairings do not generalize to the reals, because factorization is not unique.

Observe that a single differential equation system is *not* sufficient for defining real pairing functions as their solutions are differentiable, and yet, as a consequence of Morayne's theorem [Mor87], there is no differentiable surjection $\mathbb{R} \to \mathbb{R}^2$, nor to any part of $\mathbb{R}^2$ of positive measure. We show that real sequences can be encoded nevertheless by chaining the effects of solutions of multiple (but finitely many!) differential equations and quantifiers.

Lemma 13 ($\mathbb{R}$-Gödel encoding)**.** *The formula at(Z, n, j, z), which holds iff Z is a real number that represents a Gödel encoding of a sequence of n real numbers with real value z at position j (for $1 \leq j \leq m$), is definable in FOD. For a formula $\phi(z)$ we abbreviate $\exists z\,(\text{at}(Z, n, j, z) \wedge \phi(z))$ by $\phi(Z_j^{(n)})$.*

Proof. The basic idea of the $\mathbb{R}$-Gödel encoding is to interleave the bits of real numbers as depicted in Figure 7 (for a pairing of $n = 2$ numbers a and b). For defining at(Z, n, j, z), we use several

$$
\begin{aligned}
\mathrm{at}(Z,n,j,z) &\leftrightarrow \forall i:\mathbb{Z}\ \mathrm{digit}(z,i) = \mathrm{digit}(Z,n(i-1)+j) \wedge \mathit{nat}(n) \wedge \mathit{nat}(j) \wedge n>0\\
\mathrm{digit}(a,i) &= \mathrm{intpart}(2\,\mathrm{frac}(2^{i-1}a))\\
\mathrm{intpart}(a) &= a - \mathrm{frac}(a)\\
\mathrm{frac}(a)=z &\leftrightarrow \exists i:\mathbb{Z}\ z = a-i \wedge -1<z\wedge z<1 \wedge az\geq 0\\
2^i=z &\leftrightarrow i\geq 0 \wedge \exists x\,\exists t\,(x=1\wedge t=0\wedge \langle x'=x\ln 2, t'=1\rangle(t=i\wedge x=z))\\
&\vee\ i<0\wedge\exists x\,\exists t\,(x=1\wedge t=0\wedge\langle x'=-x\ln 2,t'=-1\rangle(t=i\wedge x=z))\\
\ln 2=z &\leftrightarrow \exists x\,\exists t\,(x=1\wedge t=0\wedge\langle x'=x,t'=1\rangle(x=2\wedge t=z))
\end{aligned}
$$

Figure 8: FOD definition characterizing Gödel encoding of $\mathbb{R}$-sequences in one real number

auxiliary functions to improve readability; see Figure 8. Note that these definitions need no recursion. Hence, as in the notation $\phi(Z_j^{(n)})$, we can consider occurrences of the function symbols as syntactic abbreviations for quantified variables satisfying the respective definitions.

The function symbol $\mathrm{digit}(a,i)$ gives the ith bit of $a\in\mathbb{R}$ when represented with basis 2. For $i>0$, $\mathrm{digit}(a,i)$ yields fractional bits, and, for $i\leq 0$, it yields bits of the integer part. For instance, $\mathrm{digit}(a,1)$ yields the first fractional bit, $\mathrm{digit}(a,0)$ is the least-significant bit of the integer part of a. The function $\mathrm{intpart}(a)$ represents the integer part of $a\in\mathbb{R}$. The function $\mathrm{frac}(a)$ represents the fractional part of $a\in\mathbb{R}$, which drops all integer bits. The last constraint in its definition implies that $\mathrm{frac}(a)$ keeps the sign of a (or 0). Consequently, $\mathrm{intpart}(a)$ and $\mathrm{digit}(a,i)$ also keep the sign of a (or 0). Exponentiation 2^i is definable using differential equations, using an auxiliary characterization of the natural logarithm $\ln 2$. The definition of 2^i splits into the case of exponential growth when $i\geq 0$ and a symmetric case of exponential decay when $i<0$. □

B.3 Expressibility and Rendition of Hybrid Program Semantics

$$
\begin{aligned}
\mathcal{S}_{x_i:=\theta}(\vec{x},\vec{v}) &\equiv v_i=\theta\wedge\bigwedge_{j\neq i} v_j=x_j\\
\mathcal{S}_{x'=\theta}(\vec{x},\vec{v}) &\equiv \langle x'=\theta\rangle\vec{v}=\vec{x}\\
\mathcal{S}_{x'=\theta\,\&\,\chi}(\vec{x},\vec{v}) &\equiv \exists t\,\big(t=0\wedge\langle x'=\theta,t'=1\rangle\big(\vec{v}=\vec{x}\wedge[x'=-\theta,t'=-1](t\geq 0\to\chi)\big)\big)\\
\mathcal{S}_{?\chi}(\vec{x},\vec{v}) &\equiv \vec{v}=\vec{x}\wedge\chi\\
\mathcal{S}_{\beta\cup\gamma}(\vec{x},\vec{v}) &\equiv \mathcal{S}_\beta(\vec{x},\vec{v})\vee\mathcal{S}_\gamma(\vec{x},\vec{v})\\
\mathcal{S}_{\beta;\gamma}(\vec{x},\vec{v}) &\equiv \exists\vec{z}\,(\mathcal{S}_\beta(\vec{x},\vec{z})\wedge\mathcal{S}_\gamma(\vec{z},\vec{v}))\\
\mathcal{S}_{\beta^*}(\vec{x},\vec{v}) &\equiv \exists Z\,\exists n:\mathbb{N}\,\big(Z_1^{(n)}=\vec{x}\wedge Z_n^{(n)}=\vec{v}\wedge\forall i:\mathbb{N}\,(1\leq i<n\to\mathcal{S}_\beta(Z_i^{(n)},Z_{i+1}^{(n)}))\big)
\end{aligned}
$$

Figure 9: Explicit rendition of hybrid program transition semantics in FOD

In order to show that $\mathsf{d}\mathcal{L}$ is sufficiently expressive to state the invariants and variants that are

needed for proving valid statements about loops with axioms I and C, we prove an expressibility result. We give a constructive proof that the state transition relation of hybrid programs is definable in FOD, i.e., there is a FOD formula $\mathcal{S}_\alpha(\vec{x},\vec{v})$ characterizing the state transitions of hybrid program α from the state characterized by the vector $\vec{x}$ of variables to the state characterized by vector $\vec{v}$.

For this, we need to characterize hybrid programs equivalently by differential equations in FOD. Observe that the existence of such characterizations does *not* follow from results embedding Turing machines into differential equations [Bra95,GCB07], because, unlike Turing machines, hybrid programs are not restricted to discrete values on a grid (such as $\mathbb{N}^k$) but work with continuous real values. Furthermore, Turing machines only have repetitions of discrete transitions on discrete data (e.g., $\mathbb{N}$). For hybrid programs, in contrast, we have to characterize repetitive interactions of interacting discrete and continuous transitions in continuous space (some $\mathbb{R}^k$).

Lemma 14 (Hybrid program rendition)**.** *For every hybrid program α with variables among $\vec{x} = x_1,\dots,x_k$, there is a FOD formula $\mathcal{S}_\alpha(\vec{x},\vec{v})$ with variables among the $2k$ distinct variables $\vec{x} = x_1,\dots,x_k$ and $\vec{v} = v_1,\dots,v_k$ such that*

$$\vDash \mathcal{S}_\alpha(\vec{x},\vec{v}) \leftrightarrow \langle\alpha\rangle\vec{x} = \vec{v}$$

Proof. By the Coincidence Lemma 9, interpretations of the vectors $\vec{x}$ and $\vec{v}$ completely characterize the input and output states, respectively, as far as α is concerned. These vectors are finite because α is finite. Vectorial equalities like $\vec{x} = \vec{v}$ or quantifiers $\exists\vec{v}$ are to be understood componentwise. The program rendition is defined inductively in Figure 9.

The (vectorial) differential equation case $x' = \theta$ (we avoid the notation $\vec{x}' = \theta$) gives FOD formulas; no further reduction is needed. Evolution along differential equations with evolution domain restrictions is definable in terms of differential equations by the soundness of axiom [&]. Formula $\mathcal{S}_{x'=\theta\,\&\,\chi}(\vec{x},\vec{v})$ is obtained by duality from the right-hand side of axiom [&]. We add a clock t to x explicitly. Unlike all other cases, this case in Figure 9 uses nested FOD modalities, which can be avoided altogether when using the following equivalent FOD formula instead (cf. Figure 2 on p. 5):

$$\begin{aligned}&\exists t\,\exists r\,\big(t = 0 \wedge \langle x' = \theta, t' = 1\rangle(\vec{v} = \vec{x} \wedge r = t) \wedge\\ &\quad \forall\vec{x}\,\forall t\,(\vec{x} = \vec{v} \wedge t = r \rightarrow [x' = -\theta, t' = -1](t \geq 0 \rightarrow \chi))\big).\end{aligned}$$

With a finite formula, the characterization of repetition $\mathcal{S}_{\beta^*}(\vec{x},\vec{v})$ in FOD needs to capture arbitrarily long sequences of intermediate real-valued states and the correct transition between successive states of such a sequence. To achieve this with first-order quantifiers, we use the real Gödel encoding from Lemma 13 in Figure 9 to map unbounded sequences of real-valued states reversibly to a single real number Z, which can be quantified over in first-order logic. □

Using the program rendition from Lemma 14 to characterize modalities, we prove that every dL formula can be expressed equivalently in FOD.

Lemma 15 (dL expressibility)**.** *Logic* dL *is* expressible *in FOD: for each* dL *formula ϕ there is a FOD formula $\phi^\flat$ that is equivalent, i.e., $\vDash \phi \leftrightarrow \phi^\flat$. The converse holds trivially.*

Proof. The proof follows an induction on the structure of formula ϕ for which it is imperative to find an equivalent $\phi^\flat$ in FOD. Observe that the construction of $\phi^\flat$ from ϕ is effective.

0) If ϕ is a first-order formula, then $\phi^\flat := \phi$ already is a FOD formula such that nothing has to be shown.

1. If ϕ is of the form $\varphi \vee \psi$, then by the induction hypothesis there are FOD formulas $\varphi^\flat, \psi^\flat$ such that $\vDash \varphi \leftrightarrow \varphi^\flat$ and $\vDash \psi \leftrightarrow \psi^\flat$, from which we can conclude by congruence that

$$\vDash (\varphi \vee \psi) \leftrightarrow (\varphi^\flat \vee \psi^\flat)$$

giving $\vDash \phi \leftrightarrow \phi^\flat$ by choosing $\varphi^\flat \vee \psi^\flat$ for $\phi^\flat$. Similar reasoning addresses the other propositional connectives or quantifiers by congruence.

2. The case where ϕ is of the form $\langle\alpha\rangle\psi$ is a consequence of the characterization of the semantics of hybrid programs in FOD. Expressibility follows from the induction hypothesis using the equivalence of explicit hybrid program renditions from Lemma 14:

$$\vDash \langle\alpha\rangle\psi \leftrightarrow \exists\vec{v}\,(\mathcal{S}_\alpha(\vec{x},\vec{v}) \wedge \psi^{\flat\vec{v}}_{\vec{x}}).$$

3. The case where ϕ is $[\alpha]\psi$ is again a consequence of Lemma 14:

$$\vDash [\alpha]\psi \leftrightarrow \forall\vec{v}\,(\mathcal{S}_\alpha(\vec{x},\vec{v}) \rightarrow \psi^{\flat\vec{v}}_{\vec{x}})$$

□

Observe that the construction of $\phi^\flat$ out of ϕ is effective.

B.4 First-Order Continuous Relative Completeness

As special cases of Theorem 1, we first prove relative completeness for first-order assertions about hybrid programs. These first-order cases constitute the basis for the general completeness proof for arbitrary formulas of $\mathsf{d}\mathcal{L}$. We use the notation $\vdash_{\mathcal{D}} \phi$ to indicate that a $\mathsf{d}\mathcal{L}$ formula ϕ is derivable in the $\mathsf{d}\mathcal{L}$ calculus (Figure 1) from FOD tautologies. The following formula derives[7] from K by duality

$$(\mathrm{K}_{\langle\rangle})\quad [\alpha](\phi \rightarrow \psi) \rightarrow (\langle\alpha\rangle\phi \rightarrow \langle\alpha\rangle\psi)$$

Proposition 16 (Relative completeness of first-order safety)**.** *For every hybrid program α and all FOD formulas F, G*

$$\vDash F \rightarrow [\alpha]G \textit{ implies } \vdash_{\mathcal{D}} F \rightarrow [\alpha]G.$$

[7] $[\alpha](\neg\psi \rightarrow \neg\phi) \rightarrow ([\alpha]\neg\psi \rightarrow [\alpha]\neg\phi)$ by K. Thus, propositionally, $[\alpha](\neg\psi \rightarrow \neg\phi) \rightarrow (\neg[\alpha]\neg\phi \rightarrow \neg[\alpha]\neg\psi)$. By duality $\langle\alpha\rangle\phi \equiv \neg[\alpha]\neg\phi$, this is $[\alpha](\neg\psi \rightarrow \neg\phi) \rightarrow (\langle\alpha\rangle\phi \rightarrow \langle\alpha\rangle\psi)$. Thus, $[\alpha](\phi \rightarrow \psi) \rightarrow (\langle\alpha\rangle\phi \rightarrow \langle\alpha\rangle\psi)$ derives as follows. From the propositional tautology $(\phi \rightarrow \psi) \rightarrow (\neg\psi \rightarrow \neg\phi)$ we derive $[\alpha]((\phi \rightarrow \psi) \rightarrow (\neg\psi \rightarrow \neg\phi))$ with G, from which K derives $[\alpha](\phi \rightarrow \psi) \rightarrow [\alpha](\neg\psi \rightarrow \neg\phi)$, from which propositional reasoning yields the result.

Proof. We generalize the relative completeness proof by Cook [Coo78] and Harel et al. [HMP77] to $\mathsf{d}\mathcal{L}$ and follow an induction on the structure of program α. In the following, *IH* is short for the induction hypothesis.

1. The cases where α is of the form $x := \theta$, $?\chi$, $\beta \cup \gamma$, or $\beta;\gamma$ are consequences of the soundness of the equivalence rules [;],[∪],[?],[:=]. Whenever their respective left-hand side is valid, their right-hand side is valid and of smaller complexity (the programs get simpler), and hence derivable by IH. Thus, we can derive $F \to [\alpha]G$ by applying the respective rule. We explicitly show the proof for $\beta;\gamma$ as it contains an extra twist.

2. $\vDash F \to [\beta;\gamma]G$, which implies $\vDash F \to [\beta][\gamma]G$. By Lemma 15, there is a FOD formula $G^\flat$ such that $\vDash G^\flat \leftrightarrow [\gamma]G$. From that validity we conclude by IH that $\vdash_{\mathcal{D}} F \to [\beta]G^\flat$ is derivable. Similarly, due to $\vDash G^\flat \to [\gamma]G$, we conclude $\vdash_{\mathcal{D}} G^\flat \to [\gamma]G$ by IH. Extending the latter by G, we derive $\vdash_{\mathcal{D}} [\beta](G^\flat \to [\gamma]G)$. Thus, K derives $\vdash_{\mathcal{D}} [\beta]G^\flat \to [\beta][\gamma]G$. Combining the above derivations propositionally (cut with $[\beta]G^\flat$), we derive $\vdash_{\mathcal{D}} F \to [\beta][\gamma]G$, from which [;] derives $\vdash_{\mathcal{D}} F \to [\beta;\gamma]G$.

3. $\vDash F \to [x' = \theta]G$ is a FOD formula and hence provable by assumption.

4. $\vDash F \to [x' = \theta \,\&\, \chi]G$, then this formula is, by axiom [&], provably equivalent to a formula without evolution domain restrictions. This is definable in FOD by Lemma 14, which we use as an abbreviation in FOD. Later, in the proof of Theorem 1, axiom [&] also directly gives a provably equivalent but structurally simpler formula, which is, thus, provable by induction hypothesis. That part is like the case for [&] in the proof of Theorem 8.

5. $\vDash F \to [\beta^*]G$ can be derived by induction as follows. Formula $[\beta^*]G$, which expresses that all post-states of β^* satisfy G, is an invariant of β^* by ind, because $[\beta^*]G \to [\beta][\beta^*]G$ is valid, even provable by $[^{*n}]$. Thus, its equivalent FOD encoding according to Lemma 15 is an invariant:
$$\phi \equiv ([\beta^*]G)^\flat \equiv \forall \vec{v}\,(\mathcal{S}_{\beta^*}(\vec{x},\vec{v}) \to G^{\vec{v}}_{\vec{x}}).$$
$F \to \phi$ and $\phi \to G$ are valid FOD formulas, hence derivable by assumption. By G the latter derivation extends to $\vdash_{\mathcal{D}} [\beta^*](\phi \to G)$, from which K derives $\vdash_{\mathcal{D}} [\beta^*]\phi \to [\beta^*]G$. As above, $\phi \to [\beta]\phi$ is valid by the semantics of repetition, and thus derivable by IH, since β is less complex. Thus, G derives $\vdash_{\mathcal{D}} [\beta^*](\phi \to [\beta]\phi)$, from which I derives $\vdash_{\mathcal{D}} \phi \to [\beta^*]\phi$. The above derivations combine propositionally (cut with $[\beta^*]\phi$ and ϕ) to $\vdash_{\mathcal{D}} F \to [\beta^*]G$. □

Proposition 17 (Relative completeness of first-order liveness)**.** *For each hybrid program α and all FOD formulas F, G*
$$\vDash F \to \langle\alpha\rangle G \textit{ implies } \vdash_{\mathcal{D}} F \to \langle\alpha\rangle G.$$

Proof. Most cases of this proof follow directly from the equivalence axioms used in Proposition 16, just by the duality $\langle\alpha\rangle G \equiv \neg[\alpha]\neg G$. What is different is that axiom I for repetitions is no equivalence, and thus, does not give dual arguments. We generalize the arithmetic completeness proof by Harel [HMP77] to the hybrid case. Assume that $\vDash F \to \langle\beta^*\rangle G$. To derive this formula by C, we

use a FOD formula $\varphi(n)$ as a variant expressing that, after n iterations, β can lead to a state satisfying G. This formula is obtained from Lemmas 14 and 15 as $(\langle\beta^*\rangle G)^\flat \equiv \exists\vec{v}\,(\mathcal{S}_{\beta^*}(\vec{x},\vec{v}) \wedge G^{\vec{v}}_{\vec{x}})$, *except* that the quantifier on the repetition count n is removed such that n becomes a free variable (plus index shifting to count repetitions). We define $\varphi(n-1)$ to be

$$\exists\vec{v}\,\exists Z\,\big(G^{\vec{v}}_{\vec{x}} \wedge Z^{(n)}_1 = \vec{x} \wedge Z^{(n)}_n = \vec{v} \\ \wedge \forall i:\mathbb{N}\,(1 \leq i < n \rightarrow \mathcal{S}_\beta(Z^{(n)}_i, Z^{(n)}_{i+1}))\big).$$

By Lemma 13, $\varphi(n)$ can only hold true if n is a natural number. Now $\vDash \varphi(n) \wedge n > 0 \rightarrow \langle\beta\rangle\varphi(n-1)$ is valid by construction according to the loop semantics: If $n > 0$ is a natural number then so is $n-1$, and if β can reach G after n repetitions, then, after executing β once, $n-1$ repetitions of β can reach G. By IH, this formula is derivable, since β contains less loops. From this, $\forall$-generalization and G derive $\vdash_{\mathcal{D}} [\beta^*]\forall n{>}0\,(\varphi(n) \rightarrow \langle\beta\rangle\varphi(n-1))$. Thus, C derives

$$\vdash_{\mathcal{D}} \forall v\,(\varphi(v) \rightarrow \langle\beta^*\rangle\exists v{\leq}0\,\varphi(v))$$

Standard first-order reasoning extends the latter to $\vdash_{\mathcal{D}} \exists v\,\varphi(v) \rightarrow \langle\beta^*\rangle\exists v{\leq}0\,\varphi(v)$. It only remains to show that the antecedent is derivable from F and $\langle\beta^*\rangle G$ is derivable from the succedent. The following formulas are valid FOD formulas, hence derivable by assumption:

- $\vDash F \rightarrow \exists v\,\varphi(v)$, because $\vDash F \rightarrow \langle\beta^*\rangle G$, and
- $\vDash (\exists v{\leq}0\,\varphi(v)) \rightarrow G$, because $v \leq 0$, and the fact, that by Lemma 13, $\varphi(v)$ only holds true for natural numbers, imply $\varphi(0)$. Further, $\varphi(0)$ entails G, because zero repetitions of β have no effect.

By G, the latter extends to $\vdash_{\mathcal{D}} [\beta^*](\exists v{\leq}0\,\varphi(v) \rightarrow G)$. From this, the dual ($\mathrm{K}_{\langle\rangle}$) of K directly derives $\vdash_{\mathcal{D}} \langle\beta^*\rangle\exists v{\leq}0\,\varphi(v) \rightarrow \langle\beta^*\rangle G$. The above derivations combine propositionally to

$$\vdash_{\mathcal{D}} F \rightarrow \langle\beta^*\rangle G$$

(by a cut with $\langle\beta^*\rangle\exists v{\leq}0\,\varphi(v)$ and with $\exists v\,\varphi(v)$). □

B.5 Continuous Relative Completeness of dL

Having succeeded with the proofs of the above results we can finish the proof of Theorem 1.

Proof of Theorem 1. The proof follows a basic structure analogous to that of Harel et al.'s proof for the discrete case [HMP77]. We have to show that every valid dL formula ϕ can be proven from FOD axioms within the dL calculus: from $\vDash \phi$ we have to prove $\vdash_{\mathcal{D}} \phi$. The proof proceeds as follows: By propositional recombination, we inductively identify fragments of ϕ that correspond to $\phi_1 \rightarrow [\alpha]\phi_2$ or $\phi_1 \rightarrow \langle\alpha\rangle\phi_2$ logically. Next, we express subformulas ϕ_i equivalently in FOD by Lemma 15, and use Propositions 16 and 17 to resolve these first-order safety or liveness assertions. Finally, we prove that the original dL formula can be re-derived from the subproofs in the dL calculus.

We can assume ϕ to be given in conjunctive normal form by appropriate propositional reasoning. In particular, we assume that negations are pushed inside over modalities using the dualities $\neg[\alpha]\phi \equiv \langle\alpha\rangle\neg\phi$ and $\neg\langle\alpha\rangle\phi \equiv [\alpha]\neg\phi$. The remainder of the proof follows an induction on a measure $|\phi|$ defined as the number of modalities in ϕ. For a simple and uniform proof, we assume quantifiers to be abbreviations for modal formulas: $\exists x\, \phi \equiv \langle x' = 1\rangle\phi \vee \langle x' = -1\rangle\phi$ and $\forall x\, \phi \equiv [x' = 1]\phi \wedge [x' = -1]\phi$.

0) $|\phi| = 0$; then ϕ is a (quantifier-free) first-order formula; hence provable by assumption (even decidable [Tar51]).

1. ϕ is of the form $\neg\phi_1$; then ϕ_1 is first-order, as we assumed negations to be pushed inside. Hence, $|\phi| = 0$ and Case 0 applies.

2. ϕ is of the form $\phi_1 \wedge \phi_2$, then individually deduce simpler proofs for $\vdash_{\mathcal{D}} \phi_1$ and $\vdash_{\mathcal{D}} \phi_2$ by IH, which combine propositionally to a proof for $\vdash_{\mathcal{D}} \phi_1 \wedge \phi_2$.

3. ϕ is a disjunction and—without loss of generality—has one of the following forms (otherwise use associativity and commutativity to select a different order for the disjunction):

$$\phi_1 \vee [\alpha]\phi_2$$
$$\phi_1 \vee \langle\alpha\rangle\phi_2$$

   As a unified notation for those cases we use $\phi_1 \vee \langle\!\langle\alpha\rangle\!\rangle\phi_2$. Then, $|\phi_2| < |\phi|$, since ϕ_2 has less modalities. Likewise, $|\phi_1| < |\phi|$ because $\langle\!\langle\alpha\rangle\!\rangle\phi_2$ contributes one modality to $|\phi|$ that is not part of ϕ_1.

   According to Lemma 15 there are FOD formulas $\phi_1^\flat, \phi_2^\flat$ with $\vDash \phi_i \leftrightarrow \phi_i^\flat$ for $i = 1, 2$. By congruence, the validity $\vDash \phi$ yields $\vDash \phi_1^\flat \vee \langle\!\langle\alpha\rangle\!\rangle\phi_2^\flat$, which directly implies $\vDash \neg\phi_1^\flat \to \langle\!\langle\alpha\rangle\!\rangle\phi_2^\flat$. Then by Propositions 16 or 17, respectively, we derive

$$\vdash_{\mathcal{D}} \neg\phi_1^\flat \to \langle\!\langle\alpha\rangle\!\rangle\phi_2^\flat. \tag{6}$$

   Further $\vDash \phi_1 \leftrightarrow \phi_1^\flat$ implies $\vDash \neg\phi_1 \to \neg\phi_1^\flat$, which is derivable by IH, because $|\phi_1| < |\phi|$. We combine $\vdash_{\mathcal{D}} \neg\phi_1 \to \neg\phi_1^\flat$ with (6) (cut with $\neg\phi_1^\flat$) to

$$\vdash_{\mathcal{D}} \neg\phi_1 \to \langle\!\langle\alpha\rangle\!\rangle\phi_2^\flat. \tag{7}$$

   Likewise $\vDash \phi_2 \leftrightarrow \phi_2^\flat$ implies $\vDash \phi_2^\flat \to \phi_2$, which is derivable by IH, as $|\phi_2| < |\phi|$. From $\vdash_{\mathcal{D}} \phi_2^\flat \to \phi_2$ we derive $\vdash_{\mathcal{D}} [\alpha](\phi_2^\flat \to \phi_2)$ by G. Thus, by K or the dual ($\mathrm{K}_{\langle\rangle}$) of K, we derive[8] $\vdash_{\mathcal{D}} \langle\!\langle\alpha\rangle\!\rangle\phi_2^\flat \to \langle\!\langle\alpha\rangle\!\rangle\phi_2$. Finally we combine the latter derivation propositionally with (7) by a cut with $\langle\!\langle\alpha\rangle\!\rangle\phi_2^\flat$ to derive $\vdash_{\mathcal{D}} \neg\phi_1 \to \langle\!\langle\alpha\rangle\!\rangle\phi_2$, from which $\vdash_{\mathcal{D}} \phi_1 \vee \langle\!\langle\alpha\rangle\!\rangle\phi_2$ can be obtained propositionally to complete the proof.

4. The case where ϕ is of the form $[\alpha]\phi_2$ or $\langle\alpha\rangle\phi_2$ is included in case 3 with $\phi_1 \equiv \mathit{false}$.

This completes the proof of Theorem 1. □

[8] We consider quantifiers as abbreviations. Otherwise, we would use a derivable variant of Hilbert's $\forall$-generalization rule: From $\phi \to \psi$ conclude $\forall x\, \phi \to \forall x\, \psi$ (dually conclude $\exists x\, \phi \to \exists x\, \psi$).