Deductive Verification for Ordinary Differential Equations: Safety, Liveness, and Stability

Yong Kiam Tan

CMU-CS-22-114

June 2022

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
André Platzer, Chair
Jeremy Avigad
Stefan Mitsch
Frank Pfenning
Joël Ouaknine (MPI-SWS, Saarland University)

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

This research was sponsored by the National Science Foundation under NSF CAREER Award CNS-1054246 and Grant No. CNS-1739629, the AFOSR under grant number FA9550-16-1-0288, and the Alexander von Humboldt Foundation. The author was supported by A*STAR, Singapore. 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, the U.S. government or any other entity.

Keywords: deductive verification, differential dynamic logic, ordinary differential equations, switched systems, safety, invariance, liveness, existence, stability

Auspicium Melioris Aevi

Abstract

Ordinary differential equations (ODEs) are quintessential models of real-world continuous behavior in the physical and engineering sciences. They also feature prominently in *hybrid system* models that combine discrete and continuous dynamics, and interactions thereof. Formal verification of ODEs and hybrid systems is of increasing practical importance because the real-world systems they model, such as control systems and cyber-physical systems, are often required to operate in safety- and mission-critical settings—obtaining *comprehensive* and *trustworthy* verification results for continuous and hybrid systems gives a strong measure of confidence that the real-world systems they model operate correctly.

This thesis studies *deductive verification for ordinary differential equations* with a focus on proofs of their *i) safety*, *ii) liveness*, and *iii) stability* properties. These proofs are *compositionally* extended to obtain proofs of *iv) stability for hybrid (switched) systems*. The combination of safety, liveness, and stability is crucial for comprehensive correctness of real-world systems: *i) safety* of a system model ensures that it always stays within a prescribed set of safe states throughout its operation, *ii) liveness* ensures that the modeled system will eventually reach its specified goal or complete its mission, and *iii) & iv) stability* ensures that the idealized models are robust to real-world perturbations, which is important for control system designs.

The overarching thesis insight is the use of *deductive reasoning* as a basis for understanding the aforementioned properties and for developing their proofs. Specifically, this thesis uses *differential dynamic logic* (dL), a logic for deductive verification of hybrid systems, as a trustworthy logical foundation upon which all reasoning principles for safety, liveness, and stability are rigorously derived. The thesis first shows how ODE invariance, a key ingredient in proofs of ODE safety, can be completely axiomatized and reasoned about syntactically in dL. Then, ODE liveness and existence properties are formally proved through refinement-based reasoning in dL, where each refinement step is justified by proving an ODE safety property. Finally, stability properties for ODEs and hybrid systems are specified using dL's ability to nest safety and liveness modalities with first-order quantification. Proofs of those stability specifications build on ODE safety and liveness (sub-)proofs by compositionally adding dL reasoning for the first-order quantifiers and hybrid systems.

Formal dL specifications elucidate the logical relationships between the properties studied in this thesis. Indeed, these relationships are reflected in the thesis structure outlined above because they yield chapter-by-chapter identification, buildup, and generalization of the deductive building blocks underlying proof methods for the respective properties. The deductive approach enables such generalizations while retaining utmost confidence in the correctness of the resulting proofs because every step is soundly and syntactically justified using dL's parsimonious axiomatization. The derived proof principles and insights are put into practice by implementing them in the KeYmaera X theorem prover for hybrid systems based on dL.

Acknowledgments

I would like to start by thanking my family and friends for their support throughout my studies, especially during the COVID-19 pandemic. My undergraduate and graduate studies were generously funded by A*STAR, Singapore.

I was first introduced to formal methods research through an undergraduate internship working on the CakeML project at Cambridge, which was made possible by Mike Gordon and Alan Mycroft. My internship advisors, Ramana Kumar and Magnus Myreen, were instrumental in sparking my interest in verification research. I am also thankful for numerous past and ongoing collaboration opportunities with members of the wider CakeML community.

At CMU, I am very fortunate to be advised by André Platzer, who taught me most of what I know about scientific research, writing, speaking, teaching, and much more. Most of the results and ideas presented in this thesis were sparked by initial conversations with André and significantly polished over the course of our in-person/Zoom meetings. The thesis material is also significantly shaped by the insights I gained from working with members and visitors of the Logical Systems Lab, both past and present, including: Stefan Mitsch, Khalil Ghorbal, Fabian Immler, Jean-Baptiste Jeannin, William Simmons, Andrew Sogokon, Noah Abou El Wafa, Rose Bohrer, Nathan Fulton, Aditi Kabra, Katherine Kosaian, Jonathan Laurent, Samuel Teuber, and James Gallicchio.

In the broader CMU community, I am grateful for the opportunity to learn about neighboring research areas from the Principles of Programming group and the Lean working group. The administrative staff, especially Deb Cavlovich and Jessica Packer, helped make my time at CMU a smooth sailing one. The Pittsburgh Symphony Orchestra provided many evenings of quality entertainment; readers interested in classical music are strongly recommended to check out their subscription packages.

Finally, I would like to thank all of my thesis committee members for their time and for their invaluable advice on the contents of this thesis.

Contents

Chapter 1

Introduction

Ordinary differential equations (ODEs) are quintessential mathematical models in the engineering and physical sciences. Their importance stems from the fact that they provide a succinct and tractable way of describing the complex dynamics of real-world continuous systems. However, many real-world systems of practical interest, such as self-driving cars, autonomous factory robots, or airplane autopilots, cannot be modeled purely continuously because they are also controlled by discrete software components. For example, the continuous motion of an autonomous vehicle may be controlled by software that makes abrupt, discrete changes to its steering, braking, or other control inputs. These systems are examples of so-called *cyber-physical systems* (CPSs) that feature an interaction between discrete computational control and continuous real-world physics. The increasing prevalence of CPSs and the fact that they often operate in safety- and mission-critical settings entail the need to ensure that those systems operate safely and correctly. *Hybrid systems* [28, 45, 66, 75] are mathematical models that combine discrete and continuous dynamics, and interactions thereof; this combination makes them natural models to use in the study of the discrete-continuous dynamics present in CPSs [6, 144].

Deductive reasoning for hybrid systems is an emergent approach for the formal verification of CPSs [45, 135, 144, 164, 188, 208, 213]. Broadly speaking, deductive reasoning refers to a class of logic- and proof-based verification techniques [45] where a given system is shown to have the desired correctness properties through a series of syntactic deduction steps. Such techniques are attractive for hybrid systems because their flexibility allows for proofs of comprehensive correctness specifications, while their sound logical foundations ensure trustworthiness of the resulting conclusions. Compositionality is the key for scaling deductive techniques to larger models and realizing the aforementioned benefits in practical applications. In particular, deductive proofs for hybrid systems formally decompose those systems into their constituent dynamics and then prove properties of the resulting simpler subsystems separately [45]. Conversely, compositionality also lifts insights for discrete and continuous dynamics to corresponding insights for hybrid systems. The goal of this thesis is to push the boundaries of formal verification for hybrid systems by furthering the understanding of deductive reasoning for their constituent continuous dynamics described by ODEs. The pursuit of this goal is guided by the thesis statement:

> *Deductive reasoning provides a powerful, uniform, and foundational way of proving properties of ordinary differential equations. This logical foundation, in turn, yields new insights towards the verification of continuous and hybrid systems.*

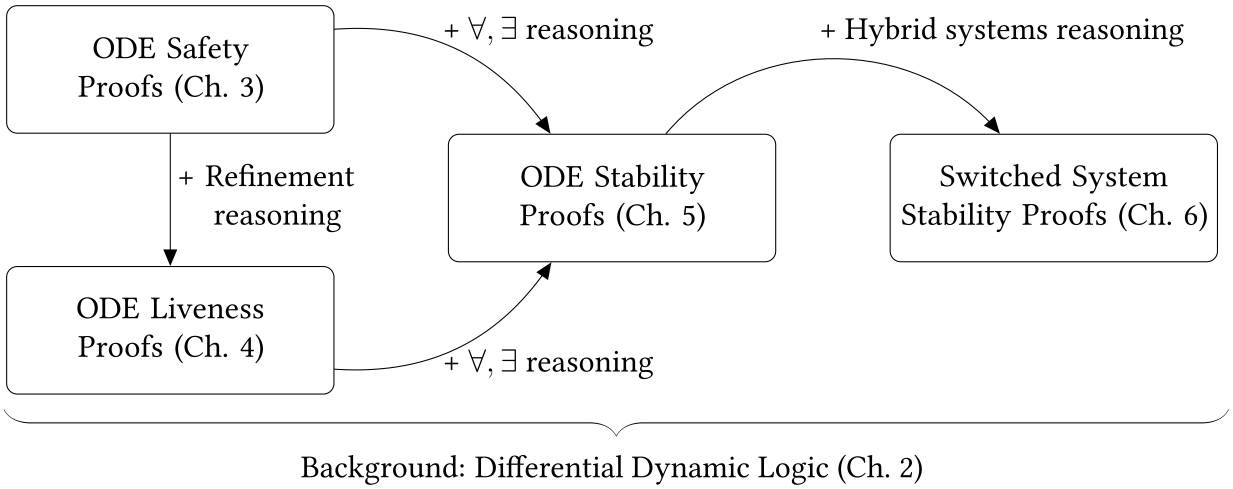
ODE Safety Proofs (Ch. 3)
+ ∀, ∃ reasoning
+ Hybrid systems reasoning
+ Refinement reasoning
ODE Stability Proofs (Ch. 5)
Switched System Stability Proofs (Ch. 6)
ODE Liveness Proofs (Ch. 4)
+ ∀, ∃ reasoning
Background: Differential Dynamic Logic (Ch. 2)

Figure 1.1: An overview of the thesis chapters and dependencies between those chapters. An arrow Chapter A $\longrightarrow$ Chapter B indicates that Chapter B builds on material from Chapter A by adding new proof ideas and deduction techniques.

1.1 Thesis Overview

This thesis studies deductive verification for three key classes of ODE properties: *safety* (Chapter 3), *liveness* (Chapter 4), and *stability* (Chapter 5). Briefly, an ODE is *safe* if its continuous solutions from a given set of initial states *always* stay within a prescribed set of safe states. Dually, an ODE is *live* if its solutions *eventually* enter a specified goal region from the given initial states. Proofs of ODE safety and liveness specifications provide mathematical guarantees that a given continuous system operates safely and reaches its goals successfully. Proofs of *stability* show that the ODE models under consideration are suitably robust to small, mathematically-defined perturbations of the system. Such stability properties are complementary to safety and liveness properties, especially for control systems which must be designed to operate robustly in the presence of real-world perturbations. Figure 1.1 provides an overview of the thesis chapters and illustrates the key role that compositionality plays throughout this thesis—later chapters build on the deductive proofs from earlier chapters by adding various reasoning aspects. This compositional arrangement culminates in the study of deductive stability proofs for *switched systems* (Chapter 6), a class of hybrid systems featuring discrete switching between a family of continuous modes. Hybrid switching designs can be used to achieve control objectives that cannot otherwise be achieved by purely continuous means, but they can also exhibit subtle stability behavior (see Section 1.1.4) which makes rigorous proofs of their stability desirable.

All of the aforementioned properties are formalized in this thesis using *differential dynamic logic* (dL), a logic for deductive verification of hybrid systems [135, 139, 142, 144] (Chapter 2). The use of dL as a uniform logical foundation throughout this thesis has several key benefits:

1. The dL proof calculus enables *compositional reasoning*, as illustrated by the logical dependencies in Fig. 1.1. Proof insights for ODEs are lifted to corresponding results for hybrid systems through dL's *hybrid program* modeling language and reasoning principles [144].

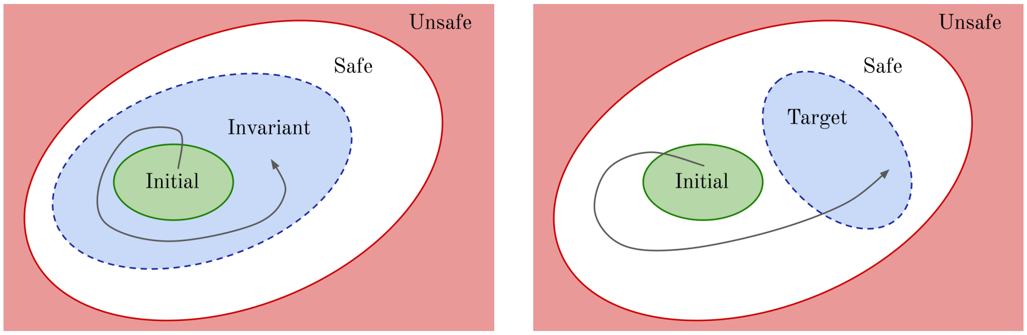

Figure 1.2: (Left) An appropriate invariant set shows that ODE evolutions from the initial states always stays in the invariant and cannot enter the unsafe states. (Right) Trajectories from the initial state eventually enter the target region without leaving the domain of safe states.

2. All reasoning principles for safety, liveness, and stability are syntactically *derived* from a parsimonious set of dL axioms with well-established semantic and axiomatic foundations [139, 142] which yields utmost confidence in the correctness of the resulting proofs. Moreover, the dL foundation provides a basis for investigating proof-theoretical questions, such as the completeness and deductive power of the proof calculus [139, 140].

3. The dL proof calculus is implemented in the KeYmaera family of hybrid system theorem provers [115, 147] and related tools [55, 180]. Insights from this thesis are put into practice through the KeYmaera X prover [54] which implements dL's uniform substitution calculus [142]. Thanks to the shared dL foundation, the thesis implementations and case studies require *minimal* extensions to KeYmaera X's soundness-critical axiomatic core and thereby directly inherit the trustworthiness of KeYmaera X's microkernel design [54, 115].

Informally, dL is also an excellent setting in which familiar intuitions from discrete program verification can be analogously applied to continuous and hybrid systems. This connection is exploited throughout the thesis, for example, the classical technique of introducing auxiliary *ghost* state to a program for the sake of its verification [127] is crucially used to aid differential equations reasoning in Chapters 3 and 4 through dL's continuous *differential ghost* principle [139, 140]. The same idea reappears under a slightly different guise for stability specifications in Chapter 5 and yet again in Chapter 6 for a switched system stability verification case study.

1.1.1 Safety and Invariance for Ordinary Differential Equations

An *invariant* (also called *positive invariant*) of an ODE is a set of states that cannot be left by evolving the ODE forward in time starting from any state in that set. Such invariants play an important role in proofs of safety for ODEs [144], as illustrated in Fig. 1.2 (left). Suppose the system under consideration is evolved forward in time from a set of initial states (in green) and it is undesirable for the system to enter the unsafe states (in red). The invariant set (in blue, with dashed boundary) contains the initial set and is a subset of the safe states (unshaded). Consequently, all forward evolutions of the ODE from the initial set are trapped within the invariant and can therefore never enter the set of unsafe states.

Chapter 3 presents a dL *differential equation axiomatization* [148, 149] that is complete for reasoning about ODE invariants. The key result is that, given an ODE and a candidate invariant characterized by a formula of arithmetic, one can always syntactically prove (or disprove) invariance of the candidate invariant from the parsimonious axiomatization of differential equations reasoning principles in dL. The dL logical foundation is crucial for this result and, in fact, Chapter 3 [149] proves the completeness result more generally for all ODEs and invariants in any extended dL term language that meets three extended term conditions because the dL axiomatization remains sound for those extensions [149]. The rich class of *Noetherian* functions [12, 56, 57, 201] is shown to meet those criteria and therefore automatically inherits the completeness result. Extension of dL's term language with Noetherian functions is practically useful because the Noetherian class includes many (non-polynomial) functions of use in models of continuous and hybrid systems, e.g., the exponential and trigonometric functions. The chapter also proves related completeness results for the special case of *equational* invariants, where dL's hybrid program axioms are used to compositionally lift those results to hybrid systems.

1.1.2 Liveness and Existence for Ordinary Differential Equations

An ODE is *live* if its solutions from the prescribed initial states eventually enter a given goal or target region. For reach-avoid type applications [18, 22], these solutions may also be required to avoid certain unsafe states before the goal or target region is reached. This liveness property is illustrated in Fig. 1.2 (right), where trajectories from the initial set (in green) enter the target region (in blue, dashed boundary) without entering the set of unsafe states (in red). Liveness arguments are subtle in the continuous setting, e.g., ODE solutions may cease to exist after a short time period before they reach the target region. These subtleties are the source of several soundness errors identified in liveness arguments from the literature (see Table 4.1, page 76).

Chapter 4 [192, 195] presents an approach for deducing liveness properties of ODEs through systematic, step-by-step *refinements*, where each step is justified using an ODE safety (or invariance) property proved using the results of Chapter 3. The idea of using safety properties in proofs of liveness properties is reminiscent of deductive proofs of liveness for discrete (concurrent) systems [109, 128] and Chapter 4's key insight is that this idea generalizes to the continuous setting—as long as the new technical subtleties, such as sufficient duration existence of solutions, are appropriately identified and handled. Refinement reasoning in dL provides a sound and uniform basis for navigating these subtleties, in contrast to earlier ad hoc (often unsound) approaches from the literature (surveyed in Table 4.1). Refinements also naturally generalize to new ODE liveness arguments by soundly mix-and-matching or generalizing previously identified refinement steps in new ways. As a special case, Chapter 4 applies the refinement approach to deductive proofs of sufficient duration existence for ODEs which are a key hypothesis behind ODE liveness arguments. These insights are put into practice through an implementation of ODE existence and liveness proofs in KeYmaera X.

1.1.3 Stability for Ordinary Differential Equations

At a high level, stability properties can be understood as a combination and/or variation of two underlying properties [165]:

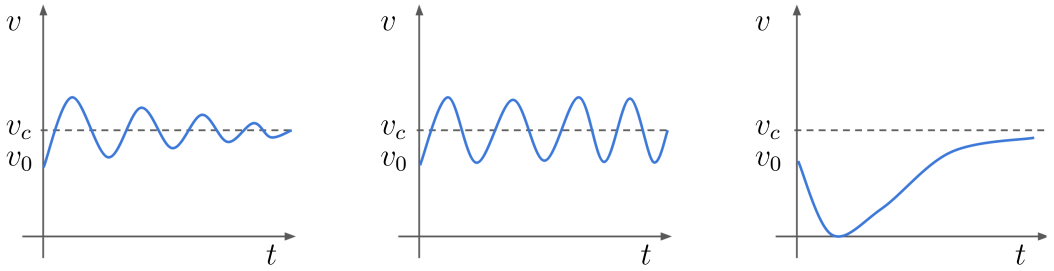

Figure 1.3: The leftmost plot shows a stable and attractive cruise controller driving a car's speed v back to the desired cruising speed v_c over time t from a small initial perturbation v_0. Note that v remains near v_c throughout (stability) and converges to v_c as $t \to \infty$ (attractivity). The middle plot shows a stable controller lacking attractivity, where v oscillates around v_c but never converges to v_c. The rightmost plot shows an attractive controller that lacks stability; although v converges to v_c, the car is transiently driven to a stop $v = 0$ before speeding up again.

> *Stability.*[1] A *stable* system always stays close to its desired operating state(s) when initially slightly perturbed from those operating state(s).
>
> *Attractivity.* An *attractive* system dissipates initial perturbations and eventually returns to a desired operating state.

A familiar example illustrating these properties is a cruise controller that is attempting to keep a car at its desired cruising speed [6]. Stability of the cruise controller ensures that small perturbations to the car's cruising speed, e.g., from minor bumps in the road, do not result in large and potentially unsafe changes in its speed. Attractivity of the controller ensures that those perturbations are eventually dissipated so that the car successfully returns to its cruising speed and does not, e.g., remain in uncomfortable oscillations that might be dangerous for the car's passengers. The behavior of various cruise controllers are illustrated in Fig. 1.3.

The informal descriptions of stability and attractivity are, respectively, similar to the descriptions of ODE safety and liveness, with the added twist that stability and attractivity are concerned with the *local* and *long-term* (asymptotic) behavior of a system near its desired operating state(s). Chapter 5 [194] shows how ODE stability and attractivity can be formally specified as *quantified* and *nested* ODE safety and liveness properties, where the first-order quantifiers $\forall, \exists$ are used to express local and/or long-term asymptotic behaviors of the ODE solutions. Accordingly, deductive proofs of stability specifications are built by formalizing classical *Lyapunov function*-based techniques [98] for stability through the combination of quantifier reasoning and dL's ODE safety and liveness reasoning principles studied in the preceding Chapters 3 and 4. The flexibility afforded by dL's formula syntax and proof calculus is crucial because there are a number of stability variations, e.g., exponential or set stability, that may be of interest for any given system with different specifications (see Chapter 5 for further examples). Unlike existing stability verification approaches [3, 61, 88, 104, 170], the deductive approach rigorously proves every step of a stability argument as opposed to arithmetic conditions that imply a given stability

[1] The word "stability" is often used as a one-size-fits-all term to refer to various related stability notions. The description of stability here is in the sense of the mathematical definition of Lyapunov stability for ODEs [71, 89, 165].

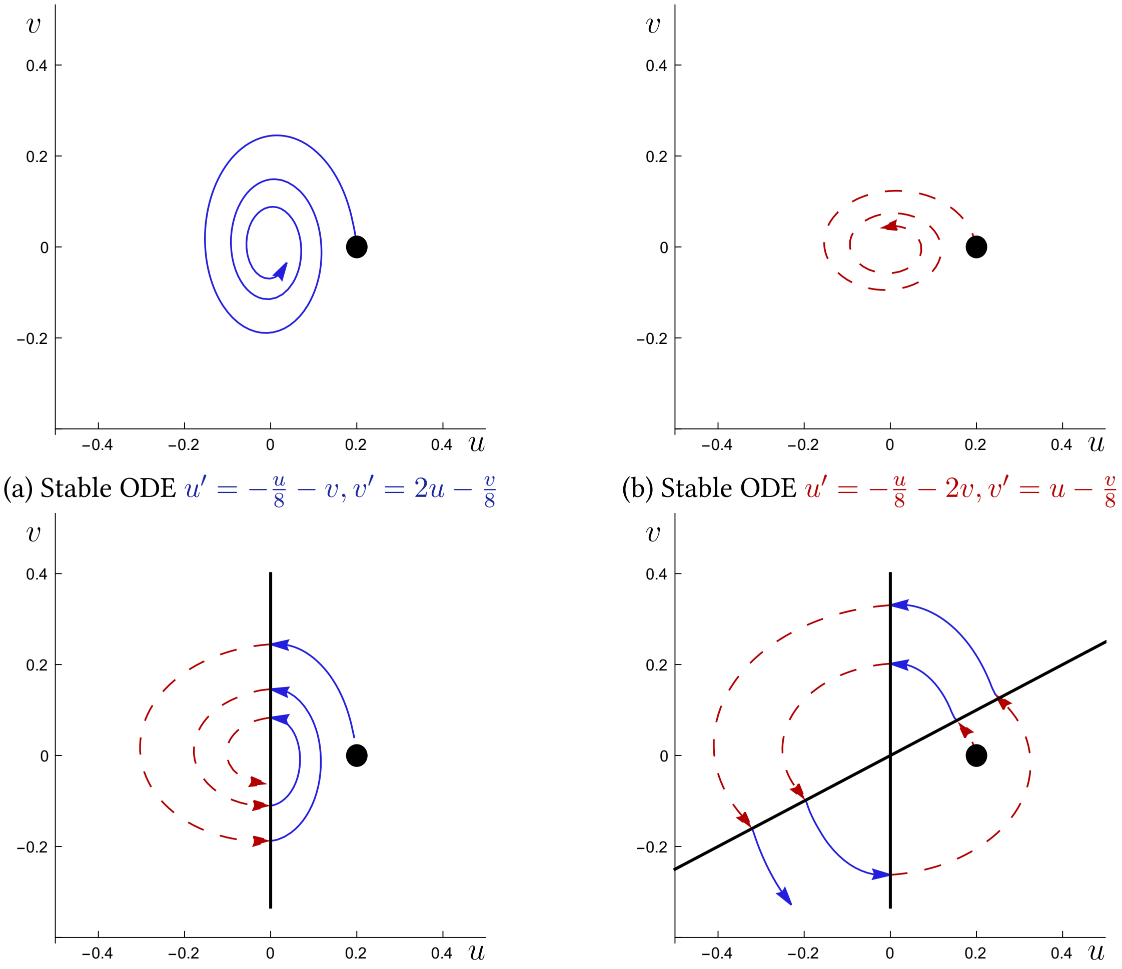

(a) Stable ODE $u' = -\frac{u}{8} - v, v' = 2u - \frac{v}{8}$

(b) Stable ODE $u' = -\frac{u}{8} - 2v, v' = u - \frac{v}{8}$

(c) Stable switching along $u = 0$

(d) Unstable switching along $u = 0$ and $u = 2v$

Figure 1.4: Figures 1.4a and 1.4b in the top row show trajectories for ODEs in blue and dashed red that spiral towards the stable origin $u = 0, v = 0$. Figures 1.4c and 1.4d in the bottom row show two different switching designs that produce opposite stability outcomes for the overall system: a stable spiral towards the origin for Fig. 1.4c but an unstable trajectory that diverges from the origin for Fig. 1.4d. The alternating colors show the ODE that is being followed at each point along the trajectory and the solid black lines indicate switching boundaries.

notion. Thus, instead of building an entirely new verification tool or proof approach for each new stability notion, the deductive approach formalizes them all within dL, which enables the use of KeYmaera X as a single trustworthy tool for stability verification.

1.1.4 Stability for Switched Systems

Stability is also an important design objective when the real-world systems of interest are modeled by hybrid systems instead of purely continuous ODEs [44, 99, 118, 189]. The added complication is that stability of a given hybrid system is a function of *both* its continuous ODEs and its discrete dynamics, so both dynamics must be adequately accounted for in the overall stability proof. This subtlety is present even for the sub-class of *switched systems*, i.e., hybrid systems that discretely switch between a family of continuous modes but without discrete jumps in their system state [28]. For example, Fig. 1.4 shows a system that switches between two ODEs which are individually stable (Figs. 1.4a and 1.4b), but where the overall system becomes unstable when subjected to a destabilizing switching signal (Fig. 1.4d). It is desirable to design discrete switching mechanisms that maintain stability of the overall system (Fig. 1.4c) and, in some cases, it may even be possible to use switching to stabilize otherwise unstable ODEs or to use appropriately designed switching control to achieve control goals that cannot be achieved by purely continuous means [99]. Rigorous proofs of switched system stability are important to ensure that a given switching design achieves the intended effect of stabilizing a system, especially for complicated designs where pen-and-paper proofs become error-prone.

Switched systems and their stability questions provide a practically useful proving grounds for demonstrating the latter part of the thesis statement, i.e., that the deductive approach to ODEs *yields new insights towards the verification of hybrid systems.* Chapter 6 [193, 196] applies the results of the preceding Chapters 3–5 to the study of switched system stability.

Switched Systems as Hybrid Programs [193]. Various classes of switched systems are modeled as *looping* hybrid programs in dL, where each loop iteration models a discrete switching step followed by continuous evolution of the chosen mode. This bridge between formalisms—switched systems from control theory and hybrid programs from verification—leads to fruitful cross-pollination of ideas from their respective fields. For example, the completeness results for ODE invariants in Chapter 3 are compositionally extended to complete invariance proof rules for various classes of switched systems, which facilitates their effective safety verification in dL.

Switched System Stability Verification [196]. Switched system stability is proved for looping hybrid programs by blending classical ideas from the controls and verification literature using dL. From controls, standard stability notions for various classes of switching mechanisms are used [65, 66, 99], along with their corresponding Lyapunov function-based analysis techniques [27, 89, 99]. From verification, properties of looping hybrid programs (modeling switched systems) are verified by finding appropriate *loop invariants*, i.e., properties that are preserved across each loop iteration [144]. This blend of ideas enables a trustworthy implementation of switched system stability verification in KeYmaera X providing fully automated stability proofs for standard classes of switching mechanisms, including automatically searching for suitable Lyapunov functions. The generality of the dL approach also allows for verification of switching control laws that require non-standard stability arguments because users can design loop invariants that suitably express specific intuitions behind those control laws. This flexibility is demonstrated on several case studies in Chapter 6.

1.1.5 Chapter Layout

Chapters 3–6 each have an associated appendix containing omitted details and proofs (Appendix A–D). It is recommended to read the thesis chapters *in order*, because later thesis chapters build significantly on results developed in earlier chapters (see Fig. 1.1).

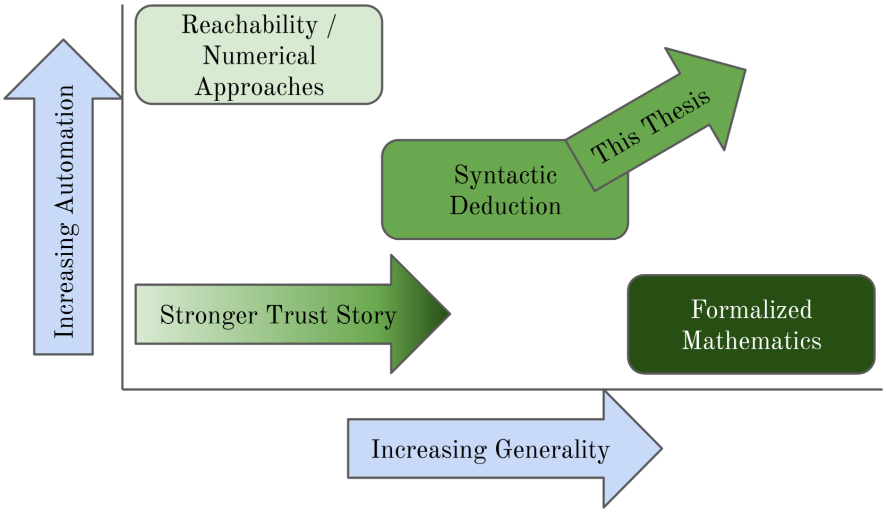

Figure 1.5: An overview of approaches for the verification of continuous and hybrid systems. The underlying trust story for each approach is shown in green, with darker shades corresponding to a comparatively stronger trust story behind the approach's verification results.

1.2 Related Work

This related work discussion broadly examines approaches for the formal verification of continuous and hybrid systems and explains how this thesis relates to the broader hybrid systems verification landscape. Chapters 3–6 contain further related work discussion specific to the material of the respective chapters. The discussion here is structured along a high-level categorization of formal verification approaches along the *automation*, *generality*, and *trust story* axes, as shown in Fig. 1.5. The *automation* axis examines the tools that implement a given category of approaches; the *generality* axis examines the types of specifications that can be tackled, especially the safety, liveness, and stability properties studied in this thesis; and the *trust story* axis examines what needs to be trusted by a user in order to trust a verification result produced by the tools. This thesis seeks to improve syntactic deduction along the automation and generality axes while retaining its strong trust story. Of course, the boundaries between categories are not always clear-cut and approaches that overlap several categories are highlighted as well.

1.2.1 Reachability Approaches

There is an extensive literature on model checking and reachability analysis for hybrid systems [7, 31, 35, 48, 52, 53], often based on analyzing hybrid automata [7, 75], see Doyen et al. [45] for a comprehensive survey. Briefly, these approaches compute an overapproximation of the image of a set of initial states under the dynamics of a hybrid system. Safety specifications are verified by checking that the overapproximate images do not intersect unsafe states.

- *Automation.* There are a number of tools for automated hybrid system reachability analysis that handle different sub-classes of hybrid dynamics, e.g., SpaceEx [53] handles piecewise affine hybrid automata while Flow* [31] supports automata with nonlinear dynamics, see Doyen et al. [45, Chapter 30.7] for a survey of other reachability tools [5, 15, 48, 52, 172]. A key ingredient for reachability computations is the use of efficient internal representations of state sets that provide different trade-offs in runtime and accuracy for computing overapproximations [45]. For example, SpaceEx converts between template polyhedra and support function representations to apply different operations [53] while Flow* is based on Taylor model overapproximations [31]. Several tools provide modular libraries [5, 15, 172] so that users and developers can flexibly experiment on different combinations of representations and algorithms.

- *Generality.* Reachability analysis tools are often limited to *i)* sub-classes of hybrid systems, e.g., with piecewise affine dynamics in SpaceEx [53], *ii)* analysis over sufficiently small bounded initial sets and for a finite time horizon to avoid error growth in overapproximations [45], and *iii)* fixed (or bounded) parameters for the hybrid system models. The use of reachability analysis for liveness-type properties has not been as extensively explored [122], although specialized tools have used, e.g., (under-approximate) reachability for liveness [32, 67] and stability verification [151], see related work discussion in Chapters 4–6. Overall, reachability approaches have powerful automation but apply to restricted subsets of the safety, liveness, and stability properties considered in this thesis.

- *Trust story.* To trust the output of each verification tool mentioned above, one must trust the correctness of the underlying mathematical justification *intrinsic* to each approach and the *extrinsic* tool implementation details. Intrinsic errors in mathematical justification nullify verification results since there is no guarantee that those results correctly imply the desired property. Extrinsic errors are mismatches between the tool implementation and underlying justification. This also leads to gap in the trust story, e.g., the use of floating point rather than exact arithmetic for efficiency in implementations [53]. One way to eliminate both forms of error is to formally prove their absence within a general purpose proof assistant, e.g., Immler [79] verified an ODE solver in Isabelle/HOL [124] using rigorous Runge-Kutta methods [199, Section 2.7] that can produce verified enclosures for solutions of ODEs. However, the verified tool's performance is not competitive with other (unverified) tools [84]. Another trust gap is that a disparate combination of tools must be used in concert to achieve a comprehensive verification result of multiple specifications, e.g., of safety, liveness, and stability for a given system, which enlarges the trusted base of the verification result. Notwithstanding the trust story, the bag-of-tools approach is advantageous because it applies complementary tools with different strengths to tackle sub-classes of specifications. Combined approaches using reachability tools and others (discussed below) have been successfully used for several case studies [101, 186, 212].

1.2.2 Numerical and Certificate-Based Approaches

An alternative class of automated verification techniques is based on generating certificates that imply the properties of interest [45, Chapter 30.6]. For example, *barrier certificates* [155] are a popular technique for certifying safety of a hybrid system, where the trajectories of the hybrid system are proved to never enter unsafe states by showing that a given barrier function (typically a polynomial function) satisfies certain arithmetical conditions that imply safety. Such barriers are often found by numerically solving a sum-of-squares problem [129] with constraints that encode the appropriate arithmetical conditions. Similar techniques for ODE and hybrid system safety are based on finding *invariants* of those systems [146, 161, 169] (see Chapter 3). Other (numerical) certificates include *variants* that imply liveness [156, 157, 191] (see Chapter 4) and *Lyapunov functions* that imply stability [3, 61, 88, 170] (see Chapters 5 and 6).

- *Automation.* Approaches based on automatically generating certificates are implemented in various tools [3, 61, 88, 91, 155, 156, 157, 170, 178, 180, 191]. The main difference between tools is how they encode different (sufficient) arithmetical conditions into an appropriate input format for different numerical solvers. For example, sufficient conditions for barrier certificates come in various flavors, including strict [155], Darboux [210], exponential [91], and vector barrier certificates [178]. These conditions have also been encoded (and solved) differently, e.g., with linear programming [209], sum-of-squares programming [91, 155], or other novel combinations of encoding and/or solving techniques [88, 132, 205, 210].

- *Generality.* Compared to reachability approaches, barrier certificates and invariants are useful for certifying safety with respect to unbounded initial sets and over infinite time horizons. However, tools for generating these certificates are highly-dependent on the success of numerically solving the encoded problems and they suffer from scalability issues for higher dimensional systems [180]. Some numerical generation techniques also require bounded domains for the input model's state or parameter values [3, 61, 170]. Thus, problems can have safety certificates *in theory* but those certificates may not be easily found by tools *in practice.* Certificates for liveness and stability can also be numerically generated [49, 104, 130, 131, 156, 157, 200], subject to similar caveats, see related work discussion in Chapters 4–6. Existing certificate-based methods form the basis for further study in this thesis, e.g., Chapter 4 surveys and generalizes liveness arguments from the literature while Chapters 5 and 6 formalize Lyapunov functions for stability.

- *Trust story.* Intrinsic errors in the mathematical justification for certificate-based approaches have been reported in the literature and in this thesis, e.g., for barrier certificates [45, Chapter 30.6.1], in the survey of liveness arguments (Table 4.1, page 76), and for stability (Appendix C.2). A key contribution of this thesis is to soundly and syntactically justify proof rules for certificates by deriving them from a trustworthy logical basis. Another important trust gap for certificate-based approaches arises when numerically generated results are invalid because the results contain numerical errors and so do not satisfy all of the required arithmetical conditions. Numerical issues have been highlighted in the literature [3, 40, 88, 167, 170, 180] and they are often handled by *not* trusting the output of numerical solvers but instead separately checking the correctness of any *untrusted* candidate certificates with a separate *trusted* tool. Certificate-based approaches can be fruitfully integrated with deductive tools (discussed next) by using the latter tools to soundly check untrusted certificates. For example, untrusted ODE invariant candidates generated by the Pegasus tool [180] are formally checked by KeYmaera X.

1.2.3 Syntactic Deduction

A number of proof calculi have been proposed for syntactic, deductive reasoning for hybrid systems [45, 135, 164, 188, 213], including differential dynamic logic (dL) [135, 139, 142, 144] which is used as the logical foundation throughout this thesis (Chapter 2).

- *Automation.* Hybrid Hoare Logic (HHL) [102, 213] extends Hoare logic to support hybrid systems and it is implemented in the HHL prover [206] embedded in Isabelle/HOL. The dL proof calculus is implemented in the KeYmaera [147] and KeYmaera X [54] provers for hybrid systems [115]. A Hoare-style logic with dL-inspired proof rules has also been implemented in Isabelle/HOL [50, 51]. These tools provide users with a high degree of hybrid system-specific automation [54], e.g., ODE solving, invariant generation [180], and automated application of canonical reasoning steps. However, unlike the reachability and certificate-based approaches, syntactic deduction approaches are *semi-automatic* and often require some manual user input to supply insights for more difficult parts of proof, such as supplying loop invariants or proving properties of (complicated) ODEs [144, 213].

  This thesis improves automation: Chapter 3 furnishes a means of automatically proving ODE invariance; Chapter 4 provides an implementation of ODE liveness; Chapters 5 and 6 derive high-level stability proof rules in dL and implement them in KeYmaera X.

- *Generality.* Proof calculi are syntactically limited by the properties that can be expressed in their specification language and by the reasoning principles available for proving those specifications. For example, in dL, users can model and specify first-order dynamical properties of hybrid system models [73, 144]. Deductive calculi can be compositionally extended to accommodate, e.g., (continuous) differential-algebraic [137], (constructive) adversarial [18, 141, 143], doxastic [110], distributed [136], stochastic [138], and concurrent [164, 213] dynamics. The base logic of dL has also been extended with temporal [85, 133], relational [90], hybrid-logical [17], and refinement [19, 106] reasoning. Compositionality can also be exploited through contract-based reasoning to scale verification to larger hybrid systems consisting of interacting components [87, 100, 119].

  This thesis improves generality: Chapter 3 extends the base term language of dL while retaining complete invariance reasoning; Chapter 4 develops reasoning principles for ODE liveness; Chapters 5 and 6 uses dL to tackle stability and its reasoning principles. The fundamental advantage of compositional reasoning is that all results of this thesis can be applied to answer ODE safety, liveness, and stability sub-questions arising in any of the aforementioned extensions of dL *without* compromising the soundness of their logics.

- *Trust story.* To trust a syntactically proved result, one must trust its logical foundations. This thesis uses dL, whose foundations are well established in the literature [135, 139, 142, 144]. Furthermore, soundness of dL's *uniform substitution* calculus has been formally verified in the general purpose proof assistants Isabelle/HOL and Coq [20]. Both KeYmaera and KeYmaera X implement dL but with different design principles [115]: KeYmaera is built on the KeY prover and inherits its significant visualization, automation, and user interaction features; KeYmaera X is a clean-slate prover designed around a minimal, soundness-critical core kernel implementing the aforementioned uniform substitution calculus for dL. Crucially, KeYmaera X's microkernel ($\approx$2000 lines of Scala) [54, 115] is the *only* part of its implementation that must be trusted to trust its axiomatic verification results. The remaining trust question is how to handle *arithmetic* sub-questions arising in proofs: KeYmaera and KeYmaera X both provide tools for proving simple arithmetic questions and interfaces to various external arithmetic solvers [115]. Verified arithmetic is an active area of research, including formally verified provers [36, 38, 113, 120, 121, 171], certificate-based techniques [74, 150] (also available in KeYmaera), and combinations thereof [97].

This thesis retains the strong trust story of dL and KeYmaera X: it is based on syntactically deriving high-level reasoning principles for ODE safety, liveness, and stability from a core dL axiomatization. Such a syntactic approach fits particularly well with an implementation in KeYmaera X because those high-level reasoning principles are implemented as (untrusted) tactics in KeYmaera X [55], with its sound kernel as a safeguard against implementation errors or mistakes in pen-and-paper proof rule derivations.

1.2.4 Formalized Mathematics

General purpose proof assistants have been used to formalize specifications and theorems for ODEs and hybrid systems. Examples of such efforts include: foundational theory for ODEs and libraries for their verified numerical analysis [80, 83, 108], LaSalle's invariance principle for stability analysis [37, 166], and the Poincaré-Bendixson theorem [82].

- *Automation.* Formalizing mathematics in a proof assistant is often beset by significant manual user effort in proofs, e.g., in dynamical systems formalization, Immler and Tan [82] required $\approx$7000 lines of proof script to formalize the Poincaré-Bendixson theorem, building on significant existing libraries; Rouhling [166] required $\approx$1000 lines of proof to define and prove the stability of an inverted pendulum controller after extending an earlier formalization of LaSalle's invariance principle [37]. This effort contrasts with syntactic deduction (e.g., in KeYmaera X) because users have to reason over the underlying mathematical semantics of the properties being formalized. Addition of domain specific automation, e.g., in the style of dL for hybrid systems [51], can help reduce user effort.

- *Generality.* In principle, general purpose proof assistants allow for formalization of complicated mathematical properties of ODEs and hybrid systems, as long as the required background libraries are available (or sufficient time and resources are devoted to developing those libraries). Examples of such libraries include the Archive of Formal Proofs (AFP) for Isabelle/HOL,[2] the Lean mathematical library [111], and the Mathematical Components library for Coq.[3] The AFP ODE library [81] has been used in a number of ODE and hybrid system formalizations by various authors [20, 82, 145, 208], while the Mathematical Components library is used by Cohen and Rouhling [37]. Libraries for real analysis are available in most proof assistants [23].

[2] https://www.isa-afp.org

[3] https://math-comp.github.io

- *Trust story.* To trust a formalized result, one must trust: *i)* the logical foundations of the proof assistant used, *ii)* the proof assistant implementation, and *iii)* the correct definition of mathematical concepts involved in the formalization. There is ongoing research towards improving trust in the former two points, e.g., by mechanizing the logical foundations [95, 184] and developing verified kernel implementations [2, 42]. The latter point may be alleviated by careful inspection, or by re-using well-accepted definitions from other formalization efforts available in common libraries, such as the ones mentioned above.

Proof assistants can also be used to strengthen the trust story underlying the preceding categories of approaches. Several deductive calculi for hybrid systems have been mechanized, e.g., dL and its game logic extension [20, 145], modal Kleene algebra for hybrid systems [50, 51, 208], and Hybrid Hoare Logic [206]. The VeriDrone [163] and ROSCoq [8] projects both provide mechanized frameworks within Coq which can be used to model and reason about hybrid systems. Full mechanical verification is an intriguing future endpoint for the results of this thesis, although the proof effort required may be prohibitive. In this respect, an important contribution of this thesis is to identify *what* results are possible and useful to mechanize. For example, a key property involving how ODE solutions can enter or exit semialgebraic sets [64, 103] which is used in the complete analysis of ODE invariants (see Chapter 3) has been formalized in PVS for real analytic functions [174].

Chapter 2

Background: Differential Dynamic Logic

This chapter reviews the syntax, semantics, and axiomatics of differential-form differential dynamic logic (dL), which is the hybrid systems specification and verification logic used throughout this thesis. The exposition herein is necessarily abbreviated and adapted for the purposes of this thesis; interested readers are referred to the literature [135, 139, 142, 144] for more complete expositions of dL. Notational conventions used in this thesis are also established in this chapter.

2.1 Syntax

The *syntax* of dL is the language in which hybrid systems are modeled and specified.

2.1.1 Terms

Terms are generated by the following grammar, where $x \in \mathbb{V}$ is a variable from the set of all variables $\mathbb{V}$ and $c \in \mathbb{Q}$ is a rational constant:

$$e, \tilde{e} ::= x \mid c \mid e + \tilde{e} \mid e \cdot \tilde{e} \mid (e)'$$

Terms generated using only the first 4 clauses of this grammar correspond to polynomials over the variables $\mathbb{V}$ under consideration. *Differentials* $(e)'$ are used in dL for sound differential equations reasoning [142], where the value of $(e)'$ relates to how the value of term e changes with each of its variables as a function also of how those variables themselves change. The fundamental insight is that, along the evolution of a differential equation, the semantic value of differential $(e)'$ coincides with the analytic time derivative of e [142, Lem. 35], so that proofs about *equations of differentials* yield proofs about *differential equations*. It is crucial for soundness and compositionality that differentials have a local semantics defined in any state, so that they can be used correctly in any context to draw sound conclusions from syntactic manipulations mixing dynamic statements about differential equations and static statements about differentials. The precise semantics of differentials is elaborated in Section 2.2, while Section 2.3.3 explains how they can be used to obtain a syntactic representation of (semantic) time derivatives along solutions to differential equations. Syntactically, every variable $x \in \mathbb{V}$ is assumed to have a corresponding differential variable $x' \in \mathbb{V}$ which, like differential terms, are syntactic representations of the semantic time derivative of x along ODE solutions. Section 2.3.3 shows that, in the context of an ODE, differential terms (and variables) can be provably turned into terms that do not contain any differentials or differential variables. Thus, $e, \tilde{e}$ is exclusively used in this thesis to refer to differential-free terms, i.e., $(e)'$ is the differential of e, where e is a differential-free term.

Notational Conventions (Terms). Results are often proved vectorially in this thesis, e.g., if they hold for a differential equation system over any (finite) number of state dimensions. Thus, for convenience, x will often be used to abbreviate a vector of state variables $x = (x_1, \dots, x_n)$. Terms $e(x), \tilde{e}(x)$ are written with explicit arguments (x) for emphasis when it is important that the terms depend only on variables x free. When this dependency is unimportant, terms $e, \tilde{e}$ are written as usual without any arguments. The notation p, q is reserved for emphasis when the terms p, q are polynomials, with dependencies $p(x), q(x)$ added when necessary. For n-dimensional vectors of terms $e = (e_1, \dots, e_n), \tilde{e} = (\tilde{e}_1, \dots, \tilde{e}_n)$, the dot product is $e \cdot \tilde{e} \overset{\text{def}}{=} \sum_{i=1}^{n} e_i \tilde{e}_i$ and the squared Euclidean norm is $\|e\|_2^2 \overset{\text{def}}{=} \sum_{i=1}^{n} e_i^2$. Other norms are explicitly defined when they are used.

2.1.2 Formulas

Formulas are generated by the following grammar, where $\sim \in \{=, \neq, \geq, >, \leq, <\}$ is a comparison operator and α is a hybrid program (defined in Section 2.1.3).

$$\phi, \psi ::= \overbrace{e \sim \tilde{e} \mid \phi \land \psi \mid \phi \lor \psi \mid \neg \phi \mid \forall x \, \phi \mid \exists x \, \phi}^{\text{First-order formulas of real arithmetic } P, Q} \mid [\alpha]\phi \mid \langle \alpha \rangle \phi$$

This grammar extends the first-order language of real arithmetic ($\text{FOL}_{\mathbb{R}}$) with the box ($[\alpha]\phi$) and diamond ($\langle \alpha \rangle \phi$) modality formulas, which express that *all* or *some* runs of hybrid program α satisfy the *postcondition* ϕ, respectively. Note that the modalities $[\cdot], \langle \cdot \rangle$ can be freely nested with first-order and modal connectives, which is crucial for the dL axiomatization (Section 2.3) and the specification of stability properties in Chapters 5 and 6.

Notational Conventions (Formulas). In real arithmetic, formulas can be equivalently normalized such that every atomic comparison $e \sim \tilde{e}$ has 0 on the right-hand side and $\sim \in \{=, \geq, >\}$. The notation $e \succcurlyeq \tilde{e}$ (resp. $e \preccurlyeq \tilde{e}$) is used when the comparison operator can be either $\geq$ or $>$ (resp. $\leq$ or $<$). Other logical connectives, e.g., $\rightarrow, \leftrightarrow$, are definable as usual in classical logic. As with terms, variable dependencies for formulas $\phi(x), \psi(x)$ are added when necessary. Formulas *not* containing the modal connectives, i.e., formulas of $\text{FOL}_{\mathbb{R}}$, are written as P, Q. This convention for formulas P, Q is modified for extended term languages introduced in Section 3.2 (page 33).

2.1.3 Hybrid Programs

Hybrid programs are generated by the following grammar, where x is a variable, e is a dL term and Q is a dL formula.

$$\alpha, \beta ::= \overbrace{x' = f(x) \,\&\, Q}^{\text{Continuous program}} \mid \overbrace{x := e \mid ?Q \mid \alpha ; \beta \mid \alpha \cup \beta \mid \alpha^*}^{\text{Discrete programs and connectives}}$$

Continuous Programs. Chapters 3–5 focus on deductive verification for the *continuous program* $x' = f(x) \,\&\, Q$, which is a differential equation system $x_1' = f_1(x), \dots, x_n' = f_n(x)$ over variables $x = (x_1, \dots, x_n)$, where the LHS x_i' is the time derivative of x_i and the RHS $f_i(x)$ is a dL term over variables x. Following the notational convention, term $f_i(x)$ is differential-free so the ODE system $x' = f(x)$ is given in explicit form [142]. The *autonomous* ODEs $x' = f(x)$ do not depend explicitly on time on the RHS. A standard transformation is to add a clock variable t to the system with $x' = f(x,t), t' = 1$ if time dependency on the RHS is desired. The *evolution domain constraint* formula Q restricts the set of states in which the ODE is allowed to evolve continuously; the ODE is simply written as $x' = f(x)$ when there is no domain constraint, i.e., $Q \equiv true$. The following example ODE α_e over variables u, v is illustrated in Fig. 2.1:

$$\alpha_e \equiv u' = -v + \frac{u}{4}(1 - u^2 - v^2), v' = u + \frac{v}{4}(1 - u^2 - v^2) \tag{2.1}$$

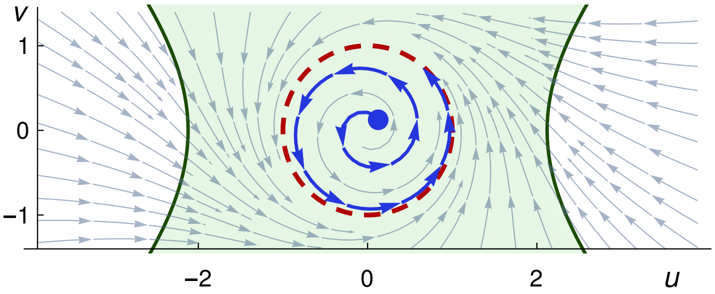

Figure 2.1: The red dashed circle $u^2 + v^2 = 1$ is approached by solutions of α_e from all initial points except the origin, e.g., the blue trajectory starting from $(\frac{1}{8}, \frac{1}{8})$ spirals towards the circle. The red circle, green region $u^2 \leq v^2 + \frac{9}{2}$, and origin are invariants of the system.

To intuitively understand continuous evolution, it is useful to draw an analogy between ODEs and discrete program loops: solutions of an ODE must continuously (locally) follow its RHS at each time instant; analogously, a looping program must repeat its local description given by its loop body on each iteration.[1] The continuous evolution of α_e is visualized in Fig. 2.1 with directional arrows corresponding to the RHS of α_e evaluated at points on the plane. Observe that, even though the RHS of α_e are polynomials in variables u, v, its solutions, which must locally follow the arrows, trace out trajectories in the u, v plane that exhibit complex global behavior. For example, Fig. 2.1 suggests that all points on the u, v plane (except the origin) globally evolve towards the unit circle.

Returning to the analogy, the complex behavior of ODE solutions is unsurprising: even though the body of a loop may be simple, it is almost always impractical to reason about the global behavior of loops by unfolding all possible iterations. Instead, the premier reasoning technique for loops is to study their loop invariants, i.e., inductive properties that are preserved across each execution of the loop body. Similarly, invariants of ODEs describe subsets of the state space from which solutions of the ODEs cannot escape. Invariance reasoning principles for ODEs form the basis upon which all further ODE reasoning is built in this thesis.

Discrete and Hybrid Programs. *Hybrid* dynamics arise in hybrid programs through the combination of continuous ODEs with discrete programming constructs: discrete assignment $x := e$ sets the value of variable x to that of term e in the current state; test $?Q$ checks that formula

[1] In fact, this analogy can be made precise: dL also has a converse relative completeness theorem [139, Theorem 2] that reduces hybrid systems and their ODEs completely to discrete Euler approximation loops.

Q is true in the current state and aborts the run otherwise; the sequence program $\alpha;\beta$ runs program β after α; the choice program $\alpha \cup \beta$ nondeterministically chooses to run either α or β; and the loop program α^* repeats α for $n \in \mathbb{N}$ iterations where n is chosen nondeterministically. The nondeterminism inherent in hybrid programs is useful for abstractly modeling real-world behaviors [144], such as the discrete switching behavior between continuous ODEs studied in Chapter 6.

Notational Conventions (Hybrid Programs). Conditional branching programs (`if-else`) are defined as $\texttt{if}(\phi)\{\alpha\}\texttt{else}\{\beta\} \equiv (?\phi;\alpha)\cup(?\neg\phi;\beta)$. Single-sided conditionals (`if`) are defined as $\texttt{if}(\phi)\{\alpha\} \equiv (?\phi;\alpha) \cup (?\neg\phi)$. Nondeterministic assignments $x := * \equiv x' = 1 \cup x' = -1$ model the assignment of an arbitrary value for x [144, Appendix 12.9.2]. Nondeterministic choice over a finite family of hybrid programs α_p for $p \in \mathcal{P}$, $\mathcal{P} \equiv \{1,\dots,m\}$ is denoted $\bigcup_{p\in\mathcal{P}} \alpha_p \equiv \alpha_1 \cup \alpha_2 \cup \ldots \cup \alpha_m$.

Notational Conventions (Operator Precedence). This thesis uses the standard operator precedences for dL, see [144, Expedition 4.3]. Briefly, all unary operators ($\neg$, $\forall x$, $\exists x$, $[\alpha]$, $\langle\alpha\rangle$, and $(\cdot)^*$) bind tighter than binary operators, sequential composition ; binds tighter than choice $\cup$, conjunction $\land$ binds tighter than disjunction $\lor$, and both $\land$, $\lor$ bind tighter than (bi)implication $\rightarrow$, $\leftrightarrow$. Arithmetic operators are left-associative, while logical and program operators are right-associative, except $\rightarrow$, $\leftrightarrow$ which require explicit parentheses.

2.2 Semantics

The *semantics* of dL gives a formal, mathematical meaning to the elements of its syntax.

2.2.1 Terms

A dL state $\omega : \mathbb{V} \to \mathbb{R}$ assigns a real value to each variable in $\mathbb{V}$; the set of all states is written $\mathbb{S}$. The semantics of term e in state ω is written as $\omega[\![e]\!] \in \mathbb{R}$ and it is defined as usual for the standard arithmetic operators [144, Definition 2.4], e.g., $\omega[\![x]\!] = \omega(x)$ and $\omega[\![e + \tilde{e}]\!] = \omega[\![e]\!] + \omega[\![\tilde{e}]\!]$. The semantics of differentials [142] is the sum of partial derivatives $\frac{\partial\omega[\![e]\!]}{\partial x}$ by all variables $x \in \mathbb{V}$ multiplied by the values of their associated differential variables x', where $\omega(x')$ selects the direction in which x evolves locally and $\frac{\partial\omega[\![e]\!]}{\partial x}$ describes how the value of e changes with a change in the value of x:

$$\omega[\![(e)']\!] = \sum_{x\in\mathbb{V}} \omega(x') \frac{\partial\omega[\![e]\!]}{\partial x} \tag{2.2}$$

Note that the semantics of differentials $(e)'$ is well-defined for isolated states ω, independent of any ODEs. Their importance for differential equations reasoning in dL stems from the (upcoming) semantics of dL formulas and hybrid programs.

2.2.2 Formulas and Hybrid Programs

Formulas. The semantics of comparison operations and first-order logical connectives are defined as usual [144, Definition 2.5], with $[\![\phi]\!] \subseteq \mathbb{S}$ being the set of states where formula ϕ is true, e.g., $\omega \in [\![e \leq \tilde{e}]\!]$ iff $\omega[\![e]\!] \leq \omega[\![\tilde{e}]\!]$, and $\omega \in [\![\phi \land \psi]\!]$ iff $\omega \in [\![\phi]\!]$ and $\omega \in [\![\psi]\!]$. The semantics of the modal connectives are defined over hybrid program semantics $[\![\alpha]\!] \subseteq \mathbb{S} \times \mathbb{S}$ (below):

$$\begin{aligned}
&\omega \in [\![[\alpha]\phi]\!] \text{ iff for all states } \nu \text{ such that } (\omega, \nu) \in [\![\alpha]\!], \nu \in [\![\phi]\!] \\
&\omega \in [\![\langle\alpha\rangle\phi]\!] \text{ iff there is a state } \nu \text{ such that } (\omega, \nu) \in [\![\alpha]\!] \text{ and } \nu \in [\![\phi]\!]
\end{aligned}$$

The semantics of hybrid programs are transition relations $[\![\alpha]\!] \subseteq \mathbb{S} \times \mathbb{S}$, where $(\omega, \nu) \in [\![\alpha]\!]$ iff state ν is reachable from state ω by running hybrid program α. The semantics is defined inductively over the syntax of hybrid programs as follows.

Continuous Programs. The semantics of an ODE, $[\![x' = f(x) \,\&\, Q]\!]$, is the set of all pairs of states that are connected by some solution of the ODE [142, Definition 7]:

$$\begin{aligned}
(\omega, \nu) \in [\![x' = f(x) \,\&\, Q]\!] \text{ iff there is a duration } 0 \leq T \in \mathbb{R} \text{ and a function } \varphi : [0, T] \to \mathbb{S} \\
\text{with } \varphi(0) = \omega \text{ on } \{x'\}^{\complement}, \varphi(T) = \nu, \text{ and } \varphi \models x' = f(x) \land Q
\end{aligned}$$

The condition $\varphi \models x' = f(x) \land Q$ checks that the differential equations and domain are satisfied $\varphi(\zeta) \in [\![x' = f(x) \land Q]\!]$, with $\varphi(0) = \varphi(\zeta)\ \{x, x'\}^{\complement}$ for all times $0 \leq \zeta \leq T$, and, if $T > 0$, then $\frac{\mathrm{d}\varphi(t)(x)}{\mathrm{d}t}(\zeta)$ exists, and is equal to $\varphi(\zeta)(x')$ for all $0 \leq \zeta \leq T$. In other words, φ is a solution of the differential equations $x' = f(x)$ that always stays in the evolution domain constraint Q. It is required to hold all variables other than x, x' constant and, importantly, the values of the differential variables x' are required to match the value of the RHS $f(x)$ of the differential equation along the solution.

Discrete and Hybrid Programs. The remaining cases for hybrid programs α are as follows, where $\circ$ denotes relational composition and $\alpha^0 \equiv ?\mathit{true}$, $\alpha^{n+1} \equiv \alpha^n; \alpha$, see [144, Definition 3.2].

$$\begin{array}{ll}
(\omega, \nu) \in [\![x := e]\!] & \text{iff } \nu = \omega \text{ except } \nu(x) = \omega[\![e]\!] \\
(\omega, \nu) \in [\![?Q]\!] & \text{iff } \nu = \omega \text{ and } \omega \in [\![Q]\!] \\
(\omega, \nu) \in [\![\alpha; \beta]\!] & \text{iff } (\omega, \nu) \in [\![\alpha]\!] \circ [\![\beta]\!] \\
& \text{i.e., there exists state } \mu \text{ such that } (\omega, \mu) \in [\![\alpha]\!] \text{ and } (\mu, \nu) \in [\![\beta]\!] \\
(\omega, \nu) \in [\![\alpha \cup \beta]\!] & \text{iff } (\omega, \nu) \in [\![\alpha]\!] \text{ or } (\omega, \nu) \in [\![\beta]\!] \\
(\omega, \nu) \in [\![\alpha^*]\!] & \text{iff } (\omega, \nu) \in \bigcup\limits_{n \in \mathbb{N}} [\![\alpha^n]\!], \text{i.e., } (\omega, \nu) \in [\![\alpha^n]\!] \text{ for some } n \in \mathbb{N}
\end{array}$$

Formula ϕ is *valid* iff it is true in all states, i.e., $[\![\phi]\!] = \mathbb{S}$. If formula $P \to [x' = f(x) \,\&\, Q]P$ is valid, then the formula P is called an *invariant* of the ODE, $x' = f(x) \,\&\, Q$. Unfolding the semantics, this means that from any initial state $\omega \in [\![P]\!]$, any solution φ of $x' = f(x)$ starting from ω, which does not leave the evolution domain $[\![Q]\!]$, stays in $[\![P]\!]$ for its *entire duration*. Fig. 2.1 suggests several invariants for the ODE α_e from (2.1). The unit circle, $u^2 + v^2 = 1$, is an equational invariant because the direction of flow on the circle is always tangential to it. The open unit disk, $u^2 + v^2 < 1$, is also invariant because trajectories within the disk spiral towards the circle but never reach it. The green region described by $u^2 \leq v^2 + \frac{9}{2}$ is invariant but needs a careful proof. Similarly, if the formula $P \rightarrow [\alpha^*]P$ is valid, then the formula P is called a *loop invariant* for the looping hybrid program α^*.

Notational Conventions (Semantics). Variables $y \in \mathbb{V} \setminus \{x\}$ that do not occur on the LHS of ODE $x' = f(x)$ remain constant along solutions $\varphi : [0, T] \rightarrow \mathbb{S}$ of the ODE, with $\varphi(\tau)(y) = \varphi(0)(y)$ for all $\tau \in [0, T]$. Since only the values of $x = (x_1, \ldots, x_n)$ change along the solution φ, the solution may also be viewed geometrically as a trajectory in $\mathbb{R}^n$, dependent on the initial values of the constant *parameters* y. Similarly, the values of terms and formulas depend only on the values of their free variables [142]. Thus, terms (or formulas) whose free variables are all parameters for $x' = f(x)$ also have provably constant (truth) values along solutions of the ODE. For formulas $\phi(x)$ that only mention free variables x, $[\![\phi]\!]$ can also be viewed geometrically as a subset of $\mathbb{R}^n$. Such a formula is said to *characterize* a (topologically) open (resp. closed, bounded, compact) set with respect to variables x iff the set $[\![\phi]\!] \subseteq \mathbb{R}^n$ is topologically open (resp. closed, bounded, compact) with respect to the Euclidean topology. In Appendix B.1.3, a more general definition of these side conditions is given for formulas ϕ that mention parameters y. These side conditions are decidable [14, 197] when ϕ is a formula of first-order real arithmetic and there are simple syntactic criteria for checking if they hold (Appendix B.1.3). Analogously, by projecting dL states onto the state variables x of interest for a hybrid program α, the transition semantics of α can be equivalently viewed as a relation in Euclidean space $[\![\alpha]\!] \subseteq \mathbb{R}^n \times \mathbb{R}^n$. Formulas $\mathring{P}$, $\overline{P}$ and ∂P are the syntactically definable topological interior, closure, and boundary of the set characterized by P, respectively [14]. For example, the closure formula for $P(x)$ is defined as, $\overline{P(x)} \equiv \forall t \exists y\, (P(x) \wedge (\|y - x\|_2^2 < t^2 \vee t = 0))$, with fresh variables y, t not appearing in $P(x)$ [14, Proposition 2.2.2]. The interior formula is defined as $\mathring{P} \equiv \neg(\overline{\neg P})$ and the boundary formula is defined as $\partial P \equiv \overline{P} \wedge \neg\mathring{P}$.

2.3 Axiomatics

The *axiomatics* of dL are the sound axioms and proof rules by which dL syntax can be soundly and syntactically manipulated to derive new, valid conclusions.

2.3.1 Sequent Calculus

This thesis uses a standard, classical sequent calculus [134] with the usual rules for manipulating logical connectives and sequents. The semantics of *sequent*[2] $\Gamma \vdash \phi$ is equivalent to the formula $(\bigwedge_{\psi \in \Gamma} \psi) \rightarrow \phi$ and the sequent is valid iff its corresponding formula is valid. Formulas Γ are called *antecedents* of the sequent, while formula ϕ is its *succedent*. Proofs are written as a

[2] For notational simplicity, this thesis uses sequents with exactly one succedent $\Gamma \vdash \phi$ because all of the thesis results focus on proving validity of an ODE or hybrid system specification ϕ from assumptions Γ. Presentations of dL with sequents $\Gamma \vdash \Delta$ where Δ is a set of succedents are available in the literature [135, 139, 142, 144].

sequence of deduction steps, where the axiom or proof rule used in each step is annotated to the left, as shown in the following illustrative proof outline:

Deduction

$$\uparrow \qquad \text{cut}\,\dfrac{\dfrac{\vdots}{\Gamma \vdash \psi} \qquad {[;]}\,\dfrac{\dfrac{\vdots}{\psi \vdash [\alpha][\beta]\phi}}{\Gamma, \psi \vdash [\alpha;\beta]\phi}}{\Gamma \vdash [\alpha;\beta]\phi}$$

Starting from the desired *conclusion* (below the rule bar), application of a proof rule, like the propositional cut rule, yields its *premises* (above the rule bar). When an implicational or equivalence axiom is used, like the [;] axiom $[\alpha;\beta]\phi \leftrightarrow [\alpha][\beta]\phi$ for sequential compositions, propositional sequent manipulation steps are omitted and the proof step is directly labeled with the axiom, giving the resulting premises accordingly [142]. Weakening steps that drop irrelevant assumptions are also omitted. Completed branches are marked with $*$ above the rule bar.

An axiom is *sound* iff all of its instances are valid and a proof rule is *sound* iff validity of all of its premises imply validity of its conclusion. Axioms and proof rules are *derivable* iff all of their instances can be deduced from sound dL axioms and proof rules. Soundness of the dL axiomatization ensures that all axioms and proof rules that are syntactically derived from the axiomatization are sound [142, 144] and can thereby be soundly used in subsequent deductions.

Arithmetic. First-order real arithmetic (over polynomial terms) is decidable [14, 197] so access to such a decision procedure is assumed. Proof steps are labeled with $\mathbb{R}$ whenever they follow as a substitution instance of a valid formula of first-order real arithmetic.

Propositional and First-Order Proof Rules. The following is an excerpt of sound propositional and first-order sequent calculus proof rules that are used in this thesis. Presentations of the entire calculus, e.g., with the $\vee$R rule(s), are available in the literature [135, 139, 142, 144]. In rules $\forall$L, $\exists$R, an arbitrary dL term e can be used to instantiate the respective quantifiers. In rules $\exists$L, $\forall$R, the Skolem variable y is fresh, i.e., does not occur free, in the conclusion:

$$\begin{array}{lll}
\neg\text{L}\ \dfrac{\Gamma \vdash \phi}{\Gamma, \neg\phi \vdash \mathit{false}} & \wedge\text{L}\ \dfrac{\Gamma, \phi_1, \phi_2 \vdash \psi}{\Gamma, \phi_1 \wedge \phi_2 \vdash \psi} & \rightarrow\text{L}\ \dfrac{\Gamma \vdash \phi_1 \quad \Gamma, \phi_2 \vdash \psi}{\Gamma, \phi_1 \rightarrow \phi_2 \vdash \psi} \\
\neg\text{R}\ \dfrac{\Gamma, \phi \vdash \mathit{false}}{\Gamma \vdash \neg\phi} & \wedge\text{R}\ \dfrac{\Gamma \vdash \phi_1 \quad \Gamma \vdash \phi_2}{\Gamma \vdash \phi_1 \wedge \phi_2} & \rightarrow\text{R}\ \dfrac{\Gamma, \phi_1 \vdash \phi_2}{\Gamma \vdash \phi_1 \rightarrow \phi_2} \\
\vee\text{L}\ \dfrac{\Gamma, \phi_1 \vdash \psi \quad \Gamma, \phi_2 \vdash \psi}{\Gamma, \phi_1 \vee \phi_2 \vdash \psi} & \forall\text{L}\ \dfrac{\Gamma, \phi(e) \vdash \psi}{\Gamma, \forall x\, \phi(x) \vdash \psi} & \exists\text{L}\ \dfrac{\Gamma, \phi(y) \vdash \psi}{\Gamma, \exists x\, \phi(x) \vdash \psi} \\
\text{cut}\ \dfrac{\Gamma \vdash \psi \quad \Gamma, \psi \vdash \phi}{\Gamma \vdash \phi} & \forall\text{R}\ \dfrac{\Gamma \vdash \phi(y)}{\Gamma \vdash \forall x\, \phi(x)} & \exists\text{R}\ \dfrac{\Gamma \vdash \phi(e)}{\Gamma \vdash \exists x\, \phi(x)}
\end{array}$$

2.3.2 Base Axioms and Proof Rules

This section presents the subset of dynamic logic and hybrid program axioms of dL [139, 142, 144] used in this thesis, except the axiomatization of differential equations which is deferred to the subsequent sections.

Theorem 2.1 (Base axioms and proof rules [139, 142])**.** *The following are sound axioms and proof rules of* dL*.*

$$\langle\cdot\rangle \quad \langle\alpha\rangle\phi \leftrightarrow \neg[\alpha]\neg\phi \qquad\qquad \text{G} \quad \frac{\vdash \phi}{\Gamma \vdash [\alpha]\phi}$$

$$\text{K} \quad [\alpha](\phi \to \psi) \to ([\alpha]\phi \to [\alpha]\psi) \qquad \text{V} \quad \phi \to [\alpha]\phi \quad \text{(no free variable of } \phi \text{ is bound by } \alpha)$$

The three axioms $\langle\cdot\rangle$, K, V and proof rule G are usual reasoning principles for dynamic logics [139] and they apply generally for any hybrid program α. Axiom $\langle\cdot\rangle$ expresses the duality between the diamond and box modalities, allowing conversion between the two with a double negation. Kripke axiom K is the modal modus ponens for postconditions of the box modality. Vacuous axiom V says if no free variable of ϕ is bound by hybrid program α, then the truth value of ϕ is also unchanged by a run of α. The free and bound variables of dL terms, formulas, and hybrid programs are defined as usual [142, 144]. The Gödel generalization rule G reduces proofs of $[\alpha]\phi$ to a proof of ϕ but must discard all assumptions in the antecedents Γ for soundness.

Theorem 2.2 (Hybrid program axioms [142])**.** *The following are sound axioms of* dL*.*

$$[:=] \quad [x := e]\phi(x) \leftrightarrow \phi(e) \quad (e \text{ free for } x \text{ in } \phi) \qquad [;] \quad [\alpha;\beta]\phi \leftrightarrow [\alpha][\beta]\phi$$

$$[?] \quad [?Q]\phi \leftrightarrow (Q \to \phi) \qquad\qquad [^*] \quad [\alpha^*]\phi \leftrightarrow \phi \land [\alpha][\alpha^*]\phi$$

$$[\cup] \quad [\alpha \cup \beta]\phi \leftrightarrow [\alpha]\phi \land [\beta]\phi \qquad\qquad \text{I} \quad [\alpha^*]\phi \leftrightarrow \phi \land [\alpha^*](\phi \to [\alpha]\phi)$$

Axioms $[:=]$, $[?]$, $[\cup]$, $[;]$, $[^*]$, I unfold box modalities of their respective hybrid programs according to their semantics (see Section 2.2.2). Assignment axiom $[:=]$ says to show postcondition $\phi(x)$ after an assignment to variable x, it suffices to prove $\phi(e)$ in the initial state; test axiom $[?]$ assumes that the test Q succeeded when proving postcondition ϕ; nondeterministic choice axiom $[\cup]$ proves the box modality for both branches α and β of the choice program $\alpha \cup \beta$ separately; sequential composition axiom $[;]$ says to prove postcondition ϕ for all runs of $\alpha;\beta$, one can equivalently prove the nested box modality postcondition $[\beta]\phi$ for all runs of program α; loop iteration axiom $[^*]$ says in order for postcondition ϕ to be true for all iterations of a loop α^*, ϕ must be true in the initial state before running the loop and it must remain true after running at least one iteration $[\alpha][\alpha^*]\phi$; and loop induction axiom I says to prove ϕ for all runs of loop α^*, ϕ must be true in the initial state and it must be inductively true along the loop $[\alpha^*](\phi \to [\alpha]\phi)$. The soundness proofs for Theorems 2.1 and 2.2 are available elsewhere [139, 142], along with an in-depth textbook exposition of the axioms [144].

Notational Conventions (Axioms). Key subformulas of (derived) equivalence or implication axioms are marked in blue to indicate that the equivalence or implication is typically applied to syntactically rewrite the marked subformula in proofs. For example, axiom $[;]$ is typically used to equivalently rewrite the key formula $[\alpha;\beta]\phi$ to $[\alpha][\beta]\phi$.

Derived Axioms and Proof Rules

The dynamic logic and hybrid program axioms of dL can be fruitfully combined to derive new axioms and proof rules that are useful in subsequent proofs. For example, the loop induction rule loop with a chosen loop invariant *Inv* (below) derives from induction axiom I using Gödel's generalization rule G [139, 144]; the monotonicity rule M[·], which derives from axiom K and rule G, strengthens the postcondition ψ of a box modality to ϕ provided that formula ϕ implies ψ; axiom [·]∧, which is also derived from axiom K [144], says to prove a conjunctive postcondition for a box modality it suffices to prove those postconditions for the box modality separately.

$$\text{loop}\ \dfrac{\Gamma \vdash \mathit{Inv} \quad \mathit{Inv} \vdash [\alpha]\,\mathit{Inv} \quad \mathit{Inv} \vdash \phi}{\Gamma \vdash [\alpha^*]\phi} \qquad \text{M}[\cdot]\ \dfrac{\phi \vdash \psi \quad \Gamma \vdash [\alpha]\phi}{\Gamma \vdash [\alpha]\psi} \qquad [\cdot]\wedge\ \ [\alpha](\phi \wedge \psi) \leftrightarrow [\alpha]\phi \wedge [\alpha]\psi$$

Axiom $\langle\cdot\rangle$ is useful for deriving diamond modality versions of dL axioms and proof rules. For example, the diamond Kripke axiom K$\langle\cdot\rangle$ (below) derives from K by dualizing its inner implication with $\langle\cdot\rangle$ [139] and the monotonicity rule for diamond modality postconditions M$\langle\cdot\rangle$ derives from axiom K$\langle\cdot\rangle$ by rule G. Similarly, dualizing $[^*]$ yields the diamond loop unfolding axiom $\langle^*\rangle$, which expresses that a loop can be repeated to reach postcondition ϕ iff either ϕ is already true in the initial state or it can be reached by repeating the loop for at least one iteration. Duality is exploited further for ODEs in Section 2.3.4 and in the refinement approach to ODE liveness introduced in Chapter 4.

$$\text{K}\langle\cdot\rangle\ \ [\alpha](\phi \rightarrow \psi) \rightarrow (\langle\alpha\rangle\phi \rightarrow \langle\alpha\rangle\psi) \qquad \text{M}\langle\cdot\rangle\ \dfrac{\phi \vdash \psi \quad \Gamma \vdash \langle\alpha\rangle\phi}{\Gamma \vdash \langle\alpha\rangle\psi}$$

$$\langle^*\rangle\ \ \langle\alpha^*\rangle\phi \leftrightarrow \phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi$$

Axiom V is particularly useful when working with *constant assumptions* [144]. If formula $\phi(y)$ is true initially and variable y is not bound in hybrid program α, then it remains true for all states reachable by running α because the value of y is unchanged by α and the truth value of $\phi(y)$ depends only on the (unchanged) value of its free variables y [142]. Axiom V proves this for box modalities in succedents and, by duality, for diamond modalities in antecedents. Conversely, if a constant assumption $\phi(y)$ is true in a final state reachable by hybrid program α, then it must already be true initially. This is shown formally by the derivation below which uses a classical case split with a cut on whether the formula $\phi(y)$ is already true initially:

$$\text{cut}\ \dfrac{\vee\text{L}\ \dfrac{\dfrac{*}{\Gamma, \phi(y), \langle\alpha\rangle(P \wedge \phi(y)) \vdash \phi(y)} \qquad \Gamma, \neg\phi(y), \langle\alpha\rangle(P \wedge \phi(y)) \vdash \phi(y)}{\Gamma, \phi(y) \vee \neg\phi(y), \langle\alpha\rangle(P \wedge \phi(y)) \vdash \phi(y)}}{\Gamma, \langle\alpha\rangle(P \wedge \phi(y)) \vdash \phi(y)}$$

The left premise closes trivially. For the right premise, a contradiction is derived with $\langle\cdot\rangle$ as follows, where the V, M[·] step uses the propositional tautology $\neg\phi(y) \rightarrow \neg(P \wedge \phi(y))$:

$$\langle\cdot\rangle, \neg\text{L}\ \dfrac{\text{V}, \text{M}[\cdot]\ \dfrac{*}{\neg\phi(y) \vdash [\alpha]\neg(P \wedge \phi(y))}}{\neg\phi(y), \langle\alpha\rangle(P \wedge \phi(y)) \vdash \mathit{false}}$$

Notational Conventions (Constant Assumptions). In the sequel, routine steps to maintain constant assumptions in contexts are omitted (or simply labeled with V for emphasis). For example, in the loop induction rule, all constant assumptions $\Gamma_c \subseteq \Gamma$ for α can be soundly kept across rule application of loop by adding those assumptions to the loop invariant and proving them with an internal use of V as follows [144]:

$$ {\scriptstyle\text{cut}}\dfrac{\Gamma \vdash \phi \qquad {\scriptstyle\text{loop}}\dfrac{{\scriptstyle [\cdot]\wedge,\,\mathrm{V}}\dfrac{\Gamma_c, \phi \vdash [\alpha]\phi}{\Gamma_c \wedge \phi \vdash [\alpha](\Gamma_c \wedge \phi)}}{\Gamma, \phi \vdash [\alpha^*]\phi}}{\Gamma \vdash [\alpha^*]\phi} $$

Additional constant contexts, like Γ_c above, are useful when working with assumptions on symbolic parameters, e.g., $v > 0$ to model a (constant) positive velocity unchanged by α.

2.3.3 Differentials and Lie Derivatives

ODEs $x' = f(x)$ precisely specify equations on time derivatives that their solutions must obey. The deduction of properties of ODE solutions from their differential equations therefore relates to the study of time derivatives of quantities mentioned in those properties. However, directly using time derivatives leads to numerous subtle sources of unsoundness because they are semantic objects that only make sense when a "time" axis even exists at all. Such a continuous time axis is furnished by the domain of definition of an ODE solution, but time derivatives are not otherwise well-defined in arbitrary contexts, e.g., in isolated states or across discrete transitions.

It is of utmost importance for soundness that, unlike time derivatives, differentials have a local semantics (2.2) that is well-defined in single states which enables their use in arbitrary contexts for sound syntactic manipulations [142]. The crucial differential lemma [142, Lem. 35] shows that, along a solution of the ODE $x' = f(x)$, the value of the differential term $(e)'$ coincides with the time derivative $\frac{\mathrm{d}}{\mathrm{d}t}$ of the value of term e. This relationship allows conclusions to be drawn about the differential equations directly from syntactic dL proofs involving differentials. The latter syntactic manipulation of differentials is achieved using the *differential axioms* of dL, which are given below. In axiom DE, $x' = f(x)$ is understood vectorially, i.e., x is a vector of variables $x_1, \ldots, x_n$, x' the corresponding vector of differential variables $x'_1, \ldots, x'_n$, and $f(x)$ a vector of terms $f_1(x), \ldots, f_n(x)$.

Theorem 2.3 (Differential axioms [142])**.** *The following are sound axioms of* dL*:*

$$\text{DE}\quad [x' = f(x) \,\&\, Q(x)]P(x, x') \leftrightarrow [x' = f(x) \,\&\, Q(x)][x' := f(x)]P(x, x')$$

$$\begin{array}{llll} c' & (c)' = 0 & x' & (x)' = x' \\ +' & (e + \tilde{e})' = (e)' + (\tilde{e})' & \cdot' & (e \cdot \tilde{e})' = (e)' \cdot \tilde{e} + e \cdot (\tilde{e})' \end{array}$$

The differential effect axiom (DE) says that the differential variables x' take on the values of the RHS along solutions to an ODE. This is expressed on its RHS with an assignment $x' := f(x)$ to the differential variable x'. Together, axioms DE, $[:=]$ allow replacing free occurrences of x' in the postcondition $P(x, x')$ yielding postcondition $P(x, f(x))$. However, proofs usually need to work with differentials of terms $(e)'$ rather than differential variables directly. This is where the differential axioms (c', x', $+'$, $\cdot'$) are used. Axiom c' says that the differential of a constant is 0, while axiom x' says the differential of a variable $(x)'$ is the corresponding differential variable x'. Axioms $+'$, $\cdot'$ are the sum and product rules of differentiation respectively. Soundness of these axioms allows differential terms to be rewritten equationally in all contexts, including in the postcondition of an ODE and within sub-terms. The differential axioms enable sound syntactic differentiation because differential terms $(e)'$ can be rewritten according to these equational axioms until no further differential sub-terms occur; any remaining differential variables are substituted away using DE, [:=] under an ODE. Such *exhaustive use of differential axioms* is simply labeled as ( )′ in proofs. The following example shows such a derivation concretely using a polynomial term from the example in Fig. 2.1:

Example 2.4 (Syntactic differentiation)**.** The following dL derivation syntactically differentiates the polynomial term $v^2 - u^2 + \frac{9}{2}$ along the ODE α_e from (2.1) for any comparison operator $\sim$:

$$\text{DE}\;\dfrac{(\ )'\;\dfrac{[:=]\;\dfrac{\mathrm{M}[\cdot],\mathbb{R}\;\dfrac{\vdash [\alpha_e]4uv + \frac{1}{2}(1-u^2-v^2)(v^2-u^2) \sim 0}{\vdash [\alpha_e]2v(u+\frac{v}{4}(1-u^2-v^2)) - 2u(-v+\frac{u}{4}(1-u^2-v^2)) \sim 0}}{\vdash [\alpha_e][u' {:=} -v+\frac{u}{4}(1-u^2-v^2)][v' {:=} u+\frac{v}{4}(1-u^2-v^2)]2vv' - 2uu' \sim 0}}{\vdash [\alpha_e][u' {:=} -v+\frac{u}{4}(1-u^2-v^2)][v' {:=} u+\frac{v}{4}(1-u^2-v^2)](v^2-u^2+\frac{9}{2})' \sim 0}}{\vdash [\alpha_e](v^2-u^2+\frac{9}{2})' \sim 0}$$

The first DE step makes available assignments on variables u', v' in the postcondition. The ( )′ step is then used to syntactically simplify $(v^2 - u^2 + \frac{9}{2})'$ yielding $2vv' - 2uu'$. A subsequent use of [:=] replaces the resulting differential variables u', v' with their respective RHS along the ODE. Finally, rules M[·], $\mathbb{R}$ are used to rearrange the calculated derivative arithmetically, which results in a (simplified) polynomial term. △

The exhaustive use of differential axioms means that all differential terms $(e)'$ under an ODE $x' = f(x)$ can be axiomatically rewritten to another term not mentioning differentials and differential variables. This resulting term is called the *Lie derivative* of term e along ODE $x' = f(x)$, succinctly written as follows:

$$\mathcal{L}_{f(x)}(e) \overset{\text{def}}{=} \sum_{i=1}^{n} \frac{\partial e}{\partial x_i} \cdot f_i(x)$$

Unlike (semantic) time derivatives, Lie derivatives can be written down syntactically in the syntactic term language. Like time derivatives though, Lie derivatives still depend on the ODE context in which they are used, so they do not give a compositional means of defining syntactic differentiation. The use of differentials in dL solves this problem by giving a compositional term semantics that is defined independently of any hybrid programs or formulas. Along an ODE $x' = f(x)$, the value of Lie derivative $\mathcal{L}_{f(x)}(e)$ coincides with that of the differential $(e)'$ and dL allows transformation between the two by proof with the differential axioms [142]. The Lie derivative $\mathcal{L}_{f(x)}(e)$ is written as $\dot{e}$ when the ODE $x' = f(x)$ is clear from the context. The

i-th *higher Lie derivative* $\overset{\bullet}{e}^{(i)}$ of term e along the ODE $x' = f(x)$ is defined by iterating the Lie derivation operator:

$$\overset{\bullet}{e}^{(0)} \overset{\text{def}}{=} e, \quad \overset{\bullet}{e}^{(i+1)} \overset{\text{def}}{=} \mathcal{L}_{f(x)}(\overset{\bullet}{e}^{(i)}), \quad \overset{\bullet}{e} \overset{\text{def}}{=} \overset{\bullet}{e}^{(1)}$$

2.3.4 Differential Equation Axiomatization

Having enabled meaningful syntactic differentiation with differentials, it remains to give axioms for working with differential equations. The following are dL axioms for differential equations [142, Figure 3], where each axiom DI, DC, DG progressively strengthens dL's deductive power for differential equation invariants [140]. All axioms are understood vectorially for differential equations as described in Theorem 2.3 for axiom DE.

Theorem 2.5 (Differential equation axiomatization [142]). *The following are sound axioms of* dL.[3] *In axiom DG, the* $\exists$ *quantifier can be replaced with a* $\forall$ *quantifier.*

$$\text{DI}_= \ [x' = f(x) \,\&\, Q](e)' = 0 \rightarrow \big([x' = f(x) \,\&\, Q]e = 0 \leftrightarrow (Q \rightarrow e = 0)\big)$$

$$\text{DI}_{\succcurlyeq} \ [x' = f(x) \,\&\, Q](e)' \geq 0 \rightarrow \big([x' = f(x) \,\&\, Q]e \succcurlyeq 0 \leftrightarrow (Q \rightarrow e \succcurlyeq 0)\big) \quad (\succcurlyeq \text{ either } \geq \text{ or } >)$$

$$\text{DC} \ [x' = f(x) \,\&\, Q]R \rightarrow \big([x' = f(x) \,\&\, Q]P \leftrightarrow [x' = f(x) \,\&\, Q \wedge R]P\big)$$

$$\text{DG} \ [x' = f(x) \,\&\, Q]P \leftrightarrow \exists y \, [x' = f(x), y' = a(x)y + b(x) \,\&\, Q]P$$

Differential invariants (DI) reduce questions about invariance of $e = 0, e \succcurlyeq 0$ (globally, along solutions of the ODE) to local questions about differentials. Only two instances ($\text{DI}_=$, $\text{DI}_{\succcurlyeq}$) of the more general DI axiom [142] are needed here. Axiom $\text{DI}_=$ says that the value of term e always stays zero if its differential $(e)'$ is always zero along the solution, while axiom $\text{DI}_{\succcurlyeq}$ says that e stays non-negative (or strictly positive) if its differential stays non-negative. Note that axiom $\text{DI}_{\succcurlyeq}$ only requires $(e)' \geq 0$ in its premise even for the $e > 0$ case. These axioms internalize the mean value theorem (see Corollary 2.7). *Differential cut* (DC) expresses that, if the system never leaves R while staying in Q (the outer assumption), then R may be additionally assumed in the domain constraint when proving the postcondition P (the RHS of the inner equivalence). Axiom DC increases dL's deductive power for invariants over DI [140] and the deductive power increases even further [140] with the *differential ghost* axiom (DG) which adds a *fresh* variable y to the system of ODEs for the sake of the proof. Since y is fresh, its initial value can be either existentially (DG) or universally ($\text{DG}_\forall$) quantified [142]. The syntactic restriction of DG is that the new ODE must be linear (or affine) in y, hence $a(x), b(x)$ are not allowed to mention y. This restriction prevents the newly added equation from unsoundly restricting the duration of existence for solutions to the differential equations [140], e.g., the (unsound) differential ghost $y' = y^2$ may cause finite-time (or early) blowup of solutions [204] (see Chapter 4). The added differential ghost variable y co-evolves along solutions and crucially enables the expression

[3] For simplicity, this thesis only uses ODE axioms, e.g., DC, DG, with postconditions and domains involving formulas without modal quantifiers, P, Q, R, following the notational conventions from Section 2.1.2.

of new (integral) relationships between variables along the differential equations. These new relationships are then used to syntactically deduce properties of interest in the original system.

To use axioms DI, DC, DG in proofs, additional differential equation axioms syntactically internalize temporal and first-order reasoning over differential equations.

Theorem 2.6 (Differential equation axiomatization (continued) [142])**.** *The following are sound axioms of* dL.

$$\mathrm{B}'\quad \langle x' = f(x) \,\&\, Q(x)\rangle \exists y P(x,y) \leftrightarrow \exists y\, \langle x' = f(x) \,\&\, Q(x)\rangle P(x,y) \quad (y \notin x)$$

$$\mathrm{DW}\quad [x' = f(x) \,\&\, Q]Q$$

$$\mathrm{DX}\quad [x' = f(x) \,\&\, Q]P \leftrightarrow (Q \rightarrow P \land [x' = f(x) \,\&\, Q]P) \qquad (x' \notin P, Q)$$

$$\mathrm{D}[;]\quad [x' = f(x) \,\&\, Q]P \leftrightarrow [x' = f(x) \,\&\, Q][x' = f(x) \,\&\, Q]P$$

$$\mathrm{DMP}\quad [x' = f(x) \,\&\, Q](Q \rightarrow R) \rightarrow ([x' = f(x) \,\&\, R]P \rightarrow [x' = f(x) \,\&\, Q]P)$$

Proof. The soundness of all axioms and proof rules in Theorem 2.6 are proved elsewhere [142], except DX, D[;], DMP, which are proved here since they are written differently elsewhere.

DX Let ω be an initial state. Classically, either $\omega \in [\![Q]\!]$ or $\omega \in [\![\neg Q]\!]$. If $\omega \in [\![Q]\!]$, then, propositionally, it suffices to assume $\omega \in [\![[x' = f(x) \,\&\, Q]P]\!]$ and show $\omega \in [\![P]\!]$. Since $\omega \in [\![Q]\!]$, there is a trivial solution $\varphi : [0,0] \rightarrow \mathbb{S}$ where $\varphi \models x' = f(x) \land Q$ and $\varphi(0) = \omega$ on $\{x'\}^\complement$. By assumption, $\varphi(0) \in [\![P]\!]$. Since $x' \notin P$, coincidence for formulas [142] implies $\omega \in [\![P]\!]$. Conversely, if $\omega \in [\![\neg Q]\!]$, then, propositionally, it suffices to show $\omega \in [\![[x' = f(x) \,\&\, Q]P]\!]$. The box modality is vacuous because, by definition, no solution $\varphi : [0,T] \rightarrow \mathbb{S}$ can exist for any $T \geq 0$ with $\varphi \models x' = f(x) \land Q$. Any such solution would require $\varphi(0) \in [\![Q]\!]$ by definition. However, because $x' \notin Q$, coincidence for formulas [142] with state ω gives $\omega \in [\![Q]\!]$, contradiction.

D[;] Let ω be an initial state and let $\varphi : [0,T) \rightarrow \mathbb{S}, 0 < T \leq \infty$ be the unique, right-maximal solution [33] to the ODE $x' = f(x)$ with initial value $\varphi(0) = \omega$. Unfolding the semantics of the outer box modality, the RHS of axiom D[;] is true in state ω iff for all times $0 \leq \tau < T$ such that $\varphi(\zeta) \in [\![Q]\!]$ for all $0 \leq \zeta \leq \tau$, the solution at time τ satisfies $\varphi(\tau) \in [\![[x' = f(x) \,\&\, Q]P]\!]$. Unfolding the semantics again, by uniqueness of ODE solutions [33], this means that for all times $\tau \leq t < T$, such that $\varphi(\zeta) \in [\![Q]\!]$ for all $\tau \leq \zeta \leq t$, the solution at time t satisfies $\varphi(t) \in [\![P]\!]$. Thus, the RHS is true in state ω iff for all times $0 \leq \tau < T$ such that $\varphi(\zeta) \in [\![Q]\!]$ for all $0 \leq \zeta \leq \tau$, the solution at time τ satisfies $\varphi(\tau) \in [\![P]\!]$, which is the unfolded semantics of the LHS of axiom D[;].

DMP Let ω be an initial state satisfying both formulas on the left of the implications in DMP, i.e., ① $\omega \in [\![[x' = f(x) \,\&\, Q](Q \rightarrow R)]\!]$ and ② $\omega \in [\![[x' = f(x) \,\&\, R]P]\!]$. Consider any solution $\varphi : [0,T] \rightarrow \mathbb{S}$ where $\varphi(0) = \omega$ on $\{x'\}^\complement$, and $\varphi \models x' = f(x) \land Q$. By definition, $\varphi(\zeta) \in [\![Q]\!]$ for all $\zeta \in [0,T]$, and so by ①, $\varphi(\zeta) \in [\![Q \rightarrow R]\!]$ for all $\zeta \in [0,T]$. Therefore, $\varphi(\zeta) \in [\![R]\!]$ for all $\zeta \in [0,T]$, and thus $\varphi \models x' = f(x) \land R$. By ②, $\varphi(T) \in [\![P]\!]$. □

The ODE Barcan axiom B$'$ specializes the Barcan axiom of dynamic logic [139] to ODEs in the diamond modality. It commutes an existential quantifier $\exists y$ with the diamond modality, where the variables y are required to be fresh in the ODE $x' = f(x)$ (i.e., $y \notin x$). The *differential weakening* axiom DW expresses that domain constraints are always obeyed along ODE solutions. The *differential skip* axiom DX expresses a reflexivity property of differential equation solutions. If domain constraint Q is false in an initial state ω, then the formula $[x' = f(x) \,\&\, Q]P$ is trivially true in ω because no solution of the ODE starting from ω stays in the domain constraint. Conversely, if Q is true in ω, then the postcondition P must already be true in ω because of the trivial solution of duration zero. The condition $x' \notin P, Q$ of axiom DX is met as P, Q are differential-free (Section 2.1). The *differential composition* axiom D[;] is a transitivity property of differential equation solutions which says that any state reachable from two sequential runs of the same ODE is reachable in a single run of that ODE. Axiom DMP is the modus ponens for domain constraints of ODEs which underlies differential cuts DC [142].

Derived Axioms and Proof Rules

Similar to Section 2.3.2, additional ODE axioms and proof rules can be derived from the ODE axioms of Theorems 2.5 and 2.6 for use in subsequent deductive proofs. Rule dI$_\succcurlyeq$ (below) derives by combining DI$_\succcurlyeq$ with the differential axioms $(\ )'$ to provably transform differentials to Lie derivatives in its premise. Rules dC, dW derive from their underlying axioms DC, DW respectively [144].

$$\text{dI}_\succcurlyeq\ \frac{Q \vdash \mathcal{L}_{f(x)}(e) \geq \mathcal{L}_{f(x)}(\tilde{e})}{\Gamma, e \succcurlyeq \tilde{e} \vdash [x' = f(x) \,\&\, Q] e \succcurlyeq \tilde{e}} \qquad \text{(where } \succcurlyeq \text{ is either } \geq \text{ or } >\text{)}$$

$$\text{dC}\ \frac{\Gamma \vdash [x' = f(x) \,\&\, Q]R \quad \Gamma \vdash [x' = f(x) \,\&\, Q \land R]P}{\Gamma \vdash [x' = f(x) \,\&\, Q]P}$$

$$\text{dW}\ \frac{Q \vdash P}{\Gamma \vdash [x' = f(x) \,\&\, Q]P}$$

Axiom $\langle\cdot\rangle$ also yields dual diamond readings for the ODE axioms and proof rules. For example, the DI$_\succcurlyeq$ axiom internalizes a version of the mean value theorem [204, Appendix B.I]. This is shown in derived axiom MVT below which says that, if the value of term e is non-negative initially and eventually becomes negative along the ODE $x' = f(x) \,\&\, Q$, then its differential $(e)'$ must be negative somewhere along the solution to that ODE.

Corollary 2.7 (Mean value theorem). *The mean value theorem axiom MVT derives from DI$_\succcurlyeq$:*

$$\text{MVT}\ \ e \geq 0 \land \langle x' = f(x) \,\&\, Q\rangle e < 0 \rightarrow \langle x' = f(x) \,\&\, Q\rangle (e)' < 0$$

Proof. The derivation takes contrapositives (dualizing with $\langle\cdot\rangle$) before DI$_\succcurlyeq$ finishes it.

$$\dfrac{\dfrac{*}{e \geq 0, [x' = f(x) \,\&\, Q](e)' \geq 0 \vdash [x' = f(x) \,\&\, Q]e \geq 0}\,\text{DI}_\succcurlyeq}{e \geq 0, \langle x' = f(x) \,\&\, Q\rangle e < 0 \vdash \langle x' = f(x) \,\&\, Q\rangle (e)' < 0}\,\langle\cdot\rangle, \neg\text{L}, \neg\text{R} \qquad \square$$

Other useful proof rules are derived by combining the differential equation axioms with the base axioms and proof rules. For example, the following monotonicity rules (for both box and diamond modalities) are derived from axiom K and dW. Compared to the monotonicity rules M[$\cdot$], M$\langle\cdot\rangle$ which apply for general hybrid programs α, the ODE monotonicity rules M[$'$], M$\langle'\rangle$ additionally assume domain constraint Q when strengthening the postcondition from P to R for an ODE $x' = f(x) \,\&\, Q$:

$$\text{M}['] \; \frac{Q, R \vdash P \quad \Gamma \vdash [x' = f(x) \,\&\, Q]R}{\Gamma \vdash [x' = f(x) \,\&\, Q]P} \qquad \text{M}\langle'\rangle \; \frac{Q, R \vdash P \quad \Gamma \vdash \langle x' = f(x) \,\&\, Q\rangle R}{\Gamma \vdash \langle x' = f(x) \,\&\, Q\rangle P}$$

The "$\leftarrow$" direction of DX allows domain constraint Q to be assumed true initially when proving $[x' = f(x) \,\&\, Q]P$ (shown below, on the left). The "$\rightarrow$" direction has the following equivalent contrapositive reading using $\langle\cdot\rangle$ and propositional simplification: $Q \land P \rightarrow \langle x' = f(x) \,\&\, Q\rangle P$, i.e., if the domain constraint Q and postcondition P are both true initially, then $\langle x' = f(x) \,\&\, Q\rangle P$ is true because of the trivial solution of duration zero. When proving the liveness property $\langle x' = f(x) \,\&\, Q\rangle P$, one can therefore always additionally assume $\neg(Q \land P)$ because, by DX, $\langle\cdot\rangle$, there is nothing to prove otherwise (shown below, on the right).

$$\text{DX} \frac{\Gamma, Q \vdash [x' = f(x) \,\&\, Q]P}{\Gamma \vdash [x' = f(x) \,\&\, Q]P} \qquad \text{DX, } \langle\cdot\rangle \frac{\Gamma, \neg(Q \land P) \vdash \langle x' = f(x) \,\&\, Q\rangle P}{\Gamma \vdash \langle x' = f(x) \,\&\, Q\rangle P}$$

Chapter 3

Safety and Invariance for Ordinary Differential Equations

This chapter begins the study of deductive verification for ordinary differential equations (ODEs) by examining dL proofs of ODE safety and invariance, and their proof theory. Classically, differential equations are studied by analyzing their solutions, which is at odds with the fact that those solutions are often much more complicated than the differential equations themselves. This stark difference between the simple local description as differential equations and the complex global behavior exhibited by their solutions is fundamental to the descriptive power of differential equations. Poincaré's qualitative study of differential equations [152] calls for the exploitation of this difference by deducing properties of solutions *directly from the differential equations*. This chapter completes an important step in Poincaré's enterprise by identifying the *logical foundations for proving invariance properties of differential equations* described by Noetherian functions [12, 56, 57, 201]. These invariance proof principles further serve as foundational building blocks for subsequent thesis chapters, e.g., as stepping stones in refinement proofs of liveness (Chapter 4) and as powerful black box proof steps for ODE safety (sub-)questions arising in stability proofs (Chapters 5 and 6). Consequently, the generality of all results presented in this chapter yields corresponding generality in subsequent thesis chapters.

3.1 Introduction

The ODE safety specification $\Gamma \vdash [x' = f(x) \,\&\, Q]\phi$ says that, from initial states satisfying assumptions Γ, all states reached by following the ODE $x' = f(x) \,\&\, Q$ from those states satisfy the safety postcondition ϕ. These ODE safety questions may arise, for example, when proving that a continuous control law always keeps a system within safe bounds throughout its operation, or as sub-questions within a larger hybrid system safety proof [144]. They can also come from the system designer's insights into physically meaningful quantities of the continuous system, for example, conservation of an energy quantity E along an ODE is expressed by the formula $[x' = f(x)]E = E_0$, which says that E always stays at its initial value E_0 along the system's continuous evolution. Such a conserved quantity (if proved) may, in turn, be used as part of safety proofs or to provide physical insights into a given system.

Unfortunately, it is rarely the case that ODEs $x' = f(x)$ have closed form solutions that can be mathematically analyzed to prove safety properties of its trajectories directly (recall Section 2.1.3, page 16). Instead, the premier technique for proving ODE safety is to find a suitable *invariant* P of the ODE such that: *i)* it contains the initial states $\Gamma \vdash P$, *ii)* it is safe $P \vdash \phi$, and *iii)* solutions of the ODEs cannot escape it $P \vdash [x' = f(x) \,\&\, Q]P$. The use of an invariant P is illustrated by the following partial derivation with cut (to prove *i*) and M[·] (to prove *ii*):

$$\text{cut}\,\dfrac{\Gamma \vdash P \qquad \text{M}[\cdot]\,\dfrac{\dfrac{\cdots}{P \vdash [x' = f(x) \,\&\, Q]P} \qquad P \vdash \phi}{P \vdash [x' = f(x) \,\&\, Q]\phi}}{\Gamma \vdash [x' = f(x) \,\&\, Q]\phi}$$

This chapter shows that one can *always* fill in the remaining “$\cdots$” steps in the derivation because the invariance question for *any* formula P and ODE $x' = f(x) \,\&\, Q$ is provably equivalent to an arithmetic question within dL. Thus, the remaining practical challenge for proofs of ODE safety is to find suitable and succinct ODE invariants P; the invariance of any candidate formula P can be *proved* or *disproved* by checking validity of the resulting arithmetic question after applying the derived dL equivalence. In fact, a stronger result is proved for equational safety properties, such as the energy conservation formula $[x' = f(x)]E = E_0$ above: dL completely reduces such questions to arithmetic *without* the need for an intermediary invariant. These completeness results are proved more generally for an *extended* dL term language, which enables the use of extended (non-polynomial) functions in models and proofs of continuous and hybrid systems while retaining sound and complete dL ODE invariance reasoning.

Formally, this chapter presents a *differential equation invariance axiomatization* based on dL. For extended term languages (and ODEs) meeting three extended term conditions, the chapter proves the following results:

1. *All* analytic invariants, i.e., finite conjunctions and disjunctions of equations between extended terms, are provable using only the three dL ODE axioms DI, DC, DG. This result generalizes to analytic hybrid programs with dL's hybrid program axioms.

2. With axioms internalizing the existence and uniqueness theorems for solutions of differential equations, all *local progress* properties of ODEs are provable for all semianalytic formulas, i.e., propositional combinations of inequalities between extended terms.

3. With a real induction axiom that reduces invariance to local progress, the dL calculus is complete for proving *all* semianalytic invariants of differential equations.

4. These are *axiomatic completeness* results: all (semi)analytic invariance and local progress questions are provably *equivalent* in dL to questions about the underlying arithmetic. This equivalence also yields disproofs when the resulting arithmetic questions are refuted.

These results are proved constructively with syntactic derivations for each equivalence, thus yielding practical and purely logical proof-producing procedures for reducing ODE invariance questions to arithmetical questions in dL. The axiomatic approach crucially enables these contributions because the syntactic dL axioms internalize basic properties of ODEs and thus remain sound and complete for *all* extended term languages meeting the extended term conditions.

Furthermore, the identification of a parsimonious yet complete ODE axiomatization provides the best of both worlds: *parsimony* minimizes effort required in implementation and verification of the proof calculus while *completeness* guarantees that all ODE invariance reasoning is possible using only syntactic proofs from the foundational axioms.

The most subtle step in the proofs is the construction of suitable differential ghosts that simplify the analysis as a function of both the differential equations and desired invariant. Just as discrete ghosts can make a program logic relatively complete [127], differential ghosts achieve completeness for algebraic (and analytic) invariants in dL.

5. *Scalar* differential ghosts DG are used to derive (complete) reasoning principles for analytic invariants, Darboux (in)equalities, and *barrier certificates* [155] for ODEs.

The result (5) is significant for implementation and proof-theoretical purposes because it uncovers classes of invariants that are provable efficiently using only a *constant* number of differential ghost reasoning steps. In practice, these classes cover many automatic invariant generation techniques, such as those used by the Pegasus tool [180], which enables an efficient combination of invariant generation and sound checking in KeYmaera X [54].

Finally, since Noetherian functions from real analytic geometry [12, 56, 57, 201] generate Noetherian rings closed under partial derivatives, they meet all of the extended term conditions and thus provide an ideal setting for extending dL's term language. Many functions of practical interest for modeling hybrid systems are Noetherian, e.g., the real exponential and trigonometric functions, which are implicitly definable in dL [137, 139] but do not come with effective reasoning principles.[1] Making them first-class members of the term language enables their explicit use in hybrid systems models and proofs, especially in descriptions of ODE invariants.

6. Noetherian functions are shown to meet the extended term conditions. Any such extension automatically inherits all of the aforementioned completeness results.

Thanks to the common dL logical foundation of this thesis, result (6) gives license to all subsequent chapters to freely use extended Noetherian function terms, with the assurance that the underlying ODE invariance reasoning can be handled completely and compositionally, in isolation from other continuous and hybrid systems reasoning [139].

Contribution. *The material for this chapter is drawn from Platzer and Tan [148, 149].*

3.2 Differential Dynamic Logic with Extended Terms

This section introduces a generic extended term language for dL which enables models and proofs of hybrid systems featuring non-polynomial terms. Of course, such a syntactic extension cannot be completely arbitrary, e.g., adding functions whose interpretations are nowhere differentiable would fundamentally break the enterprise of studying ODEs directly by their local behavior. These unsuitable syntactic extensions are ruled out by a set of extended term conditions, which

[1] The relative decidability theorem for dL [139, Theorem 11] needs either an oracle for (continuous) differential equation properties or an oracle for discrete program properties.

are developed and motivated along the way, with a summary in Section 3.2.4. The class of Noetherian functions, which meets all the extended term conditions, is introduced in Section 3.7.

3.2.1 Syntax

The dL term language is *extended* with a finite number of new k-ary fixed function symbols, $h \in \{h_1, \dots, h_r\}$, with fixed interpretations.

$$e, \tilde{e} ::= x \mid c \mid e + \tilde{e} \mid e \cdot \tilde{e} \mid \underline{h(e_1, \dots, e_k)} \mid (e)'$$

As a running example of such an extended term language, consider the unary function symbols $\exp, \sin, \cos$ which are always interpreted as the real exponential and trigonometric functions respectively:

$$e, \tilde{e} ::= x \mid c \mid e + \tilde{e} \mid e \cdot \tilde{e} \mid \exp(e) \mid \sin(e) \mid \cos(e) \mid (e)' \tag{3.1}$$

Notational Conventions (Extended Term Language). Polynomial terms are useful as familiar illustrative examples and they also enjoy special properties not necessarily shared by extended term languages. As usual (Section 2.1.1), the notation p, q is reserved for polynomial terms, with dependencies $p(x), q(x)$ added when necessary. Formulas over extended term languages that do not contain the first-order quantifiers nor the modal connectives are called *semianalytic* formulas and are written as P, Q. The word "analytic" refers to the (semantic) real analyticity [94] of terms extended with Noetherian functions in Section 3.7. Every semianalytic formula can be normalized to one that is formed from only conjunctions and disjunctions of atomic comparison formulas. Formulas P, Q that are formed from only conjunctions and disjunctions of equalities are called *analytic* formulas. When all atomic comparisons in (semi)analytic formulas are restricted to only occur between polynomial terms $p \sim q$, the resulting formulas are also known as *(semi)algebraic* formulas [14]. Compared to Section 2.1.2 the notational convention for semianalytic formulas P, Q disallows first-order quantification. No expressiveness is lost by disallowing first-order quantifiers for polynomial terms because the first-order theory of the reals with polynomial terms (and with quantifiers) admits quantifier elimination [14, 197], so every first-order formula of real arithmetic (over polynomial terms) is provably equivalent to a quantifier-free semialgebraic formula. Unfortunately, quantifier elimination is impossible even for simple term language extensions like the exponential function [202].

3.2.2 Semantics

The dL term semantics $\omega[\![e]\!] \in \mathbb{R}$ is extended as follows, where the semantics of each k-ary fixed function symbol h is given by a corresponding real-valued function $h : \mathbb{R}^k \to \mathbb{R}$ (using the same symbol h for the LHS syntactic function symbol and its RHS semantic interpretation by a slight abuse of notation):

$$\omega[\![h(e_1, \dots, e_k)]\!] = h(\omega[\![e_1]\!], \dots, \omega[\![e_k]\!])$$

There are two subtleties to highlight. First, the real-valued interpretations h are required to be defined on the domain $\mathbb{R}^k$ so that the term semantics are well-defined in all states. It is possible to extend dL with terms that are only defined within an open domain of definition rather than the entire real domain [21] and this would allow, e.g., rational functions to be added to the term language. Such an extension will not be pursued in this thesis although the Noetherian functions from Section 3.7 and Proposition 3.34 give an implicit way of working with quotients of extended terms. Second, the semantics of differentials implicitly requires that the partial derivatives $\frac{\partial \omega[\![e]\!]}{\partial x}$ exist for any term e. In fact, partial derivatives of any order for the semantics of any term must exist because their differentials (which provably reduce to differential-free terms by Section 2.3.3), in turn, have differentials that must also have well-defined semantics. Following the dL interpretation of function symbols [142], it suffices to require that the fixed function symbols h are interpreted as smooth C^∞ functions, i.e., $h : \mathbb{R}^k \to \mathbb{R}$ with partial derivatives of any order. Since the C^∞ functions are closed under addition, multiplication and function composition, the resulting term semantics are also smooth [142], as required.

3.2.3 Axiomatics

The soundness proofs for the axiomatization of dL in Section 2.3 carry over unchanged for extended term languages because fixed function symbols h are interpreted as smooth C^∞ functions [142]. The subtleties are in arithmetic reasoning using rule $\mathbb{R}$ and the introduction of differential axioms for the fixed function symbols.

Arithmetic. Even for the extended term language (3.1) with trigonometric functions, arithmetic questions are already undecidable [162]. Therefore, special care must be taken to distinguish first-order properties of the real closed fields, i.e., those described by (semi)algebraic formulas [14], from those properties described by (semi)analytic formulas, as illustrated next.

Example 3.1 (Proving with $\mathbb{R}$)**.** Consider the following two proofs in extended term language (3.1) with the real exponential function exp. The left sequent proves (as a substitution instance) by $\mathbb{R}$ because the negation of a real number is its additive inverse. In contrast, the right sequent does not prove by $\mathbb{R}$ alone (indicated by the subscript on rule $\mathbb{R}_{\exp}$) because it uses the fact that the real exponential function is strictly positive.

$$\mathbb{R}\frac{*}{\vdash \exp(x) + (-\exp(x)) = 0} \qquad\qquad \mathbb{R}_{\exp}\frac{*}{\vdash \exp(x) > 0}$$

An alternative understanding is that rule $\mathbb{R}$ can be used to conclude valid arithmetic properties that follow *only* from first-order properties of the real closed fields [14, 197]. △

Differential Axioms. Section 2.3.3 showed how differential terms under an ODE $x' = f(x)$ can be axiomatically rewritten to the Lie derivative using the differential axioms $(\,)'$. This is the case for differentials of polynomial terms $(p)'$, but differential axioms are still needed for the fixed function symbols. Consider the case of a unary fixed function symbol h which is semantically interpreted as the function $h(y) : \mathbb{R} \to \mathbb{R}$. Expanding the semantics of term $(h(e))'$ and applying the chain rule:

$$\omega[\![(h(e))']\!] = \sum_{x \in \mathbb{V}} \omega(x') \frac{\partial \omega[\![h(e)]\!]}{\partial x} = \sum_{x \in \mathbb{V}} \omega(x') \frac{\partial h}{\partial y}(\omega[\![e]\!]) \frac{\partial \omega[\![e]\!]}{\partial x}$$

$$= \frac{\partial h}{\partial y}(\omega[\![e]\!]) \sum_{x \in \mathbb{V}} \omega(x') \frac{\partial \omega[\![e]\!]}{\partial x} = \frac{\partial h}{\partial y}(\omega[\![e]\!]) \omega[\![(e)']\!]$$

The RHS product between $\frac{\partial h}{\partial y}(\omega[\![e]\!])$ and $\omega[\![(e)']\!]$ can be represented syntactically provided that the partial derivative of h with respect to its argument y is representable as a term. Assume (suggestively) that such a term is written as $\frac{\partial h}{\partial y}(e)$. The easiest case is to think of $\frac{\partial h}{\partial y}$ as another unary fixed function symbol and $\frac{\partial h}{\partial y}(e)$ as function application, hence the suggestive notation. This is not strictly necessary: $\frac{\partial h}{\partial y}$ can be another term that mentions variable y free, in which case $\frac{\partial h}{\partial y}(e)$ corresponds to substituting e for y in that term. The differential axiom h' for fixed function symbol h uses a product of the (syntactic) partial derivative $\frac{\partial h}{\partial y}$ and differential $(e)'$ of e:

$$h' \quad (h(e))' = \frac{\partial h}{\partial y}(e) \cdot (e)'$$

Example 3.2 (Unary extended differential axioms)**.** For the extended term language (3.1), the extended terms for the partial derivatives are as usual from calculus:

$$\frac{\partial \exp(y)}{\partial y} = \exp(y) \qquad \frac{\partial \sin(y)}{\partial y} = \cos(y) \qquad \frac{\partial \cos(y)}{\partial y} = -\sin(y)$$

Following the axiom schema h', the differential axioms for these fixed function symbols are:

$$exp' \quad (\exp(e))' = \exp(e) \cdot (e)' \quad sin' \quad (\sin(e))' = \cos(e) \cdot (e)' \quad cos' \quad (\cos(e))' = -\sin(e) \cdot (e)'$$

Axioms sin', cos' illustrate a syntactic subtlety: fixed function symbols must be introduced in a syntactically complete way with respect to differentials. The unary function symbol sin for the trigonometric sine function cannot be added without also adding one for the cosine function because there would otherwise be no way to express the differential of sin syntactically.[2] △

The following lemma generalizes the syntactic representation condition from the above example and gives sound differential axioms for k-ary fixed function symbols.

Lemma 3.3 (Extended differential axioms)**.** *Let the k-ary fixed function symbol h be semantically interpreted as a differentiable function $h : \mathbb{R}^k \to \mathbb{R}$. Suppose its partial derivative $\frac{\partial h}{\partial y_i}(y_1, \dots, y_k)$ at $y_1, \dots, y_k$ is syntactically represented by the term $\frac{\partial h}{\partial y_i}$ for each i such that $\omega[\![\frac{\partial h}{\partial y_i}(y_1, \dots, y_k)]\!] = \frac{\partial h}{\partial y_i}(\omega(y_1), \dots, \omega(y_k))$ for all states ω. Then the differential axiom schema h' for h is sound:*

$$h' \quad (h(e_1, \dots, e_k))' = \sum_{i=1}^{k} \frac{\partial h}{\partial y_i}(e_1, \dots, e_k) \cdot (e_i)'$$

Proof. The terms $\frac{\partial h}{\partial y_i}(e_1, \dots, e_k)$ appearing on the RHS of axiom h' are understood as (syntactic) function application of $\frac{\partial h}{\partial y_i}$ to the arguments $e_1, \dots, e_k$. Soundness of this axiom follows from the (multivariate) chain rule and the semantics of differential terms. For any given state ω:

$$\omega[\![(h(e_1, \dots, e_k))']\!] = \sum_{x \in \mathbb{V}} \omega(x') \frac{\partial \omega[\![h(e_1, \dots, e_k)]\!]}{\partial x} = \sum_{x \in \mathbb{V}} \omega(x') \sum_{i=1}^{k} \frac{\partial h}{\partial y_i}(\omega[\![e_1]\!], \dots, \omega[\![e_k]\!]) \frac{\partial \omega[\![e_i]\!]}{\partial x}$$

[2] Technically, π could be added and $\cos(x)$ encoded as $\sin(x + \pi)$ but that also requires another 0-ary function symbol π.

$$= \sum_{i=1}^{k} \frac{\partial h}{\partial y_i}(\omega[\![e_1]\!], \dots, \omega[\![e_k]\!]) \sum_{x \in \mathbb{V}} \omega(x') \frac{\partial \omega[\![e_i]\!]}{\partial x} = \sum_{i=1}^{k} \frac{\partial h}{\partial y_i}(\omega[\![e_1]\!], \dots, \omega[\![e_k]\!]) \omega[\![(e_i)']\!]$$

$$= \sum_{i=1}^{k} \omega[\![\frac{\partial h}{\partial y_i}(e_1, \dots, e_k)]\!] \omega[\![(e_i)']\!] = \omega[\![\sum_{i=1}^{k} \frac{\partial h}{\partial y_i}(e_1, \dots, e_k)(e_i)']\!]$$

The penultimate step uses the fact that all partial derivatives are syntactically represented in the term language to replace a semantic function application with its syntactic representation. □

Lemma 3.3 allows derivations to freely replace differential terms of extended term languages $(e)'$ under an ODE $x' = f(x)$ with their corresponding Lie derivative $\mathcal{L}_{f(x)}(e)$ by proof with the differential axioms [142].

3.2.4 Extended Term Conditions

Two natural conditions on the fixed function symbols $h \in \{h_1, \dots, h_r\}$ and their semantics have been uncovered thus far. For this chapter's completeness results, a third condition is needed. All three *extended term conditions* are assumed throughout this chapter:

(S) *Smoothness.* All fixed function symbols $h \in \{h_1, \dots h_r\}$ in the extended term language are interpreted as smooth C^∞ functions $h : \mathbb{R}^k \to \mathbb{R}$.

(P) *Syntactic partial derivatives.* Each partial derivative $\frac{\partial h}{\partial y_i}$ of $h(y_1, \dots, y_k)$ has a syntactic representation in the extended term language in the sense of Lemma 3.3.

(R) *Computable differential radicals.* The extended term language has computable differential radicals, i.e., for each extended term e and ODE $x' = f(x)$ with extended terms in its RHS $f(x)$, there must computably exist a natural number $N \geq 1$ and N extended terms $(g_0, g_1, \dots, g_{N-1})$ such that the higher Lie derivatives of e along $x' = f(x)$ provably satisfy the following *differential radical identity* [63]:

$$\overset{\bullet}{e}^{(N)} = \sum_{i=0}^{N-1} g_i \overset{\bullet}{e}^{(i)} \tag{3.2}$$

Condition (S) ensures that the semantics are well-defined, while conditions (P) and (R) enable (complete) *syntactic* analysis of differential equations invariance by their local (differential) behavior. The C^∞ smoothness required by (S) is subtly weaker than *real analyticity* [94]. This chapter often gives brief but intuitive (semantic) explanations of results and explicitly indicates when those arguments only apply *in the real analytic setting*. None of the actual proofs given in this chapter require real analyticity. Condition (R) requires an algorithm that computes and proves the identity (3.2). This identity is crucially used for completeness in Section 3.4 and Section 3.6, where it is also motivated logically. Intuitively, it yields a finiteness property on the number of Lie derivatives that need to be analyzed for any given term e and ODE, i.e., from identity (3.2), the first $N-1$ Lie derivatives will turn out to suffice for completely determining the local behavior of extended term e along the ODE $x' = f(x)$. All three extended term conditions are met by the polynomial term language without extensions.

Proposition 3.4. *Polynomial term languages satisfy the extended term conditions.*

Proof Sketch. A full proof is omitted because this is a corollary of a later result (Theorem 3.37). Briefly, conditions (S) and (P) are met because polynomial functions are smooth (even real analytic) and the polynomials are closed under partial derivatives. Condition (R) generalizes different flavors of results that have been proved in the literature [63, 103, 125]. The proofs rely on the fact that polynomials form a Noetherian ring [24] so that the ascending chain of ideals[3] formed by successive (polynomial) Lie derivatives stabilizes. The polynomial identity (3.2) is computable by successive ideal membership checks [63, 64, 103]. Moreover, it is a formula of real arithmetic and can therefore always be proved by the rule $\mathbb{R}$ for decidable real arithmetic. □

It is less straightforward to show that an extended term language like (3.1) meets these conditions. Indeed, even the simple language extension (3.1) already features exponential rings which are not Noetherian [198, Remark 1.4.2] and undecidable arithmetic over the trigonometric functions [162]. In the interest of a general presentation, the question of how to determine if a candidate term language extension $\{h_1, \ldots, h_r\}$ meets the extended term conditions is deferred to Section 3.7. Until then, the only assumption about the extended term language is that it satisfies those three conditions. This suffices for the chapter's completeness results, which the next section begins to show.

3.3 Darboux Invariants

This section exploits differential ghosts for proving an important class of invariance properties. These are called *Darboux invariants* because they are inspired by Darboux polynomials [41]. The derived proof rule for Darboux equalities corresponds to the case $N = 1$ in the differential radical identity (3.2), while the subsequent rule for Darboux *in*equalities is a crucial step for the completeness result in Section 3.4. Their derivations also show how analytic and geometric notions from the theory of differential equations, such as Darboux polynomials [41] and Grönwall's lemma [70, 204, §29.VI] can be internalized syntactically with differential ghost arguments without extension to any dL axiom.

3.3.1 Darboux Equalities

Assume that the extended term e satisfies the differential radical identity (3.2) with $N = 1$ and extended term cofactor g, i.e., $\dot{e} = ge$. Taking Lie derivatives on both sides gives:

$$\dot{e}^{(2)} = \mathcal{L}_{f(x)}(\dot{e}) = \mathcal{L}_{f(x)}(ge) = \dot{g}e + g\dot{e} = (\dot{g} + g^2)e$$

By repeatedly taking Lie derivatives, note that *all* higher Lie derivatives of e can be written as a product between e and some extended term cofactor. Now, consider an initial state ω where e evaluates to $\omega[\![e]\!] = 0$, then:

$$\omega[\![\dot{e}]\!] = \omega[\![ge]\!] = \omega[\![g]\!] \cdot \omega[\![e]\!] = 0$$

[3] The *ideal* [14] generated by polynomials $p_1, \ldots, p_s \in \mathbb{R}[x]$ is the set of all their linear combination with polynomial cofactors $g_i \in \mathbb{R}[x]$, denoted by $(p_1, \ldots, p_s) \overset{\text{def}}{=} \{\sum_{i=1}^{s} g_i p_i \ : \ g_i \in \mathbb{R}[x]\}$.

Similarly, because every higher Lie derivative is a product with e, all of them evaluate simultaneously to 0 in state ω. Thus, *in the real analytic setting*, $e = 0$ stays invariant along solutions to the ODE starting at ω because all its derivatives are 0. This motivates the following proof rule for invariance of $e = 0$:

$$\text{dbx}\ \frac{Q \vdash \dot{e} = ge}{e = 0 \vdash [x' = f(x) \,\&\, Q]e = 0}$$

Rule dbx derives using differential ghosts, which provides a first hint at their deductive power for equational invariants. A special case of dbx proves invariance for *Darboux polynomials*, which are polynomials p satisfying the polynomial identity $\dot{p} = gp$ for some polynomial cofactor g. These polynomials are of significant interest in the study of (polynomial) ODEs [41] and invariant generation for continuous and hybrid systems [169, 180]. In Section 3.4, dbx is generalized vectorially to yield proofs of *all* analytic invariants with differential ghosts. Although the rule can be derived from DG directly, this section follows a detour through a proof rule for Darboux *in*equalities instead, which is crucially used for this vectorial generalization.

3.3.2 Darboux Inequalities

Assume that the extended term e satisfies the Darboux *inequality* $\dot{e} \geq ge$ for some extended term cofactor g. Semantically, in an initial state ω where $\omega[\![e]\!] \geq 0$, Grönwall's lemma [70, 204, §29.VI] implies that $e \geq 0$ stays invariant along solutions starting at ω because the semantic value of extended term e is bounded below by a (typically decaying) non-negative exponential solution of the non-autonomous linear differential equation $e' = g(t)e$ for the variable e. Here, $g(t)$ is the time-dependent function corresponding to the value of term g evaluated along the solution to the differential equations $x' = f(x)$ from ω, see Fig. 3.1a for an illustration. Indeed, if e satisfies the Darboux equality $\dot{e} = ge$ with cofactor g, then it satisfies both Darboux inequalities $\dot{e} \geq ge$ and $\dot{e} \leq ge$, giving an alternative semantic argument for the invariance of $e = 0$ in rule dbx.

Differentials and Lie derivatives along differential equations $x' = f(x)$ provably coincide (Section 2.3.3), so axiomatic Darboux inequalities assume $[x' = f(x) \,\&\, Q](e)' \geq ge$ and Darboux equalities assume $[x' = f(x) \,\&\, Q](e)' = ge$ instead of $\dot{e} \geq ge$ and $\dot{e} = ge$, respectively. The use of differentials in the axioms yield particularly efficient proofs within dL's uniform substitution calculus [142] because they derive once-and-for-all, independently of the ODE $x' = f(x)$. Subsequently substituting [142] for specific ODE instances means that only the final Lie derivative calculation steps DE, $(\ )'$, $[:=]$ are needed for each concrete derived instance.

Lemma 3.5 (Darboux (in)equalities are diff. ghosts). *The Darboux equality DBX and Darboux inequality DBX$_{\succcurlyeq}$ axioms derive from DG (and DI, DC) for any extended term cofactor g.*

$$\text{DBX}\ \ [x' = f(x) \,\&\, Q](e)' = ge \rightarrow (e = 0 \rightarrow [x' = f(x) \,\&\, Q]e = 0)$$

$$\text{DBX}_{\succcurlyeq}\ \ [x' = f(x) \,\&\, Q](e)' \geq ge \rightarrow (e \succcurlyeq 0 \rightarrow [x' = f(x) \,\&\, Q]e \succcurlyeq 0) \qquad (\succcurlyeq \text{ either } \geq \text{ or } >)$$

Proof. Axiom DBX$_{\succcurlyeq}$ is derived first before axiom DBX is derived as a corollary. After propositional normalization, the derivation starts with a DG, DG$_\forall$ step, introducing a new ghost variable y satisfying a carefully chosen differential equation $y' = -gy$. Next, ∃R, ∀L pick an initial value for y. It suffices to pick any $y > 0$. The augmented ODE is abbreviated with $\alpha_y \equiv x' = f(x), y' = -gy$ in the derivation.

$$
{\scriptstyle \text{DG, DG}_\forall}\dfrac{{\scriptstyle \exists\text{R}, \forall\text{L}}\dfrac{[\alpha_y \,\&\, Q](e)' \geq ge, e \succcurlyeq 0, y > 0 \vdash [\alpha_y \,\&\, Q]e \succcurlyeq 0}{\forall y\, [\alpha_y \,\&\, Q](e)' \geq ge, e \succcurlyeq 0 \vdash \exists y\, [\alpha_y \,\&\, Q]e \succcurlyeq 0}}{[x' = f(x) \,\&\, Q](e)' \geq ge, e \succcurlyeq 0 \vdash [x' = f(x) \,\&\, Q]e \succcurlyeq 0}
$$

The augmented ODE α_y has a new provable invariant relationship $ey \succcurlyeq 0$ (see Fig. 3.1 and discussion after this proof). To deduce the original property of interest ($e \succcurlyeq 0$) from this new relationship, it suffices to prove $y > 0$ invariant because the formula $ey \succcurlyeq 0 \land y > 0 \to e \succcurlyeq 0$ is provable by $\mathbb{R}$. Axiom DC is used to prove $y > 0$ separately (right premise abbreviated with ①) and assume it in the evolution domain constraints of the left premise. Subsequently, monotonicity rule M['] and $\mathbb{R}$ strengthen the postcondition to $ey \succcurlyeq 0$ using the added domain constraint $y > 0$.

$$
{\scriptstyle \text{DC}}\dfrac{{\scriptstyle \text{M}['], \mathbb{R}}\dfrac{[\alpha_y \,\&\, Q \land y > 0](e)' \geq ge, e \succcurlyeq 0, y > 0 \vdash [\alpha_y \,\&\, Q \land y > 0]ey \succcurlyeq 0}{[\alpha_y \,\&\, Q \land y > 0](e)' \geq ge, e \succcurlyeq 0, y > 0 \vdash [\alpha_y \,\&\, Q \land y > 0]e \succcurlyeq 0} \qquad \text{①}}{[\alpha_y \,\&\, Q](e)' \geq ge, e \succcurlyeq 0, y > 0 \vdash [\alpha_y \,\&\, Q]e \succcurlyeq 0}
$$

From the left premise, a cut, $\mathbb{R}$ step adds $ey \succcurlyeq 0$ to the assumptions using the provable arithmetic formula $e \succcurlyeq 0 \land y > 0 \to ey \succcurlyeq 0$. Axiom $\text{DI}_\succcurlyeq$ is used to prove the inequational invariant $ey \succcurlyeq 0$ and the resulting differential $(ey)'$ simplifies with $(\,)'$ from Section 2.3.3 (page 24). An additional DE, [:=] step replaces the differential variable y' according to the augmented ODE α_y, before a monotonicity M['] (with a cut) and $\mathbb{R}$ step closes the derivation using the domain constraint $y > 0$. The differential ghost $y' = -gy$ is specifically crafted so that this final arithmetic step proves with $\mathbb{R}$.

$$
{\scriptstyle \text{cut}, \mathbb{R}}\dfrac{{\scriptstyle \text{DI}_\succcurlyeq}\dfrac{{\scriptstyle (\,)'}\dfrac{{\scriptstyle \text{DE}, [:=]}\dfrac{{\scriptstyle \text{M}[']}\dfrac{{\scriptstyle \mathbb{R}}\dfrac{*}{(e)' \geq ge, y > 0 \vdash (e)'y + e(-gy) \geq 0}}{[\alpha_y \,\&\, Q \land y > 0](e)' \geq ge \vdash [\alpha_y \,\&\, Q \land y > 0](e)'y + e(-gy) \geq 0}}{[\alpha_y \,\&\, Q \land y > 0](e)' \geq ge \vdash [\alpha_y \,\&\, Q \land y > 0](e)'y + ey' \geq 0}}{[\alpha_y \,\&\, Q \land y > 0](e)' \geq ge \vdash [\alpha_y \,\&\, Q \land y > 0](ey)' \geq 0}}{[\alpha_y \,\&\, Q \land y > 0](e)' \geq ge, ey \succcurlyeq 0 \vdash [\alpha_y \,\&\, Q \land y > 0]ey \succcurlyeq 0}}{[\alpha_y \,\&\, Q \land y > 0](e)' \geq ge, e \succcurlyeq 0, y > 0 \vdash [\alpha_y \,\&\, Q \land y > 0]ey \succcurlyeq 0}
$$

The derivation continues from premise ① with a second differential ghost $z' = \frac{g}{2}z$ analogously:

$$
{\scriptstyle \text{DG}}\dfrac{{\scriptstyle \exists\text{R}, \text{M}['], \mathbb{R}}\dfrac{{\scriptstyle \text{DI}_=, (\,)'}\dfrac{{\scriptstyle \text{DE}, [:=]}\dfrac{{\scriptstyle \text{dW}}\dfrac{{\scriptstyle \mathbb{R}}\dfrac{*}{Q \vdash (-gy)z^2 + 2yz(\frac{g}{2}z) = 0}}{\vdash [x' = f(x), y' = -gy, z' = \frac{g}{2}z \,\&\, Q](-gy)z^2 + 2yz(\frac{g}{2}z) = 0}}{\vdash [x' = f(x), y' = -gy, z' = \frac{g}{2}z \,\&\, Q]y'z^2 + 2yzz' = 0}}{yz^2 = 1 \vdash [x' = f(x), y' = -gy, z' = \frac{g}{2}z \,\&\, Q]yz^2 = 1}}{y > 0 \vdash \exists z\, [x' = f(x), y' = -gy, z' = \frac{g}{2}z \,\&\, Q]y > 0}}{y > 0 \vdash [x' = f(x), y' = -gy \,\&\, Q]y > 0}
$$

In the ∃R, M['], $\mathbb{R}$ step, observe that if $y > 0$ initially, then there exists z such that $yz^2 = 1$. Moreover, $yz^2 = 1$ is sufficient to imply $y > 0$ in the postcondition. Rule $\mathbb{R}$ again applies here since both of these are properties of real arithmetic. The differential ghost $z' = \frac{g}{2}z$ is specifically constructed so that $yz^2 = 1$ can be proved invariant along the differential equation.

Axiom DBX derives using the derived axiom $[\cdot]\wedge$, the equivalence $e = 0 \leftrightarrow e \geq 0 \wedge -e \geq 0$ by $\mathbb{R}$, and the equality $(-e)' = -(e)'$ provable by $(\,)'$. The ODE is abbreviated in the derivation below with $\alpha_x \equiv x' = f(x) \,\&\, Q$.

$$\begin{array}{rr}
 & * \\ \cline{2-2}
\text{DBX}_{\succcurlyeq}, \wedge\text{L}, \wedge\text{R} & [\alpha_x](e)' \geq ge \wedge [\alpha_x](-e)' \geq g(-e), e \geq 0 \wedge -e \geq 0 \vdash [\alpha_x]e \geq 0 \wedge [\alpha_x]-e \geq 0 \\ \cline{2-2}
[\cdot]\wedge & [\alpha_x]((e)' \geq ge \wedge (-e)' \geq g(-e)), e \geq 0 \wedge -e \geq 0 \vdash [\alpha_x](e \geq 0 \wedge -e \geq 0) \\ \cline{2-2}
\text{M}['], \mathbb{R} & [x' = f(x) \,\&\, Q](e)' = ge, e = 0 \vdash [x' = f(x) \,\&\, Q]e = 0 \qquad \Box
\end{array}$$

The first two syntactic deduction steps in the derivation of $\text{DBX}_{\succcurlyeq}$ do not appear to have changed the sequent much, but they correspond to a significant geometric transformation of the problem, as illustrated in Fig. 3.1b and Fig. 3.1c. In the system extended with differential ghost y, there is now a *new* invariant $ey \succcurlyeq 0$ which can be observed along solutions! While the value of e decays (dangerously) towards 0, the chosen differential equation $y' = -gy$ yields an (integral) value for y that counteracts this change, ensuring that their product still always stays non-negative along all solutions. In fact, the value of ey even remains constant when the extended term e satisfies the equational identity $\dot{e} = ge$. The second differential ghost $z' = \frac{g}{2}z$ in the proof is similarly constructed so that $yz^2 = 1$ can be proved invariant *along the differential equation*. The geometric transformation from this second syntactic differential ghost is illustrated in Fig. 3.1d. Since the first differential ghost y satisfies a differential *equation*, the second ghost z exactly balances it out with the value of yz^2 remaining (provably) constant and positive at 1 along solutions (similarly to Fig. 3.1b).

The derivation of $\text{DBX}_{\succcurlyeq}$ illustrates how the ODE axioms of dL (DI, DC, DG) complement each other in proofs of ODE invariance. For brevity, the same derivation is used for both $\geq$ and $>$ cases of $\text{DBX}_{\succcurlyeq}$ even though the latter only needs one ghost (using $y' = -\frac{g}{2}y$ and invariant $ey^2 > 0$ instead). Axiom DBX also derives directly (similarly to $\text{DBX}_{\succcurlyeq}$, using the invariant $ey = 0$ instead) using just two differential ghosts rather than the four incurred with $[\cdot]\wedge$.

Corollary 3.6 (Darboux (in)equality rules)**.** *The Darboux equality dbx and Darboux inequality* $\textit{dbx}_{\succcurlyeq}$ *proof rules derive from DG (and DI, DC) for any extended cofactor term* g.

$$\text{dbx}\ \frac{Q \vdash \dot{e} = ge}{e = 0 \vdash [x' = f(x) \,\&\, Q]e = 0} \qquad \text{dbx}_{\succcurlyeq}\ \frac{Q \vdash \dot{e} \geq ge}{e \succcurlyeq 0 \vdash [x' = f(x) \,\&\, Q]e \succcurlyeq 0} \quad (\succcurlyeq \text{ either } \geq \text{ or } >)$$

Proof. The dbx proof rule derives from axiom DBX (and rule $\text{dbx}_{\succcurlyeq}$ from axiom $\text{DBX}_{\succcurlyeq}$) using an additional $(\,)'$, DE step to differentials $(e)'$ into Lie derivatives $\dot{e}$, followed by dW:

$$\begin{array}{rr}
 & Q \vdash \dot{e} = ge \\ \cline{2-2}
\text{dW} & \vdash [x' = f(x) \,\&\, Q]\dot{e} = ge \\ \cline{2-2}
(\,)', \text{DE}, [:=] & \vdash [x' = f(x) \,\&\, Q](e)' = ge \\ \cline{2-2}
\text{DBX} & e = 0 \vdash [x' = f(x) \,\&\, Q]e = 0
\end{array}
\qquad
\begin{array}{rr}
 & Q \vdash \dot{e} \geq ge \\ \cline{2-2}
\text{dW} & \vdash [x' = f(x) \,\&\, Q]\dot{e} \geq ge \\ \cline{2-2}
(\,)', \text{DE}, [:=] & \vdash [x' = f(x) \,\&\, Q](e)' \geq ge \\ \cline{2-2}
\text{DBX}_{\succcurlyeq} & e \succcurlyeq 0 \vdash [x' = f(x) \,\&\, Q]e \succcurlyeq 0
\end{array}
\qquad \Box$$

The following example shows a concrete proof utilizing the newly derived proof rules.

Example 3.7 (Proving ODE properties in dL)**.** Judging by the plot (Fig. 2.1) of the ODE α_e from (2.1), trajectories from within the open (or closed) disk stay trapped within the disk. Rather

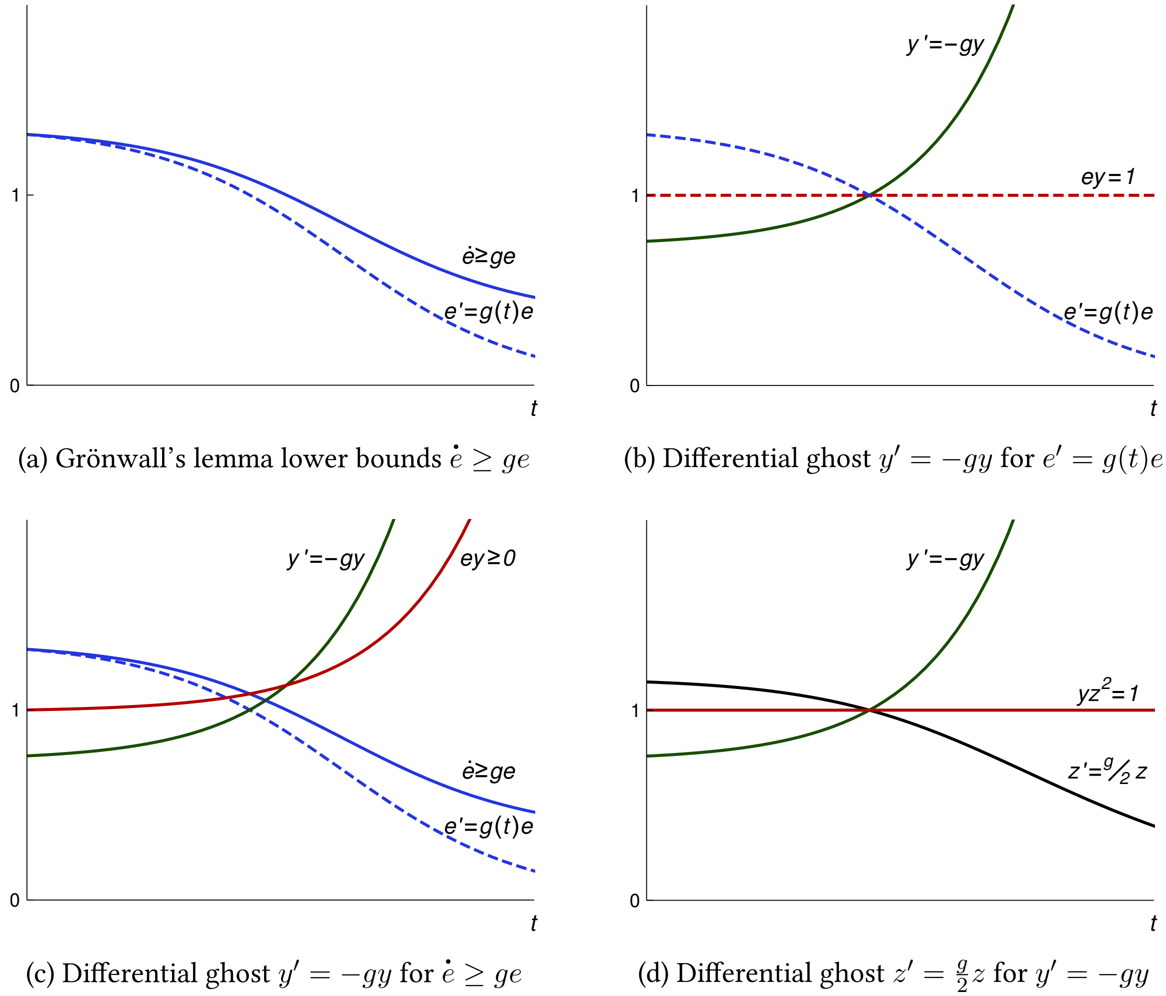

(a) Grönwall's lemma lower bounds $\dot{e} \geq ge$

(b) Differential ghost $y' = -gy$ for $e' = g(t)e$

(c) Differential ghost $y' = -gy$ for $\dot{e} \geq ge$

(d) Differential ghost $z' = \frac{g}{2}z$ for $y' = -gy$

Figure 3.1: The horizontal axis tracks the evolution of time t along solutions. Dashed lines indicate steps based on semantical arguments while solid lines indicate constructions used in the syntactical proof of Lemma 3.5 (lines are labeled above with their respective equations or inequalities). In Fig. 3.1a, solutions of $\dot{e} \geq ge$ (solid blue) are bounded below by those of the non-autonomous linear differential equation $e' = g(t)e$ (dashed blue) by Grönwall's lemma. In Fig. 3.1b, the differential ghost $y' = -gy$ (solid green) balances out $e' = g(t)e$ so that the value of ey (dashed red) remains constant at 1. In Fig. 3.1c, the same ghost $y' = -gy$ also balances out $\dot{e} \geq ge$, where the value of ey (solid red) remains non-negative but not necessarily constant. In Fig. 3.1d, a second differential ghost $z' = \frac{g}{2}z$ (solid black) balances out $y' = -gy$ so that the value of yz^2 (solid red) remains constant at 1. The constant 1 in the RHS of $ey = 1$ and $yz^2 = 1$ in Figs. 3.1b and 3.1d respectively is chosen for simplicity. Any positive constant suffices with appropriate initial values of the differential ghosts.

than relying (informally) on a potentially incorrect plot though, this fact can be shown formally by proving that $e \succcurlyeq 0$, with $e = 1 - u^2 - v^2$, is an invariant of α_e. The Lie derivative of e along α_e is: $\mathcal{L}_{\alpha_e}(e) = -2u(-v + \frac{u}{4}(1 - u^2 - v^2)) - 2v(u + \frac{v}{4}(1 - u^2 - v^2)) = -\frac{1}{2}(u^2 + v^2)e$. Thus, the following derivation with dbx$_{\succcurlyeq}$ proves invariance of $1 - u^2 - v^2 \succcurlyeq 0$, since e satisfies the (polynomial) inequality $\dot{e} \geq ge$ with polynomial cofactor $g = -\frac{1}{2}(u^2 + v^2)$.

$$\text{dbx}_{\succcurlyeq}\,\dfrac{\mathbb{R}\,\dfrac{*}{\vdash \mathcal{L}_{\alpha_e}(1 - u^2 - v^2) \geq -\frac{1}{2}(u^2 + v^2)(1 - u^2 - v^2)}}{1 - u^2 - v^2 \succcurlyeq 0 \vdash [\alpha_e]1 - u^2 - v^2 \succcurlyeq 0}$$

In fact, the term e obeys the special equational case $\dot{e} = ge$ (Fig. 3.1b) in which the seemingly innocuous syntactic introduction of a differential ghost $y' = -gy$ even *exactly* balances out the complicated (decaying) evolution of e geometrically. Indeed, in this case, $e = 0$ can also be proved invariant for the ODE α_e using rule dbx. This proves the observation from Fig. 2.1 that the unit circle is also invariant for α_e. △

The derivations of axioms DBX, DBX$_{\succcurlyeq}$ give constructive choices of differential ghosts when the invariant is a Darboux (in)equality. The derived rule dbx$_{\succcurlyeq}$ already exceeds the deductive power of DI, DC because the formula $y > 0 \rightarrow [y' = -y]y > 0$ is easily provable by dbx$_{\succcurlyeq}$ using the Darboux equality $\dot{y} = -y$, but is *not* provable with DI, DC alone [140].

3.3.3 Barrier Certificates

Barrier certificates [155, 158] are certificates of safety for ODEs and hybrid systems. Briefly, given an ODE safety question $\Gamma \vdash [x' = f(x) \,\&\, Q]\phi$, a barrier certificate term b is such that formula $b \succcurlyeq 0$ is a suitable invariant for proving that safety question, i.e., the sequents $\Gamma \vdash b \succcurlyeq 0$, $b \succcurlyeq 0 \vdash \phi$, and $b \succcurlyeq 0 \vdash [x' = f(x) \,\&\, Q]b \succcurlyeq 0$ are all valid. In practice, such certificates are computationally attractive because they can (sometimes) be found by suitably encoding the above validity questions, e.g., with sum-of-squares programming [155, 158], and solving the resulting problem through numerical methods. However, numerical inaccuracies in the results, especially when polynomials of high degree are involved, means that the generated candidates are often unsound [40] and they *must* be carefully checked for soundness when used in safety *proofs*. A useful application of rule dbx$_{\succcurlyeq}$ is to derive a *sound* proof rule[4] that checks the invariance for $b \succcurlyeq 0$ according to the barrier certificates condition.

Corollary 3.8 (Strict barrier certificates are differential ghosts)**.** *The barrier certificates Barr proof rule derives from DG (and DI, DC) for any **polynomial** term p and **polynomial** ODE $x' = f(x)$ in a closed semialgebraic domain Q, i.e., Q is formed from conjunctions and disjunctions of non-strict inequalities over **polynomial** terms.*

$$\text{Barr}\ \dfrac{Q, p = 0 \vdash \dot{p} > 0}{p \succcurlyeq 0 \vdash [x' = f(x) \,\&\, Q]p \succcurlyeq 0} \qquad (\textit{where} \succcurlyeq \textit{is} \geq \textit{or} >)$$

Proof. The proof starts by considering the cofactor "term" $\hat{g} = \frac{\dot{p}p}{\max(\dot{p}, p^2)}$ which, crucially for soundness, is *not* included in the syntax of extended dL terms because it contains division by

[4] The sound justification of this rule is important, e.g., the earlier presentation of barrier certificates [155] with $\dot{p} \geq 0$ instead of $\dot{p} > 0$ in the succedent of the premise of rule Barr is unsound [45, Example 2].

the non-smooth $\max$ function. This cofactor is modified later to be a polynomial term but, for the moment, note the following valid arithmetic inequality:

$$\begin{aligned} \max(\dot{p}, p^2)\dot{p} &= \begin{cases} \dot{p}^2 & \text{if } \dot{p} \geq p^2 \\ \dot{p}p^2 & \text{otherwise } \dot{p} < p^2 \end{cases} \\ &\geq \dot{p}p^2 \qquad \text{(in both cases)} \end{aligned} \tag{3.3}$$

In the former case, inequality (3.3) is justified by multiplying both sides of the case assumption $\dot{p} \geq p^2$ by $\dot{p}$ (which is non-negative because it is bounded below by a squared term). From the premise of rule Barr, the "term" $\max(\dot{p}, p^2)$ is strictly positive in domain Q, so dividing both sides of (3.3) by $\max(\dot{p}, p^2)$ proves the inequality $\dot{p} \geq \hat{g}p$. This argument justifies the LHS "derivation" below which, however, uses an illegal cofactor "term" $\hat{g}$.

$$\text{dbx}_{\succcurlyeq}\,\dfrac{\dfrac{Q, p = 0 \vdash \dot{p} > 0}{Q \vdash \dot{p} \geq \hat{g}p}}{p \succcurlyeq 0 \vdash [x' = f(x) \,\&\, Q]p \succcurlyeq 0} \qquad\qquad \text{dbx}_{\succcurlyeq}\,\dfrac{\mathbb{R}\,\dfrac{Q, p = 0 \vdash \dot{p} > 0}{Q \vdash \dot{p} \geq gp}}{p \succcurlyeq 0 \vdash [x' = f(x) \,\&\, Q]p \succcurlyeq 0}$$

It remains to identify a cofactor g that satisfies the inequality $\hat{g}p \geq gp$ in Q so that the RHS derivation above is justified by rule $\mathbb{R}$ (because all entries of the sequent are formulas in $\text{FOL}_{\mathbb{R}}$). The identification of g uses a bound on the rate of growth of continuous semialgebraic functions from real algebraic geometry [14, Prop 2.6.2]. By an abuse of notation, all polynomial terms below, e.g., $\dot{p}, p$, denote their respective polynomial functions over x.

The function $\sigma(x) = \frac{\dot{p}(x)}{\max(\dot{p}(x), p(x)^2)}$ is a *semialgebraic function* on the closed domain $[\![Q]\!] \subseteq \mathbb{R}^n$ because, within $[\![Q]\!]$, its denominator is strictly positive and the graph relation $\sigma(x) = y$ is characterized by the following semialgebraic formula, where the disjuncts case split on the $\max$ function appearing in the denominator of $\sigma(x)$:

$$\sigma(x) = y \equiv \big(\dot{p}(x) \geq p(x)^2 \land \dot{p}(x) = y\dot{p}(x)\big) \lor \big(\dot{p}(x) < p(x)^2 \land \dot{p}(x) = yp(x)^2\big)$$

Thus, by [14, Prop 2.6.2], the function σ is bounded in norm with $|\sigma(x)| \leq c\left(1 + \|x\|_2^2\right)^r$ for some (positive) constant $c \in \mathbb{R}$ and power $r \in \mathbb{N}$. This bound justifies the choice of (polynomial) cofactor $g = -c\left(1 + \|x\|_2^2\right)^r p(x)$ because for all $x \in [\![Q]\!]$,

$$\hat{g}(x)p(x) = \sigma(x)p(x)^2 \geq -c\left(1 + \|x\|_2^2\right)^r p(x)^2 = g(x)p(x) \qquad \square$$

Corollary 3.8 generalizes an earlier result [178], which showed that many flavors of barrier certificates in the literature can be understood using a comparison principle (in particular, with $\text{dbx}_{\succcurlyeq}$) to the case of *strict* barrier certificates [155, 158]. Together with dbx, $\text{dbx}_{\succcurlyeq}$, rule Barr enables sound checking of various invariant candidates using (at most) two DG steps. However, Corollary 3.8 is restricted to polynomials because it uses a bound from real algebraic geometry [14, Prop 2.6.2] to select an appropriate polynomial cofactor. This restriction is immaterial in practice because techniques for generating barrier certificates use optimization over polynomials anyway [155, 158]. Nevertheless, the general barrier certificates proof rule for arbitrary extended terms e shown below is a special case of the complete proof rule for semianalytic invariants derived in Theorem 3.29 with axiomatic extensions.

$$\text{Barr}\;\dfrac{Q, e = 0 \vdash \dot{e} > 0}{e \succcurlyeq 0 \vdash [x' = f(x) \,\&\, Q]e \succcurlyeq 0} \qquad (\text{where } \succcurlyeq \text{ is } \geq \text{ or } >)$$

The next section builds on these constructions, showing that the deductive power afforded by axiom DG extends to *all* true analytic invariants.

3.4 Analytic Invariants

Analytic formulas are formed from finite conjunctions and disjunctions of equalities, but, over $\mathbb{R}$, can be normalized to a single equality $e = 0$ using the provable real arithmetic equivalences: $e = 0 \land \tilde{e} = 0 \leftrightarrow e^2 + \tilde{e}^2 = 0$ and $e = 0 \lor \tilde{e} = 0 \leftrightarrow e\tilde{e} = 0$. Thus, it suffices to restrict attention to equational formulas $e = 0$ when proving completeness for analytic invariants.

The key to completeness is the differential radical identity (3.2) for e with arbitrary rank $N \geq 1$, which analyzes *all* higher Lie derivatives simultaneously. Suppose that extended term e satisfies identity (3.2) with rank N and some cofactors g_i. Taking Lie derivatives on both sides of identity (3.2) yields:

$$\begin{aligned}
\dot{e}^{(N+1)} &= \mathcal{L}_{f(x)}(\dot{e}^{(N)}) = \mathcal{L}_{f(x)}\left(\sum_{i=0}^{N-1} g_i\dot{e}^{(i)}\right) = \sum_{i=0}^{N-1} \mathcal{L}_{f(x)}(g_i\dot{e}^{(i)}) = \sum_{i=0}^{N-1} \left(\dot{g_i}\dot{e}^{(i)} + g_i\dot{e}^{(i+1)}\right)\\
&= \sum_{i=0}^{N-1} \left(\dot{g_i}\dot{e}^{(i)}\right) + \sum_{i=0}^{N-2} \left(g_i\dot{e}^{(i+1)}\right) + g_{N-1}\dot{e}^{(N)}\\
&= \sum_{i=0}^{N-1} \left(\dot{g_i}\dot{e}^{(i)}\right) + \sum_{i=0}^{N-2} \left(g_i\dot{e}^{(i+1)}\right) + g_{N-1}\left(\sum_{i=0}^{N-1} g_i\dot{e}^{(i)}\right)
\end{aligned}$$

The last step follows using (3.2) to expand $\dot{e}^{(N)}$. Observe that the resulting expression for $\dot{e}^{(N+1)}$ is again a sum over the lower Lie derivatives $\dot{e}^{(i)}$ for $i = 0, \ldots, N-1$ multiplied by appropriate cofactors. By repeatedly taking Lie derivatives on both sides, the higher Lie derivatives $\dot{e}^{(N)}, \dot{e}^{(N+1)}, \ldots$ can all be written as sums over these lower Lie derivatives with appropriate cofactors. Thus, *in the real analytic setting*, for initial states ω where $\omega[\![e]\!], \omega[\![\dot{e}]\!], \ldots, \omega[\![\dot{e}^{(N-1)}]\!]$ all simultaneously evaluate to 0, formula $e = 0$ (*and similarly for all its higher Lie derivatives*) stays invariant along solutions to the ODE.

This suggests that rule dbx should be generalized by considering higher Lie derivatives. The canonical technique for generalizing to higher derivatives comes from the study of ODEs. All (explicit form) ordinary differential equations involving higher derivatives can be transformed into vectorial systems of differential equations involving only first derivatives but possibly over a vector of variables [204, §11.I]. This transformation can be done syntactically and is precisely the idea used to derive the (complete) proof rule for analytic invariants by reduction to a suitable vectorial generalization of rule dbx. This crucial vectorial generalization is derived first.

3.4.1 Vectorial Darboux Equalities

Suppose that the m-dimensional vector of extended terms $\mathbf{e} = (\mathbf{e}_1, \ldots, \mathbf{e}_m)$ satisfies the vectorial identity $\dot{\mathbf{e}} = G\mathbf{e}$, where G is an $m \times m$ matrix of extended terms and $\dot{\mathbf{e}}$ denotes component-wise Lie derivatives of vector $\mathbf{e}$ along $x' = f(x)$ and $(\mathbf{e})'$ denotes component-wise differentials. If all components of $\mathbf{e}$ evaluate to 0 in an initial state, then they all always stay at 0 along $x' = f(x)$ because their component-wise Lie derivatives *all* evaluate to 0 in that initial state.

Lemma 3.9 (Vectorial Darboux equalities are differential ghosts)**.** *The vectorial Darboux axiom VDBX derives from DG (and DI, DC), where G is an $m \times m$ cofactor matrix of extended terms and $\mathbf{e}$ is an m-dimensional vector of extended terms.*

$$\text{VDBX}\ \ [x' = f(x) \,\&\, Q](\mathbf{e})' = G\mathbf{e} \to (\mathbf{e} = 0 \to [x' = f(x) \,\&\, Q]\mathbf{e} = 0)$$

Proof. First, observe that the formula $\mathbf{e} = 0$ is provably equivalent in real arithmetic to the formula $-\|\mathbf{e}\|_2^2 \geq 0$; recall that the term $\|\mathbf{e}\|_2^2 \overset{\text{def}}{=} \sum_{i=1}^{m} \mathbf{e}_i^2$ is the *squared* Euclidean norm of vector $\mathbf{e}$. The derivation starts with M[′], cut and $\mathbb{R}$ to rephrase $\mathbf{e} = 0$ using this equivalence. Thanks to this rephrasing, the sequent no longer contains vectorial quantities and the derivation is completed using a (scalar) DBX$_{\succcurlyeq}$ step with the extended term cofactor $g = \|G\|_F^2 + 1$, where the term $\|G\|_F^2 \overset{\text{def}}{=} \sum_{i=1}^{m} \sum_{j=1}^{m} G_{ij}^2$ is the squared Frobenius norm of matrix G.

$$\scriptstyle{\mathrm{M}['],\,\mathrm{cut},\,\mathbb{R}}\ \dfrac{\scriptstyle{\mathrm{DBX}_{\succcurlyeq}}\ \dfrac{\scriptstyle{\mathrm{M}[']}\ \dfrac{\scriptstyle{(\,)',\,\mathbb{R}}\ \dfrac{*}{(\mathbf{e})' = G\mathbf{e} \vdash (-\|\mathbf{e}\|_2^2)' \geq g(-\|\mathbf{e}\|_2^2)}}{[x' = f(x) \,\&\, Q](\mathbf{e})' = G\mathbf{e} \vdash [x' = f(x) \,\&\, Q](-\|\mathbf{e}\|_2^2)' \geq g(-\|\mathbf{e}\|_2^2)}}{[x' = f(x) \,\&\, Q](\mathbf{e})' = G\mathbf{e},\ -\|\mathbf{e}\|_2^2 \geq 0 \vdash [x' = f(x) \,\&\, Q]-\|\mathbf{e}\|_2^2 \geq 0}}{[x' = f(x) \,\&\, Q](\mathbf{e})' = G\mathbf{e},\ \mathbf{e} = 0 \vdash [x' = f(x) \,\&\, Q]\mathbf{e} = 0}$$

All that remains is to justify the final ( )′, $\mathbb{R}$ step after M[′] by showing that the following arithmetic formula is provable:

$$(\mathbf{e})' = G\mathbf{e} \ \to\ (-\|\mathbf{e}\|_2^2)' \geq g(-\|\mathbf{e}\|_2^2) \tag{3.4}$$

The differential $(-\|\mathbf{e}\|_2^2)'$ is calculated (and proved via ( )′ from Section 2.3.3, page 24) as follows, where $\mathbf{u} \cdot \mathbf{v}$ denotes the dot product of vectors $\mathbf{u}, \mathbf{v}$. The last step uses $(\mathbf{e})' = G\mathbf{e}$:

$$(-\|\mathbf{e}\|_2^2)' = -(\sum_{i=1}^{m} \mathbf{e}_i^2)' = -2\sum_{i=1}^{m} \mathbf{e}_i(\mathbf{e}_i)' = -2(\mathbf{e} \cdot (\mathbf{e})') = -2(\mathbf{e} \cdot (G\mathbf{e}))$$

Thus, it suffices to prove the validity of formula $-2(\mathbf{e} \cdot (G\mathbf{e})) \geq g(-\|\mathbf{e}\|_2^2)$, i.e., its truth in all states ω. Validity is first shown semantically. For ease of notation, let $\omega[\![\mathbf{e}]\!], \omega[\![G]\!]$ stand for the respective real vector and matrix values of $\mathbf{e}$ and G evaluated component-wise in state ω. The notation $\|\cdot\|_2, \|\cdot\|_F$ denotes the (real-valued) Euclidean and Frobenius norms for vectors and matrices respectively. By the Cauchy-Schwarz inequality [204, §28.I], the dot product between vectors $\omega[\![\mathbf{e}]\!]$ and $\omega[\![G]\!]\omega[\![\mathbf{e}]\!]$ is bounded by the product of their norms:

$$\omega[\![\mathbf{e}]\!] \cdot (\omega[\![G]\!]\omega[\![\mathbf{e}]\!]) \leq \|\omega[\![\mathbf{e}]\!]\|_2 \ \|\omega[\![G]\!]\omega[\![\mathbf{e}]\!]\|_2$$

The norm $\|\omega[\![G]\!]\omega[\![\mathbf{e}]\!]\|_2$ of this matrix-vector product is bounded by the product of their matrix and vector norms because the Euclidean and Frobenius norms are compatible [204, §14.II]:

$$\|\omega[\![G]\!]\omega[\![\mathbf{e}]\!]\|_2 \leq \|\omega[\![G]\!]\|_F \ \|\omega[\![\mathbf{e}]\!]\|_2$$

Expanding the (square) inequality $0 \leq (\|\omega[\![G]\!]\|_F - 1)^2$ yields an upper bound on the Frobenius norm $\|\omega[\![G]\!]\|_F$ by its squared value:

$$2\,\|\omega[\![G]\!]\|_F \leq \|\omega[\![G]\!]\|_F^2 + 1$$

Chaining these (in)equalities yields:

$$\begin{aligned}\omega[\![-2(\mathbf{e}\cdot(G\mathbf{e}))]\!] &= -2(\omega[\![\mathbf{e}]\!]\cdot(\omega[\![G]\!]\omega[\![\mathbf{e}]\!])) \geq -2\,\|\omega[\![\mathbf{e}]\!]\|_2\,\|\omega[\![G]\!]\omega[\![\mathbf{e}]\!]\|_2\\ &\geq -2\,\|\omega[\![\mathbf{e}]\!]\|_2\,\|\omega[\![G]\!]\|_F\,\|\omega[\![\mathbf{e}]\!]\|_2 = -2\,\|\omega[\![G]\!]\|_F\,\|\omega[\![\mathbf{e}]\!]\|_2^2\\ &\geq (\|\omega[\![G]\!]\|_F^2 + 1)(-\,\|\omega[\![\mathbf{e}]\!]\|_2^2) = \omega[\![g(-\,\|\mathbf{e}\|_2^2)]\!]\end{aligned}$$

where $\|\omega[\![G]\!]\|_F^2 + 1$ is precisely the semantic value of cofactor g in state ω. Since this semantic argument for the validity of implication (3.4) only depends on first-order properties of the real closed fields, which is decidable [14, 197], formula (3.4) is provable syntactically by $(\,)'$, $\mathbb{R}$. □

Corollary 3.10 (Vectorial Darboux equality rule)**.** *The vectorial Darboux equality proof rule vdbx derives from DG (and DI, DC), where G is an $m \times m$ cofactor matrix of extended terms and $\mathbf{e}$ is an m-dimensional vector of extended terms.*

$$\text{vdbx}\ \frac{Q \vdash \dot{\mathbf{e}} = G\mathbf{e}}{\mathbf{e} = 0 \vdash [x' = f(x)\,\&\,Q]\mathbf{e} = 0}$$

Proof. Rule vdbx derives from derived axiom VDBX using $(\,)'$, DE, $[:=]$ to provably transform between $(\mathbf{e})'$ and $\dot{\mathbf{e}}$, just like rule dbx derives from derived axiom DBX in Corollary 3.6. □

The use of axiom $\text{DBX}_{\succcurlyeq}$ in the derivation of axiom VDBX corresponds to an application of Grönwall's lemma [70, 204, §29.VI], as illustrated in Fig. 3.1a. In case e starts with value 0 initially and satisfies the Darboux inequality $\dot{e} \geq ge$, the constant zero solution of the differential equation $e' = g(t)e$ bounds it from below. In Fig. 3.1a, this corresponds to the case where both blue lines lie exactly on the horizontal axis. The proof uses the (squared) Euclidean and Frobenius norms to reduce a vectorial equality ($\mathbf{e} = 0$) to a scalar inequality ($-\,\|\mathbf{e}\|_2^2 \geq 0$), which enables further analysis using *scalar* differential ghosts. The convenient choice of compatible norms ensures that all syntactic proof steps are done within the extended term language. Since all norms are equivalent on finite-dimensional vector spaces [204, §10.III], this reduction can also be done using other norms with suitable syntactic representations. Convenient choices of norms are a common technique in the study of differential equations [204].

An alternative derivation of rule vdbx is given in an earlier result [148] based on Liouville's formula [204, §15.III]. That alternative derivation has a geometric interpretation as a continuous change of basis that is expressed purely syntactically [148] but requires the use of *vectorial* differential ghosts and a number of ghost variables that is quadratic in the dimension. The new derivation in Lemma 3.9 uses exactly 2 scalar differential ghosts in the $\text{DBX}_{\succcurlyeq}$ step independent of dimension and relies only on basic properties of real arithmetic. In fact, just like the scalar Darboux axioms, axiom VDBX for m-dimensional extended terms $\mathbf{e}$ derives once-and-for-all so no differential ghosts are needed for its subsequent use.

3.4.2 Completeness for Analytic Invariants

Returning to extended terms e satisfying the differential radical identity (3.2), the following proof rule for invariance of $e = 0$ based on *higher* Lie derivatives derives as a direct instance of derived rule vdbx:

Theorem 3.11 (Differential radical invariants are vectorial Darboux). *The differential radical invariant proof rule dRI derives from vdbx (which in turn derives from DG).*

$$\text{dRI}\ \frac{\Gamma, Q \vdash \bigwedge_{i=0}^{N-1} \overset{\bullet}{e}^{(i)} = 0 \quad Q \vdash \overset{\bullet}{e}^{(N)} = \sum_{i=0}^{N-1} g_i \overset{\bullet}{e}^{(i)}}{\Gamma \vdash [x' = f(x) \,\&\, Q]e = 0}$$

Proof Summary (Appendix A.2.3). Rule dRI derives from rule vdbx by transforming identity (3.2) involving higher Lie derivatives of e into a vectorial Darboux equality involving only first Lie derivatives of the extended term vector $\mathbf{e}$, using the following choice of cofactor matrix G:

$$G = \begin{pmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ 0 & 0 & \dots & 0 & 1 \\ g_0 & g_1 & \cdots & g_{N-2} & g_{N-1} \end{pmatrix}, \quad \mathbf{e} = \begin{pmatrix} e \\ \overset{\bullet}{e}^{(1)} \\ \vdots \\ \overset{\bullet}{e}^{(N-2)} \\ \overset{\bullet}{e}^{(N-1)} \end{pmatrix}$$

The matrix G has 1 on its superdiagonal and the g_i cofactors in the last row. The left premise of dRI is used to prove formula $\mathbf{e} = 0$ true initially, while the right premise of dRI is used to show the premise of vdbx. □

For any extended term e in the LHS of normalized equation $e = 0$, the computable differential radicals condition (R) requires that the differential radical identity (3.2) (computably) exists and proves with associated rank N and cofactors g_i for e. The resulting (provable) identity (3.2) proves the right premise of dRI.[5] The succedent in the remaining left premise of dRI thus gives a finitary characterization for when *all* Lie derivatives of e evaluate to zero in the initial state. This motivates the following definition of a *finite* formula summarizing that *all* higher Lie derivatives of e are zero:

Definition 3.12 (Differential radical formula). The *differential radical formula* $\overset{\bullet}{e}^{(*)} = 0$ for extended term e of rank $N \geq 1$ from identity (3.2) with Lie derivatives along $x' = f(x)$ is defined to be the formula $\overset{\bullet}{e}_f^{(*)} = 0$ below (left), where the dependency on ODE $x' = f(x)$ is dropped when it is clear from the context (shown on the right):

$$\overset{\bullet}{e}_f^{(*)} = 0 \overset{\text{def}}{\equiv} \bigwedge_{i=0}^{N-1} \mathcal{L}_{f(x)}^{(i)}(e) = 0 \qquad \overset{\bullet}{e}^{(*)} = 0 \overset{\text{def}}{\equiv} \bigwedge_{i=0}^{N-1} \overset{\bullet}{e}^{(i)} = 0$$

[5] Theorem 3.11 shows Q can be assumed when proving the right premise of dRI. A finite rank must exist either way, but assuming Q may reduce the number of higher Lie derivatives of e that need to be considered for the proof (as in Example 3.41).

The finiteness of $\dot{e}^{(*)} = 0$ depends on Lie derivatives along the particular differential equation $x' = f(x)$ of interest, because, without considering the ODE, *no* corresponding chain of higher-order differentials would stabilize. The rest of this chapter uses Lie derivatives for this finiteness property, but relies under the hood on dL's axiomatic proof transformation from differentials.

The completeness of derived rule dRI can be proved semantically by extending earlier arguments [63] to extended term languages. Even better: the following equivalent characterization in arithmetic of the truth of analytic formulas along forward evolutions of differential equations derives axiomatically using the extensions developed in Section 3.5.[6] In contrast to the semantic completeness argument, this syntactic characterization enables complete proofs and *complete disproofs* of analytic invariance within the dL calculus. In other words, *disproving* the RHS of the characterization under assumptions Γ, yields a dL proof of $\Gamma \vdash \neg[x' = f(x) \,\&\, Q]e = 0$.

Theorem 3.13 (Analytic completeness). *The differential radical invariant axiom DRI derives in* dL *when Q is a semianalytic formula formed from conjunctions and disjunctions of strict inequalities:*

$$\text{DRI}\;\; [x' = f(x) \,\&\, Q]e = 0 \leftrightarrow \big(Q \rightarrow \dot{e}^{(*)} = 0\big)$$

Proof Summary (Appendix A.2.3). The "←" direction derives (for any Q) by an application of derived rule dRI, whose right premise closes by (3.2). The "→" direction relies on existence and uniqueness of solutions to differential equations, which are internalized later as axioms in Section 3.5. □

For the proof of Theorem 3.13, the additional axioms are *only required* for syntactically deriving the "→" direction (completeness) of DRI. The "←" direction (soundness) derives using dRI, which, by Theorem 3.11, can be derived using only DI, DC, DG. Thus, the base dL axiomatization with differential ghosts is *complete* for proving properties of the form $[x' = f(x) \,\&\, Q]e = 0$ because dRI provably reduces all such questions to $Q \rightarrow \dot{e}^{(*)} = 0$. The validity of this resulting semianalytic formula is a purely arithmetical question. In fact, the base dL axiomatization *decides* $[x' = f(x) \,\&\, Q]p = 0$ in the case where $x' = f(x)$ is polynomial and Q is semialgebraic, because the resulting RHS of DRI is semialgebraic, and hence, a formula of decidable real arithmetic [14, 197]. The same applies for the next result, which is a corollary of Theorem 3.13 but applies beyond the continuous fragment, thanks to the compositional application of dL's hybrid program axioms.

Corollary 3.14 (Analytic hybrid program completeness). *For analytic formulas P and analytic hybrid programs α, i.e., whose tests and domain constraints are negations of analytic formulas, it is possible to compute an extended term e such that the equivalence $[\alpha]P \leftrightarrow e = 0$ is derivable in* dL*, provided that the term language is Noetherian.*[7]

[6] With these axiomatic extensions, the requirement in Theorem 3.13 that Q is formed from strict inequalities is not necessary. A derived equivalence axiom for analytic invariance with arbitrary semianalytic domain constraint Q is given in Theorem 3.30.

[7] The set of extended terms always forms a ring because the $+, \cdot$ operations are interpreted as the usual real-valued addition and multiplication and hence obey the ring axioms [24] for $\mathbb{R}$. An extended term language is said to be *Noetherian* if its corresponding ring of extended terms is Noetherian. Like the computable differential radicals condition (R), an algorithm is assumed that decides (and proves) ideal membership in the ring of extended terms.

Proof. The analytic formula P is provably equivalent in real arithmetic to a formula $e = 0$ for some extended term e by normalizing P with the provable real arithmetic equivalences $e = 0 \land \tilde{e} = 0 \leftrightarrow e^2 + \tilde{e}^2 = 0$ and $e = 0 \lor \tilde{e} = 0 \leftrightarrow e\tilde{e} = 0$. Assume without loss of generality that it is already written in this form, and accordingly for the negated analytic formulas in α so that analytic hybrid programs are generated by the grammar (3.5).

$$\alpha, \beta ::= x := e \mid ?d \neq 0 \mid x' = f(x) \,\&\, d \neq 0 \mid \alpha \cup \beta \mid \alpha; \beta \mid \alpha^* \tag{3.5}$$

The proof proceeds by structural induction on the fragment of dL programs generated by the grammar (3.5), showing that for some (computable) extended term $\tilde{e}$, the equivalence $[\alpha]e = 0 \leftrightarrow \tilde{e} = 0$ is derivable in dL.

- Case $x' = f(x) \,\&\, d \neq 0$. The formula $d \neq 0$ is a strict inequality so Theorem 3.13 derives $[x' = f(x) \,\&\, d \neq 0]e = 0 \leftrightarrow (d \neq 0 \rightarrow \dot{e}^{(*)} = 0)$. Let N be the rank of e so that $\dot{e}^{(*)} = 0$ expands to $\bigwedge_{i=0}^{N-1} \dot{e}^{(i)} = 0$ and let $\tilde{e} = d(\sum_{i=0}^{N-1} (\dot{e}^{(i)})^2)$, giving the provable real arithmetic equivalence $(d \neq 0 \rightarrow \dot{e}^{(*)} = 0) \leftrightarrow \tilde{e} = 0$. Rewriting with this derives the equivalence:

$$[x' = f(x) \,\&\, d \neq 0]e = 0 \leftrightarrow \tilde{e} = 0$$

- Case $x := e$. Axiom [:=] derives the equivalence $[x := e]\tilde{e}(x) = 0 \leftrightarrow \tilde{e}(e) = 0$, where $\tilde{e}(e)$ is an extended term.

- Case $?d \neq 0$. Axiom [?] derives the equivalence $[?d \neq 0]e = 0 \leftrightarrow (d \neq 0 \rightarrow e = 0)$. Rewriting with the provable real arithmetic equivalence $(d \neq 0 \rightarrow e = 0) \leftrightarrow de = 0$ derives the equivalence:

$$[?d \neq 0]e = 0 \leftrightarrow de = 0$$

- Case $\alpha \cup \beta$. Axiom [$\cup$] derives the equivalence $[\alpha \cup \beta]e = 0 \leftrightarrow [\alpha]e = 0 \land [\beta]e = 0$. By the induction hypothesis on α, β, the equivalences $[\alpha]e = 0 \leftrightarrow \tilde{e}_1 = 0$ and $[\beta]e = 0 \leftrightarrow \tilde{e}_2 = 0$ derive for some extended terms $\tilde{e}_1, \tilde{e}_2$. Moreover, $\tilde{e}_1 = 0 \land \tilde{e}_2 = 0 \leftrightarrow \tilde{e}_1^2 + \tilde{e}_2^2 = 0$ is provable in real arithmetic. Rewriting with the derived equivalences derives the equivalence:

$$[\alpha \cup \beta]e = 0 \leftrightarrow \tilde{e}_1^2 + \tilde{e}_2^2 = 0$$

- Case $\alpha; \beta$. Axiom [;] derives the equivalence $[\alpha; \beta]e = 0 \leftrightarrow [\alpha][\beta]e = 0$. By the induction hypothesis on β, the equivalence $[\beta]e = 0 \leftrightarrow \tilde{e}_2 = 0$ derives for some extended term $\tilde{e}_2$. Rewriting with this equivalence derives $[\alpha; \beta]e = 0 \leftrightarrow [\alpha]\tilde{e}_2 = 0$. By the induction hypothesis on α, the equivalence $[\alpha]\tilde{e}_2 = 0 \leftrightarrow \tilde{e}_1 = 0$ derives for some extended term $\tilde{e}_1$. Rewriting with the chain of derived equivalences derives the equivalence:

$$[\alpha; \beta]e = 0 \leftrightarrow \tilde{e}_1 = 0$$

- Case α^*. This case crucially requires that the extended term language is Noetherian. First, construct the sequence of terms $\tilde{e}_i$ defined inductively with $\tilde{e}_0 \overset{\text{def}}{=} e$ and $\tilde{e}_{i+1}$ is the term satisfying the derived equivalence $[\alpha]\tilde{e}_i = 0 \leftrightarrow \tilde{e}_{i+1} = 0$ obtained by applying the induction hypothesis on α with postcondition $\tilde{e}_i = 0$ for $i = 0, 1, 2, \ldots$. Since the term language is assumed to be Noetherian, the following ascending chain of ideals stabilizes:

$$(\tilde{e}_0) \subseteq (\tilde{e}_0, \tilde{e}_1) \subseteq (\tilde{e}_0, \tilde{e}_1, \tilde{e}_2) \subseteq \cdots$$

By decidable ideal membership for the extended term language, there is a (smallest) computable k such that $\tilde{e}_k$ satisfies the provable identity (3.6), with cofactor terms g_i:

$$\tilde{e}_k = \sum_{i=0}^{k-1} g_i \tilde{e}_i \tag{3.6}$$

The equivalence $\sum_{i=0}^{k-1} \tilde{e}_i^2 = 0 \leftrightarrow \bigwedge_{i=0}^{k-1} \tilde{e}_i = 0$ is provable by real arithmetic so, to derive the equivalence $[\alpha^*]e = 0 \leftrightarrow \sum_{i=0}^{k-1} \tilde{e}_i^2 = 0$, it suffices to show that the equivalence $[\alpha^*]e = 0 \leftrightarrow \bigwedge_{i=0}^{k-1} \tilde{e}_i = 0$ is derivable. The two directions of this latter equivalence are shown separately:

“$\rightarrow$” This direction is straightforward using the iteration axiom $[^*]$ k times together with derived axiom $[\cdot]\land$. By construction, the formulas $[\alpha]\tilde{e}_i = 0$ are provably equivalent to $\tilde{e}_{i+1} = 0$ using derived equivalences, which derives the required implication:

$$
{\scriptsize [^*]}\;\dfrac{{\scriptsize [^*],[\cdot]\land}\;\dfrac{{\scriptsize [^*],[\cdot]\land}\;\dfrac{{\scriptsize [^*],[\cdot]\land}\;\dfrac{\dfrac{*}{\tilde{e}_0 = 0 \land \tilde{e}_1 = 0 \land \tilde{e}_2 = 0 \land \cdots \land \tilde{e}_{k-1} = 0 \vdash \bigwedge_{i=0}^{k-1} \tilde{e}_i = 0}}{e = 0 \land [\alpha]e = 0 \land [\alpha][\alpha]e = 0 \land \cdots \land \underbrace{[\alpha]\ldots[\alpha]}_{k-1 \text{ times}} e = 0 \vdash \bigwedge_{i=0}^{k-1} \tilde{e}_i = 0}}{\cdots}}{e = 0 \land [\alpha]e = 0 \land [\alpha][\alpha][\alpha^*]e = 0 \vdash \bigwedge_{i=0}^{k-1} \tilde{e}_i = 0}}{e = 0 \land [\alpha][\alpha^*]e = 0 \vdash \bigwedge_{i=0}^{k-1} \tilde{e}_i = 0}}{[\alpha^*]e = 0 \vdash \bigwedge_{i=0}^{k-1} \tilde{e}_i = 0}
$$

“$\leftarrow$” The postcondition of the succedent box modality is strengthened to $\bigwedge_{i=0}^{k-1} \tilde{e}_i = 0$ by monotonicity with M[·], recalling that $\tilde{e}_0 \overset{\text{def}}{=} e$, so $\bigwedge_{i=0}^{k-1} \tilde{e}_i = 0 \rightarrow e = 0$ is a propositional tautology. Subsequently, the loop rule is used to prove that $\bigwedge_{i=0}^{k-1} \tilde{e}_i = 0$ is a loop invariant of α^*:

$$
{\scriptsize \mathrm{M}[\cdot]}\;\dfrac{{\scriptsize \text{loop}}\;\dfrac{\bigwedge_{i=0}^{k-1} \tilde{e}_i = 0 \vdash [\alpha] \bigwedge_{i=0}^{k-1} \tilde{e}_i = 0}{\bigwedge_{i=0}^{k-1} \tilde{e}_i = 0 \vdash [\alpha^*] \bigwedge_{i=0}^{k-1} \tilde{e}_i = 0}}{\bigwedge_{i=0}^{k-1} \tilde{e}_i = 0 \vdash [\alpha^*]e = 0}
$$

By axiom $[\cdot]\land$ and $\land$R, each conjunct of the postcondition (indexed by $0 \le i \le k-1$) of the premise is proved separately. By construction, each $[\alpha]\tilde{e}_i = 0$ is provably equivalent to $\tilde{e}_{i+1} = 0$ so the premises for indices $0 \le i < k-1$ all close trivially because $\tilde{e}_{i+1} = 0$ is already in the antecedent. The last premise for index $i = k-1$ has succedent $\tilde{e}_k = 0$. However, this follows (by construction and $\mathbb{R}$) from the antecedent using the provable identity (3.6).

$$
{\scriptsize [\cdot]\land}\;\dfrac{\dfrac{{\scriptsize \mathbb{R}, \land\mathrm{R}}\;\dfrac{*}{\bigwedge_{i=0}^{k-1} \tilde{e}_i = 0 \vdash \bigwedge_{i=0}^{k-1} \tilde{e}_{i+1} = 0}}{\bigwedge_{i=0}^{k-1} \tilde{e}_i = 0 \vdash \bigwedge_{i=0}^{k-1} [\alpha]\tilde{e}_i = 0}}{\bigwedge_{i=0}^{k-1} \tilde{e}_i = 0 \vdash [\alpha] \bigwedge_{i=0}^{k-1} \tilde{e}_i = 0}
$$

□

The Noetherian condition of Corollary 3.14 implies the computable differential radicals condition (R). Polynomial term languages are Noetherian so Corollary 3.14 shows that dL decides $[\alpha]P$ where P and α are both algebraic. However, extended term languages are not necessarily Noetherian, for example, extended term language (3.1), even only with $\exp$, is not Noetherian [198, Remark 1.4.2].[8] Nevertheless, the stronger Noetherian condition is only required when the analytic hybrid program α contains loops. Otherwise, the weaker condition (R) suffices for loop-free α in Corollary 3.14. A version of Corollary 3.14 is proved for Noetherian functions in Corollary 3.39 with loop-free α or with assignment-free α (with loops).

3.5 Extended Axiomatization

This section presents an axiomatic extension to dL whose purpose is to internalize standard properties of differential equations, such as existence and uniqueness [204, §10.VI], as *syntactic* reasoning principles. The extension requires that the ODE system $x' = f(x)$ *locally evolves* x, i.e., it has no fixpoint at which $f(x)$ is the 0 vector. This can be ensured syntactically, e.g., by requiring that the system contains a clock variable $x_1' = 1$ that tracks the passage of time. In proofs, such a clock can always be added using axiom DG if necessary.

3.5.1 Existence, Uniqueness, and Continuity

The differential equations of dL are smooth so the Picard-Lindelöf theorem [204, §10.VI] guarantees that for any initial state ω, a *unique* solution of the system $x' = f(x)$, i.e., $\varphi : [0,T] \to \mathbb{S}$ with $\varphi(0) = \omega$, *exists* for some duration $T > 0$. The solution φ can be extended (uniquely) to its maximal open interval of existence [204, §10.IX] and $\varphi(\zeta)$ is smooth with respect to ζ.

Lemma 3.15 (Continuous existence, uniqueness, and differential adjoints)**.** *The following axioms are sound. In Cont and Dadj, y are fresh variables (not in $x' = f(x) \,\&\, Q(x)$ or extended term e).*

Uniq $\langle x' = f(x) \,\&\, Q_1 \land Q_2\rangle P \leftrightarrow \big(\langle x' = f(x) \,\&\, Q_1\rangle P\big) \land \big(\langle x' = f(x) \,\&\, Q_2\rangle P\big)$

Cont $x = y \to \big(\langle x' = f(x) \,\&\, e > 0\rangle x \neq y \leftrightarrow e > 0\big)$

Dadj $\langle x' = f(x) \,\&\, Q(x)\rangle\, x = y \leftrightarrow \langle y' = -f(y) \,\&\, Q(y)\rangle\, y = x$

Proof Summary (Appendix A.1.1). Uniq internalizes uniqueness, Cont internalizes continuity of the values of e and existence of solutions, and Dadj internalizes differential adjoints by the group action of time on ODE solutions, which is another consequence of existence and uniqueness. □

The *uniqueness axiom* Uniq says that if a state has two solutions φ_1, φ_2 respectively staying in evolution domains Q_1, Q_2 and whose endpoints satisfy P, then, by uniqueness, one of φ_1 or φ_2 is a prefix of the other, and therefore, that prefix stays in both evolution domains $Q_1 \land Q_2$ and satisfies P at its endpoint. The *continuous existence axiom* Cont expresses a notion of *local*

[8] The ring of all extended terms (including function composition) is not to be confused with the ring of Noetherian functions generated by a *single* Noetherian chain, the latter of which is indeed Noetherian (see Section 3.7).

progress for differential equations. It says that from an initial state satisfying $x = y$, the system can locally evolve to another state satisfying $x \neq y$ while still staying in the *open set* of states characterized by $e > 0$ iff the initial state is already in that open set. This uses the assumption that the system locally evolves x at all. The *differential adjoints* axiom Dadj expresses that x can flow forward to y iff y can flow backward to x along the negated ODE. It is at the heart of the "there and back again" axiom that equivalently expresses properties of differential equations with evolution domains in terms of properties of forward and backward differential equations without evolution domains [139].

Although all three axioms are stated as (conditional[9]) equivalences to support intuition, the main properties of interest are their "$\leftarrow$" directions. For example, the "$\rightarrow$" direction of Uniq derives from domain constraint monotonicity for the diamond modality (derived rule dRW$\langle\cdot\rangle$ below). Diamond modality monotonicity principles are given below because they are useful for working with the newly introduced axioms. They are derived duals of the usual dL box modality principles using axiom $\langle\cdot\rangle$.

Corollary 3.16 (Derived diamond modality domain rules and axioms)**.** *The following axiom and its corollary proof rule derive in* dL*:*

$$\text{DR}\langle\cdot\rangle \quad [x' = f(x) \,\&\, R]Q \rightarrow \big(\langle x' = f(x) \,\&\, R\rangle P \rightarrow \langle x' = f(x) \,\&\, Q\rangle P\big)$$

$$\text{dRW}\langle\cdot\rangle \quad \frac{R \vdash Q \quad \Gamma \vdash \langle x' = f(x) \,\&\, R\rangle P}{\Gamma \vdash \langle x' = f(x) \,\&\, Q\rangle P}$$

Proof. Axiom DR$\langle\cdot\rangle$ derives from DMP (the roles of Q and R are flipped) by dualizing with the $\langle\cdot\rangle$ axiom. The final K, dW steps use the propositional tautology $Q \rightarrow (R \rightarrow Q)$.

$$\begin{array}{rr}
 & * \\ \hline
\text{dW} & \vdash [x' = f(x) \,\&\, R](Q \rightarrow (R \rightarrow Q)) \\ \hline
\text{K} & [x' = f(x) \,\&\, R]Q \vdash [x' = f(x) \,\&\, R](R \rightarrow Q) \\ \hline
\text{DMP} & [x' = f(x) \,\&\, R]Q, [x' = f(x) \,\&\, Q]\neg P \vdash [x' = f(x) \,\&\, R]\neg P \\ \hline
\langle\cdot\rangle, \neg\text{R}, \neg\text{L} & [x' = f(x) \,\&\, R]Q, \langle x' = f(x) \,\&\, R\rangle P \vdash \langle x' = f(x) \,\&\, Q\rangle P
\end{array}$$

Rule dRW$\langle\cdot\rangle$ derives from DR$\langle\cdot\rangle$ by simplifying its outer (leftmost) assumption with rule dW. □

3.5.2 Real Induction

The final axiomatic extension is based on the real induction principle [34], briefly:

Definition 3.17 (Inductive subset [34])**.** The subset $S \subseteq [a, b]$ is called an *inductive* subset of the compact interval $[a, b]$ iff for all $a \leq \zeta \leq b$ such that $[a, \zeta) \subseteq S$,

① $\zeta \in S$ and

② if $\zeta < b$ then $(\zeta, \zeta + \varepsilon] \subseteq S$ for some $\varepsilon > 0$.

Here, $[a, a)$ is the empty interval, hence ① requires $a \in S$.

[9] Axiom Cont is sound even without the condition from assumption $x = y$. It is stated conditionally to align with the intuition of local evolution from an initial state satisfying $x = y$.

Proposition 3.18 (Real induction [34])**.** *The subset $S \subseteq [a,b]$ is inductive iff $S = [a,b]$.*

Proof. In the "$\Leftarrow$" direction, $S = [a,b]$ is inductive by definition. For the "$\Rightarrow$" direction, let $S \subseteq [a,b]$ be inductive. Suppose $S \neq [a,b]$, so that the complement set $S^{\complement} = [a,b] \setminus S$ is nonempty. Let ζ be the infimum of $S^{\complement}$, then $\zeta \in [a,b]$ since $[a,b]$ is left-closed. First, note that $[a,\zeta) \subseteq S$. Otherwise, ζ is not an infimum of $S^{\complement}$, because there would exist $a \leq \tau < \zeta$, such that $\tau \in S^{\complement}$. By ①, $\zeta \in S$. Next, if $\zeta = b$, then $S = [a,b]$, contradiction. Thus, $\zeta < b$, and by ②, $(\zeta, \zeta+\varepsilon] \subseteq S$ for some $\varepsilon > 0$. Since $\zeta \in S$, this implies that $\zeta + \varepsilon$ is a greater lower bound of $S^{\complement}$ than ζ, contradiction. □

Proposition 3.18 is based on the completeness of the reals [34] for compact intervals $[a,b]$ of $\mathbb{R}$. Applying it to the time axis of ODE solutions yields real induction *along* solutions of differential equations. For brevity, only the real induction axiom for systems without evolution domain constraints is presented here, leaving the general version to Appendix A.1.1, since evolution domains are definable in dL [139].

Lemma 3.19 (Real induction)**.** *The real induction axiom RI is sound, where variable y is fresh in formula $[x' = f(x)]P$.*

$$\text{RI}\quad [x' = f(x)]P \leftrightarrow \forall y\, [x' = f(x)\,\&\,P \lor x = y]\big(x = y \to P \land \langle x' = f(x)\,\&\,P \lor x = y\rangle x \neq y\big)$$

Proof Summary (Appendix A.1.1). The RI axiom follows from the real induction principle [34] and the Picard-Lindelöf theorem [204, §10.VI]. □

Real induction axiom RI can be understood in relation to Def. 3.17: its RHS is true in a state iff the subset of times at which the solution satisfies P is inductive. First, $\forall y\,[\ldots]\big(x = y \to \ldots\big)$ can be understood as quantifying over all final states ($x = y$) reached by trajectories staying within P except possibly at the endpoint $x = y$. This corresponds to $[a,\zeta) \subseteq S$ in Def. 3.17. The left conjunct (P) under the box modality expresses that P is still true at such an endpoint, corresponding to ① in Def. 3.17. The right conjunct ($\langle x' = f(x)\,\&\,P \lor x = y\rangle x \neq y$) expresses that P continues to remain true locally when following the ODE for a short time, corresponding to ② in Def. 3.17.

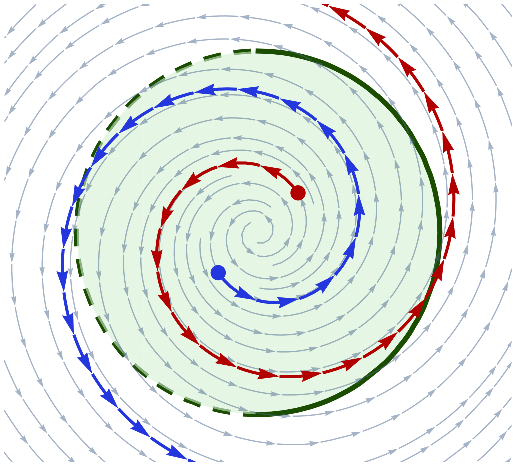

Figure 3.2: The half-open green disk is not invariant for the ODE α_e from (2.1) because the red and blue trajectories spiral out of it at a closed (solid green) or open (dashed green) boundary, respectively.

To see the topological significance of RI, recall the ODE α_e from (2.1) (page 17) and consider a set of points that is *not invariant.* Figure 3.2 illustrates two trajectories that leave the half-open disk characterized by the disjunctive formula: $u^2 + v^2 < \frac{1}{4} \lor u^2 + v^2 = \frac{1}{4} \land u \geq 0$. Trajectories starting in the disk leave it through its boundary but only in one of two ways: either at a point which is also in the disk (red trajectory exiting right) or which is not in the disk (blue trajectory exiting left). The left conjunct of RI rules out trajectories like the blue one exiting left in Fig. 3.2, while the right conjunct rules out trajectories like the red trajectory exiting right. The right conjunct of axiom RI also suggests a way to use it—axiom RI reduces proofs of invariance to local progress properties under the box modality. This motivates the following syntactic modality abbreviation for *local progress* into an evolution domain Q:

$$\langle x' = f(x) \,\&\, Q \rangle \bigcirc \overset{\text{def}}{\equiv} \langle x' = f(x) \,\&\, Q \lor x = y \rangle\, x \neq y$$

All proofs in this chapter use the $\bigcirc$ modality with an initial assumption $x = y$, where y is fresh. In this case, where $\omega[\![x]\!] = \omega[\![y]\!]$, since the ODE locally evolves x, the $\bigcirc$ modality has the following semantics:

> $\omega \in [\![\langle x' = f(x) \,\&\, Q \rangle \bigcirc]\!]$ iff there is a function $\varphi : [0, T] \to \mathbb{S}$ with $T > 0, \varphi(0) = \omega$,
> φ solves the ODE $x' = f(x)$ and $\varphi(\zeta) \in [\![Q]\!]$ for all ζ in the half-open interval $(0, T]$

Thus, the abbreviation $\bigcirc$ is a continuous-time version of the *next* modality of temporal logic [109] for differential equations. Conventionally, such a next state operator is excluded from continuous-time generalizations of temporal logic [76] because there is no unique “next” state in the continuous setting. The local progress modality $\bigcirc$ overcomes this by instead quantifying over *some* time interval $(0, T]$, with $T > 0$, of states along the solution. Intuitively, the exclusion of time 0 is because the $\bigcirc$ modality describes what solutions will do *next* (or *locally*) rather than what they are doing *now*. A precise (topological) explanation is provided in Appendix A.2.2 and a complete characterization of local progress for all semianalytic formulas is derived in Section 3.6. As a corollary, Section 3.6 shows that, like its discrete counterpart, the $\bigcirc$ modality is self-dual for semianalytic Q.

The final derived rule rI shows what the added axioms and local progress provide—axiom RI reduces global invariance properties of ODEs to local progress properties. These local progress properties are provable using Cont, Uniq and the dL axioms, as shown in the next section.

Corollary 3.20 (Real induction rule)**.** *The real induction proof rule rI derives from RI, Dadj. Variables y are fresh in the ODE $x' = f(x)$ and formula P.*

$$\text{rI}\ \frac{x{=}y, P \vdash \langle x' = f(x) \,\&\, P \rangle \bigcirc \qquad x{=}y, \neg P \vdash \langle x' = -f(x) \,\&\, \neg P \rangle \bigcirc}{P \vdash [x' = f(x)]P}$$

Proof Summary (Appendix A.1.2). Rule rI derives from axiom RI, where the left/right premises of the rule correspond respectively to the right/left conjunct of the RHS of RI. Axiom Dadj is used to syntactically flip signs in the right premise. □

3.6 Semianalytic Invariants

This section makes the simplifying assumption that domain constraint $Q \equiv \mathit{true}$ since it is definable in dL [139] and not central to the core idea of the section. Using the generalizations of RI, rI from Appendix A.1, the case of semianalytic invariants for ODEs with arbitrary semianalytic evolution domain Q is given in Appendix A.2.

The first step in invariance proofs for semianalytic P is to use derived rule rI, which yields premises of the form $x{=}y, P \vdash \langle x' = f(x) \,\&\, P\rangle \bigcirc$ (modulo sign changes and negation). These premises express local progress properties of the ODE $x' = f(x)$. Analogously to the equivalent arithmetic reduction of equational properties of differential equations in Theorem 3.13 using the *finite* differential radical formula (Def. 3.12), the key insight is that local progress for any semianalytic formula is also (provably) completely characterized by a corresponding *finite* semianalytic progress formula.

3.6.1 Local Progress

This section shows how an arithmetical characterization of local progress can be derived syntactically in dL for extended term languages and proves the completeness of this characterization. This characterization was previously used implicitly for semialgebraic invariants [64, 103]. The derivation is built up systematically, starting from the base case of atomic inequalities before moving on to the full semianalytic case. Interesting properties of this characterization, e.g., self-duality, are also observed.

Atomic Inequalities. Consider the atomic inequality $e \succcurlyeq 0$. To show *local progress* into such an inequality, it is sufficient to locally consider the *first* (significant) Lie derivative of e because the sign of a smooth function is locally dominated by the sign of its first non-zero derivative, if one exists. The key to a syntactic rendition uses the following lemma for non-strict inequalities.

Lemma 3.21 (Local progress step). *The local progress step axiom LPi$_{\geq}$ derives from Cont. Variables y are fresh in the ODE $x' = f(x)$ and extended term e.*

$$\text{LPi}_{\geq}\quad x{=}y \to \Big(e \geq 0 \land \big(e = 0 \to \langle x' = f(x) \,\&\, \dot{e} \geq 0\rangle x{\neq}y\big) \to \langle x' = f(x) \,\&\, e \geq 0\rangle x{\neq}y\Big)$$

Proof. The proof starts with a $\lor$L case split since the antecedent formula $e \geq 0$ is equivalent to formula $e > 0 \lor e = 0$ by $\mathbb{R}$. The resulting premises are respectively abbreviated ① for $e > 0$ and ② for $e = 0$ and continued below.

$${\scriptstyle \mathbb{R},\lor\mathrm{L}}\dfrac{\text{①} \qquad \text{②}}{x{=}y, e \geq 0, e = 0 \to \langle x' = f(x) \,\&\, \dot{e} \geq 0\rangle x{\neq}y \vdash \langle x' = f(x) \,\&\, e \geq 0\rangle x{\neq}y}$$

From premise ①, since the value of e is already positive initially, it must *locally* stay positive. Using dRW$\langle\cdot\rangle$, the non-strict inequality in the domain constraint of the succedent is strengthened to a strict one, after which axiom Cont finishes the derivation.

$${\scriptstyle \mathrm{dRW}\langle\cdot\rangle}\dfrac{{\scriptstyle \mathrm{Cont}}\dfrac{*}{x{=}y, e > 0 \vdash \langle x' = f(x) \,\&\, e > 0\rangle x{\neq}y}}{x{=}y, e > 0 \vdash \langle x' = f(x) \,\&\, e \geq 0\rangle x{\neq}y}$$

From premise ②, the local sign of e cannot be determined from its initial value alone. The proof looks to the Lie derivative of e, which is assumed to be locally non-negative (in the implication $e = 0 \to \ldots$). Axiom DR$\langle\cdot\rangle$ reduces the succedent to a box modality question, after which axiom DI finishes the proof; this *refinement*-style technique is further explored for proving ODE liveness properties in Chapter 4.

$$\text{→L}\;\dfrac{\text{DR}\langle\cdot\rangle\;\dfrac{\text{DI, ( )}',\text{ DE, [:=]}\;\dfrac{*}{e=0\vdash[x'=f(x)\,\&\,\dot{e}\geq 0]e\geq 0}}{e=0,\langle x'=f(x)\,\&\,\dot{e}\geq 0\rangle x{\neq}y\vdash\langle x'=f(x)\,\&\,e\geq 0\rangle x{\neq}y}}{e=0,e=0\rightarrow\langle x'=f(x)\,\&\,\dot{e}\geq 0\rangle x{\neq}y\vdash\langle x'=f(x)\,\&\,e\geq 0\rangle x{\neq}y}\qquad\square$$

Similar to DBX, $\text{DBX}_{\succcurlyeq}$, and VDBX, a version of $\text{LPi}_{\geq}$ derives once-and-for-all using differentials ($(e)'\geq 0$) in domain constraints. This presentation is omitted to keep with the notational convention that domain constraints are always differential-free formulas (Section 2.1). Mathematically, to conclude that e is locally non-negative, it is important that the Lie derivative $\dot{e}$ is assumed to be *locally* non-negative rather than just *initially* non-negative. Just as the local sign of e cannot be determined (directly) when its initial value is zero, the same is true for $\dot{e}$. This difference drives the use of *higher* Lie derivatives when $\text{LPi}_{\geq}$ is generalized below. Syntactically, this difference manifests in both $\text{DR}\langle\cdot\rangle$, DI proof steps which crucially rely on the formula $\dot{e}\geq 0$ appearing in their respective domain constraints rather than simply as an initial assumption.

Observe that $\text{LPi}_{\geq}$ allows derivations to pass from reasoning about local progress[10] for $e\geq 0$ to local progress for its (first) Lie derivative $\dot{e}\geq 0$ whilst accumulating $e=0$ in the antecedent. This is reminiscent of derivative tests from elementary calculus used for testing the local behavior around a given stationary point of a (sufficiently) smooth function. The difference is that (syntactic) Lie derivatives have to be used for soundness instead of analytic time derivatives, but these notions are provably equal along ODEs using differentials [142, Lem. 35]. Similar to derivative tests, if the first Lie derivative is indeterminate as well, then the derivation can look to the second higher Lie derivative, and so on. Deductively, this is done by repeated use of derived axiom $\text{LPi}_{\geq}$ until the k-th derivative (shown as . . . in the outline below):

$$\text{LPi}_{\geq}\;\dfrac{\Gamma\vdash e\geq 0\qquad \text{LPi}_{\geq}\;\dfrac{\Gamma,e{=}0\vdash\dot{e}\geq 0\qquad\dfrac{\Gamma,x{=}y,e{=}0,\ldots,\dot{e}^{(k-1)}=0\vdash\langle x'=f(x)\,\&\,\dot{e}^{(k)}{\geq}0\rangle x{\neq}y}{\ldots}}{\Gamma,x{=}y,e{=}0\vdash\langle x'=f(x)\,\&\,\dot{e}\geq 0\rangle x{\neq}y}}{\Gamma,x{=}y\vdash\langle x'=f(x)\,\&\,e\geq 0\rangle x{\neq}y}$$

Notice that the rightmost premise closes whenever the (strict) inequality $\dot{e}^{(k)}>0$ can be proved from the accumulated antecedents. In that case, the local sign of e (and of all its Lie derivatives below the k-th one) is dominated by that of $\dot{e}^{(k)}$ because all of the lower (Lie) derivatives have indeterminate sign. The use of Cont, $\text{dRW}\langle\cdot\rangle$ finishes the proof because the solution must then locally enter $\dot{e}^{(k)}>0$, for example, with:

$$\text{cut}\;\dfrac{\Gamma,x{=}y,e=0,\ldots,\dot{e}^{(k-1)}=0\vdash\dot{e}^{(k)}>0\qquad\text{dRW}\langle\cdot\rangle\;\dfrac{\text{Cont}\;\dfrac{*}{x{=}y,\dot{e}^{(k)}>0\vdash\langle x'=f(x)\,\&\,\dot{e}^{(k)}>0\rangle x{\neq}y}}{x{=}y,\dot{e}^{(k)}>0\vdash\langle x'=f(x)\,\&\,\dot{e}^{(k)}\geq 0\rangle x{\neq}y}}{\Gamma,x{=}y,e=0,\ldots,\dot{e}^{(k-1)}=0\vdash\langle x'=f(x)\,\&\,\dot{e}^{(k)}\geq 0\rangle x{\neq}y}$$

[10] The local progress property used in $\text{LPi}_{\geq}$ is syntactically simpler than for the $\bigcirc$ modality (no $x=y$ in the domain constraints). For non-strict inequalities, the two are equivalent but the syntactic simplification in $\text{LPi}_{\geq}$ allows its re-use as a lemma in proving $\bigcirc$ local progress for both non-strict and strict inequalities.

The extended term conditions for smoothness (S) and syntactic partial derivatives (P) guarantee that all of the (infinitely many) higher Lie derivatives of e are well-defined semantically and syntactically. Derivations, on the other hand, are finite syntactic objects and can only mention *finitely many* Lie derivatives. Thus, one might suspect they are insufficient (hence incomplete) when, e.g., none of the higher Lie derivatives has a definite sign or if (infinitely many) different choices of k are needed in the proof depending on the initial state that satisfies assumptions Γ.

This is where the third, computable differential radicals condition (R) is crucially used. When N is the rank of e according to identity (3.2), then once the derivation has gathered $e = 0, \dots, \dot{e}^{(N-1)} = 0$, i.e., $\dot{e}^{(*)} = 0$ in the antecedents, derived rule dRI proves the invariant $e = 0$ and ODEs always locally progress in invariants. Furthermore, this argument shows (mathematically) that it is unnecessary to analyze higher Lie derivatives of e beyond N when proving local progress for $e > 0$ because none of those higher Lie derivatives will be sign-definite. Thus, the rank from (3.2) provides a uniform and finite bound for the number of Lie derivatives of e that need to be analyzed in any state, regardless of assumptions Γ. This finiteness property motivates the following definition, which gathers the above open premises to obtain a finite formula characterizing the *first significant Lie derivative* of e:

Definition 3.22 (First significant Lie derivative)**.** The *progress formula* $\dot{e}_f^{(*)} > 0$ for extended term e of rank $N \geq 1$ from identity (3.2) with Lie derivatives along $x' = f(x)$ is defined to be:

$$\begin{aligned}\dot{e}_f^{(*)} > 0 \overset{\text{def}}{\equiv} & e \geq 0 \land \big(e = 0 \rightarrow \mathcal{L}_{f(x)}(e) \geq 0\big) \land \big(e = 0 \land \mathcal{L}_{f(x)}(e) = 0 \rightarrow \mathcal{L}^{(2)}_{f(x)}(e) \geq 0\big) \\ & \land \cdots \land \big(e = 0 \land \mathcal{L}_{f(x)}(e) = 0 \land \cdots \land \mathcal{L}^{(N-3)}_{f(x)}(e) = 0 \rightarrow \mathcal{L}^{(N-2)}_{f(x)}(e) \geq 0\big) \\ & \land \big(e = 0 \land \mathcal{L}_{f(x)}(e) = 0 \land \cdots \land \mathcal{L}^{(N-2)}_{f(x)}(e) = 0 \rightarrow \mathcal{L}^{(N-1)}_{f(x)}(e) > 0\big)\end{aligned}$$

The dependency on ODE $x' = f(x)$ is dropped with $\dot{e}^{(*)} > 0$ when it is clear from the context. The *progress formula* $\dot{e}^{(*)} \geq 0$ is defined to be $\dot{e}^{(*)} > 0 \lor \dot{e}^{(*)} = 0$. The formulas $\dot{e}^{-(*)} > 0$ (or $\dot{e}^{-(*)} \geq 0$) are identical except their Lie derivatives are along the negated ODE $x' = -f(x)$.

Lemma 3.23 (Local progress $\succcurlyeq$)**.** *The local progress inequality axioms LP*$_{\geq^*}$*, LP*$_{>^*}$ *derive from LPi*$_{\geq}$ *and thus from Cont. Variables y are fresh in the ODE $x' = f(x)$ and extended term e.*

$$\text{LP}_{\geq^*} \quad x{=}y \rightarrow \big(\dot{e}^{(*)} \geq 0 \rightarrow \langle x' = f(x) \,\&\, e \geq 0 \rangle \bigcirc\big)$$

$$\text{LP}_{>^*} \quad x{=}y \rightarrow \big(\dot{e}^{(*)} > 0 \rightarrow \langle x' = f(x) \,\&\, e > 0 \rangle \bigcirc\big)$$

Proof Summary (Appendix A.2.2). Both axioms derive after unfolding the syntactic abbreviation of the $\bigcirc$ modality. Axiom LP$_{\geq^*}$ derives by the preceding discussion with iterated use of derived axioms LPi$_{\geq}$ and dRI. Axiom LP$_{>^*}$ derives similarly, but with an additional tweak to weaken the strict inequality $e > 0$ so that axiom LPi$_{\geq}$ can be used. □

The difference between the derivations of LP$_{\geq^*}$ and LP$_{>^*}$ is mainly technical and boils down to the handling of the assumptions about the initial state, and in particular, $x{=}y$ (see Appendix A.1.2 and A.2.2). Intuitively, the difference arises from the fact that the formula $e \geq 0$ characterizes a topologically closed set while $e > 0$ characterizes an open set. To locally progress into a set from initial state ω, the state ω must already be in the topological closure of that set. Closed sets are equal to their closure so, e.g., ω must already satisfy $e \geq 0$ in order to locally progress into it. Sets that are not closed (e.g., open sets) are not equal to their closure as they lack points on their topological boundary. An example of this is the half-open disk illustrated in Fig. 3.2. Thus, it is possible to locally progress into such sets from their topological boundary without ω already starting in the set.

Semianalytic Formulas. Semianalytic formulas P normalize propositionally to the following disjunctive *normal form* with extended terms $e_{ij}, \tilde{e}_{ij}$:

$$P \equiv \bigvee_{i=0}^{M} \Big(\bigwedge_{j=0}^{m(i)} e_{ij} \geq 0 \land \bigwedge_{j=0}^{n(i)} \tilde{e}_{ij} > 0 \Big) \tag{3.7}$$

Progress formulas are lifted homomorphically to semianalytic formulas in normal form:

Definition 3.24 (Semianalytic progress formula)**.** The *semianalytic progress formula* $\dot{P}^{(*)}_f$ for a semianalytic formula P in normal form (3.7) and Lie derivatives along $x' = f(x)$ is defined as:

$$\dot{P}^{(*)}_f \overset{\text{def}}{\equiv} \bigvee_{i=0}^{M} \Big(\bigwedge_{j=0}^{m(i)} (\dot{e_{ij}})^{(*)}_f \geq 0 \land \bigwedge_{j=0}^{n(i)} (\dot{\tilde{e}_{ij}})^{(*)}_f > 0 \Big)$$

The dependency on ODE $x' = f(x)$ is dropped with $\dot{P}^{(*)}$ when it is clear from the context. The formula $\dot{P}^{-(*)}$ takes Lie derivatives along ODE $x' = -f(x)$ instead. A mention of the notation $\dot{P}^{(*)}$ is understood as the progress formula for semianalytic formula P after it is rewritten propositionally to any equivalent normal form (3.7).

Lemma 3.25 (Semianalytic local progress)**.** *The local progress formula axiom LPℝ derives from Cont, Uniq. Variables y are fresh in the ODE $x' = f(x)$ and semianalytic formula P.*

$$\text{LP}\mathbb{R}\ \ x{=}y \rightarrow \big(\dot{P}^{(*)} \rightarrow \langle x' = f(x) \,\&\, P \rangle \bigcirc\big)$$

Proof Summary (Appendix A.2.2). The shape of the semianalytic progress formula $\dot{P}^{(*)}$ guides the proof. The derivation is sketched at a high level here for the representative example formula:

$$\begin{aligned} P &\equiv (e_1 \geq 0 \land \tilde{e}_1 > 0) \lor (e_2 \geq 0 \land \tilde{e}_2 > 0) \\ \dot{P}^{(*)} &\equiv (\dot{e_1}^{(*)} \geq 0 \land \dot{\tilde{e}_1}^{(*)} > 0) \lor (\dot{e_2}^{(*)} \geq 0 \land \dot{\tilde{e}_2}^{(*)} > 0) \end{aligned}$$

To show local progress into a *disjunction*, it suffices to show local progress into either disjunct. The derivation starts by decomposing $\dot{P}^{(*)}$ according to its (outermost) disjunction and accordingly decomposing P in the local progress succedent with dRW$\langle\cdot\rangle$. The premise for the second disjunct resulting from the $\lor$L step is symmetric and omitted here.

$$\vee\text{L}\ \dfrac{\text{dRW}\langle\cdot\rangle\ \dfrac{x{=}y, \dot{e_1}^{(*)} \geq 0 \land \dot{\tilde{e}_1}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, e_1 \geq 0 \land \tilde{e}_1 > 0 \rangle \bigcirc}{x{=}y, \dot{e_1}^{(*)} \geq 0 \land \dot{\tilde{e}_1}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, P \rangle \bigcirc}}{x{=}y, \dot{P}^{(*)} \vdash \langle x' = f(x) \,\&\, P \rangle \bigcirc}$$

To show local progress into a *conjunction*, by Uniq, it suffices to show local progress into both conjuncts separately. The derivation continues using Uniq, $\wedge$R to split the conjunctive local progress succedent before the derived axioms LP$_{\geq^*}$, LP$_{>^*}$ are used to finish the proofs in the resulting atomic cases for inequalities $\geq, >$, respectively.

$$\text{Uniq}, \wedge\text{R}\;\dfrac{\text{LP}_{\geq^*}\dfrac{*}{x{=}y, \dot{e_1}^{(*)} \geq 0 \vdash \langle x'{=}f(x) \,\&\, e_1{\geq}0\rangle\bigcirc} \quad \text{LP}_{>^*}\dfrac{*}{x{=}y, \dot{\tilde{e}_1}^{(*)} > 0 \vdash \langle x'{=}f(x) \,\&\, \tilde{e}_1{>}0\rangle\bigcirc}}{x{=}y, \dot{e_1}^{(*)} \geq 0, \dot{\tilde{e}_1}^{(*)} > 0 \vdash \langle x'{=}f(x) \,\&\, e_1 \geq 0 \wedge \tilde{e}_1 > 0\rangle\bigcirc} \qquad \square$$

Completeness could potentially be lost in several steps of the proof of Lemma 3.25, e.g., the use of $\vee$L at the start of the derivation, or the implicational axioms LP$_{\geq^*}$, LP$_{>^*}$. The converse (completeness) direction of axiom LP$\mathbb{R}$ therefore does not follow immediately from Lemma 3.25. Instead, the axiom LP$\mathbb{R}$ can be re-used to derive its own strengthening to an equivalence. This equivalence justifies the syntactic abbreviation $\bigcirc$, recalling that the $\bigcirc$ modality of temporal logic is self-dual. It also shows that the progress formulas are congruent over equivalences.

Theorem 3.26 (Local progress completeness)**.** *The local progress axiom LP derives from Cont, Uniq. Variables y are fresh in the ODE $x' = f(x)$ and semianalytic formula P.*

$$\text{LP}\quad x{=}y \rightarrow \big(\langle x' = f(x) \,\&\, P\rangle\bigcirc \leftrightarrow \dot{P}^{(*)}\big)$$

Corollary 3.27 (Duality and congruence)**.** *The duality axiom $\neg\bigcirc$ and congruence proof rule for progress formulas CLP derive from LP and thus from Cont, Uniq. Variables y are fresh in the ODE $x' = f(x)$ and semianalytic formulas P, R.*

$$\neg\bigcirc\quad x{=}y \rightarrow \big(\langle x' = f(x) \,\&\, P\rangle\bigcirc \leftrightarrow \neg\langle x' = f(x) \,\&\, \neg P\rangle\bigcirc\big)$$

$$\text{CLP}\quad \dfrac{P \leftrightarrow R}{\dot{P}^{(*)} \leftrightarrow \dot{R}^{(*)}} \quad (\text{for ODE } x' = f(x))$$

Proof Summary for Theorem 3.26 and Corollary 3.27 (Appendix A.2.2). The derivation of axioms LP, $\neg\bigcirc$ and proof rule CLP use the homomorphic definition of semianalytic progress formulas which implies that any semianalytic formula P in normal form (3.7) has a corresponding normal form for $\neg P$ such that the equivalence $\neg(\dot{P}^{(*)}) \leftrightarrow (\dot{\neg P})^{(*)}$ is provable. Classically, in any state, either formula $\dot{P}^{(*)}$ or $\neg(\dot{P}^{(*)})$ is true. Therefore, by LP$\mathbb{R}$, the ODE must (exclusively, by uniqueness) either locally progress into P or $\neg P$ from this state. Both axioms LP, $\neg\bigcirc$ are derivable consequences of this fact, as shown syntactically in Appendix A.2.2. Rule CLP follows from LP by congruential equivalence [142]. $\square$

Congruence rule CLP shows that *any* equivalent choice of normal form (3.7) for semianalytic formula P gives a local progress formula that is (provably) equivalent to $\dot{P}^{(*)}$. The rule works for all (semianalytic) equivalences, including arithmetical ones over extended term languages, e.g., $\exp(x) = 1 \leftrightarrow x = 0$ from (3.1), so CLP does not follow immediately from the homomorphic definition of progress formulas.

3.6.2 Completeness for Semianalytic Invariants

Combining derived axiom LP and derived rule rI yields an effective proof rule which reduces a semianalytic invariance question to questions involving purely arithmetic formulas.

Theorem 3.28 (Semianalytic invariants)**.** *The semianalytic invariant proof rule sAI derives from RI, Dadj, Cont, Uniq for semianalytic formula P.*

$$\text{sAI}\ \frac{P \vdash \dot{P}^{(*)} \qquad \neg P \vdash \dot{(\neg P)}^{-(*)}}{P \vdash [x' = f(x)]P}$$

Proof. This follows immediately by rewriting the premises of rule rI with the equivalence LP. □

Completeness of sAI was first proved semantically for polynomial terms languages [103], making crucial use of semialgebraic sets and real analytic solutions to polynomial ODE systems. The proof rule sAI *derives* syntactically in dL and generalizes to semianalytic invariants for extended term languages. Its completeness derives syntactically too, which yields dL disproofs of semianalytic invariance when arithmetic counterexamples can be found.

Theorem 3.29 (Semianalytic invariant completeness)**.** *The semianalytic invariant axiom SAI derives from RI, Dadj, Cont, Uniq for semianalytic formula P.*

$$\text{SAI}\ \forall x\,(P \rightarrow [x' = f(x)]P) \leftrightarrow \forall x\,\big(P \rightarrow \dot{P}^{(*)}\big) \land \forall x\,\big(\neg P \rightarrow \dot{(\neg P)}^{-(*)}\big)$$

Proof in Appendix A.2.4.

In Appendix A.2, a generalization of Theorem 3.29 is proven that handles semianalytic evolution domains Q using LP and a corresponding generalization of axiom RI. The same appendix proves the following generalization of Theorem 3.13 for semianalytic evolution domains:

Theorem 3.30 (Analytic completeness with semianalytic domains)**.** *The differential radical invariant axiom DRI& derives from Cont, Uniq for semianalytic formula Q.*

$$\text{DRI\&}\ [x' = f(x)\,\&\,Q]e = 0 \leftrightarrow \big(Q \rightarrow e = 0 \land (\dot{Q}^{(*)} \rightarrow \dot{e}^{(*)} = 0)\big)$$

Proof in Appendix A.2.4.

Theorems 3.29 and 3.30 show that dL is complete for proving invariance of *all* (semi)analytic P of differential equations because it reduces all such questions equivalently to first-order formulas, e.g., on the RHS of derived axiom SAI. In addition, dL decides invariance properties for all first-order real arithmetic formulas P, because quantifier elimination [14, 197] can equivalently rewrite P to (semialgebraic) normal form (3.7) first. Unlike for Theorem 3.13 and its generalization Theorem 3.30, which equivalently reduce safety properties with analytic postconditions to arithmetic directly, Theorem 3.29 and its generalized version in Appendix A.2 are only equivalences for invariants P; the search for suitable invariants in proofs of ODE safety is the remaining practical challenge [180].

Of course, the complete proof rule sAI can be used to prove all of the suggested invariants for the ODE α_e from (2.1). However, Example 3.7 gives a significantly simpler proof for the invariance of $1 - u^2 - v^2 \succcurlyeq 0$ with $\text{dbx}_{\succcurlyeq}$. This has implications for implementations of sAI because simpler proofs help minimize dependence on real arithmetic decision procedures. For semianalytic formulas (that are not semialgebraic), proof rules resulting in simpler arithmetic premises might even be preferable because validity of the arithmetic premises is undecidable in general [162]. Logically, when P is formed from only strict (resp. non-strict) inequalities then the left (resp. right) premise of sAI closes trivially. This logical fact corresponds to the topological fact that the set P characterizes is topologically open (resp. closed) so only one of the two exit trajectories in Section 3.5.2 can occur.

3.7 Noetherian Functions

This section studies the class of Noetherian functions which meets all of the extended term conditions required in Section 3.2.4 and therefore inherits all soundness and completeness results of the preceding sections, including Theorems 3.29 and 3.30.

3.7.1 Mathematical Preliminaries

The following definition of Noetherian functions is standard, although the parameters that are used for studying the complexity of these functions [12, 56, 57] have been omitted. The notation $h : H \subseteq \mathbb{R}^k \to \mathbb{R}$ is used for real-valued functions with domain H, i.e., an open, connected subset of $\mathbb{R}^k$. With a slight abuse of notation, polynomials $p \in \mathbb{R}[x]$ over indeterminates $x = (x_1, \ldots, x_n)$ and their corresponding polynomial functions $p(x_1, \ldots, x_n)$ in $\mathbb{R}^n \to \mathbb{R}$ are used interchangeably.

Definition 3.31 (Noetherian chain and Noetherian function)**.** A *Noetherian chain* is a sequence of real analytic functions $h_1, \ldots, h_r : H \subseteq \mathbb{R}^k \to \mathbb{R}$ such that all partial derivatives in H for all $i = 1, \ldots, k$ and $j = 1, \ldots, r$ have the following form, where each $q_{ij} \in \mathbb{R}[y, z]$ is a polynomial in $k + r$ indeterminates with $y = (y_1, \ldots, y_k), z = (z_1, \ldots, z_r)$:

$$\frac{\partial h_j}{\partial y_i}(y) = q_{ij}(y, h_1(y), \ldots, h_r(y)) \tag{3.8}$$

The function $h : H \subseteq \mathbb{R}^k \to \mathbb{R}$ is *Noetherian* iff it can be written as $h(y) = p(y, h_1(y), \ldots, h_r(y))$, where $p \in \mathbb{R}[y, z]$ is a polynomial in $k + r$ indeterminates and $h_1, \ldots, h_r$ is a Noetherian chain. In that case, h is said to be *generated* by this polynomial and Noetherian chain respectively but the choice of generating chain and polynomial for h is not unique.

For the term language extension (3.1), $\exp$ is a 1-element Noetherian chain because its derivative is $\frac{\partial \exp(y)}{\partial y} = \exp(y)$, while $\sin, \cos$ form a 2-element Noetherian chain. All three functions together form a 3-element Noetherian chain. More generally, the union of any (finite) number of Noetherian chains is a Noetherian chain. By definition, any element of a Noetherian chain is itself a Noetherian function so $\exp, \sin, \cos$ are also Noetherian functions. It is often useful to consider Noetherian functions over a larger domain than the generating chain, e.g.,

$h(x, y) = \exp(y) + \sin(x)$ with $h : \mathbb{R}^2 \to \mathbb{R}$. In this case, the domain of definition of the generating chain is implicitly extended by treating them as functions over the dimensionally larger domain, e.g., with $\exp(x, y), \sin(x, y) : \mathbb{R}^2 \to \mathbb{R}$ which ignore their first and second argument respectively. This is compatible with Def. 3.31 because the partial derivatives with respect to the ignored arguments is trivially zero. Proposition 3.32 gives important closure properties of the Noetherian functions generated by the same Noetherian chain, which are crucial later and explain the name *Noetherian* function [57, 201].

Proposition 3.32 ([12, 56, 57])**.** *The set R of Noetherian functions generated by a given Noetherian chain $h_1, \dots, h_r : H \subseteq \mathbb{R}^k \to \mathbb{R}$ is a Noetherian ring that is closed under partial derivatives.*

Proof. Let $y = (y_1, \dots, y_k), z = (z_1, \dots, z_r)$ abbreviate indeterminates as in Def. 3.31. The set R is a ring under the usual addition and multiplication of real-valued functions because the corresponding generating polynomials form a ring. Now, R is Noetherian because it is a finitely generated algebra [24, §2.11, Corollary 3] over the Noetherian polynomial ring $\mathbb{R}[y]$. The following constructive proof yields a computational method that is used later.

Consider an ascending chain of ideals $I_0 \subseteq I_1 \subseteq \cdots$ in R. For each I_i, associate the set of generating polynomials $J_i \overset{\text{def}}{=} \{p \mid p(y, h_1(y), \dots, h_r(y)) \in I_i\} \subseteq \mathbb{R}[y, z]$ with respect to the generating Noetherian chain. Each J_i is an ideal in $\mathbb{R}[y, z]$ because the corresponding I_i are themselves ideals. By construction, $J_i \subseteq J_{i+1}$ because $I_i \subseteq I_{i+1}$ for all i. Since $J_0 \subseteq J_1 \subseteq \cdots$ is an ascending chain of ideals in $\mathbb{R}[y, z]$, which is a Noetherian polynomial ring, it must stabilize at some N with $J_N = J_{N+1} = \cdots$. Correspondingly, the chain of ideals I_i stabilizes (at the latest) at N so R is Noetherian. The chain $I_0 \subseteq I_1 \subseteq \cdots$ may stabilize earlier than the corresponding J_i chain but that is not important here.

To show that R is closed under partial derivatives, let $h(y) = p(y, h_1(y), \dots, h_r(y)) \in R$ with $p \in \mathbb{R}[y, z]$. For the partial derivative of h with respect to y_i, applying the chain rule yields:

$$\begin{aligned}\frac{\partial h}{\partial y_i}(y) &= \frac{\partial p(y, h_1(y), \dots, h_r(y))}{\partial y_i} \\ &= \frac{\partial p}{\partial y_i}(y, h_1(y), \dots, h_r(y)) + \sum_{j=1}^{r} \frac{\partial p}{\partial z_j}(y, h_1(y), \dots, h_r(y)) \frac{\partial h_j}{\partial y_i}(y)\end{aligned}$$

By definition, $\frac{\partial p}{\partial y_i}(y, h_1(y), \dots, h_r(y)) \in R$ since $\frac{\partial p}{\partial y_i}$ is a polynomial in $\mathbb{R}[y, z]$. Each summand $\frac{\partial p}{\partial z_j}(y, h_1(y), \dots, h_r(y)) \in R$ since $\frac{\partial p}{\partial z_j}$ is a polynomial in $\mathbb{R}[y, z]$, and $\frac{\partial h_j}{\partial y_i}(y) \in R$ by definition because $h_1, \dots, h_r$ is a Noetherian chain. Hence, all RHS sub-terms are in R, and so $\frac{\partial h}{\partial y_i}(y) \in R$. □

Proposition 3.32 implies that adding Noetherian functions to their generating chains yields another Noetherian chain generating the same Noetherian ring R of Noetherian functions because R is closed under ring addition and multiplication. Beyond closure properties for a single Noetherian chain, the class of all Noetherian functions is also closed under other mathematical operations, including function composition, multiplicative inverses, and function inverses (with appropriate assumptions) [12]. Closure under function composition is proved constructively in Proposition 3.33 as it is used later.

Proposition 3.33 ([12]). *If $h : H \subseteq \mathbb{R}^k \to \mathbb{R}$ is Noetherian and $\upsilon : \Upsilon \subseteq \mathbb{R}^l \to \mathbb{R}^k$ has a compatible image $\upsilon(\Upsilon) \subseteq H$ where each component $\upsilon_i : \Upsilon \subseteq \mathbb{R}^l \to \mathbb{R}$ for $i = 1, \dots, k$ is Noetherian, then the function composition $f = h(\upsilon_1, \dots, \upsilon_k) : \Upsilon \subseteq \mathbb{R}^l \to \mathbb{R}$ is Noetherian.*

Proof. Let $y = (y_1, \dots, y_k), z = (z_1, \dots, z_r), \gamma = (\gamma_1, \dots, \gamma_l)$ abbreviate indeterminates. The composed function f is well-defined on Υ since $\upsilon(\Upsilon) \subseteq H$. By assumption, $h(y) = p(y, h_1(y), \dots, h_r(y))$ for some generating Noetherian chain $h_1, \dots, h_r : H \subseteq \mathbb{R}^k \to \mathbb{R}$ and polynomial $p \in [y, z]$. Since the union of Noetherian chains is Noetherian and by Proposition 3.32, assume without loss of generality, that the Noetherian functions υ_i for $i = 1, \dots, k$ are members of the same generating Noetherian chain $\upsilon_1, \dots, \upsilon_s : \Upsilon \subseteq \mathbb{R}^l \to \mathbb{R}$ with $k \leq s$. Putting these together, f can be written as: $f = p(\upsilon, f_1, \dots, f_r)$, where $\upsilon = (\upsilon_1, \dots, \upsilon_k)$ and the function compositions $f_i \overset{\text{def}}{=} h_i(\upsilon_1, \dots, \upsilon_k)$ for $i = 1, \dots, r$. From this representation, f is generated by polynomial p over the sequence:

$$\upsilon_1, \dots, \upsilon_s, f_1, \dots, f_r \tag{3.9}$$

In order to show that (3.9) is a Noetherian chain, it suffices to check that $f_1, \dots, f_r$ obey the condition on partial derivatives (3.8) because $\upsilon_1, \dots, \upsilon_s$ is already a Noetherian chain. For each $f_i(\gamma) : \Upsilon \subseteq \mathbb{R}^l \to \mathbb{R}$, taking the partial derivative with respect to γ_j and applying the chain rule:

$$\frac{\partial f_i}{\partial \gamma_j}(\gamma) = \frac{\partial h_i(\upsilon_1(\gamma), \dots, \upsilon_k(\gamma))}{\partial \gamma_j} = \sum_{l=1}^{k} \frac{\partial h_i}{\partial y_l}(\upsilon_1(\gamma), \dots, \upsilon_k(\gamma)) \frac{\partial \upsilon_l}{\partial \gamma_j}(\gamma)$$

It suffices to check that each sub-term appearing on the RHS sum are generated as polynomials over the sequence (3.9). The case for each $\frac{\partial \upsilon_l}{\partial \gamma_j}(\gamma)$ follows immediately because $\upsilon_1, \dots, \upsilon_s$ is a Noetherian chain. Since $h_1, \dots, h_r$ is a Noetherian chain, each $\frac{\partial h_i}{\partial y_l}$ is a polynomial combination $\frac{\partial h_i}{\partial y_l} = t_{il}(y, h_1, \dots, h_r)$ for some polynomial $t_{il} \in \mathbb{R}[y, z]$ and, thus, $\frac{\partial h_i}{\partial y_l}$ is generated by chain (3.9):

$$\frac{\partial h_i}{\partial y_l}(\upsilon_1, \dots, \upsilon_k) = t_{il}(\upsilon, h_1(\upsilon), \dots, h_r(\upsilon)) = t_{il}(\upsilon, f_1, \dots, f_r) \qquad \square$$

Before turning to the study of Noetherian functions in dL, it is helpful to first understand how they help with its differential equations reasoning. Polynomial ODEs are very expressive and earlier results [69, 105, 137] make use of polynomial ODEs to implicitly characterize (and thus, eliminate) real analytic functions appearing in initial value problems (IVPs). IVPs are specified by a system of ODEs, $x' = f(x)$, defined over domain D with RHS $f(x) : D \subseteq \mathbb{R}^n \to \mathbb{R}^n$ and real initial value $X_0 \in D \subseteq \mathbb{R}^n$. The IVP is called Noetherian (resp. polynomial) when all components of the RHS $f(x)$ are Noetherian functions (resp. polynomials). Both Noetherian and polynomial functions are analytic and therefore continuously differentiable. Under the assumption of continuously differentiable RHS, the Picard-Lindelöf theorem [204, §10.VI] guarantees that the IVP has a unique maximal solution $\varphi(t) : (\alpha, \beta) \to \mathbb{R}^n$ with $-\infty \leq \alpha < 0 < \beta \leq \infty$ such that $\varphi(0) = X_0$ and $\frac{\mathrm{d}\varphi(t)}{\mathrm{d}t} = f(\varphi(t))$. Uniqueness and maximality here means that every solution of the IVP is a truncation of φ to a smaller existence interval. The following generalizes aforementioned results [69, 105, 137] to the Noetherian setting:

Proposition 3.34. *Function $\varphi : (\alpha, \beta) \to \mathbb{R}^n$ with $-\infty \le \alpha < 0 < \beta \le \infty$ is the (coordinate-projected) solution of a Noetherian IVP iff it is the (coordinate-projected) solution of a polynomial IVP.*

Proof. In the (trivial) converse "$\Leftarrow$" direction, suppose function φ solves the polynomial IVP $x' = p(x, y), y' = q(x, y)$ with initial values $X_0 \in \mathbb{R}^n, Y_0 \in \mathbb{R}^r$. Let φ_x, φ_y denote the projection onto the x and y coordinates of φ respectively. Every solution of polynomial ODEs is a univariate Noetherian function [57]. Therefore, the Noetherian IVP given by $x' = p(\varphi_x(\tau), \varphi_y(\tau)), \tau' = 1$ with the same initial value for x and 0 for τ trivially has the (unique) solution $(\varphi_x(t), t) : (\alpha, \beta) \to \mathbb{R}^n \times \mathbb{R}$.

In the (nontrivial) "$\Rightarrow$" direction, suppose that φ is the solution to the Noetherian IVP $x' = f(x)$ where each $f_i(x) : D \subseteq \mathbb{R}^n \to \mathbb{R}$ is Noetherian, and with initial value $X_0 \in D$. By uniqueness of solutions, it suffices to construct a polynomial IVP so that $\varphi(t)$ solves it in the x coordinates. Since the union of Noetherian chains is itself a Noetherian chain, assume without loss of generality that the functions $f_1, \ldots, f_n$ are generated by the same Noetherian chain $h_1, \ldots, h_r$ and that $f_i = p_i(x, h_1(x), \ldots, h_r(x))$ for some polynomials $p_i \in \mathbb{R}[x, y]$ in $n + r$ indeterminates for $i = 1, \ldots, n$. Introduce new variables y_j for $j = 1, \ldots, r$ which are meant to take on the respective value of h_j along solutions to the ODE. Accordingly, the RHS of the Noetherian ODE is rewritten by replacing each f_i with $p_i(x, y)$, i.e., the desired polynomial ODEs for x is $x' = p(x, y)$.

It remains to ensure that each of these newly introduced variables y_j take on their intended values $h_j(\varphi(t))$ along the solution φ. By (3.8), the partial derivatives for each h_j can be written as polynomials $q_{ij} \in \mathbb{R}[x, y]$ over the generating Noetherian chain. By the chain rule:

$$\frac{\mathrm{d}h_j(\varphi(t))}{\mathrm{d}t} = \sum_{i=1}^{n} \frac{\partial h_j(x)}{\partial x_i}(\varphi(t)) \frac{\mathrm{d}\varphi_i(t)}{\mathrm{d}t} = \sum_{i=1}^{n} q_{ij}\big(\varphi(t), h_1(\varphi(t)), \ldots, h_r(\varphi(t))\big) \frac{\mathrm{d}\varphi_i(t)}{\mathrm{d}t}$$

Back-substituting into the RHS of this equation using the intended values for y_j and the new ODEs for x, yields the following additional ODEs for y:

$$y_j' = \sum_{i=1}^{n} q_{ij}(x, y) p_i(x, y)$$

The RHS of these additional ODEs are polynomials in $\mathbb{R}[x, y]$, which completes the desired polynomial IVP with the initial values $Y_0 \stackrel{\text{def}}{=} (h_1(x_0), \ldots, h_r(x_0)) \in \mathbb{R}^r$ for y. The construction of this polynomial IVP is correct-by-construction because of the mechanical chain rule computation. In particular, a solution to this IVP is given by the pair $(\varphi(t), y(t)) : (\alpha, \beta) \to \mathbb{R}^n \times \mathbb{R}^r$ where $y_j(t) \stackrel{\text{def}}{=} h_j(\varphi(t)) : (\alpha, \beta) \to \mathbb{R}$. By uniqueness of solutions, this completes the proof. □

In the "$\Rightarrow$" direction of Proposition 3.34, the constructed polynomial IVP may involve additional ODEs over the variables y (with their respective initial values). The number of additional equations required in this construction is the length of the shortest Noetherian chain that generates the RHS of the input Noetherian ODE. The polynomial IVP may have a larger maximal interval of existence than the input Noetherian IVP if it leaves the domain D of the input RHS. In the "$\Leftarrow$" direction, only one additional time variable τ is required. Consequently, the solution of *any* n-dimensional IVP that is the coordinate projection of the solution of a polynomial IVP (of potentially much larger dimension) is the coordinate projection of the solution of an $(n+1)$-dimensional Noetherian IVP.

The constructive proof of the "$\Rightarrow$" direction in Proposition 3.34 yields an approach for transforming input Noetherian IVPs to polynomial IVPs assuming that the Noetherian functions can be effectively associated with generating Noetherian chains and polynomials.

Example 3.35 (Flight dynamics [137, Equation 1])**.** A simple planar model of curved aircraft motion is given by the following ODE system, where (x, y) are the aircraft's planar coordinates, θ its angular orientation, and ν, ω its linear and angular velocity respectively [137, Equation 1]:

$$x' = \nu \cos(\theta), \quad y' = \nu \sin(\theta), \quad \theta' = \omega$$

Consider an IVP for this ODE with initial values $x = X_0, y = Y_0, \theta = \Theta_0 \in \mathbb{R}$. The linear and angular velocities ν, ω are left as symbolic constants in this model. The RHS of the ODE is generated by the Noetherian chain: $\sin(\theta), \cos(\theta)$. Introducing additional variables z_1, z_2 for the elements of this chain, and replacing the RHS for x', y' according to their generating polynomials with respect to the chain gives:

$$x' = \nu z_1, \quad y' = \nu z_2, \quad \theta' = \omega$$

A symbolic calculation (see Proposition 3.34) yields the following ODEs that z_1, z_2 must obey:

$$z_1' = \omega z_2, \quad z_2' = -\omega z_1$$

To finish constructing the polynomial IVP, set the initial values $z_1 = \sin(\Theta_0), z_2 = \cos(\Theta_0)$. The resulting ODE has higher dimension but a polynomial RHS. △

Proposition 3.34 shows that the utility of adding Noetherian functions to dL is *not* an increase in expressiveness of the differential equations. Rather, extended term languages allow Noetherian ODEs to be written down naturally instead of relying on implicit polynomial characterization such as in Example 3.35. More importantly, they make it possible to use formulas as ODE invariants that are *not* semialgebraic. By Theorems 3.29 and 3.30, the dL ODE axiomatization provides an effective and complete calculus for (dis)proving the resulting semianalytic ODE invariants involving Noetherian functions. This requires Noetherian functions to meet the extended term conditions from Section 3.2.4, which is shown next.

3.7.2 Extended Term Conditions for Noetherian Functions

Assume from now on that the fixed k-ary function symbols $h_1, \ldots, h_r$ are interpreted semantically as members of a Noetherian chain $h_1, \ldots, h_r : \mathbb{R}^k \to \mathbb{R}$ respectively. Recall that extended dL terms are formed syntactically from these function symbols according to the grammar (Section 3.2.1). The first two extended term conditions are straightforward to check:

(S) All Noetherian functions are, by definition, C^∞ smooth (even real analytic) so the semantics of differentials are well-defined.

(P) The partial derivative of each $h_j(y_1, \dots, y_k) : \mathbb{R}^k \to \mathbb{R}$ with respect to y_i satisfies (3.8) for some polynomial $q_{ij} \in \mathbb{R}[y, z]$. Since polynomials are generated by addition and multiplication, these partial derivatives $\frac{\partial h}{\partial y_i}(y_1, \dots, y_k)$ are syntactically represented by the extended term:

$$q_{ij}\big(y_1, \dots, y_k, h_1(y_1, \dots, y_k), \dots, h_r(y_1, \dots, y_k)\big)$$

Thus, Lemma 3.3 adds the (sound) differential axioms for each fixed function symbol h_j and therefore, all Lie derivatives are representable in the extended term language.

The final condition (R) is more involved and relies crucially on closure properties of Noetherian functions. A syntactic subtlety arises for extended terms with nested function applications such as $\exp(\exp(x))$. Its semantics is the iterated real exponential function generated by the 2-element Noetherian chain $\exp(x), \exp(\exp(x))$. Thus, even though the fixed function symbols $h_1, \dots, h_r$ form a Noetherian chain, the extended term grammar could produce extended terms that *do not* correspond to Noetherian functions generated by that chain. The following lemma resolves this issue by computing another (syntactic) Noetherian chain that generates it instead:

Lemma 3.36. *The semantics of every extended term e over Noetherian functions is a Noetherian function and e can be effectively associated with a (syntactic) Noetherian chain that generates it.*

Proof. By structural induction on extended dL term e. The cases for variables and constants are obvious, while the cases for addition and multiplication follow inductively from closure under ring operations (Proposition 3.32) and the fact that finite unions of Noetherian chains are Noetherian chains. The only difficult case is when e is a function composition $h(e_1, \dots, e_k)$, where $e_1, \dots, e_k$ are extended terms. Inductively, each $e_1, \dots, e_k$ semantically is a Noetherian function. Moreover, h is (semantically) a Noetherian function by the assumption that the interpretation of all fixed function symbols is an element of some Noetherian chain. Thus, Proposition 3.33 implies that the semantics of their composition is also a Noetherian function. Let $v_1, \dots, v_s$ be the union of Noetherian chains obtained inductively for $e_1, \dots, e_k$. The (constructive) proof of Proposition 3.33 shows that the Noetherian chain (3.9) given by $v_1, \dots, v_s, f_1, \dots, f_r$ generates $h(e_1, \dots, e_k)$, where $v_1, \dots, v_s$ are syntactically represented by extended terms using the induction hypothesis on $e_1, \dots, e_k$ and each $f_i \overset{\text{def}}{=} h_i(e_1, \dots, e_k)$ is an extended term by the extended term grammar (Section 3.2.1). □

Lemma 3.36 makes it possible to unambiguously refer to "the" generating Noetherian chain and polynomial for any extended term e by giving an effective procedure for finding a syntactic representation of such a generating chain in the extended term language. Together with Proposition 3.32, this suffices to prove that the extended term language has the computable differential radicals condition (R).

Theorem 3.37. *Term languages with Noetherian functions satisfy the extended term conditions.*

Proof. Conditions (S) and (P) have already been shown above. It remains to show condition (R), i.e., any ODE $x' = f(x)$ and extended term e has a computable (and provable) differential radical identity (3.2). By Lemma 3.36, the terms $f(x), e$ are (semantically) Noetherian and so, by taking the union of Noetherian chains, are defined by the same generating Noetherian chain $\upsilon_1, \ldots, \upsilon_s$. The ring R generated by this chain is Noetherian by Proposition 3.32 and is closed under partial derivatives. Recall that the Lie derivative of e along $x' = f(x)$ is given by:

$$\mathcal{L}_{f(x)}(e) \overset{\text{def}}{=} \sum_{i=1}^{n} \frac{\partial e}{\partial x_i} \cdot f_i(x)$$

Every sub-term on the RHS of the Lie derivative of e is contained in the ring R because the ring already contains the RHS of the ODEs, $f(x)$, and is closed under the partial derivatives of e. Inductively, all higher Lie derivatives $\overset{\bullet}{e}^{(i)}$ for $i = 0, 1, \ldots$ are contained in R and are therefore generated by the chain $\upsilon_1, \ldots, \upsilon_s$ with $\overset{\bullet}{e}^{(i)} = p_i(x, \upsilon_1, \ldots, \upsilon_s)$ for polynomials $p_i \in \mathbb{R}[x, y]$ and $y = (y_1, \ldots, y_s)$. Following the proof of Proposition 3.32, consider this ascending chain of polynomial ideals:

$$(p_0) \subseteq (p_0, p_1) \subseteq (p_0, p_1, p_2) \subseteq \cdots$$

This chain stabilizes with the provable polynomial identity $p_N = \sum_{i=0}^{N-1} q_i p_i$ for some polynomial cofactors $q_i \in \mathbb{R}[x, y]$ and $N \geq 1$. The rank N and the polynomial cofactors q_i are computable by successive ideal membership checks [63, 64, 103]. Mapping this back into elements of R gives the required provable differential radical identity for the Lie derivatives of e by choosing cofactors $g_i \overset{\text{def}}{=} q_i(x, \upsilon_1, \ldots, \upsilon_s)$:

$$\overset{\bullet}{e}^{(N)} = \sum_{i=0}^{N-1} g_i \overset{\bullet}{e}^{(i)} \qquad \square$$

An immediate corollary is that term language extensions with Noetherian functions inherit all earlier soundness and completeness results, e.g., from Sections 3.4.2 and 3.6.2.

Corollary 3.38 (Noetherian invariant completeness)**.** *The* dL *proof calculus is complete for (semi)analytic invariants of ODEs for term languages extended with Noetherian functions.*

Proof. Completeness for Noetherian functions follows immediately from the extended term conditions for Noetherian functions (Theorem 3.37) and the completeness theorems for term languages meeting those conditions (Theorems 3.29 and 3.30). □

Similarly, analytic hybrid programs with Noetherian functions inherit Corollary 3.14 with an additional restriction on the shape of analytic hybrid programs.

Corollary 3.39 (Noetherian analytic hybrid program completeness)**.** *For term languages extended with Noetherian functions, it is possible to compute an extended term e such that the equivalence $[\alpha]P \leftrightarrow e = 0$ is derivable in* dL *for any analytic formula P and analytic hybrid programs α, where α is either **loop-free** or **assignment-free**.*

Proof. The proof for loop-free α follows immediately from the proof of Corollary 3.14 because the only case of the structural induction that requires the Noetherian property of the extended term language is loops α^*. Otherwise, only (R) is required, which is shown in Theorem 3.37.

The proof for assignment-free α follows by strengthening the inductive hypothesis in the proof of Corollary 3.14. Firstly, by Lemma 3.36, the (finite) set of all terms appearing in α

are (semantically) Noetherian and so, by taking the union of Noetherian chains, are defined by the same generating Noetherian chain. The ring R generated by this chain is Noetherian by Proposition 3.32 and is closed under partial derivatives. The structural induction in Corollary 3.14 is strengthened to show that the extended term $\tilde{e}$ computed with derivable equivalence $[\alpha]e = 0 \leftrightarrow \tilde{e} = 0$ in each case (except assignments, which are assumed to not occur) is also contained in ring R. The cases for test $?d \neq 0$, choice $\alpha \cup \beta$, and sequential composition $\alpha;\beta$ are omitted because they follow immediately from the corresponding cases in the proof of Corollary 3.14 since the computed extended term $\tilde{e}$ in those cases are (inductively) in R. The non-trivial cases are ODEs $x' = f(x) \,\&\, d \neq 0$ and loops α^*.

- Case $x' = f(x) \,\&\, d \neq 0$. As with Corollary 3.14, the following equivalence is derived for $\tilde{e} = d(\sum_{i=0}^{N-1}(\dot{e}^{(i)})^2)$, where N is the rank of e.

  $$[x' = f(x) \,\&\, d \neq 0]e = 0 \leftrightarrow \tilde{e} = 0$$

  Note that d is in R by definition of R and each subterm $\dot{e}^{(i)}$ for $0 \leq i \leq N-1$ in the summand of $\tilde{e}$ is also in R by the closure of R under partial derivatives and because the RHS of ODE $x' = f(x)$ are in R by construction. Thus, $\tilde{e}$ is also in the ring R.

- Case α^*. As with Corollary 3.14, construct the sequence of terms $\tilde{e}_i$ defined inductively with $\tilde{e}_0 \overset{\text{def}}{=} e$ and $\tilde{e}_{i+1}$ is the term satisfying the derived equivalence $[\alpha]\tilde{e}_i = 0 \leftrightarrow \tilde{e}_{i+1} = 0$ obtained by applying the induction hypothesis on α with postcondition $\tilde{e}_i = 0$ for $i = 0, 1, 2, \ldots$. By the *strengthened* induction hypothesis, each $\tilde{e}_i$ in the sequence is in the Noetherian ring R. Thus, the following ascending chain of ideals in R stabilizes:

  $$(\tilde{e}_0) \subseteq (\tilde{e}_0, \tilde{e}_1) \subseteq (\tilde{e}_0, \tilde{e}_1, \tilde{e}_2) \subseteq \cdots$$

  In particular, for some computable k, the sequence satisfies: $\tilde{e}_k = \sum_{i=0}^{k-1} g_i \tilde{e}_i$. The equivalence $[\alpha^*]e = 0 \leftrightarrow \sum_{i=0}^{k-1} \tilde{e}_i^2 = 0$ is derived identically to the loop case for Corollary 3.14 using the iteration axiom $[^*]$ for the "$\rightarrow$" direction and the loop induction rule loop for the "$\leftarrow$" direction (derivation omitted). Observe that term $\sum_{i=0}^{k-1} \tilde{e}_i^2$ is in R because each subterm $\tilde{e}_i$ for $0 \leq i \leq k-1$ is in the ring R (shown above, from inductive hypothesis). □

Additional cases of hybrid programs enjoying analytic completeness can be proved on a case-by-case basis by building on Corollary 3.39, for example, given an analytic postcondition P for a choice program $[\alpha \cup \beta]P$ between loop-free analytic program α and assignment free analytic program β, a suitable equivalence can be derived by first using equivalence axiom $[\cup]$ and then separately characterizing the resulting conjuncts in $[\alpha]P \wedge [\beta]P$.

3.7.3 Extended Term Language Example

This section illustrates the constructions from Sections 3.7.1 and 3.7.2 using the extended term language (3.1). The first example shows the computations from Lemma 3.36 and Theorem 3.37:

Example 3.40 (Syntactic manipulation of Noetherian functions). Consider the extended term ODE $x' = \exp(\sin(x))$ and the polynomial term $e = x + x^2$. The Noetherian chain for e is empty because it is already a polynomial while the Noetherian chain associated with $\exp(\sin(x))$ is $\upsilon_1 = \sin(x), \upsilon_2 = \cos(x), \upsilon_3 = \exp(\sin(x))$. The higher Lie derivatives of e are all extended terms generated by the chain $\upsilon_1, \upsilon_2, \upsilon_3$:

$$\begin{aligned}
\overset{\bullet}{e}^{(1)} &= \upsilon_3 + 2\upsilon_3 x \\
\overset{\bullet}{e}^{(2)} &= \upsilon_3^2\upsilon_2 + 2(\upsilon_3^2\upsilon_2 x + \upsilon_3^2) = (2\upsilon_3 + \upsilon_2\upsilon_3 + 2\upsilon_2\upsilon_3 x)((1+2x)\overset{\bullet}{e}^{(1)} - 4\upsilon_3 e)
\end{aligned}$$

The (polynomial) identity for $\overset{\bullet}{e}^{(2)}$ in terms of $\overset{\bullet}{e}^{(1)}, e$ and their cofactors is obtained computationally by ideal membership checks for the polynomial ring $\mathbb{R}[x, y_1, y_2, y_3]$ (the indeterminate y_i corresponds to υ_i for $i = 1, 2, 3$), following Proposition 3.32. △

The next example illustrates how the extended term language allows effective proofs of more invariants than possible with polynomial term languages.

Example 3.41 (Expressivity of Noetherian invariants)**.** The polynomial invariant $1-u^2-v^2 = 0$ was proved for the ODE α_e from (2.1) in Example 3.7. With respect to Fig. 2.1, this means that a trajectory starting at the point $(1, 0)$ stays on the circle. However, this invariant yields no information about how fast the trajectory loops around the circle or whether it revolves clockwise or anti-clockwise. In the extended term language, the most precise invariant can be proved, namely the solution to the ODEs from this initial point. The solution is a trigonometric function of time (given below), and so cannot be expressed as a polynomial (or semialgebraic) invariant [14]. The precise solution also shows that the motion is anti-clockwise, as suggested by Fig. 2.1.

The following derivation uses a DC to add the known polynomial invariant $1 - u^2 - v^2 = 0$ which proves by dbx as in Example 3.7. The abbreviated premise after the differential cut assumes $1 - u^2 - v^2 = 0$ in the ODE's domain constraint. It is abbreviated ① and continued below.

$$\text{DC}\,\dfrac{\text{dbx},\mathbb{R}\,\dfrac{*}{u = 1, v = 0, t = 0 \vdash [\alpha_e, t'{=}1]1 - u^2 - v^2 = 0} \qquad \text{①}}{u = 1, v = 0, t = 0 \vdash [\alpha_e, t'{=}1](u{-}\cos(t) = 0 \land v{-}\sin(t) = 0)}$$

From ①, first calculate the Lie derivatives, abbreviating $c = u - \cos(t), s = v - \sin(t)$:

$$\begin{aligned}
\mathcal{L}_{\alpha_e,t'=1}(c) &= \mathcal{L}_{\alpha_e,t'=1}(u - \cos(t)) = -v + \frac{u}{4}(1 - u^2 - v^2) + \sin(t) = -s + \frac{u}{4}(1 - u^2 - v^2) \\
\mathcal{L}_{\alpha_e,t'=1}(s) &= \mathcal{L}_{\alpha_e,t'=1}(v - \sin(t)) = u + \frac{v}{4}(1 - u^2 - v^2) - \cos(t) = c + \frac{v}{4}(1 - u^2 - v^2)
\end{aligned}$$

Under the domain constraint assumption $1 - u^2 - v^2 = 0$, the additional $1 - u^2 - v^2$ term in both Lie derivatives simplifies to 0. The derivation starts with a cut of the postcondition $c = 0 \land s = 0$. This arithmetic premise, abbreviated ②, is discussed afterwards. Continuing on the right premise, the vdbx step closes successfully using real arithmetic manipulations only:

$$\text{cut}\,\dfrac{\text{②} \qquad \text{vdbx}\,\dfrac{\mathbb{R}\,\dfrac{*}{1 - u^2 - v^2 = 0 \vdash \begin{pmatrix}\overset{\bullet}{c} \\ \overset{\bullet}{s}\end{pmatrix} = \begin{pmatrix}0 & -1 \\ 1 & 0\end{pmatrix}\begin{pmatrix}c \\ s\end{pmatrix}}}{c = 0 \land s = 0 \vdash [\alpha_e, t'{=}1 \,\&\, 1 - u^2 - v^2 = 0](c{=}0 \land s{=}0)}}{u = 1, v = 0, t = 0 \vdash [\alpha_e, t'{=}1 \,\&\, 1 - u^2 - v^2 = 0](c{=}0 \land s{=}0)}$$

The premise ② is valid, but it requires properties of the trigonometric functions ($\cos(0) = 1$, $\sin(0) = 0$) so it cannot be proved using $\mathbb{R}$. Instead, extended arithmetic $\mathbb{R}_{\exp,\sin,\cos}$ is needed:

$$\mathbb{R}_{\exp,\sin,\cos}\frac{*}{u = 1, v = 0, t = 0 \vdash u - \cos(t) = 0 \land v - \sin(t) = 0}$$

The extended arithmetic theory is undecidable in general [162] and so, unlike $\mathbb{R}$, rule $\mathbb{R}_{\exp,\sin,\cos}$ cannot be implemented via an underlying decision procedure. Yet, simple arithmetic questions such as ② which just involve the evaluation of trigonometric functions can be easily checked.

The above derivation takes advantage of a known Darboux equality for $1 - u^2 - v^2$ to simplify the proof using vdbx. The proof could have instead directly made use of Theorem 3.13 by encoding $c = 0 \land s = 0$ as $c^2 + s^2 = 0$ and then calculating the rank of $c^2 + s^2$ (which involve trigonometric functions) according to Theorem 3.37. This also works, but $c^2 + s^2$ has rank 3, and the resulting cofactors are too large to even fit on this page. △

The final example below highlights an important insight from Proposition 3.34: even though this chapter only considers extended term languages with terms that are defined everywhere, it is possible to use logical formulas to implicitly characterize more terms, making use of closure properties of the Noetherian functions [12]. The following example illustrates implicit characterization of quotients which are defined everywhere in the domain of interest:

Example 3.42 (Implicit characterization of quotients)**.** The trigonometric tangent function $\tan(x)$ is Noetherian and defined on the interval $(-\frac{\pi}{2}, \frac{\pi}{2})$. Consider the following "formula" where x is restricted in the domain constraint so that the RHS $\tan(x)$ is always defined:

$$x = \frac{1}{2} \rightarrow [x' = \tan(x) \,\&\, -1 \leq x \leq 1]x \geq \frac{1}{2}$$

This "formula" is not formally in the syntax of dL formulas because tan is not defined everywhere. However, Proposition 3.34 can be used to ask an equivalent question in dL. Recall from calculus:

$$\tan(x) = \frac{\sin(x)}{\cos(x)} \qquad \frac{\partial \frac{1}{\cos(x)}}{\partial x} = \frac{\sin(x)}{(\cos(x))^2}$$

Thus, $\sin(x), \cos(x), \frac{1}{\cos x}$ forms a 3-element Noetherian chain that generates $\tan(x)$. For brevity, by partially following the IVP construction of Proposition 3.34, the "formula" is rephrased as an actual dL formula with y representing $\frac{1}{\cos x}$ along the ODE. After replacing $x' = \sin(x)y$, the required differential equation for y is calculated with $y' = \sin(x)y^2(\sin(x)y) = \sin^2(x)y^3$.

$$x = \frac{1}{2} \land \cos(x)y - 1 = 0 \rightarrow [x' = \sin(x)y, y' = \sin^2(x)y^3 \,\&\, -1 \leq x \leq 1]x \geq \frac{1}{2}$$

For non-zero denominator, the initial value $\frac{1}{\cos(x)}$ of y is logically characterized by the formula $\cos(x)y - 1 = 0$. The following Lie derivative calculation shows that $\cos(x)y - 1$ satisfies a Darboux equality and so $\cos(x)y - 1 = 0$ can be proven invariant along the ODE (abbreviated as α) by dbx.

$$\mathcal{L}_\alpha(\cos(x)y - 1) = -\sin(x)(\sin(x)y)y + \cos(x)(\sin^2(x)y^3) = \sin^2(x)y^2(\cos(x)y - 1)$$

The rephrased formula proves after a DC with this Darboux invariant for y using the ODE invariant $x \geq \frac{1}{2}$ and rule sAI (or its generalization with domain constraints from Appendix A.2). Briefly, the invariance of $x \geq \frac{1}{2}$ provably reduces to the following arithmetic premise which is valid and falls within a *decidable* fragment of arithmetic with trigonometric functions [112]:

$$-1 \leq x \leq 1, \cos(x)y - 1 = 0, x = \frac{1}{2} \vdash \sin(x)y > 0 \qquad \triangle$$

3.8 Related Work

This related work discussion focuses on deductive safety and invariance verification for differential equations. Readers interested in ODEs [204], real analysis [94, 204], algebra [24], and real algebraic geometry [14] are referred to the cited textbooks. The orthogonal task of efficiently generating invariants is investigated elsewhere [63, 103, 169].

Proof Rules for ODE Invariants. Numerous useful but incomplete proof rules for ODE invariants [140, 155, 169, 190] are surveyed elsewhere [64]. The soundness and completeness theorems for dRI and sAI were previously proved semantically [63, 103]. These earlier results are limited to (semi)algebraic invariants as they depend on specific semantic properties limited to polynomials. The extended term conditions (Section 3.2.4) and Noetherian functions (Section 3.7) generalize these results, showing that all (semi)analytic invariance questions reduce completely to arithmetic. In their original presentation [63, 103], dRI and sAI are *algorithmic procedures* for checking invariance of semialgebraic sets, requiring e.g., checking ideal membership for all polynomials in the semialgebraic decomposition. This makes them difficult to implement soundly within a small, trusted axiomatic core [54]. This chapter shows that, by relying on the logic dL, these rules can be *derived* from a small set of axiomatic principles. Although these derivations also leverage ideal computations, they are only used in *derived rules*. With the aid of a theorem prover like KeYmaera X, derived rules can be implemented as tactics that crucially remain *outside* its soundness-critical axiomatic core.

Deductive Power and Proof Theory. The derivations shown in this chapter are fully general, which is necessary for completeness of the resulting derived rules. The number of conjuncts in the progress and differential radical formula for an extended term e is equal to the rank of e. Known upper bounds for the rank, even in the case of polynomials in n variables, are doubly exponential in $n^2 \ln n$ [125]. Many simpler classes of invariants can be proved using simpler derivations, as exemplified by Examples 3.7 and 3.41. This is where a study of the deductive power of sound, but incomplete, proof rules [64] is essential. For ODE invariants of a simpler class, it suffices to use a proof rule that is complete for just that class, for example, with the proof rules in Section 3.3 that derive from DG. This intuition is echoed in an earlier study [140] of the relative deductive power of differential invariants (DI), differential cuts (DC), and differential ghosts (DG). The first completeness result (Theorem 3.13) shows that dL with DG is complete for algebraic and analytic invariants. Other proof-theoretical studies of dL [139] reveal surprising correspondences between its hybrid, continuous, and discrete aspects in the sense that each aspect can be axiomatized completely and effectively relative to any other aspect.

Noetherian Functions. This chapter only touched on basic properties of Noetherian functions. The model-theoretic study of Noetherian functions and the related Pfaffian functions is fascinating in its own right [12, 56, 57]. Pfaffian functions are generated by chains satisfying (3.8) except with triangular dependencies in their partial derivatives [57], and notably, the expansion of the real field with Pfaffian functions is o-minimal [185, 203, 207]. Such o-minimal expansions have been studied in reachability analysis for *o-minimal hybrid systems* [92, 96] because they admit the construction of finite bisimulations for reachability analysis algorithms. In contrast, expansions with (more general) Noetherian functions, e.g., (unrestricted) trigonometric sine and cosine, are not o-minimal because they can be used to characterize the natural numbers. This is a barrier to the construction of finite bisimulations [96] but not for deductive approaches, as long as the relevant arithmetic is provable.

Undecidability of arithmetic is a delicate issue [162], but this chapter's completeness results show that ODE invariance verification *completely* reduces to arithmetic! Many (necessarily incomplete) approaches and tools for handling special functions are available, e.g., resolution with upper and lower bounds as implemented in MetiTarski [4], δ-decidability as implemented in dReal [59, 60], or heuristic inference-based approaches as implemented in Polya [10]. Specialized decision procedures may also be applicable for restricted fragments of arithmetic [112], as in Examples 3.40 and 3.42. Even in settings where all of these automated tools fail to verify an arithmetic question, the system designer can provide further mathematical intuition with an interactive proof in KeYmaera X [54].

The findings of this chapter identify Noetherian functions as a more general unifying theme behind earlier results in continuous/hybrid systems verification. Besides the completeness results for invariants, Proposition 3.34 also generalizes earlier results [69, 105, 137] to the Noetherian setting. This idea is called *differential axiomatization* [137] because it axiomatizes ODEs involving special functions that have undecidable arithmetic using polynomial ODEs. Similarly, [105, Proposition 1] gives an algorithm for replacing a fixed set of functions appearing in IVPs with polynomials ones. The result from [69, Theorem 4] only applies in the case of univariate Noetherian functions.

3.9 Discussion

This chapter demonstrates the impressive deductive power of differential ghosts: they prove *all* Darboux invariants and, as a consequence, *all* analytic invariants for extended term languages with the extended term conditions. Even *scalar* differential ghosts suffice for this result, but the question of whether their deductive power extends to even larger classes of invariants is left open. The chapter then introduces extensions to the dL axiomatization and shows how they can be used to extend completeness to semianalytic invariance. The case of (semi)algebraic invariants is even decidable, but the results prove completeness for much larger classes of (semi)analytic invariants. Table 3.1 gives an instructive overview of the key mathematical properties of solutions and terms that the soundness of each differential equation axiom rests on. With these axioms, mathematical reasoning for differential equations can be carried out *syntactically* and *axiomatically* within the dL proof calculus. This concise and foundational axiomatization of mathematical properties is precisely what enables generalizations of earlier results [148] to the (semi)analytic setting with

Table 3.1: Properties of ODE solutions underlying the differential equation axioms of dL.

ODE Axiom	**Mathematical Property**
DI	Mean value theorem
DC	Prefix-closure of solutions
DG	Picard-Lindelöf theorem
Cont	Existence of solutions
Uniq	Uniqueness of solutions
Dadj	Group action on solutions
RI	Completeness of field $\mathbb{R}$

Noetherian functions. A subtle question is left open: the extended term conditions in Section 3.2.4 do not require real analyticity for the fixed function symbols h even if Noetherian functions are always real analytic. This suggests that there may still be a gap between the extended term conditions and Noetherian functions. Are there C^∞ smooth (or even real analytic) functions that meet the extended term conditions but are *not* Noetherian functions? In other words, are Noetherian functions exactly the class of functions for which completeness results are possible? Certainly, this chapter's completeness results continue to hold for any functions meeting those conditions, which would make both positive and negative results interesting.

Chapter 4

Liveness and Existence for Ordinary Differential Equations

This chapter turns to deductive verification for *liveness* and *existence* properties of ordinary differential equations (ODEs), i.e., the question whether an ODE solution exists for long enough to reach a given region without leaving its domain of evolution. Numerous subtleties complicate the generalization of discrete liveness verification techniques, such as loop variants, to the continuous setting. For example, ODE solutions may blow up in finite time or their progress towards the goal may converge to zero. These subtleties are handled in dL by successively *refining* ODE liveness properties using ODE safety properties, thereby building on the complete dL axiomatization of invariants from Chapter 3. A special case of this approach is used to deduce (global) existence of solutions for ODEs which are fundamental hypotheses behind ODE liveness arguments. Proofs of global existence of solutions also help to justify the adequacy of ODE models for real-world systems because those systems (typically) do not simply cease to exist after a short time. These liveness and existence derivations are put into practice through an implementation in KeYmaera X. Together with Chapter 3's ODE safety proofs, this implementation provides the practical basis for proving dL stability specifications involving nested, first-order, and modal quantification over continuous and hybrid programs in the subsequent chapters.

4.1 Introduction

The ODE liveness specification $\Gamma \vdash \langle x' = f(x) \,\&\, Q \rangle \phi$ says that, for each initial state satisfying assumptions Γ, *some* state reached by following the ODE $x' = f(x) \,\&\, Q$ from that initial state satisfies the postcondition ϕ. Such liveness questions are dual to safety questions from Chapter 3: recall that ODE safety questions arise when proving that a continuous or hybrid system *always* satisfies a desired safety property; dually, ODE liveness questions arise when proving that those systems *eventually* reach a desired goal or target region. This form of ODE liveness is in the sense of Owicki and Lamport [128] for concurrent programs within their (linear) temporal logic. Liveness for ODEs has sometimes been called *eventuality* [157, 176] and *reachability* [191]. To minimize ambiguity, this chapter refers to the diamond modality formula $\langle x' = f(x) \,\&\, Q \rangle \phi$ as ODE *liveness*, while other related notions are discussed in Section 4.7.

Table 4.1: Surveyed ODE liveness arguments with highlighting in blue for soundness-critical corrections identified in this chapter. The applications (and corrections, if any) for each surveyed argument is briefly described here. The referenced corollaries are corresponding derived proof rules with details of the corrections.

Application	Without Domain Constraints	
Hybrid systems verification [137]	OK	(Cor. 4.16)
Automated ODE verification [156, 157]	[157, Remark 3.6] is incorrect	
Finding basin of attraction [159]	if chosen set is compact	(Cor. 4.23)
Staging set-based liveness proofs [176]	OK	(Cor. 4.20)
Switching logic synthesis [191]	if ODE solutions assumed or proved global	(Cor. 4.18)
Application	**With Domain Constraints**	
Hybrid systems verification [137]	if domain open/closed, target initially false	(Cor. 4.25)
Automated ODE verification [156, 157]	if arithmetical conditions checked globally	(Cor. 4.31)
Finding basin of attraction [159]	if chosen set is compact	(Cor. 4.27)
Staging set-based liveness proofs [176]	OK	(Cor. 4.28)
Switching logic synthesis [191]	if ODE solutions assumed or proved global	(Cor. 4.26)

For discrete systems, methods for proving liveness are well-known: loop variants show that discrete loops eventually reach a desired goal [73], while temporal logic is used to specify and study liveness properties in concurrent and infinitary settings [109, 128]. However, the deduction of (continuous) ODE liveness properties is hampered by several difficulties: *i)* solutions of ODEs may converge towards a goal without ever reaching it, *ii)* solutions of nonlinear ODEs may blow up in finite time leaving insufficient time for the goal to be reached, and *iii)* the goal may be reachable but only by illegally leaving the evolution domain constraint. Motivated by these difficulties, this chapter uses dL to perform systematic, step-by-step *refinement* of ODE liveness properties, where each refinement step is justified using an ODE safety (and invariance) property. Notably, liveness proofs focus on high-level refinement arguments while their underlying ODE safety justifications are handled transparently using Chapter 3's *complete* ODE invariance proof rules. Indeed, using safety to deduce liveness is a well-known proof technique for (discrete) concurrent systems [109, 128] and this chapter shows that those techniques generalize to the continuous setting—as long as the aforementioned difficulties are appropriately handled.

To demonstrate the applicability of the deductive refinement approach, this chapter surveys several arguments from the literature and derives them all as (corrected) dL proof rules, see Table 4.1. This logical presentation has two key benefits:

- The proof rules are *syntactically derived* from sound axioms of dL, which guarantees their correctness. Many of the surveyed arguments contain subtle soundness errors; rather than diminishing the surveyed work, these errors emphasize the need for an axiomatic, sound, and uniform way of analyzing ODE liveness instead of relying on ad hoc approaches.

- The approach identifies common refinement steps that form a basis for the surveyed liveness arguments drawn from various applications. This library of building blocks enables sound development and justification of new ODE liveness proof rules, e.g., by generalizing individual refinement steps or by exploring different combinations of those steps, e.g., in Corollaries 4.19, 4.21, and 4.30.

Another key insight is that all of the surveyed liveness arguments are based on reducing liveness properties of ODEs to assumptions about sufficient existence duration for their solutions. In fact, many of those arguments become significantly simpler (and sound) when the ODEs of concern are assumed to have global solutions, i.e., they do not blow up in finite time. It is reasonable and commonplace to make such an assumption for the continuous dynamics in models of CPSs [6, Section 6]. For example, control systems are designed to always operate near stable equilibria and they always have global solutions near those equilibria [71, Theorem 3.1]. Logically though, making an *a priori* assumption of global existence for ODEs means that the correctness of any subsequent verification results for the ODEs and hybrid system models are conditional on an unproved existence duration hypothesis. While global existence is known to hold for linear systems, even the simplest nonlinear ODEs (see Section 4.3) fail to meet the hypothesis of having global solutions without further assumptions. This chapter therefore adopts the view that (global) existence should be *proved* rather than *assumed* for the continuous dynamics in hybrid system models.

- Section 4.3 presents deductive dL proofs of global existence for ODE solutions. Together with the liveness proofs of Sections 4.4 and 4.5, this yields *unconditional* proofs of ODE liveness properties within the refinement framework, without existence presuppositions.
- Section 4.6 draws further practical insights from Sections 4.3–4.5 by implementing their liveness proof rules as tactics in KeYmaera X. This includes: *i)* the design of proof rules that are practically useful and well-suited for implementation (Section 4.6.1) and *ii)* the design of proof support to aid users in existence and liveness proofs (Section 4.6.2).

The liveness proofs of Sections 4.3–4.5 fit particularly well with an implementation in KeYmaera X because axiomatic refinement closely mirrors KeYmaera X's design principles. KeYmaera X implements dL's uniform substitution calculus [142] with a minimal, soundness-critical trusted kernel; on top of this, KeYmaera X's non-soundness-critical tactics framework [55] adds support and automation for proofs. Liveness proofs are similarly based on a series of small refinement steps which are, in turn, implemented as tactics based on a small basis of derived refinement axioms. More complicated liveness arguments, such as those from Table 4.1 or from new user insights, are implemented by piecing those tactics together using tactic combinators [55]. The implementation required minor changes to $\approx$155 lines of soundness-critical code in KeYmaera X, while the remaining $\approx$1500 lines implement ODE existence and liveness proof rules as non-soundness-critical tactics. These additions suffice to prove all of the examples in this chapter and in ODE models elsewhere [22, 176] (Section 4.6.2).

Reminder (Extended Term Language). This chapter uses an extended dL term language following the extended term conditions and notational conventions of Section 3.2 because the dL axiomatization remains sound for all extended term languages meeting those conditions.

Contribution. *The material for this chapter is drawn from Tan and Platzer [192, 195].*

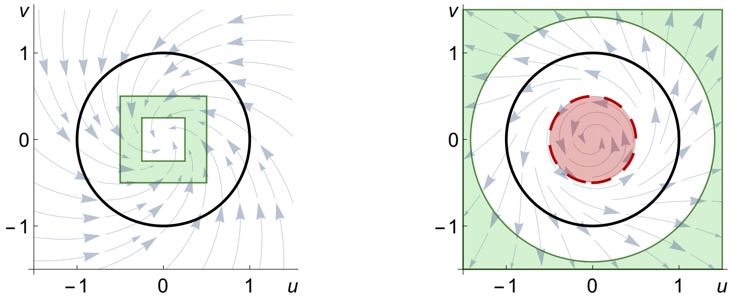

Figure 4.1: Visualization of α_l (left) and α_n (right). Solutions of α_l globally spiral towards the origin. In contrast, solutions of α_n spiral inwards within the inner red disk (dashed boundary), but spiral outwards otherwise. For both ODEs, solutions starting on the black unit circle eventually enter their respective shaded green goal regions. The ODE α_n also exhibits finite-time blow up of solutions from all initial states outside the red disk.

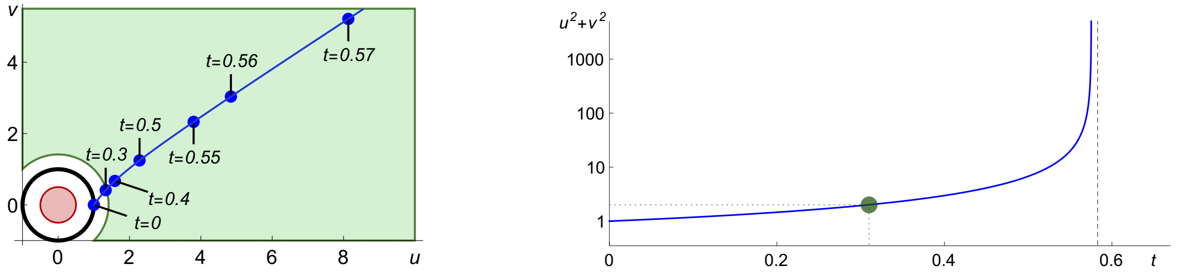

Figure 4.2: Two views of the ODE α_n evolving from initial state $u = 1, v = 0$ over time t. The left plot shows its trajectory in the u, v plane (cf. Fig. 4.1) while the right plot shows the squared Euclidean norm $u^2 + v^2$ evolving over time t (with logarithmic scaling for the vertical axis). The solution blows up in finite time with norm approaching ∞ as t approaches 0.58 (rounded up, black dashed asymptote). Nevertheless, the solution reaches the green goal region $u^2 + v^2 \geq 2$ from Fig. 4.1 at $t \approx 0.31$ (rounded up, green dot) before blowing up.

4.2 ODE Liveness via Box Refinements

This section explains step-by-step refinement for proving ODE liveness properties in dL. The following two running example ODEs α_l and α_n are visualized in Fig. 4.1 with directional arrows corresponding to their RHS evaluated at points on the plane:

$$\alpha_l \equiv u' = -v - u, v' = u - v \tag{4.1}$$

$$\alpha_n \equiv u' = -v - u(\frac{1}{4} - u^2 - v^2), v' = u - v(\frac{1}{4} - u^2 - v^2) \tag{4.2}$$

The ODE α_l is *linear* because its RHS depends linearly on variables u and v while α_n is *nonlinear* because of the cubic terms in its RHS. The nonlinearity of α_n results in more complex behavior for its solutions, e.g., the difference in spiraling behavior inside or outside the red disk shown in Fig. 4.1. In fact, solutions of α_n blow up in finite time iff they start outside the disk characterized by $u^2 + v^2 \leq \frac{1}{4}$, whereas finite-time blow up is impossible for linear ODEs like

α_l [33, 204]. An illustration of finite-time blow up for α_n from an initial state outside the red disk is shown in Fig. 4.2. This phenomenon is precisely defined and investigated in Section 4.3, which enables formal proofs of the aforementioned (absence of) finite-time blow up.

Figure 4.1 suggests that formulas[1] $\langle\alpha_l\rangle\big(\frac{1}{4} \leq \|(u,v)\|_\infty \leq \frac{1}{2}\big)$ and $\langle\alpha_n\rangle u^2 + v^2 \geq 2$ are true for initial states ω on the unit circle. These liveness properties are rigorously proved in Examples 4.17 and 4.22 respectively, using the refinement approach discussed next.

4.2.1 Liveness Refinement

Suppose that an initial liveness property $\langle x' = f(x) \,\&\, Q_0\rangle P_0$ is known for the ODE $x' = f(x)$. How could this be used to prove a desired liveness property $\langle x' = f(x) \,\&\, Q\rangle P$ for that ODE? Logically, this amounts to proving the following implication:

$$\langle x' = f(x) \,\&\, Q_0\rangle P_0 \rightarrow \langle x' = f(x) \,\&\, Q\rangle P \tag{4.3}$$

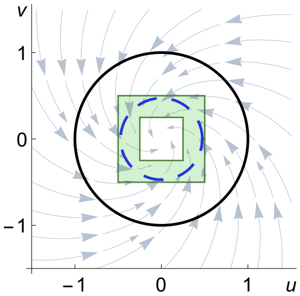

Proving implication (4.3) *refines* knowledge of the initial liveness property to the desired liveness property. As an example of such a refinement, consider the desired liveness property $\langle\alpha_l\rangle\big(\frac{1}{4} \leq \|(u,v)\|_\infty \leq \frac{1}{2}\big)$ for ODE α_l (4.1) starting from the initial circle $u^2 + v^2 = 1$ (cf. Fig. 4.1). Suppose the initial liveness property $\langle\alpha_l\rangle u^2 + v^2 = \frac{1}{4}$ is already proved, e.g., using the techniques of Section 4.4. As visualized on the right, ODE solutions starting from the black circle $u^2 + v^2 = 1$ eventually reach the dashed blue circle $u^2 + v^2 = \frac{1}{4}$. Since the blue circle is entirely contained in the green goal region, solutions that reach it must (trivially) also reach the goal region. Formally, the following instance of implication (4.3) is provable by monotonicity M$\langle\cdot\rangle$ because the implication $P_0 \rightarrow P$ between their respective postconditions is provable by $\mathbb{R}$.

$$\langle\alpha_l\rangle \underbrace{\left(u^2 + v^2 = \frac{1}{4}\right)}_{P_0} \rightarrow \langle\alpha_l\rangle \underbrace{\left(\frac{1}{4} \leq \|(u,v)\|_\infty \leq \frac{1}{2}\right)}_{P} \tag{4.4}$$

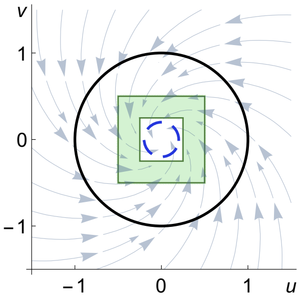

Similarly, if the implication between domain constraints $Q_0 \rightarrow Q$ is provable, then implication (4.3) is proved by monotonicity, because any solution staying in the smaller domain Q_0 must also stay in the larger domain Q. However, neither of these monotonicity-based arguments are sufficiently powerful for liveness proofs because they do not account for the specific ODE $x' = f(x)$ under consideration at all. Returning to the ODE α_l, suppose instead that the initial (known) liveness property is $\langle\alpha_l\rangle u^2 + v^2 = \frac{1}{25}$. This is visualized on the right with a smaller dashed blue circle. The following instance of implication (4.3) is also valid for solutions starting from the black circle $u^2 + v^2 = 1$, but it does *not* follow from a straightforward monotonicity argument

[1] $\|\cdot\|_\infty$ denotes the supremum norm, with $\|x\|_\infty \equiv \max_{i=1}^n |x_i|$ for an n-dimensional vector x. The inequality $\|(u,v)\|_\infty \leq \frac{1}{2}$ is expressible in first-order real arithmetic as $u^2 \leq \frac{1}{4} \wedge v^2 \leq \frac{1}{4}$. Similarly, $\frac{1}{4} \leq \|(u,v)\|_\infty$ is expressible as $\frac{1}{16} \leq u^2 \vee \frac{1}{16} \leq v^2$.

because the smaller dashed blue circle $u^2 + v^2 = \frac{1}{25}$ is not contained in the green goal region, i.e., implication $P_0 \to P$ is not valid.

$$\langle \alpha_l \rangle \underbrace{\left(u^2 + v^2 = \frac{1}{25}\right)}_{P_0} \to \langle \alpha_l \rangle \underbrace{\left(\frac{1}{4} \leq \|(u,v)\|_\infty \leq \frac{1}{2}\right)}_{P} \tag{4.5}$$

Instead, a proof of implication (4.5) requires additional information about solutions of the ODE α_l, namely, that they are continuous and the system α_l is planar. Informally, observe that it is impossible to draw a line (without lifting your pen off the page) that connects the black circle to the (smaller) dashed blue circle without crossing the green goal region. The continuous solutions of α_l are analogous to such lines and therefore must enter the green goal region before reaching the blue circle. To formalize such reasoning, this chapter's approach is built on refinement axioms that conclude instances of implication (4.3), like (4.4) and (4.5), from box modality formulas involving the ODE $x' = f(x)$.

4.2.2 Liveness Refinement Axioms

The following are four ODE refinement axioms of dL that are used for the approach. Crucially, these axioms are *derived* from their corresponding box modality axioms by exploiting the logical duality between the box and diamond modalities of dL. This makes it possible to build liveness arguments from dL's sound and parsimonious logical foundation.

Lemma 4.1 (Diamond ODE refinement axioms)**.** *The following $\langle\cdot\rangle$ ODE refinement axioms are derivable in* dL*. In axioms BDG$\langle\cdot\rangle$, DDG$\langle\cdot\rangle$, $y = (y_1, \ldots, y_m)$ is an m-dimensional vector of fresh variables (not appearing in x) and $g(x,y)$ is a corresponding m-dimensional vector of terms. Terms $e(x), L(x), M(x)$ and formulas $P(x), Q(x)$ are dependent only on free variables x (and not y).*

K$\langle\&\rangle$ $[x' = f(x) \,\&\, Q \wedge \neg P]\neg G \to \big(\langle x' = f(x) \,\&\, Q\rangle G \to \langle x' = f(x) \,\&\, Q\rangle P\big)$

DR$\langle\cdot\rangle$ $[x' = f(x) \,\&\, R]Q \to \big(\langle x' = f(x) \,\&\, R\rangle P \to \langle x' = f(x) \,\&\, Q\rangle P\big)$

BDG$\langle\cdot\rangle$ $[x' = f(x), y' = g(x,y) \,\&\, Q(x)]\ \|y\|_2^2 \leq e(x)$
$\to \big(\langle x' = f(x) \,\&\, Q(x)\rangle P(x) \to \langle x' = f(x), y' = g(x,y) \,\&\, Q(x)\rangle P(x)\big)$

DDG$\langle\cdot\rangle$ $[x' = f(x), y' = g(x,y) \,\&\, Q(x)]\, 2y \cdot g(x,y) \leq L(x)\,\|y\|_2^2 + M(x)$
$\to \big(\langle x' = f(x) \,\&\, Q(x)\rangle P(x) \to \langle x' = f(x), y' = g(x,y) \,\&\, Q(x)\rangle P(x)\big)$

Proof Summary (Appendix B.1.2). The axioms are all derived by duality using $\langle\cdot\rangle$, except DDG$\langle\cdot\rangle$ which derives from BDG$\langle\cdot\rangle$ by bounding solutions of the ghost ODE $y' = g(x,y)$ using the affine bound $2y \cdot g(x,y) \leq L(x)\,\|y\|_2^2 + M(x)$ on the Lie derivative of $\|y\|_2^2$. □

Axiom K$\langle\&\rangle$ is best understood in the contrapositive. Formula $[x' = f(x) \,\&\, Q \wedge \neg P]\neg G$ says G never happens along the solution while $\neg P$ holds. Thus, the solution cannot get to G unless it gets to P first. Axiom K$\langle\&\rangle$ formalizes the informal reasoning used for implication (4.5) above in the contrapositive, with $G \equiv u^2 + v^2 = \frac{1}{25}$ and $P \equiv \left(\frac{1}{4} \leq \|(u,v)\|_\infty \leq \frac{1}{2}\right)$. In the (partial) derivation shown below, the left premise requires a proof that the dashed blue circle

G cannot be reached while staying outside the green goal region P while the right premise requires a proof of the initial liveness property $\langle\alpha_l\rangle\big(u^2+v^2=\frac{1}{25}\big)$ for α_l. In a sequent calculus proof, refinement steps are naturally read from top-to-bottom (downwards), while deduction steps, i.e., axiom or rule applications, are read bottom-to-top (upwards).

Deduction ↑

$$\mathrm{K}\langle\&\rangle\ \dfrac{\dfrac{\vdots}{u^2+v^2=1 \vdash [\alpha_l \,\&\, \neg\big(\frac{1}{4} \le \|(u,v)\|_\infty \le \frac{1}{2}\big)]\neg\big(u^2+v^2=\frac{1}{25}\big)} \qquad \dfrac{\vdots}{u^2+v^2=1 \vdash \langle\alpha_l\rangle\big(u^2+v^2=\frac{1}{25}\big)}}{u^2+v^2=1 \vdash \langle\alpha_l\rangle\big(\frac{1}{4} \le \|(u,v)\|_\infty \le \frac{1}{2}\big)}$$

↓ **Refinement**

Refinement axiom DR$\langle\cdot\rangle$ is used in Chapter 3 but it is repeated here for clarity. In the axiom, formula $[x'=f(x) \,\&\, R]Q$ says that the ODE solution never leaves Q while staying in R, so if the solution gets to P within R, then it also gets to P within Q. The latter two refinement axioms BDG$\langle\cdot\rangle$, DDG$\langle\cdot\rangle$ are both derived from axiom BDG below, a new vectorial generalization of axiom DG that allows differential ghosts with provably bounded ODEs to be added.

Lemma 4.2 (Bounded differential ghosts)**.** *The following bounded differential ghosts axiom BDG is sound, where $y=(y_1,\dots,y_m)$ is a m-dimensional vector of fresh variables (not appearing in x), $g(x,y)$ is a corresponding m-dimensional vector of terms, and $\|y\|_2^2$ is the squared Euclidean norm of y. Term $e(x)$ and formulas $P(x), Q(x)$ are dependent only on free variables x (and not y).*

$$\text{BDG}\quad \begin{aligned}&[x'=f(x), y'=g(x,y) \,\&\, Q(x)]\ \|y\|_2^2 \le e(x)\\ &\to \big([x'=f(x) \,\&\, Q(x)]P(x) \leftrightarrow [x'=f(x), y'=g(x,y) \,\&\, Q(x)]P(x)\big)\end{aligned}$$

Proof Summary (Appendix B.1.1). The proof of BDG is similar to that for the differential ghosts axiom DG [142], but generalizes it to support vectorial, nonlinear ODEs by adding a syntactic precondition on boundedness of solutions of the added differential ghosts $y'=g(x,y)$. □

Like axiom DG, axiom BDG allows an arbitrary vector of ghost ODEs $y'=g(x,y)$ to be added syntactically to the ODEs. However, it places no syntactic restriction on the RHS of the ODE (such as linearity in axiom DG). For soundness, BDG instead adds a new precondition with a bound $\|y\|_2^2 \le e(x)$ in terms of x on the squared norm of y along solutions of the augmented ODE. This syntactic precondition ensures that y cannot blow up before x, so that solutions of $x'=f(x), y'=g(x,y)$ have existence intervals as long as those of the solutions of $x'=f(x)$. Section 4.3 shows how to prove these preconditions in order to use axiom BDG in ODE existence proofs through the refinement approach.

Returning to axioms BDG$\langle\cdot\rangle$, DDG$\langle\cdot\rangle$, the (nested) refinement in both axioms say that, if the ODE $x'=f(x)$ can reach $P(x)$, then the ODE $x'=f(x), y'=g(x,y)$, with the added variables y, can also reach $P(x)$. Axiom BDG$\langle\cdot\rangle$ is the derived diamond version of BDG, obtained by directly dualizing the inner equivalence of BDG with $\langle\cdot\rangle$ and propositional simplification. The intuition behind BDG$\langle\cdot\rangle$ is identical to BDG: if the added ghost ODEs y never blow up in norm, then they do not affect whether the solution of the original ODEs $x'=f(x)$ can reach $P(x)$.

Derived axiom DDG$\langle\cdot\rangle$ is a differential version of BDG$\langle\cdot\rangle$. Instead of bounding the squared norm $\|y\|_2^2$ explicitly, DDG$\langle\cdot\rangle$ instead limits the rate of growth of the ghost ODEs by bounding the Lie derivative $\mathcal{L}_{x'=f(x),y'=g(x,y)}(\|y\|_2^2) = 2y \cdot g(x,y)$ of the squared norm. This derivative bound in turn implicitly bounds the squared norm of the ghost ODEs by the solution of the linear differential equation $z' = L(x)z + M(x)$, with dependency on the value of x along solutions of the ODE $x' = f(x)$, which ensures that premature blow-up of y before x itself blows up is impossible. Any refinement step using axiom DDG$\langle\cdot\rangle$ can also use axiom BDG$\langle\cdot\rangle$ since the former is derived from the latter. The advantage of DDG$\langle\cdot\rangle$ is it builds in canonical differential reasoning steps once-and-for-all (see Lemma 4.1 and Section 4.3) which simplifies the proofs.

Axioms K$\langle\&\rangle$, DR$\langle\cdot\rangle$, BDG$\langle\cdot\rangle$, DDG$\langle\cdot\rangle$ all prove implication (4.3) in just one refinement step. Logical implication is transitive though, so a sequence of such steps can be chained together to prove implication (4.3). This is shown in (4.6), with neighboring implications informally chained together for illustration:

$$\begin{aligned}\langle x' = f(x) \,\&\, Q_0\rangle P_0 &\overset{\text{DR}\langle\cdot\rangle\text{ with }[x'=f(x)\,\&\,Q_0]Q_1}{\overbrace{\longrightarrow}} \langle x' = f(x) \,\&\, Q_1\rangle P_0 \overset{\text{K}\langle\&\rangle\text{ with }[x'=f(x)\,\&\,Q_1\wedge\neg P_1]\neg P_0}{\overbrace{\longrightarrow}} \langle x' = f(x) \,\&\, Q_1\rangle P_1\\ &\longrightarrow \cdots \\ &\longrightarrow \langle x' = f(x) \,\&\, Q\rangle P\end{aligned} \tag{4.6}$$

The box modality formulas annotated above each implication in (4.6) are side conditions to be proved for the chain of refinements (4.6) in order to prove the desired implication (4.3). However, an *unconditional* proof of the liveness property $\langle x' = f(x) \,\&\, Q\rangle P$ at the end of the chain still needs a proof of the hypothesis $\langle x' = f(x) \,\&\, Q_0\rangle P_0$ at the beginning of the chain. Typically, this hypothesis is a (simple) existence assumption for the differential equation. Formalizing and proving such existence properties is the focus of Section 4.3. Those proofs are also based on refinements and make use of axioms BDG$\langle\cdot\rangle$, DDG$\langle\cdot\rangle$.

Refinement with axiom DR$\langle\cdot\rangle$ requires proving the formula $[x' = f(x) \,\&\, R]Q$. Naïvely, one might expect that adding $\neg P$ to the domain constraint should also work, i.e., the solution only needs to be in Q while it has not yet gotten to P:

$$\text{DR}\langle\cdot\rangle\lightning\quad [x' = f(x) \,\&\, R \wedge \neg P]Q \rightarrow \big(\langle x' = f(x) \,\&\, R\rangle P \rightarrow \langle x' = f(x) \,\&\, Q\rangle P\big)$$

This conjectured axiom is unsound (indicated by $\lightning$) as the solution could sneak out of Q exactly when it crosses from $\neg P$ into P. In continuous settings, the language of topology makes precise what this means (recall Section 2.2.2, page 19). The following topological refinement axioms soundly restrict what happens at the crossover point:

Lemma 4.3 (Topological ODE refinement axioms)**.** *The following topological* $\langle\cdot\rangle$ *ODE refinement axioms are sound. In axiom COR, formulas* P, Q *either both characterize topologically open or both characterize topologically closed sets over variables* x*.*

$$\text{COR}\quad \neg P \wedge [x' = f(x) \,\&\, R \wedge \neg P]Q \rightarrow \big(\langle x' = f(x) \,\&\, R\rangle P \rightarrow \langle x' = f(x) \,\&\, Q\rangle P\big)$$

$$\text{SAR}\quad [x' = f(x) \,\&\, R \wedge \neg(P \wedge Q)]Q \rightarrow \big(\langle x' = f(x) \,\&\, R\rangle P \rightarrow \langle x' = f(x) \,\&\, Q\rangle P\big)$$

Proof in Appendix B.1.2.

Axiom COR is the more informative topological refinement axiom. Like the (unsound) axiom candidate DR$\langle\cdot\rangle\lightning$, it allows formula $\neg P$ to be assumed in the domain constraint when proving the box refinement. For soundness though, axiom COR has crucial topological side conditions on formulas P, Q so it can only be used when those conditions are met. Several variations of COR are possible (with similar soundness proofs), but they require alternative topological restrictions and additional topological notions. One useful variation involving the topological interior is given in Lemma 4.35. For the sake of generality, this chapter gives semantic topological side conditions with associated semantic soundness proofs in Appendix B.1.2. Axiom SAR applies more generally than COR but only assumes the less informative formula $\neg(P \land Q)$ in the domain constraint for the box modality. Its proof crucially relies on Q being a semianalytic formula so that the set it characterizes has tame topological behavior [14], see the proof in Appendix B.1.2 for more details. By topological considerations, axiom SAR is also sound if formula P (or resp. Q) characterizes a topologically closed (resp. open) set over the ODE variables x. These additional cases are also proved in Appendix B.1.2 without relying on the fact that Q is semianalytic.

4.3 Finite-Time Blow Up and Global Existence

This section explains how global existence properties can be proved for a given ODE $x' = f(x)$, subject to assumptions Γ about the initial states for the ODE. The existence and uniqueness theorems for ODEs [33, 204] guarantee that (sufficiently smooth) ODEs $x' = f(x)$ always have a unique, right-maximal solution from any initial state, $\varphi : [0, T) \to \mathbb{S}$ for some $0 < T \leq \infty$. However, these theorems give no guarantees about the precise duration T. In particular, ODEs can exhibit a technical phenomenon known as *finite-time blow up of solutions* [33], where φ is only defined on a bounded time interval $[0, T)$ with $T < \infty$. Intuitively, this happens when the solution *blows up* because its norm tends to ∞ as time t tends to T, as shown in Fig. 4.2. Additionally, it is possible that such finite-time blow up phenomena only happen for *some* initial conditions (and corresponding solutions) of the ODE. These initial conditions (with finite-time blow up) are typically irrelevant to the model of concern, especially when the dynamics of the corresponding real-world system is controlled to stay away from the blow up. For example, α_n (4.2) exhibits finite blow up of solutions only outside the red disk as shown in Fig. 4.1 and the blow up occurs well after its solutions have reached the target region, see Fig. 4.2. As an additional example, consider the following nonlinear ODE:

$$\alpha_b \equiv v' = -v^2 \tag{4.7}$$

The solution to this ODE is $v(t) = \frac{v_0}{v_0+t}$, where $v_0 \neq 0$ is the initial value of v at time $t = 0$ (if $v_0 = 0$, then $v(t) = 0$ for all t). If $v_0 < 0$ initially, then this solution is only defined to the right for the finite time interval $[0, -v_0)$, because the denominator $v_0 + t$ is 0 at $t = -v_0$. On the other hand, for $v_0 \geq 0$, the existence interval to the right is $[0, \infty)$. Thus, α_b exhibits finite-time blow up of solutions, but only for $v_0 < 0$.

4.3.1 Global Existence Proofs

The discussion above uses the mathematical solution $v(t)$ of the ODE α_b (4.7) as a function of time. For deductive proofs, the (global) existence of solutions can be expressed in dL as a special form of an ODE liveness property. The first step is to add a fresh variable t with $t' = 1$ that tracks the progress of time.[2] Then, using a fresh variable τ not in x, t, the following formula syntactically expresses that the ODE has a global solution because its solutions exceeds time τ, for any arbitrary τ:

$$\forall \tau \, \langle x' = f(x), t' = 1 \rangle \, t > \tau \tag{4.8}$$

Proving formula (4.8) shows global existence of solutions for the ODE $x' = f(x)$. The simplest instance of (4.8) is for ODE $t' = 1$ by itself without any ODE $x' = f(x)$. The formula (4.8) is valid because $t' = 1$ is an ODE with constant RHS, as shown below in axiom TEx.

Lemma 4.4 (Time existence)**.** *The following axiom is derivable in* dL.

TEx $\forall \tau \, \langle t' = 1 \rangle t > \tau$

Proof. Axiom TEx is derived directly from dL's solution axiom [142] but it also has an easy semantic soundness proof which is given here. Consider an initial state ω and the corresponding modified state ω_τ^d where the value of variable τ is replaced by an arbitrary $d \in \mathbb{R}$. The (right-maximal) solution of ODE $t' = 1$ from state ω_τ^d is given by the function $\varphi : [0, \infty) \to \mathbb{S}$, where $\varphi(\zeta)(t) = \omega_\tau^d(t) + \zeta = \omega(t) + \zeta$, and $\varphi(\zeta)(y) = \omega_\tau^d(y)$ for all other variables y. In particular, $\varphi(\zeta)(\tau) = d$. Thus, at any time $\zeta > d - \omega(t)$, $\varphi(\zeta)(t) = \omega(t) + \zeta > d = \varphi(\zeta)(\tau)$. This time ζ witnesses $\langle t' = 1 \rangle t > \tau$. □

Other instances of (4.8) can be proved by refining axiom TEx using axioms BDG$\langle\cdot\rangle$, DDG$\langle\cdot\rangle$ with appropriate assumptions about the initial conditions for the additional ODEs $x' = f(x)$. This is exemplified for the ODE α_b next.

Example 4.5 (Velocity of particle with air resistance)**.** The ODE α_b can be viewed as a model of the velocity of a particle that is slowing down due to air resistance. Of course, it does not make physical sense for the velocity of such a particle to "blow up". However, the solution of α_b only exists globally if the particle starts with positive initial velocity $v > 0$, otherwise, it only has short-lived solutions. The reason is that α_b only makes physical sense for positive velocities $v > 0$, so that the air resistance term $-v^2$ slows the particle down instead of speeding it up. Indeed, global existence (4.8) can be proved for α_b if its initial velocity is positive, i.e., the dL formula $v > 0 \to \forall \tau \, \langle v' = -v^2, t' = 1 \rangle t > \tau$ is valid.

$$\text{→R, ∀R}\,\dfrac{\text{DDG}\langle\cdot\rangle\,\dfrac{\text{M}[\cdot]\,\dfrac{\text{dbx}_{\succcurlyeq}\,\dfrac{*}{v > 0 \vdash [v' = -v^2, t' = 1]\, v > 0}}{v > 0 \vdash [v' = -v^2, t' = 1]\, 2v \cdot (-v^2) \leq 0} \qquad \text{TEx}\,\dfrac{*}{\vdash \langle t' = 1 \rangle t > \tau}}{v > 0 \vdash \langle v' = -v^2, t' = 1 \rangle t > \tau}}{\vdash v > 0 \to \forall \tau \, \langle v' = -v^2, t' = 1 \rangle t > \tau}$$

The derivation shown above starts with basic propositional steps (→R, ∀R), after which axiom DDG$\langle\cdot\rangle$ is used with $v' = -v^2$ as the differential ghost equation with the trivial choice of bounds $L = 0, M = 0$. This yields two premises, the right of which is proved by TEx. The resulting left premise requires proving the formula $2v \cdot (-v^2) \leq 0$ along the ODE. Mathematically,

[2] For consistency, the ODE $x' = f(x)$ is assumed to not mention t even if this is not always strictly necessary.

this says that the derivative of the squared norm v^2 is non-negative along α_b, so that v^2 is non-increasing and cannot blow up.[3] An M[·] step strengthens the postcondition to $v > 0$ since $v > 0$ implies $2v \cdot (-v^2) \leq 0$ in real arithmetic. The resulting premise is an invariance property for $v > 0$ which is proved using rule $\text{dbx}_{\succcurlyeq}$ from Corollary 3.6 (page 41) with cofactor $g = -v$. The initial assumption $v > 0$ is crucially used in this invariance proof step, as expected. △

The idea of refinements from Section 4.2 offers another view of the derivation in Example 4.5 as a single refinement step in the chain (4.6), recall that refinement steps are read from top-to-bottom. Here, an initial existence property for the ODE $t' = 1$ is refined to the desired existence property for the ODE $v' = -v^2, t' = 1$. The refinement step is justified using DDG⟨·⟩ with the box modality formula $[v' = -v^2, t' = 1]\, 2v \cdot (-v^2) \leq 0$.

$$\langle t' = 1\rangle t > \tau \xrightarrow{\text{DDG}\langle\cdot\rangle} \langle v' = -v^2, t' = 1\rangle t > \tau$$

This chain can be extended to prove global existence for more complicated ODEs $x' = f(x)$ in a stepwise fashion, and (possibly) alternating between uses of DDG⟨·⟩ or BDG⟨·⟩ for the refinement step. To do this, note that any ODE $x' = f(x)$ can be written in *dependency order*, where each group y_i is a vector of variables and each g_i corresponds to the respective vectorial RHS of the ODE for y_i for $i = 1, \dots, k$. The RHS of each y_i' is only allowed to depend on the preceding vectors of variables (inclusive) $y_1, \dots, y_i$.

$$\underbrace{y_1' = g_1(y_1), y_2' = g_2(y_1, y_2), y_3' = g_3(y_1, y_2, y_3), \dots, y_k' = g_k(y_1, y_2, y_3, \dots, y_k)}_{x'=f(x) \text{ written in dependency order}} \tag{4.9}$$

Corollary 4.6 (Dependency order existence). *Let the ODE $x' = f(x)$ be in dependency order* (4.9)*, and τ be a fresh variable not in x, t. The following rule with k stacked premises is derived from BDG⟨·⟩, DDG⟨·⟩ and TEx, where the postcondition of each premise P_i for $1 \leq i \leq k$ can be chosen to be either of the form:*

Ⓑ $P_i \equiv \|y_i\|_2^2 \leq e_i(t, y_1, \dots, y_{i-1})$ *for some term e_i with the indicated dependencies, or,*

Ⓓ $P_i \equiv 2y_i \cdot g_i(y_1, \dots, y_i) \leq L_i(t, y_1, \dots, y_{i-1})\, \|y_i\|_2^2 + M_i(t, y_1, \dots, y_{i-1})$ *for some terms L_i, M_i with the indicated dependencies.*

$$\text{DEx}\ \frac{\begin{array}{l}\Gamma \vdash [y_1' = g_1(y_1), t' = 1]P_1 \\ \Gamma \vdash [y_1' = g_1(y_1), y_2' = g_2(y_1, y_2), t' = 1]P_2 \\ \quad\vdots \\ \Gamma \vdash [y_1' = g_1(y_1), \dots, y_k' = g_k(y_1, \dots, y_k), t' = 1]P_k\end{array}}{\Gamma \vdash \forall \tau\, \langle x' = f(x), t' = 1\rangle t > \tau}$$

Proof Summary (Appendix B.2.1). The derivation proceeds (backward) by successive refinements using either BDG⟨·⟩ for premises corresponding to the form Ⓑ or DDG⟨·⟩ for those corresponding to Ⓓ, with the ghost equations for g_i and the respective bounds e_i or L_i, M_i at each step for $i = k, \dots, 1$. □

[3] The fact that v^2 is non-increasing can also be used in an alternative derivation with axiom BDG⟨·⟩ and the bound $e = v_0^2$, where v_0 syntactically stores the initial value of v.

Rule DEx corresponds to a refinement chain (4.6) of length k, with successive BDG$\langle\cdot\rangle$ or DDG$\langle\cdot\rangle$ refinement steps, e.g.:

$$\begin{aligned}\langle t'=1\rangle t>\tau \overset{\mathrm{BDG}\langle\cdot\rangle}{\longrightarrow} \langle y_1'=g_1(y_1), t'=1\rangle t>\tau \overset{\mathrm{DDG}\langle\cdot\rangle}{\longrightarrow} \cdots \\ \longrightarrow \langle y_1'=g_1(y_1),\dots,y_k'=g_k(y_1,\dots,y_k), t'=1\rangle t>\tau\end{aligned}$$

In rule DEx any choice of the shape of premises (Ⓑ and Ⓓ) is sound as these correspond to an underlying choice of axiom BDG$\langle\cdot\rangle$, DDG$\langle\cdot\rangle$ to apply at each refinement step, respectively. Another source of flexibility arises when choosing the dependency ordering (4.9) for the ODE $x'=f(x)$, as long as the requisite dependency requirements are met. For example, one can always choose the coarsest dependency order $y_1\equiv x, g_1\equiv f(x)$ to directly prove global existence in one step using appropriate choice of bounds L_1, M_1. The advantage of using finer dependency orders in DEx is it allows the user to choose the bounds L_i, M_i in a step-by-step manner for $i=1,\dots,k$. On the other hand, the flexibility of rule DEx can also be a drawback because it relies on manual effort from users to choose the partition and to prove the resulting premises. Section 4.3.2 explains useful recipes for using the flexibility behind rule DEx, e.g., Corollaries 4.8 and 4.12, while Section 4.6.2 shows how to automate those proofs.

The discussion thus far proves global existence for ODEs with an explicit time variable t. This is not a restriction for the liveness proofs in later sections of this chapter because such a fresh time variable can always be added using the rule dGt below, which is derived from DG. The rule also adds the assumption $t=0$ initially without loss of generality for ease of proof.

$$\text{dGt}\ \frac{\Gamma, t=0 \vdash \langle x'=f(x), t'=1\,\&\,Q\rangle P}{\Gamma \vdash \langle x'=f(x)\,\&\,Q\rangle P}$$

The derivation of rule dGt is shown below, using axiom $\langle\cdot\rangle$ to switch between the box and diamond modalities and axiom DG$_\forall$ to introduce a universally quantified time variable t which is then instantiated by $\forall$L to $t=0$.

$$\begin{array}{rr}
 & \Gamma, t=0 \vdash \langle x'=f(x), t'=1\,\&\,Q\rangle P \\ \cline{2-2}
{\scriptstyle \langle\cdot\rangle,\neg\mathrm{L}} & \Gamma, t=0, [x'=f(x), t'=1\,\&\,Q]\neg P \vdash \mathit{false} \\ \cline{2-2}
{\scriptstyle \forall\mathrm{L}} & \Gamma, \forall t\,[x'=f(x), t'=1\,\&\,Q]\neg P \vdash \mathit{false} \\ \cline{2-2}
{\scriptstyle \mathrm{DG}_\forall} & \Gamma, [x'=f(x)\,\&\,Q]\neg P \vdash \mathit{false} \\ \cline{2-2}
{\scriptstyle \langle\cdot\rangle,\neg\mathrm{R}} & \Gamma \vdash \langle x'=f(x)\,\&\,Q\rangle P
\end{array}$$

4.3.2 Derived Existence Axioms

For certain classes of ODEs and initial conditions, there are well-known mathematical techniques to prove global existence of solutions. These techniques have purely syntactic renderings in dL as special cases of BDG$\langle\cdot\rangle$, DDG$\langle\cdot\rangle$, and DEx. In particular, this section shows how axioms GEx, BEx (shown below), which were proved semantically in an earlier presentation [192], can be derived syntactically. The refinement approach also yields natural generalizations of these axioms.

Globally Lipschitz ODEs

A function $f : \mathbb{R}^m \to \mathbb{R}^n$ is *globally Lipschitz continuous* if there is a (positive) Lipschitz constant $C \in \mathbb{R}$ such that the inequality $\|f(x) - f(y)\| \leq C \|x - y\|$ holds for all $x, y \in \mathbb{R}^m$, where $\|\cdot\|$ are appropriate norms. Since norms are equivalent on finite dimensional vector spaces [204, §5.V], without loss of generality, the Euclidean norm is used for the following discussion. An ODE $x' = f(x)$ is *globally Lipschitz* if its RHS $f(x)$ is globally Lipschitz continuous. Solutions of globally Lipschitz ODEs always exist globally for all time [204, §10.VII]. Global Lipschitz continuity is satisfied, e.g., by α_l (4.1), and more generally by linear (or even affine) ODEs of the form $x' = Ax$, where A is a matrix of (constant) parameters [204] because of the following (mathematical) inequality with Lipschitz constant $\|A\|_F$, the Frobenius norm of A:

$$\|Ax - Ay\|_2 = \|A(x - y)\|_2 \leq \|A\|_F \|x - y\|_2$$

This calculation uses norms $\|\cdot\|_2, \|\cdot\|_F$, which are not terms in dL because they are not polynomials (nor extended terms of Section 3.2, e.g., $\|x\|_2$ is not differentiable at $x = 0$). Thus, a subtle technical challenge in proofs is to appropriately rephrase mathematical inequalities, typically involving norms, into ones that can be reasoned about soundly also in the presence of differentiation. In this respect, the Euclidean norm is useful, because expanding the valid arithmetic inequality $0 \leq (1 - \|x\|_2)^2$ and rearranging yields:

$$2\|x\|_2 \leq 1 + \|x\|_2^2 \tag{4.10}$$

Notice that, unlike the Euclidean norm $\|x\|_2$, the RHS of the square inequality (4.10) can be represented syntactically. Indeed, the squared Euclidean norm is already used in axiom BDG and its derived versions BDG$\langle\cdot\rangle$, DDG$\langle\cdot\rangle$. To support intuition, the proof sketches below continue to use mathematical inequalities involving Euclidean norms, while the proofs in the appendix rephrase them with (4.10) instead. The following corollary shows how global existence for globally Lipschitz ODEs is derived using a norm inequality as a special case of rule DEx.

Corollary 4.7 (Global existence)**.** *The following global existence axiom is derived from DDG$\langle\cdot\rangle$ in* dL*, where τ is a fresh variable not in x, t, and $x' = f(x)$ is globally Lipschitz.*

GEx $\forall \tau \, \langle x' = f(x), t' = 1 \rangle t > \tau$

Proof Summary (Appendix B.2.1). Let C be the Lipschitz constant for f. The proof uses DDG$\langle\cdot\rangle$ and two (mathematical) inequalities. The first inequality (4.11) bounds $\|f(x)\|_2$ linearly in $\|x\|_2$. The constant 0 is chosen here to simplify the resulting arithmetic.

$$\begin{aligned}\|f(x)\|_2 &= \|f(x) - f(0) + f(0)\|_2 \leq \|f(x) - f(0)\|_2 + \|f(0)\|_2\\ &\leq C\|x - 0\|_2 + \|f(0)\|_2 = C\|x\|_2 + \|f(0)\|_2\end{aligned} \tag{4.11}$$

The next inequality (4.12) below uses bound (4.11) on $\|f(x)\|_2$ to further bound $2x \cdot f(x)$ linearly in $\|x\|_2^2$ along the ODE with appropriate choices of L, M that only depend on the (positive) Lipschitz constant C and $\|f(0)\|_2$.

$$\begin{aligned}
2x \cdot f(x) &\leq 2\|x\|_2 \|f(x)\|_2 \overset{(4.11)}{\leq} 2\|x\|_2 \big(C\|x\|_2 + \|f(0)\|_2\big) \\
&= 2C\|x\|_2^2 + 2\|x\|_2\|f(0)\|_2 \overset{(4.10)}{\leq} 2C\|x\|_2^2 + (1+\|x\|_2^2)\|f(0)\|_2 \qquad (4.12)\\
&= \underbrace{\big(2C + \|f(0)\|_2\big)}_{L}\|x\|_2^2 + \underbrace{\|f(0)\|_2}_{M}
\end{aligned}$$
□

The derivation of axiom GEx uses DDG$\langle\cdot\rangle$, but global existence extends to more complicated ODEs with the aid of DEx as long as appropriate choices of L, M can be made. A useful example of such an extension is global existence for ODEs that have an *affine dependency order* (4.9), i.e., each $y_i' = g_i(y_1, \ldots, y_i)$ is affine in y_i with $y_i' = A_i(y_1, \ldots, y_{i-1})y_i + b_i(y_1, \ldots, y_{i-1})$ where A_i, b_i are respectively matrix and vector terms with appropriate dimensions and the indicated variable dependencies.

Corollary 4.8 (Affine dependency order global existence). *Axiom GEx is derivable from DDG$\langle\cdot\rangle$ in* dL *for ODEs $x' = f(x)$ with affine dependency order.*

Proof Summary (Appendix B.2.1). The proof is similar to Corollary 4.7 but uses DEx to prove global existence step-by-step for the dependency order. It uses the following (mathematical) inequality and corresponding choices of L_i, M_i (shown below) for $i = 1, \ldots, k$ at each step:

$$\begin{aligned}
2y_i \cdot (A_i y_i + b_i) &= 2(y_i \cdot (A_i y_i) + y_i \cdot b_i) \leq 2\|A_i\|_F \|y_i\|_2^2 + 2\|y_i\|_2\|b_i\|_2 \\
&\leq 2\|A_i\|_F\|y_i\|_2^2 + (1+\|y_i\|_2^2)\|b_i\|_2 \\
&= \underbrace{(2\|A_i\|_F + \|b_i\|_2)}_{L_i}\|y_i\|_2^2 + \underbrace{\|b_i\|_2}_{M_i} \qquad (4.13)
\end{aligned}$$

This inequality is very similar to the one used for Corollary 4.7, where $\|A_i\|_F$ corresponds to C, and $\|b_i\|_2$ corresponds to $\|f(0)\|_2$. The difference is that terms L_i, M_i are allowed to depend on the preceding variables $y_1, \ldots, y_{i-1}$. Importantly for soundness, both terms meet the appropriate variable dependency requirements of DDG$\langle\cdot\rangle$ because the terms A_i, b_i are not allowed to depend on y_i in the affine dependency order. □

With the extended refinement chain underlying DEx, Corollary 4.8 enables more general proofs of global existence for certain multi-affine ODEs that are not necessarily globally Lipschitz.

Example 4.9 (Multi-affine ODE). Consider the multi-affine ODE $u' = u, v' = uv$. The RHS of this ODE is given by the function $\begin{pmatrix} u \\ v \end{pmatrix} \mapsto \begin{pmatrix} u \\ uv \end{pmatrix}$ which is not globally Lipschitz.[4] Nevertheless, the ODE meets the dependency requirements of Corollary 4.8 and has provable global solutions.

[4] For the function to be globally Lipschitz, there must exist a constant $C \in \mathbb{R}$ such that for all $\begin{pmatrix} u_1 \\ v_1 \end{pmatrix}, \begin{pmatrix} u_2 \\ v_2 \end{pmatrix} \in \mathbb{R}^2$, the norm inequality $\left\|\begin{pmatrix} u_1 - u_2 \\ u_1v_1 - u_2v_2 \end{pmatrix}\right\|_2 \leq C\left\|\begin{pmatrix} u_1 - u_2 \\ v_1 - v_2 \end{pmatrix}\right\|_2$ is satisfied. No such C exists because the $u_1v_1 - u_2v_2$ component on the LHS grows quadratically while the corresponding component $v_1 - v_2$ on the RHS grows linearly (consider setting $u_i = v_i$ for $i = 1, 2$).

The following derivation illustrates the proof of Corollary 4.8. In the first step, rule DEx is used with dependency order $y_1 \equiv u, y_2 \equiv v$ and Lipschitz constants $L_1(t) = 2, L_2(u,t) = 2u, M_1(t) = 0, M_2(u,t) = 0$. The dependency requirements of the Lipschitz constants, notably for L_2, are satisfied by these choices and the resulting premises are proved by dW, $\mathbb{R}$ because the postconditions are valid real arithmetic formulas.

$$\mathrm{DEx}\,\dfrac{\mathrm{dW}\,\dfrac{\mathbb{R}\,\dfrac{*}{\vdash 2(u)(u) \leq (2)u^2}}{\vdash [u' = u, t' = 1]2(u)(u) \leq (2)u^2} \qquad \mathrm{dW}\,\dfrac{\mathbb{R}\,\dfrac{*}{\vdash 2(v)(uv) \leq (2u)v^2}}{\vdash [u' = u, v' = uv, t' = 1]2(v)(uv) \leq (2u)v^2}}{\vdash \forall\tau\, \langle u' = u, v' = uv, t' = 1\rangle t > \tau}$$

Observe that the premises of DEx remove the ODEs for u, v in a step-by-step fashion. This is the key for generalizing global existence for globally Lipschitz ODEs [204, §10.VII] to more general classes of ODEs. △

Bounded Existence

Returning to the example ODEs α_n (4.2) and α_b (4.7), note that axiom GEx applies to neither of those ODEs because they do not have affine dependency order. This should be unsurprising—as observed earlier in Fig. 4.1 and Example 4.5 respectively, neither α_n nor α_b have global solutions from all initial states. Although Example 4.5 shows how global existence for α_b can be proved from assumptions motivated by physics, it is also useful to have general axioms (similar to GEx) corresponding to well-known mathematical techniques for proving global existence of solutions for nonlinear ODEs under particular assumptions.

Suppose that the solution of ODE $x' = f(x)$ is trapped within a bounded set (whose compact closure is contained in the domain of the ODE), then, the ODE solution exists globally [71, Corollary 2.5][89, Theorem 3.3]. In control theory, this principle is used to show the global existence of solutions near stable equilibria [71, 89]. It also applies in case the model of interest has state variables that are *a priori* known to range within a bounded set [6, Section 6].

This discussion suggests that the following formula is valid for any ODE $x' = f(x)$, where $B(x)$ characterizes a bounded set over the variables x so the assumption $[x' = f(x)]B(x)$ says that the ODE solution is always trapped within the bounded set characterized by $B(x)$.

$$[x' = f(x)]B(x) \rightarrow \forall\tau\, \langle x' = f(x), t' = 1\rangle t > \tau \tag{4.14}$$

Formula (4.14) is (equivalently) rewritten succinctly in the following corollary by negating the box modality.

Corollary 4.10 (Bounded existence)**.** *The following bounded existence axiom derives from BDG$\langle\cdot\rangle$ in* dL*, where τ is a fresh variable not in x, t, and formula $B(x)$ characterizes a bounded set over variables x.*

BEx $\forall\tau\, \langle x' = f(x), t' = 1\rangle (t > \tau \vee \neg B(x))$

Proof Summary (Appendix B.2.1). The squared norm $\|x\|_2^2$ function is continuous in x so it is bounded above by a constant D on the compact closure of the set characterized by $B(x)$. The proof uses axiom BDG$\langle\cdot\rangle$ with $e(x) = D$ and rephrases formula (4.14) with axiom $\langle\cdot\rangle$. □

Example 4.11 (Trapped solutions)**.** Axiom BEx proves global existence for α_n (4.2) within the compact disk $u^2 + v^2 \le \frac{1}{4}$ by showing that solutions starting in the disk are trapped in it. After the first ∀R step, a K⟨&⟩ step adds a disjunction to the postcondition. On the resulting right premise, axiom BEx finishes the proof. The left premise abbreviated ① is an ODE invariance property $u^2 + v^2 \le \frac{1}{4} \vdash [\alpha_n, t' = 1 \,\&\, \neg(t > \tau)](u^2 + v^2 \le \frac{1}{4})$, whose elided invariance proof is easy using the techniques of Chapter 3.

$$\forall\text{R}\,\dfrac{\text{K}\langle\&\rangle\,\dfrac{\textcircled{1} \qquad \text{BEx}\,\dfrac{*}{\vdash \langle \alpha_n, t' = 1\rangle (t > \tau \lor \neg(u^2 + v^2 \le \frac{1}{4}))}}{u^2 + v^2 \le \frac{1}{4} \vdash \langle \alpha_n, t' = 1\rangle t > \tau}}{u^2 + v^2 \le \frac{1}{4} \vdash \forall \tau\, \langle \alpha_n, t' = 1\rangle t > \tau} \qquad \triangle$$

Axiom BEx removes the global Lipschitz (or affine dependency) requirement of GEx but weakens the postcondition to say that solutions must either exist for sufficient duration or blow up and leave the bounded set characterized by formula $B(x)$. Like axiom GEx, axiom BEx is derived by refinement using axiom BDG⟨·⟩. This commonality yields a more general version of BEx, which also incorporates ideas from GEx.

Corollary 4.12 (Dependency order bounded existence)**.** *Consider the ODE $x' = f(x)$ in dependency order* (4.9)*, and where τ is a fresh variable not in x, t. The following axiom is derived from BDG⟨·⟩, DDG⟨·⟩ in* dL*, where the indices $i = 1 \ldots, k$ are partitioned into two disjoint index sets L, N such that:*

- *For each $i \in L$, $y_i' = g_i(y_1, \ldots, y_i)$ is affine in y_i.*
- *For each $i \in N$, $B_i(y_i)$ characterizes a bounded set over the variables y_i.*

GBEx $\forall \tau\, \langle x' = f(x), t' = 1\rangle \big(t > \tau \lor \bigvee_{i \in N} \neg B_i(y_i)\big)$

Proof Summary (Appendix B.2.1). The derivation is similar to rule DEx, with an internal DDG⟨·⟩ step (similar to GEx) for $i \in L$ and an internal BDG⟨·⟩ step (similar to BEx) for $i \in N$. □

The index set L in Corollary 4.12 indicates those variables of $x' = f(x)$ whose solutions are guaranteed to exist globally (with respect to the other variables). On the other hand, the index set N indicates the variables that may cause finite-time blow up of solutions. The postcondition of axiom GBEx says that solutions either exist for sufficient duration or they blow up and leave one of the bounded sets indexed by N. An immediate modeling application of Corollary 4.12 is to identify which of the state variables in a model must be proved (or assumed) to take on bounded values [6, Section 6]. This idea underlies the automated existence proof support discussed in Section 4.6.2.

4.3.3 Completeness for Global Existence

The derivation of the existence axioms GEx, BEx, GBEx and rule DEx illustrate the use of liveness refinement for proving existence properties. Moreover, BDG⟨·⟩ is the sole ODE diamond refinement axiom underlying these derivations (recall DDG⟨·⟩ is derived from BDG⟨·⟩). This raises a natural question: are there ODEs whose solutions exist globally, but whose global existence *cannot* be proved syntactically using BDG$\langle\cdot\rangle$? The next completeness result gives a conditional completeness answer: *all* global existence properties can be proved using BDG$\langle\cdot\rangle$, if the corresponding ODE solutions are *syntactically representable* by proof in dL.

Proposition 4.13 (Global existence completeness). *If the ODE $x' = f(x)$ has a global solution representable in the (extended)* dL *term language, then the global existence formula* (4.8) *is derivable for $x' = f(x)$ from axiom BDG$\langle\cdot\rangle$.*

Proof Summary (Appendix B.2.1). Suppose that ODE $x' = f(x)$ has a global solution syntactically represented in dL as term $X(t)$ dependent only on the free variable t, the (symbolic) initial values x_0 of variables x, and the (constant) parameters for the ODE. The equality $x = X(t)$ is provable along the ODE $x' = f(x), t' = 1$ by the complete proof rule dRI from Theorem 3.11 (page 48) because solutions are equational invariants. The proof uses BDG$\langle\cdot\rangle$ with the bounding term $e = \|X(t)\|_2^2$, so that the required hypothesis of BDG$\langle\cdot\rangle$, i.e., $[x' = f(x), t' = 1]\, \|x\|_2^2 \leq \|X(t)\|_2^2$ proves trivially using the equality $x = X(t)$. □

Notably, the proof of Proposition 4.13 actually only uses a syntactically representable and provable upper bound $X(t)$ with $\|x\|_2^2 \leq \|X(t)\|_2^2$, rather than an equality. Such an upper bound, if syntactically representable in dL, also suffices for proving global existence. The following remarks illustrate the usage and limitations of Proposition 4.13 (even with an upper bound).

Remark 4.14 (Syntactically representable solutions). Consider the example ODE $u' = u, v' = uv$ proved to have global solutions in Example 4.9. Mathematically, its solution is given by the following functions (defined for all $t \in \mathbb{R}$), where u_0, v_0 are the initial values of u, v at time $t = 0$ and $\exp$ is the real exponential function.

$$u(t) = u_0 \exp(t), v(t) = \frac{v_0}{\exp(u_0)} \exp(u_0 \exp(t)) \tag{4.15}$$

Since the solution (4.15) is defined globally, Proposition 4.13 seemingly provides an alternative way to prove global existence for the ODE. The caveat is that Proposition 4.13 only applies when a bound on the solution is *syntactically representable* as a term $X(t)$ in the term language. In this case, (4.15) requires an extended term language with the real exponential function and arithmetic over those terms.

Remark 4.15 (Global solutions with no syntactic bound). A further caveat is that there is a fixed *polynomial* ODE $y' = p(y)$, constructed by Bournez and Pouly [25, Theorem 1.3], such that for *any* term language extension, there is an initial value for the ODE whose solution exists globally but cannot be bounded above by any syntactic tower of function compositions that can be written down in the extension.[5]

More precisely, consider a term language extension with a unary fixed function symbol that has smooth interpretation $h : \mathbb{R} \to \mathbb{R}$. Without loss of generality,[6] assume that h grows at least linearly, i.e., $t \leq h(t)$ for all $t \in \mathbb{R}$. Define the function $g : \mathbb{N} \to \mathbb{R}$ with $g(n) = h^{[n]}(n)$ which diagonalizes h on natural number inputs, where iterated function compositions are defined with

[5] Thanks to Jeremy Avigad (personal communication) for help with this result.

[6] Otherwise, replace h with $\hat{h}(t) = h(t)^2 + t^2 + 1$.

$h^{[0]}(n) = n$ and $h^{[i+1]}(n) = h(h^{[n]}(n))$ for all $n \in \mathbb{N}$ using the usual injection from $\mathbb{N}$ to $\mathbb{R}$. The function g is constructed such that $g(n) \geq h^{[i]}(n)$ for all $n \geq i, n, i \in \mathbb{N}$, which is proved by repeated use of the inequality $t \leq h(t)$ for all $t \in \mathbb{R}$:

$$h^{[i]}(n) \leq h^{[i+1]}(n) \leq h^{[i+2]}(n) \leq \cdots \leq h^{[n]}(n) = g(n) \tag{4.16}$$

Define the (linear) interpolation $H : \mathbb{R} \to \mathbb{R}$ of g as follows, where $\lceil \cdot \rceil$, $\lfloor \cdot \rfloor$ denote the ceiling and floor functions respectively:

$$H(t) = \begin{cases} g(0) & \text{if } t < 0 \\ g(\lfloor t \rfloor) + \Big((g(\lceil t \rceil) - g(\lfloor t \rfloor)) \cdot (t - \lfloor t \rfloor) \Big) & \text{otherwise} \end{cases} \tag{4.17}$$

The function $\hat{H}(t) = H(t) + 2$ is continuous because H is continuous so, by Bournez and Pouly [25, Theorem 1.3], the ODE $y' = p(y)$ has an initial value y_0 such that the resulting unique and global solution $y(t)$ with $y(0) = y_0$ satisfies the bound $\hat{H}(t) = H(t) + 2 \leq \|y(t)\| + 1$ for all $t \in \mathbb{R}$. For any $i \in \mathbb{N}$ and sufficiently large $n \geq i$, the ODE solution $y(t)$ satisfies the following lower bound on its norm at time n:

$$h^{[i]}(n) \underbrace{\leq}_{(4.16)} g(n) \underbrace{=}_{(4.17)} H(n) < H(n) + 1 \leq \|y(n)\|$$

Thus, a bound on the ODE solution $y(t)$ *cannot* be syntactically represented by the i-th fold composition $h^{[i]}(t)$ for *any* i because the bound is violated at the points $n \in \mathbb{N}$ shown above.

The complicated closed form solution (4.15) and the lack of explicit, representable bounds on solutions in term language extensions highlight the advantage of axioms BDG$\langle\cdot\rangle$, DDG$\langle\cdot\rangle$ and their use in the derived axioms of Corollaries 4.7–4.12 because they implicitly deduce global existence *without* needing an explicitly representable solution for the ODEs.

4.4 Liveness Without Domain Constraints

This section presents proof rules for liveness properties of ODEs $x' = f(x)$ without domain constraints, i.e., where Q is the formula true. Errors and omissions in the surveyed techniques are highlighted in blue.

Notational Conventions (Global Existence). The rest of this chapter develops ODE liveness proof rules that rely on the global existence proofs from Section 4.3. In all subsequent proof rules, the ODE $x' = f(x)$ is said to have *provable global solutions* if the global existence formula (4.8) for $x' = f(x)$ is provable. For example, if $x' = f(x)$ were globally Lipschitz (or, as a special case, linear), then its global existence can be proven using axiom GEx from Corollaries 4.7 and 4.8. For uniformity, all proof steps utilizing this assumption are marked with GEx, although proofs of global existence could use various other techniques described in Section 4.3. All proof rules can also be soundly presented with explicit sufficient duration assumptions like $\mathrm{dV}^{\Gamma}_{\succcurlyeq}$ below, but those are omitted for brevity.

4.4.1 Differential Variants

The fundamental technique for verifying liveness of discrete loops are loop variants, i.e., well-founded quantities that always increase (or always decrease) on each loop iteration. *Differential variants* [137] are their continuous analog, where the value of a given term e is shown to increase along ODE solutions by showing that its rate of change is bounded below by a positive constant $\varepsilon() > 0$ along those solutions. Recall the notational convention (Section 2.1.1, page 15) for variable dependencies, so term $\varepsilon()$ is not allowed to depend on any of the free variables $x_1, \ldots, x_n$ appearing in the ODE and must therefore remain constant along the ODE solution.

Corollary 4.16 (Atomic differential variants [137])**.** *The following proof rules (where $\succcurlyeq$ is either $\geq$ or $>$) are derivable in* dL*. Terms $\varepsilon(), e_0()$ are constant for ODE $x' = f(x), t' = 1$. In rule dV$_{\succcurlyeq}$, the ODE $x' = f(x)$ has provable global solutions.*

$$\mathrm{dV}^{\Gamma}_{\succcurlyeq}\ \frac{\neg(e \succcurlyeq 0) \vdash \dot{e} \geq \varepsilon()}{\Gamma, e = e_0(), t = 0, \langle x' = f(x), t' = 1\rangle\big(e_0() + \varepsilon()t > 0\big) \vdash \langle x' = f(x), t' = 1\rangle e \succcurlyeq 0}$$

$$\mathrm{dV}_{\succcurlyeq}\ \frac{\neg(e \succcurlyeq 0) \vdash \dot{e} \geq \varepsilon()}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle e \succcurlyeq 0}$$

Proof Summary (Appendix B.2.2). Rule dV$_{\succcurlyeq}$ is derived in Appendix B.2.2 as a corollary of rule dV$^{\Gamma}_{\succcurlyeq}$ because the ODE $x' = f(x)$ is assumed to have solutions which (provably) exist globally.

Rule dV$^{\Gamma}_{\succcurlyeq}$ is derived from axiom K$\langle\&\rangle$ with the choice of formula $G \equiv \big(e_0() + \varepsilon()t > 0\big)$:

$$\mathrm{K}\langle\&\rangle\ \frac{\Gamma, e = e_0(), t = 0 \vdash [x' = f(x), t' = 1 \,\&\, \neg(e \succcurlyeq 0)]\big(e_0() + \varepsilon()t \leq 0\big)}{\Gamma, e = e_0(), t = 0, \langle x' = f(x), t' = 1\rangle\big(e_0() + \varepsilon()t > 0\big) \vdash \langle x' = f(x), t' = 1\rangle e \succcurlyeq 0}$$

Monotonicity M[$'$] strengthens the postcondition to $e \geq e_0() + \varepsilon()t$ with the domain constraint $\neg(e \succcurlyeq 0)$. A subsequent use of dI$_{\succcurlyeq}$ completes the derivation:

$$\mathrm{M}[']\ \frac{\mathrm{dI}_{\succcurlyeq}\ \dfrac{\neg(e \succcurlyeq 0) \vdash \dot{e} \geq \varepsilon()}{\Gamma, e = e_0(), t = 0 \vdash [x' = f(x), t' = 1 \,\&\, \neg(e \succcurlyeq 0)]\big(e \geq e_0() + \varepsilon()t\big)}}{\Gamma, e = e_0(), t = 0 \vdash [x' = f(x), t' = 1 \,\&\, \neg(e \succcurlyeq 0)]\big(e_0() + \varepsilon()t \leq 0\big)} \qquad \square$$

In both rules dV$^{\Gamma}_{\succcurlyeq}$, dV$_{\succcurlyeq}$, the lower bound $\varepsilon() > 0$ on the Lie derivative $\dot{e}$ ensures that the value of e strictly increases along solutions to the ODE. Geometrically, as illustrated in Fig. 4.3a, the value of e is bounded below over time t by the line $e_0() + \varepsilon()t$ with offset $e_0()$ and positive slope $\varepsilon()$. Since $e_0() + \varepsilon()t$ is non-negative for sufficiently large values of t, the (lower bounded) value of e is also eventually non-negative.

Two key subtleties underlying rules dV$^{\Gamma}_{\succcurlyeq}$, dV$_{\succcurlyeq}$ are illustrated in Figs. 4.3c and 4.3d. The first subtlety, shown in Fig. 4.3c, is that ODE solutions must exist for sufficiently long for e or, more precisely its lower bound, to become non-negative. This is usually left as a soundness-critical side condition in liveness proof rules [137, 176], but any such side condition is antithetical to approaches for minimizing the soundness-critical core in implementations [142] because it requires checking the (semantic) condition that solutions exist for sufficient duration. The conclusion of rule dV$^{\Gamma}_{\succcurlyeq}$ formalizes this side condition as an assumption. In contrast, rule dV$_{\succcurlyeq}$ requires provable global existence for the ODEs (provable as in Section 4.3). The second subtlety,

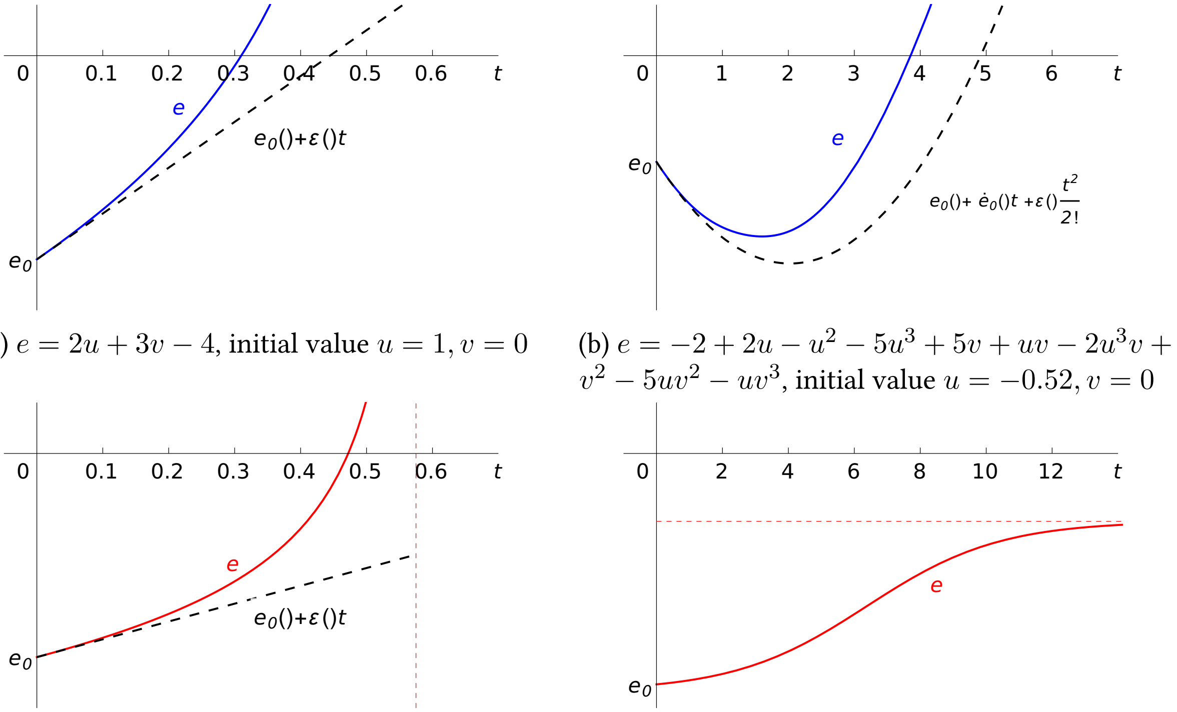

(a) $e = 2u + 3v - 4$, initial value $u = 1, v = 0$

(b) $e = -2 + 2u - u^2 - 5u^3 + 5v + uv - 2u^3v + v^2 - 5uv^2 - uv^3$, initial value $u = -0.52, v = 0$

(c) $e = u + v - 3$, initial value $u = 1, v = 0$

(d) $e{=}-10(u^2{+}v^2){-}1$, initial value $u{=}0.49, v{=}0$

Figure 4.3: The solid blue and red curves show the value of various terms e evaluated along solutions of the ODE α_n (4.2) from respective initial values e_0 over time t. The blue curves in Figs. 4.3a and 4.3b are respectively bounded below by a dashed black line (corresponding to Corollary 4.16) and a dashed quadratic curve (corresponding to Corollary 4.19) which imply that e is eventually non-negative along their respective ODE solutions. The red curve in Fig. 4.3c is also bounded below by the dashed black line, but its solution only exists for 0.575 time units (vertical red dashed asymptote) so the linear bound from Corollary 4.16 does not suffice for proving that e is eventually non-negative along the solution. The red curve in Fig. 4.3d has strictly positive derivative $\dot{e} > 0$ but the derivative tends to zero as t approaches ∞ so the value of e asymptotically increases towards a negative value (horizontal red dashed asymptote).

shown in Fig. 4.3d, is that rules $\mathrm{dV}^{\Gamma}_{\succcurlyeq}$, $\mathrm{dV}_{\succcurlyeq}$ crucially need a *constant* positive lower bound on the Lie derivative $\dot{e} \geq \varepsilon()$ for soundness [137] instead of merely requiring $\dot{e} > 0$. In the latter case, even though the value of e is strictly increasing along solutions, it is not guaranteed to become non-negative in finite time because the rate of increase can itself converge to zero. In fact, as Fig. 4.3d shows, e may stay negative by asymptotically increasing towards a negative value as t approaches ∞.

The following example shows how rule dV can be applied in combination with refinements.

Example 4.17 (Linear liveness)**.** The liveness property that Fig. 4.1 suggested for the linear ODE α_l (4.1) is proved by rule $\mathrm{dV}_{\succcurlyeq}$. The proof is shown below and visualized on the right. The first monotonicity step $\mathrm{M}\langle'\rangle$ strengthens the postcondition to the inner blue circle $u^2 + v^2 = \frac{1}{4}$ contained within the green goal region, see refinement (4.4). Next, since solutions satisfy

$u^2+v^2=1$ initially (black circle), the K$\langle\&\rangle$ refinement step requires a proof of the box modality formula $[\alpha_l \,\&\, u^2+v^2 \neq \frac{1}{4}]u^2+v^2 > \frac{1}{4}$ (omitted below). Intuitively, this formula expresses an intermediate value property: the *continuous* solution cannot reach $u^2+v^2 \leq \frac{1}{4}$ unless it crosses $u^2+v^2=\frac{1}{4}$. The postcondition is rearranged before dV$_{\succcurlyeq}$ is used with $\varepsilon()=\frac{1}{2}$. Its premise is proved by $\mathbb{R}$ because the Lie derivative of $\frac{1}{4}-(u^2+v^2)$ with respect to α_l is $2(u^2+v^2)$, which is bounded below by $\frac{1}{2}$ under the assumption $\frac{1}{4}-(u^2+v^2)<0$. This Lie derivative calculation also shows that the value of u^2+v^2 decreases along solutions of α_l with rate (at least) $\frac{1}{2}$ per unit time, which is visualized by the shrinking (dashed) circles with radii eventually smaller than $\frac{1}{4}$. Since the initial states satisfy $u^2+v^2=1$, a concrete upper bound on the time required for the solution to satisfy $u^2+v^2 \leq \frac{1}{4}$ is given by $(1-\frac{1}{4}) \,/\, \frac{1}{2} = \frac{3}{2}$ time units.

$$\begin{array}{ll}
 & \qquad\qquad * \\ \hline
\mathbb{R} & \frac{1}{4} < u^2+v^2 \vdash 2(u^2+v^2) \geq \frac{1}{2} \\ \hline
 & \frac{1}{4}-(u^2+v^2) < 0 \vdash 2(u^2+v^2) \geq \frac{1}{2} \\ \hline
\mathrm{dV}_{\succcurlyeq} & u^2+v^2=1 \vdash \langle\alpha_l\rangle \frac{1}{4}-(u^2+v^2) \geq 0 \\ \hline
 & u^2+v^2=1 \vdash \langle\alpha_l\rangle u^2+v^2 \leq \frac{1}{4} \\ \hline
\mathrm{K}\langle\&\rangle & u^2+v^2=1 \vdash \langle\alpha_l\rangle u^2+v^2 = \frac{1}{4} \\ \hline
\mathrm{M}\langle'\rangle & u^2+v^2=1 \vdash \langle\alpha_l\rangle \big(\frac{1}{4} \leq \|(u,v)\|_\infty \leq \frac{1}{2}\big)
\end{array}$$

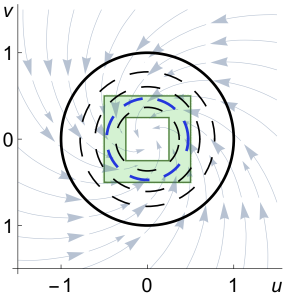

It is also instructive to examine the chain of refinements (4.6) underlying the proof above. Since α_l is a linear ODE, the first dV$_{\succcurlyeq}$ step refines the initial liveness property from GEx, i.e., that solutions exist globally (so for at least $\frac{3}{2}$ time units), to the property $u^2+v^2 \leq \frac{1}{4}$. Subsequent refinement steps can be read off from the steps above from top-to-bottom:

$$\langle\alpha_l, t'=1\rangle t > \frac{3}{2} \xrightarrow{\mathrm{dV}_{\succcurlyeq}} \langle\alpha_l\rangle u^2+v^2 \leq \frac{1}{4} \xrightarrow{\mathrm{K}\langle\&\rangle} \langle\alpha_l\rangle u^2+v^2 = \frac{1}{4} \xrightarrow{\mathrm{M}\langle'\rangle} \langle\alpha_l\rangle\Big(\frac{1}{4} \leq \|(u,v)\|_\infty \leq \frac{1}{2}\Big) \quad \triangle$$

The latter two steps in Example 4.17 illustrate the idea behind the next two surveyed proof rules. In their original presentation [191], the ODE $x'=f(x)$ is only assumed to be locally Lipschitz continuous, which is insufficient for global existence of solutions, making the original rules unsound, see Appendix B.3 for counterexamples. Compared to Corollary 4.16, Corollary 4.18 below uses the fact that the value of differential variant e evolves continuously along an ODE solution so it changes from $e \leq 0$ to $e > 0$ by crossing $e = 0$.

Corollary 4.18 (Equational differential variants [191])**.** *The following proof rules are derivable in* dL*. Term $\varepsilon()$ is constant for ODE $x'=f(x)$ and the ODE has provable global solutions.*

$$\mathrm{dV}_{=}\ \frac{e<0 \vdash \dot{e} \geq \varepsilon()}{\Gamma, \varepsilon()>0, e \leq 0 \vdash \langle x'=f(x)\rangle e=0} \qquad \mathrm{dV}_{=}^{M}\ \frac{e=0 \vdash P \quad e<0 \vdash \dot{e} \geq \varepsilon()}{\Gamma, \varepsilon()>0, e \leq 0 \vdash \langle x'=f(x)\rangle P}$$

Proof. Rule dV$_{=}^{M}$ is derived directly from dV$_{=}$ with a M$\langle'\rangle$ monotonicity step:

$$\mathrm{M}\langle'\rangle\ \cfrac{e=0 \vdash P \qquad \mathrm{dV}_{=}\ \cfrac{e<0 \vdash \dot{e} \geq \varepsilon()}{\Gamma, \varepsilon()>0, e \leq 0 \vdash \langle x'=f(x)\rangle e=0}}{\Gamma, \varepsilon()>0, e \leq 0 \vdash \langle x'=f(x)\rangle P}$$

Rule $\text{dV}_=$ derives using axiom $\text{K}\langle\&\rangle$ with $G \equiv e \geq 0$ and rule $\text{dV}_{\succcurlyeq}$ (with $\succcurlyeq$ being $\geq$) on the resulting right premise, which yields the sole premise of $\text{dV}_=$ (on the right, after $\text{dV}_{\succcurlyeq}$):

$$\text{K}\langle\&\rangle\,\dfrac{e \leq 0 \vdash [x' = f(x) \,\&\, e \neq 0]e < 0 \qquad \text{dV}_{\succcurlyeq}\dfrac{e < 0 \vdash \dot{e} \geq \varepsilon()}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle e \geq 0}}{\Gamma, \varepsilon() > 0, e \leq 0 \vdash \langle x' = f(x)\rangle e = 0}$$

From the left premise after using $\text{K}\langle\&\rangle$, axiom DX allows the domain constraint $e \neq 0$ to be assumed true initially, which strengthens the antecedent $e \leq 0$ to $e < 0$. Rule Barr proves the invariance of formula $e < 0$ for the ODE $x' = f(x) \,\&\, e \neq 0$ because the antecedents $e \neq 0, e = 0$ in its resulting premise are contradictory.

$$\text{DX}\,\dfrac{\text{Barr}\,\dfrac{\mathbb{R}\,\dfrac{*}{e \neq 0, e = 0 \vdash \dot{e} < 0}}{e < 0 \vdash [x' = f(x) \,\&\, e \neq 0]e < 0}}{e \leq 0 \vdash [x' = f(x) \,\&\, e \neq 0]e < 0}$$

□

Rule $\text{dV}_=$ extends the refinement chain of $\text{dV}_{\succcurlyeq}$ with an additional $\text{K}\langle\&\rangle$ step:

$$\langle x' = f(x), t' = 1\rangle t > e() \xrightarrow{\text{dV}_{\succcurlyeq}} \langle x' = f(x)\rangle e \geq 0 \xrightarrow{\text{K}\langle\&\rangle} \langle x' = f(x)\rangle e = 0$$

The refinement behind this additional step is an intermediate value property: if $e \leq 0$ is true initially then the (continuous) solution can never reach states satisfying $e \geq 0$ without first reaching one that satisfies $e = 0$. The view of $\text{dV}_{\succcurlyeq}$ as a refinement of GEx in Example 4.17 also yields generalizations of $\text{dV}_{\succcurlyeq}$ to higher Lie derivatives. Indeed, it suffices that *any* higher Lie derivative $\dot{e}^{(k)}$ is bounded below by a positive constant $\varepsilon()$ rather than just the first. Geometrically, this guarantees that e is bounded below by a degree k polynomial in time variable t that is non-negative for large enough t, see Fig. 4.3b for an illustration with $k = 2$.

Corollary 4.19 (Atomic higher differential variants)**.** *The following proof rule (where $\succcurlyeq$ is either $\geq$ or $>$) is derivable in* dL*. Term $\varepsilon()$ is constant for ODE $x' = f(x)$, $k \geq 1$ is a freely chosen natural number, and the ODE has provable global solutions.*

$$\text{dV}^k_{\succcurlyeq}\;\dfrac{\neg(e \succcurlyeq 0) \vdash \dot{e}^{(k)} \geq \varepsilon()}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle e \succcurlyeq 0}$$

Proof Summary (Appendix B.2.2). Since $\dot{e}^{(k)}$ is strictly positive outside the goal ($e \succcurlyeq 0$), all lower Lie derivatives $\dot{e}^{(i)}$ of e for $i < k$, including $e = \dot{e}^{(0)}$, eventually become positive. The derivation uses a sequence of dC, $\text{dI}_{\succcurlyeq}$ steps to prove a (polynomial) lower bound in t. □

4.4.2 Staging Sets

The *staging sets* [176] proof rule adds flexibility to rules such as $\text{dV}^M_=$ above by allowing users to choose a staging set formula S that *the ODE can only leave by entering the goal region P*. Staging sets are leaky invariants in the sense that they are almost invariant, except that they can be left by reaching the goal P. This staging property is expressed in the contrapositive by the box modality formula $[x' = f(x) \,\&\, \neg P]S$.

Corollary 4.20 (Staging sets [176])**.** *The following proof rule is derivable in* dL*. Term $\varepsilon()$ is constant for ODE $x' = f(x)$ and the ODE has provable global solutions.*

$$\text{SP}\ \dfrac{\Gamma \vdash [x' = f(x)\,\&\,\neg P]S \qquad S \vdash e \le 0 \land \overset{\bullet}{e} \ge \varepsilon()}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle P}$$

Proof Summary (Appendix B.2.2). The derivation starts by using refinement axiom K$\langle\&\rangle$ with $G \equiv \neg S$. The rest of the derivation is similar to $\mathrm{dV}_{\succcurlyeq}^{\Gamma}$, $\mathrm{dV}_{\succcurlyeq}$. □

The added choice of staging set formula S allows users to choose a staging set that, e.g., enables a liveness proof that uses a simpler differential variant e because e only needs to satisfy derivative bound $\overset{\bullet}{e} \ge \varepsilon()$ in the staging set S as opposed to everywhere outside the goal, as in $\mathrm{dV}_{\succcurlyeq}^{\Gamma}$, $\mathrm{dV}_{\succcurlyeq}$. Furthermore, proof rules can be significantly simplified by choosing S with desirable topological properties. For example, all of the liveness proof rules derived so far either have an explicit sufficient duration assumption (like $\mathrm{dV}_{\succcurlyeq}^{\Gamma}$) or assume that the ODEs have provable global solutions (like $\mathrm{dV}_{\succcurlyeq}$ using axiom GEx). An alternative is to use axiom BEx, by choosing the staging set formula $S(x)$ to characterize a bounded or compact set over the variables x as in the following corollary. The advantage of such a choice is the resulting staging set proof rules show (implicitly) that solutions must exist for long enough to reach the goal.

Corollary 4.21 (Bounded/compact staging sets)**.** *The following proof rules are derivable in* dL*. Term $\varepsilon()$ is constant for $x' = f(x)$. In rule SP$_b$, formula S characterizes a bounded set over variables x. In rule SP$_c$, it characterizes a compact, i.e., closed and bounded, set over those variables.*

$$\text{SP}_b\ \dfrac{\Gamma \vdash [x' = f(x)\,\&\,\neg P]S \qquad S \vdash \overset{\bullet}{e} \ge \varepsilon()}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle P} \qquad\qquad \text{SP}_c\ \dfrac{\Gamma \vdash [x' = f(x)\,\&\,\neg P]S \qquad S \vdash \overset{\bullet}{e} > 0}{\Gamma \vdash \langle x' = f(x)\rangle P}$$

Proof Summary (Appendix B.2.2). Rule SP$_b$ is derived using axiom BEx with differential variant e to establish a time bound. Rule SP$_c$ is an arithmetical corollary of SP$_b$, using the fact that continuous functions on compact domains attain their extrema. □

Example 4.22 (Nonlinear liveness)**.** The liveness property that Fig. 4.1 suggested for the nonlinear ODE α_n (4.2) is proved using rule SP$_c$ by choosing the staging set formula $S \equiv 1 \le u^2+v^2 \le 2$ and the differential variant $e = u^2 + v^2$. The proof is shown on the left and visualized on the right below; the goal $u^2 + v^2 \ge 2$ is shown in green while S is shown as a blue annulus.

$$\text{SP}_c\ \dfrac{\text{cut},\mathbb{R}\ \dfrac{\dfrac{*}{S \vdash [\alpha_n\,\&\,\neg(u^2+v^2 \ge 2)]S}}{u^2+v^2=1 \vdash [\alpha_n\,\&\,\neg(u^2+v^2 \ge 2)]S} \qquad \mathbb{R}\ \dfrac{*}{S \vdash \overset{\bullet}{e} > 0}}{u^2+v^2=1 \vdash \langle\alpha_n\rangle u^2+v^2 \ge 2}$$

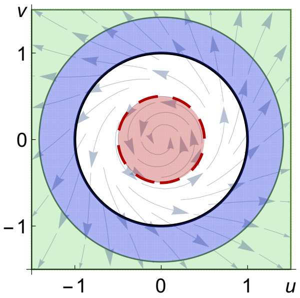

The Lie derivative $\overset{\bullet}{e}$ with respect to α_n is $2(u^2+v^2)(u^2+v^2-\frac{1}{4})$, which is bounded below by $\frac{3}{2}$ in S. Thus, the right premise of SP$_c$ closes trivially. The left premise requires proving that S is an invariant within the domain constraint $\neg(u^2+v^2 \ge 2)$. Intuitively, this is true because the ODE can only leave the blue annulus by entering the goal. The elided invariance proof for S is easy using the techniques of Chapter 3.

This proof exploits the flexibility provided by staging sets in two ways. First, the formula S is chosen to characterize a compact set (as required by rule SP_c). As explained in Section 4.3, solutions of α_n can blow up in finite time which necessitates the use of BEx for proving its liveness properties. Second, S cleverly *excludes* the red disk (dashed boundary) characterized by $u^2 + v^2 \leq \frac{1}{4}$. Solutions of α_n behave differently in this region, e.g., the Lie derivative $\dot{e}$ is *non-positive* in this disk. The chain of refinements (4.6) behind this proof can be seen from the derivation of rules SP_b, SP_c in Appendix B.2.2. The chain starts from the initial liveness property BEx with concrete[7] time bound $\frac{2}{3}$. The first $\text{K}\langle\&\rangle$ step shows that the staging set is ultimately exited ($\langle\alpha_n\rangle\neg S$), while the latter shows the desired liveness property:

$$\langle \alpha_n, t' = 1\rangle\left(t > \frac{2}{3} \lor \neg S\right) \xrightarrow{\text{K}\langle\&\rangle} \langle\alpha_n\rangle\neg S \xrightarrow{\text{K}\langle\&\rangle} \langle\alpha_n\rangle u^2 + v^2 \geq 2 \qquad \triangle$$

The need to use axiom BEx (or otherwise, assume global existence) is subtle and is often overlooked in the surveyed liveness arguments. An example of this is an incorrect claim [157, Remark 3.6] that an associated liveness argument [157, Theorem 3.5] works without assuming that the relevant sets are bounded. This chapter's axiomatic approach can be used to find and fix errors involving these subtleties from dL's sound reasoning foundations. As another example, the following *set Lyapunov function* proof rule adapts ideas from the literature [159, Theorem 2.4, Corollary 2.5] for proving liveness when the postcondition P characterizes an open set. The latter assumption on P enables a convenient choice of staging set in rule SP_c because $\neg P$ characterizes a closed set.

Corollary 4.23 (Set Lyapunov functions [159]). *The following proof rule is derivable in* dL. *Formula K characterizes a compact set over variables x, while formula P characterizes an open set over those variables.*

$$\text{SLyap}\quad \frac{e \geq 0 \vdash K \qquad \neg P, K \vdash \dot{e} > 0}{\Gamma, e \succcurlyeq 0 \vdash \langle x' = f(x)\rangle P}$$

Proof. Rule SLyap is derived from rule SP_c with $S \equiv \neg P \land K$, since the intersection of a closed set (characterized by $\neg P$) with a compact set (characterized by K) is compact. The resulting right premise from using SP_c is the right premise of SLyap:

$$\text{SP}_c\,\frac{\Gamma, e \succcurlyeq 0 \vdash [x' = f(x) \,\&\, \neg P](\neg P \land K) \qquad \neg P, K \vdash \dot{e} > 0}{\Gamma, e \succcurlyeq 0 \vdash \langle x' = f(x)\rangle P}$$

Continuing from the left premise, a monotonicity step with the premise $e \geq 0 \vdash K$ turns the postcondition to $e \succcurlyeq 0$. Rule Barr is used, which, along with the premise $e \geq 0 \vdash K$ results in the premises of rule SLyap:

$$\text{M}[']\,\dfrac{\mathbb{R}\,\dfrac{e \geq 0 \vdash K}{\neg P, e \succcurlyeq 0 \vdash \neg P \land K} \qquad \text{Barr}\,\dfrac{\text{cut}\,\dfrac{\neg P, K \vdash \dot{e} > 0 \qquad \mathbb{R}\,\dfrac{e \geq 0 \vdash K}{\neg P, e = 0 \vdash K}}{\neg P, e = 0 \vdash \dot{e} > 0}}{e \succcurlyeq 0 \vdash [x' = f(x) \,\&\, \neg P]e \succcurlyeq 0}}{\Gamma, e \succcurlyeq 0 \vdash [x' = f(x) \,\&\, \neg P](\neg P \land K)} \qquad \square$$

[7] The value of $u^2 + v^2$ grows at rate $\frac{3}{2}$ per time unit along solutions and the initial states satisfy $u^2 + v^2 = 1$. Thus, a lower bound on time required to leave the staging set (when $u^2 + v^2 > 2$) is $(2 - 1)\ /\ \frac{3}{2} = \frac{2}{3}$ time units.

Rule SLyap was claimed [159, Theorem 2.4, Corollary 2.5] to hold for any closed set K, when, in fact, K crucially needs to be compact as assumed implicitly in the associated proofs [159].

4.5 Liveness With Domain Constraints

This section presents proof rules for liveness properties ODEs $x' = f(x) \,\&\, Q$ with non-trivial domain constraints Q. These properties are significantly more subtle than liveness without domain constraints, because the limitation to a domain constraint Q may make it impossible for an ODE solution to reach a desired goal region before leaving Q.

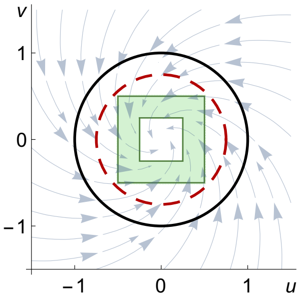

Consider the following liveness property for α_l (4.1) (shown on the right), which adds domain constraint $Q \equiv u^2 + v^2 \neq \frac{9}{16}$ restricting solutions from crossing the red dashed circle before reaching the green goal region.

$$\langle \alpha_l \,\&\, u^2 + v^2 \neq \frac{9}{16} \rangle \big(\frac{1}{4} \leq \|(u,v)\|_\infty \leq \frac{1}{2}\big) \tag{4.18}$$

As proved in Example 4.17, solutions starting from the black circle $u^2+v^2 = 1$ reach the green goal region. However, the continuous solutions must cross the red dashed circle $u^2 + v^2 = \frac{9}{16}$ to reach the goal, see discussion of implication (4.5). This violates the domain constraint and falsifies (4.18) for initial states on the black circle.

Axiom DR$\langle\cdot\rangle$ with $R \equiv \mathit{true}$ provides one way of soundly and directly generalizing the proof rules from Section 4.4, as shown in the following example.

Example 4.24 (Nonlinear liveness with domain)**.** The ODE liveness property $u^2 + v^2 = 1 \rightarrow \langle \alpha_n \rangle u^2 + v^2 \geq 2$ was proved in Example 4.22 for the nonlinear ODE α_n (4.2). The following derivation proves a stronger liveness property with the added domain constraint $1 \leq u^2 + v^2$ by extending the proof from Example 4.22 with a DR$\langle\cdot\rangle$ refinement step. The resulting left premise is an invariance property of the ODE whose proof is elided (see Chapter 3); intuitively, solutions starting from $u^2 + v^2 = 1$ grow outwards, and so they remain in the domain $1 \leq u^2 + v^2$ (see Fig. 4.1). The resulting right premise is proved in Example 4.22.

$$\text{DR}\langle\cdot\rangle\frac{\dfrac{*}{u^2+v^2=1 \vdash [\alpha_n] 1 \leq u^2+v^2} \qquad \dfrac{*}{u^2+v^2=1 \vdash \langle \alpha_n \rangle u^2+v^2 \geq 2}}{u^2+v^2=1 \vdash \langle \alpha_n \,\&\, 1 \leq u^2+v^2 \rangle u^2+v^2 \geq 2} \qquad \triangle$$

More generally, proof rules from Section 4.4 can be used from the right premise after refinement with DR$\langle\cdot\rangle$:

$$\text{DR}\langle\cdot\rangle\frac{\Gamma \vdash [x' = f(x)]Q \qquad \Gamma \vdash \langle x' = f(x) \rangle P}{\Gamma \vdash \langle x' = f(x) \,\&\, Q \rangle P}$$

This derivation extends all chains of refinements (4.6) from Section 4.4 with a DR$\langle\cdot\rangle$ step:

$$\cdots \longrightarrow \langle x' = f(x) \rangle P \xrightarrow{\text{DR}\langle\cdot\rangle} \langle x' = f(x) \,\&\, Q \rangle P$$

However, liveness arguments become much more intricate when attempting to generalize beyond domain constraint refinement with DR$\langle\cdot\rangle$, e.g., recall the unsound conjecture DR$\langle\cdot\rangle$↯. Indeed, unlike the technical glitches of Section 4.4, this chapter uncovers several subtle soundness-critical errors in the literature. With dL's deductive approach, these intricacies are isolated to the topological axioms (Lemma 4.3) which have been proved sound once-and-for-all. Errors and omissions in the surveyed techniques are again highlighted in blue.

The following proof rule generalizes differential variants dV$_\succcurlyeq$ to handle domain constraints. Like rule dV$_\succcurlyeq$, the differential variant e is guaranteed to eventually become non-negative along solutions with constant positive lower bound $\dot{e} \geq \varepsilon()$ on its Lie derivative. The additional twist is that the domain constraint Q must be proved to hold as long as e is still negative, i.e., while the goal has not been reached. This is expressed in the contrapositive by the formula $[x' = f(x) \,\&\, \neg(e \succcurlyeq 0)]Q$ in the left premise of the rule.

Corollary 4.25 (Atomic differential variants with domains [137]). *The following proof rule (where $\succcurlyeq$ is either $\geq$ or $>$) is derivable in* dL. *Term $\varepsilon()$ is constant for the ODE $x' = f(x)$ and the ODE has provable global solutions. Formula Q characterizes a closed (resp. open) set when $\succcurlyeq$ is $\geq$ (resp. $>$).*

$$\text{dV}_\succcurlyeq\& \quad \frac{\Gamma \vdash [x' = f(x) \,\&\, \neg(e \succcurlyeq 0)]Q \qquad \neg(e \succcurlyeq 0), Q \vdash \dot{e} \geq \varepsilon()}{\Gamma, \varepsilon() > 0, \neg(e \succcurlyeq 0) \vdash \langle x' = f(x) \,\&\, Q\rangle e \succcurlyeq 0}$$

Proof Summary (Appendix B.2.3). The derivation uses axiom COR choosing $R \equiv \mathit{true}$, noting that $e \geq 0$ (resp. $e > 0$) characterizes a topologically closed (resp. open) set so the appropriate topological requirements of COR are satisfied. The highlighted $\neg(e \succcurlyeq 0)$ assumption is crucial for soundly using axiom COR:

$$\text{COR}\ \frac{\Gamma \vdash [x' = f(x) \,\&\, \neg(e \succcurlyeq 0)]Q \qquad \dfrac{\dfrac{\neg(e \succcurlyeq 0), Q \vdash \dot{e} \geq \varepsilon()}{\cdots}}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle e \succcurlyeq 0}}{\Gamma, \varepsilon() > 0, \neg(e \succcurlyeq 0) \vdash \langle x' = f(x) \,\&\, Q\rangle e \succcurlyeq 0}$$

The derivation steps on the right premise are similar to the ones used in dV$_\succcurlyeq$ although an intervening dC step is used to additionally assume Q in the antecedents. □

Rule dV$_\succcurlyeq$& uses the topological refinement axiom COR to extend the refinement chain for dV$_\succcurlyeq$ as follows:

$$\cdots \xrightarrow{\text{dV}_\succcurlyeq} \langle x' = f(x)\rangle e \succcurlyeq 0 \xrightarrow{\text{COR}} \langle x' = f(x) \,\&\, Q\rangle e \succcurlyeq 0 \tag{4.19}$$

A subtle advantage of placing the refinement COR at the end of the refinement chain (4.19) is that it decouples reasoning about domain constraint Q from earlier refinement steps. Notably, earlier refinement steps like dV$_\succcurlyeq$ in the chain above can focus on handling other subtleties, such as sufficient duration existence of solutions (Section 4.3), *without* worrying about domain constraints. The original presentation of rule dV$_\succcurlyeq$& [137] omits the highlighted $\neg(e \succcurlyeq 0)$ assumption, but the rule is unsound without it. In addition, the original presentation uses a form of syntactic weak negation [137], which is unsound for open postconditions, as pointed out earlier [176], see Appendix B.3 for counterexamples.

The proofs of the next two corollaries also make use of axiom COR to derive the proof rule $\mathrm{dV}_=^M\&$ [191] and the adapted rule SLyap& [159]. These rules respectively generalize $\mathrm{dV}_=^M$ and SLyap from Section 4.4 to handle domain constraints. The soundness issues in their original presentations [159, 191], which were identified in Section 4.4, remain highlighted here. Like rule $\mathrm{dV}_\succcurlyeq\&$, rules $\mathrm{dV}_=\&$, $\mathrm{dV}_=^M\&$ below have an additional premise requiring that the domain constraint Q provably holds while the goal has not yet been reached $[x' = f(x)\,\&\, e < 0]Q$.

Corollary 4.26 (Equational differential variants with domains [191]). *The following proof rules are derivable in* dL*. Term $\varepsilon()$ is constant for the ODE $x' = f(x)$ and the ODE has provable global solutions for both rules. Formula Q characterizes a closed set over variables x.*

$$\mathrm{dV}_=\&\ \frac{\Gamma \vdash [x' = f(x)\,\&\, e < 0]Q \qquad e < 0, Q \vdash \dot{e} \geq \varepsilon()}{\Gamma, \varepsilon() > 0, e \leq 0, Q \vdash \langle x' = f(x)\,\&\, Q\rangle e = 0}$$

$$\mathrm{dV}_=^M\&\ \frac{Q, e = 0 \vdash P \qquad \Gamma \vdash [x' = f(x)\,\&\, e < 0]Q \qquad e < 0, Q \vdash \dot{e} \geq \varepsilon()}{\Gamma, \varepsilon() > 0, e \leq 0, Q \vdash \langle x' = f(x)\,\&\, Q\rangle P}$$

Proof Summary (Appendix B.2.3). Rules $\mathrm{dV}_=\&$, $\mathrm{dV}_=^M\&$ are both derived from rule $\mathrm{dV}_\succcurlyeq\&$ with $\geq$ for $\succcurlyeq$, since Q characterizes a closed set. Their derivations are respectively similar to the derivation of $\mathrm{dV}_=$, $\mathrm{dV}_=^M$ from $\mathrm{dV}_\succcurlyeq$ and require the provable global solutions assumption for soundly applying rule $\mathrm{dV}_\succcurlyeq\&$. □

Rule SLyap& below has identical premises to the corresponding SLyap rule (without domain constraints). The additional insight is that, assuming $e > 0$ is true initially, those same premises can be used to conclude the stronger liveness property $\langle x' = f(x)\,\&\, e > 0\rangle P$ because e can be additionally proved to stay positive along the solutions using the premises. This stronger conclusion can be used with a monotonicity step to prove more general liveness properties with an arbitrary domain constraint Q as exemplified by rule $\mathrm{SLyap}^M\&$.

Corollary 4.27 (Set Lyapunov functions with domains [159]). *The following proof rules are derivable in* dL*. Formula K characterizes a compact set over variables x, while formula P characterizes an open set over those variables.*

$$\mathrm{SLyap}\&\ \frac{e \geq 0 \vdash K \qquad \neg P, K \vdash \dot{e} > 0}{\Gamma, e > 0 \vdash \langle x' = f(x)\,\&\, e > 0\rangle P}$$

$$\mathrm{SLyap}^M\&\ \frac{e \geq 0 \vdash K \qquad \neg P, K \vdash \dot{e} > 0 \qquad e > 0 \vdash Q}{\Gamma, e > 0 \vdash \langle x' = f(x)\,\&\, Q\rangle P}$$

Proof Summary (Appendix B.2.3). Rule $\mathrm{SLyap}^M\&$ is derived from rule SLyap& by monotonicity on the domain constraints with the additional premise $e > 0 \vdash Q$. Rule SLyap& is derived from SLyap after a refinement step with COR since both formulas $e > 0$ and P characterize open sets as sketched below.

$$\text{COR}\ \dfrac{\dfrac{\dfrac{e \geq 0 \vdash K \qquad \neg P, K \vdash \dot{e} > 0}{\dots}}{\Gamma, e > 0 \vdash [x' = f(x)\,\&\, \neg P]e > 0} \qquad \text{SLyap}\ \dfrac{e \geq 0 \vdash K \qquad \neg P, K \vdash \dot{e} > 0}{\Gamma, e > 0 \vdash \langle x' = f(x)\rangle P}}{\Gamma, e > 0 \vdash \langle x' = f(x)\,\&\, e > 0\rangle P}$$

The left premise proves the invariance of $e > 0$ for ODE $x' = f(x)$ with domain constraint P. The elided derivation (see proof) reduces to two premises which are identical to those of rule SLyap. The right premise uses rule SLyap, which necessitates the compactness assumption for formula K for soundness. □

The following staging sets with domain constraints proof rule SP& [176] generalizes rule SP using axiom SAR. Notably, unlike the preceding rules, rule SP& requires no topological assumptions[8] about the domain constraint Q nor of the goal region P so it can be used in proofs of more general liveness properties.

Corollary 4.28 (Staging sets with domains [176])**.** *The following proof rule is derivable in* dL*. Term $\varepsilon()$ is constant for ODE $x' = f(x)$ and the ODE has provable global solutions.*

$$\text{SP\&}\ \frac{\Gamma \vdash [x' = f(x) \,\&\, \neg(P \land Q)]S \qquad S \vdash Q \land e \le 0 \land \dot{e} \ge \varepsilon()}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x) \,\&\, Q\rangle P}$$

Proof Summary (Appendix B.2.3). The derivation starts with a SAR refinement step. On the resulting left premise, an M['] monotonicity step yields the left premise and first (leftmost) conjunct of the right premise of rule SP&. On the resulting right premise, rule SP is used with a similar (see full proof) monotonicity step, which yields the remaining conjuncts of the right premise of rule SP&.

$$\text{SAR}\ \frac{\text{M[']}\ \dfrac{\Gamma \vdash [x' = f(x) \,\&\, \neg(P \land Q)]S \qquad S \vdash Q}{\Gamma \vdash [x' = f(x) \,\&\, \neg(P \land Q)]Q} \qquad \text{SP}\ \dfrac{\dfrac{S \vdash e \le 0 \land \dot{e} \ge \varepsilon()}{\dots}}{\Gamma \vdash \langle x' = f(x)\rangle P}}{\Gamma \vdash \langle x' = f(x) \,\&\, Q\rangle P} \quad □$$

The rules derived in Corollaries 4.25–4.28 demonstrate the flexibility of dL's refinement approach for deriving the surveyed liveness arguments as proof rules. Indeed, their derivations are mostly straightforward adaptations of the corresponding domain-free rules presented in Section 4.5, with the appropriate addition of either a COR or SAR axiomatic refinement step. Moreover, the derived rules are sound, in contrast to (most of) the liveness arguments which were missing subtle assumptions in the literature (summarized in Table 4.1). The flexibility and soundness of this chapter's approach is not limited to the surveyed liveness arguments because refinement steps can also be freely mixed-and-matched for specific liveness questions.

Example 4.29 (Strengthening)**.** The liveness property $u^2 + v^2 = 1 \to \langle\alpha_n\rangle u^2 + v^2 \ge 2$ for α_n (4.2) was proved in Example 4.22 using the staging set formula $S \equiv 1 \le u^2 + v^2 \le 2$, and provably strengthened in Example 4.24 by adding the domain constraint $u^2 + v^2 \ge 1$ with a DR$\langle\cdot\rangle$ refinement. Since S and $u^2 + v^2 \ge 2$ characterize closed sets, the refinement axiom COR proves an even stronger liveness property with the strengthened domain S, as shown in the derivation below. The derivation starts with axiom COR which yields three premises. The leftmost premise is proved by $\mathbb{R}$ since it is a real arithmetic fact; the middle premise

[8] Aside from the key notational convention (Section 3.2.1) that P, Q are semianalytic formulas which is crucial for the soundness of axiom SAR.

$u^2+v^2=1 \vdash [\alpha_n \,\&\, \neg(u^2+v^2 \geq 2)]S$ (abbreviated ①, proof elided) proves because S is an invariant of the ODE α_n (see Chapter 3), and the rightmost premise is proved in Example 4.22.

$$\text{COR}\,\dfrac{\mathbb{R}\,\dfrac{*}{u^2+v^2=1 \vdash \neg(u^2+v^2 \geq 2)} \qquad ① \qquad \dfrac{*}{u^2+v^2=1 \vdash \langle\alpha_n\rangle u^2+v^2 \geq 2}}{u^2+v^2=1 \vdash \langle\alpha_n \,\&\, S\rangle u^2+v^2 \geq 2}$$

Axiom COR extends the chain of refinements (4.6) from Example 4.22 as follows:

$$\langle\alpha_n, t'=1\rangle(t > \frac{2}{3} \vee \neg S) \xrightarrow{\text{K}\langle\&\rangle} \langle\alpha_n\rangle\neg S \xrightarrow{\text{K}\langle\&\rangle} \langle\alpha_n\rangle u^2+v^2 \geq 2 \xrightarrow{\text{COR}} \langle\alpha_n \,\&\, S\rangle u^2+v^2 \geq 2$$

The alternative staging set formula $\widetilde{S} \equiv 1 \leq u^2+v^2 < 2$ can also be used to prove Example 4.22 with a similar refinement chain (using SP_b instead of SP_c), but $\widetilde{S}$ does *not* characterize a closed set. The topological restriction of axiom COR crucially prevents its unsound use (indicated by $\lightning$ in the chain below):

$$\underbrace{\langle\alpha_n, t'=1\rangle(t > \frac{2}{3} \vee \neg\widetilde{S}) \xrightarrow{\text{K}\langle\&\rangle} \langle\alpha_n\rangle\neg\widetilde{S} \xrightarrow{\text{K}\langle\&\rangle} \langle\alpha_n\rangle u^2+v^2 \geq 2}_{\text{Similar to Example 4.22}} \underbrace{\xrightarrow{\text{COR}\lightning}}_{\text{Unsound step!}} \langle\alpha_n \,\&\, S\rangle u^2+v^2 \geq 2$$

The liveness property $\langle\alpha_n \,\&\, \widetilde{S}\rangle u^2+v^2 \geq 2$ is unsatisfiable because $\widetilde{S}$ does not overlap with $u^2+v^2 \geq 2$. Notice that the weakening of an inequality between domain constraints S and $\widetilde{S}$ leads to a wholly different conclusion! △

The refinement approach also enables the discovery of new, general liveness proof rules by combining the underlying refinement steps in alternative ways. As an example, the following chimeric proof rule combines ideas from Corollaries 4.19, 4.21, and 4.28:

Corollary 4.30 (Combination proof rule)**.** *The following proof rule is derivable in* dL*. Formula S characterizes a compact set over variables x.*

$$\text{SP}_c^k\& \quad \dfrac{\Gamma \vdash [x'=f(x) \,\&\, \neg(P \wedge Q)]S \qquad S \vdash Q \wedge \dot{e}^{(k)} > 0}{\Gamma \vdash \langle x'=f(x) \,\&\, Q\rangle P}$$

Proof Summary (Appendix B.2.3). The derivation combines refinement steps used in the derivations of $\text{dV}_{\succcurlyeq}^k$ (generalizing $\text{dV}_{\succcurlyeq}$ to higher derivatives), SP_c (compact staging sets), and SP& (refining domain constraints). □

The logical approach of dL derives complicated proof rules like $\text{SP}_c^k\&$ from a small set of sound logical axioms, which ensures their correctness. The proof rule $\text{E}_c\&$ below is derived from rule $\text{SP}_c^k\&$ (for $k=1$) and is adapted from the literature [157, Theorem 3.5], where additional restrictions were imposed on the sets characterized by Γ, P, Q, and different conditions were given compared to the left premise of $\text{E}_c\&$ (highlighted below). These original conditions were overly permissive as they are checked on sets that are smaller than necessary for soundness, see Appendix B.3 for counterexamples to those original conditions.

Corollary 4.31 (Compact eventuality [157]). *The following proof rule is derivable in* dL. *Formula $Q \land \neg P$ characterizes a compact set over variables x.*

$$\text{E}_c\&\ \dfrac{\Gamma \vdash [x' = f(x) \,\&\, \neg(P \land Q)]Q \qquad Q, \neg P \vdash \dot{e} > 0}{\Gamma \vdash \langle x' = f(x) \,\&\, Q\rangle P}$$

Proof. Rule E$_c$& is derived from SP$_c^k$& with $S \equiv Q \land \neg P$ and $k = 1$ because formula $Q \land \neg P$ is assumed to characterize a compact set, as required by rule SP$_c^k$&:

$$\text{SP}_c\ \dfrac{\text{M}[']\ \dfrac{\Gamma \vdash [x' = f(x) \,\&\, \neg(P \land Q)]Q}{\Gamma \vdash [x' = f(x) \,\&\, \neg(P \land Q)](Q \land \neg P)} \qquad \dfrac{Q, \neg P \vdash \dot{e} > 0}{Q, \neg P \vdash Q \land \dot{e} > 0}}{\Gamma \vdash \langle x' = f(x) \,\&\, Q\rangle P}$$

The M['] step uses the propositional tautology $\neg(P \land Q) \land Q \to Q \land \neg P$. □

4.6 ODE Liveness Proofs in Practice

The preceding sections show how axiomatic refinement can be used to fruitfully navigate and understand the zoo of ODE existence and liveness arguments from various applications (Table 4.1). The generality of the approach enables the sound and foundational derivation of those arguments from a parsimonious basis of refinement steps. This section provides a complementary study of how the refinement approach and its derived ODE existence and liveness proof rules are best implemented in practice. There are two canonical approaches for such an implementation:

1. Implement the foundational refinement steps and let users build their own arguments using those steps, e.g., by following the derivations and proofs from Sections 4.3–4.5.

2. Implement the zoo of proof rules from Sections 4.3–4.5 directly and let users pick from those rules for their particular ODE liveness applications.

The low-level flexibility of Approach 1 is also its drawback in practice because users need to tediously reconstruct high-level ODE liveness arguments from basic refinements for each proof. Approach 2 provides users with those high-level arguments but limits users to proof rules that have been implemented, which squanders the generality of the refinement approach. Moreover, users would still need to navigate the redundancies and tradeoffs among the zoo of proof rules to select one that is best-suited for their proof. To account for these drawbacks, this section advocates for a middle ground between those two extremes: implementations should provide users with the basic refinement steps, bundled with a set of carefully curated, high-level proof rules (Section 4.6.1) and associated proof support (Section 4.6.2) that help users navigate the common cases in their liveness proofs.

These ideas are put into practice through an implementation of ODE existence and liveness proof rules in KeYmaera X [54]. Proof rules and proof support are implemented as *tactics* in KeYmaera X [55], which are not soundness-critical. Such an arrangement allows for the implementation of useful ODE liveness proof rules and their associated proof support with KeYmaera X's sound kernel as a safeguard against implementation errors or mistakes in their derivations and side conditions. This core design decision underlying KeYmaera X is discussed elsewhere [54, 55, 142]. All of the ODE liveness examples in this chapter have been formally proved in KeYmaera X (Section 4.6.2). By leveraging existing infrastructure for hybrid programs in KeYmaera X, the implementation can also be used as part of liveness proofs for hybrid systems. For example, it is used for the liveness proofs of a case study involving a robot model driving along circular arcs in the plane [22].

The basic refinements steps from Section 4.2 and the proof rules in Sections 4.3–4.5 are mostly straightforward to implement by following their respective proofs. Thus, Sections 4.6.1 and 4.6.2 focus on a select number of new proof rules and proof support that are beneficial in the implementation. For the sake of completeness, syntactic derivations of all liveness proof rules presented in these sections are given in Appendix B.2.4.

4.6.1 Liveness Proof Rules

Atomic differential variants dV$_{\succcurlyeq}$ is a useful primitive proof rule to implement in KeYmaera X because many ODE liveness proof rules, e.g., dV$_{=}^{M}$, SP, derive from it. From a practical perspective though, rule dV$_{\succcurlyeq}$ as presented in Corollary 4.16 still requires users to provide a choice of the constant $\varepsilon()$, e.g., the proof in Example 4.17 uses $\varepsilon() = \frac{1}{2}$. The following slight rephrasing of dV$_{\succcurlyeq}$ enables a more automated implementation.

Corollary 4.32 (Existential atomic differential variants [137])**.** *The following proof rule (where $\succcurlyeq$ is either $\geq$ or $>$) is derivable in* dL*, where ε is a fresh variable and ODE $x' = f(x)$ has provable global solutions.*

$$\text{dV}_{\succcurlyeq}^{\exists}\ \dfrac{\Gamma \vdash \exists \varepsilon{>}0\, \forall x\, \big(\neg(e \succcurlyeq 0) \rightarrow \dot{e} \geq \varepsilon\big)}{\Gamma \vdash \langle x' = f(x) \rangle e \succcurlyeq 0}$$

Proof. The derivation starts with a cut of the sole premise of dV$_{\succcurlyeq}^{\exists}$ (the left premise abbreviated ① below). The existentially bound variable is renamed to δ throughout the derivation for clarity. After Skolemizing (with ∃L), rule dV$_{\succcurlyeq}$ is used with $\varepsilon() = \delta$. The universally quantified antecedent is constant for the ODE $x' = f(x)$ so it is soundly kept across the application of dV$_{\succcurlyeq}$. The proof is completed propositionally ∀L, →L.

$$\text{cut}\ \dfrac{① \qquad \text{∃L, ∧L}\ \dfrac{\text{dV}_{\succcurlyeq}\ \dfrac{\text{∀L, →L}\ \dfrac{*}{\forall x\, \big(\neg(e \succcurlyeq 0) \rightarrow \dot{e} \geq \delta\big), \neg(e \succcurlyeq 0) \vdash \dot{e} \geq \delta}}{\delta > 0, \forall x\, \big(\neg(e \succcurlyeq 0) \rightarrow \dot{e} \geq \delta\big) \vdash \langle x' = f(x) \,\&\, Q \rangle e \succcurlyeq 0}}{\exists \delta > 0\, \forall x\, \big(\neg(e \succcurlyeq 0) \rightarrow \dot{e} \geq \delta\big) \vdash \langle x' = f(x) \,\&\, Q \rangle e \succcurlyeq 0}}{\Gamma \vdash \langle x' = f(x) \,\&\, Q \rangle e \succcurlyeq 0}$$

□

Just like rule dV$_{\succcurlyeq}$, rule dV$_{\succcurlyeq}^{\exists}$ requires a positive lower bound $\varepsilon > 0$ on the derivative of e along solutions. The difference is that the premise of rule dV$_{\succcurlyeq}^{\exists}$ is rephrased to ask a purely arithmetical question about the existence of a suitable choice for ε. This can be decided automatically to save user effort in identifying ε, but such automation comes at added computational cost because the decision procedure must *find* a suitable instance of ε for the $\exists$ quantifier (or decide that none exist) rather than simply *check* a user-provided instance. Thus, the implementation gives users control over the desired degree of automation in their proof by giving them the option of either invoking an arithmetic decision procedure $\mathbb{R}$ on the premise of $\text{dV}^{\exists}_{\succcurlyeq}$ or manually instantiating the existential quantifier with a specific term for ε by rule ∃R.

Another useful variation of rule $\text{dV}_{\succcurlyeq}$ is its *semianalytic* generalization, i.e., where the goal region is described by a formula P formed from conjunctions and disjunctions of (in)equalities. Rules $\text{dV}^{M}_{=}$, SP provide examples of such a generalization, but they are indirect generalizations because users must still identify an underlying (atomic) differential variant e as input when applying either rule. In contrast, the new semianalytic generalization of $\text{dV}_{\succcurlyeq}$ below directly examines the syntactic structure of the goal region described by formula P. Its implementation is enabled by KeYmaera X's ODE invariance proving capabilities based on Chapter 3.

Corollary 4.33 (Semianalytic differential variants)**.** *Let b be a fresh variable, and term $\varepsilon()$ be constant for ODE $x' = f(x), t' = 1$. Let P be a semianalytic formula in the following normal form* (3.7) *and G_P be its corresponding ε-progress formula, also in normal form* (3.7)*:*

$$P \equiv \bigvee_{i=0}^{M} \Big(\bigwedge_{j=0}^{m(i)} e_{ij} \geq 0 \land \bigwedge_{j=0}^{n(i)} \tilde{e}_{ij} > 0 \Big)$$

$$G_P \equiv \bigvee_{i=0}^{M} \Big(\bigwedge_{j=0}^{m(i)} e_{ij} - (b + \varepsilon()t) \geq 0 \land \bigwedge_{j=0}^{n(i)} \tilde{e}_{ij} - (b + \varepsilon()t) \geq 0 \Big)$$

The following proof rule is derivable in dL*, where the ODE $x' = f(x)$ has provable global solutions, and $(\dot{\neg P})^{(*)}$, $(\dot{G_P})^{(*)}$ are semianalytic progress formulas Def. 3.24 with respect to $x' = f(x), t' = 1$.*

$$\text{dV}\ \frac{\neg P, (\dot{\neg P})^{(*)}, G_P \vdash (\dot{G_P})^{(*)}}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x) \rangle P}$$

$$\text{dV}^{\exists}\ \frac{\Gamma \vdash \exists \varepsilon{>}0\, \forall b\, \forall t\, \forall x\, \big(\neg P \land (\dot{\neg P})^{(*)} \land G_P \rightarrow (\dot{G_P})^{(*)}\big)}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x) \rangle P}$$

Proof Summary (Appendix B.2.4). Rule $\text{dV}^{\exists}$ is derived from dV like the derivation of rule $\text{dV}^{\exists}_{\succcurlyeq}$ from $\text{dV}_{\succcurlyeq}$. The derivation of dV is similar to rules $\text{dV}^{\Gamma}_{\succcurlyeq}$, $\text{dV}_{\succcurlyeq}$, but replaces the use of rule $\text{dI}_{\succcurlyeq}$ with complete ODE invariance reasoning (Chapter 3). The fresh variable b is used as a lower bound along solutions of the ODE of the value of all terms $e_{ij}, \tilde{e}_{ij}$ appearing in P. □

The intuition behind rule dV is similar to rule $\text{dV}_{\succcurlyeq}$, as long as the solution has not yet reached the goal P, it grows towards P at "rate" $\varepsilon()$. The technical challenge is how to formally phrase the "rate" of growth for a semianalytic formula P, which does not have a well-defined notion of derivative. Rule dV uses the ε-progress formula G_P, together with the semianalytic progress formulas $(\dot{\neg P})^{(*)}$, $(\dot{G_P})^{(*)}$ and dL's complete ODE invariance reasoning from Chapter 3 for this purpose. These formulas give sufficient, although implicit, arithmetical conditions for proving liveness for P. Rule $\text{dV}^{\exists}$ rephrases dV with an arithmetical premise, similar to how $\text{dV}^{\exists}_{\succcurlyeq}$ rephrases $\text{dV}_{\succcurlyeq}$, to give users the added flexibility of choosing between invoking an automated decision procedure or manually instantiating the existential quantifier for ε and reasoning about the resulting progress formulas. More explicit arithmetical premises for dV, $\text{dV}^{\exists}$ can be obtained by unfolding the definitions of $(\dot{\neg P})^{(*)}$, $(\dot{G_P})^{(*)}$ as exemplified below.

Example 4.34 (Non-differentiable progress functions [176]). Consider the following liveness formula with two inequalities in its postcondition:

$$\langle u' = -u \rangle (-1 \leq u \leq 1) \tag{4.20}$$

Formula (4.20) can be written equivalently with an atomic inequality using the $\min$ function:

$$\langle u' = -u \rangle \min(1-u, u+1) \geq 0 \tag{4.21}$$

However, the postcondition of (4.21) is not a formula of real arithmetic and it does not have well-defined dL semantics. Indeed, rule dV$_{\succcurlyeq}$ does not prove (4.21) because the Lie derivative of its postcondition is not well-defined. One possible solution is to generalize dV$_{\succcurlyeq}$ by considering directional derivatives of continuous (but non-differentiable) functions such as $\min, \max$ [176, Section 5.2]. However, justifying the correctness of this option would require delicate changes to dL semantics [21, 142, 143]. Rule dV instead proves (4.20) directly without requiring rephrasing, nor complications associated with directional derivatives. The proof is as follows, with $\varepsilon() = 1$ and $P \equiv u+1 \geq 0 \land 1-u \geq 0, G_P \equiv u+1-(b+t) \geq 0 \land 1-u-(b+t) \geq 0$. The left conjunct in the succedent is abbreviated with $R \equiv u+1-(b+t) = 0 \to -u-1 > 0$ and the right conjunct is omitted for brevity since the subsequent argument given below is symmetric.

$$
\text{dV}\,\dfrac{\dfrac{\dfrac{\mathbb{R}\,\dfrac{*}{u+1<0 \lor 1-u<0, u+1-(b+t) \geq 0 \land 1-u-(b+t) \geq 0 \vdash R \land \dots}}{\neg P, G_P \vdash \dot{(G_P)}^{(*)}}}{\neg P, \dot{(\neg P)}^{(*)}, G_P \vdash \dot{(G_P)}^{(*)}}}{\vdash \langle u' = -u \rangle (-1 \leq u \leq 1)}
$$

The proof starts by using rule dV, where the assumption $\dot{(\neg P)}^{(*)}$ in its premise is weakened as it is unnecessary for the proof. Unfolding the definition of $\dot{(G_P)}^{(*)}$ and simplifying leaves an arithmetical question in the succedent with two conjuncts. The left conjunct R in the succedent is proved by $\mathbb{R}$ because the assumptions $u+1-(b+t) = 0$ and $u+1-(b+t) \geq 0 \land 1-u-(b+t) \geq 0$ imply $1-u \geq u+1$. This, in turn, implies $-u-1 > 0$ using the assumption $u+1 < 0 \lor 1-u < 0$.

More generally, for a liveness postcondition comprising a conjunction of atomic inequalities $e \succcurlyeq 0 \land \tilde{e} \succcurlyeq 0$ (where $\succcurlyeq$ is either $\geq$ or $>$ in either conjunct), the premise resulting from applying dV can be simplified in real arithmetic to the following arithmetical premise:

$$\neg(e \succcurlyeq 0 \land \tilde{e} \succcurlyeq 0) \vdash (e < \tilde{e} \to \dot{e} > \varepsilon()) \land (e > \tilde{e} \to \dot{\tilde{e}} > \varepsilon()) \land (e = \tilde{e} \to \dot{e} > \varepsilon() \land \dot{\tilde{e}} > \varepsilon()) \tag{4.22}$$

The arithmetical premise (4.22) is equivalent to the arithmetical progress conditions for $\min(p,q) \geq 0$ [176, Example 14], and both are decidable in real arithmetic. The intuition behind (4.22) is that whenever e is further from the goal than $\tilde{e}$, then e is required to make ε progress towards the goal (symmetrically when $\tilde{e}$ is further than e from the goal). A similar simplification of dV for a disjunctive postcondition $e \succcurlyeq 0 \lor \tilde{e} \succcurlyeq 0$ is shown in (4.23), which asks for the term closer to the goal to make ε progress towards the goal instead. Further simplifications for semianalytic formulas P are obtained as nested combinations of (4.22) and (4.23).

$$\neg(e \succcurlyeq 0 \lor \tilde{e} \succcurlyeq 0) \vdash (e < \tilde{e} \to \dot{\tilde{e}} > \varepsilon()) \land (e > \tilde{e} \to \dot{e} > \varepsilon()) \land (e = \tilde{e} \to \dot{e} > \varepsilon() \lor \dot{\tilde{e}} > \varepsilon()) \tag{4.23}$$

This example shows the intricate definition of semianalytic progress formulas, even for the simple-looking conjunctive postcondition $-1 \leq u \leq 1$, which highlights the need for a careful and trustworthy implementation of rules dV, $\mathrm{dV}^{\exists}$, as provided by KeYmaera X. △

The variations of $\mathrm{dV}_{\succcurlyeq}$ shown in Corollaries 4.32 and 4.33 (and their implementation) allow users to focus on high-level liveness arguments in KeYmaera X rather than low-level derivation steps. Another key usability improvement afforded by an implementation is the sound and automatic enforcement of the appropriate side conditions for every proof rule. The common side conditions for ODE liveness proof rules presented in this chapter can be broadly classified as follows:

1. Freshness side conditions on variables, e.g., in rules $\mathrm{dV}_{\succcurlyeq}$, $\mathrm{dV}^{\exists}_{\succcurlyeq}$, dV, $\mathrm{dV}^{\exists}$. These are automatically enforced in the implementation because KeYmaera X's kernel insists on fresh names when required for soundness. Various forms of renaming with fresh variables are automatically supported [142].

2. Global existence of ODE solutions. These are semi-automatically proved (Section 4.6.2).

3. Topological side conditions, e.g., in axiom COR and rules $\mathrm{dV}_{\succcurlyeq}\&$, $\mathrm{dV}^{M}_{=}\&$. These conditions are important to correctly enforce because they may otherwise lead to subtle soundness errors (Section 4.5). The implementation uses syntactic criteria for checking these side conditions (Appendix B.1.3).

An example topological refinement axiom (Lemma 4.3) and its corresponding proof rule implemented in KeYmaera X with syntactic topological side conditions is given next.

Lemma 4.35 (Closed domain refinement axiom)**.** *The following topological $\langle \cdot \rangle$ ODE refinement axiom is sound, where formula Q characterizes a topologically closed set over variables x, and formula $\mathring{Q}$ characterizes the topological interior of the set characterized by Q.*

$$\text{CR}\ \neg P \land [x' = f(x) \,\&\, R \land \neg P]\mathring{Q} \rightarrow \big(\langle x' = f(x) \,\&\, R\rangle P \rightarrow \langle x' = f(x) \,\&\, Q\rangle P\big)$$

Proof in Appendix B.1.2.

Corollary 4.36 (Closed domain refinement rule)**.** *The following proof rule is derivable in* dL*, where formula Q is formed from finite conjunctions and disjunctions of non-strict inequalities $\geq, \leq$, and formula $Q^{>}_{\geq}$ is identical to Q but with strict inequalities $>, <$ in place of $\geq, \leq$ respectively.*

$$\text{cR}\ \frac{\Gamma \vdash Q \quad \Gamma \vdash [x' = f(x) \,\&\, R \land \neg P \land Q]Q^{>}_{\geq} \quad \Gamma \vdash \langle x' = f(x) \,\&\, R\rangle P}{\Gamma \vdash \langle x' = f(x) \,\&\, Q\rangle P}$$

Proof in Appendix B.2.4.

Axiom CR is a variant of axiom COR with different topological conditions. It says that if the ODE solution can reach goal P while staying in domain R then it can also reach that goal while staying in the new (closed) domain Q, provided that it stays within the *interior* $\mathring{Q}$ of the new domain while it has not yet reached P. Solutions cannot sneak out of the topologically open interior $\mathring{Q}$ as it enters the goal because, by definition of an open set, the solution must locally remain in $\mathring{Q}$ for a short time as it enters the goal (see the proof for a detailed explanation). In contrast to the semantical conditions of CR, its corresponding derived rule cR gives syntactic side conditions for the formulas $Q, Q^{>}_{\geq}$ which are easily checked in an implementation. In particular, formula $Q^{>}_{\geq}$, which syntactically underapproximates the interior $\mathring{Q}$, can be automatically generated from Q through its syntactic structure. Another advantage of the derived rule cR is that the closed domain constraint Q can be additionally assumed when proving that solutions stay within $Q^{>}_{\geq}$ in its middle premise. This addition makes rule cR a powerful primitive for refining domain constraints amongst other options such as axiom DR$\langle\cdot\rangle$.

4.6.2 Proof Support

Beyond enabling the sound implementation of complex ODE liveness proof rules such as those in Section 4.6.1, tactics can also provide substantial proof support for users.

Automatic Dependency Ordering

Recall derived axiom GBEx from Corollary 4.12, which proves (global) existence of solutions for an ODE $x' = f(x)$. Users of the axiom must still identify precisely which dependency order (4.9) to use, and provide the sequence of bounded sets B_i for each group of variables y_i involving nonlinear ODEs. The canonical choice of such a dependency order can be automatically produced by a tactic using a topological sort of the *strongly connected components* (SCCs)[9] of the dependency graph of the ODE.

More precisely, to prove global existence for an ODE $x' = f(x)$, consider the dependency graph G where each variable x_i is a vertex and with a directed edge $x_i \longrightarrow x_j$ if the RHS $f_j(x)$ for x'_j depends on free variable x_i. First, compute the SCCs of G, and then topologically sort the SCCs. The groups of variables y_i in dependency order can be chosen according to the vertices in each SCC in topological order. An illustrative dependency graph with four SCCs for the following 8-dimensional ODE is shown in Fig. 4.4.

$$\begin{aligned}&x'_1 = x_5, x'_2 = x_3 + x_6^2, x'_3 = x_3^2, x'_4 = x_1 + x_3^2 + x_6^2,\\ &x'_5 = x_4, x'_6 = x_2^2, x'_7 = x_8, x'_8 = -x_7\end{aligned} \tag{4.24}$$

After finding the appropriate SCC-induced dependency order (as in Fig. 4.4), the global existence tactic can prove global existence for the variable clusters y_i that have affine dependencies within the cluster automatically. For example, the SCC $y_4 \equiv \{x_1, x_4, x_5\}$ has affine dependencies because the RHS of the ODEs x'_1, x'_4, x'_5 are affine in x_1, x_4, x_5, so the solution of ODE (4.24) is automatically proved to be global in those variables following the proof of Corollary 4.8. The generated dependency order enables such a proof even though the RHS of x'_4 depends nonlinearly on variables x_3, x_6 from earlier clusters. For the SCC $y_3 \equiv \{x_2, x_6\}$ which has nonlinear dependencies on x_2, x_6, users are prompted to input a bounded set (or a bound on derivatives) over variables x_2, x_6 in order to prove global existence for those variables. This

[9] A strongly connected component of a directed graph is a maximal subset of vertices that are pairwise connected by paths.

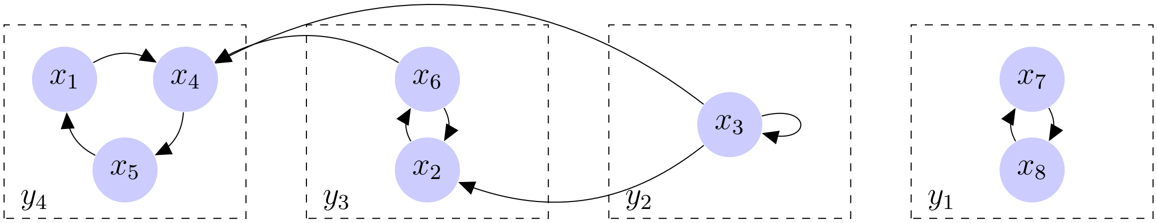

Figure 4.4: A dependency graph for the ODE (4.24) over the variables $x_1, \ldots, x_8$. There is a directed edge $x_i \longrightarrow x_j$ if the RHS for x_j' depends on free variable x_i. Each dashed rectangle is a strongly connected component. Topologically sorting these components (according to the order induced by the edges) yields one possible grouping of the variables $y_1, \ldots, y_4$ in dependency order. The vertices in y_1 are not connected to those in y_2, y_3, y_4, so the order between these groups can be chosen arbitrarily.

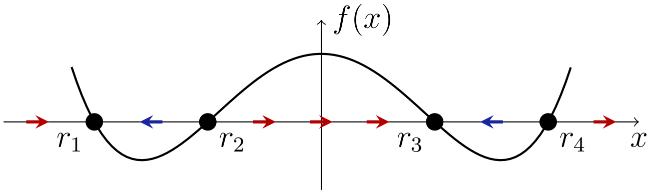

Figure 4.5: The univariate ODE $x' = x^4 - 5x^2 + 4$ is illustrated by plotting its RHS $f(x) = x^4 - 5x^2 + 4$ (vertical axis) against x (horizontal axis). Points on the horizontal axis evolve towards the right (red arrow) when $f(x) \geq 0$ and towards the left (blue arrow) when $f(x) \leq 0$. The fixed points r_1, r_2, r_3, r_4 are roots of the polynomial RHS where $f(x) = 0$. These fixed points either attract trajectories (like r_1, r_3), or repel them (like r_2, r_4). All points on the horizontal axis evolve asymptotically towards exactly one fixed point or approach ∞.

continues similarly for the SCCs y_2 (nonlinear dependency) and y_1 (affine dependency) until global existence is proved for the full ODE. This semi-automated proof support minimizes the manual effort required of the user in proving global existence by focusing their attention on the nonlinear parts of the ODE that may cause finite-time blowup of solutions.

To drive global existence proof automation further, key special cases can be added to the method described above. One such special case for univariate ODEs is shown below.

Remark 4.37 (Global existence for univariate ODEs). Consider the case where a variable group has just one variable and no incoming dependencies, e.g., $y_2 \equiv \{x_3\}$ in Fig. 4.4 or α_b (4.7). Global existence for such univariate *polynomial* ODEs is decidable [68], even if the RHS is highly nonlinear, because all of its solutions either asymptotically approach one of the (finitely many) roots of the polynomial RHS or diverge to infinity.

This result is best illustrated through the dynamical systems view of ODEs shown in Fig. 4.5 for the ODE $x' = x^4 - 5x^2 + 4$. This example ODE has global solutions from all initial states satisfying $x \leq r_4$ because the solution from all such states are globally attracted to one of the fixed points. Conversely, for all other initial conditions ($x > r_4$), the ODE blows up in finite time because the RHS is quartic in x.

More generally, for a nonlinear univariate polynomial ODE $x' = f(x)$ and initial assumptions Γ, it suffices to check validity of the following arithmetical sequent to decide global existence from a set of initial assumptions Γ on the state variable x:

$$\Gamma \vdash \exists r \left(f(r) = 0 \land (\underbrace{f(x) \geq 0 \land r \geq x}_{\textcircled{a}} \lor \underbrace{f(x) \leq 0 \land r \leq x}_{\textcircled{b}})\right)$$

The existentially quantified variable r corresponds to a fixed point (a root with $f(r) = 0$). Disjunct $\textcircled{a}$ checks whether the solution approaches r from the left, e.g., the points between r_2 and r_3 in Fig. 4.5 approach r_3 from the left. Alternatively, disjunct $\textcircled{b}$ checks whether the solution approaches r from the right. The implementation checks validity of this sequent for univariate nonlinear ODEs and then proves global existence using BDG$\langle\cdot\rangle$ because the solution is provably trapped between the initial value of x and the fixed point r.

Differential Cuts for Liveness Proofs

Differential cuts dC provide a convenient way to structure and stage safety proofs for ODEs in dL. An in-depth discussion is available elsewhere [144, Part II], but the idea is illustrated by the following derivation outline:

$$\text{dC}\,\dfrac{\Gamma \vdash [x' = f(x) \,\&\, Q]C_1 \qquad \text{dC}\,\dfrac{\cdots \qquad \text{dC}\,\dfrac{\text{dW}\,\dfrac{Q \land C_1 \land C_2 \land \cdots \land C_n \vdash P}{\vdots}}{\Gamma \vdash [x' = f(x) \,\&\, Q \land C_1 \land C_2]P}}{\Gamma \vdash [x' = f(x) \,\&\, Q \land C_1]P}}{\Gamma \vdash [x' = f(x) \,\&\, Q]P}$$

The outline uses a sequence of differential cut steps to progressively add cuts $C_1, C_2, \ldots, C_n$ to the domain constraint. A final dW step completes the proof when the postcondition P is already implied by the (now strengthened) domain constraint. Intuitively, the differential cuts are akin to dynamical lemmas in this derivation. For example, by proving the premise $\Gamma \vdash [x' = f(x) \,\&\, Q]C_1$, the cut C_1 can now be assumed in the domain constraints of subsequent steps. Just like the cut rule from sequent calculus, differential cuts dC allow safety proofs for ODEs to be staged through a sequence of lemmas about those ODEs.

For proof modularity and maintainability, it is desirable to enable a similar kind of staging for ODE liveness proofs. Suppose that the formula $[x' = f(x) \,\&\, Q]C$ has been proved as a cut:

$$\text{cut}\,\dfrac{\Gamma \vdash [x' = f(x) \,\&\, Q]C \qquad \cdots}{\Gamma \vdash \langle x' = f(x) \,\&\, Q\rangle P}$$

The challenge is how to (soundly) use this lemma in subsequent derivation steps (shown as $\cdots$). Naïvely replacing Q with $Q \land C$ in the domain constraint of the succedent is sound but may even do more harm than good because the resulting ODE liveness question becomes more difficult (Section 4.5). The refinement-based approach to ODE liveness provides a natural answer: recall that each refinement step in the chain (4.6) requires the user to prove an additional box modality formula. The insight is that, for these box modality formulas, any relevant lemmas that have been proved can be soundly added to the domain constraint. For example, suppose that rule K$\langle\&\rangle$ is used to continue the proof after the cut. The left premise of K$\langle\&\rangle$ can now be strengthened to include C in its domain constraint:

$$\text{K}\langle\&\rangle\;\dfrac{\text{dC}\;\dfrac{\Gamma, [x'{=}f(x)\,\&\,Q]C \vdash [x'{=}f(x)\,\&\,Q{\land}\lnot P{\land}C]\lnot G}{\Gamma, [x'{=}f(x)\,\&\,Q]C \vdash [x'{=}f(x)\,\&\,Q{\land}\lnot P]\lnot G} \qquad \Gamma, [x'{=}f(x)\,\&\,Q]C \vdash \langle x'{=}f(x)\,\&\,Q\rangle G}{\Gamma, [x'{=}f(x)\,\&\,Q]C \vdash \langle x'{=}f(x)\,\&\,Q\rangle P}$$

Users could manually track and apply lemmas using dC as shown above, but this becomes tedious in larger liveness proofs. The implementation instead provides users with tactics that automatically search the antecedents Γ for compatible assumptions that can be used to strengthen the domain constraints. These tactics also use a form of *ODE unification* when determining compatibility. More precisely, consider the sequent $\Gamma \vdash [x' = f(x) \,\&\, Q]P$, which may arise as a box refinement during a liveness proof. An antecedent formula $[y' = g(y) \,\&\, R]C$ in Γ is called a *compatible assumption* for the succedent $[x' = f(x) \,\&\, Q]P$ if:

1. The set of ODEs $y' = g(y)$ is a subset of the set of ODEs $x' = f(x)$ and $g(y)$ does not mention free variables in $x \setminus y$. This is order-agnostic, e.g., the ODE $u' = v, v' = u$ is a subset of the ODE $v' = u, u' = v, w' = u + v + w$.

2. The domain constraint Q implies domain constraint R, i.e., $Q \to R$ is valid.

Under these conditions, the ODE $y' = g(y) \,\&\, R$ permits more trajectories than the ODE $x' = f(x) \,\&\, Q$. Thus, if formula C is always true along solutions of the former ODE, then it also stays true along solutions of the latter. Combining compatible assumptions with implementations of liveness proof rules yields turbo-charged versions of those rules. For example, in rule dV$^{\exists}_{\succcurlyeq}$, instead of simply assuming the negation of the postcondition ($\neg(e \succcurlyeq 0) \to \cdots$) when determining the existence of suitable ε, *all* postconditions of compatible assumptions can be assumed, e.g., with $\neg(e \succcurlyeq 0) \land C \to \cdots$ for postcondition C of a compatible assumption.

Microbenchmarks

The KeYmaera X implementation is used to formally prove all of the ODE liveness examples from this chapter and elsewhere [22, 176]. Table 4.2 provides a summary of statistics from these proofs, where all experiments were run on an Ubuntu 18.04 laptop with a 2.70 GHz Intel Core i7-6820HQ CPU and 16GB memory. None of the proofs have been optimized to favor any specific metric. The specific timings and proof steps are naturally subject to change on different hardware and as various aspects of the KeYmaera X theorem prover are improved. Nevertheless, the key takeaways from these microbenchmarks remain broadly applicable.

(M)anual and (A)utomatic Proofs. The implementation provides users with powerful proof support but also exposes low-level primitives for users who prefer more fine-grained control over (parts of) their proofs. Both types of proofs are shown for this chapter's examples in Table 4.2. Proofs that heavily exploit the proof support and automation are more convenient for users and require fewer manual tactic invocations. Indeed, all of the automated proofs require fewer tactic invocations than the corresponding manual proofs (where both proofs are available for comparison). An example of this gap is Example 4.34, where the 28 step manual proof uses

Table 4.2: Proof statistics for ODE existence and liveness properties proved using the implementation. For this chapter's Examples 4.5–4.34, two proofs are presented: (M)anual proofs closely follow the pen-and-paper derivations shown in this chapter, while (A)utomatic proofs make extensive use of the implemented proof support. The cells in **bold** font indicate lower (more desirable) values. The stronger ODE liveness property proved in Example 4.29 implies those from Examples 4.22 and 4.24. Examples 11, 12 and 15 refer to the correspondingly numbered examples from Sogokon and Jackson [176]. "Dimension" is the number of continuously evolving state variables in the ODEs; "Parameters" is the number of parameters (non-state variables) in the liveness specification; "Max Degree" is the maximum degree of polynomials with respect to the state variables in the liveness specification; "Tactic Steps" counts the number of (manual) user proof steps; "Kernel Steps" counts the number of internal steps taken by the soundness-critical KeYmaera X kernel; and "Proof Time" measures the time taken (in seconds, averaged over 5 runs, rounded to 3 decimal places) for the proof to execute in KeYmaera X.

Liveness Property	Dimension	Parameters	Max Degree
Example 4.5	1	1	2
Example 4.9	2	0	2
Example 4.11	2	0	3
Example 4.17	2	0	2
Example 4.29 (Examples 4.22 and 4.24)	2	0	3
Example 4.34	1	0	1
Sogokon and Jackson [176, Example 11]	2	0	4
Sogokon and Jackson [176, Example 12]	2	0	2
Sogokon and Jackson [176, Example 15]	2	0	1
Bohrer et al. [22, Goal Position Reachable]	3	4	2
Bohrer et al. [22, Velocity Bounds Reachable]	3	4	2

Liveness Property	Tactic Steps		Kernel Steps		Proof Time (s)	
	(M)	(A)	(M)	(A)	(M)	(A)
Example 4.5	7	**3**	**2040**	12156	**1.778**	3.716
Example 4.9	8	**2**	962	**898**	0.220	**0.203**
Example 4.11	7	**3**	**1562**	1580	**0.551**	0.759
Example 4.17	29	**5**	3958	**3501**	3.286	**3.034**
Example 4.29 (Examples 4.22 and 4.24)	50	**20**	**5549**	6141	**1.714**	1.952
Example 4.34	28	**1**	**1747**	2571	**0.575**	0.990
Sogokon and Jackson [176, Example 11]	-	50	-	11272	-	9.090
Sogokon and Jackson [176, Example 12]	-	19	-	4818	-	1.388
Sogokon and Jackson [176, Example 15]	-	1	-	4781	-	1.730
Bohrer et al. [22, Goal Position Reachable]	-	34	-	8159	-	2.182
Bohrer et al. [22, Velocity Bounds Reachable]	-	37	-	10521	-	3.042

a refinement step K⟨&⟩ to equivalently replace postcondition $-1 \leq u \leq 1$ by $u^2 \leq 1$ and then completes the proof by manually following the derivation of rule $\mathrm{dV}_{\succcurlyeq}$. In contrast, the automated proof requires just one dV step.

On the other hand, the automated proofs are slower than their manual counterparts on four out of six examples. Most of this overhead arises when there is significant proof search in the automation. In particular, the automated proof of Example 4.5 is significantly slower and requires almost six times more kernel steps compared to its manual counterpart. This gap arises because the automated proof uses the decision procedure for univariate global existence outlined in Remark 4.37 while the manual proof uses the direct argument in Example 4.5. However, the latter proof required user insight about the physical system as a model of air resistance (see Example 4.5). This illustrates the need for a flexible implementation that lets users navigate the convenience and efficiency tradeoff according to their needs and proof insights.

Finally, the automated proofs are in fact *faster* for Examples 4.9 and 4.17, which both involve linear ODEs. The speedups here can be attributed to the well-tuned implementation of global existence proofs for affine systems and to the use of rule $\mathrm{dV}_{\succcurlyeq}^{\exists}$ for the latter example. Thus, the aforementioned tradeoff can be further skewed towards favoring user convenience by tuning the implemented automation.

Trusted Kernel with Untrusted Tactics. All of the proofs in Table 4.2 make extensive use of KeYmaera X's existing tactics framework [55] to handle low-level interactions with KeYmaera X soundness-critical kernel, as shown by the large number of kernel steps that each proof requires. The soundness guarantee provided by the KeYmaera X kernel makes this implementation effort a worthy tradeoff because it ensures that the proved results in Table 4.2 are trustworthy *without* needing to trust the implementation of the tactics.

Applicability. The insights of this section are not limited to this chapter's examples and they scale to larger case studies (Table 4.2). Notably, the example from Sogokon and Jackson [176, Example 11] consists of two liveness sub-properties for the same ODE, which makes it the largest (and slowest) microbenchmark. The examples from Bohrer et al. [22] are liveness properties drawn from a larger case study with a hybrid system model of a robot driving along circular arcs in the plane [22]. Proof automation is indispensable for handling the scale of these proofs.

4.7 Related Work

Existence and Liveness Proof Rules. The ODE liveness arguments surveyed in this chapter were originally presented in various notations, ranging from proof rules [137, 176, 191] to other mathematical notation [156, 157, 159, 176]. All of them were justified directly through semantical or mathematical means. This chapter unifies and corrects all of these arguments, and presents them as dL proof rules which are syntactically derived by refinement from dL axioms.

This chapter is also the first to present a deductive approach for syntactic proofs of existence properties for ODEs. In the surveyed liveness arguments [137, 156, 157, 159, 176, 191], sufficient existence duration is either assumed explicitly or is implicitly used in the correctness proofs. Such a hypothesis is unsatisfactory, since the global existence of solutions for (nonlinear) ODEs is a non-trivial question; in fact, it is undecidable even for polynomial ODEs [68]. Formal proofs of any underlying existence assumptions thus yield stronger (unconditional) ODE liveness proofs. Of course, such existence properties are an additional proof burden, but Section 4.6 also shows that proof support can help by automating easy existence questions, e.g., for affine systems where global existence is well-known. A related problem arising in the study of hybrid systems is *Zeno phenomena* [75, 214], where a trajectory of a hybrid model makes infinitely many (discrete) transitions in finite (continuous) time. Like finite-time blow up, Zeno phenomena typically occur as abstraction artifacts of hybrid systems models, and they do not occur in real systems. Thus, analogous to the question of global existence, absence of Zeno phenomena must either be assumed (or Zeno trajectories explicitly excluded) [75, 137], or proved when specifying and verifying properties of such systems [214].

The refinement-based approach to ODE existence and liveness proofs underlies this chapter's implementation described in Section 4.6. Compared to an earlier implementation in KeYmaera [147], where rules like $\mathrm{dV}_{\succcurlyeq}\&$ are implemented monolithically, this chapter's approach and implementation build those rules from smaller building blocks which yields a flexible implementation together with powerful (core-checked) proof support. The high-level lessons discussed in Section 4.6 are also broadly applicable to other deductive tools for ODEs and hybrid systems [50, 206] that all currently lack support for ODE liveness proofs.

Other Liveness Properties. The liveness properties studied in this chapter are the continuous analogues of *eventually* [109] or *eventuality* [157, 176] from temporal logics. In discrete settings, temporal logic specifications give rise to a zoo of other liveness properties [109]. In continuous settings, *weak eventuality* (requiring *almost all* initial states to reach the goal region) and *eventuality-safety* have been studied [156, 157]. In adversarial settings, *differential game variants* [143] enable proofs of winning strategies for differential games. In dynamical systems and controls, the study of *asymptotic stability* requires both stability (an invariance property) with asymptotic attraction towards a fixed point or periodic orbit (an eventuality-like property) [33, 159]. For hybrid systems, various authors have proposed generalizations of classical asymptotic stability, such as *persistence* [179], *stability* [151], and *inevitability* [46]. *Controlled* versions of these properties are also of interest, e.g., *(controlled) reachability and attractivity* [1, 191]. Eventuality(-like) properties are fundamental to all of these advanced liveness properties. The formal understanding of eventuality in this chapter is therefore a key step towards enabling formal analysis of more advanced liveness properties.

Automated Liveness Proofs. Automated reachability analysis tools [31, 53] can also be used to answer certain liveness verification questions. For an ODE and initial set $\mathcal{X}_0$, computing an over-approximation $\mathcal{O}$ of the reachable set $\mathcal{X}_t \subseteq \mathcal{O}$ at time t shows that *all* states in $\mathcal{X}_0$ reach $\mathcal{O}$ at time t [179] (if solutions do not blow up). Similarly, an under-approximation $\mathcal{U} \subseteq \mathcal{X}_t$ shows that *some* state in $\mathcal{X}_0$ eventually reaches $\mathcal{U}$ [67] (if $\mathcal{U}$ is non-empty). Neither approach handles domain constraints [67, 179] and, unlike deductive approaches, the use of reachability tools limits them to proving liveness specifications with concrete time bounds t and bounded initial sets $\mathcal{X}_0$. Deductive liveness approaches can also be (partially) automated, as shown in Section 4.6. Lyapunov functions guaranteeing (asymptotic) stability can be found by sum-of-squares (SOS) optimization [130]. Liveness arguments can be similarly combined with SOS optimization to find suitable differential variants [156, 157]. Other approaches are possible, e.g., a constraint solving-based approach can be used for finding the so-called *set Lyapunov functions* [159] (e.g., the term e used in SLyap, SLyap$\&$). Crucially, automated approaches must ultimately be based on sound underlying liveness arguments. The correct justification of these arguments is precisely what this chapter enables.

Refinement Calculi. This chapter's view of ODE liveness arguments as step-by-step refinements is closely related to *refinement proof calculi* [11, 93]. The shared idea is that the proof of a complex property, like ODE liveness or program correctness, should be broken down into (simpler) step-by-step refinements. The key difference is that, for refinement calculi, refinement typically takes place between programs (or implementations) and their specification. For example, a concrete implementation β is said to *refine* its abstract specification α if the set of transitions of β is a subset of those of α [11]. Proving such a refinement for hybrid programs α, β would, for example, prove the implication:

$$\langle\beta\rangle P \rightarrow \langle\alpha\rangle P \tag{4.25}$$

Program refinement is not directly applicable to this chapter's focus on proving liveness for specific ODEs. Instead, as hinted by (4.25), program refinement plays an important role for generalizing this chapter's results beyond ODEs to hybrid systems, where, e.g., one may use implications like (4.25) as part of a refinement chain (4.6). There are a number of refinement calculi for hybrid systems [30, 50, 106, 164]. Notably, *differential refinement logic* [106] formally extends dL with a refinement operator $\beta \leq \alpha$, and can be used together with this chapter's results. Another direction for generalizing this chapter's results is to consider larger classes of continuous dynamics, such as differential inclusions, differential-algebraic constraints [137], and differential games [143]. These open up the possibility of proving refinements between concrete (ODE) descriptions and their more abstract continuous counterparts [47, 50, 137, 143].

4.8 Discussion

This chapter presents a refinement-based approach for proving liveness and, as a special case, global existence properties for ODEs in dL. The associated KeYmaera X implementation demonstrates the utility of this approach for formally proving concrete ODE liveness questions. Beyond the particular proof rules derived in this chapter, the exploration of new and more general ODE liveness proof rules is enabled by simply piecing together more refinement steps in dL, or in the KeYmaera X implementation of those steps. Given its wide applicability and correctness guarantees, this approach is a suitable framework for justifying ODE liveness arguments, even for readers less interested in the logical aspects.

Chapter 5

Stability for Ordinary Differential Equations

This chapter studies deductive *stability* verification for ordinary differential equations (ODEs) in dL. Stability is required for real-world controlled systems as it ensures that those systems can tolerate small, real-world perturbations around their desired operating states. In contrast to the previous chapters, the question of *how* to formally specify stability is interesting because there are numerous variations of stability properties of interest in the literature, each with subtly different specifications. The key insight is to specify ODE stability by suitably nesting dL's dynamic modalities with first-order logic quantifiers. Elucidating the logical structure of stability properties in this way has three key benefits: *i)* it provides a flexible means of formally specifying various stability properties of interest in the common language of dL, *ii)* it yields rigorous proofs of those stability properties using dL's ODE safety and liveness proof principles from Chapters 3 and 4, and *iii)* it enables formal analysis of the relationships between various stability properties which, in turn, inform proofs of those properties. These ODE stability proofs lay the groundwork for the study of hybrid (switched) system stability proofs in Chapter 6.

5.1 Introduction

The study of stability has its roots in efforts to understand mechanical systems, particularly those arising in celestial mechanics [77, 98, 153]. Today, it is an important part of numerous applications in dynamical systems [187] and control theory [71, 89]. For example, in feedback control systems [71, 89], stability of continuous controllers modeled by ODEs is a key correctness requirement [6] that deserves fully rigorous proofs *alongside* proofs of other key properties such as the ODE safety and liveness properties studied in Chapters 3 and 4. Despite this, formal stability verification has received less attention compared to proofs of safety and liveness, e.g., through reachability or deductive techniques [45].

Stability for a continuous system (or ODEs) requires that *i)* its system state always stays close to some desired operating state(s) when initially slightly perturbed from those operating state(s) and *ii)* those perturbations are eventually dissipated so that the system returns to a desired operating state. These properties are especially crucial for engineered systems because those systems must be robust to real-world perturbations deviating from idealized models.

Simple pendulums provide canonical examples of stability phenomena: they are always observed to settle in the rest position of Fig. 5.1 (bottom) after some time regardless of how they are initially released. In contrast, the inverted pendulum in Fig. 5.1 (top) is *theoretically* also at a resting position but can only be observed transiently in practice because the slightest real-world perturbation will cause the pendulum to fall due to gravity. Stability explains these observations—the resting position is (asymptotically) stable while the inverted position is unstable and requires active control to ensure its stability. Proofs of safety and liveness properties are still required for the inverted pendulum under control, e.g., its controller must never generate unsafe amounts of torque and the pendulum must eventually reach the inverted position. The *triumvirate* of safety, liveness, and stability is required for holistic correctness of the inverted pendulum controller.

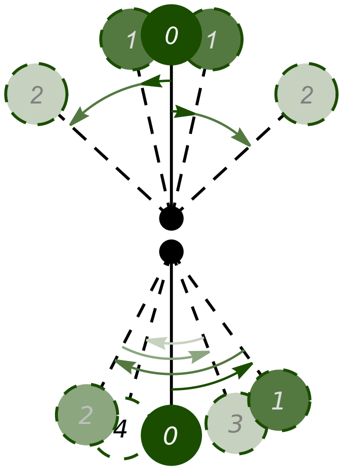

Figure 5.1: A pendulum (in green) hung by a rigid rod from a pivot (in black) perturbed from its resting state (bottom) and from its inverted, upright position (top). Perturbed states (with dashed boundaries and lines) are faded out to indicate the progression of time.

The classical way of distinguishing the aforementioned stability situations is by designing a *Lyapunov function* [98], i.e., an energy-like auxiliary measure satisfying certain *arithmetical conditions* [71, 89, 165] which implies that the auxiliary energy decreases along system trajectories towards local minima at the stable resting state(s), see Fig. 5.2. Prior approaches [3, 61, 88, 104, 170] have emphasized the need to formally verify those arithmetical conditions in order to guarantee that a conjectured Lyapunov function correctly implies stability for a given system.

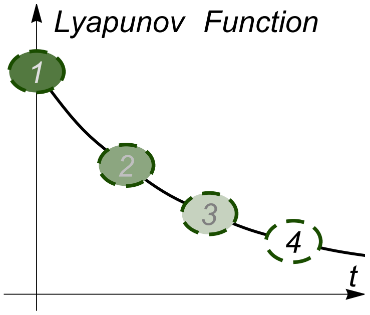

Figure 5.2: A Lyapunov function that decreases along the pendulum trajectory shown in Fig. 5.1 (bottom).

This chapter shows how deductive proofs of ODE stability can be carried out in dL. The key insight is to specify stability properties by suitably nesting the dynamic modalities of dL with quantifiers of first-order logic. This makes it possible to *syntactically derive* stability for a given system by combining the ODE safety and liveness proof principles of Chapters 3 and 4 with arithmetic and first-order quantifier reasoning in dL. This combination enables trustworthy implementation of stability proofs in KeYmaera X [54, 142]. Notably, the approach directly verifies *stability specifications*, which goes beyond verifying arithmetic that imply those specifications [3, 61, 88, 104, 170]. This is crucial for advanced stability notions because those variations generally require subtle twists to the required arithmetical conditions on their Lyapunov functions [71]. Proofs of stability specifications alleviate the onus on system designers to correctly pick and check the appropriate conditions for their applications.

Section 5.2 shows how various stability properties for ODE equilibria can be formally specified and proved in dL with Lyapunov function techniques. Section 5.3 then generalizes those stability specifications, yielding unambiguous formal specifications of advanced stability properties from the literature [71, 89], along with their *derived* proof rules. These specifications also provide rigorous insights into the logical relationship between various stability notions, which are used to inform their respective proofs. Section 5.4 illustrates the practicality of the dL approach through several stability case studies formalized in KeYmaera X. Section 5.5 examines *input-to-state* stability, a form of stability with respect to perturbation of the system dynamics [71, 89], and discusses the syntactic limitations of dL for specifying and analyzing this form of stability.

Reminder (Extended Term Language). This chapter uses an extended dL term language following the extended term conditions and notational conventions of Section 3.2 because the dL axiomatization remains sound for all extended term languages meeting those conditions.

Contribution. *The material for this chapter is drawn from Tan and Platzer [194].*

5.2 Asymptotic Stability of an Equilibrium Point

This section presents Lyapunov's classical notion of asymptotic stability [98] and its formal specification in dL. This formalization enables the derivation of dL stability proof rules using *Lyapunov functions* [71, 89, 98, 165]. Several related stability concepts are formalized in dL, along with their relationships and proof rules. The following parametric ODE model of a simple pendulum is used as a running example.

Example 5.1 (Pendulum model). The ODE $\alpha_p \equiv \theta' = \omega,\ \omega' = -\frac{g}{L}\sin(\theta) - b\omega$ models a pendulum (illustrated below, right) suspended from a pivot by a rod of length L, where θ is the angle of displacement, ω is the angular velocity of the pendulum, and $g > 0$ is the gravitational constant. Parameter $a = \frac{g}{L}$ is a positive scaling constant and parameter $b \geq 0$ is the coefficient of friction for angular velocity. The symbolic parameters a, b make analysis of α_p apply to a range of concrete values, e.g., pendulums that are suspended by a long rod (with large L) are modeled by small positive values of a, while frictionless pendulums have $b = 0$.

For illustrative purposes, a simplification of α_p is used because stability analyses often concern the behavior of the pendulum near its resting (or inverted) state where $\theta = 0$. For such nearby states with $\theta \approx 0$, the small angle approximation $\sin(\theta) \approx \theta$ yields a linear ODE α_l.[1]

$$\alpha_l \equiv \theta' = \omega,\ \omega' = -a\theta - b\omega \tag{5.1}$$

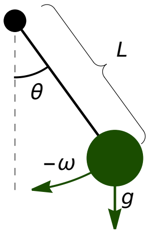

An *inverted* pendulum is modeled by a similar ODE (illustrated on the right) under a change of coordinates. Such a pendulum requires an external torque input $u(\theta, \omega)$ to maintain its stability. An appropriate input $u(\theta, \omega)$ is determined and proved correct in Section 5.4.

$$\alpha_i \equiv \theta' = \omega,\ \omega' = a\theta - b\omega - u(\theta, \omega) \tag{5.2}$$

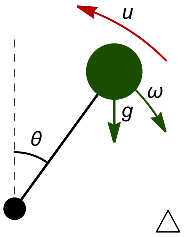

△

[1] Mathematically, this linearization is justified by the Hartman-Grobman theorem [33]. A nonlinear polynomial approximation, such as $\sin(\theta) \approx \theta - \frac{\theta^3}{6}$, can also be used.

5.2.1 Mathematical Preliminaries

An *equilibrium point* of ODE $x' = f(x)$ is a point $x_0 \in \mathbb{R}^n$ where $f(x_0) = 0$, so a system that starts at x_0 stays at x_0 along its continuous evolution. Such points are often interesting in real-world systems, e.g., the equilibrium point $\theta = 0, \omega = 0$ for α_l from (5.1) is the resting state of a pendulum. For a controlled system, equilibrium points often correspond to desired steady system states where no further continuous control input (modeled as part of $f(x)$) is required [89]. For brevity, assume the origin $0 \in \mathbb{R}^n$ is an equilibrium point of interest. Any other equilibrium point(s) of interest $x_0 \in \mathbb{R}^n$ can be translated to the origin with the change of coordinates $x \mapsto x - x_0$ for the ODE, see Lemma C.1. The following definition of asymptotic stability is standard [71, 89, 165], where the Euclidean norm $\|\cdot\|_2$ is used throughout without loss of generality because norms are equivalent on finite dimensional vector spaces [204, §5.V].[2]

Definition 5.2 (Asymptotic stability [71, 89, 165])**.** The origin $0 \in \mathbb{R}^n$ of ODE $x' = f(x)$ is

- **stable** if, for all $\varepsilon > 0$, there exists $\delta > 0$ such that for all initial states $x = x(0)$ with $\|x\|_2 < \delta$, the right-maximal ODE solution $x(t) : [0, T) \to \mathbb{R}^n$ satisfies $\|x(t)\|_2 < \varepsilon$ for all times $0 \leq t < T$,
- **attractive** if there exists $\delta > 0$ such that for all $x = x(0)$ with $\|x\|_2 < \delta$, the right-maximal ODE solution $x(t) : [0, T) \to \mathbb{R}^n$ satisfies $\lim_{t \to T} x(t) = 0$, and
- **asymptotically stable** if it is stable and attractive.

These definitions can be understood using the resting state of the pendulum from Fig. 5.1 (bottom) which is asymptotically stable. When the pendulum is given a light push from its bottom resting state (formally, $\|x\|_2 < \delta$), it gently oscillates near that resting state (formally, $\|x(t)\|_2 < \varepsilon$). In the presence of friction, these oscillations eventually dissipate so the pendulum asymptotically returns to its resting state (formally, $\lim_{t \to T} x(t) = 0$). This behavior is *local*, i.e., for any given $\varepsilon > 0$, there *exists* a sufficiently small $\delta > 0$ perturbation of the initial state that results in gentle oscillations with $\|x(t)\|_2 < \varepsilon$, see Fig. 5.3 (left). A strong push, e.g., with $\delta > \varepsilon$, could instead cause the pendulum to spin around on its pivot.

Remark 5.3. Stability and attractivity *do not* imply each other [165, Chapter I.2.7], see example α_u in Section 5.4. However, if the origin is stable, attractivity can be equivalently defined in a simpler way. This is proved in dL, after characterizing stability and attractivity syntactically.

5.2.2 Formal Specification

The formal specification of asymptotic stability in dL combines *i)* the dynamic modalities of dL, which are used to quantify over the dynamics of the ODE and *ii)* the first-order logic quantifiers, which are used to express combinations of (topologically) local and asymptotic properties of those dynamics. For a formula P, the ε-neighborhood of P with respect to x is defined as the

[2] Some definitions of asymptotic stability in the literature require, or implicitly assume, right-maximal solutions $x(t)$ to be global, i.e., with $T = \infty$, see [89, Definition 4.1] and associated discussion. The definition given here is better suited for subsequent generalizations.

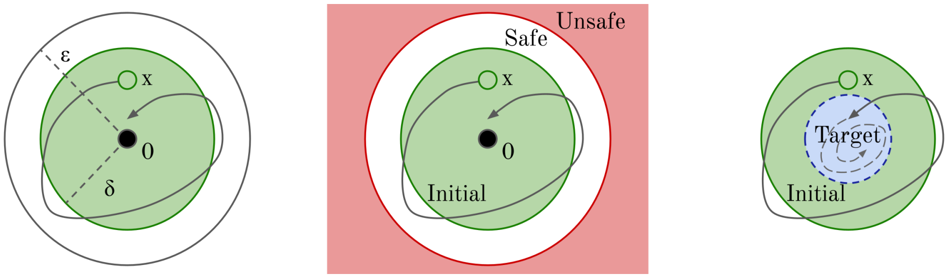

Figure 5.3: Solutions from points in the δ ball around the origin, like the green initial point x, remain within the ε ball around the origin $0 \in \mathbb{R}^n$ (black dot) and asymptotically approach the origin. The latter two plots illustrate how asymptotic stability for an ODE can be broken down into a pair of (quantified) ODE safety and liveness properties.

formula $\mathcal{U}_\varepsilon(P) \overset{\text{def}}{\equiv} \exists y \left(\|x-y\|_2^2 < \varepsilon^2 \land P(y)\right)$, where the existentially quantified variables y are fresh in P. The neighborhood formula $\mathcal{U}_\varepsilon(P)$ characterizes the set of states within (set) distance ε from P, with respect to the dynamically evolving variables x, which is useful for syntactically expressing small ε perturbations, e.g., appearing in Def. 5.2. For formulas P of first-order real arithmetic (over polynomial terms), the ε-neighborhood, $\mathcal{U}_\varepsilon(P)$, can be equivalently expressed in quantifier-free form by quantifier elimination [14, 197]. For example, the neighborhood formula $\mathcal{U}_\varepsilon(x=0)$ is equivalent to the formula $\|x\|_2^2 < \varepsilon^2$.

Lemma 5.4 (Asymptotic stability in dL). *The origin of ODE $x' = f(x)$ is, respectively, i)* ***stable****, ii)* ***attractive****, and iii)* ***asymptotically stable*** *iff the* dL *formulas i)* $\mathrm{Stab}(x' = f(x))$*, ii)* $\mathrm{Attr}(x' = f(x))$*, and iii)* $\mathrm{AStab}(x' = f(x))$ *, respectively, are valid. Variables ε, δ are fresh, i.e., not in $x, f(x)$.*

$$\begin{aligned}
\mathrm{Stab}(x' = f(x)) &\equiv \forall \varepsilon{>}0\, \exists \delta{>}0\, \forall x \left(\mathcal{U}_\delta(x=0) \rightarrow [x' = f(x)]\,\mathcal{U}_\varepsilon(x=0)\right)\\
\mathrm{Attr}(x' = f(x)) &\equiv \exists \delta{>}0\, \forall x \left(\mathcal{U}_\delta(x=0) \rightarrow \mathrm{Asym}(x' = f(x), x=0)\right)\\
\mathrm{AStab}(x' = f(x)) &\equiv \mathrm{Stab}(x' = f(x)) \land \mathrm{Attr}(x' = f(x))
\end{aligned}$$

Formula $\mathrm{Asym}(x' = f(x), P) \equiv \forall \varepsilon{>}0\, \langle x' = f(x)\rangle [x' = f(x)]\,\mathcal{U}_\varepsilon(P)$ *characterizes the set of states that asymptotically approach P along ODE solutions.*

Proof. The correctness of these specifications follows directly from the semantics of dL formulas [142, 144] because they syntactically express the logical connectives and quantifiers from Def. 5.2 in dL. The open neighborhood formulas $\mathcal{U}_\delta(x=0)$ and $\mathcal{U}_\varepsilon(x=0)$ are true in states where $\|x\|_2 < \delta$ and $\|x\|_2 < \varepsilon$ respectively. The main subtlety is formula $\mathrm{Attr}(x' = f(x))$ which characterizes the limit $\lim_{t\to T} x(t) = 0$ using its subformula $\mathrm{Asym}(x' = f(x), x=0)$ as follows. Unfolding the semantics, formula $\mathrm{Asym}(x' = f(x), P)$ is true in an initial state iff for any $\varepsilon > 0$, the right-maximal ODE solution to $x' = f(x)$ (restricted to variables x) denoted $x(t) : [0,T) \to \mathbb{R}^n$ has a time $\tau \in [0,T)$ where, because of uniqueness of ODE solutions [33, Theorem 1.2], for all future times t with $\tau \le t < T$, the solution at $x(t)$ satisfies formula $\mathcal{U}_\varepsilon(P)$. For $P \equiv x = 0$, this implies the bound $\|x(t)\|_2 < \varepsilon$ at those future times, which is the real analytic definition of the limit $\lim_{t\to T} x(t) = 0$ [168, Definition 4.1]. □

Formula $\mathrm{Stab}(x' = f(x))$ is a syntactic dL rendering of the corresponding quantifiers from Def. 5.2. The safety property $\mathcal{U}_\delta(x=0) \rightarrow [x' = f(x)]\,\mathcal{U}_\varepsilon(x=0)$ expresses that solutions starting from the δ-neighborhood of the origin always (for all times) stay safely in the ε-neighborhood, as visualized in Fig. 5.3 (middle). Formula $\mathrm{Attr}(x' = f(x))$ uses the subformula $\mathrm{Asym}(x' = f(x), x = 0)$ which characterizes the limit in Def. 5.2. Recall $\lim_{t \to T} x(t) = 0$ iff for all $\varepsilon > 0$ there exists a time τ with $0 \leq \tau < T$ such that for all times t with $\tau \leq t < T$, the solution satisfies $\|x(t)\|_2 < \varepsilon$, i.e., the limit requires for all distances $\varepsilon > 0$, the ODE solution will *eventually always* be within distance ε of the origin, as visualized in Fig. 5.3 (right). This limit is characterized using nested $\langle\cdot\rangle[\cdot]$ modalities, together with first-order quantification according to Def. 5.2. More generally, formula $\mathrm{Asym}(x' = f(x), P)$ characterizes the set of initial states where the right-maximal ODE solution asymptotically approaches the set characterized by formula P; this set is known as the *region of attraction* of P [89]. Thus, attractivity requires that the region of attraction of the origin contains an open neighborhood $\mathcal{U}_\delta(x = 0)$ of the origin.

Proving validity of the formula $\mathrm{AStab}(x' = f(x))$ yields a rigorous proof of asymptotic stability for $x' = f(x)$. Indeed, the syntactic shape of the formulas from Lemma 5.4 immediately suggests how such a proof can be carried out using earlier thesis chapters: $\mathrm{Stab}(x' = f(x))$ needs ODE safety reasoning (Chapter 3) for its inner box modality, while, at a first glance, $\mathrm{Attr}(x' = f(x))$ needs more complicated ODE liveness reasoning (Chapter 4) for the inner, nested diamond-box $\langle\cdot\rangle[\cdot]$ modalities. However, if the origin is stable, then Corollary 5.5 below simplifies the syntactic characterization of the region of attraction for the stable equilibrium from a nested $\langle\cdot\rangle[\cdot]$ formula to a $\langle\cdot\rangle$ formula, which is then directly amenable to liveness reasoning (Chapter 4). This corollary is used to simplify proofs of asymptotic stability in the next section.

Corollary 5.5 (Stable attractivity)**.** *The following axiom is derivable in* dL.

$$\mathrm{SAttr}\ \ \mathrm{Stab}(x' = f(x)) \to \big(\mathrm{Asym}(x' = f(x), x = 0) \leftrightarrow \forall \varepsilon{>}0\, \langle x' = f(x) \rangle\, \mathcal{U}_\varepsilon(x = 0)\big)$$

Proof. A full proof is omitted as axiom SAttr is an instance of the more general axiom SetSAttr derived in Corollary 5.21 (where $P \equiv x = 0$). Briefly, the "$\rightarrow$" direction of the inner equivalence is valid even without assuming stability because postcondition $[x' = f(x)]\,\mathcal{U}_\varepsilon(x = 0)$ monotonically implies postcondition $\mathcal{U}_\varepsilon(x = 0)$. The more interesting "$\leftarrow$" direction of the inner equivalence uses the stability assumption by choosing $\delta > 0$ sufficiently small so that solutions reaching $\mathcal{U}_\delta(x = 0)$ must stay in $\mathcal{U}_\varepsilon(x = 0)$ thereafter because of stability. □

Notational Conventions (Abbreviations). All derived axioms and proof rules are presented directly using the respective stability formula abbreviations, e.g., as listed in Lemma 5.4.

5.2.3 Lyapunov Functions

Lyapunov functions are the standard tool for showing stability of general, non-linear ODEs [71, 89, 165] and finding suitable Lyapunov functions is an important problem in its own right [3, 49, 61, 88, 104, 130, 131, 170, 200]. This section shows how a candidate Lyapunov function, once found, can be used to rigorously prove stability. The following derived proof rules formalize Lyapunov stability arguments [71, 89, 165] syntactically in dL.

Lemma 5.6 (Lyapunov functions)**.** *The following Lyapunov function proof rules are derivable in* dL*, where the Lyapunov function V is a* dL *term.*

$$\text{Lyap}_{\geq}\ \dfrac{\vdash f(0) = 0 \land V(0) = 0 \qquad \vdash \exists \gamma{>}0\, \forall x\, \big(0 < \|x\|_2^2 \leq \gamma^2 \rightarrow V > 0 \land \dot{V} \leq 0\big)}{\vdash \text{Stab}(x' = f(x))}$$

$$\text{Lyap}_{>}\ \dfrac{\vdash f(0) = 0 \land V(0) = 0 \qquad \vdash \exists \gamma{>}0\, \forall x\, \big(0 < \|x\|_2^2 \leq \gamma^2 \rightarrow V > 0 \land \dot{V} < 0\big)}{\vdash \text{AStab}(x' = f(x))}$$

Proof Summary (Appendix C.1.1). Rule Lyap$_{\geq}$ is derived by showing that, for carefully chosen (symbolic) constants $0 < \gamma \leq \varepsilon$ and $W > 0$, the formula $\|x\|_2^2 < \gamma^2 \land V < W$ is an invariant of the ODE $x' = f(x)$. Here, γ is chosen sufficiently small so that the left conjunct $\|x\|_2^2 < \gamma^2$ implies $\|x\|_2^2 < \varepsilon^2$ as required in the stability postcondition, while the right conjunct $V < W$ characterizes a smaller invariant set in the ball $\|x\|_2^2 < \gamma^2$ (see proof).

The derivation of rule Lyap$_{>}$ uses Lyap$_{\geq}$ as a stepping stone and the derived logical relationship SAttr to simplify the proof. The derivation starts with a cut that proves stability using rule Lyap$_{\geq}$ because the premises of Lyap$_{>}$ are identical to those of Lyap$_{\geq}$ except for a stronger, strict inequality requirement on the Lie derivative of V. The right conjunct of the right premise of Lyap$_{>}$ is cut and Skolemized with $\exists$L; the resulting antecedent is abbreviated with $\text{ⓐ} \equiv \forall x\, \big(0 < \|x\|_2^2 \leq \gamma^2 \rightarrow \dot{V} < 0\big)$ below. Next, instantiating $\varepsilon = \gamma$ in the stability antecedent with $\forall$L and Skolemizing yields an initial disturbance $\delta > 0$ so that the ODE solutions from states satisfying $\|x\|_2^2 < \delta^2$ always stay within the γ ball $\|x\|_2^2 < \gamma^2$. The resulting antecedent is abbreviated with $\text{ⓑ} \equiv \forall x\, \big(\, \|x\|_2^2 < \delta^2 \rightarrow [x' = f(x)]\ \|x\|_2^2 < \gamma^2\big)$.

$$\begin{array}{rr}
 & \text{Stab}(x' = f(x)), \text{ⓐ}, \delta > 0, \text{ⓑ} \vdash \text{Attr}(x' = f(x)) \\ \cline{2-2}
\scriptstyle \forall\text{L}, \exists\text{L} & \text{Stab}(x' = f(x)), \gamma > 0, \text{ⓐ} \vdash \text{Attr}(x' = f(x)) \\ \cline{2-2}
\scriptstyle \text{cut}, \exists\text{L} & \text{Stab}(x' = f(x)) \vdash \text{Attr}(x' = f(x)) \\ \cline{2-2}
\scriptstyle \text{cut}, \text{Lyap}_{\geq} & \vdash \text{AStab}(x' = f(x))
\end{array}$$

The existential quantifier in the succedent is witnessed by δ from the antecedents with $\exists$R and the resulting sequent is simplified by Skolemization and instantiation of ⓑ before axiom SAttr is used to further simplify the succedent using the stability antecedent.

$$\begin{array}{rc}
 & \text{ⓐ}, [x'{=}f(x)]\ \|x\|_2^2 < \gamma^2, \varepsilon{>}0 \vdash \langle x'{=}f(x)\rangle\, \|x\|_2^2 < \varepsilon^2 \\ \hline
\scriptstyle \forall\text{R} & \text{ⓐ}, [x'{=}f(x)]\ \|x\|_2^2 < \gamma^2 \vdash \forall \varepsilon{>}0\, \langle x'{=}f(x)\rangle\, \|x\|_2^2 < \varepsilon^2 \\ \hline
\scriptstyle \text{SAttr} & \text{Stab}(x'{=}f(x)), \text{ⓐ}, [x'{=}f(x)]\ \|x\|_2^2 < \gamma^2 \vdash \text{Asym}(x'{=}f(x), x{=}0) \\ \hline
\scriptstyle \forall\text{L}, \rightarrow\text{L} & \text{Stab}(x'{=}f(x)), \text{ⓐ}, \text{ⓑ}, \|x\|_2^2 < \delta^2 \vdash \text{Asym}(x'{=}f(x), x{=}0) \\ \hline
\scriptstyle \forall\text{R}, \rightarrow\text{R} & \text{Stab}(x'{=}f(x)), \text{ⓐ}, \text{ⓑ} \vdash \forall x\, \big(\, \|x\|_2^2 < \delta^2 \rightarrow \text{Asym}(x'{=}f(x), x{=}0)\big) \\ \hline
\scriptstyle \exists\text{R} & \text{Stab}(x'{=}f(x)), \text{ⓐ}, \delta > 0, \text{ⓑ} \vdash \text{Attr}(x'{=}f(x))
\end{array}$$

The remaining open premise is a liveness property which is proved using rule SP$_c$ from Corollary 4.21 (page 97) with the choice of compact staging set $S \equiv \varepsilon^2 \leq \|x\|_2^2 \leq \gamma^2$ and $e \equiv V$.

$$\text{SP}_c\ \dfrac{\text{dC}\ \dfrac{\text{dW},\mathbb{R}\ \dfrac{*}{\vdash [x' = f(x)\,\&\, \|x\|_2^2 \geq \varepsilon^2 \land \|x\|_2^2 {<} \gamma^2]S}}{[x' = f(x)]\ \|x\|_2^2 {<} \gamma^2 \vdash [x' = f(x)\,\&\, \|x\|_2^2 \geq \varepsilon^2]S} \qquad \dfrac{*}{\text{ⓐ}, \varepsilon{>}0, S \vdash \dot{V} < 0}}{\text{ⓐ}, [x' = f(x)]\ \|x\|_2^2 {<} \gamma^2, \varepsilon{>}0 \vdash \langle x' = f(x)\rangle\, \|x\|_2^2 {<} \varepsilon^2}$$

The left premise proves with a differential cut dC of the antecedent and dW, $\mathbb{R}$ from the resulting strengthened domain constraint $\|x\|_2^2 \geq \varepsilon^2 \land \|x\|_2^2 {<} \gamma^2$. The right premise proves using ⓐ with antecedents S and $\varepsilon{>}0$ to prove its implication LHS. □

Rule Lyap$_{\geq}$ uses the Lyapunov function V as an auxiliary, energy-like function near the origin which has non-positive derivative $\dot{V} \leq 0$. Intuitively, this guarantees that the energy of the system is *non-increasing* along ODE solutions near the origin so those solutions (with sufficiently low energy) must stay close to the stable origin (energy $V(0) = 0$). Rule Lyap$_{>}$ is similar, except the derivative is strictly negative $\dot{V} < 0$ so the energy *decreases* along system trajectories towards 0 at the asymptotically stable origin, see Fig. 5.2. The right premise of both rules use $\exists \gamma{>}0\, \forall x\, \big(0{<}\, \|x\|_2^2 \,{\leq}\gamma^2 \rightarrow \cdots\big)$ to require that the Lyapunov function conditions are true in a γ-neighborhood of the origin. The subtle difference in sign condition for $\dot{V}$ between rules Lyap$_{\geq}$, Lyap$_{>}$ is illustrated for the pendulum in the following example.

Example 5.7 (Pendulum asymptotic stability)**.** For ODE α_l from (5.1), a suitable Lyapunov function for proving its stability [89] is $V = a\frac{\theta^2}{2} + \frac{(b\theta+\omega)^2+\omega^2}{4}$, where the Lie derivative of V along α_l is $\dot{V} = -\frac{b}{2}(a\theta^2 + \omega^2)$. Stability[3] is formally proved in dL for *any* parameter values $a > 0, b \geq 0$ using rule Lyap$_{\geq}$ because both of its resulting arithmetical premises are provable by $\mathbb{R}$. The derivation is shown below, where the left premise resulting from rule Lyap$_{\geq}$ is omitted as it proves trivially by evaluation. The existentially quantified right premise proves by $\mathbb{R}$ since V is positive except at the origin and its Lie derivative is non-positive.

$$\text{Lyap}_{\geq}\dfrac{\mathbb{R}\,\dfrac{*}{a > 0, b \geq 0 \vdash \exists \tau{>}0\, \forall\theta\, \forall\omega\, \big(0 < \theta^2 + \omega^2 \leq \tau^2 \rightarrow V > 0 \land -\frac{b}{2}(a\theta^2 + \omega^2) \leq 0\big)}}{a > 0, b \geq 0 \vdash \mathrm{Stab}(\alpha_l)}$$

When $b > 0$, i.e., friction is non-negligible, an identical derivation with Lyap$_{>}$ instead of Lyap$_{\geq}$ proves asymptotic stability because $-\frac{b}{2}(a\theta^2 + \omega^2)$ is negative except at the origin. Indeed, displacements to the pendulum's resting state can only be dissipated in the presence of friction and not when $b = 0$. △

5.2.4 Asymptotic Stability Variations

Asymptotic stability is a strong guarantee about the local behavior of ODE solutions near equilibrium points of interest. In certain applications, stronger stability guarantees may be needed for those equilibria [89]. This section examines two standard stability variations, shows how they can be proved in dL, and analyzes their logical relationship with asymptotic stability.

Exponential Stability

As its name suggests, the first stability variation, exponential stability, guarantees an exponential rate of convergence towards the equilibrium point from an initial displacement.

[3] Recall α_l is a linearization of the trigonometric pendulum ODE model α_p from Example 5.1 for simplicity. The Lyapunov function $V = a(1 - \cos(\theta)) + \frac{(b\theta+\omega)^2+\omega^2}{4}$ with Lie derivative $\dot{V} = -\frac{b}{2}(a\theta\sin(\theta) + \omega^2)$ proves stability for α_p [89] but requires arithmetic reasoning over trigonometric functions.

Definition 5.8 (Exponential stability [71, 89, 165]). The origin $0 \in \mathbb{R}^n$ of ODE $x' = f(x)$ is **exponentially stable** if there are positive constants $\alpha, \beta, \delta > 0$ such that for all initial states $x = x(0)$ with $\|x\|_2 < \delta$, the right-maximal ODE solution $x(t) : [0, T) \to \mathbb{R}^n$ satisfies $\|x(t)\|_2 \leq \alpha \|x(0)\|_2 \exp(-\beta t)$ for all times $0 \leq t < T$.

Exponential stability bounds the norm of solutions to ODE $x' = f(x)$ near the origin by a decaying exponential. It is specified in dL as follows.

Lemma 5.9 (Exponential stability in dL). *The origin of ODE $x' = f(x)$ is **exponentially stable** iff the following* dL *formula is valid. Variables α, β, δ, y are fresh, i.e., not in $x, f(x)$.*

$$\begin{aligned}\mathrm{EStab}(x' = f(x)) \equiv \exists \alpha{>}0\, \exists \beta{>}0\, \exists \delta{>}0\, \forall x\, \big(&\mathcal{U}_\delta(x = 0) \rightarrow \\ &[y := \alpha^2 \|x\|_2^2\, ; x' = f(x), y' = -2\beta y]\ \|x\|_2^2 \leq y\big)\end{aligned}$$

Proof. The quantifiers $\exists \alpha{>}0\, \exists \beta{>}0\, \exists \delta{>}0\, \forall x\, \big(\mathcal{U}_\delta(x = 0) \rightarrow \cdots\big)$ syntactically express the respective quantifiers in the definition of exponential stability from Def. 5.8 in dL. For an initial state satisfying $\mathcal{U}_\delta(x = 0)$, i.e., with $\|x(0)\|_2 < \delta$, the assignment $y := \alpha^2 \|x\|_2^2$ sets the initial value of fresh variable y (before the ODE) to $\alpha^2 \|x(0)\|_2^2$. Let $x(t) : [0, T) \to \mathbb{R}^n$ and $y(t) : [0, T) \to \mathbb{R}$ respectively be the x and y projections of the unique, right-maximal solution of the ODE $x' = f(x), y' = -2\beta y$. By construction, the unique ODE solution for the y-coordinates is $y(t) = \alpha^2 \|x(0)\|_2^2 \exp(-2\beta t)$, so the postcondition $\|x\|_2^2 \leq y$ of the box modality expresses that for all times $0 \leq t < T$, $\|x(t)\|_2^2 \leq \alpha^2 \|x(0)\|_2^2 \exp(-2\beta t)$ or, equivalently, $\|x(t)\|_2 \leq \alpha \|x(0)\|_2 \exp(-\beta t)$, as required. □

Formula $\mathrm{EStab}(x' = f(x))$ uses a fresh ghost variable y with ODE $y' = -2\beta y$ and initialized to $\alpha^2 \|x\|_2^2$ so that y differentially axiomatizes (Proposition 3.34) the *squared* decaying exponential function $\alpha^2 \|x(0)\|_2^2 \exp(-2\beta t)$ along ODE solutions which bounds the *squared* norm term $\|x\|_2^2$. An alternative (explicit) specification of exponential stability can also be given in an extended term language with the exponential function exp, using fresh variables y to syntactically store the initial norm value and t to syntactically track the progression of time with $t' = 1$:

$$\begin{aligned}\mathrm{EStabE}(x' = f(x)) \equiv \exists \alpha{>}0\, \exists \beta{>}0\, \exists \delta{>}0\, \forall x\, \big(&\mathcal{U}_\delta(x = 0) \rightarrow \\ &[y := \alpha^2 \|x\|_2^2\, ; t := 0; x' = f(x), t' = 1]\ \|x\|_2^2 \leq y \exp(-2\beta t)\big)\end{aligned}$$

Corollary 5.10 (Exponential stability characterizations). *The following axiom is derivable in* dL*, where variables y, t are fresh in ODE $x' = f(x)$.*

$$\text{EStabE}\quad \begin{aligned}&[y := \alpha^2 \|x\|_2^2\, ; t := 0; x' = f(x), t' = 1]\ \|x\|_2^2 \leq y \exp(-2\beta t) \\ &\leftrightarrow [y := \alpha^2 \|x\|_2^2\, ; x' = f(x), y' = -2\beta y]\ \|x\|_2^2 \leq y\end{aligned}$$

Proof in Appendix C.1.1.

The explicit characterization $\mathrm{EStabE}(x' = f(x))$ corresponds more directly to Def. 5.8 but needs explicit reasoning over the exponential function in the postcondition of the box modality. In contrast, the implicit polynomial characterization of exponential decay given in $\mathrm{EStab}(x' = f(x))$ allows syntactic proof steps to use decidable real arithmetic reasoning [14, 197] if the provided Lyapunov functions V are also polynomial terms. Corollary 5.10 shows that the two characterizations can be used interchangeably in proofs.

Lemma 5.11 (Lyapunov function for exponential stability). *The following Lyapunov function proof rule for exponential stability is derivable in* dL*, where $k_1, k_2, k_3 \in \mathbb{Q}$ are positive constants.*

$$\text{Lyap}_{\text{E}} \; \frac{\vdash \exists \gamma{>}0 \, \forall x \left(\|x\|_2^2 \leq \gamma^2 \rightarrow k_1^2 \|x\|_2^2 \leq V \leq k_2^2 \|x\|_2^2 \wedge \dot{V} \leq -2k_3 V \right)}{\vdash \text{EStab}(x' = f(x))}$$

Proof in Appendix C.1.1.

Rule Lyap_{E} enables proofs of exponential stability in dL where the Lyapunov function derivative condition $\dot{V} \leq -2k_3 V$ guarantees that the (auxiliary) system energy V is bounded above by the decaying exponential $\exp(-2k_3 t)$ as required in $\text{EStab}(x' = f(x))$ (with $\beta = k_3$); the factor 2 arises from the use of the squared norm in the specification. In fact, the proof of Lemma 5.11 yields *quantitative* bounds, where $\text{EStab}(x' = f(x))$ is explicitly witnessed with scaling constant $\alpha = \frac{k_2}{k_1}$ and decay rate $\beta = k_3$. These can be used to calculate time bounds when the system state will return sufficiently close to the origin. Similarly, the disturbance δ in $\text{EStab}(x' = f(x))$ is quantitatively witnessed by $\frac{k_1}{k_2}\gamma$ for any γ witnessing validity of the premise of rule Lyap_{E}. This yields a provable estimate of the region around the origin where exponential stability holds; this latter estimate is explored next.

Region of Attraction

Formulas $\text{Attr}(x' = f(x))$ and $\text{EStab}(x' = f(x))$ both feature a subformula of the form $\exists \delta{>}0 \, \forall x \, (\mathcal{U}_\delta(x = 0) \rightarrow \cdots)$ which expresses that attractivity (or exponential stability) is locally true in *some* δ neighborhood of the origin. In many applications, it is useful to find and rigorously prove that a given set is attractive or exponentially stable with respect to the origin [89, Chapter 8.2]. The second stability variation yields *provable* subsets of the region of attraction, including the special case where it is the entire state space. This is formalized using the following variants of $\text{Attr}(x' = f(x))$ and $\text{EStab}(x' = f(x))$ within a region characterized by the formula P.

$$\begin{aligned}
\text{Attr}^{\text{P}}(x' = f(x), P) &\equiv \forall x \left(P \rightarrow \text{Asym}(x' = f(x), x = 0) \right) \\
\text{EStab}^{\text{P}}(x' = f(x), P) &\equiv \exists \alpha{>}0 \, \exists \beta{>}0 \, \forall x \, \big(P \rightarrow \\
&\qquad [y := \alpha^2 \|x\|_2^2 \, ; x' = f(x), y' = -2\beta y] \; \|x\|_2^2 \leq y \big)
\end{aligned}$$

The formula $\text{Attr}^{\text{P}}(x' = f(x), P)$ is valid iff the set characterized by P is a subset of the origin's region of attraction [89]. For example, $\text{Attr}(x' = f(x))$ is equivalent to an existentially quantified region of attraction $\exists \delta > 0 \; \text{Attr}^{\text{P}}(x' = f(x), \mathcal{U}_\delta(x = 0))$. The flexibility afforded by the choice of P is useful for formalizing stronger notions of stability in dL, such as the following *global* stability notions [71, 89].

Definition 5.12 (Global stability [71, 89, 165]). The origin $0 \in \mathbb{R}^n$ of ODE $x' = f(x)$ is **globally asymptotically stable** if it is stable and its region of attraction is the entire state space, i.e., for all $x = x(0) \in \mathbb{R}^n$, the right-maximal ODE solution $x(t) : [0, T) \rightarrow \mathbb{R}^n$ satisfies $\lim_{t \rightarrow T} x(t) = 0$. The origin is **globally exponentially stable** if there are positive constants $\alpha, \beta > 0$ such that for all initial states $x = x(0) \in \mathbb{R}^n$, the right-maximal ODE solution $x(t) : [0, T) \rightarrow \mathbb{R}^n$ satisfies $\|x(t)\|_2 \leq \alpha \|x(0)\|_2 \exp(-\beta t)$ for all times $0 \leq t < T$.

Lemma 5.13 (Global stability in dL)**.** *The origin of ODE $x' = f(x)$ is **globally asymptotically stable** iff the* dL *formula* $\mathrm{Stab}(x' = f(x)) \land \mathrm{Attr}^{\mathrm{P}}(x' = f(x), \mathit{true})$ *is valid. The origin is **globally exponentially stable** iff the* dL *formula* $\mathrm{EStab}^{\mathrm{P}}(x' = f(x), \mathit{true})$ *is valid.*

Proof. The proof is identical to Lemmas 5.4 and 5.9, respectively, except the existential quantification over a local neighborhood of the origin $\exists \delta > 0\, \forall x\, (\mathcal{U}_\delta(x = 0) \to \cdots)$ is replaced by universal quantification over all initial states (where $P \equiv \mathit{true}$), i.e., $\forall x\, (\mathit{true} \to \cdots)$, as required by the global stability definitions. □

Global stability ensures that *all* perturbations to the system state are eventually dissipated. Their proof rules are similar to $\mathrm{Lyap}_>$ and $\mathrm{Lyap}_{\mathrm{E}}$ respectively.

Lemma 5.14 (Lyapunov function for global stability)**.** *The following Lyapunov function proof rules for global asymptotic and exponential stability are derivable in* dL*. In rule* $\mathrm{Lyap}^{\mathrm{G}}_{\mathrm{E}}$*,* $k_1, k_2, k_3 \in \mathbb{Q}$ *are positive constants.*

$$\mathrm{Lyap}^{\mathrm{G}}_{>}\ \frac{\vdash f(0) = 0 \land V(0) = 0 \quad x{\neq}0 \vdash V{>}0 \land \overset{\bullet}{V}{<}0 \quad \vdash \forall b\, \exists \gamma{>}0\, \forall x\, \big(V{\leq}b \to \mathcal{U}_\gamma(x = 0)\big)}{\vdash \mathrm{Stab}(x' = f(x)) \land \mathrm{Attr}^{\mathrm{P}}(x' = f(x), \mathit{true})}$$

$$\mathrm{Lyap}^{\mathrm{G}}_{\mathrm{E}}\ \frac{\vdash k_1^2\, \|x\|_2^2 \leq V \leq k_2^2\, \|x\|_2^2 \land \overset{\bullet}{V} \leq -2k_3 V}{\vdash \mathrm{EStab}^{\mathrm{P}}(x' = f(x), \mathit{true})}$$

Proof in Appendix C.1.1.

Example 5.15 (Pendulum global exponential stability)**.** For simplicity, instantiate Example 5.7 with parameters $a = 1, b = 1$. The Lyapunov function then simplifies to $V = \frac{\theta^2}{2} + \frac{(\theta+\omega)^2+\omega^2}{4}$ with Lie derivative $\overset{\bullet}{V} = -\frac{(\theta^2+\omega^2)}{2}$, which satisfies the real arithmetic inequalities $\frac{\theta^2+\omega^2}{4} \leq V \leq \theta^2{+}\omega^2$ and $\overset{\bullet}{V} \leq -\frac{1}{2}V$. Thus, rule $\mathrm{Lyap}^{\mathrm{G}}_{\mathrm{E}}$ proves global exponential stability of α_l with $k_1 = \frac{1}{2}$, $k_2 = 1$, and $k_3 = \frac{1}{4}$. An important caveat is that Example 5.7 used a local small angle approximation, so this global phenomenon does *not* hold for a real-world pendulum (nor for α_p). △

Logical Relationships

With the proliferation of stability variations just introduced, it is useful to take stock of their logical relationships. An important example is shown in the following corollary.

Corollary 5.16 (Exponential stability implies asymptotic stability)**.** *The following axioms are derivable in* dL*.*

$$\mathrm{EStabStab}\quad \mathrm{EStab}(x' = f(x)) \to \mathrm{Stab}(x' = f(x))$$

$$\mathrm{EStabAttr}\quad \mathrm{EStab}^{\mathrm{P}}(x' = f(x), P) \to \mathrm{Attr}^{\mathrm{P}}(x' = f(x), P)$$

Proof in Appendix C.1.1.

Derived axioms EStabStab, EStabAttr show that exponential stability implies asymptotic stability. In proofs, EStabAttr allows the region of attraction to be estimated using the region where solutions are exponentially bounded.

5.3 General Stability

This section provides stability definitions and proof rules that generalize stability for an equilibrium point from Section 5.2 to the stability of sets. These definitions are useful when the desired stable system state(s) is not modeled by a single equilibrium point, but may instead, e.g., lie on a periodic trajectory [89], a hyperplane, or a continuum of equilibrium points within the state space [71]. The generalized definition is used to formalize two stability notions from the literature [71, 89] and to justify their Lyapunov function proof rules.

5.3.1 General Stability and General Attractivity

The following *general stability* formula defines stability in dL with respect to an ODE $x' = f(x)$ and formulas P, R. The quantified variables ε, δ are assumed to be fresh by bound renaming, i.e., do not appear in $x, f(x), P$ or R.

$$\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, R) \equiv \forall \varepsilon{>}0\, \exists \delta{>}0\, \forall x\, \big(\mathcal{U}_\delta(P) \rightarrow [x' = f(x)]\, \mathcal{U}_\varepsilon(R)\big)$$

This formula generalizes stability of the origin $\mathrm{Stab}(x' = f(x))$ by adding two logical tuning knobs that can be intuitively understood as follows. The *precondition* P characterizes the initial states from which the system state is expected to be disturbed by some disturbance δ. The *postcondition* R characterizes the set of desired operating states that the system must remain close (within the ε neighborhood of R) after being disturbed from its initial states.

The *general attractivity* formula similarly generalizes $\mathrm{Attr}^{\mathrm{P}}(x' = f(x), P)$ with a postcondition R towards which the ODE solutions from initial states satisfying precondition P are asymptotically attracted.

$$\mathrm{Attr}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, R) \equiv \forall x\, \big(P \rightarrow \mathrm{Asym}(x' = f(x), R)\big)$$

Lemma 5.17 (General Lyapunov functions)**.** *The following Lyapunov function proof rule for general stability with two stacked premises is derivable in* dL*, where formulas $\partial(\mathcal{U}_\gamma(R))$ and $\overline{\mathcal{U}_\gamma(R)}$ characterize the topological boundary and closure of the set characterized by $\mathcal{U}_\gamma(R)$, respectively (see Section 2.2.2 and Appendix B.1.3).*

$$\text{GLyap}\ \dfrac{\begin{array}{l} \vdash P \rightarrow R \\ \vdash \forall \varepsilon{>}0\, \exists 0{<}\gamma{\leq}\varepsilon\, \exists W \left(\begin{array}{l} \forall x\, (\partial(\mathcal{U}_\gamma(R)) \rightarrow V \geq W) \land \\ \exists 0{<}\delta{\leq}\gamma\, \forall x\, (\mathcal{U}_\delta(P) \rightarrow R \lor V{<}W) \land \\ \forall x\, \big(R \lor V{<}W \rightarrow [x'{=}f(x)\, \&\, \overline{\mathcal{U}_\gamma(R)}](R \lor V{<}W)\big) \end{array} \right) \end{array}}{\vdash \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x'{=}f(x), P, R)}$$

Proof Summary (Appendix C.1.2). The derivation of rule GLyap generalizes the ideas behind the derivation of rule $\mathrm{Lyap}_{\geq}$, where the second (lower) premise of the rule gives an (unsimplified) condition on the Lyapunov function V for proving general stability. □

Rule GLyap proves general stability for precondition P and postcondition R. It generalizes the Lyapunov function reasoning underlying rule $\mathrm{Lyap}_{\geq}$ to support arbitrary pre- and postconditions. The conjunct $\forall x\, (\partial(\mathcal{U}_\gamma(R)) \rightarrow V \geq W)$ requires $V \geq W$ on the boundary of $\mathcal{U}_\gamma(R)$

while the middle conjunct requires $V < W$ for some small neighborhood of P excluding R. The conjunct $\forall x\,\big(R \lor V < W \rightarrow \cdots\big)$ asserts that $R \lor V < W$ is an invariant of the ODE *within* closed domain $\overline{\mathcal{U}_\gamma(R)}$. This invariance question is provably equivalent in dL to a formula of arithmetic (Chapter 3), so the premises of rule GLyap are, *in theory*, even decidable by $\mathbb{R}$ for a given candidate (polynomial) Lyapunov function V and for semialgebraic formulas P, R (Section 3.2.1). In practice, it is prudent to consider specialized stability notions, for which the premise of rule GLyap can be arithmetically simplified. Proof rules for generalized attractivity are also derivable for specialized instances.

5.3.2 Specialization

General stability specializes to several stability notions in the literature.

Set Stability

An important special case of $\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, R)$ is when the desired operating states are exactly the states from which disturbances are expected, i.e., $R \equiv P$. This leads to the notion of **set stability** of the set characterized by P [71, 89]. The following set stability definitions are standard [71, 89], except (compared to the literature) the following definitions do *not* assume any topological properties of the set characterized by formula P. The motivation for additional topological restrictions is explained in Corollary 5.21.

Definition 5.18 (Set stability [71, 89])**.** Let $\mathrm{dist}(x, P)$ denote the distance of a point $x \in \mathbb{R}^n$ to the set characterized by formula P. The set characterized by formula P is

- **stable** if, for all $\varepsilon{>}0$, there exists $\delta{>}0$ such that for all initial states $x = x(0)$ with $\mathrm{dist}(x, P) < \delta$, the right-maximal ODE solution $x(t) : [0, T) \rightarrow \mathbb{R}^n$ satisfies the distance bound $\mathrm{dist}(x(t), P) < \varepsilon$ for all times $0 \leq t < T$,
- **attractive** if there exists $\delta{>}0$ such that for all initial states $x = x(0)$ with $\mathrm{dist}(x, P) < \delta$, the right-maximal ODE solution $x(t) : [0, T) \rightarrow \mathbb{R}^n$ approaches the set characterized by P asymptotically with $\lim_{t \rightarrow T} \mathrm{dist}(x(t), P) = 0$,
- **asymptotically stable** if it is stable and attractive, and
- **globally asymptotically stable** if it is stable and for all initial states $x = x(0) \in \mathbb{R}^n$, the right-maximal ODE solution $x(t) : [0, T) \rightarrow \mathbb{R}^n$ satisfies the limit $\lim_{t \rightarrow T} \mathrm{dist}(x(t), P){=}0$.

Lemma 5.19 (Set Stability in dL)**.** *For the ODE $x' = f(x)$, the set characterized by formula P is i) **stable**, ii) **attractive**, iii) **asymptotically stable**, and iv) **globally asymptotically stable** iff the following* dL *formulas are valid, respectively:*

i) $\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P)$,

ii) $\exists \delta{>}0\ \mathrm{Attr}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), \mathcal{U}_\delta(P), P)$,

iii) $\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P) \land \exists \delta{>}0\ \mathrm{Attr}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), \mathcal{U}_\delta(P), P)$*, and*

iv) $\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P) \land \mathrm{Attr}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), \mathit{true}, P)$*.*

Proof. Like Lemma 5.4, the correctness of these definitions is immediate from the semantics of dL formulas because these definitions directly syntactically express the definitions in dL. For $\varepsilon > 0$, the neighborhood formula $\mathcal{U}_\varepsilon(P)$ characterizes the set of points $x \in \mathbb{R}^n$ within distance ε from P, i.e., $\mathrm{dist}(x, P) < \varepsilon$. Formula $\mathrm{Asym}(x' = f(x), P)$ syntactically expresses the limit $\lim_{t \to T} \mathrm{dist}(x(t), P) = 0$ for the right-maximal ODE solution $x(t) : [0, T) \to \mathbb{R}^n$ in dL, as shown in the proof of Lemma 5.4. □

The intuition for Lemma 5.19 is similar to Lemmas 5.4 and 5.13, except formula P (instead of the origin) characterizes the set of desirable states. An application of set stability is shown in the following example.

Example 5.20 (Tennis racket theorem [9])**.** The following system of ODEs models the rotation of a 3D rigid body [33, 71], where x_1, x_2, x_3 are angular velocities and $I_1 > I_2 > I_3 > 0$ are the principal moments of inertia along the respective axes.

$$\alpha_r \equiv x_1' = \frac{I_2 - I_3}{I_1} x_2 x_3, \quad x_2' = \frac{I_3 - I_1}{I_2} x_3 x_1, \quad x_3' = \frac{I_1 - I_2}{I_3} x_1 x_2$$

When such a rigid object is spun or rotated on each of its axes, a well-known physical curiosity [9] is that the rotation is stable in the first and third axes, whilst additional (unstable) twisting motion is observed for the intermediate axis. Mathematically, a perfect rotation, e.g., about x_1, corresponds to a (large) initial value for x_1 with no rotation in the other axes, i.e., $x_2 = 0$, $x_3 = 0$. Accordingly the real-world observation of stability for rotations about the first principal axis is explained by stability with respect to small perturbations in x_2, x_3, as formally specified by formula (5.3) below. Note that the set characterized by formula $x_2 = 0 \land x_3 = 0$ is the entire x_1 axis, not just a single point. Similarly, rotations are stable about the third principal axis iff formula (5.4) is valid.

$$\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(\alpha_r, x_2 = 0 \land x_3 = 0, x_2 = 0 \land x_3 = 0) \tag{5.3}$$

$$\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(\alpha_r, x_1 = 0 \land x_2 = 0, x_1 = 0 \land x_2 = 0) \tag{5.4}$$

The validity of formulas (5.3) and (5.4) are proved in Example 5.23. △

The formal specification of set stability yields three provable logical consequences which are important stepping stones for the set stability proof rules.

Corollary 5.21 (Set stability properties)**.** *The following axioms are derivable in* dL. *In axiom SClosure, formula* $\overline{P}$ *characterizes the topological closure of formula* P *(see Section 2.2.2 and Appendix B.1.3). In axiom SClosed, formula* P *characterizes a closed set.*

SetSAttr $\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P) \to \big(\mathrm{Asym}(x' = f(x), P) \leftrightarrow \forall \varepsilon{>}0\, \langle x' = f(x) \rangle\, \mathcal{U}_\varepsilon(P)\big)$

SClosure $\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P) \leftrightarrow \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), \overline{P}, \overline{P})$

SClosed $\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P) \to \forall x \big(P \to [x' = f(x)]P\big)$

Proof in Appendix C.1.2.

Axiom SetSAttr generalizes SAttr and provides a syntactic simplification of the region of attraction for formula P when P is stable. Axiom SClosure says that stability of P is equivalent to stability of its closure $\overline{P}$, because for any perturbation $\delta > 0$, the neighborhoods $\mathcal{U}_\delta(P)$ and $\mathcal{U}_\delta(\overline{P})$ are provably equivalent in real arithmetic. Axiom SClosed says that for closed formulas P, invariance of P is a necessary condition for stability of P. Without loss of generality, it suffices to develop proof rules for stability of formulas characterizing closed (using SClosure) and invariant (using SClosed) sets. Indeed, standard definitions of set stability [71, 89] usually assume that the set of concern is closed and invariant.

Lemma 5.22 (Set stability Lyapunov functions)**.** *The following Lyapunov function proof rules for set stability are derivable in* dL*. In derived rules SLyap*$_\geq$ *and SLyap*$_>$*, formula P characterizes a compact (i.e., closed and bounded) set. In derived rule SLyap*$^*_\geq$*, the two premises are stacked and there are* ***no*** *topological restrictions on formula P.*

$$\text{SLyap}_{\geq}\ \frac{P \vdash [x' = f(x)]P \quad \neg P \vdash V > 0 \land \dot{V} \leq 0 \qquad \partial P \vdash V \leq 0}{\vdash \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P)}$$

$$\text{SLyap}_{>}\ \frac{P \vdash [x' = f(x)]P \quad \neg P \vdash V > 0 \land \dot{V} < 0 \qquad \partial P \vdash V \leq 0}{\vdash \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P) \land \exists \delta{>}0\ \mathrm{Attr}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), \mathcal{U}_\delta(P), P)}$$

$$\text{SLyap}^*_{\geq}\ \frac{\begin{array}{l} P \vdash [x' = f(x)]P \\ \vdash \forall \varepsilon{>}0\, \exists 0{<}\gamma{\leq}\varepsilon \left(\begin{array}{l} \exists W \left(\begin{array}{l} \forall x\, (\partial(\mathcal{U}_\gamma(P)) \to V \geq W) \land \\ \exists 0{<}\delta{\leq}\gamma\, \forall x\, (\mathcal{U}_\delta(P) \land \neg P \to V < W) \end{array} \right) \land \\ \forall x\, (\overline{\mathcal{U}_\gamma(P)} \land \neg P \to \dot{V} \leq 0) \end{array} \right) \end{array}}{\vdash \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P)}$$

Proof in Appendix C.1.2.

All three proof rules have the necessary premise $P \vdash [x' = f(x)]P$ which says that formula P is an invariant of the ODE $x' = f(x)$. Rules SLyap$_\geq$, SLyap$_>$ are slight generalizations of Lyapunov function proof rules for set stability [71] and they respectively generalize rules Lyap$_\geq$, Lyap$_>$ to prove stability for an invariant P. Importantly, both rules assume that P characterizes a compact, i.e., closed and bounded set, which simplifies the arithmetical conditions on V in their premises. The rule *without* the boundedness requirement on P suggested in the remark after [89, Definition 8.1], is unsound, see Counterexample C.2. For asymptotic stability (in rule SLyap$_>$), boundedness also guarantees that perturbed ODE solutions always exist for sufficient duration, which is a fundamental step in the ODE liveness proofs (Section 4.3). Rule SLyap$^*_\geq$ is derived from rule GLyap using invariance of P by the first premise; it provides a means of formally proving the set stability properties (5.3) and (5.4) from Example 5.20.

Example 5.23 (Stability of rigid body motion)**.** The proof for (5.3) uses the Lyapunov function $V = \frac{1}{2}(\frac{I_1 - I_2}{I_3} x_2^2 - \frac{I_3 - I_1}{I_2} x_3^2)$, whose Lie derivative is $\dot{V} = 0$, and rule SLyap$^*_\geq$ with formula $P \equiv x_2 = 0 \land x_3 = 0$. The proof for (5.4) is symmetric. For the top premise of rule SLyap$^*_\geq$, formula P is a provable invariant of the ODE α_r by DRI. The bottom premise, although arithmetically complicated, can be simplified by choosing $\gamma = \varepsilon$ and deciding the resulting formula by $\mathbb{R}$.

Recall that the x_1 axis is *not* a compact set so neither of the standard proof rules for set stability SLyap$_{\geq}$, SLyap$_{>}$ would be sound for this proof. $\triangle$

Epsilon-Stability

Motivated by numerical robustness of proofs of stability, Gao et al. [61] define ε-stability for ODEs as follows, except the first quantification over $\gamma > \varepsilon$ below is strict whereas in the original definition [61] it is $\gamma \geq \varepsilon$. This difference is immaterial for the purpose of ε-stability as ε is a numerical parameter for the radius of a ball around which disturbances to the origin are to be ignored. In particular, an ODE is ε-stable by the following definition is $\alpha\varepsilon$-stable for any $\alpha \in (0, 1)$ by its original definition [61].

Definition 5.24 (Epsilon-Stability [61])**.** The origin $0 \in \mathbb{R}^n$ of ODE $x' = f(x)$ is ε**-stable** for a positive constant $\varepsilon > 0$ if, for all $\gamma > \varepsilon$, there exists $\delta > 0$ such that for points $x = x(0)$ with $\|x\|_2 < \delta$, the right-maximal ODE solution $x(t) : [0, T) \to \mathbb{R}^n$ satisfies the norm bound $\|x(t)\|_2 < \gamma$ for all times $0 \leq t < T$.

Lemma 5.25 (ε-Stability in dL)**.** *The origin of ODE $x' = f(x)$ is **ε-stable** for constant $\varepsilon > 0$ iff the* dL *formula* $\mathrm{Stab}_{\mathrm{R}}^{\mathrm{P}}(x' = f(x), x = 0, \mathcal{U}_\varepsilon(x = 0))$ *is valid.*

Proof. The formula $\mathrm{Stab}_{\mathrm{R}}^{\mathrm{P}}(x' = f(x), x = 0, \mathcal{U}_\varepsilon(x = 0))$ is valid iff for all $\gamma > 0$, there exists $\delta > 0$ such that for points $x = x(0)$ with $\|x\|_2 < \delta$, the right-maximal ODE solution $x(t) : [0, T) \to \mathbb{R}^n$ satisfies $\|x(t)\|_2 < \varepsilon + \gamma$ for all times $0 \leq t < T$, where the neighborhood $\mathcal{U}_\gamma(\mathcal{U}_\varepsilon(x = 0))$ is equivalently characterized by $\|x\|_2^2 < \gamma + \varepsilon$. This unfolded semantics is equivalent to the mathematical definition of ε-stability in Def. 5.24 by reindexing the universal quantifier with $\gamma \mapsto \gamma + \varepsilon$ instead. $\square$

Unlike set stability, ε-stability is an instance of general stability where the pre- and postconditions differ. In ε-stability, systems are perturbed from the precondition $x = 0$ (the origin), but the postcondition enlarges the set of desired states to a $\varepsilon > 0$ neighborhood of the origin, which is considered indistinguishable from the origin itself [61]. An immediate consequence of Lemma 5.25 is that rule GLyap can be used to prove ε-stability, as shown in the next section.

5.4 Stability in KeYmaera X

This section puts the dL stability specifications and derivations from the preceding sections into practice through proofs for several case studies in the KeYmaera X theorem prover [54].[4] Examples 5.7, 5.15, 5.20, 5.23 have also been formalized. The insights from these proofs are discussed after an overview of the case studies.

[4] See `https://github.com/LS-Lab/KeYmaeraX-projects/blob/master/stability/ODE`.
Git hash: `c856fddb383232adbd86679ef65567f9b90190bf`

Inverted Pendulum. The stability of the resting state of the pendulum is investigated in Examples 5.7 and 5.15. For the inverted pendulum α_i from (5.2), the controlled torque $u(\theta, \omega)$ must be designed and rigorously proved to ensure *feedback stabilization* [89] of the inverted position. A standard PD (Proportional-Derivative) feedback controller can be used for stabilization, where the continuous control input has the form $u(\theta, \omega) = k_1\theta + k_2\omega$ for tuning parameters k_1, k_2, based on the values of the state variables θ (proportional term) and ω (derivative term). Asymptotic stability of the inverted position is achieved for any control parameter choice where $k_1 > a$ and $k_2 > -b$. The sequent $a > 0, b \geq 0, k_1 > a, k_2 > -b \vdash \mathrm{AStab}(\alpha_i)$ is proved in KeYmaera X using the Lyapunov function $\frac{(k_1-a)\theta^2}{2} + \frac{((b+k_2)\theta+\omega)^2+\omega^2}{4}$.

Frictional Tennis Racket Theorem. The stability of a 3D rigid body is investigated for α_r in Examples 5.20 and 5.23. The following ODEs model additional frictional forces that oppose the rotational motion in each axis of the rigid body, where $\alpha_1, \alpha_2, \alpha_3 > 0$ are coefficients of friction:

$$\alpha_f \equiv x_1' = \frac{I_2 - I_3}{I_1}x_2x_3 - \alpha_1x_1,\ x_2' = \frac{I_3 - I_1}{I_2}x_3x_1 - \alpha_2x_2,\ x_3' = \frac{I_1 - I_2}{I_3}x_1x_2 - \alpha_3x_3$$

In the presence of friction, rotations of the rigid body are globally asymptotically stable in the first and third principal axes, as proved in KeYmaera X.

$$\Gamma \equiv I_1 > I_2, I_2 > I_3, I_3 > 0, \alpha_1 > 0, \alpha_2 > 0, \alpha_3 > 0$$
$$\Gamma \vdash \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(\alpha_f, x_2 = 0 \land x_3 = 0, x_2 = 0 \land x_3 = 0) \land \mathrm{Attr}^{\mathrm{P}}_{\mathrm{R}}(\alpha_f, \mathit{true}, x_2 = 0 \land x_3 = 0)$$
$$\Gamma \vdash \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(\alpha_f, x_1 = 0 \land x_2 = 0, x_1 = 0 \land x_2 = 0) \land \mathrm{Attr}^{\mathrm{P}}_{\mathrm{R}}(\alpha_f, \mathit{true}, x_1 = 0 \land x_2 = 0)$$

Both asymptotic stability properties are proved using $\mathrm{SLyap}^*_{\geq}$ and the liveness property (Chapter 4) that the kinetic energy $I_1x_1^2 + I_2x_2^2 + I_3x_3^2$ of the system tends to zero over time. The latter property implies that solutions of α_f exist globally and that the values of x_1, x_2, x_3 asymptotically tend to zero, which proves global asymptotic stability with the aid of SetSAttr. Even though a proof rule for (global) asymptotic stability of general nonlinear ODEs and unbounded sets is not available (Section 5.3), this example shows that formalized stability properties can still be proved on a case-by-case basis using dL's ODE reasoning principles.

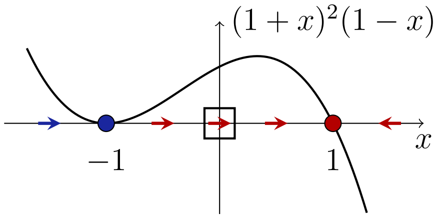

Attractive but Unstable System. Consider the ODE $\alpha_u \equiv x' = (1+x)^2(1-x)$ which has equilibrium points at $x = \pm 1$ as visualized on the right. The set characterized by formula $P \equiv x = 1 \lor x = -1$ is globally attractive but *not* stable, i.e., the formulas $\mathrm{Attr}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), \mathit{true}, P)$ and $\neg\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P)$ are valid. Intuitively, all points to the left of $x = -1$ are attracted towards the blue equilibrium point while all other points are attracted to the red equilibrium point. However, the set consisting of both equilibria is *not* set stable because states arbitrarily close to the right of the blue point are attracted towards the red point, so those points attain a maximal distance of 1 from *both* equilibria at $x = 0$ (unfilled black square). The proof of global attractivity in KeYmaera X follows the above intuition by case splitting on $x \leq -1 \lor -1 < x < 1 \lor x \geq 1$. The proof of *instability* shows that points close to the right of $x = -1$ eventually reach $x = 0$ (violating set stability) on their way towards $x = 1$.

Table 5.1: Proof statistics for ODE stability properties proved in Section 5.4. Examples 5–11 refer to the correspondingly numbered examples from Ahmed et al. [3]. "Dim." is the number of continuously evolving state variables in the ODEs; "Param." is the number of parameters (non-state variables) in the stability specification; "Deg." is the maximum degree of polynomials with respect to the state variables in the ODEs; "Tactic Steps" counts the number of (manual) user proof steps; and "Proof Time" measures the time taken (in seconds, averaged over 5 runs, rounded to 3 decimal places) for the proof to execute in KeYmaera X.

Stability Property	Dim.	Param.	Deg.	Tactic Steps	Proof Time (s)
Example 5.7	2	2	1	119	6.757
Example 5.15	2	2	1	20	4.533
Inverted Pendulum	2	4	1	119	31.427
Examples 5.20 and Example 5.23	3	3	2	58	5.801
Frictional Tennis Racket Theorem	3	6	2	162	11.919
Attractive but Unstable System	1	0	3	93	21.911
Moore-Greitzer Jet Engine [61]	2	0	3	34	6.258
Ahmed et al. [3, Example 5]	6	0	3	86	26.440
Ahmed et al. [3, Example 6]	2	0	3	31	0.673
Ahmed et al. [3, Example 7]	3	0	3	33	1.835
Ahmed et al. [3, Example 8]	2	0	5	31	0.978
Ahmed et al. [3, Example 9]	4	0	3	54	4.694
Ahmed et al. [3, Example 10]	2	1	1	31	1.126
Ahmed et al. [3, Example 11]	2	4	3	31	11.460

Moore-Greitzer Jet Engine [61]. The origin of the ODE modeling a simplified jet engine $\alpha_m \equiv x_1' = -x_2 - \frac{3}{2}x_1^2 - \frac{1}{2}x_1^3,\ x_2' = 3x_1 - x_2$ is ε-stable for $\varepsilon = 10^{-10}$ [61]. The sequent $\varepsilon = 10^{-10} \vdash \mathrm{Stab}_{\mathrm{R}}^{\mathrm{P}}(\alpha_m, x_1^2 + x_2^2 = 0, x_1^2 + x_2^2 < \varepsilon^2)$ is proved in KeYmaera X. The key proof ingredients are a ε-Lyapunov function which can be automatically generated [61] and manual arithmetic steps, e.g., instantiating existential quantifiers appearing in the specification of ε-stability with appropriate values [61].

Other Examples [3]. Stability for several ODEs with Lyapunov functions generated by an inductive synthesis technique [3, Examples 5–11] were successfully verified in KeYmaera X. The proof for the largest, 6-dim. nonlinear ODE [3, Example 5] required substantial manual arithmetic reasoning in KeYmaera X.[5] The arithmetical conditions in [3, Equation 1] are identical to the premises of rule Lyap$_{\geq}$ except [3, Equation 1] unsoundly omits the condition $V(0) = 0$, see Counterexample C.3. The generated Lyapunov functions remain correct because the inductive synthesis technique [3] implicitly guarantees this omitted condition.

[5] The Lyapunov function as given in [3, Example 5] does *not* work for its associated ODE. It works if the ODE is corrected with $\dot{x}_1 = -x_1^3 + 4x_2^3 - 6x_3x_4$, as in the literature [130].

Summary. These case studies demonstrate the feasibility of carrying out proofs of various (advanced) stability properties within KeYmaera X using this chapter's stability specifications. Table 5.1 provides a summary of statistics from these proofs, where all experiments were run on an Ubuntu 18.04 laptop with a 2.70 GHz Intel Core i7-6820HQ CPU and 16GB memory. The proofs share similar high-level proof structure, which suggests that proof automation could significantly reduce proof effort [55]. Such automation should also support user input of key insights for difficult reasoning steps, e.g., real arithmetic reasoning with nested, alternating quantifiers. These insights are used in the implementation of (more general) switched system stability automation in Chapter 6.

5.5 Input-to-State Stability

This chapter has, thus far, focused on stability of ODEs with respect to perturbations of the *system state*. This section takes a brief detour to examine stability properties of ODEs under perturbations of the *system dynamics*, e.g., for a system under continuous feedback control, the system designer may wish to account for noisy or unexpected perturbations to the continuous inputs, in addition to state disturbances. A key barrier to formalizing stability under continuous perturbations using ODEs in dL, i.e., *without* extensions like differential games [143] or differential-algebraic programs [137], is those properties often need higher-order quantification over functions that model perturbations of ODEs. As an example, consider the notion of input-to-state stability [71, 89, 181, 182] for an ODE $x' = f(x, u(t))$ with time-dependent control $u(t)$ defined below.

Definition 5.26 (Input-to-state stability [71, 89, 181, 182])**.** The ODE $x' = f(x, u(t))$ with time-dependent input $u(t)$ and $f(0,0) = 0$ is **input-to-state stable** iff

- the unforced ODE (with 0 input) $x' = f(x, 0)$ is globally asymptotically stable and
- the ODE $x' = f(x, u(t))$ has the **asymptotic gain** property, i.e., for all $E > 0$, there exists $\Delta > 0$, such that for all bounded functions $u(t)$ with supremum norm $\|u(t)\|_\infty \leq \Delta$ and for all $x = x(0)$, the right-maximal ODE solution $x(t) : [0, T) \to \mathbb{R}^n$ satisfies $\limsup_{t \to T} \|x(t)\| \leq E$.

The notion of input-to-state stability says that, without the forcing perturbation $u(t)$, the ODE $x' = f(x, 0)$ has the usual asymptotic stability with respect to perturbations of the system state at the origin. In addition, for a bounded perturbation modeled by function $u(t)$, the ODE remains asymptotically close to a ball around the origin with radius dependent on the magnitude of the perturbing input. In particular, sufficiently small bounded perturbations to the continuous dynamics *cannot* force the system arbitrarily far from the origin. Note that Def. 5.26 differs slightly from the controls literature [71, 89, 181, 182] for uniformity with the rest of this chapter. It is common to define input-to-state stability using so-called comparison functions, see the literature [182, Section A]. The definition of input-to-state stability as a combination of global asymptotic stability and asymptotic gain is one of many mathematically equivalent definitions of input-to-state stability, as shown by Sontag and Wang [182, Theorem 1].

Turning to dL specification, by Lemma 5.13, global asymptotic stability of $x' = f(x, 0)$ is specified by the formula $\mathrm{Stab}(x' = f(x, 0)) \land \mathrm{Attr}^{\mathrm{P}}(x' = f(x, 0), \mathit{true})$. The asymptotic gain property, however, presents a challenge because the first-order quantifier syntax $\forall x\,/\exists x$ in dL can only directly quantify over real numbers x. The following "formula" almost captures the definition of asymptotic gain, where the required asymptotic bound on the norm of $x(t)$ is expressed by formula $\mathrm{Asym}(\{x' = f(x, u(t)), t' = 1\}, \mathcal{U}_E(x = 0))$ from Lemma 5.4. However, the "formula" is not syntactically allowed in dL because it uses higher-order (dependent) quantification over all Δ-bounded functions u, as highlighted in red.

$$\forall E{>}0\, \exists \Delta{>}0\, \forall \|u\|_\infty {\leq} \Delta\, \forall x\, (\mathrm{Asym}(\{x' = f(x, u(t)), t' = 1\}, \mathcal{U}_E(x = 0)))$$

An alternative is to use dynamical extensions of dL that can express continuous perturbations directly. For example, with the differential game [143] $\{x' = f(x, u) \&^d \|u\|_\infty \leq \Delta\}$, where the input u is modeled by a continuous adversary that is constrained to output values u within the domain $\|u\|_\infty \leq \Delta$, or the differential-algebraic program [137] $\{\exists u\, x' = f(x, u) \land \|u\|_\infty \leq \Delta\}$, which models an input u that is chosen nondeterministically at each (continuous) time instant.

The base hybrid program language (without extensions) can be used to directly model some sub-classes of perturbations $u(t)$. As a preview of Chapter 6, consider the case where the perturbations are piecewise constant, so they *discretely switch* between different values on each time interval. Formally, the perturbations $u(t)$ of interest are restricted to piecewise constant functions with finitely many pieces on each finite time interval, so each $u(t)$ is defined by a sequence of *switching times* $0 = \tau_0 < \tau_1 < \tau_2 < \dots$ with $\tau_i \to \infty$ and a sequence $p_1, p_2, \dots \in \mathbb{R}$, such that $u(t) = p_i$ for all $\tau_{i-1} \leq t \leq \tau_i, 1 \leq i$. The looping hybrid program $\alpha_{\mathtt{piece}}$ below nondeterministically chooses a value with $\|u\|_\infty \leq \Delta$ on each loop iteration, and then follows the ODE $x' = f(x, u)$ for a nondeterministic duration.

$$\alpha_{\mathtt{piece}} \equiv (u := *; ?\, \|u\|_\infty \leq \Delta; x' = f(x, u))^* \tag{5.5}$$

Proposition 5.27. *A state is reachable by hybrid program $\alpha_{\mathtt{piece}}$ iff it is reachable in finite time by the time-dependent ODE $x' = f(x, u(t))$ for some piecewise constant function $u(t)$ with $\|u(t)\|_\infty \leq \Delta, \Delta \in \mathbb{R}$ and finitely many pieces on each finite time interval.*

Proof Summary (Appendix D.1.1). The full proof is deferred to Chapter 6 because it is similar to the proofs of the adequacy theorems for switched systems in that chapter. In the "$\Rightarrow$" direction, a piecewise constant function $u(t)$ is constructed from the sequence of choices for $u := *$ from the semantics of the loop in $\alpha_{\mathtt{piece}}$. In the "$\Leftarrow$" direction, a trace of hybrid program $\alpha_{\mathtt{piece}}$ is constructed from the defining sequences $0 = \tau_0 < \tau_1 < \tau_2 < \dots$ and $p_1, p_2, \dots \in \mathbb{R}$ for $u(t)$, where each time interval $[\tau_i, \tau_{i+1}]$ corresponds to one loop iteration in $\alpha_{\mathtt{piece}}$. □

Thanks to Proposition 5.27, the box modality formula $[\alpha_{\mathtt{piece}}]P$ expresses that for piecewise constant $u(t)$ (bounded with $\|u(t)\|_\infty \leq \Delta$), the solution of ODE $x' = f(x, u(t))$ satisfies postcondition P at all times. Hence, safety properties of ODEs with piecewise constant (bounded) perturbations can be specified and reasoned about in dL. Hybrid program models for other switching mechanisms and their corresponding safety (and stability) specifications are explained in more detail in Chapter 6.

Dually, the *diamond* modality formula $\langle\alpha_{\mathtt{piece}}\rangle P$ expresses that for *some* piecewise constant perturbations $u(t)$, the solution of ODE $x' = f(x, u(t))$ eventually satisfies postcondition P. This leads to an important subtlety for stability specification because the formula

$\forall \varepsilon{>}0\, \langle \alpha_{\texttt{piece}} \rangle [\alpha_{\texttt{piece}}]\, \mathcal{U}_\varepsilon(P)$ (similar to Lemma 5.4) characterizes that $\alpha_{\texttt{piece}}$ asymptotically approaches P for all $u(t)$ with *some* (in red) piecewise constant prefix. A similar subtlety arises for the differential game $\{x' = f(x,u) \&^d \|u\|_\infty \leq \Delta\}$ and differential-algebraic program $\{\exists u\, x' = f(x,u) \land \|u\|_\infty \leq \Delta\}$ discussed above, where, intuitively, the quantification over all or some adversarial choice (or nondeterministic input $\exists u \cdots$) depends on the modality in which the game appears (or differential-algebraic program) [143]. The question of how to specify asymptotic gain for these models is subtle and left out of scope for this thesis. Nevertheless, Chapter 6 shows how to sidestep these subtleties for switched systems using the notion of pre-attractivity from the controls literature [65, 66] which is specified as a quantified safety property of switched system models.

5.6 Related Work

Stability is a fundamental property of interest across many different fields of mathematics [33, 77, 98, 153, 165, 187] and engineering [71, 89, 99]. This related work discussion focuses on formal approaches to stability of ODEs.

Logical Specification of Stability. Rouche et al. [165] provide a pioneering example of using logical notation to specify and classify stability properties of ODEs. Alternative logical frameworks have also been used to specify stability and related properties: stability is expressed in HyperSTL [123] as a hyperproperty relating the trace of an ODE against two constant traces; ε-stability is studied in the context of δ-complete reasoning over the reals [61]; region stability for hybrid systems [151] is discussed using CTL*; the syntactic specification of $\mathrm{Asym}(x' = f(x), P)$ resembles the limit definition using filters [78]. This chapter uses dL as a *sweet spot* logical framework, general enough to specify various stability properties of interest, e.g., asymptotic or exponential stability, and the stability of sets, while simultaneously enabling syntactic, formal *proofs* of those properties within the logic.

Formal Verification of Stability. There is a vast literature on finding Lyapunov functions for stability, e.g., through numerical [130, 131, 200] and algebraic methods [49, 104]. Formal approaches are often based on finding Lyapunov function candidates and *certifying* the correctness of those generated candidates [3, 61, 88, 170]. This chapter's approach directly proves stability specifications by step-by-step derivation using dL's parsimonious axiomatization [139, 142, 144], which goes beyond certification of arithmetic conditions on Lyapunov functions that imply stability. This dL formalization yields logical insights into relationships between stability properties (e.g., Section 5.2.4 and Corollary 5.21). Furthermore, the practical application of the stability derivations in KeYmaera X [54, 142] enables highly trustworthy certification of Lyapunov function candidates generated by the aforementioned approaches (Section 5.4). Sections 5.3 and 5.4 also show that this chapter's approach supports verification of advanced stability properties [61, 71, 89] within the same dL framework. New stability proof rules like GLyap can also be soundly and *syntactically* justified in dL without the need for (low-level) semantic reasoning about the underlying ODE mathematics. As an example of the latter, semantic approach, LaSalle's invariance principle is formalized in Coq [37] and used to verify the correctness of an inverted pendulum controller [166].

5.7 Discussion

George Box's famous quote—"All models are wrong, but some are useful"—is an important adage to keep in mind whenever one is working with a model of a real-world system. Proofs of stability for a model give increased confidence that conclusions drawn about that model *are* useful because the models are sufficiently robust to real-world deviations. Such proofs are especially important for justifying the correctness of control systems which must be designed to operate robustly in the presence of real-world perturbations. This chapter shows how ODE stability can be formalized in dL using the key idea that stability properties are $\forall/\exists$-quantified dynamical formulas. These specifications, their proof rules, and their logical relationships are all syntactically derived from dL's sound proof calculus, which enables trustworthy KeYmaera X proofs that rigorously verify *every step* in an ODE stability argument, from arithmetical premises down to dynamical reasoning for ODEs. Of course, the same adage applies to *hybrid* system models of CPSs, which must also exhibit stability from both modeling and control perspectives. The next chapter begins the exploration of hybrid system stability by extending the ideas and derivations of this chapter to systems that discretely switch between a family of ODEs.

Chapter 6

Stability for Switched Systems

This chapter applies the safety, liveness, and stability results for ordinary differential equations (ODEs) from the preceding chapters to deductive verification for *switched systems*. Discrete switching between continuous controllers is a simple yet powerful hybrid control design paradigm, but switched systems are known to exhibit subtle (in)stability behaviors so system designers must carefully analyze the stability of closed-loop systems that arise from their proposed switching control laws. This chapter begins by modeling various classes of switched systems as *looping* hybrid programs so that safety and liveness properties for switched systems can be compositionally verified by combining dL's hybrid program reasoning principles with the ODE safety and liveness reasoning developed in Chapters 3 and 4. Deductive proofs of *switched system stability* for those models further blend classical ideas from the controls and verification literature using dL: from controls, the approach uses standard stability notions for various classes of switching mechanisms and their corresponding Lyapunov function-based analysis techniques; from verification, ODE invariants underlying stability proofs from Chapter 5 are lifted to switched systems by identifying appropriate *loop invariants* for each switching mechanism, i.e., properties that are preserved across every switching loop iteration for their respective looping hybrid program models. This blend of ideas enables a trustworthy implementation of switched system stability verification in KeYmaera X, providing fully automated stability proofs for standard classes of switching mechanisms. The generality of the dL approach also allows for verification of switching control laws that require non-standard stability arguments by modifying loop invariants to suitably express specific intuitions behind those control laws. This flexibility is demonstrated on several case studies drawn from the literature.

6.1 Introduction

The study of *hybrid systems*, i.e., mathematical models that combine discrete and continuous dynamics, is motivated by the need to understand the hybrid dynamics present in many real-world systems [6, 28, 45, 66, 75, 144] (see Chapter 1). Various formalisms can be used to describe hybrid systems, for example, impulsive differential equations [72]; switched systems [99, 189]; hybrid time combinations of discrete and continuous dynamics [65, 66]; hybrid automata [75]; and language-based models [45, 135, 144, 164, 188, 208, 213]. These formalisms differ in their generality and in how the discrete-continuous dynamical combination is modeled, e.g., ranging from differential equations with discontinuous right-hand sides, to combinators that piece together discrete and continuous programs. Consequently, different formalisms may be better suited for different hybrid system applications and it is worthwhile to explore connections between different formalisms in order to exploit their various strengths for a given application. This chapter investigates the connection between switched systems and dL's hybrid programs.

A *switched system* consists of a family of continuous ODEs together with a discrete switching signal which prescribes the active ODE that the system follows at each time. Switched systems provide a powerful mathematical paradigm for the design and analysis of discontinuous (or nondifferentiable) control mechanisms [43, 99, 118, 189]. Examples of such mechanisms include: bang-bang controllers that switch between on/off modes; gain schedulers that switch between a family of locally valid linear controllers; and supervisory control, where a supervisor switches between candidate continuous controllers based on logical criteria [99, 118]. Switching control laws can be used to stabilize systems that cannot otherwise be stabilized by continuous feedback control [99] (see Section 6.5.3). However, switched systems are known to exhibit subtle (in)stability behaviors—recall Fig. 1.4 (page 6) which shows that switching between stable subsystems can lead to instability [99]—so it is important for system designers to adequately justify the stability of their proposed switching designs.

Hybrid programs model hybrid dynamics by combining discrete programming constructs with continuous ODEs (see Section 2.1.3). This combination yields a rich and flexible language for describing hybrid systems, e.g., with event- or time-triggered design paradigms [144]. Section 6.2 shows how various classes of switched systems can be fruitfully modeled using *looping* hybrid programs, as illustrated by the following snippet. The switching loop body runs a discrete controller $u := ctrl(x)$ which selects a mode u, followed by the continuous plant $x' = f_u(x)$ which evolves according to the selected mode.

$$\begin{array}{lll} \{ \quad & u := ctrl(x); & \text{// switching controller (discrete dynamics)} \\ & x' = f_u(x) & \text{// actuate decision (continuous dynamics)} \\ \}^{*}\texttt{@invariant}(\ \dots\ ) & & \text{// switching loop with invariant annotation} \end{array} \tag{6.1}$$

These models enable sound and compositional verification of switched systems in dL, e.g., the completeness results for ODE invariants from Chapter 3 are generalized to switched systems, yielding an effective technique for proving switched system safety. Subtleties associated with those models are also investigated, along with methods for detecting and avoiding those pitfalls.

Section 6.3 develops stability proofs for various classes of switching mechanisms using their hybrid program models. The key insight is that control-theoretic stability arguments for switching control can be formally justified by blending techniques from discrete program verification with analysis of continuous differential equations using dL. Similar to Chapter 5, switched system stability is specified by nesting dL's first-order quantification with its dynamic modalities, with the added twist that the specifications must quantify over *all* switching behaviors of the system. The resulting specifications are proved by combining fundamental ideas from verification and control, namely: *i)* identification of appropriate *loop invariants*, `@invariant` in (6.1), i.e., properties of the (discrete) switching loop that are preserved across all executions of the loop body, *ii) compositional verification* for separately analyzing the discrete $u := ctrl(x)$ and continuous

$x' = f_u(x)$ dynamics of the loop body, and *iii) Lyapunov functions*, i.e., auxiliary energy functions that enable stability analysis for the continuous dynamics. Crucially, these stability proofs are *syntactically derived* from dL's sound foundations for hybrid program reasoning [142, 144], *without* the need to introduce new mathematical concepts such as non-classical weak solutions or nondifferentiable Lyapunov functions [39, 65]. The remaining practical challenge is how to (automatically) find suitable Lyapunov function candidates for a given switching mechanism; the correctness of any generated candidates can be soundly checked in dL.

Section 6.4 adds support for switched systems to KeYmaera X [54], including a modeling interface for switched systems, sum-of-squares search for Lyapunov function candidates [129, 154], and fully automatic verification of stability specifications for standard switching mechanisms. Notably, the implementation requires *no extensions* to KeYmaera X's soundness-critical core and thereby directly inherits all of KeYmaera X's correctness guarantees [54, 115]. This trustworthiness is necessary for computer-aided verification of complex switching designs because the number of correctness conditions on their Lyapunov functions scales quadratically with the number of switching modes (Section 6.3.3), making pen-and-paper proofs potentially error-prone or infeasible. Section 6.5 further applies the deductive approach on three case studies, chosen because each require subtle twists to standard switched system stability arguments.

- *Longitudinal flight control* [26]: This model is parametric (5 parameters, 2 state variables) and its stability justification due to Branicky uses a "noncustomary" Lyapunov function [26, 43] with intricate arithmetic reasoning. The proof uses *ghost switching*, where virtual switching modes are introduced for the sake of stability analysis, analogous to the use of ghost variables in program verification [127] or in proofs for ODEs (Chapters 3 and 4).

- *Automatic cruise control* [126]: This hybrid automaton features switching between several modes based on specific guard conditions: standard/emergency braking, accelerating, and PI control. Lyapunov function candidates can be numerically generated with existing tools, e.g., Stabhyli [116], but the results must be corrected for soundness.

- *Brockett's nonholonomic integrator* [29]: A large class of control systems can be transformed to the nonholonomic integrator but this system is not stabilizable by continuous feedback [29, 99]. The stability argument must account for an initial control mode that drives the system into a suitable region before a stabilizing control law can be applied.

These case studies are verified semi-automatically in KeYmaera X, with user guidance to design and prove modified loop invariants that suitably capture the specific intuitions behind their respective control laws. The flexibility and generality of the deductive approach enables such (modified) stability arguments, while ensuring that every step in the argument remains rigorously justified using sound dL logical foundations.

Reminder (Extended Term Language). This chapter uses an extended dL term language following the extended term conditions and notational conventions of Section 3.2 because the dL axiomatization remains sound for all extended term languages meeting those conditions.

Contribution. *The material for this chapter is drawn from Tan and Platzer [193] and from Tan et al. [196].*

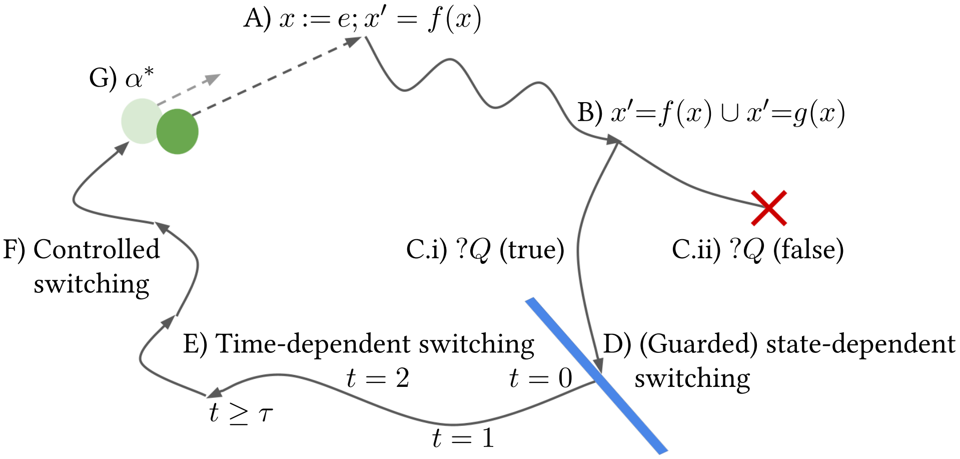

Figure 6.1: The green initial state evolves according to a hybrid program featuring (clockwise from top): A) a discrete assignment (dashed line) followed sequentially by continuous ODE evolution (solid line); B) a choice between two ODEs; C) a test that aborts (red ×) system evolutions leaving Q; D) switching when the system state crosses the thick blue switching surface; E) switching after time $t \geq \tau$ has elapsed; F) switching control that is designed to drive the system state close to its initial position; and G) a loop that repeats system evolution.

6.2 Switched Systems as Hybrid Programs

This section explains how hybrid programs are used to model various classes of switching mechanisms. Following the nomenclature from Liberzon [99], these mechanisms can be broadly categorized into: *autonomous switching* (Section 6.2.2), i.e., without an explicit control logic [28, 99] and *controlled switching* (Section 6.2.3), like (6.1) where a discrete controller $u := ctrl(x)$ decides the ODE $x' = f_u(x)$ to switch to on each switching loop iteration. The evolution of various switching mechanisms and hybrid programs are illustrated in Fig. 6.1.

6.2.1 Mathematical Preliminaries

Switching phenomena can either be modeled explicitly as a function of time, or implicitly, e.g., as a state predicate, depending on the real-world switching mechanism being modeled. Mathematically, a *switched system* is described by the following data:

1. an open, connected set $D \subseteq \mathbb{R}^n$ which is the *state space* of interest for the system,
2. a finite (non-empty) family $\mathcal{P}$ of ODEs $x' = f_p(x)$ for modes $p \in \mathcal{P}$, and,
3. for each initial state $\omega \in D$, a set of *switching signals* $\sigma : [0, \infty) \to \mathcal{P}$ prescribing the ODE $x' = f_{\sigma(t)}(x)$ to follow at time t for the system's evolution from ω.

Switching signals $\sigma : [0, \infty) \to \mathcal{P}$ are always assumed to be *well-defined* [99, 189], i.e., every σ has finitely many discontinuities on each finite time interval in its domain $[0, \infty)$, so that they model physically realizable switching. For finite $\mathcal{P}$, this means σ is a piecewise constant function with finitely many pieces on each finite time interval where, σ prescribes a mode

$p \in \mathcal{P}$ to switch to on each piece. For simplicity, σ is also assumed to be right-continuous [66], so for neighboring discontinuities at times $a < b$, σ has a constant value on the interval $[a, b)$ and its value changes at time b. With these assumptions, switching signals are equivalently defined by a sequence of *switching times* $0 = \tau_0 < \tau_1 < \tau_2 < \dots$ with $\tau_i \to \infty$ and a sequence $p_1, p_2, \dots \in \mathcal{P}$ where p_i for $i \geq 1$ specifies the value taken by σ on the time interval $[\tau_{i-1}, \tau_i)$:

$$\sigma(t) = \begin{cases} p_1 & \text{if } \tau_0 \leq t < \tau_1 \\ p_2 & \text{if } \tau_1 \leq t < \tau_2 \\ \dots & \\ p_i & \text{if } \tau_{i-1} \leq t < \tau_i \end{cases} \tag{6.2}$$

For a switching signal σ and initial state $\omega \in \mathbb{R}^n$, the *solution* φ of the switched system is the function generated inductively on the sequences τ_i and p_i as follows. Define $\varphi(0) = \omega$. For switching time τ_i with $i \geq 1$, if φ is defined at time τ_{i-1}, then the definition of φ is extended by considering the unique, right-maximal solution to the ODE $x' = f_{p_i}(x)$ starting from $\varphi(\tau_{i-1})$ [33], i.e., $\psi_i : [0, \zeta_i) \to \mathbb{R}^n$ with $\psi_i(0) = \varphi(\tau_{i-1})$, $\frac{\mathrm{d}\psi_i(t)}{\mathrm{d}t} = f_{p_i}(\psi_i(t))$, and $0 < \zeta_i \leq \infty$. If $\zeta_i \leq \tau_i - \tau_{i-1}$, then the system *blows up* before reaching the next switching time τ_i, so define $\varphi(\tau_{i-1}+t) = \psi_i(t)$ on the bounded time interval $t \in [0, \zeta_i)$. Otherwise, $\zeta_i > \tau_i - \tau_{i-1}$, then define $\varphi(\tau_{i-1}+t) = \psi_i(t)$ on the time interval $t \in [0, \tau_i - \tau_{i-1}]$. This inductive construction uniquely defines a solution $\varphi : [0, \zeta) \to \mathbb{R}^n$ associated with ω and σ for (right-maximal) time $\zeta > 0$. The switched system *reaches* $\varphi(t)$ at time $t \in [0, \zeta)$. When the system is associated with a family of domains Q_p, $p \in \mathcal{P}$, the switched system reaches $\varphi(t)$ while *obeying* the domains iff for all $i \geq 1$ and time $\gamma \in [\tau_{i-1}, \tau_i] \cap [0, t]$, the state $\varphi(\gamma)$ satisfies Q_{p_i}. The truncated solution $\varphi : [0, T_\varphi] \to \mathbb{R}^n$ for $T_\varphi < \zeta$ is *domain-obeying* if the system obeys the domains for all times $t \in [0, T_\varphi]$.

For simplicity, this chapter assumes that the state space is $D = \mathbb{R}^n$. More general definitions of switched systems are possible but are left out of scope [99]. For example, $\mathcal{P}$ can more generally be an (uncountably) infinite family, like $\alpha_{\mathtt{piece}}$ from Proposition 5.27, and some switched systems may have *impulse effects* where the system state is allowed to make instantaneous, discontinuous jumps during the system's evolution, such as the dashed jump in part A of Fig. 6.1.

6.2.2 Autonomous Switching

This section examines hybrid program models of *autonomous switching* mechanisms; these models are syntactically simpler in the sense that they do not need an explicit program to model their switching control logics [28, 99]. Nevertheless, models of these switching mechanisms are useful because they can be used to describe real-world systems where discrete switching behaviors occur nondeterministically outside of the system designer's control [99]. Safety proofs for these models account for *all* nondeterministic switching behaviors.

Arbitrary Switching

Real world systems can exhibit switching behaviors that are uncontrolled, *a priori* unknown, or too complicated to describe succinctly in a model. For example, a driving vehicle may encounter several different road conditions depending on the time of day, weather, and other unpredictable factors—given the multitude of combinations to consider, it is desirable to have a single model that exhibits and switches between all of those road conditions *without* needing explicit descriptions of exactly when and how switching between different conditions can occur.

Arbitrary switching is a useful tool for modeling such systems because it considers *all* possible switching signals and their corresponding system evolutions. It is modeled by the following hybrid program $\alpha_{\texttt{arb}}$ and illustrated in Fig. 6.2.

$$\alpha_{\texttt{arb}} \equiv \Big(\bigcup_{p \in \mathcal{P}} x' = f_p(x) \Big)^* \tag{6.3}$$

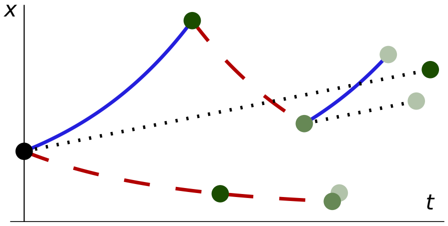

Figure 6.2: Evolution of $\alpha_{\texttt{arb}}$ that switches between ODEs $x' = x$ (solid blue), $x' = 1$ (dotted black), and $x' = -x$ (dashed red) from the initial state (black circle). Switching steps are marked by green circles and faded colors illustrate progression in loop iterations for the loop operator in $\alpha_{\texttt{arb}}$.

Intuitively, $\alpha_{\texttt{arb}}$ models arbitrary switching analogously to a computer simulation: on each loop iteration, the program makes a (discrete) nondeterministic choice of switching decision $\bigcup_{p \in \mathcal{P}} (\cdot)$ to select an ODE $x' = f_p(x)$ which it then follows continuously for an arbitrarily chosen duration before repeating the switching loop. Two subtle behaviors are illustrated by the bottom trajectory in Fig. 6.2: $\alpha_{\texttt{arb}}$ can switch to the same ODE across a loop iteration or it can *chatter* by making several discrete switches without continuously evolving its state between those switches [177]. These behaviors are harmless for safety verification because they do not change the set of reachable states of the switched system. The adequacy of $\alpha_{\texttt{arb}}$ as a model of arbitrary switching is shown in the following proposition.

Proposition 6.1. *A state is reachable by hybrid program $\alpha_{\texttt{arb}}$ iff it is reachable in finite time by a switched system $x' = f_p(x)$ for $p \in \mathcal{P}$ following a switching signal σ.*

Proof. This follows from the subsequent Proposition 6.2 with $Q_p \equiv \mathit{true}$ for all $p \in \mathcal{P}$. □

By Proposition 6.1, the dL formula $[\alpha_{\texttt{arb}}]P$ specifies safety for arbitrary switching, i.e., for *all* switching signals σ, all states reached by switching according to σ satisfy the safety postcondition P. Dually, formula $\langle \alpha_{\texttt{arb}} \rangle P$ expresses that the system eventually reaches the goal P for *some* switching signal σ. Proofs of these specifications exploit compositionality [144] by combining dL reasoning for the discrete loop and choice operators in $\alpha_{\texttt{arb}}$ with ODE safety and liveness reasoning from Chapters 3 and 4. Conversely, reasoning principles for ODEs compositionally lift to $\alpha_{\texttt{arb}}$ through dL axioms, so dL also has complete invariance reasoning principles for $\alpha_{\texttt{arb}}$. This is shown next, after a slight generalization to state-dependent switching models.

State-Dependent Switching

Arbitrary switching can be constrained by enabling switching to the ODE $x' = f_p(x)$ only when the system state belongs to a corresponding domain specified by formula Q_p. This yields the *state-dependent switching* paradigm, which is useful for modeling real systems that are either known or designed to have particular switching surfaces. State-dependent switching also provides a simple means of modeling ODEs with continuous but nondifferentiable right-hand sides, *without* extending dL's smooth term language semantics [21, 142, 143]. For example, an

ODE $x' = \max(e, \tilde{e})$ can be modeled as a system that switches between the ODEs $x' = e \,\&\, e \geq \tilde{e}$ and $x' = \tilde{e} \,\&\, e \leq \tilde{e}$ within the domains where the ODE RHS is dominated by e or $\tilde{e}$ respectively (see Section 6.5.1). For the finite family of ODEs with domains $x' = f_p(x) \,\&\, Q_p$, $p \in \mathcal{P}$, state-dependent switching is modeled as follows:

$$\alpha_{\texttt{state}} \equiv \Big(\bigcup_{p \in \mathcal{P}} x' = f_p(x) \,\&\, Q_p \Big)^* \tag{6.4}$$

Operationally, if the system is currently evolving in domain Q_i and is about to leave the domain, it must switch to another ODE with domain Q_j that is true in the current state to continue its evolution. Arbitrary switching $\alpha_{\texttt{arb}}$ is the special case of $\alpha_{\texttt{state}}$ with no domain restrictions ($Q_p \equiv \mathit{true}$ for all $p \in \mathcal{P}$). The following adequacy result generalizes Proposition 6.1 to consider only states that are reachable while obeying the specified domains.

Proposition 6.2. *A state is reachable by hybrid program α_{state} iff it is reachable in finite time by a switched system $x' = f_p(x)$ for $p \in \mathcal{P}$ following a switching signal σ while obeying the specified domains Q_p.*

Proof. Both directions of the proposition are proved separately for an initial state $\omega \in \mathbb{R}^n$ using the dL semantics of hybrid programs [135, 139, 142, 144] (see Section 2.2.2).

"⇒" Suppose $(\omega, \nu) \in [\![\alpha_{\texttt{state}}]\!]$. By the semantics of dL loops, there is a sequence of states $\omega = \omega_0, \omega_1, \dots, \omega_n = \nu$ for some $n \geq 0$ and for each $1 \leq i \leq n$, the states transition according to the loop body, i.e., $(\omega_{i-1}, \omega_i) \in [\![\bigcup_{p \in \mathcal{P}} x' = f_p(x) \,\&\, Q_p]\!]$. In particular, for each $1 \leq i \leq n$, there is a choice p_i where state ω_{i-1} reaches ω_i by following the ODE $x' = f_{p_i}(x)$ for some time $\zeta_i \geq 0$ and staying within the domain Q_{p_i} for all times $0 \leq t \leq \zeta_i$ during its evolution.
The finite sequences $(\omega_0, \omega_1, \dots, \omega_n)$, $(\zeta_1, \dots, \zeta_n)$ and $(p_1, \dots, p_n)$ correspond to a well-defined switching signal as follows. First, remove from all sequences the chattering indexes $1 \leq i \leq n$ with $\zeta_i = 0$. This yields new sequences $(\tilde{\omega}_0, \tilde{\omega}_1, \dots, \tilde{\omega}_m)$, $(\tilde{\zeta}_1, \dots, \tilde{\zeta}_m)$, and $(\tilde{p}_1, \dots, \tilde{p}_m)$ with times $\tilde{\zeta}_i > 0$. Consider the switching signal σ with switching times $\tau_i = \sum_{j=1}^{i} \tilde{\zeta}_j$ for $1 \leq i < m$ and $\tau_i = \tau_{i-1} + 1$ for $i \geq m$, so $\tau_1 < \tau_2 < \dots$ and $\tau_i \to \infty$. Furthermore, extend the sequence of switching choices with $\tilde{p}_i = \tilde{p}_m$ for $i > m$. By construction using (6.2), σ is well-defined and the solution φ associated with σ from ω reaches ν at time $\sum_{j=1}^{m} \tilde{\zeta}_j$ and obeys the domains $Q_{\tilde{p}_i}$ until that time.

"⇐" Let σ be a switching signal and $\varphi : [0, \zeta) \to \mathbb{R}^n$ be the associated switched system solution from ω. Suppose that the switched system reaches $\varphi(t)$ for $t \in [0, \zeta)$ while obeying the domains Q_p. To show $(\omega, \varphi(t)) \in [\![\alpha_{\texttt{state}}]\!]$, by the semantics of dL loops, it suffices to construct a sequence of states $\omega = \omega_0, \omega_1, \dots, \omega_n$ for some finite n, with $\omega_n = \varphi(t)$, and $(\omega_{i-1}, \omega_i) \in [\![\bigcup_{p \in \mathcal{P}} x' = f_p(x) \,\&\, Q_p]\!]$ for $1 \leq i \leq n$.
By (6.2), σ is equivalently defined by a sequence of switching times $\tau_0 < \tau_1 < \tau_2 < \dots$ and a sequence of switching choices $p_1, p_2, \dots$, where $p_i \in \mathcal{P}$. Let τ_n be the first switching time such that $t \leq \tau_n$; the index n exists since $\tau_i \to \infty$. Define the state sequence $\omega_i = \varphi(\tau_i)$ for $0 \leq i < n$ and $\omega_n = \varphi(t)$. Note that $\omega_0 = \omega$ by definition of $\varphi(0)$. It suffices to show $(\omega_{i-1}, \omega_i) \in [\![x' = f_{p_i}(x) \,\&\, Q_{p_i}]\!]$ for $1 \leq i \leq n$, but this follows by construction of

φ because ω_i is reached from ω_{i-1} by following the solution to ODE $x' = f_{p_i}(x)$, and, by assumption, $\varphi(\gamma)$ satisfies Q_{p_i} for $\gamma \in [\tau_{i-1}, \tau_i] \cap [0, t]$. □

Similar to $\alpha_{\texttt{arb}}$, Proposition 6.2 shows that dL's box and diamond modality formulas express safety (for all switching signals) and liveness (for some switching signal) properties of $\alpha_{\texttt{state}}$, respectively. The next result syntactically derives sound and complete invariance reasoning principles for state-dependent (and arbitrary) switching in dL.

Theorem 6.3. *The following axioms are derivable in* dL *where P and all ODE domains Q_p for $p \in \mathcal{P}$ are semianalytic formulas. Formulas $(\dot{Q_p})^{(*)}_{f_p}, (\dot{P})^{(*)}_{f_p}, (\dot{Q_p})^{(*)}_{-f_p}, (\dot{\neg P})^{(*)}_{-f_p}$ are the respective semianalytic progress formulas Def. 3.24 with respect to the indicated ODEs for mode $p \in \mathcal{P}$.*

$$\text{Inv}_{\text{state}} \quad \forall x\, (P \rightarrow [\alpha_{\texttt{state}}]P) \leftrightarrow \textstyle\bigwedge_{p\in\mathcal{P}} \forall x\, (P \rightarrow [x' = f_p(x) \,\&\, Q_p]P)$$

$$\text{SAI}_{\text{state}} \quad \forall x\, (P \rightarrow [\alpha_{\texttt{state}}]P) \leftrightarrow \textstyle\bigwedge_{p\in\mathcal{P}} \begin{pmatrix} \forall x\, \big(P \land Q_p \land (\dot{Q_p})^{(*)}_{f_p} \rightarrow (\dot{P})^{(*)}_{f_p}\big) \land \\ \forall x\, \big(\neg P \land Q_p \land (\dot{Q_p})^{(*)}_{-f_p} \rightarrow (\dot{\neg P})^{(*)}_{-f_p}\big) \end{pmatrix}$$

Proof Summary (Appendix D.1.1). Axiom $\text{SAI}_{\text{state}}$ derives immediately from $\text{Inv}_{\text{state}}$ by equivalently rewriting its RHS using the derived equivalent characterization of ODE invariants SAI& from Theorem A.11 (page 212). Both directions of the equivalence $\text{Inv}_{\text{state}}$ are derived separately. The "←" direction, uses dL's loop invariant rule to prove that P is a loop invariant of $\alpha_{\texttt{state}}$. The "→" direction shows that a run of ODE $x' = f_p(x) \,\&\, Q_p$, $p \in \mathcal{P}$ must also be a run of $\alpha_{\texttt{state}}$, so if formula P is true for all runs of $\alpha_{\texttt{state}}$, it must also be true for all runs of the ODEs. □

Axiom $\text{Inv}_{\text{state}}$ equivalently characterizes invariants of $\alpha_{\texttt{state}}$ with invariants of each of its constituent ODEs separately. Thus, when searching for an invariant of $\alpha_{\texttt{state}}$, it suffices to search for a *common* invariant of every constituent ODE. The invariance of any candidate (common) invariant P can be equivalently turned into an arithmetic question in dL by axiom $\text{SAI}_{\text{state}}$ which derives using dL's complete axiomatization for ODE invariance from Chapter 3. In the case where all ODEs in $\alpha_{\texttt{state}}$ and P are all described by polynomials, axiom $\text{SAI}_{\text{state}}$ shows that invariance for those state-dependent switching models is decidable because its RHS is a decidable formula of real arithmetic [14, 197]. Following Chapter 3, the completeness result (but not decidability) also applies to Noetherian functions, e.g., exponentials and trigonometric functions, that can be used to describe state-dependent switching mechanisms. Indeed, for Noetherian extensions, Corollary 3.39 shows that safety questions $[\alpha_{\texttt{state}}]P$ with analytic postconditions P are provably equivalent to arithmetic because $\alpha_{\texttt{state}}$ is assignment-free.

Modeling Subtleties

The model $\alpha_{\texttt{state}}$ as defined above makes no *a priori* assumptions about how the ODEs and their domains $x' = f_p(x) \,\&\, Q_p$ are designed, so results like Theorem 6.3 apply generally to all state-dependent switching designs. However, state-dependent switching can exhibit some well-known subtleties [99, 177] and it becomes the onus of modelers to appropriately account for these subtleties. This section examines various subtleties that can arise in $\alpha_{\texttt{state}}$ and prescribes sufficient arithmetical criteria for avoiding them; like Theorem 6.3, these arithmetical criteria are decidable for systems with polynomial terms [14, 197]. As a running example, let the line $x_1 = x_2$ be a *switching surface*, i.e., the example systems described below are intended to exhibit switching when their system state reaches this line.

Well-defined Switching. First, observe that the union of domains Q_p for $p \in \mathcal{P}$ must cover the entire state space; otherwise, there would be system states of interest where no continuous dynamics is active. This can be formally guaranteed by checking validity of the formula ①: $\bigvee_{p\in\mathcal{P}} Q_p$. Next, consider the following pair of ODEs (illustrated below, right):

$$\underbrace{x_1' = 0, x_2' = 1 \,\&\, x_1 \geq x_2}_{x'=f_A(x)\,\&\,Q_A \text{ in green}} \qquad \underbrace{x_1' = -1, x_2' = 0 \,\&\, x_1 < x_2}_{x'=f_B(x)\,\&\,Q_B \text{ in blue}}$$

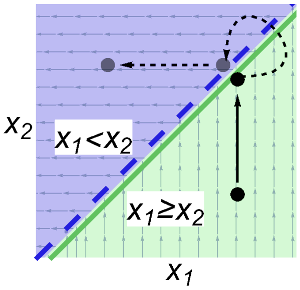

A system evolution starting in $Q_A \equiv x_1 \geq x_2$ is illustrated on the right. When the system reaches $x_1 = x_2$ (illustration offset for clarity), it is about to *locally progress* into $Q_B \equiv x_1 < x_2$ by switching to ODE $x' = f_B(x)$ but it gets stuck because it cannot make the infinitesimal jump from Q_A to enter Q_B. Augmenting the domain Q_B to $x_1 \leq x_2$ enables the switch. More generally, to avoid systems getting stuck on infinitesimal jumps, domains Q_p can be augmented to include states that locally progress into Q_p under the ODE $x' = f_p(x)$ and, symmetrically, states that locally exit Q_p [177]. Local progress (and exit) for ODEs is formalized using the local progress $\bigcirc$ modality introduced in Section 3.5: recall that formula $\langle x' = f(x) \,\&\, Q\rangle\bigcirc$ characterizes the states from which ODE $x' = f(x)$ locally progresses into Q; conversely, $\langle x' = -f(x) \,\&\, Q\rangle\bigcirc$ characterizes those from which the ODE locally exits Q. By the derived axiom LP from Theorem 3.26 (page 60), local progress and local exit are provably characterized by arithmetic formulas $(\dot{Q})^{(*)}_{f}$, $(\dot{Q})^{(*)}_{-f}$ respectively. To avoid the stuck states exemplified above for ODEs $x' = f_p(x) \,\&\, Q_p$, $p \in \mathcal{P}$ in $\alpha_{\mathtt{state}}$, it suffices to check validity of the formula ②: $(\dot{Q}_p)^{(*)}_{f_p} \lor (\dot{Q}_p)^{(*)}_{-f_p} \rightarrow Q_p$ for each $p \in \mathcal{P}$, which expresses that states locally entering or exiting Q_p for ODE $x' = f_p(x)$ are included in Q_p. Condition ② is syntactically simpler but equivalent to the domain augmentation presented in Sogokon et al. [177] for piecewise continuous models, a form of state-dependent switching.

Sliding Modes. The preceding subtlety arose from incomplete domain constraint specifications. Another subtlety that can arise because of incomplete specification of ODE dynamics, as exemplified by the following pair of ODEs (illustrated below, right):

$$\underbrace{x_1' = 0, x_2' = 1 \,\&\, x_1 \geq x_2}_{x'=f_A(x)\,\&\,Q_A \text{ in green}} \qquad \underbrace{x_1' = 1, x_2' = 0 \,\&\, x_1 \leq x_2}_{x'=f_B(x)\,\&\,Q_B \text{ in blue}}$$

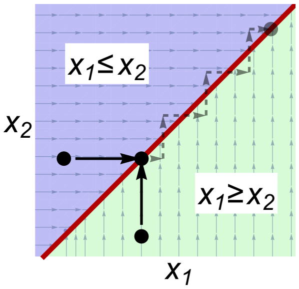

Systems starting in $Q_A \equiv x_1 \geq x_2$ or $Q_B \equiv x_1 \leq x_2$ eventually reach the line $x_1 = x_2$ (in red) but they then get stuck because the ODEs on either side of $x_1 = x_2$ drive system evolution onto the line. Mathematically, the system enters a *sliding mode* [99] along $x_1 = x_2$. This can be thought of as infinitely fast switching between the ODEs that results in a new sliding dynamics *along* the switching surface $x_1 = x_2$ (dashed grey trajectory). When the sliding dynamics can be calculated exactly, it suffices to add those dynamics explicitly to the switched system, e.g., adding the sliding dynamics $x_1' = \frac{1}{2}, x_2' = \frac{1}{2} \,\&\, x_1 = x_2$ to the example above allows stuck system states on $x_1 = x_2$ to continuously progress along the line (illustrated below, left). An alternative is *hysteresis switching* [99] which enlarges domains adjacent to the sliding mode so that a system that reaches the sliding surface is allowed to briefly continue following its current dynamics before switching. For example, for a fixed $\varepsilon > 0$, the enlarged domains $Q_A \equiv x_1 \geq x_2 - \varepsilon$ and $Q_B \equiv x_1 \leq x_2 + \varepsilon$ allows the stuck states to evolve off the line for a short distance $\varepsilon > 0$. This yields arbitrary switching in the overlapped part of both domains (illustrated below, right). For domains $Q_p, p \in \mathcal{P}$ meeting conditions ① and ②, hysteresis switching is modeled by replacing each Q_p with its closed ε-neighborhood for $\varepsilon > 0$. Another way of modeling hysteresis using an auxiliary memory variable to remember the current system mode is shown in the next section.

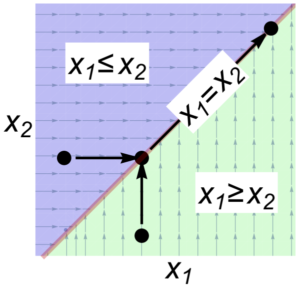

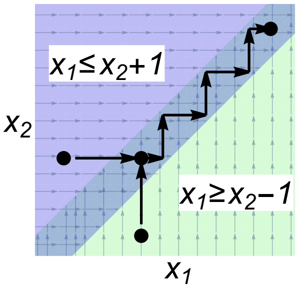

Explicit sliding: $x_1' = 0, x_2' = 1 \,\&\, x_1 \geq x_2$

$x_1' = 1, x_2' = 0 \,\&\, x_1 \leq x_2$

$x_1' = \dfrac{1}{2}, x_2' = \dfrac{1}{2} \,\&\, x_1 = x_2$

Hysteresis: $x_1' = 0, x_2' = 1 \,\&\, x_1 \geq x_2 - 1$

$x_1' = 1, x_2' = 0 \,\&\, x_1 \leq x_2 + 1$

Both approaches can be used (and can be mixed) in the hybrid program model $\alpha_{\texttt{state}}$. To guarantee the absence of stuck states, it suffices to check validity of the formula ③: $\bigvee_{p\in\mathcal{P}} (\dot{Q}_p)^{(*)}_{f_p}$, i.e., every point in the state space can switch to an ODE which locally progresses in its associated domain. Models meeting conditions ② and ③ also meet condition ①.

Zeno Behavior. Hybrid and switched system models can also exhibit *Zeno behavior*, where the model makes infinitely many discrete transitions in a finite time interval [99, 214], see also [144, Expedition 9.1]. Such behaviors are an artifact of the model and are not reflective of the real world. As such, Zeno traces are typically excluded when reasoning about hybrid system models [99, 214], e.g., all switching signals considered in this section (Section 6.2.1) are assumed to be well-defined (thus non-Zeno) and Proposition 6.2 specifies safety for all *finite* executions of state-dependent switching. The detection of Zeno behavior in switched systems is left out of scope for this thesis and the (upcoming) specification of switched system stability in Section 6.3 implicitly excludes them from consideration

6.2.3 Controlled Switching

This section turns to *controlled switching* models, where an explicit controller program is responsible for making logical switching decisions between the ODEs $x' = f_p(x), p \in \mathcal{P}$. The discrete fragment of hybrid programs can be used to flexibly model (computable) switching logics, e.g., those that combine state-dependent and time-dependent switching constraints, or make complex switching decisions based on the state of the system. Controlled switching is modeled by the hybrid program $\alpha_{\texttt{ctrl}}$ in (6.5):

$$\alpha_{\texttt{ctrl}} \equiv \underset{\substack{\downarrow\\ \text{initialization}}}{\alpha_i}; \Big(\overset{\substack{\text{switching controller}\\ \uparrow}}{\alpha_u}; \overbrace{\bigcup_{p\in\mathcal{P}} \big(?u = p; x' = f_p(x,y), y' = g_p(x,y) \,\&\, Q_p\big)}^{\alpha_p \text{ (plant, actuate decision)}}\Big)^* \tag{6.5}$$

The model $\alpha_{\texttt{ctrl}}$ resembles the shape of standard models of event-triggered and time-triggered systems in dL [144] but is adapted for controlled switching. It uses three subprograms: α_i initializes the system, then α_u (modeling the switching controller) and α_p (modeling the continuous plant dynamics) are run in a switching loop. The discrete programs α_i, α_u decide on values for the control output $u = p, p \in \mathcal{P}$ and the plant program α_p responds to this output by evolving the corresponding ODE $x' = f_p(x,y), y' = g_p(x,y) \,\&\, Q_p$. The programs α_i, α_u must not modify the system state variables x, but they may modify other auxiliaries, including *auxiliary continuous state* variables y used to model timers or integral terms used in controllers, see Section 6.5.2. This control-plant loop is a typical structure for hybrid systems modeled in dL [135, 144]. As an example, the controller α_u below models the discrete switching logic present in hybrid automata [28, 75, 135], without discrete jumps in the system state:

$$\begin{aligned} \alpha_u &\equiv \bigcup_{p\in\mathcal{P}} \Big(?u = p; (\bigcup_{q\in\mathcal{P}} (?G_{p,q}; R_{p,q}; u := q) \cup u := u)\Big) \\ R_{p,q} &\equiv y_1 := e_1; y_2 := e_2; \dots; y_k := e_k \end{aligned} \tag{6.6}$$

For each mode $p \in \mathcal{P}$, the switching controller may nondeterministically switch to mode $q \in \mathcal{P}$ if the *guard* formula $G_{p,q}$ (with free variables x, y) is true in the current state. By default, the controller can trivially choose to stay in the current mode with $u := u$. If the transition is taken, the *reset map* $R_{p,q}$ sets the values of auxiliary state variables $y_1, \dots, y_k$ respectively to the value of terms $e_1, \dots, e_k$. Two important classes of switching mechanisms are modeled as special instances of $\alpha_{\texttt{ctrl}}$ next.

Guarded State-Dependent Switching

The instance $\alpha_{\texttt{guard}}$ corresponds to the automata controller from (6.6) with $\alpha_i \equiv \bigcup_{p\in\mathcal{P}} u := p$ and guard formulas $G_{p,q} \equiv G_{p,q}(x)$. It does not use auxiliaries y nor the reset map $R_{p,q}$.

$$\alpha_{\texttt{guard}} \equiv \begin{cases} \alpha_i \equiv \bigcup\limits_{p\in\mathcal{P}} u := p \\ \alpha_u \equiv \bigcup\limits_{p\in\mathcal{P}} \Big(?u = p; (\bigcup\limits_{q\in\mathcal{P}} (?G_{p,q}; u := q) \cup u := u)\Big) \end{cases} \tag{6.7}$$

The model $\alpha_{\texttt{guard}}$ adds a form of *hysteresis* [86] to the state-dependent switching model from Section 6.2.2, so that switching decisions at each guard $G_{p,q}$ depend explicitly on memory of the current discrete mode u in addition to the continuous state. For simplicity, the system is initialized by α_i to start in any mode $p \in \mathcal{P}$, although this can be modified for different applications, e.g., if the system has known initial mode(s).

Proposition 6.4. *A state is reachable by hybrid program α_{guard} iff it is reachable in finite time by a switched system $x' = f_p(x)$ for $p \in \mathcal{P}$ following a switching signal σ while obeying the specified domains Q_p and **guards** $G_{p,q}$ for modes $p, q \in \mathcal{P}$.*

Proof Summary (Appendix D.1.1). The proof is similar to Proposition 6.2; in the "⇒" direction, the proof constructs a suitable switching signal from the execution of hybrid program $\alpha_{\texttt{guard}}$; in the "⇐" direction, the proof shows that a given switching signal corresponds to a run of $\alpha_{\texttt{guard}}$. The main difference in both directions is to show that the fresh auxiliary variable u used to control the switching signal produces the intended hysteresis effect(s). □

The model $\alpha_{\texttt{piece}}$ from (5.5) in Section 5.5 (page 135) similarly uses an auxiliary variable u to model a piecewise continuous function. The deferred proof of adequacy for $\alpha_{\texttt{piece}}$ (Proposition 5.27, page 136) is given in Appendix D.1.1.

Time-Dependent Switching

The instance $\alpha_{\texttt{time}}$ shown below models *time-dependent switching*, where switching decisions are based on the time elapsed in each mode. Here, time is tracked by an auxiliary timer variable τ with $\tau' = 1$ added to each ODE.

$$\alpha_{\texttt{time}} \equiv \begin{cases} \alpha_i \equiv \tau := 0; \bigcup\limits_{p \in \mathcal{P}} u := p \\ \alpha_u \equiv \bigcup\limits_{p \in \mathcal{P}} \Big(?u = p; (\bigcup\limits_{q \in \mathcal{P}} \big(?\theta_{p,q} \le \tau; \tau := 0; u := q \big) \cup u := u) \Big) \\ \alpha_p \equiv \bigcup\limits_{p \in \mathcal{P}} \big(?u = p; x' = f_p(x), \tau' = 1 \,\&\, \tau \le \Theta_p \big) \end{cases} \tag{6.8}$$

The controller program α_u enables switching from mode $p \in \mathcal{P}$ to $q \in \mathcal{P}$ after a *minimum* dwell time $0 \le \theta_{p,q} \le \tau$ has elapsed and resets the timer whenever such a switch occurs. Conversely, the plant α_p restricts modes with a *maximum* dwell time $\tau \le \Theta_p$ for $\Theta_p > 0, p \in \mathcal{P}$; an unbounded dwell time $\Theta_p = \infty$ is represented by the domain constraint true. A special case of $\alpha_{\texttt{time}}$ is *slow switching*, where the system is allowed to switch arbitrarily between ODEs but there is a global minimum dwell time $\theta > 0$ where $\theta_{p,q} = \theta$ for all $p, q \in \mathcal{P}$. A sufficiently large minimum dwell time θ can be used to stabilize switching between any family of stable linear ODEs [99]. Dwell time restrictions can also be used to stabilize systems that switch between stable *and unstable* modes [211]. Intuitively, the system should stay in stable modes for sufficient duration ($\theta_{p,q} \le \tau$) while it should avoid staying in unstable modes for too long ($\tau \le \Theta_p$).

Proposition 6.5. *A state is reachable by hybrid program α_{time} iff it is reachable in finite time by a switched system $x' = f_p(x)$ for $p \in \mathcal{P}$ following a switching signal σ that spends **at least** time $\theta_{p,q} \ge 0$ for each switch $p \in \mathcal{P}$ to $q \in \mathcal{P}$ along the solution and **at most** time $\Theta_p > 0$ for each mode $p \in \mathcal{P}$ entered along the solution.*

Proof in Appendix D.1.1.

Building on Propositions 6.1–6.5, the next section turns to the study of switched system stability in dL through their hybrid program models.

6.3 Switched System Stability

This section explains how stability for hybrid program models of switched systems is formally specified and verified using dL. Similar to ODEs, various stability notions are of interest in the continuous and hybrid systems literature [61, 66, 99, 126, 151, 189]. Stability variations for switched systems can also be formally specified (see Chapter 5) but this chapter focuses on proving one form of stability (UGpAS, defined below) for variations of switching mechanisms.

6.3.1 Stability as Quantified Loop Safety

This section studies *uniform global pre-asymptotic stability* (UGpAS) for switched systems [65, 66, 99], defined as follows:

Definition 6.6 (UGpAS [65, 66])**.** Let $\Phi(x)$ denote the set of all (domain-obeying) solutions $\varphi : [0, T_\varphi] \to \mathbb{R}^n$ for a switched system from state $x \in \mathbb{R}^n$. The origin $0 \in \mathbb{R}^n$ is:

- **uniformly stable** if, for all $\varepsilon > 0$, there exists $\delta > 0$ such that from all initial states $x \in \mathbb{R}^n$ with $\|x\|_2 < \delta$, all solutions $\varphi \in \Phi(x)$ satisfy $\|\varphi(t)\|_2 < \varepsilon$ for all times $0 \leq t \leq T_\varphi$,
- **uniformly globally pre-attractive** if, for all $\varepsilon > 0, \delta > 0$, there exists $T \geq 0$ such that from all initial states $x \in \mathbb{R}^n$ with $\|x\|_2 < \delta$, all solutions $\varphi \in \Phi(x)$ satisfy $\|\varphi(t)\|_2 < \varepsilon$ for all times $T \leq t \leq T_\varphi$, and
- **uniformly globally pre-asymptotically stable** if the system is uniformly stable and uniformly globally pre-attractive.

The UGpAS definition can be understood intuitively for a system with a given switching control mechanism analogously to stability for ODEs (e.g., Section 5.2.1 from page 120):

- *stability* means the mechanism keeps the system close to the origin if the system is initially perturbed close to the origin,
- *global pre-attractivity* means the mechanism drives the system to the origin asymptotically as $t \to \infty$, and
- *uniform* means the stability and pre-attractivity properties are independent of both the nondeterminism in the switching mechanism (e.g., arbitrary switching) and the choice of initial states satisfying $\|x\|_2 < \delta$; for brevity in subsequent sections, “uniform” is elided when describing stability properties.

Remark 6.7. Switched systems whose solutions are all uniformly bounded in time, i.e., there exists T_m such that for all solutions φ, $T_\varphi \leq T_m$, are trivially pre-attractive. Goebel et al. [65, 66] introduce the notion of *pre-attractivity* as opposed to *attractivity* for hybrid systems because it separates considerations about whether a hybrid system’s solutions are *complete*, i.e., solutions exist for all (forward) time, from conditions for stability and attractivity. Pre-attractivity also sidesteps the difficult question of whether a switched system exhibits Zeno behavior (recall modeling subtleties in Section 6.2.2) [99, 214]. Indeed, it is common in the hybrid and switched systems literature to either *ignore* incomplete solutions or *assume* the models under consideration only have complete solutions [99, 116, 214]. Instead of predicating proofs on these hypotheses, this chapter formalizes the (weaker) notion of UGpAS for switched systems, leaving proofs of completeness of solutions out of scope.

The definition of UGpAS nests alternating quantification over real numbers with temporal quantification over the solutions φ of switched systems. Accordingly, UGpAS for switched systems is formally specified by nesting dL's box modality with the first-order quantifiers; when program α models a switched system, the box modality $[\alpha](\cdot)$ quantifies (uniformly) over all times for all switching signals arising from the switching mechanism.

Notational Conventions (Norm Bounds). For notational simplicity in this chapter, Euclidean norm bound formulas are directly written with $\|x\|_2 \sim \varepsilon \overset{\text{def}}{\equiv} (\sum_{i=1}^{n} x_i^2) \sim \varepsilon^2$ (for $\varepsilon \geq 0$) for comparison operators $\sim \in \{=, \neq, \geq, >, \leq, <\}$.

Lemma 6.8 (UGpAS in differential dynamic logic). *The origin $0 \in \mathbb{R}^n$ for a switched system modeled by hybrid program α is UGpAS iff the* dL *formula* $\mathrm{UGpAS}(\alpha)$ *is valid. Variables $\varepsilon, \delta, T, t$ are fresh in α:*

$$\begin{aligned}
\mathrm{UStab}(\alpha) &\equiv \forall \varepsilon{>}0\, \exists \delta{>}0\, \forall x\, \big(\|x\|_2 < \delta \rightarrow [\alpha]\ \|x\|_2 < \varepsilon \big) \\
\mathrm{UGpAttr}(\alpha) &\equiv \forall \varepsilon{>}0\, \forall \delta{>}0\, \exists T{\geq}0\, \forall x\, \big(\|x\|_2 < \delta \rightarrow [t := 0; \alpha, t' = 1]\, (t \geq T \rightarrow \|x\|_2 < \varepsilon) \big) \\
\mathrm{UGpAS}(\alpha) &\equiv \mathrm{UStab}(\alpha) \wedge \mathrm{UGpAttr}(\alpha)
\end{aligned}$$

Here, $\mathrm{UStab}(\alpha)$ *and* $\mathrm{UGpAttr}(\alpha)$ *characterize stability and global pre-attractivity of α, respectively. In* $\mathrm{UGpAttr}(\alpha)$*, $\alpha, t' = 1$ denotes the hybrid program obtained from α by augmenting its continuous dynamics so that variable t tracks the progression of time.*

Proof. Let $\Phi(x)$ be the set of all domain-obeying solutions $\varphi : [0, T_\varphi] \rightarrow \mathbb{R}^n$ for a given switched system from state $x \in \mathbb{R}^n$ as in Def. 6.6. Hybrid program α *models* the given switched system if, for any initial state $\omega \in \mathbb{R}^n$, the state ν is reachable from state ω, i.e., $(\omega, \nu) \in [\![\alpha]\!]$ by dL's hybrid program semantics [142, 144], iff $\nu = \varphi(\tau)$ for some $\varphi \in \Phi(\omega)$ and $\tau \in [0, T_\varphi]$. For the augmented program $\alpha, t' = 1$, in particular, t syntactically tracks the progression of time so that $(\omega, \nu) \in [\![\alpha, t' = 1]\!]$ iff $\nu = \varphi(\tau)$ for some $\varphi \in \Phi(\omega)$ and $\tau = \nu(t) - \omega(t)$. The adequacy of looping hybrid program models for several switching mechanisms is proved in Section 6.2.

The formulas $\mathrm{UStab}(\alpha)$ and $\mathrm{UGpAttr}(\alpha)$ syntactically express their respective quantifiers from Def. 6.6, where the box modality $[\cdot]$ is used in both formulas to quantify over all reachable states of α (and $\alpha, t' = 1$), i.e., all times $\tau \in [0, T_\varphi]$ along all solutions $\varphi \in \Phi$. Thus, the correctness of these specifications follows directly from the definition of dL's formula semantics [142, 144]. In $\mathrm{UGpAttr}(\alpha)$, the clock variable t is set to 0 initially and has ODE $t' = 1$ so it tracks the progression of time along the continuous evolution of program α. The implication $t \geq T \rightarrow \dots$ in the postcondition of the box modality restricts temporal quantification to all times τ with $T \leq \tau \leq T_\varphi$ for all solutions $\varphi \in \Phi(\omega)$ for uniform pre-attractivity. □

Formulas $\mathrm{UStab}(\alpha)$ and $\mathrm{UGpAttr}(\alpha)$ syntactically formalize in dL the corresponding quantifiers in Def. 6.6. In $\mathrm{UGpAttr}(\alpha)$, the fresh clock variable t is initialized to 0 and syntactically tracks the progression of time along switched system solutions. The program $\alpha, t' = 1$ can, e.g., be constructed by adding a clock ODE $t' = 1$ to all ODEs in the switched system model α. Accordingly, the postcondition $t \geq T \rightarrow \|x\|_2 < \varepsilon$ expresses that the system state norm is bounded by ε after T time units along any switching trajectory, as required in Def. 6.6. The key (derived) dL proof rule used to prove UGpAS specifications is the loop rule [144], recalled below.

$$\text{loop}\ \frac{\Gamma \vdash \mathit{Inv} \quad \mathit{Inv} \vdash [\alpha]\,\mathit{Inv} \quad \mathit{Inv} \vdash \phi}{\Gamma \vdash [\alpha^*]\phi}$$

The loop rule says that, in order to prove validity of the conclusion (below the rule bar), it suffices to prove the three premises (above the rule bar), respectively from left to right: *i)* the initial assumptions Γ imply *Inv*, *ii) Inv* is preserved across the loop body α, i.e., *Inv* is a *loop invariant* for α^*, and *iii) Inv* implies the postcondition ϕ. The identification of loop invariants is crucial for formal proofs of UGpAS for *looping* models of switched systems, as illustrated by the following deductive proof skeleton for stability (a similar skeleton is used for pre-attractivity):

$$\dfrac{\dfrac{\text{loop}\ \dfrac{\dfrac{*}{\Gamma \vdash \mathit{Inv}} \quad \dfrac{\dfrac{\Gamma_1 \vdash \phi_1 \quad \cdots \quad \Gamma_k \vdash \phi_k}{\vdots \left(\begin{array}{l}\text{hybrid} \quad \text{program}\\ \text{reasoning for } \alpha\end{array}\right)}}{\mathit{Inv} \vdash [\alpha]\,\mathit{Inv}} \quad \dfrac{*}{\mathit{Inv} \vdash \|x\|_2 < \varepsilon}}{\Gamma \vdash [\alpha^*]\ \|x\|_2 < \varepsilon}}{\vdots \left(\begin{array}{l}\text{logic/arithmetic}\\ \text{reasoning for } \Gamma\end{array}\right)}}{\vdash \mathrm{UStab}(\alpha^*)}$$

The proof skeleton above syntactically *derives* a proof rule that reduces a stability proof for α^* to proofs of its top-most premises, $\Gamma_1 \vdash \phi_1 \ \cdots \ \Gamma_k \vdash \phi_k$. These correspond to required logical and arithmetical conditions on Lyapunov functions for various switching mechanisms. The choice of loop invariant (highlighted in red) crucially ties together these arithmetic conditions on Lyapunov functions with hybrid program reasoning for switched systems. Throughout this section, loop invariants are progressively tweaked to account for new design insights behind increasingly complex switching mechanisms from Section 6.2.

6.3.2 Stability for Autonomous Switching

This section identifies loop invariants for proving UGpAS under autonomous switching mechanisms with Lyapunov functions [27, 89, 99]; relevant mathematical arguments are presented briefly, see Appendix D.1.2 for more details.

Arbitrary Switching

Stability for the arbitrary switching model $\alpha_{\mathtt{arb}}$ from (6.3) in Section 6.2.2 can be verified by finding a so-called *common Lyapunov function* V for all of the ODEs $x' = f_p(x), p \in \mathcal{P}$ satisfying the following arithmetical conditions [99, 189]:

i) $V(0) = 0$ and $V(x) > 0$ for all $\|x\|_2 > 0$,

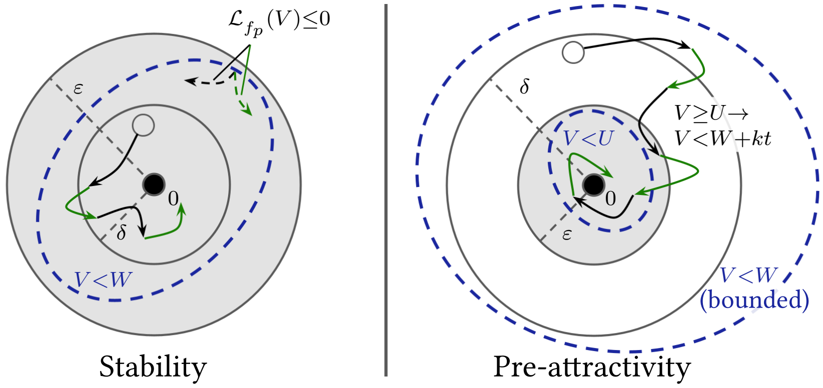

Figure 6.3: Loop invariants for UGpAS (arbitrary switching), stability (left) and pre-attractivity (right). Switching trajectories are illustrated by alternating black and green arrows.

ii) V is *radially unbounded*, i.e., for all b, there exists $\gamma > 0$ such that $\|x\|_2 < \gamma$ for all $V(x) \le b$, and

iii) for each ODE $x' = f_p(x), p \in \mathcal{P}$, the Lie derivative $\mathcal{L}_{f_p}(V)$ satisfies: $\mathcal{L}_{f_p}(V)(0) = 0$ and $\mathcal{L}_{f_p}(V)(x) < 0$ for all $\|x\|_2 > 0$.

Conditions i)–iii) are generalizations of well-known conditions for stability of ODEs [33, 89] to arbitrary switching. Intuitively, conditions i) and iii) ensure that V acts as an auxiliary energy function whose value decreases asymptotically to zero (at the origin) along all switching trajectories of the system; the radial unboundedness condition ii) ensures that this argument applies to all system states for *global* pre-attractivity [89]. Correctness of these conditions can be proved in dL using loop invariants, see Fig. 6.3 (explained below).

Stability. The specification $\text{UStab}(\alpha_{\texttt{arb}})$ requires that all trajectories of $\alpha_{\texttt{arb}}$ stay in the grey ball $\|x\|_2 < \varepsilon$, starting from a chosen ball $\|x\|_2 < \delta$, see Fig. 6.3 (left). Condition i) guarantees that the ball $\|x\|_2 < \varepsilon$ contains (a connected component of) the sublevel set $V{<}W$ for some $W{>}0$ (dashed blue curve) and this sublevel set contains a smaller ball $\|x\|_2 < \delta$ [33, 89]. Condition iii) shows that this sublevel set is invariant for each ODE $x' = f_p(x), p \in \mathcal{P}$ because $\mathcal{L}_{f_p}(V)(x) \le 0$, illustrated by the dashed black and green arrows for two switching choices $p \in \mathcal{P}$ both locally pointing inwards on the boundary of the sublevel set. Formula $\mathit{Inv}_s \equiv \|x\|_2 < \varepsilon \land V < W$, which characterizes the blue sublevel set, is an invariant for all possible switching choices in the loop body of $\alpha_{\texttt{arb}}$. Thus, Inv_s is a suitable loop invariant for $\text{UStab}(\alpha_{\texttt{arb}})$.

Pre-Attractivity. The specification $\text{UGpAttr}(\alpha_{\texttt{arb}})$ requires that all trajectories of $\alpha_{\texttt{arb}}$ stay in the grey ball $\|x\|_2 < \varepsilon$ after a chosen time T, starting from the initial ball $\|x\|_2 < \delta$, see Fig. 6.3 (right). The ball $\|x\|_2 < \delta$ is bounded, so it is contained in a sublevel set satisfying $V < W$ for some $W > 0$ (outer dashed blue curve); this sublevel set is bounded by condition ii). Like the stability argument, condition i) guarantees that there is a sublevel set $V < U$ for some $U > 0$ (inner dashed blue curve) contained in the ball $\|x\|_2 < \varepsilon$, and condition iii) shows that the sublevel sets characterized by $V < W$ and $V < U$ are both invariants for every ODE in the loop body of $\alpha_{\texttt{arb}}$. The set characterized by formula $V \ge U \land V \le W$ is compact and bounded away from the origin, which implies by condition iii) that there is a uniform bound $k < 0$ on this set, where for each ODE $x' = f_p(x), p \in \mathcal{P}$, $\mathcal{L}_{f_p}(V)(x) \leq k$. Thus, the value of Lyapunov function V decreases at rate k, regardless of switching choices in the loop body of $\alpha_{\texttt{arb}}$, *as long as* it has not entered $V < U$. The loop invariant for $\mathrm{UGpAttr}(\alpha_{\texttt{arb}})$ syntactically expresses this intuition: $Inv_a \equiv V < W \land (V \geq U \rightarrow V < W + kt)$. For sufficiently large T with $W + kT \leq U$, trajectories at time $t \geq T$ satisfy $V < U$ so they are contained in the $\|x\|_2 < \varepsilon$ ball.

The loop invariants identified above enable derivation of a formal dL stability proof rule for $\alpha_{\texttt{arb}}$ (deferred to a more general version for $\alpha_{\texttt{state}}$ in Corollary 6.9 below). In fact, since arbitrary switching is the most permissive form of switching [99], UGpAS for any switching mechanism can be soundly justified using the loop invariants above in case a suitable common Lyapunov function can be found (see Table 6.1, common Lyapunov function column).

State-Dependent Switching

State-dependent switching is modeled by hybrid program $\alpha_{\texttt{state}}$ from (6.4) in Section 6.2.2. The same loop invariants for $\alpha_{\texttt{arb}}$ are used for $\alpha_{\texttt{state}}$ to derive the following proof rule. For brevity, premises of all derived stability proof rules are implicitly conjunctively quantified over $p \in \mathcal{P}$.

Corollary 6.9 (UGpAS for state-dependent switching, CLF)**.** *The following proof rule for common Lyapunov function V with three stacked premises is derivable in* dL.

$$\text{CLF}\ \frac{\begin{array}{l}\vdash V(0) = 0 \land \forall x\, (\|x\|_2 > 0 \rightarrow V(x) > 0)\\ \vdash \forall b\, \exists \gamma\, \forall x\, (V(x) \leq b \rightarrow \|x\|_2 \leq \gamma)\\ \vdash \mathcal{L}_{f_p}(V)(0) = 0 \land \forall x\, (\|x\|_2 > 0 \land \overline{Q_p} \rightarrow \mathcal{L}_{f_p}(V)(x) < 0)\end{array}}{\vdash \mathrm{UGpAS}(\alpha_{\texttt{state}})}\quad \text{(for all } p \in \mathcal{P}\text{)}$$

Proof in Appendix D.1.2.

Corollary 6.9 syntactically derives a slight generalization of conditions i)–iii) from Section 6.3.2 for $\alpha_{\texttt{state}}$, where the Lie derivatives $\mathcal{L}_{f_p}(V)(x)$ for each $p \in \mathcal{P}$ are required to be negative on their respective domain closures[1] $\overline{Q_p}$. This generalization is justified by the same loop invariants explained in Section 6.3.2 because the ODE invariance properties are only required to hold in their respective domains.

The domain asymmetry in $\alpha_{\texttt{state}}$ suggests another way of generalizing the stability arguments, namely, through the use of *multiple Lyapunov functions*, where a (possibly) different Lyapunov function V_p is associated to each mode $p \in \mathcal{P}$ [27]. Here, the function V_p is responsible for justifying stability within domain Q_p, i.e., its value decreases along system trajectories whenever the system is within Q_p, as illustrated in Fig. 6.4. Constraints on these functions are obtained by modifying the loop invariants to account for this intuition.

Stability. The stability loop invariant is modified by case splitting disjunctively on the domains $Q_p, p \in \mathcal{P}$, and requiring that the sublevel set characterized by $V_p < W$ is invariant within its respective domain Q_p: $Inv_s \equiv \|x\|_2 < \varepsilon \land \bigvee_{p \in \mathcal{P}} \big(Q_p \land V_p < W\big)$. Like Section 6.3.2, the bound W is chosen so that each sublevel set characterized by $V_p < W$ is contained in the ball $\|x\|_2 < \varepsilon$.

[1] The topological closure $\overline{Q}$ of domain Q is needed for soundness of a technical compactness argument used in the pre-attractivity proof, see Appendix D.1.2.

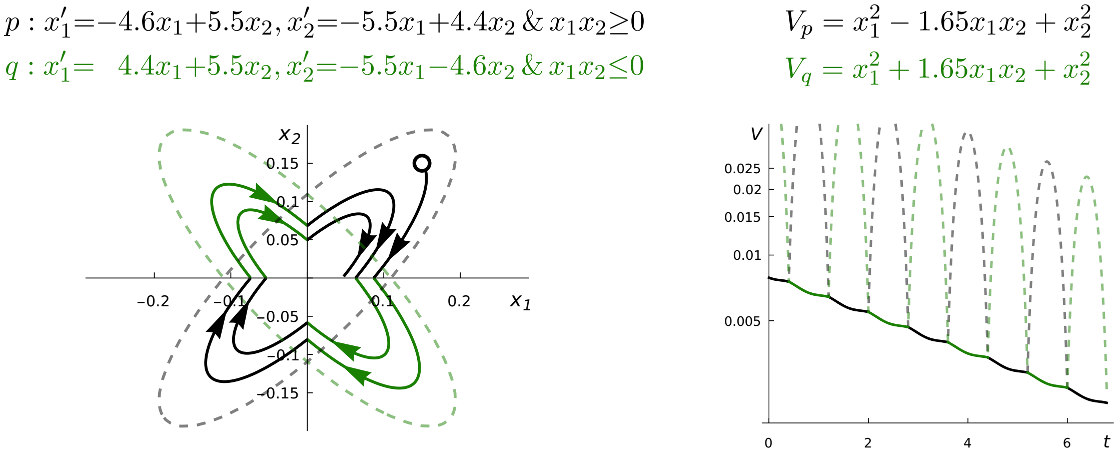

Figure 6.4: A switching trajectory for Example 7 from Section 6.4.2 with state-dependent switching (left) and the value of two Lyapunov functions along that trajectory (right, log-scale on vertical axis). Solid lines indicate the active Lyapunov function at time t. Two sublevel sets $V_p, V_q < W = 0.012$ are shown dashed on the left within which the switching trajectory is respectively trapped at any given time.

Pre-Attractivity. The pre-attractivity loop invariant is similarly modified by disjunctively requiring that each V_p decreases along system trajectories when the system is in their respective domains Q_p: $\mathit{Inv}_a \equiv \bigvee_{p\in\mathcal{P}} \big(Q_p \land V_p < W \land (V_p \geq U \rightarrow V_p < W + kt)\big)$. The constants U, W, k, T are chosen as appropriate lower or upper bounds for all the Lyapunov functions (see proof of Corollary 6.10). Arithmetical conditions for the Lyapunov functions $V_p, p \in \mathcal{P}$ are derived from the modified invariants in the following rule.

Corollary 6.10 (UGpAS for state-dependent switching, MLF)**.** *The following proof rule for multiple Lyapunov functions $V_p, p \in \mathcal{P}$ with four stacked premises is derivable in* dL*.*

$$\text{MLF}\;\frac{\begin{array}{l}\vdash V_p(0) = 0 \land \forall x\,(\|x\|_2 > 0 \rightarrow V_p(x) > 0)\\ \vdash \forall b\,\exists\gamma\,\forall x\,(V_p(x) \leq b \rightarrow \|x\|_2 \leq \gamma)\\ \vdash \mathcal{L}_{f_p}(V_p)(0) = 0 \land \forall x\,(\|x\|_2 {>} 0 \land \overline{Q_p} \rightarrow \mathcal{L}_{f_p}(V_p)(x){<}0)\\ \vdash \bigwedge_{q\in\mathcal{P}} \Big(Q_p \land Q_q \rightarrow V_p = V_q\Big)\end{array}}{\vdash \text{UGpAS}(\alpha_{\texttt{state}})}\quad(\text{for all } p \in \mathcal{P})$$

Proof in Appendix D.1.2.

The top three premises of Corollary 6.10 are similar to those of Corollary 6.9, but are now required to hold for each Lyapunov function $V_p, p \in \mathcal{P}$ separately. The (new) bottom premise corresponds to a compatibility condition between the Lyapunov functions arising from the loop invariants. For example, consider the stability loop invariant (similarly for pre-attractivity) and suppose the system currently satisfies disjunct $Q_p \land V_p < W$ with V_p justifying stability in domain Q_p. If the system switches to the ODE $x' = f_q(x)$ within domain Q_q, then Lyapunov function V_q becomes the active Lyapunov function which must satisfy $V_q < W$ to preserve the stability loop invariant. The premise $Q_p \land Q_q \rightarrow V_p = V_q$ says that the Lyapunov functions V_p, V_q take on equal values whenever such a switch is possible (in either direction), i.e., when their domains Q_p, Q_q overlap. This is illustrated in Fig. 6.4 (right), where the plotted Lyapunov functions values are equal whenever switching occurs (black to green, or vice versa).

6.3.3 Stability for Controlled Switching

Stability analysis for controlled switching proceeds by identifying suitable loop invariants *Inv* for $\alpha_{\texttt{ctrl}}$ from (6.5) in Section 6.2.3. *Compositional reasoning* [135, 144] allows for separate analysis of the discrete (α_i, α_u) and continuous (α_p) dynamics and then lifting those results to the full hybrid dynamics. This idea is exemplified by the following derived variation of the loop rule:

$$\text{loopT}\ \frac{\Gamma \vdash [\alpha_i]\,\mathit{Inv} \quad \mathit{Inv} \vdash [\alpha_u]\,\mathit{Inv} \quad \mathit{Inv} \vdash [\alpha_p]\,\mathit{Inv} \quad \mathit{Inv} \vdash \phi}{\Gamma \vdash [\alpha_i; (\alpha_u; \alpha_p)^*]\phi}$$

Rule loopT says that loop invariant *Inv* is maintained *throughout* $\alpha_{\texttt{ctrl}} \equiv \alpha_i; (\alpha_u; \alpha_p)^*$, i.e., system initialization α_i puts the system into a state satisfying the invariant *Inv* and *Inv* is compositionally preserved by *both* the discrete switching logic α_u and the continuous dynamics α_p. This rule is used to analyze the two instances of $\alpha_{\texttt{ctrl}}$ introduced in Section 6.2.3 next.

Guarded State-Dependent Switching

The instance $\alpha_{\texttt{guard}}$ from (6.7) adds hysteresis to state-dependent switching through the auxiliary memory variable u which tracks the current mode of the system. This design change is reflected in the loop invariants and in the corresponding proof rule below.

Stability. The stability loop invariant is modified (cf. Section 6.3.2) to case split on the possible discrete modes $u = p$ rather than the ODE domains: $\mathit{Inv}_s \equiv \|x\|_2 < \varepsilon \land \bigvee_{p\in\mathcal{P}} \big(u = p \land V_p < W\big)$.

Pre-Attractivity. The pre-attractivity loop invariant is modified similarly to case split on the discrete modes, with: $\mathit{Inv}_a \equiv \bigvee_{p\in\mathcal{P}} \big(u = p \land V_p{<}W \land (V_p \geq U \rightarrow V_p < W + kt)\big)$.

Corollary 6.11 (UGpAS for guarded state-dependent switching, MLF)**.** *The following proof rule for multiple Lyapunov functions $V_p, p \in \mathcal{P}$ with four stacked premises is derivable in* dL.

$$\text{MLF}_G\ \frac{\begin{array}{l} \vdash V_p(0) = 0 \land \forall x\,(\|x\|_2 > 0 \rightarrow V_p(x) > 0) \\ \vdash \forall b\,\exists\gamma\,\forall x\,(V_p(x) \leq b \rightarrow \|x\|_2 \leq \gamma) \\ \vdash \mathcal{L}_{f_p}(V_p)(0) = 0 \land \forall x\,(\|x\|_2{>}0 \land \overline{Q_p} \rightarrow \mathcal{L}_{f_p}(V_p)(x){<}0) \\ \vdash \bigwedge_{q\in\mathcal{P}} \Big(G_{p,q} \rightarrow V_q \leq V_p\Big) \end{array}}{\vdash \text{UGpAS}(\alpha_{\texttt{guard}})} \quad (\text{for all } p \in \mathcal{P})$$

Proof in Appendix D.1.2.

The premises of rule MLF_G are identical to those from MLF except the bottom premise, which derives from loopT and unfolding the controller α_u with dL's hybrid program axioms, e.g., the following proof skeleton shows the unfolding for the stability loop invariant Inv_s corresponding to a switch from mode p to mode q:

$$
\underset{\textbf{Unfold}}{\uparrow}\quad
{\scriptstyle[\cup]}\dfrac{{\scriptstyle[;],\,[?],\,[:=]}\dfrac{{\scriptstyle\mathbb{R}}\dfrac{\vdash G_{p,q} \rightarrow V_q \leq V_p}{V_p < W \vdash G_{p,q} \rightarrow V_q < W}}{u = p \land V_p < W \vdash [?G_{p,q}; u := q](u = q \land V_q < W)}}{\mathit{Inv}_s \vdash [\alpha_u]\mathit{Inv}_s}
\quad \overset{\textbf{Arithmetic}}{\uparrow}
$$

Unlike rule MLF, the bottom premise of rule MLF$_G$ only uses an inequality, because the guards $G_{p,q}$ determine permissible switching.

Time-Dependent Switching

To reason about stability for instance $\alpha_{\mathtt{time}}$ from (6.8), consider Lyapunov function conditions $\mathcal{L}_{f_p}(V_p)(x) \leq -\lambda_p V_p$, where λ_p is a constant associated with each mode $p \in \mathcal{P}$. This condition bounds the value of V_p along the solution of $x' = f_p(x)$ by either a decaying exponential for stable modes ($\lambda_p > 0$) or a growing exponential for unstable modes ($\lambda_p \leq 0$). Let $\mathcal{S} = \{p \in \mathcal{P}, \lambda_p > 0\}$ and $\mathcal{U} = \{p \in \mathcal{P}, \lambda_p \leq 0\}$ be the indexes of the stable and unstable modes in the loop invariants below and assume an extended term language (Section 3.2) where $\exp(\cdot)$ denotes the real exponential function.

Stability. The stability loop invariant expresses the required exponential bounds with a case split depending if $p \in \mathcal{S}$ or $p \in \mathcal{U}$:

$$
\mathit{Inv}_s \equiv \tau \geq 0 \land \|x\|_2 < \varepsilon \land \left(\begin{array}{l} \bigvee\limits_{p \in \mathcal{S}} \big(u = p \land V_p < W \exp(-\lambda_p \tau)\big) \lor \\ \bigvee\limits_{p \in \mathcal{U}} \big(u = p \land V_p < W \exp(-\lambda_p(\tau - \Theta_p)) \land \tau \leq \Theta_p\big) \end{array} \right)
$$

For $p \in \mathcal{S}$, $\exp(-\lambda_p \tau)$ is the accumulated decay factor for V_p after staying in the stable mode for time τ. For $p \in \mathcal{U}$, $\exp(-\lambda_p(\tau - \Theta_p))$ is a buffer factor for the growth of V_p in the unstable mode so that $V_p < W$ still holds at the maximum dwell time $\tau = \Theta_p$. In both cases, the internal timer variable is non-negative ($\tau \geq 0$).

Pre-Attractivity. The pre-attractivity loop invariant has similar exponential decay and growth bounds for each $p \in \mathcal{P}$ in the current mode. In addition, it has an overall exponential decay term $\exp(-\sigma(t - \tau))$ for some $\sigma > 0$, which ensures that the value of V_p tends to 0 as $t \rightarrow \infty$ for all switching trajectories; recall t is the global clock introduced in the specification of pre-attractivity in Lemma 6.8 while τ is the auxiliary timer used to track time in the model $\alpha_{\mathtt{time}}$.

$$
\begin{aligned}
\mathit{Inv}_a \equiv\ & \tau \geq 0 \land t \geq \tau \land \\
& \left(\begin{array}{l} \bigvee\limits_{p \in \mathcal{S}} \big(u = p \land V_p < W \exp(-\sigma(t - \tau)) \exp(-\lambda_p \tau)\big) \lor \\ \bigvee\limits_{p \in \mathcal{U}} \big(u = p \land V_p < W \exp(-\sigma(t - \tau)) \exp(-\lambda_p(\tau - \Theta_p)) \land \tau \leq \Theta_p\big) \end{array} \right)
\end{aligned}
$$

Term $\exp(-\sigma(t-\tau))$ in both $p \in \mathcal{S},\mathcal{U}$ cases is the accumulated *overall* decay factor for V_p *until* the switch to mode p at time $t-\tau$. Term $\exp(-\lambda_p\tau)$ (resp. $\exp(-\lambda_p(\tau-\Theta_p))$) is the *current* decay (resp. growth) factor *since* the switch to mode $p \in \mathcal{S}$ (resp. $p \in \mathcal{U}$).

Corollary 6.12 (UGpAS for time-dependent switching, MLF)**.** *The following proof rule for multiple Lyapunov functions $V_p, p \in \mathcal{P}$ with five stacked premises is derivable in* dL*.*

$$\mathrm{MLF}_\tau \quad \dfrac{\begin{array}{l}\vdash V_p(0)=0 \land \forall x\,(\|x\|_2 > 0 \rightarrow V_p(x) > 0)\\ \vdash \forall b\,\exists\gamma\,\forall x\,(V_p(x) \le b \rightarrow \|x\|_2 \le \gamma)\\ \vdash \mathcal{L}_{f_p}(V_p) \le -\lambda_p V_p\\ \color{red}{\mathit{Inv}_s \vdash [\alpha_u]\mathit{Inv}_s \qquad \mathit{Inv}_a \vdash [\alpha_u]\mathit{Inv}_a}\end{array}}{\vdash \mathrm{UGpAS}(\alpha_{\mathtt{time}})} \quad \text{(for all } p \in \mathcal{P}\text{)}$$

The two red premises on the bottom row are expanded to arithmetical conditions on V_p by unfolding the program structure of α_u with dL *axioms in Appendix D.1.2.*

Proof in Appendix D.1.2.

The bottom premises of MLF_τ and MLF_G exemplify a key benefit of dL stability reasoning: conditions on V_p that arise from *Inv*$_s$, *Inv*$_a$ are derived by systematically unfolding the discrete dynamics of α_u with sound dL axioms. This enables automatic, *correct-by-construction* derivation of those conditions, which is especially important for controlled switching because the number of possible transitions scales quadratically $|\mathcal{P}|^2$ with the number of modes $|\mathcal{P}|$.

6.4 KeYmaera X Implementation

This section presents a prototype implementation of a switched system modeling and proof package for KeYmaera X [54]. The implementation consists of ≈2700 lines and, crucially, does not require any extension to KeYmaera X's existing soundness-critical core. Accordingly, verification results for switched systems obtained through this implementation directly inherit the strong correctness properties guaranteed by the design of KeYmaera X [54, 115].

6.4.1 Modeling and Proof Interface

The implementation extends KeYmaera X's proof IDE [114] with a convenient interface for modeling switching mechanisms, as shown in Fig. 6.5. The interface allows users to express switching mechanisms intuitively by rendering automaton plots while abstracting away the underlying hybrid programs. It provides templates for switched systems following the switching mechanisms of Section 6.2: state-dependent (autonomous), guarded, timed, and general controlled switching (tabs "Autonomous", "Guarded", "Timed", "Generic" in Fig. 6.5). From these templates, KeYmaera X automatically generates programs and stability specifications, ensuring that they have the correct dL hybrid program and formula structure.

Switched systems are represented internally with a common interface `SwitchedSystem` which is implemented by four classes: `StateDependent` $\alpha_{\mathtt{state}}$, `Guarded` $\alpha_{\mathtt{guard}}$, `Timed` $\alpha_{\mathtt{time}}$, and `Controlled` $\alpha_{\mathtt{ctrl}}$. The `SwitchedSystem` interface provides default stability and pre-attractivity specifications, which can be adapted by users on the UI if needed. Corollaries 6.9–6.12

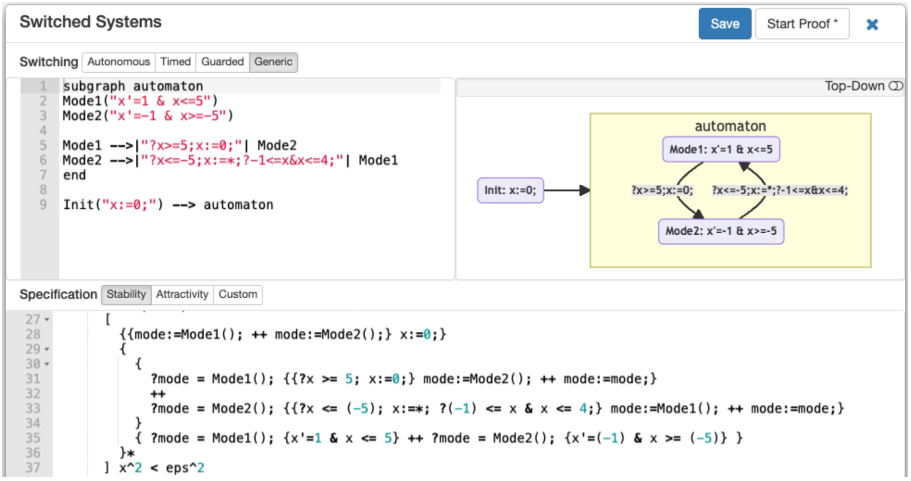

```
subgraph automaton
Mode1("x'=1 & x<=5")
Mode2("x'=-1 & x>=-5")

Mode1 -->|"?x>=5;x:=0;"| Mode2
Mode2 -->|"?x<=-5;x:=*;?-1<=x&x<=4;"| Mode1
end

Init("x:=0;") --> automaton
```

```
        [
          {{mode:=Mode1(); ++ mode:=Mode2();} x:=0;}
          {
            {
              ?mode = Mode1(); {{?x >= 5; x:=0;} mode:=Mode2(); ++ mode:=mode;}
              ++
              ?mode = Mode2(); {{?x <= (-5); x:=*; ?(-1) <= x & x <= 4;} mode:=Mode1(); ++ mode:=mode;}
            }
            { ?mode = Mode1(); {x'=1 & x <= 5} ++ ?mode = Mode2(); {x'=(-1) & x >= (-5)} }
          }*
        ] x^2 < eps^2
```

Figure 6.5: Screenshot of the KeYmaera X switched systems editor: automata input on top-left, rendered automaton top-right, generated hybrid program and specification(s) in dL at the bottom.

Table 6.1: Available tactics in KeYmaera X for switched systems stability proofs and Lyapunov function generation.

`SwitchedSystem`	Common Lyap. Proof	Common Lyap. Gen.	Multiple Lyap. Proof	Multiple Lyap. Gen.
`StateDependent` $\alpha_{\texttt{state}}$	✓	✓	✓	✓
`Guarded` $\alpha_{\texttt{guard}}$	✓	✓	✓	✓
`Timed` $\alpha_{\texttt{time}}$	✓	✓	✓	—
`Controlled` $\alpha_{\texttt{ctrl}}$	✓	✓	—	—

are implemented as UGpAS *proof tactics* in KeYmaera X's Bellerophon tactic language [55]. These tactics automate all of the reasoning steps underlying stability proofs for their respective switching mechanisms, so that users only need to input candidate Lyapunov functions for KeYmaera X to (attempt to) complete their proofs. Additionally, when candidates are not provided by the user, the implementation uses sum-of-squares programming [129, 154] to automatically generate candidate Lyapunov functions for a subset of switching designs. The generated candidates are checked for correctness by KeYmaera X so the generator does not need to be trusted for correctness of the resulting proofs. Table 6.1 summarizes the available proof tactics and Lyapunov function generation for classes of switching mechanisms.

6.4.2 Examples

The implementation is tested on a suite of examples drawn from the literature [27, 86, 154, 189] featuring various switching mechanisms.[2] Table 6.2 summarizes the proof statistics, where all experiments were run on a MacBook Pro 2019 with Intel Core i7 (6-core, 2.6GHz) and 32GB

[2] See `https://github.com/LS-Lab/KeYmaeraX-projects/blob/master/stability/UGpAS` (including case studies from Section 6.5). Git hash: `c856fddb383232adbd86679ef65567f9b90190bf`

Table 6.2: Stability proofs for examples drawn from the literature. The "Time" columns indicate time (in seconds, rounded to 1 d.p.) to run the KeYmaera X proofs for stability (Stab.) and attractivity (Attr.), × indicates incomplete proof. A ✓ in the "Gen." column indicates successful Lyapunov function(s) generation, ? indicates that a candidate was generated but with numerical issues, and — indicates inapplicability. In the latter two cases (?, —) known Lyapunov functions from the literature were used for the proofs (if available).

Example	Model	Time (Stab.)	Time (Attr.)	Gen.
1 [27, Ex. 2.1]	$\alpha_{\texttt{state}}$	2.6	3.0	✓
2 [86, Motiv. ex.]	$\alpha_{\texttt{state}}$	2.2	2.3	✓
3 [86, Ex. 1]	$\alpha_{\texttt{state}}$	3.3	4.1	✓
4 [86, Ex. 2 & 3]	$\alpha_{\texttt{guard}}$	2.8	3.8	?
5 [154, Ex. 6]	$\alpha_{\texttt{guard}}$	×	×	?
6 [189, Ex. 2.45]	$\alpha_{\texttt{arb}}$	19.4	11.1	✓
7 [189, Ex. 3.25]	$\alpha_{\texttt{state}}$	2.4	2.9	✓
8 [189, Ex. 3.49]	$\alpha_{\texttt{time}}$	4.4	5.6	—
9 [211, Ex. 1]	$\alpha_{\texttt{time}}$	4.7	5.3	—
10 [211, Ex. 2]	$\alpha_{\texttt{time}}$	256.9	×	—

memory. All examples have a 2 dimensional state space and switch between 2 modes except Example 4 (2 dimensions, 4 modes) and Example 6 (3 dimensions, 2 modes). The proof tactics successfully prove most of the examples across various switching mechanisms. For Example 5, a suitable Lyapunov function (without numerical errors) could not be found.[3] For the time-dependent switching models (Examples 8–10), KeYmaera X internally uses verified polynomial Taylor approximations to the exponential function for decidability of arithmetic [14, 197]; Example 10 needs a high degree approximation (15 terms in the polynomial) for sufficient accuracy and its attractivity proof could not be completed in reasonable time. Overall, this suite of examples shows the feasibility of (fully) automated stability verification for various classes of switched systems in KeYmaera X. Future work could improve internal arithmetic reasoning steps (e.g., Examples 8–10) to shorten proof times.

6.5 Case Studies

This section presents three case studies applying the deductive verification approach to justify various non-standard switched system stability arguments in KeYmaera X.

6.5.1 Canonical Max System

Branicky [26] investigates the longitudinal dynamics of an aircraft with an elevator controller that mediates between two control objectives: *i)* tracking potentially unsafe pilot input and

[3] Prajna and Papachristodoulou [154, Ex. 6] report that a sextic (degree 6) Lyapunov function can be generated with sum-of-squares techniques, but do not provide the generated function explicitly.

ii) respecting safety constraints on the aircraft's angle of attack. Assuming a state feedback control law, the model is transformed to the following *canonical max system* [26, Remark 5], with state variables x, y and parameters a, b, f, g, γ satisfying $a, b, a - f, b - g > 0$ and $\gamma \leq 0$.

$$x' = y, y' = -ax - by + \max(fx + gy + \gamma, 0) \tag{6.9}$$

The right-hand side of system (6.9) is non-differentiable but the equations can be equivalently rewritten with state-dependent switching between a family of two ODEs corresponding to either possibility for the $\max(fx + gy + \gamma, 0)$ term in the equation for y' as follows, where the system follows ODE Ⓐ in domain $fx + gy + \gamma \leq 0$ and ODE Ⓑ in domain $fx + gy + \gamma \geq 0$.

$$\textcircled{A} \equiv x' = y, y' = -ax - by \qquad \textcircled{B} \equiv x' = y, y' = -(a - f)x - (b - g)y + \gamma$$

Stability of this parametric system is *not* directly provable using standard techniques for state-dependent switching presented in Section 6.3.2. For example, the ODE Ⓐ stabilizes the system to the origin but the ODE Ⓑ stabilizes to the point $(-\frac{\gamma}{a-f}, 0)$, away from the origin for $\gamma < 0$. Instead, Branicky [26] proves global asymptotic stability of (6.9) with the following "noncustomary" [43] Lyapunov function involving a nondifferentiable integrand:

$$V = \frac{1}{2}y^2 + \int_0^x a\xi - \max(f\xi + \gamma, 0)d\xi \tag{6.10}$$

The key idea used to deductively prove stability here is *ghost switching*: analogous to ghost variables in program verification which are added for the sake of program proofs [127, 144] or in proofs for ODEs (Chapters 3 and 4), ghost switching modes do not change the physical dynamics of the system but are introduced for the purposes of the stability analysis. Here, ghost switching between $fx + \gamma \leq 0$ and $fx + \gamma \geq 0$ is used to obtain closed form representations for the integral in (6.10). This yields an instance of state-dependent switching $\alpha_{\texttt{state}}$ with 4 switching modes and the stability specification P_m:

$$\begin{aligned}
&\alpha_m \equiv \left(\textcircled{A}_1 \cup \textcircled{A}_2 \cup \textcircled{B}_1 \cup \textcircled{B}_2\right)^* \quad p \equiv fx + gy + \gamma \quad q \equiv fx + \gamma\\
&\textcircled{A}_1 \equiv \textcircled{A} \,\&\, p \leq 0 \land q \leq 0 \qquad \textcircled{A}_2 \equiv \textcircled{A} \,\&\, p \leq 0 \land q \geq 0\\
&\textcircled{B}_1 \equiv \textcircled{B} \,\&\, p \geq 0 \land q \leq 0 \qquad \textcircled{B}_2 \equiv \textcircled{B} \,\&\, p \geq 0 \land q \geq 0\\
&P_m \equiv a{>}0 \land b{>}0 \land a{-}f{>}0 \land b{-}g{>}0 \land f{\neq}0 \land \gamma{\leq}0 \rightarrow \mathrm{UGpAS}(\alpha_m)
\end{aligned}$$

The ghost switching modes enable a multiple Lyapunov function argument for stability using the following modified closed-form representations of Branicky's Lyapunov function (6.10), with V_1 for $\textcircled{A}_1, \textcircled{B}_1$ and V_2 for $\textcircled{A}_2, \textcircled{B}_2$.[4]

$$\begin{aligned}
V_1(x, y) &= \frac{1}{2}(bcx^2 + 2cxy + y^2) + \overbrace{\frac{a}{2}x^2}^{\int_0^x a\xi - \max(f\xi+\gamma,0)d\xi \text{ where } f\xi + \gamma \leq 0}\\
V_2(x, y) &= \frac{1}{2}(bcx^2 + 2cxy + y^2) + \underbrace{\frac{a}{2}x^2 - \frac{(fx + \gamma)^2}{2f}}_{\int_0^x a\xi - \max(f\xi+\gamma,0)d\xi \text{ where } f\xi + \gamma \geq 0}
\end{aligned}$$

[4] An important technical requirement for V_2 to be well-defined is $f \neq 0$. The case with $f = 0$ is also verified in KeYmaera X but the details are omitted here for brevity. It does not require ghost switching and uses only V_1 as its common Lyapunov function.

The Lyapunov functions V_1, V_2 are also modified from (6.10) to use a quadratic form with an additional constant c satisfying constraints $0 < c < b, c < b - g, c < \frac{(a-f)(b-g)}{a-f+g^2}, c < \frac{a(b-g)}{a+g^2}$ (such a constant always exists under the assumptions on a, b, f, g). This technical modification is required to prove UGpAS for α_m directly with the Lyapunov functions.

Another challenging aspect of this case study is verification of the *parametric* arithmetical conditions for V_1, V_2, i.e., stability is verified for *all* possible parameter values a, b, f, g, γ that satisfy the assumptions in P_m. Such questions are decidable in theory [14, 197], but are difficult for automated solvers in practice (even out of reach of solvers that require numerically bounded parameters [60]). KeYmaera X enables a user-aided proof of the required arithmetic conditions. For example, the Lie derivative of the Lyapunov function V_1 for Ⓑ$_1$ is given by:

$$\dot{V_1} = \mathcal{L}_{x'=y,y'=-(a-f)x-(b-g)y+\gamma}(V_1) = -(b-c)y^2 - acx^2 + (cx+y)(fx+gy+\gamma)$$

Here, $\dot{V_1}$ is required to be strictly negative away from the origin for stability. The arithmetical argument uses domain constraint $fx + gy + \gamma \geq 0 \land fx + \gamma \leq 0$ from Ⓑ$_1$ as follows: if $cx + y \leq 0$, then by constraint $fx + gy + \gamma \geq 0$, $\dot{V_1}$ satisfies $\dot{V_1} \leq -(b-c)y^2 - acx^2$. Otherwise, $cx + y > 0$, then by constraint $fx + \gamma \leq 0$, $\dot{V_1}$ satisfies $\dot{V_1} \leq -(b-g-c)y^2 - acx^2 + gcxy$. In either case, the RHS bound is a negative definite quadratic form by the earlier choice of parameter c and therefore, $\dot{V_1}$ is negative away from the origin. The verification of P_m and all of the required arithmetic reasoning is done in KeYmaera X.

6.5.2 Automated Cruise Control

Oehlerking [126, Sect. 4.6] verifies the stability of an automatic cruise controller modeled as a hybrid automaton with 6 operating modes and 11 transitions between them: normal proportional-integral (PI) control, acceleration, service braking (2 modes), and emergency braking (2 modes). Figure 6.6 shows an abridged version of the corresponding KeYmaera X model (using $\alpha_{\texttt{ctrl}}$) with the PI control mode, where v is the relative velocity to be controlled to $v = 0$ and x, t are auxiliary integral and timer variables used in the controller. Briefly, this controller is designed to use the PI controller near $v = 0$ for stability, while its other control modes drive the system toward $v = 0$ by accelerating or braking.

Lyapunov function candidates for this model can be successfully generated using the Stabhyli [116] stability tool for hybrid automata, but Stabhyli (with default configurations) outputs a Lyapunov function candidate for the PI control mode that is numerically unsound, see Appendix D.2 for the output and a counterexample; this is a known issue with Stabhyli for control modes at the origin [116]. For this case study, the issue is manually resolved by truncating terms with very small magnitude coefficients in the generated output and then checking in KeYmaera X that the arithmetical conditions for the PI mode are satisfied for the truncated candidate.

Insights from the controller design are used in the UGpAS proof in KeYmaera X. Since stability only concerns states and modes that are active near the origin, the stability argument only needs to mention a single Lyapunov function for the PI control mode, while choosing δ (in Def. 6.6) sufficiently small so that none of the other modes can be entered. In fact, the PI controller equations are exactly those of a linearized pendulum, which has known Lyapunov functions [89]

```
normalPI("v’ = -0.001*x-0.052*v, x’ = v, t’ = 0
        & -15 <= v & v <= 15
        & -500 <= x & x <= 500")
normalPI -->|"?(13 <= v & v <= 15 &
            -500 <= x & x <= 500);
            t := 0;"| sbrakeact
normalPI -->|"?(-15 <= v & v <= -14 &
            -500 <= x & x <= 500);"| accelerate
... // Other modes
```

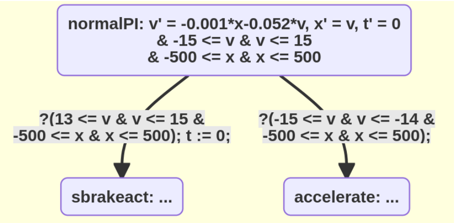

```
\forall eps ( eps > 0 -> // Abridged stability specification
  ...
  [
  ... // Initialization
  {
     { // Switching controller
       ... ++ // Transitions for other modes
       ?mode = normalPI();
       { {?13 <= v & v <= 15 & -500 <= x & x <= 500; t := 0;}
         mode := sbrakeact(); ++
         ?-15 <= v & v <= -14 & -500 <= x & x <= 500;
         mode := accelerate(); ++
         mode := mode; }
     }
     { // Plant
       ...  ++ // Plant ODEs for other modes
       ?mode = normalPI();
       { v’ = -0.001*x-0.052*v, x’ = v, t’ = 0 &
         -15 <= v & v <= 15 & -500 <= x & x <= 500 }
     }
  }*  // Switching loop
  ] v^2 < eps^2
```

Figure 6.6: Snippets of an automated cruise controller [126] modeled as a (switching) hybrid automaton. Users express the automaton within the description language (top left) and KeYmaera X visualizes the automaton on-the-fly (top left). The implementation automatically generates the appropriate hybrid program representation and UGpAS specification (bottom); `++`, `&`, `()` denote choice, conjunction, and constants in KeYmaera X’s ASCII syntax respectively.

(see Example 5.7, page 124)—it could be interesting to modify Stabhyli to accept user-provided Lyapunov function hints for certain modes. Similarly, pre-attractivity only requires reasoning about *asymptotic* convergence to the origin for the PI control mode so it suffices to show that the system leaves all other modes in finite time.

6.5.3 Brockett's Nonholonomic Integrator

Verification of stabilizing control laws for Brockett's nonholonomic integrator [29] is of significant interest because stability for a large class of models can be reduced to that of the integrator via coordinate transformations, e.g., Liberzon [99] transforms a unicycle model to the integrator and provides a stabilizing switching control law corresponding to parking of the unicycle. The nonholonomic integrator is described by the following system of differential equations with state variables x, y, z and state feedback control inputs $u = u(x, y, z), v = v(x, y, z)$ (to be determined further below).

$$x' = u, y' = v, z' = xv - yu$$

Notably, this is a classical example of a system that is not stabilizable by purely continuous feedback control. Intuitively, no choice of controls u, v can produce motion along the z-axis ($x = y = 0$). Thus, to stabilize the system to the origin, the controller must first drive the system away from the z-axis before switching to a control law that stabilizes the system from states away from the z-axis. This intuition can be realized using two different switching strategies that are analogous to the event-triggered and time-triggered CPS design paradigms respectively [144].

Event-Triggered Controller

Bloch and Drakunov [13] use the switching controller $u = -x + ay\operatorname{sign}(z), v = -y - ax\operatorname{sign}(z)$ to asymptotically stabilize the integrator in the region $\frac{a}{2}(x^2 + y^2) \geq |z|$ for any given constant $a > 0$. This controller first drives the system towards the plane $z = 0$ and, once it reaches the plane, *slides* along the plane towards the origin. The closed-loop system is modeled as an instance of state-dependent switching $\alpha_{\mathtt{state}}$ with 3 modes depending on the sign of z and specification P_e:

$$\begin{aligned}
&\textcircled{A} \equiv x' = -x + ay, y' = -y - ax, z' = -a(x^2 + y^2) \,\&\, z \geq 0\\
&\textcircled{B} \equiv x' = -x - ay, y' = -y + ax, z' = a(x^2 + y^2) \,\&\, z \leq 0\\
&\textcircled{C} \equiv x' = -x, y' = -y, z' = 0 \,\&\, z = 0 \quad \alpha_e \equiv \left(\textcircled{A} \cup \textcircled{B} \cup \textcircled{C}\right)^*\\
&P_e \equiv a > 0 \rightarrow \mathrm{UStab}(\alpha_e) \land\\
&\quad \forall \delta{>}0\, \forall \varepsilon{>}0\, \exists T{\geq}0\, \forall x, y, z \left(\|x, y, z\|_2 < \delta \land \frac{a}{2}(x^2 + y^2) \geq |z| \rightarrow \right.\\
&\qquad\qquad\qquad \left. [t := 0; \alpha_e, t' = 1](t \geq T \rightarrow \|x, y, z\|_2 < \varepsilon \right)
\end{aligned}$$

The specification P_e is identical to UGpAS except it restricts pre-attractivity to the applicable region $\frac{a}{2}(x^2 + y^2) \geq |z|$ for Bloch and Drakunov [13]'s controller.[5] It is verified using the squared norm term $V = x^2 + y^2 + z^2$ as a common Lyapunov function for $\textcircled{A}$–$\textcircled{C}$. The key modification to the pre-attractivity proof, cf. Section 6.3.2, is to use (and verify) the fact that $\frac{a}{2}(x^2 + y^2) \geq |z|$ is a loop invariant of α_e. This additional invariant corresponds to the fact that the controller keeps the system within its applicable region (if the system is initially within that region). In

[5] The applicable region is equivalently characterized by the real arithmetic formula $(z \geq 0 \rightarrow \frac{a}{2}(x^2 + y^2) \geq z) \land (z \leq 0 \rightarrow \frac{a}{2}(x^2 + y^2) \geq -z)$, omitted for brevity.

fact, α_e can be extended to a globally stabilizing controller, as modeled by $\alpha_{\hat{e}}$ below (`if`, `else` branching is supported as an abbreviation in KeYmaera X [144]):

$$\textcircled{D} \equiv x' = u, y' = v, z' = xv - yu \,\&\, \frac{a}{2}(x^2+y^2) \le |z|$$

$$\textcircled{E} \equiv x' = u, y' = v, z' = xv - yu \,\&\, \frac{a}{2}(x^2+y^2) \ge |z|$$

$$\alpha_{\hat{e}} \equiv \begin{pmatrix} \texttt{if}\Big(\frac{a}{2}(x^2+y^2) \ge |z|\Big)\ \{\textcircled{A} \cup \textcircled{B} \cup \textcircled{C}\} \\ \texttt{else}\ \Big\{ \\ \quad \texttt{if}((x-y)z \le 0)\{u := c; v := c\}\ \texttt{else}\{u := -c; v := -c\}; \\ \quad \{\textcircled{D} \cup \textcircled{E}\} \quad \Big\} \end{pmatrix}^*$$

$$P_{\hat{e}} \equiv a > 0 \land c > 0 \rightarrow \mathrm{UGpAS}(\alpha_{\hat{e}})$$

If the system is in the applicable region (outer `if` branch), then the previous controller from α_e is used. Otherwise, outside the applicable region (outer `else` branch), the system applies a constant control $c > 0$ chosen to drive the system into the applicable region. The pair of ODEs Ⓓ and Ⓔ model an event-trigger in dL [144], where the switching controller is triggered to make its next decision when the system reaches the switching surface $\frac{a}{2}(x^2+y^2) = |z|$.

The specification $P_{\hat{e}}$ is proved by modifying the loop invariants to account for an initial period where the system is outside the applicable region. For example, the stability loop invariant $\mathit{Inv}_s \equiv (\neg\frac{a}{2}(x^2+y^2) \ge |z| \rightarrow |z| < \delta) \land (\frac{a}{2}(x^2+y^2) \ge |z| \rightarrow \|x,y,z\|_2 < \varepsilon)$ expresses that the controller keeps $|z|$ sufficiently small with $|z| < \delta$ to preserve stability outside the applicable region. The pre-attractivity loop invariant is similarly split between the two cases, with an explicit time estimate on the time it takes for the system to enter the applicable region.

Time-Triggered Controller

The time-triggered switching strategy [144], modeled by α_τ below, is similar to that proposed by Liberzon [99, Section 4.2]. If the system is on the z-axis and away from the origin Ⓐ, the controller sets an internal stopwatch τ and drives the system away from the axis for maximum duration $T_0 > 0$ with $u = z, v = z$. Otherwise Ⓑ, the controller drives the system towards the origin along a parabolic curve of the form $\frac{a}{2}(x^2+y^2) = z$.

$$\alpha_\tau \equiv \begin{pmatrix} \texttt{if}(x = 0 \land y = 0 \land z \ne 0)\ \big\{\tau := 0; x' = z, y' = z, z' = xz - yz \,\&\, \tau \le T_0\big\} & \textcircled{A} \\ \texttt{else}\ \Big\{a := \dfrac{2z}{x^2+y^2}; x' = -x + ay, y' = -y - ax, z' = -a(x^2+y^2)\Big\} & \textcircled{B} \end{pmatrix}^*$$

$$P_\tau \equiv T_0 > 0 \rightarrow \mathrm{UGpAS}(\alpha_\tau)$$

The specification P_τ is again proved by analyzing both cases of the controller in the loop invariants, e.g., with the pre-attractivity invariant Inv_a:

$$\big(x = 0 \land y = 0 \land z \ne 0 \rightarrow |z| < \delta \land t = 0\big) \land$$
$$\big(\neg(x = 0 \land y = 0 \land z \ne 0) \rightarrow \|x,y,z\|_2 > \varepsilon \rightarrow \|x,y,z\|_2^2 < \delta^2(2T_0^2+1) - \varepsilon^2(t - T_0)\big)$$

The top conjunct says the system may start transiently on the z-axis (with $z \neq 0$) at time $t = 0$. The bottom conjunct gives explicit bounds on $\|x, y, z\|_2$, which, for sufficiently large $t \geq T$, implies that the system enters $\|x, y, z\|_2 < \varepsilon$ as required for pre-attractivity. The transient term $\delta^2(2T_0^2 + 1)$ upper bounds the (squared) norm of the system state after starting on the z-axis in ball $\|x, y, z\|_2 < \delta$ and following mode Ⓐ for the maximum stopwatch duration $\tau = T_0$.

6.6 Related Work

Hybrid System Formalisms. There are numerous hybrid system formalisms in the literature [65, 66, 72, 75, 99, 102, 135, 144, 164, 189]; see the cited articles and textbooks for further references. Connections between several formalisms have been examined in prior work. Platzer [135] shows how hybrid automata can be embedded into hybrid programs for their safety verification; dL can also be generalized with (disjunctive) differential-algebraic constraints that can be used to model and verify continuous dynamics with state-dependent switching [135, Chapter 3]. This chapter models switching with discrete program operators which enables compositional reasoning for the hybrid dynamics in switched systems. Sogokon et al. [177] study hybrid automata models for ODEs with piecewise continuous right-hand sides and highlight various subtleties in the resulting models; similar subtleties for state-dependent switching models are presented in Section 6.2.2. Goebel et al. [65, 66] show how impulsive differential equations, hybrid automata, and switched systems can all be understood as hybrid time models, and derive their properties using this connection. The decidability of invariance for state-dependent switched systems is proved in Theorem 6.3 using their dL hybrid program models.

Switched Systems in Control. Comprehensive introductions to the analysis and design of switching control can be found in the literature [43, 99, 189]. An important design consideration (which this chapter sidesteps, cf. Remark 6.7) is whether a given switched or hybrid system has complete solutions [65, 66, 107, 214]. Justification of such design considerations, and other stability notions of interest for switching designs, e.g., quadratic, region, or set-based stability [65, 66, 99, 151, 189], can be done in dL with appropriate formal specifications of the desired properties from the literature [135, 144]. Another complementary question is how to design a switching control law that *stabilizes* a given system. Switching design approaches are often guided by underlying stability arguments [99, 160, 189]; the loop invariants from Section 6.3 are expected to help guide correct-by-construction synthesis of such controllers.

Stability Analysis and Verification. Corollaries 6.9–6.12 formalize Lyapunov function-based stability arguments from the literature [27, 211] using loop invariants, yielding trustworthy, computer-checked stability proofs in KeYmaera X [54, 55]. Other computer-aided approaches for switched system stability analysis are based on finding Lyapunov functions that satisfy the requisite arithmetical conditions [88, 116, 126, 154, 170, 173]. Although the search for such functions can often be done efficiently with numerical techniques [116, 129, 154], various authors have emphasized the need to check that their outputs satisfy the arithmetical conditions *exactly*, i.e., without numerical errors compromising the resulting stability claims [3, 88, 167] (see, e.g., Section 6.5.2). This chapter's deductive approach goes further as it comprehensively verifies *all* steps of the stability argument down to its underlying discrete and continuous reasoning steps [142, 144]. The generality of this approach is precisely what enables verification of various classes of switching mechanisms all within a common logical framework (Section 6.3) and verification of non-standard stability arguments (Section 6.5). Alternative approaches to stability verification are based on abstraction [62, 183] and model checking [151].

6.7 Discussion

This chapter provides a blueprint for developing and verifying hybrid program models of switched systems. In particular, it shows how to deductively verify switched system stability, using dL's nested quantification over hybrid programs to specify stability, and dL's axiomatics to prove those specifications. Loop invariants—a classical technique from verification—are used to succinctly capture the desired properties of a given switching design; through deductive proofs, these invariants yield systematic, correct-by-construction derivation of the requisite arithmetical conditions on Lyapunov functions for stability arguments in implementations. An interesting direction for future work is to add other Lyapunov function generation techniques [88, 116, 126, 173] to the implementation, which—thanks to the presented approach—do not have to be trusted since their results can be checked independently by KeYmaera X. This would enable fully automated, yet sound and trustworthy verification of switched system stability based on dL's parsimonious hybrid program reasoning principles.

Chapter 7

Conclusion

An inspiration for the material of this thesis is the following snippet from the introduction to Liberzon's textbook on switched control systems [99]:

> *The field of hybrid systems has a strong interdisciplinary flavor, and different communities have developed different viewpoints. One approach, favored by researchers in computer science, is to concentrate on studying the discrete behavior of the system, while the continuous dynamics are assumed to take a relatively simple form. Basic issues in this context include well-posedness, simulation, and verification. Many researchers in systems and control theory, on the other hand, tend to regard hybrid systems as continuous systems with switching and place a greater emphasis on properties of the continuous state. The main issues then become stability analysis and control synthesis.*

This thesis takes a significant step in reconciling the two viewpoints mentioned by Liberzon. Chapters 3 and 4 show that deductive reasoning techniques from computer science scale to non-trivial safety and liveness properties of non-trivial ODEs, while Chapters 5 and 6 further show that those syntactic techniques are also well-suited for the formal study of control-theoretic stability for continuous and hybrid (switched) systems. Deductive reasoning crucially enables comprehensive and trustworthy verification results throughout this thesis—the *modus operandi* of the aforementioned chapters is to *i)* identify logical building blocks for reasoning about their respective specifications; then *ii)* generalize and/or combine those building blocks to obtain powerful new reasoning principles while retaining confidence in the correctness of the results by soundly and syntactically justifying every step from a parsimonious axiomatic foundation.

7.1 Thesis Summary

Differential dynamic logic (dL) is used as the common logical foundation throughout this thesis, which comes with three key benefits:

- Syntactic deduction in dL provides a means of formally proving a *comprehensive* set of specifications of interest—safety, liveness, and stability—for a given system, all within the same logical framework. These specifications come with powerful derived reasoning principles, such as the complete invariance reasoning principles of Chapter 3, refinement reasoning for liveness in Chapter 4, and the compositional combination of first-order and hybrid systems reasoning for stability specifications in Chapters 5 and 6.

- The uniform logical treatment yields *trustworthy* derived reasoning principles because the correctness of those principles syntactically reduce to the soundness of dL's parsimonious axiomatic foundations. Of course, an important caveat for pen-and-paper renditions of proof rules in this thesis is that one must trust the correctness of syntactic derivations presented in Chapters 3–6 (and their appendices). While every effort has been made to check the thesis results, the implementation of these derivations as (untrusted) tactics in KeYmaera X provides an additional correctness safeguard in practice because all proof steps are checked against KeYmaera X's microkernel implementation of dL [54, 55, 115, 142].
- The compositional reasoning principles of dL is used throughout this thesis to lift thesis results for ODEs to corresponding results for hybrid systems, notably Corollaries 3.14 and 3.39 which equivalently reduce analytic safety questions for analytic hybrid programs to arithmetic questions; and Chapter 6 which derives stability verification techniques for switched systems by combining (discrete) loop invariants with (continuous) Lyapunov function reasoning. The ODE reasoning principles for dL's box modality (Chapter 3) and diamond modality (Chapter 4) developed in this thesis also provide practical building blocks for proving sub-questions about continuous dynamics that arise, through compositional reasoning, as part of larger hybrid systems proofs in dL and KeYmaera X.

The thesis statement (recalled below) provides a lens with which to view these benefits and to summarize the results of Chapters 3–6.

> *Deductive reasoning provides a powerful, uniform, and foundational way of proving properties of ordinary differential equations. This logical foundation, in turn, yields new insights towards the verification of continuous and hybrid systems.*

7.2 Future Directions

Various avenues for future investigation have been discussed in Chapters 3–6. These avenues are summarized below, along with other potential directions.

Dynamical Generalizations. Ordinary differential equations (ODEs) are the quintessential model of continuous dynamics used throughout this thesis. A natural avenue for generalizing the thesis results is to investigate syntactic reasoning for alternative, or more general, descriptions of continuous dynamics, such as differential games [143] and differential-algebraic programs [137]. As suggested in Section 5.5, such dynamical extensions could even be useful for specifying properties of ODEs; recall, e.g., the continuous adversarial dynamics of a differential game can be used to model noisy perturbations of an ODE. Broadly speaking, the term language extensions from Chapter 3 are examples of practically fruitful dynamical generalizations because they expand the class of differential equations that can be directly modeled in dL while retaining sound and complete reasoning for their ODE invariants. Indeed, an implementation of Noetherian term language extensions for KeYmaera X is ongoing work [58]. Future work could examine the question posed at the end of Chapter 3, i.e., are there function classes that meet the chapter's extended term conditions but are *not* Noetherian functions? Or, more generally, is sound and complete syntactic reasoning for ODE invariants possible for function classes that do not necessarily meet those conditions? For example, dL can be extended with definite descriptions [21] which yields terms that are not necessarily smooth (not even differentiable) and it would be interesting to explore sound ODE reasoning principles for such an extension.

Syntactic Deduction and Control. Chapters 5 and 6 show how to carry out formal, syntactic proofs of stability(-like) properties for differential equations and switched systems. Various additional properties of control mechanisms are of interest in the literature [71, 89] and further work could investigate how to formally specify and verify these properties within dL. As argued in Section 5.5, such specifications may provide an avenue for exploring compositional extensions of dL with other forms of dynamics [136, 137, 138, 141, 143]; or for using extensions of dL's specification language [17, 19, 85, 90, 106, 133] to (more directly) syntactically express the control-theoretic properties of interest. This investigation is practically useful because it would identify extensions of KeYmaera X (if any) that are needed to enable its use as a powerful, semi-automated reasoning tool for trustworthy verification of continuous and hybrid control designs. Existing automation in KeYmaera X can also be improved for this purpose, e.g., the stability proof tactics for switched systems from Chapter 6 can be further complemented with a collection of methods for automatically generating Lyapunov function candidates for ODEs and switched systems, similar to the Pegasus tool [180] for generating ODE invariant candidates.

Verified Verification. From a bird's eye view, this thesis can also be seen as an instance of *verified verification*, i.e., the use of one verification technique (syntactic reasoning in dL) to study and justify other verification tools or techniques. This is best exemplified by the refinement-based approach to ODE liveness in Chapter 4, which is used to formally survey and correct existing liveness arguments from the literature. The self-evident benefit of verified verification is it adds an additional, alternative layer of correctness guarantee to the resulting verification tools, which adds to their overall trustworthiness; this is especially useful for safety-critical application domains, such as CPS verification [6, 144], where any added guarantee is worthwhile. An additional, subtle benefit is that the new perspectives afforded by alternative verification techniques may also lend themselves well to new verification insights. For example, the uniform refinement approach of Chapter 4 reveals building blocks behind ODE liveness arguments in the literature that can be generalized and pieced together to form new liveness arguments. Future work can explore verified verification in alternative domains, such as software and hardware verification, where the resulting correctness guarantees and generalizable insights are beneficial.

Bibliography

[1] Alessandro Abate, Alessandro D'Innocenzo, Maria Domenica Di Benedetto, and Shankar Sastry. Understanding deadlock and livelock behaviors in hybrid control systems. *Nonlinear Anal. Hybrid Syst.*, 3(2):150 – 162, 2009. DOI: `10.1016/j.nahs.2008.12.005`.

[2] Oskar Abrahamsson. A verified proof checker for higher-order logic. *J. Log. Algebraic Methods Program.*, 112:100530, 2020. DOI: `10.1016/j.jlamp.2020.100530`.

[3] Daniele Ahmed, Andrea Peruffo, and Alessandro Abate. Automated and sound synthesis of Lyapunov functions with SMT solvers. In Armin Biere and David Parker, editors, *TACAS*, volume 12078 of *LNCS*, pages 97–114. Springer, 2020. DOI: `10.1007/978-3-030-45190-5_6`.

[4] Behzad Akbarpour and Lawrence C. Paulson. MetiTarski: An automatic theorem prover for real-valued special functions. *J. Autom. Reason.*, 44(3):175–205, 2010. DOI: `10.1007/s10817-009-9149-2`.

[5] Matthias Althoff. An introduction to CORA 2015. In Goran Frehse and Matthias Althoff, editors, *ARCH*, volume 34 of *EPiC Series in Computing*, pages 120–151. EasyChair, 2015. DOI: `10.29007/zbkv`.

[6] Rajeev Alur. *Principles of Cyber-Physical Systems.* MIT Press, 2015.

[7] Rajeev Alur, Costas Courcoubetis, Nicolas Halbwachs, Thomas A. Henzinger, Pei-Hsin Ho, Xavier Nicollin, Alfredo Olivero, Joseph Sifakis, and Sergio Yovine. The algorithmic analysis of hybrid systems. *Theor. Comput. Sci.*, 138(1):3–34, 1995. DOI: `10.1016/0304-3975(94)00202-T`.

[8] Abhishek Anand and Ross A. Knepper. ROSCoq: Robots powered by constructive reals. In Christian Urban and Xingyuan Zhang, editors, *ITP*, volume 9236 of *LNCS*, pages 34–50. Springer, 2015. DOI: `10.1007/978-3-319-22102-1_3`.

[9] Mark S. Ashbaugh, Carmen C. Chicone, and Richard H. Cushman. The twisting tennis racket. *Journal of Dynamics and Differential Equations*, 3:67–85, 1991. DOI: `10.1007/BF01049489`.

[10] Jeremy Avigad, Robert Y. Lewis, and Cody Roux. A heuristic prover for real inequalities. *J. Autom. Reason.*, 56(3):367–386, 2016. DOI: `10.1007/s10817-015-9356-y`.

[11] Ralph-Johan Back and Joakim von Wright. *Refinement Calculus - A Systematic Introduction.* Springer, 1998. DOI: `10.1007/978-1-4612-1674-2`.

[12] Gal Binyamini. Density of algebraic points on Noetherian varieties. *Geom. Funct. Anal.*, 29(1):72–118, 2019. doi: 10.1007/s00039-019-00475-7.

[13] Anthony Bloch and Sergey Drakunov. Stabilization and tracking in the nonholonomic integrator via sliding modes. *Systems & Control Letters*, 29(2):91–99, 1996. doi: 10.1016/S0167-6911(96)00049-7.

[14] Jacek Bochnak, Michel Coste, and Marie-Françoise Roy. *Real Algebraic Geometry*. Springer, Heidelberg, 1998. doi: 10.1007/978-3-662-03718-8.

[15] Sergiy Bogomolov, Marcelo Forets, Goran Frehse, Kostiantyn Potomkin, and Christian Schilling. JuliaReach: A toolbox for set-based reachability. In Necmiye Ozay and Pavithra Prabhakar, editors, *HSCC*, pages 39–44. ACM, 2019. doi: 10.1145/3302504.3311804.

[16] Rose Bohrer. Differential dynamic logic. *Archive of Formal Proofs*, February 2017. https://isa-afp.org/entries/Differential_Dynamic_Logic.html, Formal proof development.

[17] Rose Bohrer and André Platzer. A hybrid, dynamic logic for hybrid-dynamic information flow. In Anuj Dawar and Erich Grädel, editors, *LICS*, pages 115–124. ACM, 2018. doi: 10.1145/3209108.3209151.

[18] Rose Bohrer and André Platzer. Constructive hybrid games. In Nicolas Peltier and Viorica Sofronie-Stokkermans, editors, *IJCAR*, volume 12166 of *LNCS*, pages 454–473. Springer, 2020. doi: 10.1007/978-3-030-51074-9_26.

[19] Rose Bohrer and André Platzer. Refining constructive hybrid games. In Zena M. Ariola, editor, *FSCD*, volume 167 of *LIPIcs*, pages 14:1–14:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi: 10.4230/LIPIcs.FSCD.2020.14.

[20] Rose Bohrer, Vincent Rahli, Ivana Vukotic, Marcus Völp, and André Platzer. Formally verified differential dynamic logic. In Yves Bertot and Viktor Vafeiadis, editors, *CPP*, pages 208–221. ACM, 2017. doi: 10.1145/3018610.3018616.

[21] Rose Bohrer, Manuel Fernández, and André Platzer. dl_ι: Definite descriptions in differential dynamic logic. In Pascal Fontaine, editor, *CADE*, volume 11716 of *LNCS*, pages 94–110. Springer, 2019. doi: 10.1007/978-3-030-29436-6_6.

[22] Rose Bohrer, Yong Kiam Tan, Stefan Mitsch, Andrew Sogokon, and André Platzer. A formal safety net for waypoint-following in ground robots. *IEEE Robot. Autom. Lett.*, 4(3):2910–2917, 2019. doi: 10.1109/LRA.2019.2923099.

[23] Sylvie Boldo, Catherine Lelay, and Guillaume Melquiond. Formalization of real analysis: a survey of proof assistants and libraries. *Math. Struct. Comput. Sci.*, 26(7):1196–1233, 2016. doi: 10.1017/S0960129514000437.

[24] Nicolas Bourbaki. *Commutative Algebra. Chapters 1–7.* Springer, Berlin, 1998.

[25] Olivier Bournez and Amaury Pouly. A universal ordinary differential equation. *Log. Methods Comput. Sci.*, 16(1), 2020. doi: 10.23638/LMCS-16(1:28)2020.

[26] Michael S. Branicky. Analyzing continuous switching systems: theory and examples. In *ACC*, volume 3, pages 3110–3114, 1994. DOI: 10.1109/ACC.1994.735143.

[27] Michael S. Branicky. Multiple Lyapunov functions and other analysis tools for switched and hybrid systems. *IEEE Trans. Autom. Control.*, 43(4):475–482, 1998. DOI: 10.1109/9.664150.

[28] Michael S. Branicky. Introduction to hybrid systems. In Dimitrios Hristu-Varsakelis and William S. Levine, editors, *Handbook of Networked and Embedded Control Systems*, pages 91–116. Birkhäuser, 2005. DOI: 10.1007/0-8176-4404-0_5.

[29] R. W. Brockett. Asymptotic stability and feedback stabilization. In *Differential Geometric Control Theory*, pages 181–191. Birkhauser, 1983.

[30] Michael J. Butler, Jean-Raymond Abrial, and Richard Banach. Modelling and refining hybrid systems in Event-B and Rodin. In Luigia Petre and Emil Sekerinski, editors, *From Action Systems to Distributed Systems - The Refinement Approach*, pages 29–42. Chapman and Hall/CRC, 2016. DOI: 10.1201/b20053.

[31] Xin Chen, Erika Ábrahám, and Sriram Sankaranarayanan. Flow*: An analyzer for non-linear hybrid systems. In Natasha Sharygina and Helmut Veith, editors, *CAV*, volume 8044 of *LNCS*, pages 258–263, Heidelberg, 2013. Springer. DOI: 10.1007/978-3-642-39799-8_18.

[32] Xin Chen, Sriram Sankaranarayanan, and Erika Ábrahám. Under-approximate flowpipes for non-linear continuous systems. In Koen Claessen and Viktor Kuncak, editors, *FMCAD*, pages 59–66. IEEE, 2014. DOI: 10.1109/FMCAD.2014.6987596.

[33] Carmen Chicone. *Ordinary Differential Equations with Applications.* Springer, New York, second edition, 2006. DOI: 10.1007/0-387-35794-7.

[34] Pete L. Clark. The instructor's guide to real induction. *Math. Mag.*, 92(2):136–150, 2019. DOI: 10.1080/0025570X.2019.1549902.

[35] Edmund M. Clarke, Ansgar Fehnker, Zhi Han, Bruce H. Krogh, Joël Ouaknine, Olaf Stursberg, and Michael Theobald. Abstraction and counterexample-guided refinement in model checking of hybrid systems. *Int. J. Found. Comput. Sci.*, 14(4):583–604, 2003. DOI: 10.1142/S012905410300190X.

[36] Cyril Cohen and Assia Mahboubi. Formal proofs in real algebraic geometry: from ordered fields to quantifier elimination. *Log. Methods Comput. Sci.*, 8(1), 2012. DOI: 10.2168/LMCS-8(1:2)2012.

[37] Cyril Cohen and Damien Rouhling. A formal proof in Coq of LaSalle's invariance principle. In Mauricio Ayala-Rincón and César A. Muñoz, editors, *ITP*, volume 10499 of *LNCS*, pages 148–163. Springer, 2017. DOI: 10.1007/978-3-319-66107-0_10.

[38] Katherine Cordwell, Yong Kiam Tan, and André Platzer. A verified decision procedure for univariate real arithmetic with the BKR algorithm. In Liron Cohen and Cezary Kaliszyk, editors, *ITP*, volume 193 of *LIPIcs*, pages 14:1–14:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. DOI: 10.4230/LIPIcs.ITP.2021.14.

[39] Jorge Cortes. Discontinuous dynamical systems. *IEEE Control Systems Magazine*, 28(3): 36–73, 2008. DOI: 10.1109/MCS.2008.919306.

[40] Liyun Dai, Ting Gan, Bican Xia, and Naijun Zhan. Barrier certificates revisited. *J. Symb. Comput.*, 80:62–86, 2017. DOI: 10.1016/j.jsc.2016.07.010.

[41] Jean-Gaston Darboux. Mémoire sur les équations différentielles algébriques du premier ordre et du premier degré. *Bull. Sci. Math.*, 2(1):151–200, 1878.

[42] Jared Davis and Magnus O. Myreen. The reflective Milawa theorem prover is sound (down to the machine code that runs it). *J. Autom. Reason.*, 55(2):117–183, 2015. DOI: 10.1007/s10817-015-9324-6.

[43] Raymond A. Decarlo, Michael S. Branicky, Stefan Pettersson, and Bengt Lennartson. Perspectives and results on the stability and stabilizability of hybrid systems. *Proceedings of the IEEE*, 88(7):1069–1082, 2000. DOI: 10.1109/5.871309.

[44] Raymond A. Decarlo, Michael S. Branicky, Stefan Pettersson, and Bengt Lennartson. Perspectives and results on the stability and stabilizability of hybrid systems. *Proc. IEEE*, 88(7):1069–1082, 2000. DOI: 10.1109/5.871309.

[45] Laurent Doyen, Goran Frehse, George J. Pappas, and André Platzer. Verification of hybrid systems. In Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors, *Handbook of Model Checking*, pages 1047–1110. Springer, 2018. DOI: 10.1007/978-3-319-10575-8_30.

[46] Parasara Sridhar Duggirala and Sayan Mitra. Lyapunov abstractions for inevitability of hybrid systems. In Thao Dang and Ian M. Mitchell, editors, *HSCC*, pages 115–124, New York, 2012. ACM. DOI: 10.1145/2185632.2185652.

[47] Guillaume Dupont, Yamine Aït Ameur, Marc Pantel, and Neeraj Kumar Singh. Handling refinement of continuous behaviors: A proof based approach with Event-B. In Dominique Méry and Shengchao Qin, editors, *TASE*, pages 9–16. IEEE, 2019. DOI: 10.1109/TASE.2019.00-25.

[48] Chuchu Fan, Bolun Qi, Sayan Mitra, Mahesh Viswanathan, and Parasara Sridhar Duggirala. Automatic reachability analysis for nonlinear hybrid models with C2E2. In Swarat Chaudhuri and Azadeh Farzan, editors, *CAV*, volume 9779 of *LNCS*, pages 531–538. Springer, 2016. DOI: 10.1007/978-3-319-41528-4_29.

[49] K. Forsman. Construction of Lyapunov functions using Gröbner bases. In *CDC*, volume 1, pages 798–799. IEEE, 1991. DOI: 10.1109/CDC.1991.261424.

[50] Simon Foster, Jonathan Julián Huerta y Munive, and Georg Struth. Differential Hoare logics and refinement calculi for hybrid systems with Isabelle/HOL. In Uli Fahrenberg, Peter Jipsen, and Michael Winter, editors, *RAMiCS*, volume 12062 of *LNCS*, pages 169–186. Springer, 2020. DOI: 10.1007/978-3-030-43520-2_11.

[51] Simon Foster, Jonathan Julián Huerta y Munive, Mario Gleirscher, and Georg Struth. Hybrid systems verification with Isabelle/HOL: Simpler syntax, better models, faster proofs. In Marieke Huisman, Corina S. Pasareanu, and Naijun Zhan, editors, *FM*, volume 13047 of *LNCS*, pages 367–386. Springer, 2021. DOI: 10.1007/978-3-030-90870-6_20.

[52] Goran Frehse. PHAVer: Algorithmic verification of hybrid systems past HyTech. In Manfred Morari and Lothar Thiele, editors, *HSCC*, volume 3414 of *LNCS*, pages 258–273. Springer, 2005. DOI: 10.1007/978-3-540-31954-2_17.

[53] Goran Frehse, Colas Le Guernic, Alexandre Donzé, Scott Cotton, Rajarshi Ray, Olivier Lebeltel, Rodolfo Ripado, Antoine Girard, Thao Dang, and Oded Maler. SpaceEx: Scalable verification of hybrid systems. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *CAV*, volume 6806 of *LNCS*, pages 379–395, Heidelberg, 2011. Springer. DOI: 10.1007/978-3-642-22110-1_30.

[54] Nathan Fulton, Stefan Mitsch, Jan-David Quesel, Marcus Völp, and André Platzer. KeYmaera X: An axiomatic tactical theorem prover for hybrid systems. In Amy P. Felty and Aart Middeldorp, editors, *CADE*, volume 9195 of *LNCS*, pages 527–538, Cham, 2015. Springer. DOI: 10.1007/978-3-319-21401-6_36.

[55] Nathan Fulton, Stefan Mitsch, Rose Bohrer, and André Platzer. Bellerophon: Tactical theorem proving for hybrid systems. In Mauricio Ayala-Rincón and César A. Muñoz, editors, *ITP*, volume 10499 of *LNCS*, pages 207–224, Cham, 2017. Springer. DOI: 10.1007/978-3-319-66107-0_14.

[56] Andrei Gabrielov and Askold Khovanskii. Multiplicity of a Noetherian intersection. In *Geometry of Differential Equations*, pages 119–130. Amer. Math. Soc., Providence, 1998. DOI: 10.1090/trans2/186/03.

[57] Andrei Gabrielov and Nicolai Vorobjov. Complexity of computations with Pfaffian and Noetherian functions. In *Normal Forms, Bifurcations and Finiteness Problems in Differential Equations*, pages 211–250. Kluwer Acad. Publ., Netherlands, 2004.

[58] James Gallicchio, Yong Kiam Tan, Stefan Mitsch, and André Platzer. Implicit definitions with differential equations for KeYmaera X (system description). *CoRR*, abs/2203.01272, 2022. URL http://arxiv.org/abs/2203.01272.

[59] Sicun Gao, Jeremy Avigad, and Edmund M. Clarke. Delta-decidability over the reals. In *LICS*, pages 305–314. IEEE Computer Society, 2012. DOI: 10.1109/LICS.2012.41.

[60] Sicun Gao, Soonho Kong, and Edmund M. Clarke. dReal: An SMT solver for nonlinear theories over the reals. In Maria Paola Bonacina, editor, *CADE*, volume 7898 of *LNCS*, pages 208–214, Heidelberg, 2013. Springer. DOI: 10.1007/978-3-642-38574-2_14.

[61] Sicun Gao, James Kapinski, Jyotirmoy V. Deshmukh, Nima Roohi, Armando Solar-Lezama, Nikos Aréchiga, and Soonho Kong. Numerically-robust inductive proof rules for continuous dynamical systems. In Isil Dillig and Serdar Tasiran, editors, *CAV*, volume 11562 of *LNCS*, pages 137–154. Springer, 2019. DOI: 10.1007/978-3-030-25543-5_9.

[62] Miriam García Soto and Pavithra Prabhakar. Abstraction based verification of stability of polyhedral switched systems. *Nonlinear Analysis: Hybrid Systems*, 36:100856, 2020. DOI: https://doi.org/10.1016/j.nahs.2020.100856.

[63] Khalil Ghorbal and André Platzer. Characterizing algebraic invariants by differential radical invariants. In Erika Ábrahám and Klaus Havelund, editors, *TACAS*, volume 8413 of *LNCS*, pages 279–294, Heidelberg, 2014. Springer. DOI: 10.1007/978-3-642-54862-8_19.

[64] Khalil Ghorbal, Andrew Sogokon, and André Platzer. A hierarchy of proof rules for checking positive invariance of algebraic and semi-algebraic sets. *Comput. Lang. Syst. Struct.*, 47:19–43, 2017. DOI: 10.1016/j.cl.2015.11.003.

[65] Rafal Goebel, Ricardo G. Sanfelice, and Andrew R. Teel. Hybrid dynamical systems. *IEEE Control Systems Magazine*, 29(2):28–93, 2009. DOI: 10.1109/MCS.2008.931718.

[66] Rafal Goebel, Ricardo G. Sanfelice, and Andrew R. Teel. *Hybrid Dynamical Systems: Modeling, Stability, and Robustness.* Princeton University Press, 2012.

[67] Eric Goubault and Sylvie Putot. Forward inner-approximated reachability of non-linear continuous systems. In Goran Frehse and Sayan Mitra, editors, *HSCC*, pages 1–10, New York, 2017. ACM. DOI: 10.1145/3049797.3049811.

[68] Daniel S. Graça, Jorge Buescu, and Manuel Lameiras Campagnolo. Boundedness of the domain of definition is undecidable for polynomial ODEs. *Electron. Notes Theor. Comput. Sci.*, 202:49–57, 2008. DOI: 10.1016/j.entcs.2008.03.007.

[69] Daniel S. Graça, Manuel L. Campagnolo, and Jorge Buescu. Computability with polynomial differential equations. *Adv. Appl. Math.*, 40(3):330 – 349, 2008. DOI: 10.1016/j.aam.2007.02.003.

[70] Thomas H. Grönwall. Note on the derivatives with respect to a parameter of the solutions of a system of differential equations. *Ann. Math.*, 20(4):292–296, 1919. DOI: 10.2307/1967124.

[71] Wassim M. Haddad and VijaySekhar Chellaboina. *Nonlinear Dynamical Systems and Control: A Lyapunov-Based Approach.* Princeton University Press, 2008.

[72] Wassim M. Haddad, VijaySekhar Chellaboina, and Sergey G. Nersesov. *Impulsive and Hybrid Dynamical Systems: Stability, Dissipativity, and Control.* Princeton University Press, 2006.

[73] David Harel. *First-Order Dynamic Logic*, volume 68 of *LNCS*. Springer, 1979. DOI: 10.1007/3-540-09237-4.

[74] John Harrison. Verifying nonlinear real formulas via sums of squares. In Klaus Schneider and Jens Brandt, editors, *TPHOLs*, volume 4732 of *LNCS*, pages 102–118. Springer, 2007. DOI: 10.1007/978-3-540-74591-4_9.

[75] Thomas A. Henzinger. The theory of hybrid automata. In *LICS*, pages 278–292. IEEE Computer Society, 1996. DOI: 10.1109/LICS.1996.561342.

[76] Thomas A. Henzinger. It's about time: Real-time logics reviewed. In Davide Sangiorgi and Robert de Simone, editors, *CONCUR*, volume 1466 of *LNCS*, pages 439–454, Heidelberg, 1998. Springer. DOI: 10.1007/BFb0055640.

[77] Morris W. Hirsch. The dynamical systems approach to differential equations. *Bull. Amer. Math. Soc. (N.S.)*, 11(1):1–64, 07 1984.

[78] Johannes Hölzl, Fabian Immler, and Brian Huffman. Type classes and filters for mathematical analysis in Isabelle/HOL. In Sandrine Blazy, Christine Paulin-Mohring, and David Pichardie, editors, *ITP*, volume 7998 of *LNCS*, pages 279–294. Springer, 2013. DOI: 10.1007/978-3-642-39634-2_21.

[79] Fabian Immler. A verified ODE solver and the Lorenz attractor. *J. Autom. Reason.*, 61(1-4): 73–111, 2018. DOI: 10.1007/s10817-017-9448-y.

[80] Fabian Immler and Johannes Hölzl. Numerical analysis of ordinary differential equations in Isabelle/HOL. In Lennart Beringer and Amy P. Felty, editors, *ITP*, volume 7406 of *LNCS*, pages 377–392. Springer, 2012. DOI: 10.1007/978-3-642-32347-8_26.

[81] Fabian Immler and Johannes Hölzl. Ordinary differential equations. *Archive of Formal Proofs*, April 2012. https://isa-afp.org/entries/Ordinary_Differential_Equations.html, Formal proof development.

[82] Fabian Immler and Yong Kiam Tan. The Poincaré-Bendixson theorem in Isabelle/HOL. In Jasmin Blanchette and Catalin Hritcu, editors, *CPP*, pages 338–352. ACM, 2020. DOI: 10.1145/3372885.3373833.

[83] Fabian Immler and Christoph Traut. The flow of ODEs: Formalization of variational equation and poincaré map. *J. Autom. Reason.*, 62(2):215–236, 2019. DOI: 10.1007/s10817-018-9449-5.

[84] Fabian Immler, Matthias Althoff, Luis Benet, Alexandre Chapoutot, Xin Chen, Marcelo Forets, Luca Geretti, Niklas Kochdumper, David P. Sanders, and Christian Schilling. ARCH-COMP19 category report: Continuous and hybrid systems with nonlinear dynamics. In Goran Frehse and Matthias Althoff, editors, *ARCH*, volume 61 of *EPiC Series in Computing*, pages 41–61. EasyChair, 2019. DOI: 10.29007/m75b.

[85] Jean-Baptiste Jeannin and André Platzer. dTL2: Differential temporal dynamic logic with nested temporalities for hybrid systems. In Stéphane Demri, Deepak Kapur, and Christoph Weidenbach, editors, *IJCAR*, volume 8562 of *LNCS*, pages 292–306. Springer, 2014. DOI: 10.1007/978-3-319-08587-6_22.

[86] Martin Johansson and Anders Rantzer. Computation of piecewise quadratic Lyapunov functions for hybrid systems. *IEEE Trans. Autom. Control.*, 43(4):555–559, 1998. DOI: 10.1109/9.664157.

[87] Eduard Kamburjan. From post-conditions to post-region invariants: deductive verification of hybrid objects. In Sergiy Bogomolov and Raphaël M. Jungers, editors, *HSCC*, pages 9:1–9:11. ACM, 2021. DOI: 10.1145/3447928.3456633.

[88] James Kapinski, Jyotirmoy V. Deshmukh, Sriram Sankaranarayanan, and Nikos Aréchiga. Simulation-guided Lyapunov analysis for hybrid dynamical systems. In Martin Fränzle and John Lygeros, editors, *HSCC*, pages 133–142. ACM, 2014. DOI: 10.1145/2562059.2562139.

[89] Hassan K. Khalil. *Nonlinear systems.* Macmillan Publishing Company, New York, 1992.

[90] Juraj Kolcák, Jérémy Dubut, Ichiro Hasuo, Shin-ya Katsumata, David Sprunger, and Akihisa Yamada. Relational differential dynamic logic. In Armin Biere and David Parker, editors, *TACAS*, volume 12078 of *LNCS*, pages 191–208. Springer, 2020. DOI: 10.1007/978-3-030-45190-5_11.

[91] Hui Kong, Fei He, Xiaoyu Song, William N. N. Hung, and Ming Gu. Exponential-condition-based barrier certificate generation for safety verification of hybrid systems. In Natasha Sharygina and Helmut Veith, editors, *CAV*, volume 8044 of *LNCS*, pages 242–257. Springer, 2013. DOI: 10.1007/978-3-642-39799-8_17.

[92] Margarita V. Korovina and Nicolai Vorobjov. Pfaffian hybrid systems. In Jerzy Marcinkowski and Andrzej Tarlecki, editors, *CSL*, volume 3210 of *LNCS*, pages 430–441, Heidelberg, 2004. Springer. DOI: 10.1007/978-3-540-30124-0_33.

[93] Dexter Kozen. Kleene algebra with tests. *ACM Trans. Program. Lang. Syst.*, 19(3):427–443, 1997. DOI: 10.1145/256167.256195.

[94] Steven G. Krantz and Harold R. Parks. *A Primer of Real Analytic Functions.* Birkhäuser, Boston, second edition, 2002. DOI: 10.1007/978-0-8176-8134-0.

[95] Ramana Kumar, Rob Arthan, Magnus O. Myreen, and Scott Owens. HOL with definitions: Semantics, soundness, and a verified implementation. In Gerwin Klein and Ruben Gamboa, editors, *ITP*, volume 8558 of *LNCS*, pages 308–324. Springer, 2014. DOI: 10.1007/978-3-319-08970-6_20.

[96] Gerardo Lafferriere, George J. Pappas, and Shankar Sastry. O-minimal hybrid systems. *Math. Control Signals Systems*, 13(1):1–21, 2000. DOI: 10.1007/PL00009858.

[97] Wenda Li, Grant Olney Passmore, and Lawrence C. Paulson. Deciding univariate polynomial problems using untrusted certificates in Isabelle/HOL. *J. Autom. Reason.*, 62(1): 69–91, 2019. DOI: 10.1007/s10817-017-9424-6.

[98] A. Liapounoff. Probléme général de la stabilité du mouvement. *Annales de la Faculté des sciences de Toulouse : Mathématiques*, 9:203–474, 1907.

[99] Daniel Liberzon. *Switching in Systems and Control.* Systems & Control: Foundations & Applications. Birkhäuser, 2003. DOI: 10.1007/978-1-4612-0017-8.

[100] Timm Liebrenz, Paula Herber, and Sabine Glesner. Deductive verification of hybrid control systems modeled in Simulink with KeYmaera X. In Jing Sun and Meng Sun, editors, *ICFEM*, volume 11232 of *LNCS*, pages 89–105. Springer, 2018. DOI: 10.1007/978-3-030-02450-5_6.

[101] Qin Lin, Stefan Mitsch, André Platzer, and John M. Dolan. Safe and resilient practical waypoint-following for autonomous vehicles. *IEEE Control. Syst. Lett.*, 6:1574–1579, 2022. DOI: 10.1109/LCSYS.2021.3125717.

[102] Jiang Liu, Jidong Lv, Zhao Quan, Naijun Zhan, Hengjun Zhao, Chaochen Zhou, and Liang Zou. A calculus for hybrid CSP. In Kazunori Ueda, editor, *APLAS*, volume 6461 of *LNCS*, pages 1–15. Springer, 2010. DOI: 10.1007/978-3-642-17164-2_1.

[103] Jiang Liu, Naijun Zhan, and Hengjun Zhao. Computing semi-algebraic invariants for polynomial dynamical systems. In Samarjit Chakraborty, Ahmed Jerraya, Sanjoy K. Baruah, and Sebastian Fischmeister, editors, *EMSOFT*, pages 97–106, New York, 2011. ACM. DOI: 10.1145/2038642.2038659.

[104] Jiang Liu, Naijun Zhan, and Hengjun Zhao. Automatically discovering relaxed Lyapunov functions for polynomial dynamical systems. *Math. Comput. Sci.*, 6(4):395–408, 2012. DOI: 10.1007/s11786-012-0133-6.

[105] Jiang Liu, Naijun Zhan, Hengjun Zhao, and Liang Zou. Abstraction of elementary hybrid systems by variable transformation. In Nikolaj Bjørner and Frank S. de Boer, editors, *FM*, volume 9109 of *LNCS*, pages 360–377, Cham, 2015. Springer. DOI: 10.1007/978-3-319-19249-9_23.

[106] Sarah M. Loos and André Platzer. Differential refinement logic. In Martin Grohe, Eric Koskinen, and Natarajan Shankar, editors, *LICS*, pages 505–514. ACM, 2016. DOI: 10.1145/2933575.2934555.

[107] John Lygeros, Karl Henrik Johansson, Slobodan N. Simic, Jun Zhang, and Shankar S. Sastry. Dynamical properties of hybrid automata. *IEEE Trans. Autom. Control.*, 48(1):2–17, 2003. DOI: 10.1109/TAC.2002.806650.

[108] Evgeny Makarov and Bas Spitters. The Picard algorithm for ordinary differential equations in Coq. In Sandrine Blazy, Christine Paulin-Mohring, and David Pichardie, editors, *ITP*, volume 7998 of *LNCS*, pages 463–468. Springer, 2013. DOI: 10.1007/978-3-642-39634-2_34.

[109] Zohar Manna and Amir Pnueli. *The Temporal Logic of Reactive and Concurrent Systems - Specification.* Springer, New York, 1992. DOI: 10.1007/978-1-4612-0931-7.

[110] João Martins, André Platzer, and João Leite. Dynamic doxastic differential dynamic logic for belief-aware cyber-physical systems. In Serenella Cerrito and Andrei Popescu, editors, *TABLEAUX*, volume 11714 of *LNCS*, pages 428–445. Springer, 2019. DOI: 10.1007/978-3-030-29026-9_24.

[111] The mathlib Community. The Lean mathematical library. In Jasmin Blanchette and Catalin Hritcu, editors, *CPP*, pages 367–381. ACM, 2020. DOI: 10.1145/3372885.3373824.

[112] Scott McCallum and Volker Weispfenning. Deciding polynomial-transcendental problems. *J. Symb. Comput.*, 47(1):16–31, 2012. DOI: 10.1016/j.jsc.2011.08.004.

[113] Sean McLaughlin and John Harrison. A proof-producing decision procedure for real arithmetic. In Robert Nieuwenhuis, editor, *CADE*, volume 3632 of *LNCS*, pages 295–314. Springer, 2005. DOI: 10.1007/11532231_22.

[114] Stefan Mitsch and André Platzer. The KeYmaera X proof IDE: Concepts on usability in hybrid systems theorem proving. In Catherine Dubois, Paolo Masci, and Dominique Méry, editors, *3rd Workshop on Formal Integrated Development Environment*, volume 240 of *EPTCS*, pages 67–81, 2016. DOI: 10.4204/EPTCS.240.5.

[115] Stefan Mitsch and André Platzer. A retrospective on developing hybrid system provers in the KeYmaera family - A tale of three provers. In Wolfgang Ahrendt, Bernhard Beckert, Richard Bubel, Reiner Hähnle, and Mattias Ulbrich, editors, *Deductive Software Verification: Future Perspectives - Reflections on the Occasion of 20 Years of KeY*, volume 12345 of *LNCS*, pages 21–64. Springer, 2020. DOI: 10.1007/978-3-030-64354-6_2.

[116] Eike Möhlmann and Oliver E. Theel. Stabhyli: a tool for automatic stability verification of non-linear hybrid systems. In Calin Belta and Franjo Ivancic, editors, *HSCC*, pages 107–112. ACM, 2013. DOI: 10.1145/2461328.2461347.

[117] Eike Möhlmann and Oliver E. Theel. Stabhyli, 2021. URL https://uol.de/svs/forschung/avacs/stabhyli. [Online; accessed 27-October-2021].

[118] A. S. Morse. Control using logic-based switching. In Alberto Isidori, editor, *Trends in Control*, pages 69–113, London, 1995. Springer London. DOI: 10.1007/978-1-4471-3061-1_4.

[119] Andreas Müller, Stefan Mitsch, Werner Retschitzegger, Wieland Schwinger, and André Platzer. Tactical contract composition for hybrid system component verification. *Int. J. Softw. Tools Technol. Transf.*, 20(6):615–643, 2018. DOI: 10.1007/s10009-018-0502-9.

[120] César A. Muñoz, Anthony J. Narkawicz, and Aaron Dutle. A decision procedure for univariate polynomial systems based on root counting and interval subdivision. *J. Formaliz. Reason.*, 11(1):19–41, 2018. DOI: 10.6092/issn.1972-5787/8212.

[121] Anthony Narkawicz, César A. Muñoz, and Aaron Dutle. Formally-verified decision procedures for univariate polynomial computation based on Sturm's and Tarski's theorems. *J. Autom. Reason.*, 54(4):285–326, 2015. DOI: 10.1007/s10817-015-9320-x.

[122] Eva M. Navarro-López and Rebekah Carter. Deadness and how to disprove liveness in hybrid dynamical systems. *Theor. Comput. Sci.*, 642:1–23, 2016. DOI: 10.1016/j.tcs.2016.06.009.

[123] Luan Viet Nguyen, James Kapinski, Xiaoqing Jin, Jyotirmoy V. Deshmukh, and Taylor T. Johnson. Hyperproperties of real-valued signals. In Jean-Pierre Talpin, Patricia Derler, and Klaus Schneider, editors, *MEMOCODE*, pages 104–113. ACM, 2017. DOI: 10.1145/3127041.3127058.

[124] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL - A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002. DOI: 10.1007/3-540-45949-9.

[125] Dimitri Novikov and Sergei Yakovenko. Trajectories of polynomial vector fields and ascending chains of polynomial ideals. *Ann. I. Fourier*, 49(2):563–609, 1999. DOI: 10.5802/aif.1683.

[126] Jens Oehlerking. *Decomposition of stability proofs for hybrid systems*. PhD thesis, Carl von Ossietzky University of Oldenburg, 2011.

[127] Susan S. Owicki and David Gries. Verifying properties of parallel programs: An axiomatic approach. *Commun. ACM*, 19(5):279–285, 1976. DOI: 10.1145/360051.360224.

[128] Susan S. Owicki and Leslie Lamport. Proving liveness properties of concurrent programs. *ACM Trans. Program. Lang. Syst.*, 4(3):455–495, 1982. DOI: 10.1145/357172.357178.

[129] A. Papachristodoulou, J. Anderson, G. Valmorbida, S. Prajna, P. Seiler, P. A. Parrilo, M. M. Peet, and D. Jagt. *SOSTOOLS: Sum of squares optimization toolbox for MATLAB*. http://arxiv.org/abs/1310.4716, 2021. Available from https://github.com/oxfordcontrol/SOSTOOLS.

[130] Antonis Papachristodoulou and Stephen Prajna. On the construction of Lyapunov functions using the sum of squares decomposition. In *CDC*, pages 3482–3487. IEEE, 2002. DOI: 10.1109/CDC.2002.1184414.

[131] Pablo A. Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. PhD thesis, California Institute of Technology, 2000.

[132] Andrea Peruffo, Daniele Ahmed, and Alessandro Abate. Automated and formal synthesis of neural barrier certificates for dynamical models. In Jan Friso Groote and Kim Guldstrand Larsen, editors, *TACAS*, volume 12651 of *LNCS*, pages 370–388. Springer, 2021. DOI: 10.1007/978-3-030-72016-2_20.

[133] André Platzer. A temporal dynamic logic for verifying hybrid system invariants. In Sergei N. Artëmov and Anil Nerode, editors, *LFCS*, volume 4514 of *LNCS*, pages 457–471. Springer, 2007. DOI: 10.1007/978-3-540-72734-7_32.

[134] André Platzer. Differential dynamic logic for hybrid systems. *J. Autom. Reasoning*, 41(2):143–189, 2008. DOI: 10.1007/s10817-008-9103-8.

[135] André Platzer. *Logical Analysis of Hybrid Systems - Proving Theorems for Complex Dynamics*. Springer, 2010. DOI: 10.1007/978-3-642-14509-4.

[136] André Platzer. Quantified differential dynamic logic for distributed hybrid systems. In Anuj Dawar and Helmut Veith, editors, *CSL*, volume 6247 of *LNCS*, pages 469–483. Springer, 2010. DOI: `10.1007/978-3-642-15205-4_36`.

[137] André Platzer. Differential-algebraic dynamic logic for differential-algebraic programs. *J. Log. Comput.*, 20(1):309–352, 2010. DOI: `10.1093/logcom/exn070`.

[138] André Platzer. Stochastic differential dynamic logic for stochastic hybrid programs. In Nikolaj Bjørner and Viorica Sofronie-Stokkermans, editors, *CADE*, volume 6803 of *LNCS*, pages 446–460. Springer, 2011. DOI: `10.1007/978-3-642-22438-6_34`.

[139] André Platzer. The complete proof theory of hybrid systems. In *LICS*, pages 541–550. IEEE Computer Society, 2012. DOI: `10.1109/LICS.2012.64`.

[140] André Platzer. The structure of differential invariants and differential cut elimination. *Log. Meth. Comput. Sci.*, 8(4):1–38, 2012. DOI: `10.2168/LMCS-8(4:16)2012`.

[141] André Platzer. Differential game logic. *ACM Trans. Comput. Log.*, 17(1):1:1–1:51, 2015. DOI: `10.1145/2817824`.

[142] André Platzer. A complete uniform substitution calculus for differential dynamic logic. *J. Autom. Reason.*, 59(2):219–265, 2017. DOI: `10.1007/s10817-016-9385-1`.

[143] André Platzer. Differential hybrid games. *ACM Trans. Comput. Log.*, 18(3):19:1–19:44, 2017. DOI: `10.1145/3091123`.

[144] André Platzer. *Logical Foundations of Cyber-Physical Systems*. Springer, Cham, 2018. DOI: `10.1007/978-3-319-63588-0`.

[145] André Platzer. Uniform substitution at one fell swoop. In Pascal Fontaine, editor, *CADE*, volume 11716 of *LNCS*, pages 425–441. Springer, 2019. DOI: `10.1007/978-3-030-29436-6_25`.

[146] André Platzer and Edmund M. Clarke. Computing differential invariants of hybrid systems as fixedpoints. *Formal Methods Syst. Des.*, 35(1):98–120, 2009. DOI: `10.1007/s10703-009-0079-8`.

[147] André Platzer and Jan-David Quesel. KeYmaera: A hybrid theorem prover for hybrid systems. In Alessandro Armando, Peter Baumgartner, and Gilles Dowek, editors, *IJCAR*, volume 5195 of *LNCS*, pages 171–178. Springer, 2008. DOI: `10.1007/978-3-540-71070-7_15`.

[148] André Platzer and Yong Kiam Tan. Differential equation axiomatization: The impressive power of differential ghosts. In Anuj Dawar and Erich Grädel, editors, *LICS*, pages 819–828, New York, 2018. ACM. DOI: `10.1145/3209108.3209147`.

[149] André Platzer and Yong Kiam Tan. Differential equation invariance axiomatization. *J. ACM*, 67(1):6:1–6:66, 2020. DOI: `10.1145/3380825`.

[150] André Platzer, Jan-David Quesel, and Philipp Rümmer. Real world verification. In Renate A. Schmidt, editor, *CADE*, volume 5663 of *LNCS*, pages 485–501. Springer, 2009. doi: 10.1007/978-3-642-02959-2_35.

[151] Andreas Podelski and Silke Wagner. Model checking of hybrid systems: From reachability towards stability. In João P. Hespanha and Ashish Tiwari, editors, *HSCC*, volume 3927 of *LNCS*, pages 507–521. Springer, 2006. doi: 10.1007/11730637_38.

[152] Henri Poincaré. Mémoire sur les courbes définies par une équation différentielle. *J. Math. Pures Appl.*, 1881.

[153] Henri Poincaré. *Les méthodes nouvelles de la mécanique céleste*. Gauthier-Villars, Paris, 1892–1899.

[154] S. Prajna and A. Papachristodoulou. Analysis of switched and hybrid systems - beyond piecewise quadratic methods. In *ACC*, volume 4, pages 2779–2784 vol.4, 2003. doi: 10.1109/ACC.2003.1243743.

[155] Stephen Prajna and Ali Jadbabaie. Safety verification of hybrid systems using barrier certificates. In Rajeev Alur and George J. Pappas, editors, *HSCC*, volume 2993 of *LNCS*, pages 477–492, Heidelberg, 2004. Springer. doi: 10.1007/978-3-540-24743-2_32.

[156] Stephen Prajna and Anders Rantzer. Primal-dual tests for safety and reachability. In Manfred Morari and Lothar Thiele, editors, *HSCC*, volume 3414 of *LNCS*, pages 542–556, Heidelberg, 2005. Springer. doi: 10.1007/978-3-540-31954-2_35.

[157] Stephen Prajna and Anders Rantzer. Convex programs for temporal verification of nonlinear dynamical systems. *SIAM J. Control Optim.*, 46(3):999–1021, 2007. doi: 10.1137/050645178.

[158] Stephen Prajna, Ali Jadbabaie, and George J. Pappas. A framework for worst-case and stochastic safety verification using barrier certificates. *IEEE Trans. Autom. Control.*, 52(8): 1415–1428, 2007. doi: 10.1109/TAC.2007.902736.

[159] Stefan Ratschan and Zhikun She. Providing a basin of attraction to a target region of polynomial systems by computation of Lyapunov-like functions. *SIAM J. Control Optim.*, 48(7):4377–4394, 2010. doi: 10.1137/090749955.

[160] Hadi Ravanbakhsh and Sriram Sankaranarayanan. Counter-example guided synthesis of control Lyapunov functions for switched systems. In *CDC*, pages 4232–4239. IEEE, 2015. doi: 10.1109/CDC.2015.7402879.

[161] Rachid Rebiha, Arnaldo Vieira Moura, and Nadir Matringe. Generating invariants for non-linear hybrid systems. *Theor. Comput. Sci.*, 594:180–200, 2015. doi: 10.1016/j.tcs.2015.06.018.

[162] Daniel Richardson. Some undecidable problems involving elementary functions of a real variable. *J. Symb. Log.*, 33(4):514–520, 1968. doi: 10.2307/2271358.

[163] Daniel Ricketts, Gregory Malecha, Mario M. Alvarez, Vignesh Gowda, and Sorin Lerner. Towards verification of hybrid systems in a foundational proof assistant. In *MEMOCODE*, pages 248–257. IEEE, 2015. DOI: 10.1109/MEMCOD.2015.7340492.

[164] Mauno Rönkkö, Anders P. Ravn, and Kaisa Sere. Hybrid action systems. *Theor. Comput. Sci.*, 290(1):937–973, 2003. DOI: 10.1016/S0304-3975(02)00547-9.

[165] Nicolas Rouche, P. Habets, and M. Laloy. *Stability Theory by Liapunov's Direct Method.* Springer, New York, 1977. DOI: 10.1007/978-1-4684-9362-7.

[166] Damien Rouhling. A formal proof in Coq of a control function for the inverted pendulum. In June Andronick and Amy P. Felty, editors, *CPP*, pages 28–41. ACM, 2018. DOI: 10.1145/3167101.

[167] Pierre Roux, Yuen-Lam Voronin, and Sriram Sankaranarayanan. Validating numerical semidefinite programming solvers for polynomial invariants. *Form. Methods Syst. Des.*, 53 (2):286–312, 2018. DOI: 10.1007/s10703-017-0302-y.

[168] Walter Rudin. *Principles of Mathematical Analysis.* McGraw-Hill, third edition, 1976.

[169] Sriram Sankaranarayanan, Henny B. Sipma, and Zohar Manna. Constructing invariants for hybrid systems. *Form. Methods Syst. Des.*, 32(1):25–55, 2008. DOI: 10.1007/s10703-007-0046-1.

[170] Sriram Sankaranarayanan, Xin Chen, and Erika Ábrahám. Lyapunov function synthesis using Handelman representations. In Sophie Tarbouriech and Miroslav Krstic, editors, *NOLCOS*, pages 576–581. IFAC, 2013. DOI: 10.3182/20130904-3-FR-2041.00198.

[171] Matias Scharager, Katherine Cordwell, Stefan Mitsch, and André Platzer. Verified quadratic virtual substitution for real arithmetic. In Marieke Huisman, Corina S. Pasareanu, and Naijun Zhan, editors, *FM*, volume 13047 of *LNCS*, pages 200–217. Springer, 2021. DOI: 10.1007/978-3-030-90870-6_11.

[172] Stefan Schupp, Erika Ábrahám, Ibtissem Ben Makhlouf, and Stefan Kowalewski. HyPro: A C++ library of state set representations for hybrid systems reachability analysis. In Clark W. Barrett, Misty Davies, and Temesghen Kahsai, editors, *NFM*, volume 10227 of *LNCS*, pages 288–294, 2017. DOI: 10.1007/978-3-319-57288-8_20.

[173] Zhikun She and Bai Xue. Discovering multiple Lyapunov functions for switched hybrid systems. *SIAM J. Control. Optim.*, 52(5):3312–3340, 2014. DOI: 10.1137/130934313.

[174] J. Tanner Slagel, Lauren White, and Aaron Dutle. Formal verification of semi-algebraic sets and real analytic functions. In Catalin Hritcu and Andrei Popescu, editors, *CPP*, pages 278–290. ACM, 2021. DOI: 10.1145/3437992.3439933.

[175] Andrew Sogokon. *Direct methods for deductive verification of temporal properties in continuous dynamical systems.* PhD thesis, Laboratory for Foundations of Computer Science, School of Informatics, University of Edinburgh, 2016.

[176] Andrew Sogokon and Paul B. Jackson. Direct formal verification of liveness properties in continuous and hybrid dynamical systems. In Nikolaj Bjørner and Frank S. de Boer, editors, *FM*, volume 9109 of *LNCS*, pages 514–531, Cham, 2015. Springer. DOI: 10.1007/978-3-319-19249-9_32.

[177] Andrew Sogokon, Khalil Ghorbal, and Taylor T. Johnson. Operational models for piecewise-smooth systems. *ACM Trans. Embed. Comput. Syst.*, 16(5s):185:1–185:19, 2017.

[178] Andrew Sogokon, Khalil Ghorbal, Yong Kiam Tan, and André Platzer. Vector barrier certificates and comparison systems. In Klaus Havelund, Jan Peleska, Bill Roscoe, and Erik P. de Vink, editors, *FM*, volume 10951 of *LNCS*, pages 418–437. Springer, 2018. DOI: 10.1007/978-3-319-95582-7_25.

[179] Andrew Sogokon, Paul B. Jackson, and Taylor T. Johnson. Verifying safety and persistence in hybrid systems using flowpipes and continuous invariants. *J. Autom. Reasoning*, 63(4): 1005–1029, 2019. DOI: 10.1007/s10817-018-9497-x.

[180] Andrew Sogokon, Stefan Mitsch, Yong Kiam Tan, Katherine Cordwell, and André Platzer. Pegasus: Sound continuous invariant generation. *Form. Methods Syst. Des.*, 58:5–41, 2021. DOI: 10.1007/s10703-020-00355-z.

[181] Eduardo D. Sontag and Yuan Wang. On characterizations of the input-to-state stability property. *Systems & Control Letters*, 24(5):351–359, 1995. DOI: https://doi.org/10.1016/0167-6911(94)00050-6.

[182] Eduardo D. Sontag and Yuan Wang. New characterizations of input-to-state stability. *IEEE Trans. Autom. Control.*, 41(9):1283–1294, 1996. DOI: 10.1109/9.536498.

[183] Miriam García Soto and Pavithra Prabhakar. Averist: Algorithmic verifier for stability of linear hybrid systems. In Maria Prandini and Jyotirmoy V. Deshmukh, editors, *HSCC*, pages 259–264. ACM, 2018. DOI: 10.1145/3178126.3178154.

[184] Matthieu Sozeau, Simon Boulier, Yannick Forster, Nicolas Tabareau, and Théo Winterhalter. Coq Coq correct! verification of type checking and erasure for Coq, in Coq. *Proc. ACM Program. Lang.*, 4(POPL):8:1–8:28, 2020. DOI: 10.1145/3371076.

[185] Patrick Speissegger. The Pfaffian closure of an o-minimal structure. *J. Reine Angew. Math.*, 508:189–211, 1999. DOI: 10.1515/crll.1999.026.

[186] Thomas Strathmann and Jens Oehlerking. Verifying properties of an electro-mechanical braking system. In Goran Frehse and Matthias Althoff, editors, *ARCH*, volume 34 of *EPiC Series in Computing*, pages 49–56. EasyChair, 2015. DOI: 10.29007/x87p.

[187] Steven H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering.* Westview Press, Boulder, CO, second edition, 2015.

[188] Kohei Suenaga and Ichiro Hasuo. Programming with infinitesimals: A WHILE-language for hybrid system modeling. In Luca Aceto, Monika Henzinger, and Jirí Sgall, editors, *ICALP*, volume 6756 of *LNCS*, pages 392–403. Springer, 2011. DOI: 10.1007/978-3-642-22012-8_31.

[189] Zhendong Sun and Shuzhi Sam Ge. *Stability Theory of Switched Dynamical Systems.* Communications and Control Engineering. Springer, 2011. DOI: 10.1007/978-0-85729-256-8.

[190] Ankur Taly and Ashish Tiwari. Deductive verification of continuous dynamical systems. In Ravi Kannan and K. Narayan Kumar, editors, *FSTTCS*, volume 4 of *LIPIcs*, pages 383–394, Dagstuhl, 2009. Schloss Dagstuhl. DOI: 10.4230/LIPIcs.FSTTCS.2009.2334.

[191] Ankur Taly and Ashish Tiwari. Switching logic synthesis for reachability. In Luca P. Carloni and Stavros Tripakis, editors, *EMSOFT*, pages 19–28, New York, 2010. ACM. DOI: 10.1145/1879021.1879025.

[192] Yong Kiam Tan and André Platzer. An axiomatic approach to liveness for differential equations. In Maurice H. ter Beek, Annabelle McIver, and José N. Oliveira, editors, *FM*, volume 11800 of *LNCS*, pages 371–388. Springer, 2019. DOI: 10.1007/978-3-030-30942-8_23.

[193] Yong Kiam Tan and André Platzer. Switched systems as hybrid programs. In Raphaël M. Jungers, Necmiye Ozay, and Alessandro Abate, editors, *ADHS*, volume 54 of *IFAC-PapersOnLine*, pages 247–252. Elsevier, 2021. DOI: 10.1016/j.ifacol.2021.08.506.

[194] Yong Kiam Tan and André Platzer. Deductive stability proofs for ordinary differential equations. In Jan Friso Groote and Kim Guldstrand Larsen, editors, *TACAS*, volume 12652 of *LNCS*, pages 181–199. Springer, 2021. DOI: 10.1007/978-3-030-72013-1_10.

[195] Yong Kiam Tan and André Platzer. An axiomatic approach to existence and liveness for differential equations. *Formal Aspects Comput.*, 33(4):461–518, 2021. DOI: 10.1007/s00165-020-00525-0.

[196] Yong Kiam Tan, Stefan Mitsch, and André Platzer. Verifying switched system stability with logic. In Ezio Bartocci and Sylvie Putot, editors, *HSCC*. ACM, 2022. DOI: 10.1145/3501710.3519541. To appear.

[197] Alfred Tarski. *A Decision Method for Elementary Algebra and Geometry.* RAND Corporation, Santa Monica, CA, 1951.

[198] Giuseppina Terzo. *Consequences of Schanuel's Conjecture in Exponential Algebra.* PhD thesis, University of Naples Federico II, 2007.

[199] Gerald Teschl. *Ordinary Differential Equations and Dynamical Systems.* Graduate Studies in Mathematics. American Mathematical Society, 2012.

[200] Ufuk Topcu, Andrew K. Packard, and Peter J. Seiler. Local stability analysis using simulations and sum-of-squares programming. *Autom.*, 44(10):2669–2675, 2008. DOI: 10.1016/j.automatica.2008.03.010.

[201] Jean-Claude Tougeron. Algèbres analytiques topologiquement noethériennes. Théorie de Khovanskiĭ. *Ann. I. Fourier*, 41(4):823–840, 1991. DOI: 10.5802/aif.1275.

[202] Lou van den Dries. Remarks on Tarski's problem concerning (R, +, *, exp). In Gabriele Lolli, Giuseppe Longo, and Annalisa Marcja, editors, *Logic Colloquium '82*, volume 112, pages 97–121. North-Holland, Amsterdam, 1984. doi: 10.1016/S0049-237X(08)71811-1.

[203] Lou van den Dries. *Tame Topology and O-Minimal Structures*. Cambridge University Press, Cambridge, 1998. doi: 10.1017/CBO9780511525919.

[204] Wolfgang Walter. *Ordinary Differential Equations*. Springer, New York, 1998. doi: 10.1007/978-1-4612-0601-9.

[205] Qiuye Wang, Mingshuai Chen, Bai Xue, Naijun Zhan, and Joost-Pieter Katoen. Synthesizing invariant barrier certificates via difference-of-convex programming. In Alexandra Silva and K. Rustan M. Leino, editors, *CAV*, volume 12759 of *LNCS*, pages 443–466. Springer, 2021. doi: 10.1007/978-3-030-81685-8_21.

[206] Shuling Wang, Naijun Zhan, and Liang Zou. An improved HHL prover: An interactive theorem prover for hybrid systems. In Michael J. Butler, Sylvain Conchon, and Fatiha Zaïdi, editors, *ICFEM*, volume 9407 of *LNCS*, pages 382–399. Springer, 2015. doi: 10.1007/978-3-319-25423-4_25.

[207] Alex J. Wilkie. A theorem of the complement and some new o-minimal structures. *Sel. Math. New Ser.*, 5(4):397–421, 1999. doi: 10.1007/S000290050052.

[208] Jonathan Julián Huerta y Munive and Georg Struth. Verifying hybrid systems with modal Kleene algebra. In Jules Desharnais, Walter Guttmann, and Stef Joosten, editors, *RAMiCS*, volume 11194 of *LNCS*, pages 225–243. Springer, 2018. doi: 10.1007/978-3-030-02149-8_14.

[209] Zhengfeng Yang, Chao Huang, Xin Chen, Wang Lin, and Zhiming Liu. A linear programming relaxation based approach for generating barrier certificates of hybrid systems. In John S. Fitzgerald, Constance L. Heitmeyer, Stefania Gnesi, and Anna Philippou, editors, *FM*, volume 9995 of *LNCS*, pages 721–738, 2016. doi: 10.1007/978-3-319-48989-6_44.

[210] Xia Zeng, Wang Lin, Zhengfeng Yang, Xin Chen, and Lilei Wang. Darboux-type barrier certificates for safety verification of nonlinear hybrid systems. In Petru Eles and Rahul Mangharam, editors, *EMSOFT*, pages 11:1–11:10. ACM, 2016. doi: 10.1145/2968478.2968484.

[211] Guisheng Zhai, Bo Hu, Kazunori Yasuda, and Anthony N. Michel. Stability analysis of switched systems with stable and unstable subsystems: An average dwell time approach. *Int. J. Syst. Sci.*, 32(8):1055–1061, 2001. doi: 10.1080/00207720116692.

[212] Bohua Zhan, Bin Gu, Xiong Xu, Xiangyu Jin, Shuling Wang, Bai Xue, Xiaofeng Li, Yao Chen, Mengfei Yang, and Naijun Zhan. Brief industry paper: Modeling and verification of descent guidance control of Mars lander. In *RTAS*, pages 457–460. IEEE, 2021. doi: 10.1109/RTAS52030.2021.00051.

[213] Naijun Zhan, Shuling Wang, and Hengjun Zhao. Formal modelling, analysis and verification of hybrid systems. In Zhiming Liu, Jim Woodcock, and Huibiao Zhu, editors, *Unifying Theories of Programming and Formal Engineering Methods*, volume 8050 of *LNCS*, pages 207–281. Springer, 2013. DOI: 10.1007/978-3-642-39721-9_5.

[214] Jun Zhang, Karl Henrik Johansson, John Lygeros, and Shankar Sastry. Zeno hybrid systems. *Int. J. Robust Nonlinear Control.*, 11(5):435–451, 2001. DOI: 10.1002/rnc.592.

Appendix A

Appendix: Safety and Invariance for Ordinary Differential Equations

A.1 Differential Dynamic Logic Axiomatization

A.1.1 Extended Axiomatization Soundness

This section proves the soundness of the axiomatic extension from Section 3.5. For the solution $\varphi : [0, T] \to \mathbb{S}$, its truncation to the interval $[0, t]$ for some $0 \leq t \leq T$ is denoted $\varphi|_t : [0, t] \to \mathbb{S}$, with $\varphi|_t(\zeta) = \varphi(\zeta)$ for $\zeta \in [0, t]$. The shorthand notation $\varphi([a, b]) \in [\![P]\!]$ means $\varphi(\zeta) \in [\![P]\!]$ for all $a \leq \zeta \leq b$, where the interval $[a, b]$ is required to be a closed subinterval of the interval $[0, T]$. Analogously, $\varphi((a, b))$ is used when the interval is open, and similarly for the half-open cases.

As explained in Section 3.5, the soundness of the extended axioms requires that the ODE system $x' = f(x)$ always locally evolves x. An easy syntactic way to ensure this condition is to check that the system already contains an equation $x_1' = 1$ to track the passage of time, which can be added using axiom DG if necessary before using the axioms. However, the soundness proofs below are more general and only use the assumption that the ODE system locally evolves x, whether by $x_1' = 1$ or otherwise.

The soundness proofs make use of dL's coincidence lemmas [142, Lemmas 10,11]:

Lemma A.1 (Coincidence for terms and formulas [142]). *The following coincidence properties hold for* dL*, where free variables FV*(e)*, FV*(ϕ) *are defined as expected [142, Sections 2.3 and 2.4].*

- *If the states ω, ν agree on the free variables of term e (FV(e)), then $\omega[\![e]\!] = \nu[\![e]\!]$.*
- *If the states ω, ν agree on the free variables of formula ϕ (FV(ϕ)), then $\omega \in [\![\phi]\!]$ iff $\nu \in [\![\phi]\!]$.*

Existence, Uniqueness, and Continuity. First, the axioms from Lemma 3.15 internalizing basic existence and uniqueness properties of solutions of differential equations are proved sound.

Proof of Lemma 3.15. Let ω be an arbitrary initial state. When interpreted as a function of the variables x, the RHS $f(x)$ of the ODE system $x' = f(x)$ is continuously differentiable. By the Picard-Lindelöf theorem [204, §10.VI], from ω, there is an interval $[0, \tau), \tau > 0$ on which there is a unique, continuous solution $\varphi : [0, \tau) \to \mathbb{S}$ with $\varphi(0) = \omega$ on $\{x'\}^{\complement}$. The solution can be uniquely extended in time up to its right-maximal open interval of existence [204, §10.IX].

Uniq The "$\to$" direction follows directly from monotonicity of domain constraints because of the propositional tautology $Q_1 \land Q_2 \to Q_1$ (and similarly for Q_2). For the "$\leftarrow$" direction, suppose that initial state ω satisfies both conjuncts with $\omega \in [\![\langle x' = f(x) \,\&\, Q_1\rangle P]\!]$ and $\omega \in [\![\langle x' = f(x) \,\&\, Q_2\rangle P]\!]$. Expanding the definition of the diamond modality, there exist two solutions $\varphi_1 : [0, T_1] \to \mathbb{S}$, $\varphi_2 : [0, T_2] \to \mathbb{S}$ from ω such that $\varphi_1 \models x' = f(x) \land Q_1$ and $\varphi_2 \models x' = f(x) \land Q_2$, with both $\varphi_1(T_1) \in [\![P]\!]$ and $\varphi_2(T_2) \in [\![P]\!]$. Suppose $T_1 \leq T_2$. Since $\varphi_2([0, T_2]) \in [\![Q_2]\!]$ and, by uniqueness, φ_1 is a truncation of φ_2 to a smaller existence interval, $\varphi_1 \models x' = f(x) \land (Q_1 \land Q_2)$. At time T_1, the solution satisfies $\varphi_1(T_1) \in [\![P]\!]$, so $\omega \in [\![\langle x' = f(x) \,\&\, Q_1 \land Q_2\rangle P]\!]$, as required. The case for $T_2 < T_1$ is similar, except with $\varphi_2 \models x' = f(x) \land (Q_1 \land Q_2)$ and satisfying $\varphi_2(T_2) \in [\![P]\!]$ at time T_2 instead.

Cont Assume that ω satisfies the outermost implication, i.e., $\omega \in [\![x = y]\!]$. The (inner) "$\to$" direction follows by definition because in order for there to be a solution staying in $e > 0$ at all, the initial state ω must already satisfy $e > 0$ (evolution domains are differential-free). For the (inner) "$\leftarrow$" direction, suppose further that $\omega \in [\![e > 0]\!]$. Since x' is not a free variable of term e as e is differential-free (Section 2.3.3), coincidence (Lemma A.1) implies $\varphi(0) \in [\![e > 0]\!]$. As a composition of continuous evaluation [142, Definition 5] with the continuous solution φ, $\varphi(t)[\![e]\!]$ is a continuous function of time t. Thus, $\varphi(0) \in [\![e > 0]\!]$ implies $\varphi([0, T]) \in [\![e > 0]\!]$ for some $0 < T \leq \tau$ and the truncated solution $\varphi|_T$ satisfies $\varphi|_T \models x' = f(x) \land e > 0$. Since y is constant for the ODE but $x' = f(x)$ was assumed to locally evolve (for example with $x'_1 = 1$), there is a time $0 < \varepsilon \leq T$ at which $\varphi(\varepsilon) \in [\![x \neq y]\!]$. The truncation $\varphi|_\varepsilon$ witnesses $\omega \in [\![\langle x' = f(x) \,\&\, e > 0\rangle x \neq y]\!]$.

Dadj The "$\leftarrow$" direction follows immediately from the "$\to$" direction by swapping the names x, y, because $-(-f(x)) = f(x)$. Therefore, it suffices to prove the "$\to$" direction. Suppose $\omega \in [\![\langle x' = f(x) \,\&\, Q(x)\rangle\, x = y]\!]$. Unfolding the semantics, there is a solution $\varphi : [0, T] \to \mathbb{S}$, of the system $x' = f(x)$, with $\varphi(0) = \omega$ on $\{x'\}^{\complement}$, with $\varphi([0, T]) \in [\![Q(x)]\!]$ and $\varphi(T) \in [\![x = y]\!]$. Since the variables y do not appear in the differential equations $x' = f(x)$, their values are constant along the solution φ. Consider the time- and variable-reversal $\psi : [0, T] \to \mathbb{S}$, where:

$$\psi(\tau)(z) \overset{\text{def}}{=} \begin{cases} \varphi(T - \tau)(x_i) & z = y_i \\ -\varphi(T - \tau)(x'_i) & z = y'_i \\ \omega(z) & \text{otherwise} \end{cases}$$

By construction, $\psi(0)$ agrees with ω on $\{y'\}^{\complement}$ because $\varphi(T) \in [\![x = y]\!]$. The signs of the differential variables y'_i are negated along ψ. By uniqueness, the solutions of $x' = -f(x)$ are the time-reversed solutions of $x' = f(x)$. As constructed, ψ is the time-reversed solution for $x' = f(x)$ except the x were replaced by y instead. Moreover, since $\varphi([0, T]) \in [\![Q(x)]\!]$, by construction and coincidence (Lemma A.1), $\psi([0, T]) \in [\![Q(y)]\!]$. Therefore, $\psi \models y' = -f(y) \land Q(y)$. Finally, observe that $\psi(T)(y) = \varphi(0)(x)$, but ψ holds the values of x constant, thus $\psi(T)(x) = \omega(x) = \varphi(0)(x)$ and so $\psi(T) \in [\![y = x]\!]$ and ψ witnesses $\omega \in [\![\langle y' = -f(y) \,\&\, Q(y)\rangle y = x]\!]$ □

Real Induction. The following real induction axiom with domain constraints is proved sound. Axiom RI from Lemma 3.19 follows as an instance with no domain constraint, i.e., $Q \equiv \mathit{true}$.

$$\text{RI\&}\quad [x' = f(x) \,\&\, Q]P \leftrightarrow \forall y\, [x' = f(x) \,\&\, Q \wedge (P \vee x{=}y)]\Big(x{=}y \rightarrow \underbrace{P}_{\textcircled{a}} \wedge \underbrace{\big(\langle x' = f(x) \,\&\, Q \vee x{=}y\rangle x \neq y \rightarrow \langle x' = f(x) \,\&\, P \vee x{=}y\rangle x \neq y\big)}_{\textcircled{b}}\Big)$$

Similar to axiom RI, the axiom RI& is based on the real induction principle [34] but also accounts for an arbitrary domain constraint Q. Its RHS conjuncts labeled ⓐ and ⓑ correspond to ① and ② in Def. 3.17 respectively. The quantification $\forall y\,[\dots \,\&\, Q]\big(x = y \rightarrow \dots\big)$ now only considers final states ($x = y$) reachable by trajectories that *always* stay within Q, and within P except possibly at the endpoint $x = y$. The conjunct ⓐ expresses that P is still true at such an endpoint. The conjunct ⓑ expresses that P continues to remain true locally but only when Q itself remains true locally. This added assumption for Q corresponds to the "If $\zeta < b$ then . . . " assumption in ② of Def. 3.17. The conjunct ⓑ can be rewritten succinctly with the local progress $\bigcirc$ modality as:

$$\langle x' = f(x) \,\&\, Q\rangle \bigcirc \rightarrow \langle x' = f(x) \,\&\, P\rangle \bigcirc$$

With completeness for local progress (Theorem 3.26), this gives a first hint at how RI& will be used to obtain a complete proof rule for semianalytic invariants with domain constraints in Appendix A.2.

Lemma A.2 (Real induction with domain constraints)**.** *The real induction axiom RI& is sound, where y is fresh in $[x' = f(x) \,\&\, Q]P$.*

Proof (implies Lemma 3.19). The conjuncts on the RHS of RI& are labeled as ⓐ and ⓑ respectively, as shown above. Consider an initial state ω, both directions of the axiom are proved separately.

"$\rightarrow$" Assume the LHS of RI& is true initially with ⊛ $\omega \in [\![[x' = f(x) \,\&\, Q]P]\!]$. Unfolding the quantification and box modality on the RHS, let ω_y be identical to ω except where the values for y are replaced with any arbitrary values $d \in \mathbb{R}^n$. Consider any solution $\varphi_y : [0, T] \rightarrow \mathbb{S}$ where $\varphi_y \models x' = f(x) \wedge \big(Q \wedge (P \vee x = y)\big)$, $\varphi_y(0) = \omega_y$ on $\{x'\}^\complement$, and $\varphi_y(T) \in [\![x = y]\!]$. The following similar solution $\varphi : [0, T] \rightarrow \mathbb{S}$ keeps y constant at their initial values in ω:

$$\varphi(t)(z) \overset{\text{def}}{=} \begin{cases} \varphi_y(t)(z) & z \in \{y\}^\complement \\ \omega(z) & z \in \{y\} \end{cases}$$

By construction, $\varphi(0)$ is identical to ω on $\{x'\}^\complement$ and φ is identical to φ_y on $\{y\}^\complement$. Since y is fresh in $x' = f(x) \,\&\, Q$, by coincidence (Lemma A.1) the latter implies that $\varphi \models x' = f(x) \wedge Q$. By assumption ⊛, $\varphi(T) \in [\![P]\!]$, which implies that $\varphi_y(T) \in [\![P]\!]$ by coincidence (Lemma A.1) since y is fresh in P. This proves conjunct ⓐ. Unfolding the implication and diamond modality of conjunct ⓑ, assume there is another solution $\psi_y : [0, \tau] \rightarrow \mathbb{S}$

from $\varphi_y(T)$ with $\psi_y \models x' = f(x) \land (Q \lor x = y)$ and $\psi_y(\tau) \in \llbracket x \neq y \rrbracket$. Note that $\psi_y(0) = \varphi_y(T)$ *exactly* rather than just on $\{x'\}^{\complement}$, because both states have the same values for the differential variables. To show the RHS of the implication in ⓑ, i.e., that $\varphi_y(T) \in \llbracket \langle x' = f(x) \,\&\, P \lor x = y \rangle x \neq y \rrbracket$, it suffices to show: $\psi_y \models x' = f(x) \land P$, because P propositionally implies $P \lor x = y$. In particular, since ψ_y already satisfies the requisite differential equations and $\psi_y(\tau) \in \llbracket x \neq y \rrbracket$, it remains to show that ψ_y stays in the evolution domain P for its entire duration, i.e., $\psi_y([0,\tau]) \in \llbracket P \rrbracket$. Let $0 \leq \zeta \leq \tau$ and consider the concatenated solution $\Phi : [0, T+\zeta] \to \mathbb{S}$ defined by:

$$\Phi(t)(z) \overset{\text{def}}{=} \begin{cases} \varphi_y(t)(z) & t \leq T, z \in \{y\}^{\complement} \\ \psi_y(t-T)(z) & t > T, z \in \{y\}^{\complement} \\ \omega(z) & z \in \{y\} \end{cases}$$

As with φ, the solution Φ is constructed to keep y constant at their initial values in ω. Since ψ_y must uniquely extend φ_y [204, §10.IX], the concatenated solution Φ is a solution starting from ω, solving the system $x' = f(x)$. It stays in Q for its entire duration by coincidence (Lemma A.1) because $\varphi_y(T) \in \llbracket Q \rrbracket$ and all states satisfying $x = y$ agree with $\varphi_y(T)$ on the free variables of formula Q. In other words, $\Phi \models x' = f(x) \land Q$. By ⍟, $\Phi(T+\zeta) \in \llbracket P \rrbracket$, which implies $\psi(\zeta) \in \llbracket P \rrbracket$ by coincidence (Lemma A.1) and so $\psi_y([0,\zeta]) \in \llbracket P \rrbracket$, as required.

"←" Assume the RHS of RI& is true in initial state ω and show the LHS. Consider an arbitrary solution $\varphi : [0,T] \to \mathbb{S}$ starting from ω such that $\varphi \models x' = f(x) \land Q$. To show $\varphi([0,T]) \in \llbracket P \rrbracket$, using the real induction principle (Proposition 3.18), it suffices to show that the set of times $S \overset{\text{def}}{=} \{\zeta : \varphi(\zeta) \in \llbracket P \rrbracket\}$ is an inductive subset of $[0,T]$, i.e., it satisfies properties ① and ② in Def. 3.17. So, assume that $[0,\zeta) \subseteq S$ for some time $0 \leq \zeta \leq T$.

The proof instantiates quantified variables y on the RHS of RI& to match the values of x at $\varphi(\zeta)$. Since y is constant for the ODE, this allows properties of $\varphi(\zeta)$ to be deduced using the RHS (namely ⓐ, ⓑ) by mediating between φ and its augmentation φ_y below. More precisely, consider the state ω_y identical to ω, except where the values for variables y are replaced with the corresponding values of x in $\varphi(\zeta)$. Correspondingly, consider the solution $\varphi_y : [0,\zeta] \to \mathbb{S}$ identical to φ but which keeps y constant at those initial values in ω_y rather than in ω:

$$\omega_y(z) \overset{\text{def}}{=} \begin{cases} \omega(z) & z \in \{y\}^{\complement} \\ \varphi(\zeta)(x_i) & z = y_i \end{cases} \qquad \varphi_y(t)(z) \overset{\text{def}}{=} \begin{cases} \varphi(t)(z) & z \in \{y\}^{\complement} \\ \omega_y(z) & z \in \{y\} \end{cases}$$

By construction and coincidence (Lemma A.1), φ_y is a solution from initial state ω_y, solving $\varphi_y \models x' = f(x) \land Q$ and $\varphi_y(\zeta) \in \llbracket x = y \rrbracket$. By assumption and coincidence (Lemma A.1), $\varphi_y([0,\zeta)) \in \llbracket P \rrbracket$. Therefore, $\varphi_y([0,\zeta]) \in \llbracket Q \land (P \lor x = y) \rrbracket$. Unfolding the quantification, box modality and implication on the RHS yields $\varphi_y(\zeta) \in \llbracket ⓐ \land ⓑ \rrbracket$.

① By ⓐ, $\varphi_y(\zeta) \in \llbracket P \rrbracket$ so by coincidence (Lemma A.1), $\varphi(\zeta) \in \llbracket P \rrbracket$ as required for ①.

② Further assume that $\zeta < T$ and show $\varphi((\zeta, \zeta+\varepsilon]) \in \llbracket P \rrbracket$ for some $\varepsilon > 0$. Observe that since $\zeta < T$, there is a solution that extends from state $\varphi(\zeta)$, i.e., $\psi : [0, T-\zeta] \to$

$\mathbb{S}$, where $\psi(\tau) \overset{\text{def}}{=} \varphi(\tau+\zeta)$ and with $\psi \models x' = f(x) \wedge Q$. Construct the corresponding solution $\psi_y : [0, T-\zeta] \to \mathbb{S}$ that extends from state $\varphi_y(\zeta)$ and still keeps y constant at their values in ω_y:

$$\psi_y(t)(z) \overset{\text{def}}{=} \begin{cases} \psi(t)(z) & z \in \{y\}^{\complement} \\ \varphi_y(\zeta)(z) & z \in \{y\} \end{cases}$$

By coincidence (Lemma A.1), $\psi_y \models x' = f(x) \wedge Q$, so by weakening the domain constraint, $\psi_y \models x' = f(x) \wedge (Q \vee x = y)$. Since $\varphi_y(\zeta) \in [\![x = y]\!]$ by construction and the differential equation is assumed to always locally evolve (for example with $x_1' = 1$), there must be some duration $0 < \delta < T - \zeta$ (recall $T - \zeta > 0$) after which the value of x has changed from its initial value held constant in y, i.e., $\psi_y(\delta) \in [\![x \neq y]\!]$. The truncation $\psi_y|_\delta$ witnesses the LHS of the implication in ⓑ with: $\varphi_y(\zeta) \in [\![\langle x' = f(x) \,\&\, Q \vee x = y\rangle x \neq y]\!]$. Using this with the implication in ⓑ yields $\varphi_y(\zeta) \in [\![\langle x' = f(x) \,\&\, P \vee x = y\rangle x \neq y]\!]$. Unfolding the semantics, this gives a solution which, by uniqueness, is a truncation $\psi_y|_\varepsilon$ of ψ_y, for some $\varepsilon > 0$, that satisfies $\psi_y|_\varepsilon([0,\varepsilon]) \in [\![P \vee x = y]\!]$. From ⓐ and coincidence (Lemma A.1), all states satisfying $x = y$ agree with $\varphi(\zeta)$ on the free variables of formula P thus $\psi_y|_\varepsilon([0,\varepsilon]) \in [\![P]\!]$. By construction, $\psi_y|_\varepsilon(\tau)$ coincides with $\varphi(\tau + \zeta)$ on x for all $0 \leq \tau \leq \varepsilon$, which implies $\varphi((\zeta, \zeta+\varepsilon]) \in [\![P]\!]$ by Lemma A.1. □

Conjunct ⓑ of RI& can be written as $\langle x' = f(x) \,\&\, Q\rangle x \neq y \to \langle x' = f(x) \,\&\, P\rangle x \neq y$ because Q and P can be assumed true in the context where the conjunct appears. This flexibility can be seen from its soundness proof above and will be made explicit syntactically in Corollary A.3.

A.1.2 Extended Derived Rules and Axioms

This section derives additional rules and axioms that make use of the axiomatic extensions from Section 3.5.

Local Progress Properties. The local progress modality $\bigcirc$ excludes the initial state ($x = y$) in the domain constraint when expressing local progress for formula Q. Recall:

$$\langle x' = f(x) \,\&\, Q\rangle \bigcirc \overset{\text{def}}{\equiv} \langle x' = f(x) \,\&\, Q \vee x = y\rangle\, x \neq y$$

The disjunct $x = y$ in the domain constraint makes local progress an interesting question for formulas characterizing sets that are not topologically closed (e.g., open sets as characterized by the formula $e > 0$). As axiom Cont shows, the formula $\langle x' = f(x) \,\&\, e > 0\rangle x \neq y$ which *does not* exclude $x = y$ in the evolution domain constraint is already *equivalent* to $e > 0$. A precise syntactic characterization of this difference is shown by the following derived axiom.

Corollary A.3 (Initial state inclusion). *The following axiom derives in* dL. *Variables y are fresh in the ODE $x' = f(x)$ and formula Q.*

$$\text{Init}\quad x{=}y \to \big(\langle x' = f(x) \,\&\, Q\rangle x \neq y \leftrightarrow Q \wedge \langle x' = f(x) \,\&\, Q\rangle \bigcirc\big)$$

Proof. First, by dualizing via $\langle\cdot\rangle$ both sides of axiom DX, the following equivalence is derived:

$$\langle x' = f(x) \,\&\, Q\rangle P \leftrightarrow (Q \land (P \lor \langle x' = f(x) \,\&\, Q\rangle P))$$

The derivation of Init starts by using this derived equivalence (DX, $\langle\cdot\rangle$), followed by a series of equivalent propositional rewrites that simplify the logical structure of the succedent. The propositional steps are shown below, first removing the disjunct $x \neq y$ using the assumption $x{=}y$, and then pulling out the common conjunct Q as an antecedent assumption.

$$\text{DX, }\langle\cdot\rangle\,\dfrac{\dfrac{\dfrac{x{=}y, Q \vdash \langle x' = f(x) \,\&\, Q\rangle x \neq y \leftrightarrow \langle x' = f(x) \,\&\, Q\rangle \bigcirc}{x{=}y \vdash Q \land \langle x' = f(x) \,\&\, Q\rangle x \neq y \leftrightarrow Q \land \langle x' = f(x) \,\&\, Q\rangle \bigcirc}}{x{=}y \vdash Q \land (x \neq y \lor \langle x' = f(x) \,\&\, Q\rangle x \neq y) \leftrightarrow Q \land \langle x' = f(x) \,\&\, Q\rangle \bigcirc}}{x{=}y \vdash \langle x' = f(x) \,\&\, Q\rangle x \neq y \leftrightarrow Q \land \langle x' = f(x) \,\&\, Q\rangle \bigcirc}$$

Both directions of the resulting equivalence are proved separately by unfolding the abbreviation $\bigcirc$. In the "$\rightarrow$" direction, a dRW$\langle\cdot\rangle$ step suffices using the tautology $Q \rightarrow Q \lor x{=}y$:

$$\text{dRW}\langle\cdot\rangle\,\dfrac{*}{x{=}y, Q, \langle x' = f(x) \,\&\, Q\rangle x \neq y \vdash \langle x' = f(x) \,\&\, Q \lor x{=}y\rangle x \neq y}$$

In the "$\leftarrow$" direction, the derivation starts with a DR$\langle\cdot\rangle$ step which reduces to the box modality. Since the formulas $x = y$ and Q are true initially, a V, DC step introduces the constant assumption $Q(y)$ into the domain constraint, which is Q with y in place of x. The derivation closes with dW using the strengthened domain constraint.

$$\text{DR}\langle\cdot\rangle\,\dfrac{\text{V, DC}\,\dfrac{\text{dW}\,\dfrac{*}{\vdash [x' = f(x) \,\&\, (Q \lor x{=}y) \land Q(y)]Q}}{x{=}y, Q \vdash [x' = f(x) \,\&\, Q \lor x{=}y]Q}}{x{=}y, Q, \langle x' = f(x) \,\&\, Q \lor x{=}y\rangle x \neq y \vdash \langle x' = f(x) \,\&\, Q\rangle x \neq y}$$

□

It is not possible to locally progress into both formula P and its negation $\neg P$ simultaneously, by uniqueness. This is the "$\rightarrow$" direction of the duality axiom $\neg\bigcirc$ for local progress from Corollary 3.27. The converse "$\leftarrow$" direction is more involved and relies on the characterization axiom LP to be derived later.

Corollary A.4 (Local progress duality "$\rightarrow$"). *The following axiom derives from Uniq. Variables y are fresh in the ODE $x' = f(x)$ and formula P.*

$$\neg\bigcirc_{\rightarrow}\;\; x{=}y \rightarrow \big(\langle x' = f(x) \,\&\, P\rangle \bigcirc \rightarrow \neg\langle x' = f(x) \,\&\, \neg P\rangle \bigcirc\big)$$

Proof. The derivation starts with $\neg$R, after which the resulting local progress antecedents are combined by axiom Uniq, giving a conjunction of their domain constraints because the formula $(P \lor x{=}y) \land (\neg P \lor x{=}y)$ is propositionally equivalent to $(P \land \neg P) \lor x{=}y$. The conjunction $P \land \neg P$ in the domain constraint is propositionally equivalent to *false* and, intuitively, no local progress is possible into an empty set of states.

$$\neg\text{R}\,\dfrac{\text{Uniq}\,\dfrac{\langle x' = f(x) \,\&\, P \land \neg P\rangle \bigcirc \vdash \mathit{false}}{\langle x' = f(x) \,\&\, P\rangle \bigcirc, \langle x' = f(x) \,\&\, \neg P\rangle \bigcirc \vdash \mathit{false}}}{\langle x' = f(x) \,\&\, P\rangle \bigcirc \vdash \neg\langle x' = f(x) \,\&\, \neg P\rangle \bigcirc}$$

The derivation is completed by unfolding the $\bigcirc$ syntactic abbreviation, and shifting to the box modality by $\langle\cdot\rangle$ duality. The final step after using dW is a propositional tautology:

$$\dfrac{\text{⟨·⟩, ¬L}\;\dfrac{\text{dW}\;\dfrac{\dfrac{*}{(P \land \neg P) \lor x = y \vdash x = y}}{\vdash [x' = f(x) \,\&\, (P \land \neg P) \lor x = y]x = y}}{\langle x' = f(x) \,\&\, (P \land \neg P) \lor x = y\rangle x \neq y \vdash \mathit{false}}}{\langle x' = f(x) \,\&\, P \land \neg P\rangle \bigcirc \vdash \mathit{false}} \qquad \square$$

Reflection. The next two derived axioms rfl$\langle\cdot\rangle$ and rfl internalize a mathematical property of ODE invariants, namely, the formula P is invariant for the forward ODE $x' = f(x)$ iff its negation $\neg P$ is invariant for the backward ODE $x' = -f(x)$. This invariant reflection principle is used in Appendix A.2 for proving completeness for semianalytic invariants and to flip the signs in the second premise of rule rI. It is useful in its own right as it allows freely switching between proving invariance for either the forward or backward ODEs, e.g., if one direction yields simpler arithmetic.

Corollary A.5 (Reflection)**.** *The reflection axioms rfl⟨·⟩, rfl derive from Dadj:*

$$\text{rfl}\langle\cdot\rangle\quad \exists x\, (P(x) \land \langle x' = f(x) \,\&\, Q(x)\rangle R(x)) \leftrightarrow \exists x\, (R(x) \land \langle x' = -f(x) \,\&\, Q(x)\rangle P(x))$$

$$\text{rfl}\quad \forall x\, \big(P(x) \to [x' = f(x) \,\&\, Q(x)]P(x)\big) \leftrightarrow \forall x\, \big(\neg P(x) \to [x' = -f(x) \,\&\, Q(x)]\neg P(x)\big)$$

Proof. Axiom rfl derives from rfl$\langle\cdot\rangle$ by instantiating with $R(x) \overset{\text{def}}{\equiv} \neg P(x)$ and negating both sides of the equivalence with $\langle\cdot\rangle$. The diamond reflection axiom rfl$\langle\cdot\rangle$ is derived from Dadj. Both implications are proved separately and the "$\leftarrow$" direction follows by instantiating the proof of the "$\rightarrow$" direction, since $-(-f(x)) = f(x)$. The "$\rightarrow$" direction is proved below.

In the derivation below, the formulas are bound renamed [142] for clarity. After Skolemizing, the first K$\langle\cdot\rangle$, dW step introduces an existentially quantified y under the diamond modality in the antecedent by monotonicity using the provable first-order formula $R(x) \to \exists y\, (x = y \land R(y))$.

$$\text{∃L, ∧L}\;\dfrac{\text{K⟨·⟩, dW}\;\dfrac{P(x), \langle x' = f(x) \,\&\, Q(x)\rangle \exists y\, (x = y \land R(y)) \vdash \exists y\, (R(y) \land \langle y' = -f(y) \,\&\, Q(y)\rangle P(y))}{P(x), \langle x' = f(x) \,\&\, Q(x)\rangle R(x) \vdash \exists y\, (R(y) \land \langle y' = -f(y) \,\&\, Q(y)\rangle P(y))}}{\exists x\, (P(x) \land \langle x' = f(x) \,\&\, Q(x)\rangle R(x)) \vdash \exists y\, (R(y) \land \langle y' = -f(y) \,\&\, Q(y)\rangle P(y))}$$

The ODE Barcan B′ axiom moves the existentially quantified y out of the diamond modality since y is not in $x' = f(x)$. A subsequent V step also moves the postcondition $R(y)$ out from the diamond modality into the antecedents.

$$\text{B}'\;\dfrac{\text{∃L}\;\dfrac{\text{V}\;\dfrac{P(x), R(y), \langle x' = f(x) \,\&\, Q(x)\rangle x = y \vdash \exists y\, (R(y) \land \langle y' = -f(y) \,\&\, Q(y)\rangle P(y))}{P(x), \langle x' = f(x) \,\&\, Q(x)\rangle (x = y \land R(y)) \vdash \exists y\, (R(y) \land \langle y' = -f(y) \,\&\, Q(y)\rangle P(y))}}{P(x), \exists y\, \langle x' = f(x) \,\&\, Q(x)\rangle (x = y \land R(y)) \vdash \exists y\, (R(y) \land \langle y' = -f(y) \,\&\, Q(y)\rangle P(y))}}{P(x), \langle x' = f(x) \,\&\, Q(x)\rangle \exists y\, (x = y \land R(y)) \vdash \exists y\, (R(y) \land \langle y' = -f(y) \,\&\, Q(y)\rangle P(y))}$$

The derivation continues using differential adjoints Dadj to syntactically flip the antecedent differential equations from evolving x forward to evolving y backward. The V, K$\langle\cdot\rangle$ step then strengthens the postcondition to $P(y)$ exploiting that the (negated) ODE does not modify x so that $P(x)$ remains true along the ODE. This completes the proof using y as a witness for $\exists y$.

$$\text{Dadj}\;\dfrac{\text{V, K}\langle\cdot\rangle\;\dfrac{\exists\text{R}\;\dfrac{*}{R(y), \langle y' = -f(y)\,\&\,Q(y)\rangle P(y) \vdash \exists y\,(R(y) \wedge \langle y' = -f(y)\,\&\,Q(y)\rangle P(y))}}{P(x), R(y), \langle y' = -f(y)\,\&\,Q(y)\rangle y = x \vdash \exists y\,(R(y) \wedge \langle y' = -f(y)\,\&\,Q(y)\rangle P(y))}}{P(x), R(y), \langle x' = f(x)\,\&\,Q(x)\rangle x = y \vdash \exists y\,(R(y) \wedge \langle y' = -f(y)\,\&\,Q(y)\rangle P(y))} \qquad \square$$

Real Induction Rule. The real induction rule with domain constraints corresponding to axiom RI& is derived next. It is stated with the $\bigcirc$ modality from Section 3.5. The real induction rule rI from Corollary 3.20 derives as an instance with domain constraint $Q \equiv \mathit{true}$.

Corollary A.6 (Real induction rule with domain constraints). *The real induction proof rule rI& (with two stacked premises) derives from RI&, Dadj, Uniq. Variables y are fresh in the ODE $x' = f(x)$ and formulas P, Q.*

$$\text{rI\&}\;\dfrac{\begin{array}{c} x{=}y, P, Q, \langle x' = f(x)\,\&\,Q\rangle\bigcirc \vdash \langle x' = f(x)\,\&\,P\rangle\bigcirc \\ x{=}y, \neg P, Q, \langle x' = -f(x)\,\&\,Q\rangle\bigcirc \vdash \langle x' = -f(x)\,\&\,\neg P\rangle\bigcirc \end{array}}{P \vdash [x' = f(x)\,\&\,Q]P}$$

Proof (implies Corollary 3.20). The derivation starts by rewriting the succedent with RI&, the resulting right conjunct is abbreviated with $R \overset{\text{def}}{\equiv} \langle x' = f(x)\,\&\,Q\rangle\bigcirc \rightarrow \langle x' = f(x)\,\&\,P\rangle\bigcirc$. The M[′] step rewrites the postcondition with propositional tautology $P \wedge R \leftrightarrow P \wedge (P \rightarrow R)$ which allows the left conjunct P to be assumed when proving the right conjunct R (the implication $x{=}y$ is also distributed over the conjunction). The two conjuncts are then split by $[\cdot]\wedge$, $\wedge$R, with the resulting two premises labeled ① and ② respectively. These are shown and proved below.

$$\text{RI\&}\;\dfrac{\forall\text{R}\;\dfrac{\text{M}[']\;\dfrac{[\cdot]\wedge,\wedge\text{R}\;\dfrac{\textcircled{1} \qquad \textcircled{2}}{P \vdash [x' = f(x)\,\&\,Q \wedge (P \vee x{=}y)]\big((x{=}y \rightarrow P) \wedge (x{=}y \wedge P \rightarrow R)\big)}}{P \vdash [x' = f(x)\,\&\,Q \wedge (P \vee x{=}y)](x{=}y \rightarrow P \wedge R)}}{P \vdash \forall y\,[x' = f(x)\,\&\,Q \wedge (P \vee x{=}y)](x{=}y \rightarrow P \wedge R)}}{P \vdash [x' = f(x)\,\&\,Q]P}$$

The premise ② yields the top premise of rule rI& directly (unfolding the abbreviation for R):

$$\text{dW}\;\dfrac{\rightarrow\text{R}, \wedge\text{L}\;\dfrac{x{=}y, P, Q, \langle x' = f(x)\,\&\,Q\rangle\bigcirc \vdash \langle x' = f(x)\,\&\,P\rangle\bigcirc}{Q \vdash (x{=}y \wedge P \rightarrow R)}}{P \vdash [x' = f(x)\,\&\,Q \wedge (P \vee x{=}y)](x{=}y \wedge P \rightarrow R)}$$

Continuing from premise ①, the derivation splits classically on whether $x{=}y$ is true initially, yielding two further premises labeled ③ when $x{=}y$ and ④ when $x \neq y$.

$$\text{cut}\;\dfrac{\vee\text{L}\;\dfrac{\textcircled{3} \qquad \textcircled{4}}{x{=}y \vee x \neq y, P \vdash [x' = f(x)\,\&\,Q \wedge (P \vee x{=}y)](x{=}y \rightarrow P)}}{P \vdash [x' = f(x)\,\&\,Q \wedge (P \vee x{=}y)](x{=}y \rightarrow P)}$$

From ③, the antecedents $x{=}y$ and P imply that $P(y)$ is true initially by a cut. Since y is held constant by the ODE $x' = f(x)$, a monotonicity step M['] followed by V completes the proof:

$$\text{cut, M[$'$]}\,\dfrac{\text{V}\,\dfrac{*}{P(y) \vdash [x' = f(x) \mathbin{\&} Q \land (P \lor x{=}y)]P(y)}}{x{=}y, P \vdash [x' = f(x) \mathbin{\&} Q \land (P \lor x{=}y)](x{=}y \to P)}$$

From ④, the derivation continues by dualizing to the diamond modality with $\langle\cdot\rangle$, ¬R.

$$\langle\cdot\rangle,\neg\text{R}\,\dfrac{x{\neq}y, P, \langle x' = f(x) \mathbin{\&} Q \land (P \lor x{=}y)\rangle(x{=}y \land \neg P) \vdash \mathit{false}}{x{\neq}y, P \vdash [x' = f(x) \mathbin{\&} Q \land (P \lor x{=}y)](x{=}y \to P)}$$

From the resulting premise, axiom rfl$\langle\cdot\rangle$ is used to syntactically reverse the diamond modality ODE in the antecedent from $x' = f(x)$ to $x' = -f(x)$. The dRW$\langle\cdot\rangle$ step weakens the resulting postcondition of the diamond modality with the propositional tautology $x{\neq}y \land P \to x{\neq}y$.

$$\text{rfl}\langle\cdot\rangle\,\dfrac{\text{dRW}\langle\cdot\rangle\,\dfrac{x{=}y, \neg P, \langle x' = -f(x) \mathbin{\&} Q \land (P \lor x{=}y)\rangle x{\neq}y \vdash \mathit{false}}{x{=}y, \neg P, \langle x' = -f(x) \mathbin{\&} Q \land (P \lor x{=}y)\rangle(x{\neq}y \land P) \vdash \mathit{false}}}{x{\neq}y, P, \langle x' = f(x) \mathbin{\&} Q \land (P \lor x{=}y)\rangle(x{=}y \land \neg P) \vdash \mathit{false}}$$

The diamond modality in the antecedent splits by axiom Uniq into two assumptions with domain constraints Q and $P \lor x{=}y$ respectively. With a use of derived axiom Init, all the antecedents of the bottom premise of rule rI& are gathered, leaving only its succedent.

$$\text{Uniq}\,\dfrac{\text{Init}\,\dfrac{x{=}y, \neg P, Q, \langle x' = -f(x) \mathbin{\&} Q\rangle\bigcirc, \langle x' = -f(x) \mathbin{\&} P\rangle\bigcirc \vdash \mathit{false}}{x{=}y, \neg P, \langle x' = -f(x) \mathbin{\&} Q\rangle x \neq y, \langle x' = -f(x) \mathbin{\&} P \lor x{=}y\rangle x \neq y \vdash \mathit{false}}}{x{=}y, \neg P, \langle x' = -f(x) \mathbin{\&} Q \land (P \lor x{=}y)\rangle x \neq y \vdash \mathit{false}}$$

Continuing with the derived implication $\neg\bigcirc_{\to}$ results in the bottom premise of rule rI&:

$$\neg\bigcirc_{\to}\,\dfrac{\neg\text{L}\,\dfrac{x{=}y, \neg P, Q, \langle x' = -f(x) \mathbin{\&} Q\rangle\bigcirc \vdash \langle x' = -f(x) \mathbin{\&} \neg P\rangle\bigcirc}{x{=}y, \neg P, Q, \langle x' = -f(x) \mathbin{\&} Q\rangle\bigcirc, \neg\langle x' = -f(x) \mathbin{\&} \neg P\rangle\bigcirc \vdash \mathit{false}}}{x{=}y, \neg P, Q, \langle x' = -f(x) \mathbin{\&} Q\rangle\bigcirc, \langle x' = -f(x) \mathbin{\&} P\rangle\bigcirc \vdash \mathit{false}} \qquad \square$$

The rule rI& discards any additional context in the antecedents of its premises. This is due to the use of RI& which focuses on particular states along trajectories of the ODE $x' = f(x)$. It would be unsound to keep any assumptions about the initial state that depend on x because the state being examined may not be the initial state! On the other hand, assumptions that do not depend on x remain true along the ODE. These constant assumptions can be kept with uses of V throughout the derivation above or added into Q before using rI& by a DC that proves with V. Following the notational conventions (Section 2.3.2), such additional steps are elided and rule rI& is used directly while keeping these *constant* assumptions around.

A.2 Completeness

This appendix gives completeness proofs for the derived rules dRI, sAI and the derived local progress characterization LP. Completeness of dRI is proved by showing that DRI is a derived axiom and similarly for sAI by showing that SAI is a derived axiom. This syntactic approach to proving completeness of dRI and sAI demonstrates the versatility of the dL calculus and it also enables complete *dis*proofs of invariance properties as opposed to just failing to apply a complete proof rule. To conclude that invariance is *disproved* after applying an algorithmic procedure (like the presentations of (semi)algebraic dRI and sAI [63, 103]), one would need to trust, in addition to soundness, that no completeness errors are present in the implementation.

Recall that axioms Cont, RI& have an additional syntactic requirement, e.g., the presence of $x_1' = 1$, which is assumed to be met throughout this appendix, using axiom DG if necessary.

A.2.1 Progress Formulas

Throughout this section, progress formulas are with respect to ODE $x' = f(x)$. The following are useful logical rearrangements of the progress formulas for extended term e.

Proposition A.7 (Atomic progress formula equivalences)**.** *Let N be the rank of extended term e. The following are provable equivalences on the progress and differential radical formulas for e:*

$$
\begin{aligned}
\dot{e}^{(*)} > 0 \leftrightarrow{}& e > 0 \lor (e = 0 \land \dot{e} > 0) \lor \dots \\
&\lor \big(e = 0 \land \dot{e} = 0 \land \dots \land \dot{e}^{(N-3)} = 0 \land \dot{e}^{(N-2)} > 0\big) \\
&\lor \big(e = 0 \land \dot{e} = 0 \land \dots \land \dot{e}^{(N-2)} = 0 \land \dot{e}^{(N-1)} > 0\big)
\end{aligned}
\tag{A.1}
$$

$$
\begin{aligned}
\dot{e}^{(*)} \geq 0 \leftrightarrow{}& e \geq 0 \land \big(e = 0 \rightarrow \dot{e} \geq 0\big) \land \dots \\
&\land \big(e = 0 \land \dot{e} = 0 \land \dots \land \dot{e}^{(N-3)} = 0 \rightarrow \dot{e}^{(N-2)} \geq 0\big) \\
&\land \big(e = 0 \land \dot{e} = 0 \land \dots \land \dot{e}^{(N-2)} = 0 \rightarrow \dot{e}^{(N-1)} \geq 0\big)
\end{aligned}
\tag{A.2}
$$

$$
\neg(\dot{e}^{(*)} > 0) \leftrightarrow (\dot{-e})^{(*)} \geq 0 \qquad \neg(\dot{e}^{(*)} \geq 0) \leftrightarrow (\dot{-e})^{(*)} > 0 \tag{A.3}
$$

$$
\neg(\dot{e}^{(*)} = 0) \leftrightarrow \dot{e}^{(*)} > 0 \lor (\dot{-e})^{(*)} > 0 \tag{A.4}
$$

Proof. The equivalences are proved one at a time. By linearity of Lie derivatives, $(\dot{-e})^{(i)} = -(\dot{e}^{(i)})$ proves in real arithmetic for any i. The proof also uses these provable real arithmetic equivalences:

$$
e \geq 0 \leftrightarrow e = 0 \lor e > 0 \qquad -e \geq 0 \land e \geq 0 \leftrightarrow e = 0 \qquad \neg(e > 0) \leftrightarrow -e \geq 0
$$

(A.1) This equivalence follows by real arithmetic, and simplifying with propositional rearrangement as follows (here, the remaining conjuncts of $\dot{e}^{(*)} > 0$ are abbreviated to $\dots$):

$$
\begin{aligned}
e \geq 0 \land \Big((e = 0 \rightarrow \dot{e} \geq 0) \land \dots\Big) \leftrightarrow{}& e > 0 \land \Big((e = 0 \rightarrow \dot{e} \geq 0) \land \dots\Big) \\
&\lor e = 0 \land \Big((e = 0 \rightarrow \dot{e} \geq 0) \land \dots\Big)
\end{aligned}
$$

The first disjunct on the RHS simplifies by real arithmetic to $e > 0$ since all of its implicational conjuncts contain $e = 0$ on the left of an implication. The latter disjunct simplifies to $e = 0 \land \Big(\dot{e} \geq 0 \land \dots\Big)$, yielding the provable equivalence:

$$
\dot{e}^{(*)} > 0 \leftrightarrow e > 0 \lor e = 0 \land \Big(\dot{e} \geq 0 \land \dots\Big)
$$

The equivalence (A.1) proves by iterating this expansion on the RHS of this equivalence for its nested conjuncts with higher Lie derivatives.

(A.2) This equivalence proves by expanding the formula $\dot{e}^{(*)} \geq 0$ which yields a disjunction between $\dot{e}^{(*)} > 0$ and $\dot{e}^{(*)} = 0$. The latter formula is used to relax the strict inequality in the last conjunct of $\dot{e}^{(*)} > 0$ to a non-strict inequality.

(A.3) The equivalence for $\neg(\dot{e}^{(*)} > 0)$ follows by negating both sides of equivalence (A.1) and moving negations on the RHS inwards propositionally, yielding the provable equivalence:

$$\neg(\dot{e}^{(*)} > 0) \leftrightarrow \Big(\neg(e > 0) \land (e = 0 \rightarrow \neg(\dot{e} > 0)) \land \ldots \\ \land \left(e{=}0 \land \dot{e}{=}0 \land \cdots \land \dot{e}^{(N-2)}{=}0 \rightarrow \neg(\dot{e}^{(N-1)} > 0)\right)\Big)$$

The desired equivalence derives by negating the inequalities and by equivalence (A.2) for $\dot{(-e)}^{(*)} \geq 0$. The equivalence for $\neg(\dot{e}^{(*)} \geq 0)$ derives by negating both sides of the equivalence for $\neg(\dot{e}^{(*)} > 0)$, since $-(-e) = e$.

(A.4) By (A.3), the following equivalence is provable:

$$\neg(\dot{e}^{(*)} > 0) \land \neg(\dot{(-e)}^{(*)} > 0) \leftrightarrow (\dot{(-e)}^{(*)} \geq 0) \land (\dot{e}^{(*)} \geq 0)$$

By rewriting with (A.2), the RHS of this equivalence is provably equivalent to the formula $\dot{e}^{(*)} = 0$ in real arithmetic. Negating both sides yields the provable equivalence (A.4). □

The provable equivalences (A.3) are particularly important, because they underlie the next proposition, from which the complete characterization of local progress follows:

Proposition A.8 (Negated semianalytic progress formula)**.** *Let semianalytic formula P be in normal form* (3.7)*, then $\neg P$ can be put into normal form such that $\neg(\dot{P}^{(*)}) \leftrightarrow \dot{(\neg P)}^{(*)}$ is provable:*

$$\neg P \equiv \bigvee_{i=0}^{N} \Big(\bigwedge_{j=0}^{a(i)} d_{ij} \geq 0 \land \bigwedge_{j=0}^{b(i)} \tilde{d}_{ij} > 0 \Big)$$

Proof. The proof uses the propositional tautologies $\neg(A \land B) \leftrightarrow \neg A \lor \neg B$ and $\neg(A \lor B) \leftrightarrow \neg A \land \neg B$. Formula P is negated (in normal form (3.7)) and all sub-terms are negated so the inequalities have 0 on the RHS, yielding the following provable equivalence. The resulting RHS is abbreviated by ϕ:

$$\neg P \leftrightarrow \underbrace{\bigwedge_{i=0}^{M} \Big(\bigvee_{j=0}^{m(i)} -e_{ij} > 0 \lor \bigvee_{j=0}^{n(i)} -\tilde{e}_{ij} \geq 0 \Big)}_{\phi}$$

The progress formula $\dot{P}^{(*)}$ for the normal form of P satisfies the equivalence (by definition):

$$\dot{P}^{(*)} \leftrightarrow \bigvee_{i=0}^{M} \Big(\bigwedge_{j=0}^{m(i)} \dot{e_{ij}}^{(*)} \geq 0 \land \bigwedge_{j=0}^{n(i)} \dot{\tilde{e}_{ij}}^{(*)} > 0 \Big)$$

Negating both sides of the progress formula for P and simplifying propositionally proves:

$$\neg(\dot{P}^{(*)}) \leftrightarrow \bigwedge_{i=0}^{M} \Big(\bigvee_{j=0}^{m(i)} \neg(\dot{e_{ij}}^{(*)} \geq 0) \vee \bigvee_{j=0}^{n(i)} \neg(\dot{\tilde{e}_{ij}}^{(*)} > 0) \Big)$$

Rewriting the RHS with equivalences (A.3) from Proposition A.7 yields the following provable equivalence. The resulting RHS is abbreviated by ψ:

$$\neg(\dot{P}^{(*)}) \leftrightarrow \underbrace{\bigwedge_{i=0}^{M} \Big(\bigvee_{j=0}^{m(i)} (\dot{-e_{ij}})^{(*)} > 0 \vee \bigvee_{j=0}^{n(i)} (\dot{-\tilde{e}_{ij}})^{(*)} \geq 0 \Big)}_{\psi}$$

Observe that ϕ, ψ have the same conjunctive normal form shape. Distribute the outer conjunction over the inner disjunctions in ϕ to obtain the following provable equivalence, whose RHS is a normal form for $\neg P$ (for some indices $N, a(i), b(i)$ and extended terms $d_{ij}, \tilde{d}_{ij}$):

$$\neg P \leftrightarrow \bigvee_{i=0}^{N} \Big(\bigwedge_{j=0}^{a(i)} d_{ij} \geq 0 \vee \bigwedge_{j=0}^{b(i)} \tilde{d}_{ij} > 0 \Big)$$

Distribute the disjunction in ψ following the same syntactic steps taken for ϕ to obtain the following provable equivalence:

$$\psi \leftrightarrow \bigvee_{i=0}^{N} \Big(\bigwedge_{j=0}^{a(i)} \dot{d_{ij}}^{(*)} \geq 0 \vee \bigwedge_{j=0}^{b(i)} \dot{\tilde{d}_{ij}}^{(*)} > 0 \Big)$$

Rewriting with the equivalences derived so far, and using the above normal form for $\neg P$, yields the required, provable equivalence:

$$\neg(\dot{P}^{(*)}) \leftrightarrow (\dot{\neg P})^{(*)}$$

□

A.2.2 Local Progress

This section derives the characterizations of local progress from Section 3.6.1. These characterizations are used in the completeness proofs for both analytic and semianalytic invariants.

Atomic Inequalities. The proof of Lemma 3.23 was outlined in Section 3.6.1. The case where $\succcurlyeq$ is $\geq$ is proved first, while the more technical case where $\succcurlyeq$ is $>$ is proved subsequently.

Proof of Lemma 3.23 (LP$_{\geq^*}$). Let N be the rank of extended term e with respect to $x' = f(x)$ from (3.2). For the derivation of LP$_{\geq^*}$, the additional flexibility of the $\bigcirc$ modality with a disjunct $x = y$ in the domain constraint is not needed. This disjunction is removed after unfolding the $\bigcirc$ abbreviation using a dRW$\langle\cdot\rangle$ monotonicity step. The definition of $\dot{e}^{(*)} \geq 0$ is also unfolded, with both disjuncts handled separately. The resulting premises are labeled ① (for the $\overset{\bullet}{e}^{(*)} > 0$ disjunct) and ② (for the $\overset{\bullet}{e}^{(*)} = 0$ disjunct).

$$\dfrac{{\scriptstyle \mathrm{dRW}\langle\cdot\rangle}\dfrac{{\scriptstyle \lor\mathrm{L}}\dfrac{\text{①}\qquad\text{②}}{x{=}y, \overset{\bullet}{e}^{(*)} > 0 \lor \overset{\bullet}{e}^{(*)} = 0 \vdash \langle x' = f(x) \,\&\, e \geq 0\rangle x \neq y}}{x{=}y, \overset{\bullet}{e}^{(*)} > 0 \lor \overset{\bullet}{e}^{(*)} = 0 \vdash \langle x' = f(x) \,\&\, e \geq 0 \lor x = y\rangle x \neq y}}{x{=}y, \overset{\bullet}{e}^{(*)} \geq 0 \vdash \langle x' = f(x) \,\&\, e \geq 0\rangle \bigcirc}$$

From ②, dRW$\langle\cdot\rangle$ strengthens the inequality $e \geq 0$ in the domain constraint to an equation $e = 0$. The derivation continues using DR$\langle\cdot\rangle$, because by dRI, $e = 0$ is provably invariant. The proof is completed with Cont using the trivial arithmetic fact $1 > 0$:

$${\scriptstyle \mathrm{dRW}\langle\cdot\rangle}\dfrac{{\scriptstyle \mathrm{DR}\langle\cdot\rangle}\dfrac{{\scriptstyle \mathrm{dRI}}\dfrac{*}{\overset{\bullet}{e}^{(*)} = 0 \vdash [x' = f(x) \,\&\, 1 > 0]e = 0} \qquad {\scriptstyle \mathrm{Cont}}\dfrac{*}{x{=}y \vdash \langle x' = f(x) \,\&\, 1 > 0\rangle x \neq y}}{x{=}y, \overset{\bullet}{e}^{(*)} = 0 \vdash \langle x' = f(x) \,\&\, e = 0\rangle x \neq y}}{x{=}y, \overset{\bullet}{e}^{(*)} = 0 \vdash \langle x' = f(x) \,\&\, e \geq 0\rangle x \neq y}$$

From ①, the premise is lined up for the derived step axiom LPi$_{\geq}$. The proof proceeds by closing the (left) premises obtained by iterating LPi$_{\geq}$ for higher Lie derivatives. In this way, the derivation continues until the final (rightmost) open premise which is abbreviated here with $\ldots$ and continued below:

$${\scriptstyle \mathrm{LPi}_{\geq}}\dfrac{{\scriptstyle \mathbb{R}}\dfrac{*}{\overset{\bullet}{e}^{(*)} > 0 \vdash e \geq 0} \qquad {\scriptstyle \mathrm{LPi}_{\geq}}\dfrac{{\scriptstyle \mathbb{R}}\dfrac{*}{\overset{\bullet}{e}^{(*)} > 0, e = 0 \vdash \overset{\bullet}{e} \geq 0} \qquad {\scriptstyle \mathrm{LPi}_{\geq}}\dfrac{x{=}y, \overset{\bullet}{e}^{(*)} > 0, \ldots \vdash \ldots}{\ldots}}{x{=}y, \overset{\bullet}{e}^{(*)} > 0, e = 0 \vdash \langle x' = f(x) \,\&\, \overset{\bullet}{e} \geq 0\rangle x \neq y}}{x{=}y, \overset{\bullet}{e}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, e \geq 0\rangle x \neq y}$$

The remaining open premise corresponds to the last conjunct of $\overset{\bullet}{e}^{(*)} > 0$. The implication in the conjunct uses the gathered antecedents $e = 0, \ldots, \overset{\bullet}{e}^{(N-2)} = 0$ after which Cont, dRW$\langle\cdot\rangle$ completes the proof:

$${\scriptstyle \mathrm{cut}}\dfrac{{\scriptstyle \mathrm{Cont,\ dRW}\langle\cdot\rangle}\dfrac{*}{x{=}y, \overset{\bullet}{e}^{(N-1)} > 0 \vdash \langle x' = f(x) \,\&\, \overset{\bullet}{e}^{(N-1)} \geq 0\rangle x \neq y}}{x{=}y, \overset{\bullet}{e}^{(*)} > 0, e = 0, \ldots, \overset{\bullet}{e}^{(N-2)} = 0 \vdash \langle x' = f(x) \,\&\, \overset{\bullet}{e}^{(N-1)} \geq 0\rangle x \neq y} \qquad \square$$

Unlike the non-strict case just derived for Lemma 3.23, the strict case (where $\succcurlyeq$ is $>$) crucially uses the fact that the $\bigcirc$ modality *excludes* the initial state, so that it is possible to locally progress into the strict inequality $e > 0$ without already satisfying it in the initial state. Topological considerations made this exclusion irrelevant for the non-strict case (see Section 3.6.1), as derived axiom Init explains logically. The idea behind the remaining proof of Lemma 3.23 for the strict case is to syntactically embed this difference into the derivation of LP$_{>^*}$. Moreover, this syntactic transformation reduces the proof to the non-strict case, so that the derived step axiom LPi$_{\geq}$ can again be used to progressively analyze higher Lie derivatives. The following proposition is used for the transformation:

Proposition A.9. *Let $d = e^k$ for some $k \geq 1$ and $x' = f(x)$ be an ODE with extended terms $e, d, f(x)$. For each $0 \leq i \leq k - 1$, there (computably) exists an extended term cofactor g such that the following identity is provable in real arithmetic:*

$$\overset{\bullet}{d}^{(i)} = ge$$

Proof. The proof proceeds by induction on k.

- For $k = 1$, $d = e^1$ so $\overset{\bullet}{e}^{(0)} = e$ hence the cofactor $g = 1$ suffices.
- For $d = e^{k+1}$, the j-th Lie derivative of d for $0 \leq j \leq k$ is given by Leibniz's rule:

  $$\overset{\bullet}{d}^{(j)} = \mathcal{L}^{(j)}_{f(x)}(e^k e) = \sum_{i=0}^{j} \binom{j}{i} \overset{\bullet}{(e^k)}^{(j-i)} \overset{\bullet}{e}^{(i)}$$

  The induction hypothesis implies $\overset{\bullet}{(e^k)}^{(j-i)} = g_i e$ is a provable identity for some computable extended term cofactor g_i for each $1 \leq i \leq j$. The final summand for $i = 0$ is:

  $$\binom{j}{0} \overset{\bullet}{(e^k)}^{(j)} \overset{\bullet}{e}^{(0)} = \overset{\bullet}{(e^k)}^{(j)} e$$

  Thus, the cofactor $g = \overset{\bullet}{(e^k)}^{(j)} + \sum_{i=1}^{j} \binom{j}{i} g_i \overset{\bullet}{e}^{(i)}$ yields the identity $\overset{\bullet}{d}^{(j)} = ge$. This identity is provable because it only depends on first-order properties of real arithmetic. □

For $d = e^k, k \geq 1$, Proposition A.9 shows that formula $e = 0 \rightarrow \bigwedge_{i=0}^{k-1} \overset{\bullet}{d}^{(i)} = 0$ is provable in real arithmetic for extended terms e, d, which enables the remaining proof of Lemma 3.23.

Proof of Lemma 3.23 (LP$_{>^*}$). Let $N \geq 1$ be the rank of extended term e with respect to $x' = f(x)$. This rank bounds the number of higher Lie derivatives of e that need to be considered.

The derivation starts by unfolding the syntactic abbreviation of the $\bigcirc$ modality and reducing to the non-strict case with dRW$\langle\cdot\rangle$ and the real arithmetic fact $e{-}d \geq 0 \rightarrow e > 0 \lor x = y$ for the abbreviation $d \overset{\text{def}}{=} |x - y|^{2N}$, which is a (polynomial) term: $\big((x_1 - y_1)^2 + \cdots + (x_n - y_n)^2\big)^N$.

$$\text{dRW}\langle\cdot\rangle\dfrac{\mathbb{R}\dfrac{*}{e{-}d \geq 0 \vdash e > 0 \lor x = y} \qquad x{=}y, \overset{\bullet}{e}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, e{-}d \geq 0\rangle x \neq y}{\dfrac{x{=}y, \overset{\bullet}{e}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, e > 0 \lor x = y\rangle x \neq y}{x{=}y, \overset{\bullet}{e}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, e > 0\rangle \bigcirc}}$$

Next, the initial assumption $x{=}y$ in the antecedent is used. The first cut proves using the formula of real arithmetic: $x{=}y \rightarrow |x - y|^2 = 0$. As remarked, with $|x - y|^2 = 0$ and $N \geq 1$, by Proposition A.9, $|x - y|^2 = 0 \rightarrow \bigwedge_{i=0}^{N-1} \overset{\bullet}{d}^{(i)} = 0$ is a provable real arithmetic formula. The second cut proves using this fact. The proof of the remaining open premise is continued below.

$$\text{cut, }\mathbb{R}\dfrac{\text{cut, }\mathbb{R}\dfrac{x{=}y, \bigwedge_{i=0}^{N-1} \overset{\bullet}{d}^{(i)} = 0, \overset{\bullet}{e}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, e{-}d \geq 0\rangle x \neq y}{x{=}y, |x - y|^2 = 0, \overset{\bullet}{e}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, e{-}d \geq 0\rangle x \neq y}}{x{=}y, \overset{\bullet}{e}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, e{-}d \geq 0\rangle x \neq y}$$

To continue, observe that for $0 \leq i \leq N - 1$, by linearity of the Lie derivative:

$$\mathcal{L}^{(i)}_{f(x)}(e{-}d) = \overset{\bullet}{e}^{(i)} - \overset{\bullet}{d}^{(i)}$$

Using the conjunction $\bigwedge_{i=0}^{N-1} \overset{\bullet}{d}^{(i)} = 0$ in the antecedents, the formula $\overset{\bullet}{(e{-}d)}^{(i)} = \overset{\bullet}{e}^{(i)}$ proves by a cut and real arithmetic for $0 \leq i \leq N - 1$. This justifies the next real arithmetic step from the open premise, with the assumptions $\Gamma_d \overset{\text{def}}{\equiv} \bigwedge_{i=0}^{N-1} \dot{(e{-}d)}^{(i)} = \dot{e}^{(i)}$. Intuitively, Γ_d allows the derivation to locally work with higher Lie derivatives of e instead of higher Lie derivatives of $e - d$ in subsequent steps.

$$\text{cut, }\mathbb{R}\ \dfrac{\Gamma_d, x{=}y, \dot{e}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, e{-}d \geq 0\rangle x \neq y}{x{=}y, \bigwedge_{i=0}^{N-1} \dot{d}^{(i)} = 0, \dot{e}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, e{-}d \geq 0\rangle x \neq y}$$

The derivation is completed using the same technique of iterating $\text{LPi}_{\geq}$, as shown in the earlier proof of Lemma 3.23 for the non-strict case $\text{LP}_{\geq^*}$. It starts with a single $\text{LPi}_{\geq}$ step. The left premise closes by real arithmetic because $\dot{e}^{(*)} > 0$ has the conjunct $e \geq 0$, and Γ_d provides $e - d = e$, which imply $e - d \geq 0$. The remaining open premise on the right is proved below.

$$\text{LPi}_{\geq}\ \dfrac{\mathbb{R}\dfrac{*}{\Gamma_d, \dot{e}^{(*)} > 0 \vdash e{-}d \geq 0} \qquad \Gamma_d, x{=}y, \dot{e}^{(*)} > 0, e{-}d = 0 \vdash \langle x' = f(x) \,\&\, \dot{(e{-}d)}^{(1)}{\geq}0\rangle x \neq y}{\Gamma_d, x{=}y, \dot{e}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, e{-}d \geq 0\rangle x \neq y}$$

Continuing from the open premise, local progress for the first Lie derivative of $e{-}d$ is proved. The first step simplifies formula $e{-}d{=}0$ in the antecedents using Γ_d. The derived axiom $\text{LPi}_{\geq}$, together with Γ_d, simplifies and proves the left premise. The right premise is abbreviated ① (shown and continued below).

$$\mathbb{R}\ \dfrac{\text{LPi}_{\geq}\ \dfrac{\mathbb{R}\ \dfrac{\mathbb{R}\ \dfrac{*}{e = 0 \to \dot{e} \geq 0, e = 0 \vdash \dot{e}^{(1)} \geq 0}}{\Gamma_d, \dot{e}^{(*)} > 0, e = 0 \vdash \dot{(e{-}d)}^{(1)}{\geq}0} \qquad \text{①}}{\Gamma_d, x{=}y, \dot{e}^{(*)} > 0, e = 0 \vdash \langle x' = f(x) \,\&\, \dot{(e{-}d)}^{(1)}{\geq}0\rangle x \neq y}}{\Gamma_d, x{=}y, \dot{e}^{(*)} > 0, e{-}d = 0 \vdash \langle x' = f(x) \,\&\, \dot{(e{-}d)}^{(1)}{\geq}0\rangle x \neq y}$$

The derivation continues from ① similarly for higher Lie derivatives of $e{-}d$, using Γ_d to replace $\dot{(e{-}d)}^{(i)}$ with $\dot{e}^{(i)}$, and then using the corresponding conjunct of $\dot{e}^{(*)} > 0$. The final open premise obtained from ① by iterating $\text{LPi}_{\geq}$ corresponds to the last conjunct of $\dot{e}^{(*)} > 0$:

$$\mathbb{R}\ \dfrac{\text{LPi}_{\geq}\ \dfrac{\text{LPi}_{\geq}\ \dfrac{\Gamma_d, x{=}y, \dot{e}^{(*)} > 0, e{=}0, \dots, \dot{e}^{(N-2)}{=}0 \vdash \langle x' = f(x) \,\&\, \dot{(e{-}d)}^{(N-1)}{\geq}0\rangle x \neq y}{\cdots}}{\Gamma_d, x{=}y, \dot{e}^{(*)} > 0, e = 0, \dot{e}^{(1)} \geq 0 \vdash \langle x' = f(x) \,\&\, \dot{(e{-}d)}^{(2)}{\geq}0\rangle x \neq y}}{\Gamma_d, x{=}y, \dot{e}^{(*)} > 0, e{=}0, \dot{(e{-}d)}^{(1)}{\geq}0 \vdash \langle x' = f(x) \,\&\, \dot{(e{-}d)}^{(2)}{\geq}0\rangle x \neq y}$$

The gathered antecedents $e = 0, \dots, \dot{e}^{(N-2)} = 0$ are respectively obtained from Γ_d by real arithmetic. The proof is closed with $\text{dRW}\langle\cdot\rangle$, Cont, similarly to the non-strict case.

$$\text{cut}\ \dfrac{\text{dRW}\langle\cdot\rangle\ \dfrac{\text{cut, }\mathbb{R}\ \dfrac{\text{Cont}\ \dfrac{*}{x{=}y, \dot{(e{-}d)}^{(N-1)}{>}0 \vdash \langle x' = f(x) \,\&\, \dot{(e{-}d)}^{(N-1)}{>}0\rangle x \neq y}}{\Gamma_d, x{=}y, \dot{e}^{(N-1)}{>}0 \vdash \langle x' = f(x) \,\&\, \dot{(e{-}d)}^{(N-1)}{>}0\rangle x \neq y}}{\Gamma_d, x{=}y, \dot{e}^{(N-1)}{>}0 \vdash \langle x' = f(x) \,\&\, \dot{(e{-}d)}^{(N-1)}{\geq}0\rangle x \neq y}}{\Gamma_d, x{=}y, \dot{e}^{(*)} > 0, e{=}0, .., \dot{e}^{(N-2)}{=}0 \vdash \langle x' = f(x) \,\&\, \dot{(e{-}d)}^{(N-1)}{\geq}0\rangle x \neq y} \qquad \square$$

Semianalytic Formulas. The proof in the semianalytic case is outlined in Section 3.6.1. It lifts derived axioms LP$_{\geq^*}$ and LP$_{>^*}$ according to the homomorphic definition of the semianalytic progress formula, using axiom Uniq to prove local progress into a conjunction of two formulas simultaneously.

Proof of Lemma 3.25. By congruential equivalence [142], assume, without loss of generality, that formula P is propositionally rewritten to the same normal form (3.7) as in the corresponding semianalytic progress formula $\dot{P}^{(*)}$. Throughout this proof, similar premises are collapsed in proofs and directly indexed by i, j. The i-th disjunct of P is abbreviated with

$$P_i \overset{\text{def}}{\equiv} \bigwedge_{j=0}^{m(i)} e_{ij} \geq 0 \land \bigwedge_{j=0}^{n(i)} \tilde{e}_{ij} > 0$$

The derivation starts by splitting the (outermost) disjunction in $\dot{P}^{(*)}$ with $\vee$L. For each resulting premise (indexed by i), local progress is proved for the corresponding disjunct P_i of P. The domain change with dRW$\langle\cdot\rangle$ proves because $P_i \vee x = y \to P \vee x = y$ is a propositional tautology for each i.

$$\vee\text{L}\;\dfrac{\text{dRW}\langle\cdot\rangle\;\dfrac{x{=}y, \bigwedge_{j=0}^{m(i)} \dot{e_{ij}}^{(*)} \geq 0 \land \bigwedge_{j=0}^{n(i)} \dot{\tilde{e}_{ij}}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, P_i\rangle \bigcirc}{x{=}y, \bigwedge_{j=0}^{m(i)} \dot{e_{ij}}^{(*)} \geq 0 \land \bigwedge_{j=0}^{n(i)} \dot{\tilde{e}_{ij}}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, P\rangle \bigcirc}}{x{=}y, \dot{P}^{(*)} \vdash \langle x' = f(x) \,\&\, P\rangle \bigcirc}$$

It suffices now to prove local progress in P_i. The uniqueness axiom Uniq splits conjuncts in P_i then the dRW$\langle\cdot\rangle$ step distributes $x{=}y$ in domain constraint from $\bigcirc$ over conjunctions using the propositional tautology $(R_1 \land R_2) \vee x{=}y \leftrightarrow (R_1 \vee x{=}y) \land (R_2 \vee x = y)$. This leaves premises (indexed by j) for the non-strict and strict inequalities of P_i which are closed by LP$_{\geq^*}$ and LP$_{>^*}$ respectively (labeled ① and ② respectively and shown immediately below).

$$\text{Uniq, }\land\text{R, dRW}\langle\cdot\rangle\;\dfrac{①\quad ②}{x{=}y, \bigwedge_{j=0}^{m(i)} \dot{e_{ij}}^{(*)} \geq 0 \land \bigwedge_{j=0}^{n(i)} \dot{\tilde{e}_{ij}}^{(*)} > 0 \vdash \langle x'{=}f(x) \,\&\, P_i\rangle \bigcirc}$$

From ①:

$$\text{LP}_{\geq^*}\;\dfrac{*}{x{=}y, \dot{e_{ij}}^{(*)} \geq 0 \vdash \langle x'{=}f(x) \,\&\, e_{ij} \geq 0\rangle \bigcirc}$$

From ②:

$$\text{LP}_{>^*}\;\dfrac{*}{x{=}y, \dot{\tilde{e}_{ij}}^{(*)} > 0 \vdash \langle x'{=}f(x) \,\&\, \tilde{e}_{ij} > 0\rangle \bigcirc}$$

□

The implicational semianalytic local progress axiom LP$\mathbb{R}$ from Lemma 3.25 is strengthened to an equivalent characterization of semianalytic local progress using Proposition A.8.

Proof of Theorem 3.26. By congruential equivalence [142], assume, without loss of generality, that formula P is propositionally rewritten to the same normal form (3.7) as in the corresponding semianalytic progress formula $\dot{P}^{(*)}$. By Proposition A.8, there is a normal form for $\neg P$ with the provable equivalence $\neg(\dot{P}^{(*)}) \leftrightarrow (\dot{\neg P})^{(*)}$. The "←" direction of the inner equivalence is

LPℝ. The derivation in the "→" direction of the inner equivalence starts by reducing to the contrapositive statement by propositional logic transformations. The final step rewrites the negation in the antecedents using the above normal form for $\neg P$ from Proposition A.8.

$$\text{cut, }\neg\text{L, }\neg\text{R}\;\dfrac{\mathbb{R}\;\dfrac{x{=}y, (\dot{\neg P})^{(*)} \vdash \neg\langle x' = f(x)\,\&\, P\rangle\bigcirc}{x{=}y, \neg(\dot{P}^{(*)}) \vdash \neg\langle x' = f(x)\,\&\, P\rangle\bigcirc}}{x{=}y, \langle x' = f(x)\,\&\, P\rangle\bigcirc \vdash \dot{P}^{(*)}}$$

By the derived axiom LPℝ from Lemma 3.25, the progress formula for $\neg P$ in the antecedent implies local progress for $\neg P$. The proof is completed with derived axiom $\neg\bigcirc_{\rightarrow}$ of Corollary A.4:

$$\text{LP}\mathbb{R}\;\dfrac{\neg\bigcirc_{\rightarrow}\;\dfrac{*}{x{=}y, \langle x' = f(x)\,\&\, \neg P\rangle\bigcirc \vdash \neg\langle x' = f(x)\,\&\, P\rangle\bigcirc}}{x{=}y, (\dot{\neg P})^{(*)} \vdash \neg\langle x' = f(x)\,\&\, P\rangle\bigcirc}$$

□

Proof of Corollary 3.27. Self-duality axiom $\neg\bigcirc$ derives by using LP twice together with the provable equivalence $\neg(\dot{P}^{(*)}) \leftrightarrow (\dot{\neg P})^{(*)}$ from Proposition A.8 (and double negation elimination).

$$\text{LP}\;\dfrac{\mathbb{R}\;\dfrac{\text{LP}\;\dfrac{*}{x{=}y \vdash \langle x' = f(x)\,\&\, P\rangle\bigcirc \leftrightarrow \dot{P}^{(*)}}}{x{=}y \vdash \langle x' = f(x)\,\&\, P\rangle\bigcirc \leftrightarrow \neg(\dot{\neg P})^{(*)}}}{x{=}y \vdash \langle x' = f(x)\,\&\, P\rangle\bigcirc \leftrightarrow \neg\langle x' = f(x)\,\&\, \neg P\rangle\bigcirc}$$

Local progress congruence rule CLP derives similarly by introducing an initial assumption $x{=}y$ with cut, ℝ, ∃L, equivalently rewriting with LP, and congruential equivalence [142] in the last step.

$$\text{cut, }\mathbb{R}\text{, }\exists\text{L}\;\dfrac{\text{LP}\;\dfrac{\dfrac{\vdash P \leftrightarrow R}{\vdash \langle x' = f(x)\,\&\, P\rangle\bigcirc \leftrightarrow \langle x' = f(x)\,\&\, R\rangle\bigcirc}}{x{=}y \vdash \dot{P}^{(*)} \leftrightarrow \dot{R}^{(*)}}}{\vdash \dot{P}^{(*)} \leftrightarrow \dot{R}^{(*)}}$$

□

A.2.3 Analytic Invariants

This section derives the analytic completeness axiom DRI (and its generalization DRI&), thus proving completeness for analytic (Noetherian) invariants and also for analytic postconditions.

Differential Radical Invariants. The differential radical invariants proof rule dRI derives from rule vdbx by equivalently turning the differential radical identity (3.2) into a provable vectorial Darboux equality.

Proof of Theorem 3.11 . Let e be an extended term satisfying both premises of the dRI proof rule and let $\mathbf{e}$ be the vector of extended terms with components $\mathbf{e}_i \overset{\text{def}}{=} \dot{e}^{(i-1)}$ for $i = 1, 2, \ldots, N$. The derivation starts by setting up the premise for an application of derived rule vdbx. In the first step, axiom DX is used to assume that the domain constraint Q is true initially. On the left premise after the cut, arithmetic equivalence $\bigwedge_{i=0}^{N-1} \dot{e}^{(i)} = 0 \leftrightarrow \mathbf{e} = 0$ is used to rewrite the succedent to the left premise of dRI. On the right premise, monotonicity M['] strengthens the postcondition to $\mathbf{e} = 0$:

$$\text{DX}\,\dfrac{\text{cut}\,\dfrac{\mathbb{R}\,\dfrac{\Gamma, Q \vdash \bigwedge_{i=0}^{N-1} \overset{\bullet}{e}^{(i)} = 0}{\Gamma, Q \vdash \mathbf{e} = 0} \qquad \mathrm{M}[']\,\dfrac{\mathbf{e} = 0 \vdash [x' = f(x) \,\&\, Q]\mathbf{e} = 0}{\mathbf{e} = 0 \vdash [x' = f(x) \,\&\, Q]e = 0}}{\Gamma, Q \vdash [x' = f(x) \,\&\, Q]e = 0}}{\Gamma \vdash [x' = f(x) \,\&\, Q]e = 0}$$

Continuing from the right premise, the component-wise Lie derivative of $\mathbf{e}$ is defined as $(\overset{\bullet}{\mathbf{e}})_i = \mathcal{L}_{f(x)}(\mathbf{e}_i) = \overset{\bullet}{e}^{(i)}$. The vector $\overset{\bullet}{\mathbf{e}}$ will be obtained from $\mathbf{e}$ by matrix multiplication with the following $N \times N$ extended term cofactor matrix G with 1 on its superdiagonal, and the g_i cofactors in the last row:

$$G = \begin{pmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ 0 & 0 & \dots & 0 & 1 \\ g_0 & g_1 & \dots & g_{N-2} & g_{N-1} \end{pmatrix}, \quad \mathbf{e} = \begin{pmatrix} e \\ \overset{\bullet}{e}^{(1)} \\ \vdots \\ \overset{\bullet}{e}^{(N-2)} \\ \overset{\bullet}{e}^{(N-1)} \end{pmatrix}, \quad \overset{\bullet}{\mathbf{e}} = \begin{pmatrix} \overset{\bullet}{e}^{(1)} \\ \overset{\bullet}{e}^{(2)} \\ \vdots \\ \overset{\bullet}{e}^{(N-1)} \\ \overset{\bullet}{e}^{(N)} \end{pmatrix}$$

The vectorial equation $\overset{\bullet}{\mathbf{e}} = G\mathbf{e}$ is provably equivalent to the equation $\overset{\bullet}{e}^{(N)} = \sum_{i=0}^{N-1} g_i \overset{\bullet}{e}^{(i)}$. To see this, note that for indices $1 \leq i < N$, matrix multiplication yields:

$$(\overset{\bullet}{\mathbf{e}})_i = \overset{\bullet}{e}^{(i)} = (\mathbf{e})_{i+1} = (G\mathbf{e})_i$$

Therefore, all but the final component-wise equality are trivially valid and prove by $\mathbb{R}$. The remaining (non-trivial) equation for $i = N$ is $(\overset{\bullet}{\mathbf{e}})_N = (G\mathbf{e})_N$. The LHS of this equation simplifies with $(\overset{\bullet}{\mathbf{e}})_N = \overset{\bullet}{e}^{(N)}$, while the RHS simplifies by matrix multiplication to:

$$(G\mathbf{e})_N = \sum_{i=1}^{N} g_{i-1}(\mathbf{e})_i = \sum_{i=1}^{N} g_{i-1}\overset{\bullet}{e}^{(i-1)} = \sum_{i=0}^{N-1} g_i \overset{\bullet}{e}^{(i)}$$

Hence, real arithmetic equivalently turns the formula $\overset{\bullet}{\mathbf{e}} = G\mathbf{e}$ into the succedent of the right premise of rule dRI. An application of derived rule vdbx from Corollary 3.10 with cofactor matrix G followed by rule $\mathbb{R}$ yields the remaining right premise of rule dRI, completing the derivation.

$$\text{vdbx}\,\dfrac{\mathbb{R}\,\dfrac{Q \vdash \overset{\bullet}{e}^{(N)} = \sum_{i=0}^{N-1} g_i \overset{\bullet}{e}^{(i)}}{Q \vdash \overset{\bullet}{\mathbf{e}} = G\mathbf{e}}}{\mathbf{e} = 0 \vdash [x' = f(x) \,\&\, Q]\mathbf{e} = 0}$$

□

The derivation of rule dRI uses a specific choice of cofactor matrix G in rule vdbx to prove invariance of the equation $e = 0$. This suffices for analytic completeness because analytic formulas can always be normalized to a single equation in real arithmetic. However, such normalization may not yield the computationally most efficient way of proving an analytic invariant (see Example 3.41).

Completeness for Analytic Invariants. The analytic completeness axiom with semianalytic domain constraints DRI& from Theorem 3.30 is derived next, making use of LP$\mathbb{R}$ from Lemma 3.25. The completeness argument can be summarized by taking contrapositives: if the local progress formula $\dot{e}^{(*)} = 0$ is false in an initial state, then some higher Lie derivative of e is non-zero and gives a definite (local) sign to the value of e, which, by LP$\mathbb{R}$ implies progress to $e \neq 0$. For completeness, axiom DRI& also handles the vacuous case where domain constraint Q is false initially ($Q \rightarrow \ldots$ in DRI&) and the stuck case where the domain constraint is true initially but cannot locally progress ($\dot{Q}^{(*)} \rightarrow \ldots$ in DRI&).

Proof of Theorem 3.30 (implies Theorem 3.13). For formulas Q formed from conjunctions and disjunctions of strict inequalities, $Q \rightarrow \dot{Q}^{(*)}$ is provable in real arithmetic, so axiom DRI follows as an arithmetical corollary of DRI&. The derivation of axiom DRI& starts by rewriting its LHS equivalently with axiom DX. This is followed by equivalent propositional rewrites that simplify the logical structure of the succedent. The propositional steps are shown below, first pulling out the common implication Q then the common conjunct $e = 0$ as antecedent assumptions.

$$\begin{array}{lr}
 & Q, e = 0 \vdash \big([x' = f(x) \,\&\, Q]e = 0\big) \leftrightarrow \big(\dot{Q}^{(*)} \rightarrow \dot{e}^{(*)} = 0\big) \\ \cline{2-2}
 & Q \vdash \big(e = 0 \land [x' = f(x) \,\&\, Q]e = 0\big) \leftrightarrow \big(e = 0 \land (\dot{Q}^{(*)} \rightarrow \dot{e}^{(*)} = 0)\big) \\ \cline{2-2}
 & \vdash \big(Q \rightarrow e = 0 \land [x' = f(x) \,\&\, Q]e = 0\big) \leftrightarrow \big(Q \rightarrow e = 0 \land (\dot{Q}^{(*)} \rightarrow \dot{e}^{(*)} = 0)\big) \\ \cline{2-2}
\text{DX} & \vdash [x' = f(x) \,\&\, Q]e = 0 \leftrightarrow \big(Q \rightarrow e = 0 \land (\dot{Q}^{(*)} \rightarrow \dot{e}^{(*)} = 0)\big)
\end{array}$$

Next, a cut of the first-order formula $\exists y \, x{=}y$ proves trivially in real arithmetic and Skolemizing it with $\exists$L yields an initial state assumption ($x{=}y$ for fresh variables y). To make use of this initial state assumption, the derivation continues with a classical case split on whether the semianalytic progress formula $\dot{Q}^{(*)}$ is true initially. The resulting premises are labeled ① (for the $\dot{Q}^{(*)}$ disjunct) and ② (for the $\neg(\dot{Q}^{(*)})$ disjunct) and continued below.

$$\begin{array}{lr}
 & ① \qquad ② \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \\ \cline{2-2}
\lor\text{L} & x{=}y, Q, e = 0, \dot{Q}^{(*)} \lor \neg(\dot{Q}^{(*)}) \vdash \big([x' = f(x) \,\&\, Q]e = 0\big) \leftrightarrow \big(\dot{Q}^{(*)} \rightarrow \dot{e}^{(*)} = 0\big) \\ \cline{2-2}
\text{cut} & x{=}y, Q, e = 0 \vdash \big([x' = f(x) \,\&\, Q]e = 0\big) \leftrightarrow \big(\dot{Q}^{(*)} \rightarrow \dot{e}^{(*)} = 0\big) \\ \cline{2-2}
\exists\text{L} & \exists y \, x{=}y, Q, e = 0 \vdash \big([x' = f(x) \,\&\, Q]e = 0\big) \leftrightarrow \big(\dot{Q}^{(*)} \rightarrow \dot{e}^{(*)} = 0\big) \\ \cline{2-2}
\text{cut}, \mathbb{R} & Q, e = 0 \vdash \big([x' = f(x) \,\&\, Q]e = 0\big) \leftrightarrow \big(\dot{Q}^{(*)} \rightarrow \dot{e}^{(*)} = 0\big)
\end{array}$$

The premise ② corresponds to the case where solutions are *stuck* in the initial state because no local progress in the domain constraint Q is possible. Topologically, this corresponds to the situation where initial states are on the boundary of the set characterized by Q (and also in Q)[1] but the ODE locally leaves Q. Since $e = 0$ is already true in this stuck state, it trivially remains true for all solutions staying in domain constraint Q. The derivation from ② starts with a propositional simplification of the succedent since its RHS is vacuously equivalent to $true$ by assumption $\neg(\dot{Q}^{(*)})$. The local progress characterization axiom LP equivalently rewrites the sub-formula $\dot{Q}^{(*)}$ to local progress for Q before axiom $\langle\cdot\rangle$ unfolds the $\bigcirc$ modality turning it into

[1] This situation is impossible for domain constraints Q characterizing topologically open sets, which is the semantical reason for derived axiom DRI having a simpler RHS characterization than DRI&.

a box modality formula.

$$\mathbb{R}\,\cfrac{\text{LP}\,\cfrac{\langle\cdot\rangle\,\cfrac{x{=}y, e=0, [x'=f(x)\,\&\,Q \lor x=y]x=y \vdash [x'=f(x)\,\&\,Q]e=0}{x{=}y, e=0, \neg(\langle x'=f(x)\,\&\,Q\rangle\bigcirc) \vdash [x'=f(x)\,\&\,Q]e=0}}{x{=}y, e=0, \neg(\dot{Q}^{(*)}) \vdash [x'=f(x)\,\&\,Q]e=0}}{x{=}y, e=0, \neg(\dot{Q}^{(*)}) \vdash \big([x'=f(x)\,\&\,Q]e=0\big) \leftrightarrow \big(\dot{Q}^{(*)} \rightarrow \dot{e}^{(*)}=0\big)}$$

By axiom V, the constant assumption $e(y) = 0$ (with y in place of x) strengthens the postcondition of the antecedent box modality to $e = 0$ using the provable arithmetic formula $e(y) = 0 \land x = y \rightarrow e = 0$. A subsequent DMP, dW step finishes the proof using the propositional tautology $Q \rightarrow Q \lor x = y$.

$$\text{V}\,\cfrac{\text{DMP, dW}\,\cfrac{\cfrac{*}{Q \vdash Q \lor x=y}}{[x'=f(x)\,\&\,Q \lor x=y]e=0 \vdash [x'=f(x)\,\&\,Q]e=0}}{x{=}y, e=0, [x'=f(x)\,\&\,Q \lor x=y]x=y \vdash [x'=f(x)\,\&\,Q]e=0}$$

From premise ①, the succedent propositionally simplifies to $\big([x'=f(x)\,\&\,Q]e=0\big) \leftrightarrow \dot{e}^{(*)}=0$ by assumption $\dot{Q}^{(*)}$. The two directions of this simplified succedent are proved separately. In the "←" direction, the derivation uses rule dRI by setting N to the rank of extended term e, so that the succedent of its left premise is exactly $\dot{e}^{(*)} = 0$. The right premise resulting from dRI closes by real arithmetic, since N is the rank of e, it must, by definition satisfy the provable rank identity (3.2).

$$\text{dRI}\,\cfrac{\mathbb{R}\,\cfrac{*}{\vdash \dot{e}^{(N)} = \sum_{i=0}^{N-1} g_i \dot{e}^{(i)}}}{\dot{e}^{(*)}=0 \vdash [x'=f(x)\,\&\,Q]e=0}$$

The derivation in the "→" direction starts by reducing to the contrapositive statement with duality $\langle\cdot\rangle$ and propositional logical manipulation.

$$\langle\cdot\rangle\,\cfrac{\neg\text{L}, \neg\text{R}\,\cfrac{x{=}y, Q, \dot{Q}^{(*)}, \neg(\dot{e}^{(*)}=0) \vdash \langle x'=f(x)\,\&\,Q\rangle e\neq 0}{x{=}y, Q, \dot{Q}^{(*)}, \neg\langle x'=f(x)\,\&\,Q\rangle e\neq 0 \vdash \dot{e}^{(*)}=0}}{x{=}y, Q, \dot{Q}^{(*)}, [x'=f(x)\,\&\,Q]e=0 \vdash \dot{e}^{(*)}=0}$$

By derived axiom LP, the antecedent assumption $\dot{Q}^{(*)}$ is equivalently rewritten to local progress for Q before Init (Corollary A.3) is used to strengthen it with the assumption Q. The final step rewrites the resulting negated differential radical formula in the antecedents to two progress formulas by (A.4) from Proposition A.7. Subsequent splitting with ∨L yields two premises, which are abbreviated ③ (for disjunct $\dot{e}^{(*)} > 0$) and ④ (for disjunct $(\dot{-e})^{(*)} > 0$) respectively, and continued below.

$$\text{LP}\,\cfrac{\text{Init}\,\cfrac{\mathbb{R}\,\cfrac{\lor\text{L}\,\cfrac{\text{③} \qquad \text{④}}{x{=}y, \langle x'=f(x)\,\&\,Q\rangle x\neq y, \dot{e}^{(*)}>0 \lor (\dot{-e})^{(*)}>0 \vdash \langle x'=f(x)\,\&\,Q\rangle e\neq 0}}{x{=}y, \langle x'=f(x)\,\&\,Q\rangle x\neq y, \neg(\dot{e}^{(*)}=0) \vdash \langle x'=f(x)\,\&\,Q\rangle e\neq 0}}{x{=}y, Q, \langle x'=f(x)\,\&\,Q\rangle\bigcirc, \neg(\dot{e}^{(*)}=0) \vdash \langle x'=f(x)\,\&\,Q\rangle e\neq 0}}{x{=}y, Q, \dot{Q}^{(*)}, \neg(\dot{e}^{(*)}=0) \vdash \langle x'=f(x)\,\&\,Q\rangle e\neq 0}$$

Continuing from ③, the assumption $\dot{e}^{(*)} > 0$ is rewritten with $\text{LP}_{>^*}$ to obtain local progress for $e > 0$. Unfolding the $\bigcirc$ abbreviation, the uniqueness axiom Uniq combines the two diamond modality formulas in the antecedent:

$$\text{LP}_{>^*}\,\dfrac{\dfrac{\text{Uniq}\,\dfrac{\langle x' = f(x) \,\&\, Q \land (e > 0 \lor x = y)\rangle x \neq y \vdash \langle x' = f(x) \,\&\, Q\rangle e \neq 0}{\langle x' = f(x) \,\&\, Q\rangle x \neq y, \langle x' = f(x) \,\&\, e > 0 \lor x = y\rangle x \neq y \vdash \langle x' = f(x) \,\&\, Q\rangle e \neq 0}}{\langle x' = f(x) \,\&\, Q\rangle x \neq y, \langle x' = f(x) \,\&\, e > 0\rangle \bigcirc \vdash \langle x' = f(x) \,\&\, Q\rangle e \neq 0}}{x{=}y, \langle x' = f(x) \,\&\, Q\rangle x \neq y, \dot{e}^{(*)} > 0 \vdash \langle x' = f(x) \,\&\, Q\rangle e \neq 0}$$

The succedent's domain constraint is strengthened to match the antecedent's using rule dRW$\langle\cdot\rangle$ since $Q \land (e > 0 \lor x = y) \to Q$ is a propositional tautology. The Kripke axiom K$\langle\cdot\rangle$ reduces the succedent to the box modality, after which the proof finishes with a dW step because the formula $e > 0 \lor x = y$ in the domain constraint implies the succedent by real arithmetic.

$$\text{dRW}\langle\cdot\rangle\,\dfrac{\text{K}\langle\cdot\rangle\,\dfrac{\text{dW}\,\dfrac{\mathbb{R}\,\dfrac{*}{Q \land (e > 0 \lor x = y) \vdash (x{\neq}y{\to}e{\neq}0)}}{\vdash [x'{=}f(x) \,\&\, Q{\land}(e{>}0 \lor x{=}y)](x{\neq}y{\to}e{\neq}0)}}{\langle x' = f(x) \,\&\, Q \land (e > 0 \lor x = y)\rangle x \neq y \vdash \langle x' = f(x) \,\&\, Q \land (e > 0 \lor x = y)\rangle e \neq 0}}{\langle x' = f(x) \,\&\, Q \land (e > 0 \lor x = y)\rangle x \neq y \vdash \langle x' = f(x) \,\&\, Q\rangle e \neq 0}$$

The remaining premise ④ follows similarly, except that the progress formula $(\dot{-e})^{(*)} > 0$ enables the cut $\langle x' = f(x) \,\&\, -e > 0\rangle \bigcirc$ which leads to the same postcondition $e \neq 0$ in the succedent instead. □

A.2.4 Completeness for Semianalytic Invariants with Semianalytic Evolution Domains

The following generalized version of rule sAI from Theorem 3.28 additionally handles evolution domain constraints. It derives from derived rule rI& and derived axiom LP.

Theorem A.10 (Semianalytic invariants with semianalytic domain constraints)**.** *The semianalytic invariant proof rule with semianalytic domain constraints sAI& derives from RI&, Dadj, Cont, Uniq for semianalytic formulas P, Q.*

$$\text{sAI\&}\;\; \dfrac{P, Q, \dot{Q}^{(*)} \vdash \dot{P}^{(*)} \qquad \neg P, Q, \dot{Q}^{-(*)} \vdash (\dot{\neg P})^{-(*)}}{P \vdash [x' = f(x) \,\&\, Q]P}$$

Proof (implies Theorem 3.28). Rule sAI& derives from rule rI& derived in Corollary A.6 and the characterization of semianalytic local progress LP derived in Theorem 3.26. The $x{=}y$ assumptions provided by rI& are used to convert between local progress modalities and the semianalytic progress formulas by LP, but, by weakening, $x{=}y$ can be elided again in the premises of sAI&. □

Recalling the earlier discussion for derived rule rI&, axiom V can be used, as usual, to keep *constant* context assumptions that do not depend on variables x for the ODEs $x' = f(x)$ in rule sAI&, because it immediately derives from rI&, which supports constant contexts. The proof rule sAI& is complete for invariance properties. This is proved syntactically, enabling complete disproofs of invariance. The completeness of sAI from Theorem 3.29 follows as a special case, where $Q \equiv true$.

Theorem A.11 (Semianalytic invariant completeness with semianalytic domains)**.** *The semianalytic invariant axiom with semianalytic domain constraints SAI& derives from RI&, Dadj, Cont, Uniq for semianalytic formulas P, Q.*

$$\text{SAI\&}\;\; \forall x\,(P \rightarrow [x' = f(x)\,\&\,Q]P) \leftrightarrow \begin{array}{l} \overbrace{\forall x\,\big(P \land Q \land \dot{Q}^{(*)} \rightarrow \dot{P}^{(*)}\big)}^{\textcircled{a}} \land \\ \underbrace{\forall x\,\big(\neg P \land Q \land \dot{Q}^{-(*)} \rightarrow (\dot{\neg P})^{-(*)}\big)}_{\textcircled{b}} \end{array}$$

Proof (implies Theorem 3.29). The left and right conjunct on the RHS of SAI& are abbreviated ⓐ and ⓑ respectively. The "←" direction derives by sAI&. The antecedents ⓐ and ⓑ are first-order formulas quantified over x, the variables evolved by the ODE $x' = f(x)$. They are kept as constant context in the antecedents of the premises when applying rule sAI& and later instantiated by ∀L.

$$\text{∀R, →R}\;\dfrac{\text{sAI\&}\;\dfrac{\text{∀L, →L}\dfrac{*}{\textcircled{a}, P, Q, \dot{Q}^{(*)} \vdash \dot{P}^{(*)}} \qquad \text{∀L, →L}\dfrac{*}{\textcircled{b}, \neg P, Q, \dot{Q}^{-(*)} \vdash (\dot{\neg P})^{-(*)}}}{\textcircled{a}, \textcircled{b}, P \vdash [x' = f(x)\,\&\,Q]P}}{\textcircled{a}, \textcircled{b} \vdash \forall x\,(P \rightarrow [x' = f(x)\,\&\,Q]P)}$$

In the "→" direction, the derivation proceeds by contraposition in both cases after ∧R. For ⓑ, the derived invariant reflection axiom (rfl) is used to syntactically turn the invariance assumption for the forward ODE into an invariance assumption for the backward ODE.

$$\text{∧R}\;\dfrac{\forall x\,(P \rightarrow [x' = f(x)\,\&\,Q]P) \vdash \textcircled{a} \qquad \text{rfl}\;\dfrac{\forall x\,(\neg P \rightarrow [x' = -f(x)\,\&\,Q]\neg P) \vdash \textcircled{b}}{\forall x\,(P \rightarrow [x' = f(x)\,\&\,Q]P) \vdash \textcircled{b}}}{\forall x\,(P \rightarrow [x' = f(x)\,\&\,Q]P) \vdash \textcircled{a} \land \textcircled{b}}$$

Continuing from the left premise (with ⓐ in its succedent), standard logical manipulation is used to dualize both sides of the sequent. The ∃L step Skolemizes the existential in the antecedent, with the resulting x used to witness the (then) existentially quantified succedent with ∃R:

$$\text{⟨·⟩, ¬L, ¬R}\;\dfrac{\text{∃L}\;\dfrac{\text{∃R, ∧R}\;\dfrac{P, Q, \dot{Q}^{(*)}, \neg(\dot{P}^{(*)}) \vdash \langle x' = f(x)\,\&\,Q\rangle \neg P}{P, Q, \dot{Q}^{(*)}, \neg(\dot{P}^{(*)}) \vdash \exists x\,(P \land \langle x' = f(x)\,\&\,Q\rangle \neg P)}}{\exists x\,\big(P \land Q \land \dot{Q}^{(*)} \land \neg(\dot{P}^{(*)})\big) \vdash \exists x\,(P \land \langle x' = f(x)\,\&\,Q\rangle \neg P)}}{\forall x\,(P \rightarrow [x' = f(x)\,\&\,Q]P) \vdash \forall x\,\big(P \land Q \land \dot{Q}^{(*)} \rightarrow \dot{P}^{(*)}\big)}$$

Next, an initial state assumption $x{=}y$ is introduced by a cut, $\mathbb{R}$ followed by ∃L to Skolemize the resulting existential quantifier. The antecedent assumption $\neg(\dot{P}^{(*)})$ is replaced with $(\dot{\neg P})^{(*)}$ equivalently, by Proposition A.8. Both local progress formulas in the antecedents are then replaced equivalently with the local progress modalities using the derived axiom LP.

$$
\begin{array}{lr}
 & x{=}y, P, Q, \langle x' = f(x) \,\&\, Q\rangle \bigcirc, \langle x' = f(x) \,\&\, \neg P\rangle \bigcirc \vdash \langle x' = f(x) \,\&\, Q\rangle \neg P \\ \hline
\text{LP} & x{=}y, P, Q, \dot{Q}^{(*)}, (\dot{\neg P})^{(*)} \vdash \langle x' = f(x) \,\&\, Q\rangle \neg P \\ \hline
\mathbb{R} & x{=}y, P, Q, \dot{Q}^{(*)}, \neg(\dot{P}^{(*)}) \vdash \langle x' = f(x) \,\&\, Q\rangle \neg P \\ \hline
\exists\text{L} & \exists y\, x{=}y, P, Q, \dot{Q}^{(*)}, \neg(\dot{P}^{(*)}) \vdash \langle x' = f(x) \,\&\, Q\rangle \neg P \\ \hline
\text{cut}, \mathbb{R} & P, Q, \dot{Q}^{(*)}, \neg(\dot{P}^{(*)}) \vdash \langle x' = f(x) \,\&\, Q\rangle \neg P
\end{array}
$$

By Init from Corollary A.3, local progress for Q is strengthened, while $\bigcirc$ can only be unfolded for $\neg P$. The two resulting diamond modality formulas are combined with Uniq:

$$
\begin{array}{lr}
 & \langle x' = f(x) \,\&\, Q \wedge (\neg P \vee x = y)\rangle x \neq y \vdash \langle x' = f(x) \,\&\, Q\rangle \neg P \\ \hline
\text{Uniq} & P, \langle x' = f(x) \,\&\, Q\rangle x \neq y, \langle x' = f(x) \,\&\, \neg P \vee x = y\rangle x \neq y \vdash \langle x' = f(x) \,\&\, Q\rangle \neg P \\ \hline
\text{Init} & x{=}y, P, Q, \langle x' = f(x) \,\&\, Q\rangle \bigcirc, \langle x' = f(x) \,\&\, \neg P\rangle \bigcirc \vdash \langle x' = f(x) \,\&\, Q\rangle \neg P
\end{array}
$$

With a K$\langle\cdot\rangle$, dW step, the diamond modality in the antecedent strengthens to $\neg P$ in its postcondition with the propositional tautology $Q \wedge (\neg P \vee x = y) \rightarrow (x \neq y \rightarrow \neg P)$. A dRW$\langle\cdot\rangle$ step completes the proof using the propositional tautology $Q \wedge (\neg P \vee x = y) \rightarrow Q$.

$$
\begin{array}{lr}
 & * \\ \hline
\text{dRW}\langle\cdot\rangle & \langle x' = f(x) \,\&\, Q \wedge (\neg P \vee x = y)\rangle \neg P \vdash \langle x' = f(x) \,\&\, Q\rangle \neg P \\ \hline
\text{K}\langle\cdot\rangle, \text{dW} & \langle x' = f(x) \,\&\, Q \wedge (\neg P \vee x = y)\rangle x \neq y \vdash \langle x' = f(x) \,\&\, Q\rangle \neg P
\end{array}
$$

The remaining derivation from the right premise (with ⓑ in its succedent) is similar using local progress for the already reflected backward differential equations instead. □

Appendix B

Appendix: Liveness and Existence for Ordinary Differential Equations

B.1 Proof Calculus

This appendix presents derivations and proofs for dL axioms and proof rules that are used in the refinement approach. It also gives a generalized definition of the required topological side conditions for those axioms.

B.1.1 Base Calculus

The bounded differential ghost axiom BDG from Lemma 4.2 (quoted and proved below) is a new vectorial generalization of DG which allows differential ghosts with provably bounded ODEs to be added.

$$\text{BDG}\quad \begin{array}{l}[x' = f(x), y' = g(x, y) \,\&\, Q(x)]\ \|y\|_2^2 \le e(x)\\ \quad \rightarrow \big([x' = f(x) \,\&\, Q(x)]P(x) \leftrightarrow [x' = f(x), y' = g(x, y) \,\&\, Q(x)]P(x)\big)\end{array}$$

Proof of Lemma 4.2. The proof of BDG follows the proof of the differential ghosts axiom [142], but generalizes it to support vectorial, nonlinear ODEs by adding a precondition on boundedness of solutions. Let y be a vector of m fresh variables and $y' = g(x, y)$ be its corresponding vector of ghost ODEs. Both directions of the (inner) equivalence of axiom BDG are proved separately.

“$\rightarrow$” The (easier) “$\rightarrow$” direction does not require the outer bounding assumption of BDG, i.e., the implication $[x' = f(x) \,\&\, Q(x)]P(x) \rightarrow [x' = f(x), y' = g(x, y) \,\&\, Q(x)]P(x)$ is valid for any ODE $y' = g(x, y)$ meeting the freshness condition on y. The proof for this direction is identical to the proof of soundness for differential ghosts [142, Theorem 38].

“$\leftarrow$” In the “$\leftarrow$” direction, consider an initial state $\omega \in \mathbb{S}$ and let $\varphi : [0, T) \rightarrow \mathbb{S}, 0 < T \le \infty$ be the unique, right-maximal solution [33, 204] to the ODE $x' = f(x)$ with initial value $\varphi(0) = \omega$. Similarly, let $\varphi_y : [0, T_y) \rightarrow \mathbb{S}, 0 < T_y \le \infty$ be the unique, right-maximal solution to the ODE $x' = f(x), y' = g(x, y)$ with initial value $\varphi_y(0) = \omega$. Assume that ω

satisfies both of the following assumptions in BDG:

$$\omega \in [\![[x' = f(x), y' = g(x,y) \,\&\, Q(x)]\ \|y\|_2^2 \leq e(x)]\!] \tag{B.1}$$
$$\omega \in [\![[x' = f(x), y' = g(x,y) \,\&\, Q(x)] P(x)]\!] \tag{B.2}$$

To show $\omega \in [\![[x' = f(x) \,\&\, Q(x)]P(x)]\!]$, unfold the semantics of the box modality and consider any finite time τ with $0 \leq \tau < T$ where $\varphi(\zeta) \in [\![Q(x)]\!]$ for all $0 \leq \zeta \leq \tau$. It is proved further below that τ is also in the existence interval for solution φ_y, i.e., $\circledast$: $\tau < T_y$. By uniqueness, φ, φ_y agree on the values of x on their common existence interval, which includes the time interval $[0, \tau]$ by $\circledast$. Therefore, by coincidence for terms and formulas [142], $\varphi_y(\zeta) \in [\![Q(x)]\!]$ for all $0 \leq \zeta \leq \tau$. Thus, by (B.2), $\varphi_y(\tau) \in [\![P(x)]\!]$ and by coincidence for formulas [142], $\varphi(\tau) \in [\![P(x)]\!]$.
In order to prove $\circledast$, suppose for contradiction that $T_y \leq \tau$. Let $x(\cdot) : [0, T) \to \mathbb{R}^n$ denote the projection of solution φ onto its x coordinates, and let $e(x(\cdot)) : [0, T) \to \mathbb{R}$ denote the evaluation of term e along $x(\cdot)$. Since the projection $x(\cdot)$ and its composition with the smooth term evaluation $e(x(\cdot))$ are continuous in t [142], $e(x(\cdot))$ is bounded above by (and attains) its maximum value $e_{\max} \in \mathbb{R}$ on the compact interval $[0, \tau]$.
Let $y(\cdot) : [0, T_y) \to \mathbb{R}^m$ similarly denote the projection of φ_y onto its y coordinates and $\|y(\cdot)\|_2^2$ denote the squared norm evaluated along $y(\cdot)$. Since $T_y \leq \tau < T$, note that $y(\cdot)$ must be the unique right-maximal solution of the time-dependent differential equation $y' = g(x(t), y)$. Otherwise, if there is a longer solution $\psi : [0, \zeta) \to \mathbb{R}^m$ for $y' = g(x(t), y)$ which exists for time ζ with $T_y < \zeta \leq T$, then the combined solution given by $(x(t), \psi(t)) : [0, \zeta) \to \mathbb{R}^n \times \mathbb{R}^m$ extends φ_y beyond T_y (by keeping all variables other than x, y constant at their initial values in state ω). This contradicts right-maximality of φ_y. Moreover, for all times $0 \leq \zeta < T_y$, by assumption $\zeta \leq \tau$ and $\varphi(\zeta) \in [\![Q(x)]\!]$, so the solution φ_y satisfies $\varphi_y(\zeta) \in [\![Q(x)]\!]$ by coincidence for formulas [142]. From (B.1), for all times $0 \leq \zeta < T_y$, the squared norm is bounded by $e_{\max}$, with $\|y(\zeta)\|_2^2 \leq e(x(\zeta)) \leq e_{\max}$. Hence, $y(\cdot)$ remains trapped within the compact $\mathbb{R}^m$ ball of radius $\sqrt{e_{\max}}$ on its domain of definition $[0, T_y)$. By [33, Theorem 1.4], and right-maximality of $y(\cdot)$ for the time-dependent ODE $y' = g(x(t), y)$, the domain of definition of solution $y(\cdot)$ is equal to the domain of definition of $y' = g(x(t), y)$, i.e., $T_y = T$, which contradicts $T_y \leq \tau < T$. □

The following lemma recalls derived dL ODE invariance proof rules from Chapter 3 that are used in the derivations in Appendix B.2 for ease of reference.

Lemma B.1 (ODE invariance proof rules of dL). *The following are derived ODE invariance proof rules of* dL*. In rule dbx$_\succcurlyeq$, g is any cofactor term. In rule sAI&, $\dot{Q}^{(*)}$, $\dot{P}^{(*)}$, $\dot{Q}^{-(*)}$, $(\dot{\neg P})^{-(*)}$ are progress formulas Def. 3.24 with respect to $x' = f(x)$. In rule Enc, formula P is formed from finite conjunctions and disjunctions of strict inequalities $>, <$, and formula $P^{\geq}_{>}$ is identical to P but with non-strict inequalities $\geq, \leq$ in place of $>, <$ respectively.*

$$\text{dbx}_{\succcurlyeq}\ \frac{Q \vdash \dot{e} \geq ge}{e \succcurlyeq 0 \vdash [x' = f(x) \,\&\, Q] e \succcurlyeq 0} \quad (\succcurlyeq \text{ either } \geq \text{ or } >)$$

$$\text{sAI\&}\ \frac{P, Q, \dot{Q}^{(*)} \vdash \dot{P}^{(*)} \qquad \neg P, Q, \dot{Q}^{-(*)} \vdash (\dot{\neg P})^{-(*)}}{P \vdash [x' = f(x) \,\&\, Q] P}$$

$$\text{Barr}\ \frac{Q, e = 0 \vdash \dot{e} > 0}{e \succcurlyeq 0 \vdash [x' = f(x) \,\&\, Q] e \succcurlyeq 0} \qquad \text{(where } \succcurlyeq \text{ is } \geq \text{ or } >\text{)}$$

$$\text{Enc}\ \frac{\Gamma \vdash P^{\geq}_{>} \quad \Gamma \vdash [x' = f(x) \,\&\, Q \land P^{\geq}_{>}] P}{\Gamma \vdash [x' = f(x) \,\&\, Q] P}$$

Proof. These ODE invariance proof rules are all derived from the complete dL axiomatization for ODE invariants from Chapter 3. □

Rule dbx$_{\succcurlyeq}$ is the Darboux inequality proof rule, while rule sAI& is dL's complete proof rule for ODE invariants from Chapter 3. For closed (resp. open) semianalytic formulas P, the right (resp. left) premise of rule sAI& closes trivially (see Section 3.6.2). This simplification is useful for obtaining more succinct proof rules, e.g., rule dV makes use of sAI& with a closed semianalytic formula. Rule Barr is the strict barrier certificates proof rule, which derives from DG for polynomial terms p and as a special case of rule sAI& for extended terms e. Finally, rule Enc says that, in order to prove that solutions stay in postcondition P which characterizes an open set, it suffices to prove it assuming $P^{\geq}_{>}$ in the domain constraint, where $P^{\geq}_{>}$ relaxes all strict inequalities in P and thus provides an over-approximation of the topological closure of the set characterized by P. The rule can also be understood in the contrapositive: if a continuous solution leaves P, then it either already started outside the closure (ruled out by left premise), or it starts in the closure and leaves P on its topological boundary (included in the closure). The latter case is ruled out by the right premise of Enc, which says that solutions remaining in the closure must stay in P.

B.1.2 Refinement Calculus

The following ODE liveness refinement axioms are quoted from Lemma 4.1, and their syntactic derivations in the dL proof calculus are given below.

$$\text{K}\langle\&\rangle \quad [x' = f(x) \,\&\, Q \land \neg P]\neg G \to \big(\langle x' = f(x) \,\&\, Q\rangle G \to \langle x' = f(x) \,\&\, Q\rangle P\big)$$

$$\text{DR}\langle\cdot\rangle \quad [x' = f(x) \,\&\, R] Q \to \big(\langle x' = f(x) \,\&\, R\rangle P \to \langle x' = f(x) \,\&\, Q\rangle P\big)$$

$$\text{BDG}\langle\cdot\rangle \quad \begin{aligned} &[x' = f(x), y' = g(x,y) \,\&\, Q(x)]\ \|y\|_2^2 \leq e(x) \\ &\to \big(\langle x' = f(x) \,\&\, Q(x)\rangle P(x) \to \langle x' = f(x), y' = g(x,y) \,\&\, Q(x)\rangle P(x)\big) \end{aligned}$$

$$\text{DDG}\langle\cdot\rangle \quad \begin{aligned} &[x' = f(x), y' = g(x,y) \,\&\, Q(x)]\ 2y \cdot g(x,y) \leq L(x)\, \|y\|_2^2 + M(x) \\ &\to \big(\langle x' = f(x) \,\&\, Q(x)\rangle P(x) \to \langle x' = f(x), y' = g(x,y) \,\&\, Q(x)\rangle P(x)\big) \end{aligned}$$

Proof of Lemma 4.1. The four axioms are derived in order.

K⟨&⟩ Axiom K⟨&⟩ is derived as follows, starting with ⟨·⟩, ¬L, ¬R to dualize the diamond modalities in the antecedent and succedent to box modalities. A dC step using the right antecedent completes the proof.

$$\dfrac{\text{dC}\ \dfrac{*}{[x' = f(x) \,\&\, Q \land \neg P]\neg G, [x' = f(x) \,\&\, Q]\neg P \vdash [x' = f(x) \,\&\, Q]\neg G}}{\langle\cdot\rangle, \neg\text{L}, \neg\text{R}\ \ [x' = f(x) \,\&\, Q \land \neg P]\neg G, \langle x' = f(x) \,\&\, Q\rangle G \vdash \langle x' = f(x) \,\&\, Q\rangle P}$$

DR⟨·⟩ Axiom DR⟨·⟩ is similarly derived from axiom DMP with ⟨·⟩, see Corollary 3.16.

BDG⟨·⟩ Axiom BDG⟨·⟩ is derived from axiom BDG using axiom ⟨·⟩. The leftmost antecedent is abbreviated with: $R \equiv [x' = f(x), y' = g(x,y) \,\&\, Q(x)]\ \|y\|_2^2 \leq e(x)$.

$$\scriptstyle\langle\cdot\rangle,\,\neg\text{L},\,\neg\text{R}\ \dfrac{\scriptstyle\text{BDG}\ \dfrac{*}{R, [x'{=}f(x), y'{=}g(x,y) \,\&\, Q(x)]\neg P(x) \vdash [x'{=}f(x) \,\&\, Q(x)]\neg P(x)}}{R, \langle x'{=}f(x) \,\&\, Q(x)\rangle P(x) \vdash \langle x'{=}f(x), y'{=}g(x,y) \,\&\, Q(x)\rangle P(x)}$$

DDG⟨·⟩ Axiom DDG⟨·⟩ is derived as a differential version of axiom BDG⟨·⟩ with the aid of DG. The derivation starts with ⟨·⟩, ¬L, ¬R to turn diamond modalities in the sequent to box modalities. Axiom DG then introduces a fresh ghost ODE $z' = L(x)z + M(x)$, where the antecedents are universally quantified over ghost variable z by DG, while the succedent is existentially quantified. All quantifiers are then instantiated using ∀L, ∃R, with $z = \|y\|_2^2$ so that z stores the initial value of the squared norm of y. Axiom BDG is used with $y' = g(x,y)$ as the ghost ODEs and with $e(x,z) = z$. The antecedents are abbreviated as follows and the topmost open premise is abbreviated ①:

$$\begin{aligned}
R &\equiv [x'{=}f(x), y'{=}g(x,y) \,\&\, Q(x)]\ 2y \cdot g(x,y) \leq L(x)\ \|y\|_2^2 + M(x)\\
R_z &\equiv [x'{=}f(x), y'{=}g(x,y), z'{=}L(x)z{+}M(x) \,\&\, Q(x)]\ 2y \cdot g(x,y) \leq L(x)\ \|y\|_2^2 {+} M(x)\\
S &\equiv [x'{=}f(x), y'{=}g(x,y) \,\&\, Q(x)]\neg P(x)\\
S_z &\equiv [x'{=}f(x), y'{=}g(x,y), z'{=}L(x)z + M(x) \,\&\, Q(x)]\neg P(x)\\
① &\equiv z{=}\ \|y\|_2^2, R_z \vdash [x'{=}f(x), y'{=}g(x,y), z'{=}L(x)z + M(x) \,\&\, Q(x)]\ \|y\|_2^2 \leq z
\end{aligned}$$

$$\scriptstyle\langle\cdot\rangle,\,\neg\text{L},\,\neg\text{R}\ \dfrac{\scriptstyle\text{DG}\ \dfrac{\scriptstyle\forall\text{L},\,\exists\text{R}\ \dfrac{\scriptstyle\text{BDG}\ \dfrac{①}{z = \|y\|_2^2, R_z, S_z \vdash [x' = f(x), z' = L(x)z + M(x) \,\&\, Q(x)]\neg P(x)}}{\forall z\, R_z, \forall z\, S_z \vdash \exists z\, [x' = f(x), z' = L(x)z + M(x) \,\&\, Q(x)]\neg P(x)}}{R, S \vdash [x' = f(x) \,\&\, Q(x)]\neg P(x)}}{R, \langle x' = f(x) \,\&\, Q(x)\rangle P(x) \vdash \langle x' = f(x), y' = g(x,y) \,\&\, Q(x)\rangle P(x)}$$

From the open premise ①, a dC step adds the postcondition of R_z to the domain constraint of the succedent, while M['] rearranges the postcondition into the form expected by rule $\text{dbx}_{\succcurlyeq}$. The proof is completed using $\text{dbx}_{\succcurlyeq}$ with cofactor $g = L(x)$. Its resulting arithmetical premise is proved by $\mathbb{R}$ because the Lie derivative of $z - \|y\|_2^2$ is bounded above by the following calculation, where the inequality from the domain constraint is used in the second step.

$$\begin{aligned}
\mathcal{L}_{x'=f(x), y'=g(x,y), z'=L(x)z+M(x)}(z - \|y\|_2^2) &= L(x)z + M(x) - 2y \cdot g(x,y)\\
&\geq L(x)z + M(x) - (L(x)\ \|y\|_2^2 + M(x))\\
&= L(x)(z - \|y\|_2^2)
\end{aligned}$$

The ODEs $x' = f(x), y' = g(x,y), z' = L(x)z{+}M(x)$ are abbreviated . . . in the derivation below and the premise after $\text{dbx}_{\succcurlyeq}$ is abbreviated with:

$$② \equiv 2y \cdot g(x,y) {\leq} L(x)\ \|y\|_2^2 + M(x) \vdash L(x)z + M(x) - 2y \cdot g(x,y) {\geq} L(x)(z - \|y\|_2^2)$$

$$\text{dC}\ \dfrac{\text{M}[']\ \dfrac{\text{dbx}_{\succcurlyeq}\ \dfrac{\mathbb{R}\ \dfrac{*}{\textcircled{2}}}{z = \|y\|_2^2 \vdash [\ldots \,\&\, Q(x) \land 2y \cdot g(x,y) \le L(x)\,\|y\|_2^2 + M(x)]\; z - \|y\|_2^2 \ge 0}}{z = \|y\|_2^2 \vdash [\ldots \,\&\, Q(x) \land 2y \cdot g(x,y) \le L(x)\,\|y\|_2^2 + M(x)]\; \|y\|_2^2 \le z}}{z = \|y\|_2^2, R_z \vdash [\ldots \,\&\, Q(x)]\; \|y\|_2^2 \le z} \qquad \square$$

The following topological $\langle\cdot\rangle$ ODE refinement axioms are quoted from Lemmas 4.3 and 4.35. The topological side conditions for these axioms are listed in Lemmas 4.3 and 4.35 respectively. For semianalytic postcondition P and domain constraints Q, R, these refinement axioms are derived syntactically from the real induction axiom in Section 3.5.2. For the sake of generality, the proofs below directly use the topological conditions.

COR $\neg P \land [x' = f(x) \,\&\, R \land \neg P]Q \rightarrow \big(\langle x' = f(x) \,\&\, R\rangle P \rightarrow \langle x' = f(x) \,\&\, Q\rangle P\big)$

CR $\neg P \land [x' = f(x) \,\&\, R \land \neg P]\mathring{Q} \rightarrow \big(\langle x' = f(x) \,\&\, R\rangle P \rightarrow \langle x' = f(x) \,\&\, Q\rangle P\big)$

SAR $[x' = f(x) \,\&\, R \land \neg(P \land Q)]Q \rightarrow \big(\langle x' = f(x) \,\&\, R\rangle P \rightarrow \langle x' = f(x) \,\&\, Q\rangle P\big)$

Proof of Lemmas 4.3 and 4.35. Let $\omega \in \mathbb{S}$ and $\varphi : [0, T) \rightarrow \mathbb{S}, 0 < T \le \infty$ be the unique, right-maximal solution [33, 204] to the ODE $x' = f(x)$ with initial value $\varphi(0) = \omega$. By definition, φ is differentiable, and therefore continuous. This proof uses the fact that preimages under continuous functions of open sets are open [168, Theorem 4.8]. In particular, for an open set $\mathcal{O}$, if $\varphi(t) \in \mathcal{O}$ at some time $0 < t < T$ then the preimage of a sufficiently small open ball $\mathcal{O}_\varepsilon \subseteq \mathcal{O}$ centered at $\varphi(t)$ is open. Thus, if $t > 0$ and $\varphi(t) \in \mathcal{O}$, then φ stays in the open set $\mathcal{O}$ for some open time interval[1]around t, i.e., for some $\varepsilon > 0$:

$$\varphi(\zeta) \in \mathcal{O} \text{ for all } t - \varepsilon \le \zeta \le t + \varepsilon \tag{B.3}$$

For the soundness proof of axioms COR, CR, and SAR, assume that $\omega \in [\![\langle x' = f(x) \,\&\, R\rangle P]\!]$, i.e., there is a time $\tau \in [0, T)$ such that $\varphi(\tau) \in [\![P]\!]$ and $\varphi(\zeta) \in [\![R]\!]$ for all $0 \le \zeta \le \tau$. The proofs make use of the following set $\mathbb{T}$ containing all times t such that the solution φ never enters P on the time interval $[0, t]$.

$$\mathbb{T} \equiv \{t \mid \varphi(\zeta) \notin [\![P]\!] \text{ for all } 0 \le \zeta \le t\} \tag{B.4}$$

COR For axiom COR, assume that $\omega \in [\![\neg P \land [x' = f(x) \,\&\, R \land \neg P]Q]\!]$. The set of times $\mathbb{T}$ (B.4) is non-empty since $\omega = \varphi(0) \notin [\![P]\!]$ so it has a supremum t with $0 \le t \le \tau$ and $\varphi(\zeta) \notin [\![P]\!]$ for all $0 \le \zeta < t$.

- Suppose P, Q both characterize topologically closed sets. Since P characterizes a topologically closed set, its complement formula $\neg P$ characterizes a topologically open set. If $\varphi(t) \notin [\![P]\!]$, i.e., $\varphi(t) \in [\![\neg P]\!]$, then $t < \tau$ and by (B.3), the solution stays in $\neg P$ until time $t + \varepsilon$ for some $\varepsilon > 0$, so t is not the supremum of $\mathbb{T}$, which is a contradiction. Thus, $\varphi(t) \in [\![P]\!]$ and $0 < t$ because $\varphi(0) \notin [\![P]\!]$. Hence, $\varphi(\zeta) \in [\![R \land \neg P]\!]$ for

[1] In case $t = 0$, the time interval in (B.3) is truncated to the left with $\varphi(\zeta) \in \mathcal{O}$ for all $0 \le \zeta < t + \varepsilon$.

all $0 \leq \zeta < t$, which, together with the assumption $\omega \in [\![[x' = f(x) \,\&\, R \wedge \neg P]Q]\!]$ implies $\varphi(\zeta) \in [\![Q]\!]$ for all $0 \leq \zeta < t$. Since Q characterizes a topologically closed set, this implies $\varphi(t) \in [\![Q]\!]$; otherwise, $\varphi(t) \in [\![\neg Q]\!]$ and $\neg Q$ characterizes an open set, so (B.3) implies $\varphi(\zeta) \in [\![\neg Q]\!]$ for some $0 \leq \zeta < t$, which contradicts the earlier observation that $\varphi(\zeta) \in [\![Q]\!]$ for all $0 \leq \zeta < t$. Thus, $\omega \in [\![\langle x' = f(x) \,\&\, Q\rangle P]\!]$ because $\varphi(t) \in [\![P]\!]$ and $\varphi(\zeta) \in [\![Q]\!]$ for all $0 \leq \zeta \leq t$.

- Suppose P, Q both characterize topologically open sets. Then, $\varphi(t) \notin [\![P]\!]$; otherwise, $\varphi(t) \in [\![P]\!]$ and since P characterizes an open set, by (B.3), there is a time $0 \leq \zeta < t$ where $\varphi(\zeta) \in [\![P]\!]$, which contradicts t being the supremum of $\mathbb{T}$. Note that $t < \tau$ and $\varphi(\zeta) \in [\![R \wedge \neg P]\!]$ for all $0 \leq \zeta \leq t$, which, together with the assumption $\omega \in [\![[x' = f(x) \,\&\, R \wedge \neg P]Q]\!]$ implies $\varphi(\zeta) \in [\![Q]\!]$ for all $0 \leq \zeta \leq t$. Since Q characterizes a topologically open set, by (B.3), there exists $\varepsilon > 0$ where $t + \varepsilon < \tau$ such that $\varphi(t+\zeta) \in [\![Q]\!]$ for all $0 \leq \zeta \leq \varepsilon$. By definition of the supremum, for every such $\varepsilon > 0$, there exists ζ where $0 < \zeta \leq \varepsilon$ and $\varphi(t + \zeta) \in [\![P]\!]$, which yields the desired conclusion.

CR For axiom CR, assume that $\omega \in [\![\neg P]\!]$ and

$$\omega \in [\![[x' = f(x) \,\&\, R \wedge \neg P]\mathring{Q}]\!] \tag{B.5}$$

The set of times $\mathbb{T}$ (B.4) is non-empty since $\omega = \varphi(0) \notin [\![P]\!]$ so it has a supremum t with $0 \leq t \leq \tau$ and $\varphi(\zeta) \notin [\![P]\!]$ for all $0 \leq \zeta < t$. Furthermore, $\varphi(\zeta) \in [\![R \wedge \neg P]\!]$ for all $0 \leq \zeta < t$, so by (B.5), $\varphi(\zeta) \in [\![\mathring{Q}]\!]$ for all $0 \leq \zeta < t$. By assumption, formula $\mathring{Q}$ characterizes the open topological interior of the closed formula Q so by continuity of φ, $\varphi(t) \in [\![Q]\!]$. Furthermore, the interior of a set is contained in the set itself, i.e., $[\![\mathring{Q}]\!] \subseteq [\![Q]\!]$, so $\varphi(\zeta) \in [\![Q]\!]$ for all $0 \leq \zeta \leq t$. Classically, either $\varphi(t) \in [\![P]\!]$ or $\varphi(t) \notin [\![P]\!]$.

- If $\varphi(t) \in [\![P]\!]$, then since $\varphi(\zeta) \in [\![Q]\!]$ for all $0 \leq \zeta \leq t$, by definition, $\omega \in [\![\langle x' = f(x) \,\&\, Q\rangle P]\!]$.
- If $\varphi(t) \notin [\![P]\!]$, then $t < \tau$ and furthermore, by (B.5), $\varphi(t) \in [\![\mathring{Q}]\!]$. Since the interior is topologically open, by (B.3), there exists $\varepsilon > 0$ where $t + \varepsilon < \tau$ such that $\varphi(t+\zeta) \in [\![\mathring{Q}]\!] \subseteq [\![Q]\!]$ for all $0 \leq \zeta \leq \varepsilon$. By definition of the supremum, for every such $\varepsilon > 0$, there exists ζ where $0 < \zeta \leq \varepsilon$ and $\varphi(t + \zeta) \in [\![P]\!]$, which yields the desired conclusion.

SAR For axiom SAR, assume that

$$\omega \in [\![[x' = f(x) \,\&\, R \wedge \neg(P \wedge Q)]Q]\!] \tag{B.6}$$

If $\omega \in [\![P \wedge Q]\!]$, then $\omega \in \langle x' = f(x) \,\&\, Q\rangle P$ trivially by following the solution φ for duration 0. Thus, assume $\omega \notin [\![P \wedge Q]\!]$. From (B.6), $\omega \in [\![Q]\!]$ which further implies $\omega \notin [\![P]\!]$. The set of times $\mathbb{T}$ (B.4) is non-empty since $\omega = \varphi(0) \notin [\![P]\!]$ and has a supremum t with $0 \leq t \leq \tau$ and $\varphi(\zeta) \notin [\![P]\!]$ for all $0 \leq \zeta < t$. Thus, $\varphi(\zeta) \in [\![R \wedge \neg(P \wedge Q)]\!]$ for all $0 \leq \zeta < t$. By (B.6), $\varphi(\zeta) \in [\![Q]\!]$ for all $0 \leq \zeta < t$. Classically, either $\varphi(t) \in [\![P]\!]$ or $\varphi(t) \notin [\![P]\!]$.

- Suppose $\varphi(t) \in [\![P]\!]$, if $\varphi(t) \in [\![Q]\!]$, then $\varphi(\zeta) \in [\![Q]\!]$ for all $0 \leq \zeta \leq t$ and so, by definition, $\omega \in [\![\langle x' = f(x) \,\&\, Q \rangle P]\!]$. On the other hand, if $\varphi(t) \notin [\![Q]\!]$, then $\varphi(\zeta) \in [\![R \land \neg(P \land Q)]\!]$ for all $0 \leq \zeta \leq t$, so from (B.6), $\varphi(t) \in [\![Q]\!]$, which yields a contradiction.

  If the formula P is further assumed to characterize a closed set, this sub-case (with $\varphi(t) \in [\![P]\!]$) is the only possibility. Otherwise, $\varphi(t) \in [\![\neg P]\!]$ and $\neg P$ characterizes an open set, so by (B.3), for some $\varepsilon > 0$, $\varphi(t + \zeta) \in [\![\neg P]\!]$ for all $0 \leq \zeta < \varepsilon$ which contradicts t being the supremum of $\mathbb{T}$.

- Suppose $\varphi(t) \notin [\![P]\!]$, then $t < \tau$ and $\varphi(\zeta) \in [\![R \land \neg(P \land Q)]\!]$ for all $0 \leq \zeta \leq t$, so from (B.6), $\varphi(t) \in [\![Q]\!]$. Since Q is a semianalytic formula, solutions of the ODEs either locally progress into the set characterized by Q or $\neg Q$ by Corollary 3.27,[2] i.e., there exists $\varepsilon > 0$, where $t + \varepsilon < \tau$, such that either ① $\varphi(t + \zeta) \in [\![Q]\!]$ for all $0 < \zeta \leq \varepsilon$ or ② $\varphi(t + \zeta) \notin [\![Q]\!]$ for all $0 < \zeta \leq \varepsilon$. Since t is the supremum of $\mathbb{T}$, by definition, for every such ε there exists ζ where $0 < \zeta \leq \varepsilon$ and $\varphi(t + \zeta) \in [\![P]\!]$. In case ①, since $\varphi(t + \zeta) \in [\![P]\!]$ and $\varphi(\nu) \in [\![Q]\!]$ for all $0 \leq \nu \leq t + \zeta$, then $\omega \in [\![\langle x' = f(x) \,\&\, Q \rangle P]\!]$. If the formula Q is further assumed to characterize an open set, this sub-case (①) is the only possibility, even if Q is not a formula of first-order real arithmetic, because $\varphi(t) \in [\![Q]\!]$ implies φ continues to satisfy Q for some time interval to the right of t by (B.3). In case ②, observe that $\varphi(\nu) \in [\![R \land \neg(P \land Q)]\!]$ for all $0 \leq \nu \leq t + \zeta$, from (B.6), $\varphi(t + \zeta) \in [\![Q]\!]$, which yields a contradiction. □

B.1.3 Topological Side Conditions

In Section 2.2.2, topological conditions are defined for formulas ϕ that only mention free variables x occurring in an ODE $x' = f(x)$. For example, ϕ is said to characterize an open set with respect to x iff the set $[\![\phi]\!]$ is open when considered as a subset of $\mathbb{R}^n$ by projecting its semantics over variables $x = (x_1, \ldots, x_n)$. This section defines a more general notion, where ϕ is allowed to mention additional free parameters y that do not occur in the ODE. Adopting these (parametric) side conditions makes the topological refinement axioms that use them, like COR, CR, more general. Let $(y_1, \ldots, y_r) = \mathbb{V} \setminus \{x\}$ be parameters, and $\omega \in \mathbb{S}$ be a state. For brevity, write $y = (y_1, \ldots, y_r)$ for the parameters and $\omega(y) = (\omega(y_1), \ldots, \omega(y_r)) \in \mathbb{R}^r$ for the componentwise projection, and similarly for $\omega(x) \in \mathbb{R}^n$. Given the set $[\![\phi]\!] \subseteq \mathbb{S}$ and $\gamma \in \mathbb{R}^r$, define:

$$([\![\phi]\!])_\gamma \overset{\text{def}}{=} \{\omega(x) \in \mathbb{R}^n \mid \omega \in [\![\phi]\!], \omega(y) = \gamma\}$$

The set $([\![\phi]\!])_\gamma \subseteq \mathbb{R}^n$ is the projection onto variables x of all states ω that satisfy ϕ and having values γ for the parameters y. Formula ϕ *characterizes* a (topologically) open (resp. closed, bounded, compact) set with respect to variables x iff for all $\gamma \in \mathbb{R}^r$, the set $([\![\phi]\!])_\gamma \subseteq \mathbb{R}^n$ is topologically open (resp. closed, bounded, compact) with respect to the Euclidean topology.

These topological side conditions are even decidable [14, 197] for first-order formulas of real arithmetic over polynomial terms because in Euclidean spaces they can be phrased as conditions

[2] This property is specific to sets characterized by semianalytic formulas and ODEs (and certain topologically well-behaved extensions [174, 176]) and is not true for arbitrary sets and ODEs.

involving the first-order quantifiers. The following conditions are standard [14], although special care is taken to universally quantify over the parameters y. Let $P(x, y)$ be a formula mentioning variables x and parameters y, then it is (with respect to variables x):

- *open* if the formula $\forall y\,\forall x\,\Big(P(x,y) \rightarrow \exists \varepsilon{>}0\,\forall z\,\big(\,\|x-z\|_2^2 < \varepsilon^2 \rightarrow P(z,y)\big)\Big)$ is valid, where the variables $z = (z_1, \ldots, z_n)$ are fresh for $P(x, y)$,
- *closed* if its complement formula $\neg P(x, y)$ is open,
- *bounded* if the formula $\forall y\,\exists r{>}0\,\forall x\,\big(P(x,y) \rightarrow \|x\|_2^2{<}r^2\big)$ is valid, where variable r is fresh for $P(x, y)$, and
- *compact* if it is closed and bounded, by the Heine-Borel theorem [168, Theorem 2.4.1].

There are also syntactic criteria that are sufficient (but not necessary[3]) for checking whether a formula satisfies the semantic conditions. For example, the formula $P(x, y)$ is (with respect to variables x):

- *open* if it is formed from finite conjunctions and disjunctions of strict inequalities ($\neq, >, <$),
- *closed* if it is formed from finite conjunctions and disjunctions of non-strict inequalities ($=, \geq, \leq$),
- *bounded* if it is of the form $\|x\|_2^2 \preccurlyeq e(y) \land R(x,y)$, where $e(y)$ is a term depending only on parameters y and $R(x, y)$ is a formula. This syntactic criterion uses the fact that the intersection of a bounded set (characterized by $\|x\|_2^2 \preccurlyeq e(y)$) with any set (characterized by $R(x, y)$) is bounded, and the formula $P(x, y)$ is *compact* if $\preccurlyeq$ is $\leq$ and $R(x, y)$ is closed.

The importance of these syntactic criteria is they are easily checkable by an implementation that inspects the syntactic shape of input formulas P, even if P contains extended terms with undecidable arithmetic [162]. In contrast, checking the semantic topological conditions for P requires invoking expensive real arithmetic decision procedures and those procedures are only guaranteed to work if P is a $\mathrm{FOL}_{\mathbb{R}}$ formula over polynomial terms. As an example, the syntactic side condition of rule cR from Corollary 4.36 enables its effective implementation, compared to its underlying axiom CR from Lemma 4.35 which is more general but uses requires checking semantic side conditions.

Notational Conventions (Topological Side Conditions and Arithmetic). The derivations used in this thesis often require properties of the smooth term semantics from real analysis. For example, continuous functions on compact domains attain their extrema [168, Theorem 4.16] so, for any compact formula $P(x, y)$ and term $e(x)$, the formula $\exists m\,\forall x\,(P(x,y) \rightarrow e(x) \leq m)$ is a valid arithmetic formula expressing that m is an upper bound of e on P. If e is a polynomial

[3] If there are no parameters y, these syntactic checks are “necessary” conditions for semialgebraic formulas in the sense that every open (resp. closed) semialgebraic formula P is provably equivalent in real arithmetic to a (computable) formula formed from finite conjunctions and disjunctions of strict (resp. non-strict) inequalities [14, Theorem 2.7.2].

and P is a formula of real arithmetic over polynomial terms, then $\mathbb{R}$ proves the existence of m above. For simplicity, such properties are also assumed to be provable by $\mathbb{R}$ in derivations involving extended terms and the corresponding argument from real analysis is also provided. This notational simplification is sound with the implicit understanding that any such real analytic properties are axiomatized as additional (arithmetic) axioms for extended dL terms.

B.2 Derived Existence and Liveness Proof Rules

This appendix gives all omitted syntactic derivations of the existence and liveness proof rules in Chapter 4. For ease of reference, this appendix is organized into four sections, corresponding to Sections 4.3–4.6 of Chapter 4. The high-level intuition behind these proofs is available as proof sketches in their respective sections, while motivation for important proof steps is given directly in the subsequent proofs. Further motivation for the surveyed liveness arguments can also be found in their original presentations [137, 156, 157, 159, 176, 191].

B.2.1 Proofs for Finite-Time Blow Up and Global Existence

Proof of Corollary 4.6. Assume that the ODE $x' = f(x)$ is in dependency order (4.9). The derivation successively removes the ODEs $y_k, y_{k-1}, \dots, y_1$ in reverse dependency order using either axiom BDG⟨·⟩ or DDG⟨·⟩, as shown below. This continues until all of the ODEs are removed and the rightmost premise closes by axiom TEx. The left premises arising from refinement with axioms BDG⟨·⟩, DDG⟨·⟩ are the premises of rule DEx. They are collectively labeled $\circledast$ and explained below.

$$\text{∀R}\dfrac{\text{BDG⟨·⟩, DDG⟨·⟩}\dfrac{\circledast \qquad \text{BDG⟨·⟩, DDG⟨·⟩}\dfrac{\circledast \qquad \text{BDG⟨·⟩, DDG⟨·⟩}\dfrac{\circledast \qquad \text{TEx}\dfrac{*}{\Gamma \vdash \langle t'{=}1\rangle t{>}\tau}}{\vdots}}{\Gamma \vdash \langle y_1'{=}g_1(y_1), \dots, y_{k-1}'{=}g_{k-1}(y_1, \dots, y_{k-1}), t'{=}1\rangle\, t{>}\tau}}{\Gamma \vdash \langle y_1'{=}g_1(y_1), \dots, y_{k-1}'{=}g_{k-1}(y_1, \dots, y_{k-1}), y_k'{=}g_k(y_1, \dots, y_k), t'{=}1\rangle\, t{>}\tau}}{\Gamma \vdash \forall \tau\, \langle \underbrace{y_1'{=}g_1(y_1), \dots, y_{k-1}'{=}g_{k-1}(y_1, \dots, y_{k-1}), y_k'{=}g_k(y_1, \dots, y_k)}_{x'{=}f(x) \text{ written in dependency order}}, t'{=}1\rangle\, t{>}\tau}$$

At each step $i = k, \dots, 1$, the ODE y_i in the succeedent is removed using either axiom BDG⟨·⟩ or DDG⟨·⟩, depending on the user-chosen form (Corollary 4.6) of postcondition P_i.

Ⓑ In case formula $P_i \equiv \|y_i\|_2^2 \leq e_i(t, y_1, \dots, y_{i-1})$ is of form Ⓑ (as defined in Corollary 4.6), axiom BDG⟨·⟩ is used. This yields the two stacked premises shown below, where the top premise corresponds to premise $\circledast$ above. The dependency order (4.9) enables the sound use of axiom BDG⟨·⟩ for this refinement step because the ODEs for $y_1, \dots, y_{i-1}$ are not allowed to depend on variables y_i. The term $e(t, y_1, \dots, y_{i-1})$ also meets the dependency requirements of BDG⟨·⟩ because it does not depend on y_i.

$$\text{BDG⟨·⟩}\dfrac{\begin{array}{l}\Gamma \vdash [y_1' = g_1(y_1), \dots, y_{i-1}' = g_{i-1}(y_1, \dots, y_{i-1}), y_i' = g_i(y_1, \dots, y_i), t' = 1]P_i \\ \Gamma \vdash \langle y_1' = g_1(y_1), \dots, y_{i-1}' = g_{i-1}(y_1, \dots, y_{i-1}), t' = 1\rangle\, t > \tau\end{array}}{\Gamma \vdash \langle y_1' = g_1(y_1), \dots, y_{i-1}' = g_{i-1}(y_1, \dots, y_{i-1}), y_i' = g_i(y_1, \dots, y_i), t' = 1\rangle\, t > \tau}$$

- Ⓓ In case formula $P_i \equiv 2y_i \cdot g_i(y_1,\ldots,y_i) \leq L_i(t,y_1,\ldots,y_{i-1})\left\|y_i\right\|_2^2 + M_i(t,y_1,\ldots,y_{i-1})$ is of form Ⓓ (as defined in Corollary 4.6), axiom DDG⟨·⟩ is used instead. Again, terms $L_i(t,y_1,\ldots,y_{i-1}), M_i(t,y_1,\ldots,y_{i-1})$ meet the dependency requirements of DDG⟨·⟩ because they do not depend on y_i. The top premise corresponds to premise ⍟ above, while the ODE for y_i is removed in the bottom premise.

$$\text{DDG}\langle\cdot\rangle\;\dfrac{\begin{array}{l}\Gamma \vdash [y_1' = g_1(y_1),\ldots,y_{i-1}' = g_{i-1}(y_1,\ldots,y_{i-1}), y_i' = g_i(y_1,\ldots,y_i), t'=1]P_i \\ \Gamma \vdash \langle y_1' = g_1(y_1),\ldots,y_{i-1}' = g_{i-1}(y_1,\ldots,y_{i-1}), t'=1\rangle\, t > \tau\end{array}}{\Gamma \vdash \langle y_1' = g_1(y_1),\ldots,y_{i-1}' = g_{i-1}(y_1,\ldots,y_{i-1}), y_i' = g_i(y_1,\ldots,y_i), t'=1\rangle\, t > \tau} \qquad \square$$

Proof of Corollary 4.7. The proof closely follows the proof sketch for Corollary 4.7 but with an extra step to ensure that the chosen terms L, M are within the term language of dL. Let the ODE $x' = f(x)$ be globally Lipschitz and C be the (positive) Lipschitz constant for f, i.e., $\left\|f(x) - f(y)\right\|_2 \leq C\left\|x-y\right\|_2$. Then f satisfies the following inequality, where the first step (4.12) is proved in the sketch but its RHS contains norms $\left\|\cdot\right\|_2$ which are not in the dL term syntax. The inequality (4.12) is prolonged by using inequality (4.10) to remove these non-squared norm terms, which yields corresponding choices of bounding dL terms L, M.

$$\begin{aligned} 2x \cdot f(x) &\overset{(4.12)}{\leq} \big(2C + \left\|f(0)\right\|_2\big)\left\|x\right\|_2^2 + \left\|f(0)\right\|_2 \\ &\overset{(4.10)}{\leq} \underbrace{\big(2C + \frac{1}{2}(1 + \left\|f(0)\right\|_2^2)\big)}_{L}\left\|x\right\|_2^2 + \underbrace{\frac{1}{2}(1 + \left\|f(0)\right\|_2^2)}_{M} \end{aligned} \tag{B.7}$$

The inequality (B.7) is a valid real arithmetic formula and is thus provable by rule $\mathbb{R}$. This enables the derivation below using axiom DDG⟨·⟩ because L, M satisfy the respective variable constraints of the axiom. The resulting left premise is proved, after a dW step, by $\mathbb{R}$. The resulting right premise, after the ODEs $x' = f(x)$ have been removed, is proved by axiom TEx.

$$\forall\text{R}\;\dfrac{\text{DDG}\langle\cdot\rangle\;\dfrac{\text{dW}\;\dfrac{\mathbb{R}\;\dfrac{*}{\vdash 2x \cdot f(x) \leq L\left\|x\right\|_2^2 + M}}{\vdash [x' = f(x), t'=1]\, 2x \cdot f(x) \leq L\left\|x\right\|_2^2 + M} \qquad \text{TEx}\;\dfrac{*}{\vdash \langle t'=1\rangle\, t > \tau}}{\vdash \langle x' = f(x), t'=1\rangle\, t > \tau}}{\vdash \forall\tau\, \langle x' = f(x), t'=1\rangle\, t > \tau} \qquad \square$$

Proof of Corollary 4.8. Assume that ODE $x' = f(x)$ has affine dependency order (4.9) where each ODE $y_i' = g_i(y_1,\ldots,y_i)$ is of the affine form $y_i' = A_i(y_1,\ldots,y_{i-1})y_i + b_i(y_1,\ldots,y_{i-1})$ for some matrix and vector terms A_i, b_i respectively with the indicated variable dependencies. From the proof sketch for Corollary 4.8, A_i, b_i satisfy inequality (4.13) for each $i = 1,\ldots,k$. Like the proof of inequality (B.7), inequality (4.13) is prolonged by inequality (4.10) to remove non-squared norm terms in its RHS, to obtain corresponding choices of bounding dL terms L_i, M_i.

$$\begin{aligned} 2y_i \cdot (A_i y_i + b_i) &\overset{(4.13)}{\leq} (2\left\|A_i\right\|_F + \left\|b_i\right\|_2)\left\|y_i\right\|_2^2 + \left\|b_i\right\|_2 \\ &\overset{(4.10)}{\leq} \underbrace{\big(1 + \left\|A_i\right\|_F^2 + \frac{1}{2}(1 + \left\|b_i\right\|_2^2)\big)}_{L_i}\left\|y_i\right\|_2^2 + \underbrace{\frac{1}{2}(1 + \left\|b_i\right\|_2^2)}_{M_i} \end{aligned} \tag{B.8}$$

The inequality from (B.8) is a valid real arithmetic formula, and thus provable by $\mathbb{R}$ for each $i = 1, \ldots, k$. The derivation uses rule DEx, where the postcondition of each premise is chosen to be of form Ⓓ. The resulting premises are all proved, after a dW step, by $\mathbb{R}$ with the above choice of L_i, M_i for each $i = 1, \ldots, k$.

$$\text{DEx}\,\dfrac{\text{dW}\,\dfrac{\mathbb{R}\,\dfrac{*}{\vdash 2y_1 \cdot (A_1y_1{+}b_1){\le}L_1 \left\|y_1\right\|_2^2 {+}M_1}}{\vdash [y_1'{=}g_1(y_1), t'{=}1]P_1} \quad \cdots \quad \text{dW}\,\dfrac{\mathbb{R}\,\dfrac{*}{\vdash 2y_k \cdot (A_ky_k{+}b_k){\le}L_k \left\|y_k\right\|_2^2 {+}M_k}}{\vdash [y_1'{=}g_1(y_1), \ldots, y_k'{=}g_k(y_1, \ldots, y_k), t'{=}1]P_k}}{\vdash \forall \tau \, \langle x'{=}f(x), t'{=}1\rangle \, t > \tau} \qquad \square$$

Proof of Corollary 4.10. The derivation starts by Skolemizing with ∀R, then switching the diamond modality in the succedent to a box modality in the antecedent using $\langle\cdot\rangle$, ¬R. The postcondition of the box modality is simplified using the propositional tautologies $\neg(\phi \vee \psi) \leftrightarrow \neg\phi \wedge \neg\psi$ and $\neg\neg\phi \leftrightarrow \phi$. Axiom $[\cdot]\wedge$, ∧L splits the conjunction in the antecedent, before $\langle\cdot\rangle$ is used again to flip the left antecedent to a diamond modality in the succedent. These (mostly) propositional steps recover the more verbose phrasing of BEx from (4.14).

$$\forall\text{R}\,\dfrac{\langle\cdot\rangle, \neg\text{R}\,\dfrac{[\cdot]\wedge, \wedge\text{L}\,\dfrac{\langle\cdot\rangle, \neg\text{L}\,\dfrac{[x'{=}f(x), t'{=}1]B(x) \vdash \langle x'{=}f(x), t'{=}1\rangle \, t > \tau}{[x'{=}f(x), t'{=}1]\neg(t > \tau), [x'{=}f(x), t'{=}1]B(x) \vdash \mathit{false}}}{[x'{=}f(x), t'{=}1](\neg(t > \tau) \wedge B(x)) \vdash \mathit{false}}}{\vdash \langle x'{=}f(x), t'{=}1\rangle (t > \tau \vee \neg B(x))}}{\vdash \forall\tau \, \langle x'{=}f(x), t'{=}1\rangle (t > \tau \vee \neg B(x))}$$

The formula $B(x)$ is assumed to characterize a bounded set with respect to the variables x. The closure of this set (with respect to x) is compact so the continuous norm function $\left\|\cdot\right\|_2^2$ attains its maximum value on that set. Hence, the formula $\exists D \, \forall x \, (B(x) \rightarrow \left\|x\right\|_2^2 \leq D)$ is valid in first-order real arithmetic and is thus provable by rule $\mathbb{R}$ (Appendix B.1.3). The derivation continues with a cut of this formula and Skolemizing with ∃L. Axiom BDG$\langle\cdot\rangle$ is then used to remove the ODE $x' = f(x)$ with $e(x) = D$. The resulting right premise is proved by axiom TEx, while the resulting left premise is labeled ① and continued below.

$$\text{cut}, \mathbb{R}, \exists\text{L}\,\dfrac{\text{BDG}\langle\cdot\rangle\,\dfrac{① \qquad \text{TEx}\,\dfrac{*}{\vdash \langle t' = 1\rangle \, t > \tau}}{[x' = f(x), t' = 1]B(x), \forall x \, (B(x) \rightarrow \left\|x\right\|_2^2 \leq D) \vdash \langle x' = f(x), t' = 1\rangle \, t > \tau}}{[x' = f(x), t' = 1]B(x) \vdash \langle x' = f(x), t' = 1\rangle \, t > \tau}$$

From premise ①, a dC step adds the postcondition of the leftmost antecedent, $B(x)$, to the domain constraint. Since the remaining antecedent is universally quantified over variables x, it is soundly kept across an application of a subsequent dW step and the proof is completed with ∀L, →L.

$$\text{dC}\,\dfrac{\text{dW}\,\dfrac{\forall\text{L}, \rightarrow\text{L}\,\dfrac{*}{\forall x \, (B(x) \rightarrow \left\|x\right\|_2^2 \leq D), B(x) \vdash \left\|x\right\|_2^2 \leq D}}{\forall x \, (B(x) \rightarrow \left\|x\right\|_2^2 \leq D) \vdash [x' = f(x), t' = 1 \,\&\, B(x)] \, \left\|x\right\|_2^2 \leq D}}{[x' = f(x), t' = 1]B(x), \forall x \, (B(x) \rightarrow \left\|x\right\|_2^2 \leq D) \vdash [x' = f(x), t' = 1] \, \left\|x\right\|_2^2 \leq D} \qquad \square$$

Proof of Corollary 4.12. Assume ODE $x' = f(x)$ is in dependency order (4.9) and indices $i = 1, \dots, k$ are partitioned into disjoint sets L, N as in Corollary 4.12. The first step in the derivation Skolemizes the succedent with ∀R.

$$\forall\mathrm{R}\,\dfrac{\vdash \langle x' = f(x), t' = 1\rangle\big(t > \tau \lor \bigvee_{j\in N} \neg B_j(y_j)\big)}{\vdash \forall\tau\, \langle x' = f(x), t' = 1\rangle\big(t > \tau \lor \bigvee_{j\in N} \neg B_j(y_j)\big)}$$

The derivation combines ideas from Corollaries 4.6,4.8, and 4.10 to remove the ODE $y_i' = g_i(y_1, \dots, y_i)$ at each step. The corresponding disjunct $\neg B_i(y_i)$ (if present) is also removed from the succedent when $i \in N$. More precisely, at each step i, the derivation turns a succedent of the form (B.9) to the form (B.10) below which removes the variables y_i from the formula.

$$\langle y_1'{=}g_1(y_1), \dots, y_{i-1}'{=}g_{i-1}(y_1, \dots, y_{i-1}), y_i'{=}g_i(y_1, \dots, y_i), t'{=}1\rangle\big(t > \tau \lor \bigvee_{j\in N\cap\{1,\dots,i\}} \neg B_j(y_j)\big) \tag{B.9}$$

$$\langle y_1'{=}g_1(y_1), \dots, y_{i-1}'{=}g_{i-1}(y_1, \dots, y_{i-1}), t'{=}1\rangle\big(t > \tau \lor \bigvee_{j\in N\cap\{1,\dots,i-1\}} \neg B_j(y_j)\big) \tag{B.10}$$

The derivation proceeds with two cases depending on whether $i \in L$ or $i \in N$.

- For each $i \in L$ (similarly to Corollary 4.8), the ODE $y_i'{=}A_i(y_1, \dots, y_{i-1})y_i{+}b_i(y_1, \dots, y_{i-1})$ is affine for some matrix and vector terms A_i, b_i respectively with the indicated variable dependencies. The RHS of this affine ODE satisfies the inequality (B.8) with terms L_i, M_i as given in (B.8). Axiom DDG⟨·⟩ is used with those choices of L_i, M_i, which removes the ODEs for y_i in the resulting right premise. The resulting left premise is labeled ① and explained below. Note that the freshness conditions of axiom DDG⟨·⟩ are met because the postcondition of the succedent does not mention variables y_i for $i \in L$. Similarly, the indices from $j \in N \cap \{1, \dots, i\}$ are equal to those from $j \in N \cap \{1, \dots, i-1\}$ because $i \notin N$. The preceding ODEs are abbreviated $Y_{i-1} \equiv y_1'{=}g_1(y_1), \dots, y_{i-1}'{=}g_{i-1}(y_1, \dots, y_{i-1})$.

  $$\mathrm{DDG}\langle\cdot\rangle\,\dfrac{① \qquad \vdash \langle Y_{i-1}, t'{=}1\rangle\big(t > \tau \lor \bigvee_{i\in N\cap\{1,\dots,i-1\}} \neg B_i(y_i)\big)}{\vdash \langle Y_{i-1}, y_i'{=}g_i(y_1, \dots, y_i), t'{=}1\rangle\big(t > \tau \lor \bigvee_{i\in N\cap\{1,\dots,i\}} \neg B_i(y_i)\big)}$$

  From premise ①, the proof is completed with a dW and ℝ step using inequality (B.8).

  $$\mathrm{dW}\,\dfrac{\mathbb{R}\,\dfrac{*}{\vdash 2y_i \cdot (A_iy_i + b_i) \le L_i \left\|y_i\right\|_2^2 + M_i}}{\vdash [y_1' = g_1(y_1), \dots, y_i' = g_i(y_1, \dots, y_i), t' = 1]\, 2y_i \cdot (A_iy_i + b_i) \le L_i \left\|y_i\right\|_2^2 + M_i}$$

- For each $i \in N$ (similarly to Corollary 4.10), the boundedness assumption on y_i is first extracted from the succedent, with the abbreviation $R \equiv (t > \tau \lor \bigvee_{j\in N\cap\{1,\dots,i-1\}} \neg B_j(y_j))$. The bottommost succedent is similarly abbreviated using the propositional tautology $\big(t > \tau \lor \bigvee_{j\in N\cap\{1,\dots,i\}} \neg B_j(y_j)\big) \leftrightarrow R \lor \neg B_i(y_i)$. The preceding ODEs are abbreviated $Y_i \equiv y_1' = g_1(y_1), \dots, y_i' = g_i(y_1, \dots, y_i)$.

  $$\langle\cdot\rangle,\neg\mathrm{R}\,\dfrac{[\cdot]\wedge,\wedge\mathrm{L}\,\dfrac{\langle\cdot\rangle,\neg\mathrm{L}\,\dfrac{[Y_i, t' = 1]B_i(y_i) \vdash \langle Y_i, t' = 1\rangle R}{[Y_i, t' = 1]\neg R, [Y_i, t' = 1]B_i(y_i) \vdash \mathit{false}}}{[Y_i, t' = 1]\big(\neg R \wedge B_i(y_i)\big) \vdash \mathit{false}}}{\vdash \langle Y_i, t' = 1\rangle(R \lor \neg B_i(y_i))}$$

The formula $B_i(y_i)$ is assumed to characterize a bounded set with respect to the variables y_i. Thus, like Corollary 4.10, the cut of the formula $\exists D_i \, \forall y_i \, (B_i(y_i) \rightarrow \|y_i\|_2^2 \leq D_i)$ is proved by $\mathbb{R}$. The derivation continues by Skolemizing, abbreviating $S \equiv [Y_i, t' = 1] B_i(y_i)$. Axiom BDG$\langle\cdot\rangle$ is then used with $e(y_i) = D_i$, which removes the ODEs for y_i in the resulting right premise. The resulting left premise is labeled ② and explained below.

$$\text{cut, } \mathbb{R}, \exists\text{L} \dfrac{\text{BDG}\langle\cdot\rangle \dfrac{\text{②} \qquad \vdash \langle y_1'{=}g_1(y_1), \ldots, y_{i-1}'{=}g_{i-1}(y_1, \ldots, y_{i-1}), t'{=}1\rangle R}{S, \forall y_i \, (B_i(y_i) {\rightarrow} \|y_i\|_2^2 {\leq} D_i) \vdash \langle Y_i, t'{=}1\rangle R}}{S \vdash \langle Y_i, t'{=}1\rangle R}$$

The derivation continues from premise ② identically to Corollary 4.10, with a dC step to add the postcondition of the antecedent S to the domain constraint. The proof is completed with dW and $\forall$L, $\rightarrow$L. The universally quantified antecedent $\forall y_i \, \ldots$ is soundly kept across the use of dW since it does not mention any of the bound variables $y_1, \ldots, y_i, t$ of the ODE free.

$$\text{dC} \dfrac{\text{dW} \dfrac{\forall\text{L}, \rightarrow\text{L} \dfrac{*}{\forall y_i \, (B(y_i) \rightarrow \|y_i\|_2^2 \leq D_i), B(y_i) \vdash \|y_i\|_2^2 \leq D_i}}{\forall y_i \, (B(y_i) \rightarrow \|y_i\|_2^2 \leq D_i) \vdash [Y_i, t' = 1 \,\&\, B(y_i)] \; \|y_i\|_2^2 \leq D_i}}{S, \forall y_i \, (B(y_i) \rightarrow \|y_i\|_2^2 \leq D_i) \vdash [Y_i, t' = 1] \; \|y_i\|_2^2 \leq D_i}$$

Using the steps for $i = k, \ldots, 1$ (where either $i \in L$ or $i \in N$) successively removes the ODEs for $y_k, \ldots, y_i$ from the succedent. This is shown in the derivation below with abbreviations $Y_k \equiv y_1' = g_1(y_1), \ldots, y_{k-1}' = g_{k-1}(y_1, \ldots, y_{k-1}), y_k' = g_k(y_1, \ldots, y_k)$, $Y_{k-1} \equiv y_1' = g_1(y_1), \ldots, y_{k-1}' = g_{k-1}(y_1, \ldots, y_{k-1})$. The proof is completed using TEx.

$$\dfrac{\dfrac{\dfrac{\dfrac{\text{TEx} \dfrac{*}{\vdash \langle t' = 1\rangle t > \tau}}{\vdash \langle t' = 1\rangle \big(t > \tau \vee \bigvee_{j \in N \cap \emptyset} \neg B_j(y_j)\big)}}{\vdots}}{\vdash \langle Y_{k-1}, t' = 1\rangle \big(t > \tau \vee \bigvee_{j \in N \cap \{1, \ldots, k-1\}} \neg B_j(y_j)\big)}}{\vdash \langle Y_k, t' = 1\rangle \big(t > \tau \vee \bigvee_{j \in N} \neg B_j(y_j)\big)}$$

□

Proof of Proposition 4.13. The ODE $x' = f(x)$ is assumed to have a global solution that is syntactically representable by term $X(t)$ in the term language. Formally, this representability condition means that for any initial state ω, the mathematical solution $\varphi : [0, \infty) \rightarrow \mathbb{S}$ exists globally and in addition, for each time $\tau \in [0, \infty)$, the solution satisfies $\varphi(\tau) = \omega_t^\tau [\![X(t)]\!]$, where $\omega_t^\tau [\![X(t)]\!]$ is the value of term $X(t)$ in state ω with the value of time variable t set to τ. This implies that the following formula is valid because terms $x, t - t_0$ have value $\varphi(\tau)$ and τ respectively at time $\tau \in [0, \infty)$ along the ODE $x' = f(x), t' = 1$. The variables x_0, t_0 store the initial values of x, t respectively, which may be needed for the syntactic representation $X(t)$ of the solution. Additionally, the syntactic representation $X(t)$ may mention parameters $y \notin x$ that remain constant for the ODE $x' = f(x)$.

$$t = t_0 \wedge x = x_0 \rightarrow [x' = f(x), t' = 1] \, x = X(t - t_0) \tag{B.11}$$

Validity of formula (B.11) further implies that (B.11) is provable by dRI from Section 3.4.2 because the rule is complete for equational invariants (assuming that the resulting arithmetic formula is proved). The derivation of global existence for $x' = f(x)$ first Skolemizes with ∀R, then introduces fresh variables x_0, t_0 storing the initial values of x, t with cut, $\mathbb{R}$, ∃L. Axiom BDG⟨·⟩ is used with $e(t) = \|X(t-t_0)\|_2^2$ to remove the ODEs $x' = f(x)$. The resulting right premise is proved by axiom TEx, while the resulting left premise is abbreviated ① and proved below.

$$\text{∀R}\;\dfrac{\text{cut}, \mathbb{R}, \exists\text{L}\;\dfrac{\text{BDG}\langle\cdot\rangle\;\dfrac{\text{①} \qquad \text{TEx}\;\dfrac{*}{\vdash \langle t' = 1\rangle\, t > \tau}}{t = t_0 \land x = x_0 \vdash \langle x' = f(x), t' = 1\rangle\, t > \tau}}{\vdash \langle x' = f(x), t' = 1\rangle\, t > \tau}}{\vdash \forall \tau\, \langle x' = f(x), t' = 1\rangle\, t > \tau}$$

From ①, the derivation continues with a dC, dRI step using the provable formula (B.11). The premise after dW is proved by $\mathbb{R}$ after rewriting the succedent with the equality $x = X(t-t_0)$ and by reflexivity of $\leq$.

$$\text{dC, dRI}\;\dfrac{\text{dW}\;\dfrac{\mathbb{R}\;\dfrac{*}{x = X(t-t_0) \vdash \|x\|_2^2 \leq \|X(t-t_0)\|_2^2}}{\vdash [x' = f(x), t' = 1\,\&\, x = X(t-t_0)]\; \|x\|_2^2 \leq \|X(t-t_0)\|_2^2}}{t = t_0 \land x = x_0 \vdash [x' = f(x), t' = 1]\; \|x\|_2^2 \leq \|X(t-t_0)\|_2^2}$$

Note that, instead of assuming that $X(t)$ is a syntactically representable (global) solution for the ODE $x' = f(x)$, it also suffices for this derivation to assume that premise ① is provable, i.e., that the term $\|X(t-t_0)\|_2^2$ (with free variables t, x_0, t_0 and parameters y) is a provable upper bound on the squared norm of x along solutions of the ODE. □

B.2.2 Proofs for Liveness Without Domain Constraints

Proof of Corollary 4.16. The complete derivation of rule $\mathrm{dV}^{\Gamma}_{\succcurlyeq}$ using refinement axiom K⟨&⟩ and rule $\mathrm{dI}_{\succcurlyeq}$ is already given in the proof sketch for Corollary 4.16 so it is not repeated here.

The derivation of $\mathrm{dV}_{\succcurlyeq}$ (as a corollary of $\mathrm{dV}^{\Gamma}_{\succcurlyeq}$) starts by introducing fresh variables e_0, i representing the initial values of e and the multiplicative inverse of $\varepsilon()$ respectively using arithmetic cuts (cut, $\mathbb{R}$) and Skolemizing (∃L). It then uses dGt to introduce a fresh time variable to the system of differential equations:

$$\text{cut}, \mathbb{R}\;\dfrac{\exists\text{L}\;\dfrac{\text{dGt}\;\dfrac{\Gamma, \varepsilon() > 0, e = e_0, i\varepsilon() = 1, t = 0 \vdash \langle x' = f(x), t' = 1\rangle e \succcurlyeq 0}{\Gamma, \varepsilon() > 0, e = e_0, i\varepsilon() = 1 \vdash \langle x' = f(x)\rangle e \succcurlyeq 0}}{\Gamma, \varepsilon() > 0, \exists e_0\,(e = e_0), \exists i\,(i\varepsilon() = 1) \vdash \langle x' = f(x)\rangle e \succcurlyeq 0}}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle e \succcurlyeq 0}$$

Next, an initial liveness assumption $\langle x' = f(x), t' = 1\rangle e_0 + \varepsilon() t > 0$ is cut into the antecedents after which rule $\mathrm{dV}^{\Gamma}_{\succcurlyeq}$ is used to obtain the premise of $\mathrm{dV}_{\succcurlyeq}$. Intuitively, this initial liveness assumption says that the solution exists for sufficiently long, so that the term $e_0 + \varepsilon() t$ (which is proved to lower bound e) becomes positive for sufficiently large t. This cut is abbreviated

① and proved further below.

$$\textrm{cut}\,\dfrac{\textrm{dV}^{\Gamma}_{\succcurlyeq}\,\dfrac{\neg(e \succcurlyeq 0) \vdash \dot{e} \geq \varepsilon()}{\Gamma, e = e_0, t = 0, \langle x' = f(x), t' = 1\rangle\, e_0 + \varepsilon() t > 0 \vdash \langle x' = f(x), t' = 1\rangle e \succcurlyeq 0 \quad ①}}{\Gamma, \varepsilon() > 0, e = e_0, i\varepsilon() = 1, t = 0 \vdash \langle x' = f(x), t' = 1\rangle e \succcurlyeq 0}$$

From premise ①, a monotonicity step M⟨′⟩ equivalently rephrases the postcondition of the cut in real arithmetic. The arithmetic rephrasing works using the constant assumption $\varepsilon() > 0$ and the choice of i as the multiplicative inverse of $\varepsilon()$. Since the ODE $x' = f(x)$ is assumed to have provable global solutions, axiom GEx finishes the derivation by instantiating $\tau = -ie_0$, which is constant for the ODE.

$$\mathbb{R}, \textrm{M}\langle'\rangle\,\dfrac{\textrm{GEx}\,\dfrac{*}{\Gamma \vdash \langle x' = f(x), t' = 1\rangle\, t > -ie_0}}{\Gamma, \varepsilon() > 0, i\varepsilon() = 1 \vdash \langle x' = f(x), t' = 1\rangle\, e_0 + \varepsilon() t > 0} \qquad \square$$

Proof of Corollary 4.19. Rule $\textrm{dV}^{k}_{\succcurlyeq}$ can be derived in several ways. For example, because $\dot{e}^{(k)}$ is strictly positive, one can prove that the solution successively reaches states where $\dot{e}^{(k-1)}$ is strictly positive and remains positive thereafter, followed by reaching states where $\dot{e}^{(k-2)}$ is strictly positive (and remains positive thereafter), and so on. The following derivation shows how dC can be elegantly used for this argument. The idea is to extend the derivation of rule $\textrm{dV}_{\succcurlyeq}$ to higher Lie derivatives by (symbolically) integrating with respect to the time variable t using the following sequence of inequalities, where $\dot{e}_0^{(i)}$ is a symbolic constant that represents the initial value of the i-th Lie derivative of e along $x' = f(x)$ for $i = 0, 1, \ldots, k-1$:

$$\begin{aligned}
\dot{e}^{(k)} &\geq \varepsilon() \\
\dot{e}^{(k-1)} &\geq \dot{e}_0^{(k-1)} + \varepsilon() t \\
\dot{e}^{(k-2)} &\geq \dot{e}_0^{(k-2)} + \dot{e}_0^{(k-1)} t + \varepsilon() \frac{t^2}{2} \\
&\vdots \\
\dot{e}^{(1)} &\geq \dot{e}_0^{(1)} + \cdots + \dot{e}_0^{(k-1)} \frac{t^{k-2}}{(k-2)!} + \varepsilon() \frac{t^{k-1}}{(k-1)!} \\
e &\geq \underbrace{e_0 + \dot{e}_0^{(1)} t + \cdots + \dot{e}_0^{(k-1)} \frac{t^{k-1}}{(k-1)!} + \varepsilon() \frac{t^k}{k!}}_{p(t)}
\end{aligned} \tag{B.12}$$

The RHS of the final inequality in (B.12) is a polynomial in the time variable t, denoted $p(t)$, which is positive for sufficiently large values of t because its leading coefficient $\varepsilon()$ is strictly positive, i.e., with antecedent $\varepsilon() > 0$, formula $\exists t_1\, \forall t > t_1\, p(t) > 0$ is provable in real arithmetic.

The derivation of $\textrm{dV}^{k}_{\succcurlyeq}$ starts by introducing fresh ghost variables that remember the initial values of e and the (higher) Lie derivatives $\dot{e}^{(1)}, \ldots, \dot{e}^{(k-1)}$ using cut, $\mathbb{R}$, ∃L. The resulting antecedents are abbreviated with $\Gamma_0 \equiv \big(\Gamma, e = e_0, \ldots, \dot{e}^{(k-1)} = \dot{e}_0^{(k-1)}\big)$. It also uses dGt to introduce a fresh time variable t into the system. The arithmetic fact that $p(t)$ is eventually positive for all times $t > t_1$ is introduced with cut, $\mathbb{R}$, $\exists$L.

$$\text{cut, }\mathbb{R},\exists\text{L}\;\dfrac{\text{dGt}\;\dfrac{\text{cut, }\mathbb{R},\exists\text{L}\;\dfrac{\Gamma_0, t=0, \forall t > t_1\, p(t) > 0 \vdash \langle x' = f(x), t'=1\rangle e \succcurlyeq 0}{\Gamma_0, \varepsilon() > 0, t=0 \vdash \langle x'=f(x),t'=1\rangle e\succcurlyeq 0}}{\Gamma, \varepsilon() > 0, e = e_0, \ldots, \overset{\bullet}{e}^{(k-1)} = \overset{\bullet}{e}_0^{(k-1)} \vdash \langle x'=f(x)\rangle e \succcurlyeq 0}}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle e \succcurlyeq 0}$$

Next, an initial liveness assumption, $\langle x' = f(x), t' = 1\rangle p(t) > 0$, is cut into the assumptions. Like the derivation of rule dV$_\succcurlyeq$, this initial liveness assumption says that the solution exists for sufficiently long so that the term $p(t)$ from (B.12), which is proved to lower bound e, becomes positive for sufficiently large t. The cut premise is abbreviated ① and further proved below. The derivation continues from the remaining (unabbreviated) premise by refinement axiom K$\langle\&\rangle$, with $G \equiv p(t) > 0$:

$$\text{cut}\;\dfrac{\text{K}\langle\&\rangle\;\dfrac{\Gamma_0,t=0\vdash[x'=f(x),t'=1\,\&\,\neg(e\succcurlyeq 0)]p(t)\le 0}{\Gamma_0, t=0, \langle x'=f(x),t'=1\rangle p(t)>0 \vdash \langle x'=f(x),t'=1\rangle e\succcurlyeq 0} \quad \text{①}}{\Gamma_0,t=0,\forall t>t_1\,p(t)>0\vdash\langle x'=f(x),t'=1\rangle e\succcurlyeq 0}$$

From the resulting open premise after K$\langle\&\rangle$, monotonicity M$[']$ strengthens the postcondition to $e \geq p(t)$ using the domain constraint $\neg(e \succcurlyeq 0)$ and the provable real arithmetic fact $\neg(e \succcurlyeq 0) \wedge e \geq p(t) \rightarrow p(t) \leq 0$. Notice that the resulting postcondition $e \geq p(t)$ is the final inequality from the sequence of inequalities (B.12):

$$\text{M}[']\;\dfrac{\Gamma_0,t=0\vdash[x'=f(x),t'=1\,\&\,\neg(e\succcurlyeq 0)]e\geq p(t)}{\Gamma_0,t=0\vdash[x'=f(x),t'=1\,\&\,\neg(e\succcurlyeq 0)]p(t)\le 0}$$

The derivation continues by using dC to sequentially cut in the inequality bounds outlined in (B.12). The first differential cut dC step adds $\overset{\bullet}{e}^{(k-1)} \geq \overset{\bullet}{e}_0^{(k-1)} + \varepsilon()t$ to the domain constraint. The proof of this differential cut yields the premise of dV$^k_\succcurlyeq$ after a dI$_\succcurlyeq$ step, see the derivation labeled ⍟ immediately below.

$$\text{dC}\;\dfrac{\Gamma_0,t=0\vdash[x'=f(x),t'=1\,\&\,\neg(e\succcurlyeq 0)\wedge \overset{\bullet}{e}^{(k-1)} \geq \overset{\bullet}{e}_0^{(k-1)} + \varepsilon()t]e\geq p(t) \quad \text{⍟}}{\Gamma_0,t=0\vdash[x'=f(x),t'=1\,\&\,\neg(e\succcurlyeq 0)]e\geq p(t)}$$

From ⍟, the resulting open premise is the premise of rule dV$^k_\succcurlyeq$:

$$\text{dI}_\succcurlyeq\;\dfrac{\neg(e\succcurlyeq 0)\vdash \overset{\bullet}{e}^{(k)} \geq \varepsilon()}{\Gamma_0,t=0\vdash[x'=f(x),t'=1\,\&\,\neg(e\succcurlyeq 0)]\overset{\bullet}{e}^{(k-1)} \geq \overset{\bullet}{e}_0^{(k-1)} + \varepsilon()t}$$

Subsequent dC, dI$_\succcurlyeq$ steps progressively add the inequality bounds from (B.12) to the domain constraint until the last step where the postcondition is proved invariant with dI$_\succcurlyeq$:

$$\text{dC, dI}_\succcurlyeq\;\dfrac{\text{dC, dI}_\succcurlyeq\;\dfrac{\text{dC, dI}_\succcurlyeq\;\dfrac{\text{dI}_\succcurlyeq\;\dfrac{*}{\Gamma_0,t=0\vdash[x'=f(x),t'=1\,\&\,\cdots\wedge \overset{\bullet}{e}^{(1)} \geq \overset{\bullet}{e}_0^{(1)} + \cdots + \varepsilon()\frac{t^{k-1}}{(k-1)!}]e\geq p(t)}}{\vdots}}{\Gamma_0,t=0\vdash[x'=f(x),t'=1\,\&\,\cdots\wedge \overset{\bullet}{e}^{(k-2)} \geq \overset{\bullet}{e}_0^{(k-2)} + \overset{\bullet}{e}_0^{(k-1)}t + \varepsilon()\frac{t^2}{2}]e\geq p(t)}}{\Gamma_0,t=0\vdash[x'=f(x),t'=1\,\&\,\neg(e\succcurlyeq 0)\wedge \overset{\bullet}{e}^{(k-1)} \geq \overset{\bullet}{e}_0^{(k-1)} + \varepsilon()t]e\geq p(t)}$$

From premise ①, a monotonicity step M⟨′⟩ rephrases the postcondition of the cut using the (constant) assumption $\forall t > t_1\, p(t) > 0$. Axiom GEx, with instance $\tau = t_1$, finishes the derivation because the ODE $x' = f(x)$ is assumed to have provable global solutions.

$$\text{M}\langle'\rangle\,\dfrac{\text{GEx}\,\dfrac{*}{\Gamma \vdash \langle x' = f(x), t' = 1\rangle t > t_1}}{\Gamma, \forall t > t_1\, p(t) > 0 \vdash \langle x' = f(x), t' = 1\rangle p(t) > 0}$$ □

Proof of Corollary 4.20. The derivation of rule SP begins by using axiom K⟨&⟩ with $G \equiv \neg S$. The resulting left premise is the left premise of rule SP, which is the staging property of the formula S expressing that solutions of the ODE $x' = f(x)$ can only leave S by entering P:

$$\text{K}\langle\&\rangle\,\dfrac{\Gamma \vdash [x' = f(x)\,\&\,\neg P]S \qquad \Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle\neg S}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle P}$$

The derivation continues on the right premise, similarly to $\text{dV}_{\succcurlyeq}$, by introducing fresh variables e_0, i representing the initial value of e and the multiplicative inverse of $\varepsilon()$ respectively using arithmetic cuts (cut, $\mathbb{R}$). It then uses dGt to introduce a fresh time variable:

$$\text{cut}, \mathbb{R}\,\dfrac{\exists\text{L}\,\dfrac{\text{dGt}\,\dfrac{\Gamma, \varepsilon() > 0, e = e_0, i\varepsilon() = 1, t = 0 \vdash \langle x' = f(x), t' = 1\rangle\neg S}{\Gamma, \varepsilon() > 0, e = e_0, i\varepsilon() = 1 \vdash \langle x' = f(x)\rangle\neg S}}{\Gamma, \varepsilon() > 0, \exists e_0\,(e = e_0), \exists i\,(i\varepsilon() = 1) \vdash \langle x' = f(x)\rangle\neg S}}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle\neg S}$$

The next cut introduces an initial liveness assumption where the cut premise is abbreviated ① and proved identically to the correspondingly abbreviated premise from the derivation of $\text{dV}_{\succcurlyeq}$ using axiom GEx because the ODE $x' = f(x)$ is assumed have provable global solutions.

$$\text{cut}\,\dfrac{\Gamma, e = e_0, t = 0, \langle x' = f(x), t' = 1\rangle\, e_0 + \varepsilon()t > 0 \vdash \langle x' = f(x), t' = 1\rangle\neg S \qquad \text{①}}{\Gamma, \varepsilon() > 0, e = e_0, i > 0, i\varepsilon() = 1, t = 0 \vdash \langle x' = f(x), t' = 1\rangle\neg S}$$

From the remaining open premise, axiom K⟨&⟩ is used with $G \equiv e_0 + \varepsilon()t > 0$:

$$\text{K}\langle\&\rangle\,\dfrac{\Gamma, e = e_0, t = 0 \vdash [x' = f(x), t' = 1\,\&\,S]\, e_0 + \varepsilon()t \leq 0}{\Gamma, e = e_0, t = 0, \langle x' = f(x), t' = 1\rangle\, e_0 + \varepsilon()t > 0 \vdash \langle x' = f(x), t' = 1\rangle\neg S}$$

A monotonicity step M[′] simplifies the postcondition using domain constraint S, yielding the left conjunct of the right premise of rule SP. The right premise after monotonicity is abbreviated ② and continued below.

$$\text{M}[']\,\dfrac{\mathbb{R}\,\dfrac{S \vdash e \leq 0}{S, e \geq e_0 + \varepsilon()t \vdash e_0 + \varepsilon()t \leq 0} \qquad \text{②}}{\Gamma, e = e_0, t = 0 \vdash [x' = f(x), t' = 1\,\&\,S]\, e_0 + \varepsilon()t \leq 0}$$

From ②, rule $\text{dI}_{\succcurlyeq}$ yields the right conjunct of the right premise of rule SP.

$$\text{dI}_{\succcurlyeq}\,\dfrac{S \vdash \dot{e} \geq \varepsilon()}{\Gamma, e = e_0, t = 0 \vdash [x' = f(x), t' = 1\,\&\,S]\, e \geq e_0 + \varepsilon()t}$$ □

Proof of Corollary 4.21. Rule SP_b is derived first since rule SP_c follows from SP_b as a corollary. Both proof rules make use of the fact that continuous functions on compact domains attain their extrema (see Appendix B.1.3). The derivation of SP_b is essentially similar to SP except replacing the use of the global existence axiom GEx with the bounded existence axiom BEx. It starts by using axiom K⟨&⟩ with $G \equiv \neg S$, yielding the left premise of SP$_b$:

$$\text{K}\langle\&\rangle\,\dfrac{\Gamma \vdash [x' = f(x) \,\&\, \neg P]S \qquad \Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle \neg S}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle P}$$

Continuing on the resulting right from K⟨&⟩ (similarly to SP), the derivation introduces fresh variables e_0, i representing the initial value of e and the multiplicative inverse of $\varepsilon()$ respectively using arithmetic cuts and Skolemizing (cut, $\mathbb{R}$, ∃L). Rule dGt is also used to introduce a fresh time variable t with $t = 0$ initially.

$$\text{cut, }\mathbb{R}\text{, }\exists\text{L, dGt}\,\dfrac{\Gamma, \varepsilon() > 0, e = e_0, i\varepsilon() = 1, t = 0 \vdash \langle x' = f(x), t' = 1\rangle \neg S}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle \neg S}$$

The set characterized by formula S is bounded so its closure is compact (with respect to variables x). On this compact closure, the continuous semantics of term e attains its maximum value, which implies that the value of e is bounded above in S and cannot increase unboundedly while staying in S. That is, the formula $\exists e_1\, R(e_1)$ where $R(e_1) \equiv \forall x\, (S(x) \to e \leq e_1)$ is valid in first-order real arithmetic and thus provable by $\mathbb{R}$ (Appendix B.1.3). This formula is added to the assumptions with a cut, and the existential quantifier is Skolemized with ∃L. The resulting symbolic constant e_1 represents the upper bound of e on S. Note that $R(e_1)$ is constant for the ODE $x' = f(x), t' = 1$ because it does not mention any of the variables x (nor t) free:

$$\text{cut, }\mathbb{R}\,\dfrac{\exists\text{L}\,\dfrac{\Gamma, \varepsilon() > 0, e = e_0, i\varepsilon() = 1, t = 0, R(e_1) \vdash \langle x' = f(x), t' = 1\rangle \neg S}{\Gamma, \varepsilon() > 0, e = e_0, i\varepsilon() = 1, t = 0, \exists e_1\, R(e_1) \vdash \langle x' = f(x), t' = 1\rangle \neg S}}{\Gamma, \varepsilon() > 0, e = e_0, i\varepsilon() = 1, t = 0 \vdash \langle x' = f(x), t' = 1\rangle \neg S}$$

Next, a cut introduces an initial liveness assumption saying that *either* the solution exists for sufficient time for the bound $e_0 + \varepsilon()t > e_1$ to be satisfied (at sufficiently large t) *or* the solution leaves S. This assumption is abbreviated $T \equiv \langle x' = f(x), t' = 1\rangle (e_0 + \varepsilon()t > e_1 \vee \neg S)$. The main difference from SP is that the postcondition of assumption T adds a disjunction for the possibility of leaving S (which characterizes a bounded set). This cut premise is abbreviated ① and proved further below.

$$\text{cut}\,\dfrac{\Gamma, e = e_0, t = 0, R(e_1), T \vdash \langle x' = f(x), t' = 1\rangle \neg S \qquad \text{①}}{\Gamma, \varepsilon() > 0, e = e_0, i\varepsilon() = 1, t = 0, R(e_1) \vdash \langle x' = f(x), t' = 1\rangle \neg S}$$

Continuing from the open premise on the left, axiom K⟨&⟩ is used with $G \equiv e_0 + \varepsilon()t > e_1 \vee \neg S$:

$$\text{K}\langle\&\rangle\,\dfrac{\Gamma, e = e_0, t = 0, R(e_1) \vdash [x' = f(x), t' = 1 \,\&\, S](e_0 + \varepsilon()t \leq e_1 \wedge S)}{\Gamma, e = e_0, t = 0, R(e_1), T \vdash \langle x' = f(x), t' = 1\rangle \neg S}$$

The postcondition of the resulting box modality is simplified to $e \geq e_0 + \varepsilon()t$ with a M['] monotonicity step. This step crucially uses the assumption $R(e_1)$ which is constant for the ODE. A dI$_\succcurlyeq$ step yields the remaining premise of SP$_b$ on the right, see the derivation labeled ⊛ immediately below:

$$\text{M}[']\,\dfrac{\mathbb{R}\,\dfrac{\mathbb{R}\,\dfrac{*}{S, R(e_1) \vdash e \leq e_1}}{S, R(e_1), e \geq e_0 + \varepsilon()t \vdash e_0 + \varepsilon()t \leq e_1 \wedge S} \qquad \circledast}{\Gamma, e = e_0, t = 0, R(e_1) \vdash [x' = f(x), t' = 1 \,\&\, S](e_0 + \varepsilon()t \leq e_1 \wedge S)}$$

From ⍟:

$$\text{dI}_{\succcurlyeq}\,\dfrac{S \vdash \dot{e} \geq \varepsilon()}{\Gamma, e = e_0, t = 0 \vdash [x' = f(x), t' = 1 \,\&\, S]e \geq e_0 + \varepsilon()t}$$

From premise ①, a monotonicity step M⟨′⟩ equivalently rephrases the postcondition of the cut of formula T. Axiom BEx finishes the proof because formula $S(x)$ is assumed to be bounded over variables x.

$$\mathbb{R}, \text{M}\langle'\rangle\,\dfrac{\text{BEx}\,\dfrac{*}{\vdash \langle x' = f(x), t' = 1\rangle (t > i(e_1 - e_0) \lor \neg S)}}{\varepsilon() > 0, i\varepsilon() = 1 \vdash T}$$

To derive rule SP$_c$ from SP$_b$, the compactness of the set characterized by $S(x)$ implies that formula $\exists\varepsilon{>}0\, A(\varepsilon)$ where $A(\varepsilon) \equiv \forall x\, (S(x) \rightarrow \dot{e} \geq \varepsilon)$ and formula $B \equiv \forall x\, (S(x) \rightarrow \dot{e} > 0)$ are provably equivalent in first-order real arithmetic (Appendix B.1.3). This provable real arithmetic equivalence follows from the fact that the continuous semantics of term $\dot{e}$ is bounded below by its minimum value on the compact set characterized by $S(x)$ and this minimum value is strictly positive. The derivation of SP$_c$ threads these two formulas through the use of rule SP$_b$. After Skolemizing $\exists\varepsilon{>}0\, A(\varepsilon)$ with $\exists$L, the resulting formula $A(\varepsilon)$ is constant for the ODE $x' = f(x)$ so it is kept as a constant assumption across the use of SP$_b$, leaving only the two premises of rule SP$_c$:

$$\text{cut}\,\dfrac{\exists\text{L}\,\dfrac{\text{SP}_b\,\dfrac{\Gamma \vdash [x' = f(x) \,\&\, \neg P]S \qquad \mathbb{R}\,\dfrac{*}{S, A(\varepsilon) \vdash \dot{e} \geq \varepsilon}}{\Gamma, \varepsilon > 0, A(\varepsilon) \vdash \langle x' = f(x)\rangle P}}{\Gamma, \exists\varepsilon{>}0\, A(\varepsilon) \vdash \langle x' = f(x)\rangle P} \qquad \mathbb{R}\,\dfrac{\forall\text{R}, \rightarrow\text{R}\,\dfrac{S \vdash \dot{e} > 0}{\vdash B}}{\vdash \exists\varepsilon{>}0\, A(\varepsilon)}}{\Gamma \vdash \langle x' = f(x)\rangle P}$$

□

B.2.3 Proofs for Liveness With Domain Constraints

Proof of Corollary 4.25. The derivation uses axiom COR with $R \equiv \mathit{true}$ and noting that $e \geq 0$ (resp. $e > 0$) characterizes a topologically closed (resp. open) set so the appropriate topological requirements of COR are satisfied. The resulting left premise is the left premise of dV$_{\succcurlyeq}$&:

$$\text{COR}\,\dfrac{\Gamma \vdash [x' = f(x) \,\&\, \neg(e \succcurlyeq 0)]Q \qquad \Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle e \succcurlyeq 0}{\Gamma, \varepsilon() > 0, \neg(e \succcurlyeq 0) \vdash \langle x' = f(x) \,\&\, Q\rangle e \succcurlyeq 0}$$

The proof continues from the resulting right premise (after COR) identically to the derivation of dV$_{\succcurlyeq}$ until the step where dV$^{\Gamma}_{\succcurlyeq}$ is used. The steps are repeated briefly here.

$$\text{cut}, \mathbb{R}, \exists\text{L}\,\dfrac{\text{dGt}\,\dfrac{\text{cut, GEx}\,\dfrac{\Gamma, e = e_0, t = 0, \langle x' = f(x), t' = 1\rangle\, e_0 + \varepsilon()t > 0 \vdash \langle x' = f(x), t' = 1\rangle e \succcurlyeq 0}{\Gamma, \varepsilon() > 0, e = e_0, i\varepsilon() = 1, t = 0 \vdash \langle x' = f(x), t' = 1\rangle e \succcurlyeq 0}}{\Gamma, \varepsilon() > 0, e = e_0, i\varepsilon() = 1 \vdash \langle x' = f(x)\rangle e \succcurlyeq 0}}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\rangle e \succcurlyeq 0}$$

Like the derivation of dV$^{\Gamma}_{\succcurlyeq}$, axiom K⟨&⟩ is used with $G \equiv e_0() + \varepsilon()t > 0$. The key difference is an additional dC step, which adds Q to the domain constraint.[4] The proof of this differential

[4] Notably, the differential cuts proof support from Section 4.6.2 can add such a cut automatically.

cut uses the left premise of $\text{dV}_{\succcurlyeq}\&$, it is labeled ① and shown below.

$$\text{K}\langle\&\rangle\dfrac{\text{dC}\dfrac{\Gamma, e = e_0(), t = 0 \vdash [x' = f(x), t' = 1 \,\&\, \neg(e \succcurlyeq 0) \land Q]\, e_0() + \varepsilon()t \le 0 \qquad \text{①}}{\Gamma, e = e_0(), t = 0 \vdash [x' = f(x), t' = 1 \,\&\, \neg(e \succcurlyeq 0)]\, e_0() + \varepsilon()t \le 0}}{\Gamma, e = e_0, t = 0, \langle x' = f(x), t' = 1\rangle e_0 + \varepsilon()t > 0 \vdash \langle x' = f(x), t' = 1\rangle e \succcurlyeq 0}$$

The derivation from the resulting left premise (after the cut) continues similarly to $\text{dV}_{\succcurlyeq}^{\Gamma}$ using a monotonicity step M[′] to rephrase the postcondition, followed by $\text{dI}_{\succcurlyeq}$ which results in the right premise of $\text{dV}_{\succcurlyeq}\&$:

$$\text{M}[']\dfrac{\text{dI}_{\succcurlyeq}\dfrac{\neg(e \succcurlyeq 0), Q \vdash \dot{e} \ge \varepsilon()}{\Gamma, e = e_0(), t = 0 \vdash [x' = f(x), t' = 1 \,\&\, \neg(e \succcurlyeq 0) \land Q]\, e \ge e_0() + \varepsilon()t}}{\Gamma, e = e_0(), t = 0 \vdash [x' = f(x), t' = 1 \,\&\, \neg(e \succcurlyeq 0) \land Q]\, e_0() + \varepsilon()t \le 0}$$

The derivation from ① removes the time variable t using the inverse direction of axiom DG [142, 144]. Just as rule dGt (which is derived from DG) introduces a *fresh* time variable t for the sake of proof, inverse DG simply removes the variable t since it is irrelevant for the proof of the differential cut.

$$\text{DG}\dfrac{\Gamma \vdash [x' = f(x) \,\&\, \neg(e \succcurlyeq 0)]Q}{\Gamma, e = e_0(), t = 0 \vdash [x' = f(x), t' = 1 \,\&\, \neg(e \succcurlyeq 0)]Q}$$

□

Proof of Corollary 4.26. The derivations of rules $\text{dV}_{=}\&$, $\text{dV}_{=}^{M}\&$ are similar to the derivations of rules $\text{dV}_{=}$, $\text{dV}_{=}^{M}$ respectively. Rule $\text{dV}_{=}^{M}\&$ is derived from $\text{dV}_{=}\&$ by monotonicity:

$$\text{M}\langle{}'\rangle\dfrac{Q, e = 0 \vdash P \qquad \text{dV}_{=}\&\dfrac{\Gamma \vdash [x' = f(x) \,\&\, e < 0]Q \quad e < 0, Q \vdash \dot{e} \ge \varepsilon()}{\Gamma, \varepsilon() > 0, e \le 0, Q \vdash \langle x' = f(x) \,\&\, Q\rangle e = 0}}{\Gamma, \varepsilon() > 0, e \le 0, Q \vdash \langle x' = f(x) \,\&\, Q\rangle P}$$

The derivation of rule $\text{dV}_{=}\&$ starts by using axiom K⟨&⟩ with $G \equiv e \ge 0$. The resulting box modality (right) premise is abbreviated ① and proved below. On the resulting left premise, a DX step adds the negated postcondition $e < 0$ as an assumption to the antecedents since the domain constraint Q is true initially. Following that, rule $\text{dV}_{\succcurlyeq}\&$ is used (with $\succcurlyeq$ being $\ge$, since Q characterizes a closed set). This yields the two premises of $\text{dV}_{=}\&$:

$$\text{K}\langle\&\rangle\dfrac{\text{DX}\dfrac{\text{dV}_{\succcurlyeq}\&\dfrac{\Gamma \vdash [x' = f(x) \,\&\, e < 0]Q \qquad e < 0, Q \vdash \dot{e} \ge \varepsilon()}{\Gamma, \varepsilon() > 0, e < 0 \vdash \langle x' = f(x) \,\&\, Q\rangle e \ge 0}}{\Gamma, \varepsilon() > 0, Q \vdash \langle x' = f(x) \,\&\, Q\rangle e \ge 0 \qquad \text{①}}}{\Gamma, \varepsilon() > 0, e \le 0, Q \vdash \langle x' = f(x) \,\&\, Q\rangle e = 0}$$

From premise ①, the derivation is completed similarly to $\text{dV}_{=}$ using DX and Barr:

$$\text{DX}\dfrac{\text{Barr}\dfrac{\mathbb{R}\dfrac{*}{e \ne 0, e = 0 \vdash \dot{e} < 0}}{e < 0 \vdash [x' = f(x) \,\&\, Q \land e \ne 0]e < 0}}{e \le 0 \vdash [x' = f(x) \,\&\, Q \land e \ne 0]e < 0}$$

□

Proof of Corollary 4.27. Rule $\text{SLyap}^{M}\&$ is derived from SLyap& by a DR⟨·⟩ monotonicity step followed by dW on its resulting left premise and SLyap& on its resulting right premise:

$$\text{DR}\langle\cdot\rangle\dfrac{\text{dW}\dfrac{e > 0 \vdash Q}{\Gamma, e > 0 \vdash [x' = f(x) \,\&\, e > 0]Q} \qquad \text{SLyap}\&\dfrac{e \ge 0 \vdash K \qquad \neg P, K \vdash \dot{e} > 0}{\Gamma, e > 0 \vdash \langle x' = f(x) \,\&\, e > 0\rangle P}}{\Gamma, e > 0 \vdash \langle x' = f(x) \,\&\, Q\rangle P}$$

The derivation of rule SLyap& starts by adding assumption $\neg P$ to the antecedents, because if both $e > 0$ (which is already in the antecedents) and P were true initially, then the liveness succedent is trivially true by DX. Next, axiom COR is used with $R \equiv \mathit{true}$, its topological restrictions are met since both formulas P and $e > 0$ characterize open sets. From the resulting right premise, rule SLyap yields the corresponding two premises of SLyap& because formula K (resp. P) characterizes a compact set (resp. open set):

$$\text{DX}\,\dfrac{\text{COR}\,\dfrac{\Gamma, e > 0 \vdash [x' = f(x) \,\&\, \neg P] e > 0 \qquad \text{SLyap}\,\dfrac{e \geq 0 \vdash K \qquad \neg P, K \vdash \dot{e} > 0}{\Gamma, e > 0 \vdash \langle x' = f(x) \rangle P}}{\Gamma, e > 0, \neg P \vdash \langle x' = f(x) \,\&\, e > 0 \rangle P}}{\Gamma, e > 0 \vdash \langle x' = f(x) \,\&\, e > 0 \rangle P}$$

From the leftmost open premise after COR, rule Barr is used and the resulting $e = 0$ assumption is turned into K using the left premise of SLyap&. The resulting open premises are the premises of SLyap&:

$$\text{Barr}\,\dfrac{\text{cut}\,\dfrac{\neg P, K \vdash \dot{e} > 0 \qquad \mathbb{R}\,\dfrac{e \geq 0 \vdash K}{e = 0 \vdash K}}{\neg P, e = 0 \vdash \dot{e} > 0}}{\Gamma, e > 0 \vdash [x' = f(x) \,\&\, \neg P] e > 0}$$

□

Proof of Corollary 4.28. The derivation starts with a SAR refinement step. On the resulting left premise, an M['] monotonicity step yields the left premise and first (leftmost) conjunct of the right premise of rule SP&. The derivation continues from the resulting right premise below.

$$\text{SAR}\,\dfrac{\text{M[']}\,\dfrac{\Gamma \vdash [x' = f(x) \,\&\, \neg(P \land Q)] S \qquad S \vdash Q}{\Gamma \vdash [x' = f(x) \,\&\, \neg(P \land Q)] Q} \qquad \Gamma \vdash \langle x' = f(x) \rangle P}{\Gamma \vdash \langle x' = f(x) \,\&\, Q \rangle P}$$

From the resulting right premise after using axiom SAR, rule SP yields the remaining two premises of SP&. The dW, DMP monotonicity step uses the propositional tautology $\neg P \rightarrow \neg(P \land Q)$ to weaken the domain constraint so that it matches the left premise of rule SP&.

$$\text{SP}\,\dfrac{\text{dW, DMP}\,\dfrac{\Gamma \vdash [x' = f(x) \,\&\, \neg(P \land Q)] S}{\Gamma \vdash [x' = f(x) \,\&\, \neg P] S} \qquad S \vdash e \leq 0 \land \dot{e} \geq \varepsilon()}{\Gamma \vdash \langle x' = f(x) \rangle P}$$

□

Proof of Corollary 4.30. The chimeric proof rule $\text{SP}_c^k\&$ amalgamates refinement ideas behind the derived rules SP&, $\text{dV}_{\succcurlyeq}^k$, SP_c. It is therefore unsurprising that the derivation of $\text{SP}_c^k\&$ uses various steps from the derivations of those rules. The derivation of $\text{SP}_c^k\&$ starts similarly to SP& (following Corollary 4.28) using axiom SAR:

$$\text{SAR}\,\dfrac{\Gamma \vdash [x' = f(x) \,\&\, \neg(P \land Q)] Q \qquad \Gamma \vdash \langle x' = f(x) \rangle P}{\Gamma \vdash \langle x' = f(x) \,\&\, Q \rangle P}$$

From the left premise after SAR, a monotonicity step turns the postcondition into S, yielding the left premise and first conjunct of the right premise of $\text{SP}_c^k\&$.

$$\text{M[']}\,\dfrac{\Gamma \vdash [x' = f(x) \,\&\, \neg(P \land Q)] S \qquad S \vdash Q}{\Gamma \vdash [x' = f(x) \,\&\, \neg(P \land Q)] Q}$$

From the right premise after SAR, the derivation continues using K$\langle\&\rangle$ with $G \equiv \neg S$, followed by dW, DMP. The resulting left premise is (again) the left premise of SP$_c^k$&, while the resulting right premise is abbreviated ① and continued below:

$$\text{K}\langle\&\rangle\,\cfrac{\text{dW, DMP}\,\cfrac{\Gamma \vdash [x' = f(x)\,\&\,\neg(P \wedge Q)]S}{\Gamma \vdash [x' = f(x)\,\&\,\neg P]S} \qquad \text{①}}{\Gamma \vdash \langle x' = f(x)\rangle P}$$

The derivation continues from ① by interleaving proof ideas from Corollaries 4.19 and 4.21. First, compactness of the set characterized by $S(x)$ implies that the formula $\exists \varepsilon{>}0\, A(\varepsilon)$ where $A(\varepsilon) \equiv \forall x\,(S(x) \rightarrow \dot{e}^{(k)} \geq \varepsilon)$ and the formula $B \equiv \forall x\,(S(x) \rightarrow \dot{e}^{(k)} > 0)$ are provably equivalent in first-order real arithmetic (Appendix B.1.3). These facts are added to the assumptions similarly to the derivation of SP$_c$. The resulting right open premise is the right conjunct of the right premise of SP$_c^k$&:

$$\text{cut}\,\cfrac{\exists\text{L}\,\cfrac{\Gamma, \varepsilon > 0, A(\varepsilon) \vdash \langle x' = f(x)\rangle \neg S}{\Gamma, \exists \varepsilon{>}0\, A(\varepsilon) \vdash \langle x' = f(x)\rangle \neg S} \qquad \mathbb{R}\,\cfrac{\forall\text{R}, \rightarrow\text{R}\,\cfrac{S \vdash \dot{e}^{(k)} > 0}{\vdash B}}{\vdash \exists \varepsilon{>}0\, A(\varepsilon)}}{\Gamma \vdash \langle x' = f(x)\rangle \neg S}$$

From the left premise, recall the derivation from Corollary 4.19 which introduces fresh variables for the initial values of the Lie derivatives with cut, $\mathbb{R}$, ∃L. The derivation continues similarly here, with the resulting antecedents abbreviated $\Gamma_0 \equiv \big(\Gamma, e = e_0, \dots, \dot{e}^{(k-1)} = \dot{e}_0^{(k-1)}\big)$. Rule dGt is also used to add time variable t to the system of equations with initial value $t = 0$.

$$\text{cut}, \mathbb{R}, \exists\text{L}\,\cfrac{\text{dGt}\,\cfrac{\Gamma_0, \varepsilon > 0, A(\varepsilon), t = 0 \vdash \langle x' = f(x), t' = 1\rangle \neg S}{\Gamma_0, \varepsilon > 0, A(\varepsilon) \vdash \langle x' = f(x)\rangle \neg S}}{\Gamma, \varepsilon > 0, A(\varepsilon) \vdash \langle x' = f(x)\rangle \neg S}$$

Recall from Corollary 4.21 that the formula $R(e_1) \equiv \forall x\,(S(x) \rightarrow e \leq e_1)$ can be added to the assumptions using cut, $\mathbb{R}$, ∃L, for some fresh variable e_1 symbolically representing the maximum value of e on the compact set characterized by S:

$$\text{cut}, \mathbb{R}, \exists\text{L}\,\cfrac{\Gamma_0, \varepsilon > 0, A(\varepsilon), t = 0, R(e_1) \vdash \langle x' = f(x), t' = 1\rangle \neg S}{\Gamma_0, \varepsilon > 0, A(\varepsilon), t = 0 \vdash \langle x' = f(x), t' = 1\rangle \neg S}$$

One last arithmetic cut is needed to set up the sequence of differential cuts (B.12). Recall the polynomial $p(t)$ from (B.12) is eventually positive for sufficiently large values of t because its leading coefficient is strictly positive. The same applies to the polynomial $p(t) - e_1$ so cut, $\mathbb{R}$ (and Skolemizing with ∃L) adds the formula $\forall t > t_1\,(p(t) - e_1 > 0)$ to the assumptions:

$$\text{cut}, \mathbb{R}, \exists\text{L}\,\cfrac{\Gamma_0, \varepsilon > 0, A(\varepsilon), t = 0, R(e_1), \forall t > t_1\, p(t) - e_1 > 0 \vdash \langle x' = f(x), t' = 1\rangle \neg S}{\Gamma_0, \varepsilon > 0, A(\varepsilon), t = 0, R(e_1) \vdash \langle x' = f(x), t' = 1\rangle \neg S}$$

Once all the arithmetic cuts are in place, an additional cut introduces a (bounded) sufficient duration assumption $\langle x' = f(x), t' = 1\rangle (p(t) - e_1 > 0 \vee \neg S)$ (antecedents temporarily abbreviated with $\dots$ for brevity). The cut premise, abbreviated ①, is proved further below:

$$\text{cut}\,\cfrac{\Gamma_0, \dots, \langle x' = f(x), t' = 1\rangle (p(t) - e_1 > 0 \vee \neg S) \vdash \langle x' = f(x), t' = 1\rangle \neg S \quad \text{①}}{\Gamma_0, \varepsilon > 0, A(\varepsilon), t = 0, R(e_1), \forall t > t_1\,(p(t) - e_1 > 0) \vdash \langle x' = f(x), t' = 1\rangle \neg S}$$

From the open premise on the left, axiom K$\langle\&\rangle$ is used with $G \equiv p(t) - e_1 > 0 \lor \neg S$:

$$\text{K}\langle\&\rangle\,\dfrac{\Gamma_0, \varepsilon > 0, A(\varepsilon), t{=}0, R(e_1) \vdash [x'{=}f(x), t'{=}1\,\&\,S](p(t) - e_1 \leq 0 \land S)}{\Gamma_0, \dots, \langle x'{=}f(x), t'{=}1\rangle(p(t) - e_1 > 0 \lor \neg S) \vdash \langle x'{=}f(x), t'{=}1\rangle \neg S}$$

Next, a monotonicity step M$[']$ simplifies the postcondition using the (constant) assumption $R(e_1)$ and the domain constraint S:

$$\text{M}[']\,\dfrac{\Gamma_0, t = 0, A(\varepsilon) \vdash [x' = f(x), t' = 1\,\&\,S]e \geq p(t)}{\Gamma_0, \varepsilon > 0, A(\varepsilon), t = 0, R(e_1) \vdash [x' = f(x), t' = 1\,\&\,S](p(t) - e_1 \leq 0 \land S)}$$

The derivation closes using the chain of differential cuts from (B.12). In the first dC step, the (constant) assumption $A(\varepsilon)$ is used, see the derivation labeled ⍟ immediately below:

$$\text{dC}\,\dfrac{\Gamma_0, t = 0 \vdash [x' = f(x), t' = 1\,\&\,S \land \dot{e}^{(k-1)} \geq \dot{e}_0^{(k-1)} + \varepsilon()t]e \geq p(t) \quad ⍟}{\Gamma_0, t = 0, A(\varepsilon) \vdash [x' = f(x), t' = 1\,\&\,S]e \geq p(t)}$$

From ⍟:

$$\text{dI}_{\succcurlyeq}\,\dfrac{\mathbb{R}\,\dfrac{*}{A(\varepsilon), S \vdash \dot{e}^{(k)} \geq \varepsilon()}}{\Gamma_0, t = 0, A(\varepsilon) \vdash [x' = f(x), t' = 1\,\&\,S]\dot{e}^{(k-1)} \geq \dot{e}_0^{(k-1)} + \varepsilon()t}$$

Subsequent dC, dI$_{\succcurlyeq}$ steps are similar to the derivation in Corollary 4.19:

$$\text{dC, dI}_{\succcurlyeq}\,\dfrac{\text{dC, dI}_{\succcurlyeq}\,\dfrac{\text{dI}_{\succcurlyeq}\,\dfrac{*}{\Gamma_0, t = 0 \vdash [x' = f(x), t' = 1\,\&\, \cdots \land \dot{e}^{(1)} \geq \dot{e}_0^{(1)} + \cdots + \varepsilon()\frac{t^{k-1}}{(k-1)!}]e \geq p(t)}}{\vdots}}{\Gamma_0, t = 0 \vdash [x' = f(x), t' = 1\,\&\,S \land \dot{e}^{(k-1)} \geq \dot{e}_0^{(k-1)} + \varepsilon()t]e \geq p(t)}$$

From premise ①, a monotonicity step M$\langle'\rangle$ rephrases the postcondition of the cut using the assumption $\forall t > t_1\,(p(t) - e_1 > 0)$. Axiom BEx finishes the derivation since formula $S(x)$ characterizes a compact (and hence bounded) set:

$$\text{M}\langle'\rangle\,\dfrac{\text{BEx}\,\dfrac{*}{\vdash \langle x' = f(x), t' = 1\rangle(t > t_1 \lor \neg S)}}{\forall t > t_1\,(p(t) - e_1 > 0) \vdash \langle x' = f(x), t' = 1\rangle(p(t) - e_1 > 0 \lor \neg S)} \qquad \square$$

B.2.4 Proofs for ODE Liveness Proofs in Practice

Proof of Corollary 4.33. Assume that formulas P, G_P are in normal form (3.7) as in Corollary 4.33. Rule dV is derived first since rule dV$^{\exists}$ follows from dV as a corollary. The derivation of rule dV uses variable b as a symbolic lower bound on the initial values of all terms $e_{ij}, \tilde{e}_{ij}$ appearing in formula P. The formula $\exists b \bigwedge_{i=0}^{M} \left(\bigwedge_{j=0}^{m(i)} e_{ij} \geq b \land \bigwedge_{j=0}^{n(i)} \tilde{e}_{ij} \geq b \right)$ is a valid formula of real arithmetic and is proved as a cut by $\mathbb{R}$ because P is a finite formula so there exists a lower bound b smaller than the value of all of the terms $e_{ij}, \tilde{e}_{ij}$.

The derivation starts similarly to $\text{dV}_{\succcurlyeq}$ by introducing fresh variables b (for the bound above), and i representing the multiplicative inverse of $\varepsilon()$ using arithmetic cuts cut, $\mathbb{R}$. It then Skolemizes ($\exists$L) and uses dGt to introduce a fresh time variable to the system of differential equations:

$$\begin{array}{r r}
 & \Gamma, \varepsilon() > 0, \bigwedge_{i=0}^{M} \Big(\bigwedge_{j=0}^{m(i)} e_{ij} \geq b \land \bigwedge_{j=0}^{n(i)} \tilde{e}_{ij} \geq b \Big), i\varepsilon() = 1, t = 0 \vdash \langle x' = f(x), t' = 1 \rangle P \\ \hline
\text{dGt} & \Gamma, \varepsilon() > 0, \bigwedge_{i=0}^{M} \Big(\bigwedge_{j=0}^{m(i)} e_{ij} \geq b \land \bigwedge_{j=0}^{n(i)} \tilde{e}_{ij} \geq b \Big), i\varepsilon() = 1 \vdash \langle x' = f(x) \rangle P \\ \hline
\exists\text{L} & \Gamma, \varepsilon() > 0, \exists b \bigwedge_{i=0}^{M} \Big(\bigwedge_{j=0}^{m(i)} e_{ij} \geq b \land \bigwedge_{j=0}^{n(i)} \tilde{e}_{ij} \geq b \Big), \exists i\, (i\varepsilon() = 1) \vdash \langle x' = f(x) \rangle P \\ \hline
\text{cut}, \mathbb{R} & \Gamma, \varepsilon() > 0 \vdash \langle x' = f(x) \rangle P
\end{array}$$

Next, the refinement axiom K⟨&⟩ is used with $G \equiv (b + \varepsilon()t > 0)$. This yields two premises, the right of which is proved by GEx (after monotonic rephrasing with $\mathbb{R}$, M⟨′⟩) because the ODE $x' = f(x)$ is assumed to have provable global solutions. The left premise from K⟨&⟩ is abbreviated ① and continued below.

$$\dfrac{\text{①} \qquad \dfrac{\dfrac{*}{\Gamma \vdash \langle x' = f(x), t' = 1 \rangle t > -ib}\,\text{GEx}}{\Gamma, \varepsilon() > 0, i\varepsilon() = 1 \vdash \langle x' = f(x), t' = 1 \rangle (b + \varepsilon()t > 0)}\,\mathbb{R}, \text{M}\langle ' \rangle}{\Gamma, \varepsilon() > 0, \bigwedge_{i=0}^{M} \Big(\bigwedge_{j=0}^{m(i)} e_{ij} \geq b \land \bigwedge_{j=0}^{n(i)} \tilde{e}_{ij} \geq b \Big), i\varepsilon() = 1, t = 0 \vdash \langle x' = f(x), t' = 1 \rangle P}\,\text{K}\langle \& \rangle$$

Continuing from premise ①, monotonicity strengthens the postcondition from $b + \varepsilon()t \leq 0$ to G_P under the domain constraint assumption $\neg P$. This strengthening is justified because, assuming that $\neg P$ and G_P are true in a given state, then propositionally, at least one of the following pairs (each pair listed horizontally) of sub-formulas of $\neg P$ and G_P for some indices i, j is true in that state:

$$\begin{array}{ll}
e_{ij} < 0 \qquad & e_{ij} - (b + \varepsilon()t) \geq 0 \\
\tilde{e}_{ij} \leq 0 \qquad & \tilde{e}_{ij} - (b + \varepsilon()t) \geq 0
\end{array}$$

Either pair of formulas imply that formula $b + \varepsilon()t \leq 0$ is also true in that state, so the strengthening is proved by M[′], $\mathbb{R}$. Next, a cut, $\mathbb{R}$ step adds the formula G_P to the antecedents using the assumptions $\bigwedge_{i=0}^{M} \Big(\bigwedge_{j=0}^{m(i)} e_{ij} \geq b \land \bigwedge_{j=0}^{n(i)} \tilde{e}_{ij} \geq b \Big)$ and $t = 0$. Rule sAI& yields the sole premise of rule dV because G_P characterizes a closed set.

$$\begin{array}{r r}
 & \neg P, (\dot{\neg P})^{(*)}, G_P \vdash (\dot{G_P})^{(*)} \\ \hline
\text{sAI\&} & G_P \vdash [x' = f(x), t' = 1 \,\&\, \neg P] G_P \\ \hline
\text{cut}, \mathbb{R} & \Gamma, \bigwedge_{i=0}^{M} \Big(\bigwedge_{j=0}^{m(i)} e_{ij} \geq b \land \bigwedge_{j=0}^{n(i)} \tilde{e}_{ij} \geq b \Big), t = 0 \vdash [x' = f(x), t' = 1 \,\&\, \neg P] G_P \\ \hline
\text{M}[']\text{, } \mathbb{R} & \Gamma, \bigwedge_{i=0}^{M} \Big(\bigwedge_{j=0}^{m(i)} e_{ij} \geq b \land \bigwedge_{j=0}^{n(i)} \tilde{e}_{ij} \geq b \Big), t = 0 \vdash [x' = f(x), t' = 1 \,\&\, \neg P] \big(b + \varepsilon()t \leq 0 \big)
\end{array}$$

Rule $\text{dV}^{\exists}$ is derived from rule dV similarly to the derivation of rule $\text{dV}_{\succcurlyeq}^{\exists}$ from rule $\text{dV}_{\succcurlyeq}$. The derivation starts with a cut of the sole premise of $\text{dV}^{\exists}$ (the left premise below). The existentially bound variable is renamed to δ throughout the derivation for clarity. The right premise is abbreviated ② and shown below.

$$\dfrac{\Gamma \vdash \exists \delta{>}0\, \forall b\, \forall t\, \forall x\, \big(\neg P \land (\dot{\neg P})^{(*)} \land G_P \rightarrow (\dot{G_P})^{(*)} \big) \qquad \text{②}}{\Gamma \vdash \langle x' = f(x) \,\&\, Q \rangle P}\,\text{cut}$$

From ②, after Skolemizing (with ∃L), rule dV is used with $\varepsilon() = \delta$. The universally quantified antecedent is constant for the ODE $x' = f(x)$ and the universal quantification over variables b, t ensure that those variables are fresh in the rest of the sequent so the antecedent is soundly kept across the application of rule dV. The proof is completed propositionally ∀L, →L, ∧L.

$$
{\scriptstyle \exists\mathrm{L},\, \wedge\mathrm{L}}\;\dfrac{{\scriptstyle \mathrm{dV}}\;\dfrac{{\scriptstyle \forall\mathrm{L},\, \rightarrow\mathrm{L},\, \wedge\mathrm{L}}\;\dfrac{*}{\forall b\,\forall t\,\forall x\,\big(\neg P \wedge (\dot{\neg P})^{(*)} \wedge G_P \rightarrow (\dot{G_P})^{(*)}\big), \neg P, (\dot{\neg P})^{(*)}, G_P \vdash (\dot{G_P})^{(*)}}}{\delta{>}0, \forall b\,\forall t\,\forall x\,\big(\neg P \wedge (\dot{\neg P})^{(*)} \wedge G_P \rightarrow (\dot{G_P})^{(*)}\big) \vdash \langle x' = f(x) \,\&\, Q\rangle P}}{\exists \delta{>}0\, \forall b\,\forall t\,\forall x\,\big(\neg P \wedge (\dot{\neg P})^{(*)} \wedge G_P \rightarrow (\dot{G_P})^{(*)}\big) \vdash \langle x' = f(x) \,\&\, Q\rangle P} \quad \square
$$

Proof of Corollary 4.36. The derivation of rule cR is seemingly straightforward using axiom CR followed by rule Enc on the resulting middle premise. There is a minor subtlety to address because the formula $Q^{>}_{\geq}$ (with strict inequalities replacing non-strict ones in Q) is only a syntactic *under-approximation* of the interior of the set characterized by Q, and so the axiom CR does *not* immediately apply as stated. For example, formula $x < x$ characterizes the empty set, while the formula $x \leq x$ characterizes the set of all states, whose interior is also the set of all states. Assume that the interior of Q is characterized by formula $\mathring{Q}$.[5]

The derivation starts with a cut of the formula Q which yields the leftmost premise of rule cR. This is followed with DX, which adds formula $\neg P$ to the antecedents because there is nothing to prove if both formulas Q and P are already true initially. The derivation then uses CR with the computable formula $\mathring{Q}$ characterizing the topological interior of formula Q. This yields two premises, the right of which corresponds to the rightmost premise of rule cR. From the resulting left premise (with postcondition $\mathring{Q}$), an M['], ℝ monotonicity step strengthens the postcondition because $Q^{>}_{\geq} \rightarrow \mathring{Q}$ is a provable formula of real arithmetic (Appendix B.1.3). Rule Enc completes the derivation because formula $Q^{>}_{\geq}$ is formed from finite conjunctions and disjunctions of strict inequalities, and $(Q^{>}_{\geq})^{\geq}_{>}$ is syntactically equal to Q by definition.

$$
{\scriptstyle \mathrm{cut}}\;\dfrac{\Gamma \vdash Q \qquad {\scriptstyle \mathrm{DX}}\;\dfrac{{\scriptstyle \mathrm{CR}}\;\dfrac{{\scriptstyle \mathrm{M}['],\, \mathbb{R}}\;\dfrac{{\scriptstyle \mathrm{Enc}}\;\dfrac{\Gamma \vdash [x' = f(x) \,\&\, R \wedge \neg P \wedge Q]Q^{>}_{\geq}}{\Gamma, Q \vdash [x' = f(x) \,\&\, R \wedge \neg P]Q^{>}_{\geq}}}{\Gamma, Q \vdash [x' = f(x) \,\&\, R \wedge \neg P]\mathring{Q}} \qquad \Gamma \vdash \langle x' = f(x) \,\&\, R\rangle P}{\Gamma, Q, \neg P \vdash \langle x' = f(x) \,\&\, Q\rangle P}}{\Gamma, Q \vdash \langle x' = f(x) \,\&\, Q\rangle P}}{\Gamma \vdash \langle x' = f(x) \,\&\, Q\rangle P}
$$

□

B.3 Counterexamples

This appendix gives explicit counterexamples to illustrate the soundness errors identified in Sections 4.4 and 4.5.

[5] If Q is a *semialgebraic* formula, there is a computable quantifier-free formula $\mathring{Q}$ that exactly characterizes its topological interior [14] which can be used with CR in the syntactic derivation.

B.3.1 Finite-Time Blow Up

The soundness errors identified in Section 4.4 all arise because of incorrect handling of the fact that solutions may blow up in finite time. This phenomenon is studied in detail in Section 4.3, and it is illustrated by α_n (4.2), see Fig. 4.1, or α_b (4.7), see Example 4.5. The following is a counterexample for the original presentation of $\mathrm{dV}_=$ (and $\mathrm{dV}_=^M$, $\mathrm{dV}_=\&$, $\mathrm{dV}_=^M\&$) [191]. Similar counterexamples can be constructed for [157, Remark 3.6] and for the original presentation of SLyap, SLyap& [159].

Counterexample B.2. Consider rule $\mathrm{dV}_=$ *without* the restriction that the ODE has provable global solutions. This unrestricted rule, denoted $\mathrm{dV}_=$↯, is unsound as shown by the following derivation using it with $\varepsilon() = 1$:

$$\mathrm{dV}_{=}\text{↯}\ \dfrac{\mathbb{R}\ \dfrac{*}{v - 2 < 0 \vdash 1 \geq 1}}{v - 2 \leq 0 \vdash \langle u' = u^2, v' = 1\rangle v - 2 = 0}$$

The conclusion of this derivation is not valid. Consider the initial state ω with values $\omega(u) = 1$ and $\omega(v) = 0$. The explicit solution of the ODE from ω is given by $u(t) = \frac{1}{1-t}, v(t) = t$ for $t \in [0, 1)$. This solution *does not exist* beyond the time interval $[0, 1)$ because the u-coordinate asymptotically approaches ∞, i.e., blows up, as time approaches $t = 1$. It is impossible to reach a state satisfying $v - 2 = 0$ from ω along this solution since at least 2 time units are required.

This counterexample further illustrates the difficulty in handling nonlinear ODEs. Neither the precondition ($v - 2 \leq 0$) nor postcondition ($v - 2 = 0$) mention the variable u, and the ODEs $u' = u^2, v' = 1$ do not depend on variables v, u respectively, so it is tempting to disregard the variable u entirely. Indeed, the liveness property $v - 2 \leq 0 \rightarrow \langle v' = 1\rangle v - 2 = 0$ is valid. Yet, for liveness questions about the (original) ODE, $u' = u^2, v' = 1$, the two variables are inextricably linked through the time axis of solutions to the ODE. ↯

B.3.2 Topological Considerations

The soundness errors identified in Section 4.5 arise because of incorrect topological reasoning in subtle cases where the topological boundaries of the sets characterized by the domain constraint and desired liveness postcondition intersect. The original presentation of $\mathrm{dV}_\succcurlyeq\&$ [137] gives the following proof rule for atomic inequalities $p \succcurlyeq 0$. For simplicity, assume that the ODE $x' = f(x)$ is globally Lipschitz continuous so that solutions exist for all time.

$$\mathrm{dV}_{\succcurlyeq}\&\text{↯}\ \dfrac{\Gamma \vdash [x' = f(x)\,\&\,p \leq 0]Q \qquad \neg(p \succcurlyeq 0), Q \vdash \dot{p} \geq \varepsilon()}{\Gamma, \varepsilon() > 0 \vdash \langle x' = f(x)\,\&\,Q\rangle p \succcurlyeq 0}$$

Compared to $\mathrm{dV}_\succcurlyeq\&$, the unsound rule $\mathrm{dV}_\succcurlyeq\&$↯ omits the assumption $\neg(p \succcurlyeq 0)$, makes no topological assumptions on the domain constraint Q, and uses syntactic weak negation [137] for the domain constraint of its left premise. The following two counterexamples show that the two assumptions are necessary.

Counterexample B.3. Consider the following derivation using the unsound rule $\mathrm{dV}_{\succcurlyeq}\&\lightning$ with $\varepsilon() = 1$:

$$\mathrm{dV}_{\succcurlyeq}\&\lightning\,\dfrac{\text{dW}, \mathbb{R}\,\dfrac{*}{u > 1 \vdash [u' = 1 \,\&\, u \le 0]u \le 1} \qquad \mathbb{R}\,\dfrac{*}{u < 0, u \le 1 \vdash 1 \ge 1}}{u > 1 \vdash \langle u' = 1 \,\&\, u \le 1\rangle u \ge 0}$$

The conclusion of this derivation is not valid. In states where $u > 1$ is true initially, the domain constraint is violated immediately so the diamond modality in the succedent is trivially false in those states. $\lightning$

Counterexample B.4 ([175])**.** This counterexample is adapted from [175, Example 142], which has a minor typographical error (the sign of an inequality is flipped). Consider the following derivation using the unsound rule $\mathrm{dV}_{\succcurlyeq}\&\lightning$ with $\varepsilon() = 1$:

$$\mathrm{dV}_{\succcurlyeq}\&\lightning\,\dfrac{\text{dW}, \mathbb{R}\,\dfrac{*}{\vdash [u' = 1 \,\&\, u \le 1]u \le 1} \qquad \mathbb{R}\,\dfrac{*}{u \le 1, u \le 1 \vdash 1 \ge 1}}{\vdash \langle u' = 1 \,\&\, u \le 1\rangle u > 1}$$

The conclusion of this derivation is not valid and, in fact, unsatisfiable. The domain constraint $u \le 1$ and postcondition $u > 1$ are contradictory so no solution can reach a state satisfying both simultaneously. $\lightning$

The next two counterexamples are for the liveness arguments from [156, Corollary 1] and [157, Theorem 3.5]. For clarity, the original notation from [157, Theorem 3.5] is used. The following conjecture is quoted from [157, Theorem 3.5]:

Conjecture B.5. *Consider the system $x' = f(x)$, with $f \in C(\mathbb{R}^n, \mathbb{R}^n)$. Let $\mathcal{X} \subset \mathbb{R}^n$, $\mathcal{X}_0 \subseteq \mathcal{X}$, and $\mathcal{X}_r \subseteq \mathcal{X}$ be bounded sets. If there exists a function $B \in C^1(\mathbb{R}^n)$ satisfying:*

$$B(x) \le 0 \qquad \forall x \in \mathcal{X}_0 \tag{B.13}$$

$$B(x) > 0 \qquad \forall x \in \overline{\partial\mathcal{X} \setminus \partial\mathcal{X}_r} \tag{B.14}$$

$$\frac{\partial B}{\partial x} f(x) < 0 \qquad \forall x \in \overline{\mathcal{X} \setminus \mathcal{X}_r} \tag{B.15}$$

Then the eventuality property holds, i.e., for all initial conditions $x_0 \in \mathcal{X}_0$, the trajectory $x(t)$ of the system starting at $x(0) = x_0$ satisfies $x(T) \in \mathcal{X}_r$ and $x(t) \in \mathcal{X}$ for all $t \in [0, T]$ for some $T \ge 0$. The notation $\overline{\mathcal{X}}$ (resp. $\partial\mathcal{X}$) denotes the topological closure (resp. boundary) of the set $\mathcal{X}$.

In [156, Corollary 1], stronger conditions are required. In particular, the sets $\mathcal{X}_0, \mathcal{X}_r, \mathcal{X}$ are additionally required to be topologically open, and the inequality in (B.13) is required to be strict, i.e., $B(x) < 0$ instead of $B(x) \le 0$.

The soundness errors in both of these liveness arguments stem from the condition (B.14) being too permissive. For example, notice that if the sets $\partial\mathcal{X}, \partial\mathcal{X}_r$ are equal then (B.14) is vacuously true. The first counterexample below applies for the requirements of [157, Theorem 3.5], while the second applies even for the more restrictive requirements of [156, Corollary 1].

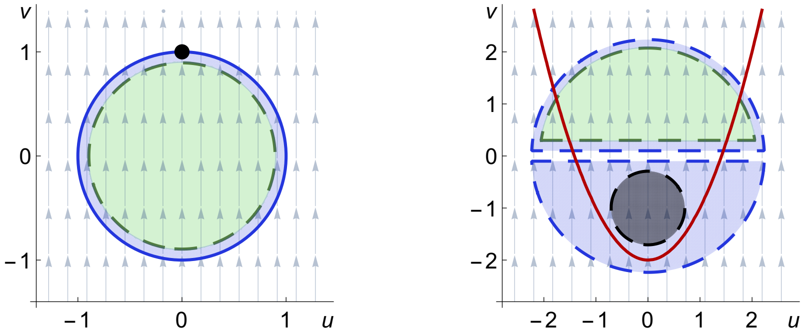

Figure B.1: **(Left)** Visualization of Counterexample B.6. The solution from initial point $u = 0, v = 1$ ($\mathcal{X}_0$, in black) leaves the domain unit disk ($\mathcal{X}$, boundary in blue) immediately without ever reaching its interior ($\mathcal{X}_r$, in green with dashed boundary). The interior is slightly shrunk for clarity in the visualization: the blue and green boundaries should actually overlap exactly. **(Right)** Visualization of Counterexample B.7. Solutions from the initial set ($\mathcal{X}_0$, in black with dashed boundary) eventually enter the goal region ($\mathcal{X}_r$, in green with dashed boundary). However, the domain ($\mathcal{X}$, in blue with dashed boundary) shares an (open) boundary with $\mathcal{X}_r$ at $v = 0$ which solutions are not allowed to cross. The sets are slightly shrunk for clarity in the visualization: the blue and green boundaries should actually overlap exactly. The level curve $B = 0$ is plotted in red. All points above the curve satisfy $B < 0$, while all points below it satisfy $B > 0$.

Counterexample B.6. Let the system $x' = f(x)$ be $u' = 0, v' = 1$. Let $\mathcal{X}_r$ be the open unit disk characterized by $u^2 + v^2 < 1$, $\mathcal{X}$ be the closed unit disk characterized by $u^2 + v^2 \leq 1$, and $\mathcal{X}_0$ be the single point characterized by $u = 0 \wedge v = 1$. All of these sets are bounded. Note that $\partial\mathcal{X} \setminus \partial\mathcal{X}_r = \emptyset$ since both topological boundaries are the unit circle $u^2 + v^2 = 1$. Let $B(u, v) = -v$, so that $\frac{\partial B}{\partial x} f(x) = \frac{\partial B}{\partial u} 0 + \frac{\partial B}{\partial v} 1 = -1 < 0$ and $B \leq 0$ on $\mathcal{X}_0$.

All conditions of [157, Theorem 3.5] are met but the eventuality property is false. The trajectory from $\mathcal{X}_0$ leaves $\mathcal{X}$ immediately and never enters $\mathcal{X}_r$, see Fig. B.1 (Left). ↯

Counterexample B.7. Let the system $x' = f(x)$ be $u' = 0, v' = 1$. Let $\mathcal{X}_r$ be the set characterized by the formula $u^2 + v^2 < 5 \wedge v > 0$, $\mathcal{X}$ be the set characterized by the formula $u^2 + v^2 < 5 \wedge v \neq 0$, and $\mathcal{X}_0$ be the set characterized by the formula $u^2 + (v + 1)^2 < \frac{1}{2}$. All of these sets are bounded and topologically open. Let $B(u, v) = -v + u^2 - 2$, so that $\frac{\partial B}{\partial x} f(x) = \frac{\partial B}{\partial u} 0 + \frac{\partial B}{\partial v} 1 = -1 < 0$, and $B < 0$ on $\mathcal{X}_0$. The set $\overline{\partial\mathcal{X} \setminus \partial\mathcal{X}_r}$ is characterized by formula $u^2 + v^2 = 5 \wedge v \leq 0$ and B is strictly positive on this set. These claims can be checked arithmetically, see Fig. B.1 (Right) for a plot of the curve $B = 0$.

All conditions of [156, Corollary 1] are met but the eventuality property is false. Solutions starting in $\mathcal{X}_0$ eventually enter $\mathcal{X}_r$ but can only do so by leaving the domain constraint $\mathcal{X}$ at $v = 0$, see Fig. B.1 (Right). ↯

Appendix C

Appendix: Stability for Ordinary Differential Equations

C.1 Derived Stability Proof Rules

This appendix provides proofs for all lemmas and corollaries that were omitted from Sections 5.2 and 5.3. For ease of reference, this appendix is organized into two sections, corresponding to proofs for Section 5.2 and Section 5.3 respectively. This appendix uses derived proof rules and axioms from Chapters 3 and 4 (and their Appendices A and B) together with an additional derived axiom:

$$\text{DCC}\quad \begin{array}{l}[x' = f(x) \,\&\, Q \land P]R \land [x' = f(x) \,\&\, Q](\neg P \rightarrow [x' = f(x) \,\&\, Q]\neg P) \\ \quad \rightarrow [x' = f(x) \,\&\, Q](P \rightarrow R)\end{array}$$

Axiom DCC to prove that an implication $P \rightarrow R$ is always true along an ODE, it suffices to prove it assuming P in the domain if, whenever the solution leaves P, then it stays in the negation $\neg P$ afterwards (so the implication is trivially true in those states). The axiom is stated as a proof rule elsewhere [90] and its axiomatic version is formally verified [16].

C.1.1 Proofs for Asymptotic Stability of an Equilibrium Point

This section concerns stability for the origin, whose ε neighborhoods $\mathcal{U}_\varepsilon(x = 0)$ are equivalently unfolded as the formula $\|x\|_2^2 < \varepsilon^2$. The following lemma formalizes the claim in Section 5.2.1 that a point x_0 of interest for the ODE $x' = f(x)$ can be rigorously translated *with proof* to the origin so that, without loss of generality, only stability of the origin needs to be considered for the stability proof rules of Section 5.2.

Lemma C.1 (Translation to origin)**.** *The following axioms are derivable in* dL*, where the ODE $y' = f(y + x_0)$ has point x_0 translated to the origin and variables y are fresh, i.e., not in ODE $x' = f(x)$ or formula $P(x)$.*

$$\text{Trans}\quad y = x - x_0 \rightarrow \big([x' = f(x)]P(x - x_0) \leftrightarrow [y' = f(y + x_0)]P(y)\big)$$

$$\text{TransStab}\quad \mathrm{Stab}(y' = f(y + x_0)) \rightarrow \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), x = x_0, x = x_0)$$

Proof. Axiom Trans is derived first before axiom TransStab is derived from it as a corollary. Only the "$\rightarrow$" direction of the inner equivalence for axiom Trans is derived since the "$\leftarrow$" direction follows from the "$\rightarrow$" direction by renaming and translation with respect to $-x_0$.

Let $\alpha_{xy} \equiv x' = f(x), y' = f(y+x_0)$ abbreviate the combined ODE for variables x and y. The derivation starts with a cut of formula $[\alpha_{xy}]\, y = x - x_0$, which says that the value of y is always equal to $x - x_0$ (component-wise) along solutions of the combined ODE α_{xy}. This cut is provable in dL using axiom DRI which is complete for analytic invariants. Subsequently, axiom BDG adds the ghost ODE $y' = f(y + x_0)$ to ODE $x' = f(x)$ in the antecedent box modality and ODE $x' = f(x)$ to ODE $y' = f(y + x_0)$ in the succedent box modality. The resulting boundedness premises from the use of BDG are respectively abbreviated ① and ② and they are both proved using the cut antecedent as shown further below.

$${\scriptstyle\text{cut, DRI}}\dfrac{{\scriptstyle\text{BDG}}\dfrac{{\scriptstyle\text{BDG}}\dfrac{\text{②} \qquad [\alpha_{xy}]y = x - x_0, [\alpha_{xy}]P(x - x_0) \vdash [\alpha_{xy}]P(y)}{\text{①} \qquad [\alpha_{xy}]y = x - x_0, [\alpha_{xy}]P(x - x_0) \vdash [y' = f(y + x_0)]P(y)}}{[\alpha_{xy}]y = x - x_0, [x' = f(x)]P(x - x_0) \vdash [y' = f(y + x_0)]P(y)}}{y = x - x_0, [x' = f(x)]P(x - x_0) \vdash [y' = f(y + x_0)]P(y)}$$

From the (unabbreviated) open right premise, a dC step adds formulas $y = x - x_0$ and $P(x - x_0)$ to the domain constraint of the succedent. A subsequent dW step completes the proof by substituting $y = x - x_0$ in the succedent $P(y)$.

$${\scriptstyle\text{dC}}\dfrac{{\scriptstyle\text{dC}}\dfrac{{\scriptstyle\text{dW}}\dfrac{{\scriptstyle\mathbb{R}}\dfrac{*}{y = x - x_0 \land P(x - x_0) \vdash P(y)}}{\vdash [\alpha_{xy} \,\&\, y = x - x_0 \land P(x - x_0)]P(y)}}{[\alpha_{xy}]P(x - x_0) \vdash [\alpha_{xy} \,\&\, y = x - x_0]P(y)}}{[\alpha_{xy}]y = x - x_0, [\alpha_{xy}]P(x - x_0) \vdash [\alpha_{xy}]P(y)}$$

For premise ①, the ghost variables y are provably bounded in (squared) norm by the term $\|x - x_0\|_2^2$ for ODE α_{xy}. The dC step adds $y = x - x_0$ from the antecedent box modality to the domain constraint, and the subsequent dW step completes the proof by substituting $y = x - x_0$ and by real arithmetic $\mathbb{R}$.

$${\scriptstyle\text{dC}}\dfrac{{\scriptstyle\text{dW}}\dfrac{{\scriptstyle\mathbb{R}}\dfrac{*}{y = x - x_0 \vdash \|y\|_2^2 \leq \|x - x_0\|_2^2}}{\vdash [\alpha_{xy} \,\&\, y = x - x_0]\, \|y\|_2^2 \leq \|x - x_0\|_2^2}}{[\alpha_{xy}]y = x - x_0 \vdash [\alpha_{xy}]\, \|y\|_2^2 \leq \|x - x_0\|_2^2}$$

The derivation for premise ② is similar, where the ghost variables x are provably bounded in (squared) norm by the term $\|y + x_0\|_2^2$ for ODE α_{xy} instead.

$${\scriptstyle\text{dC}}\dfrac{{\scriptstyle\text{dW}}\dfrac{{\scriptstyle\mathbb{R}}\dfrac{*}{y = x - x_0 \vdash \|x\|_2^2 \leq \|y + x_0\|_2^2}}{\vdash [\alpha_{xy} \,\&\, y = x - x_0]\, \|x\|_2^2 \leq \|y + x_0\|_2^2}}{[\alpha_{xy}]y = x - x_0 \vdash [\alpha_{xy}]\, \|x\|_2^2 \leq \|y + x_0\|_2^2}$$

The derivation of axiom TransStab starts by Skolemizing ε in the succedent with $\forall$R and then instantiating the antecedent with the resulting fresh Skolem variable ε using $\forall$L. This is followed by $\exists$L, $\exists$R which Skolemizes δ in the antecedent and then witnesses the succedent with

δ. Next, $\forall$R, $\rightarrow$R Skolemizes the succedent before $\forall$L instantiates y in the quantified antecedent, with the translated coordinate $y = x - x_0$. Formula $y = x - x_0 \land \|x - x_0\|_2^2 < \delta^2 \rightarrow \|y\|_2^2 < \delta^2$ is provable by substitution in real arithmetic, so $\rightarrow$L, $\mathbb{R}$ proves the LHS of the implication ($\|y\|_2^2 < \delta^2$) in the antecedent before Trans completes the proof. The formulas are abbreviated $R_y \equiv [y' = f(y + x_0)]\ \|y\|_2^2 < \varepsilon^2$ and $R \equiv [x' = f(x)]\ \|x - x_0\|_2^2 < \varepsilon^2$, respectively.

$$
\begin{array}{rc}
 & * \\ \hline
\text{Trans} & y = x - x_0, R_y \vdash R \\ \hline
\rightarrow\text{L}, \mathbb{R} & y = x - x_0, \|y\|_2^2 < \delta^2 \rightarrow R_y, \|x - x_0\|_2^2 < \delta^2 \vdash R \\ \hline
\forall\text{L} & \forall y \left(\|y\|_2^2 < \delta^2 \rightarrow R_y \right), \|x - x_0\|_2^2 < \delta^2 \vdash R \\ \hline
\forall\text{R}, \rightarrow\text{R} & \forall y \left(\|y\|_2^2 < \delta^2 \rightarrow R_y \right) \vdash \forall x \left(\|x - x_0\|_2^2 < \delta^2 \rightarrow R \right) \\ \hline
\exists\text{L}, \exists\text{R} & \exists \delta{>}0\, \forall y \left(\|y\|_2^2 < \delta^2 \rightarrow R_y \right) \vdash \exists \delta{>}0\, \forall x \left(\|x - x_0\|_2^2 < \delta^2 \rightarrow R \right) \\ \hline
\forall\text{R}, \forall\text{L} & \mathrm{Stab}(y' = f(y + x_0)) \vdash \mathrm{Stab}_{\mathrm{R}}^{\mathrm{P}}(x' = f(x), x = x_0, x = x_0)
\end{array}
\quad \square
$$

Proof of Lemma 5.6. The full derivation of rule Lyap$_>$ from rule Lyap$_\geq$ is given in Lemma 5.6. The derivation of rule Lyap$_\geq$ begins with a series of arithmetic cuts which are justified stepwise. For any $\varepsilon > 0$ and an equilibrium point at the origin with $f(0) = 0$, the second (right) premise of Lyap$_\geq$ can be equivalently strengthened to choose $\gamma \leq \varepsilon$, i.e., the second premise provably implies the following formula in real arithmetic; the universal quantifier on x is also distributed across the inner conjunction as shown in conjuncts $\textcircled{a}$ and $\textcircled{b}$ below:

$$
\exists 0{<}\gamma{\leq}\varepsilon \left(\underbrace{\forall x \left(0 < \|x\|_2^2 \leq \gamma^2 \rightarrow V > 0\right)}_{\textcircled{a}} \land \underbrace{\forall x \left(\|x\|_2^2 \leq \gamma^2 \rightarrow \dot{V} \leq 0\right)}_{\textcircled{b}} \right)
$$

The derivation begins with a cut of this formula and Skolemizing with $\exists$L, yielding antecedents $\textcircled{a}$ and $\textcircled{b}$ as indicated above. The postcondition is then monotonically strengthened to $\|x\|_2^2 < \gamma^2$ using antecedent $\gamma \leq \varepsilon$.

$$
\begin{array}{rc}
 & \gamma > 0, \textcircled{a}, \textcircled{b} \vdash \exists \delta{>}0\, \forall x \left(\|x\|_2^2 < \delta^2 \rightarrow [x' = f(x)]\ \|x\|_2^2 < \gamma^2 \right) \\ \hline
\mathrm{M}['] & \gamma > 0, \gamma \leq \varepsilon, \textcircled{a}, \textcircled{b} \vdash \exists \delta{>}0\, \forall x \left(\|x\|_2^2 < \delta^2 \rightarrow [x' = f(x)]\ \|x\|_2^2 < \varepsilon^2 \right) \\ \hline
\text{cut}, \exists\text{L} & \varepsilon > 0 \vdash \exists \delta{>}0\, \forall x \left(\|x\|_2^2 < \delta^2 \rightarrow [x' = f(x)]\ \|x\|_2^2 < \varepsilon^2 \right) \\ \hline
\forall\text{R} & \vdash \mathrm{Stab}(x' = f(x))
\end{array}
$$

From $\textcircled{a}$, the continuous Lyapunov function V is positive on the compact set characterized by $\|x\|_2^2 = \gamma^2$ and therefore is bounded below by its minimum $W > 0$ on that set. Furthermore, from premise $V(0) = 0$, by continuity, V must take values smaller than W in a ball with sufficiently small radius $0 < \delta < \gamma$ around the origin.[1] Thus, the following formula proves in real arithmetic from $\textcircled{a}$.

$$
\exists W \left(\underbrace{\forall x \left(\|x\|_2^2 = \gamma^2 \rightarrow V \geq W \right)}_{\textcircled{c}} \land \exists 0{<}\delta{<}\gamma\ \underbrace{\forall x \left(\|x\|_2^2 < \delta^2 \rightarrow V < W \right)}_{\textcircled{d}} \right)
$$

[1] Recall the notational convention (Appendix B.1.3) that rule $\mathbb{R}$ proves these real analytic properties for Lyapunov functions V over extended terms.

The derivation continues with a cut of the above formula and Skolemizing the resulting antecedent with $\exists$L, yielding antecedents $\textcircled{c}$ and $\textcircled{d}$ as indicated above.

$$\text{cut, }\mathbb{R},\exists\text{L}\ \dfrac{0{<}\delta{<}\gamma, \textcircled{b}, \textcircled{c}, \textcircled{d} \vdash \exists\delta{>}0\,\forall x\,\big(\|x\|_2^2<\delta^2 \to [x'=f(x)]\,\|x\|_2^2<\gamma^2\big)}{\gamma>0, \textcircled{a}, \textcircled{b} \vdash \exists\delta{>}0\,\forall x\,\big(\|x\|_2^2<\delta^2 \to [x'=f(x)]\,\|x\|_2^2<\gamma^2\big)}$$

The succedent existential quantifier $\exists\delta{>}0$ is instantiated with the antecedent's δ using $\exists$R, followed by simplification steps, Skolemizing and unfolding the succedent with $\forall$R, $\to$R then instantiating $\textcircled{d}$ by $\forall$L, $\to$L with the resulting $\|x\|_2^2<\delta^2$ assumption.

$$\exists\text{R}\ \dfrac{\forall\text{R},\to\text{R}\ \dfrac{\forall\text{L},\to\text{L}\ \dfrac{\delta{<}\gamma, \textcircled{b}, \textcircled{c}, \|x\|_2^2<\delta^2, V<W \vdash [x'=f(x)]\,\|x\|_2^2<\gamma^2}{\delta{<}\gamma, \textcircled{b}, \textcircled{c}, \textcircled{d}, \|x\|_2^2<\delta^2 \vdash [x'=f(x)]\,\|x\|_2^2<\gamma^2}}{\delta{<}\gamma, \textcircled{b}, \textcircled{c}, \textcircled{d} \vdash \forall x\,\big(\|x\|_2^2<\delta^2 \to [x'=f(x)]\,\|x\|_2^2<\gamma^2\big)}}{0{<}\delta{<}\gamma, \textcircled{b}, \textcircled{c}, \textcircled{d} \vdash \exists\delta{>}0\,\forall x\,\big(\|x\|_2^2<\delta^2 \to [x'=f(x)]\,\|x\|_2^2<\gamma^2\big)}$$

Since formula $\|x\|_2^2<\gamma^2$ characterizes an open ball and $\delta<\gamma$, the antecedent $\|x\|_2^2<\delta^2$ implies $\|x\|_2^2<\gamma^2$ arithmetically so rule Enc is used to assume its closure $\|x\|_2^2\le\gamma^2$ in the domain constraint of the succedent ODE. With the strengthened domain, dC adds $\dot{V}\le 0$ to the domain constraint by V using assumption $\textcircled{b}$ which is universally quantified over x.

$$\text{Enc}\ \dfrac{\text{dC, V}\ \dfrac{\textcircled{c}, V<W \vdash [x'=f(x)\,\&\,\|x\|_2^2\le\gamma^2\wedge\dot{V}\le 0]\,\|x\|_2^2<\gamma^2}{\textcircled{b}, \textcircled{c}, V<W \vdash [x'=f(x)\,\&\,\|x\|_2^2\le\gamma^2]\,\|x\|_2^2<\gamma^2}}{\delta{<}\gamma, \textcircled{b}, \textcircled{c}, \|x\|_2^2<\delta^2, V<W \vdash [x'=f(x)]\,\|x\|_2^2<\gamma^2}$$

The proof continues using dC to add invariant $V<W$ to the domain constraint; this differential cut proves by dI$_{\succcurlyeq}$ using conjunct $\dot{V}\le 0$ in the domain constraint. The subsequent dC step adds $\|x\|_2^2\ne\gamma^2$ to the domain constraint using the contrapositive direction of the universally quantified antecedent $\textcircled{c}$. Finally, a dW step completes the proof since conjuncts $\|x\|_2^2\le\gamma^2$ and $\|x\|_2^2\ne\gamma^2$ in the resulting domain constraint imply the postcondition $\|x\|_2^2<\gamma^2$ by $\mathbb{R}$.

$$\text{dI}_{\succcurlyeq},\text{dC}\ \dfrac{\text{dC}\ \dfrac{\text{dW}\ \dfrac{\mathbb{R}\ \dfrac{*}{\|x\|_2^2\le\gamma^2, \|x\|_2^2\ne\gamma^2 \vdash \|x\|_2^2<\gamma^2}}{\vdash [x'=f(x)\,\&\,\|x\|_2^2\le\gamma^2\wedge\cdots\wedge\|x\|_2^2\ne\gamma^2]\,\|x\|_2^2<\gamma^2}}{\textcircled{c} \vdash [x'=f(x)\,\&\,\|x\|_2^2\le\gamma^2\wedge\dot{V}\le 0\wedge V<W]\,\|x\|_2^2<\gamma^2}}{\textcircled{c}, V<W \vdash [x'=f(x)\,\&\,\|x\|_2^2\le\gamma^2\wedge\dot{V}\le 0]\,\|x\|_2^2<\gamma^2} \qquad \square$$

Proof of Corollary 5.10. The two directions of axiom EStabE are derived separately. In the "→" direction, the derivation starts by unfolding the discrete assignments in the sequent with [:=], [;]. In the derivation, the succedent is bound renamed with variable z [142] and the ODEs are temporarily hidden with $\cdots$ (shown further below).

$$[:=],[;]\ \dfrac{y=\alpha^2\|x\|_2^2, z=\alpha^2\|x\|_2^2, t=0, [\cdots]\,\|x\|_2^2\le y \vdash [\cdots]\,\|x\|_2^2\le z\exp(-2\beta t)}{[y:=\alpha^2\|x\|_2^2;\cdots]\,\|x\|_2^2\le y \vdash [z:=\alpha^2\|x\|_2^2; t:=0;\cdots]\,\|x\|_2^2\le z\exp(-2\beta t)}$$

Continuing from the resulting premise, the cut, $\mathbb{R}_{\exp}$ step adds formula $y=z\exp(-2\beta t)$ to the antecedent because the exponential sub-term $\exp(-2\beta t)$ simplifies to 1 when $t=0$. This step requires arithmetic over the exponential function with rule $\mathbb{R}_{\exp}$.

$$\text{cut}, \mathbb{R}_{\exp}\ \dfrac{y = z\exp(-2\beta t), [\cdots]\ \|x\|_2^2 \le y \vdash [\cdots]\ \|x\|_2^2 \le z\exp(-2\beta t)}{y = \alpha^2\|x\|_2^2, z = \alpha^2\|x\|_2^2, t = 0, [\cdots]\ \|x\|_2^2 \le y \vdash [\cdots]\ \|x\|_2^2 \le z\exp(-2\beta t)}$$

Next, the derivation uses axiom DG to add ghost ODE $t' = 1$ to the antecedent box modality and ODE $y' = -2\beta y$ to the succedent, respectively. The resulting ODEs (identical in antecedent and succedent) is abbreviated $\alpha_{yt} \equiv x'{=}f(x), y'{=}{-2\beta y}, t' = 1$. Rule dC adds the postcondition of the antecedent box modality to the domain constraint, then M['] monotonically strengthens the succedent postcondition to $y{=}z\exp(-2\beta t)$ using the strengthened domain constraint (by substitution). The proof is completed using rule dbx with cofactor $g = -2\beta$.

$$\begin{array}{rc}
 & * \\ \hline
\text{dbx} & y - z\exp(-2\beta t) = 0 \vdash [\alpha_{yt}]y - z\exp(-2\beta t) = 0 \\ \hline
 & y{=}z\exp(-2\beta t) \vdash [\alpha_{yt}]y{=}z\exp(-2\beta t) \\ \hline
\text{M}['] & y{=}z\exp(-2\beta t) \vdash [\alpha_{yt}\,\&\ \|x\|_2^2 \le y]\ \|x\|_2^2 \le z\exp(-2\beta t) \\ \hline
\text{dC} & y{=}z\exp(-2\beta t), [\alpha_{yt}]\ \|x\|_2^2 \le y \vdash [\alpha_{yt}]\ \|x\|_2^2 \le z\exp(-2\beta t) \\ \hline
\text{DG} & y{=}z\exp(-2\beta t), [x'{=}f(x), y'{=}{-2\beta y}]\ \|x\|_2^2 \le y \vdash [x'{=}f(x), t'{=}1]\ \|x\|_2^2 \le z\exp(-2\beta t)
\end{array}$$

The derivation in the "←" direction is similar and given briefly below. The derivation starts by unfolding discrete assignments with [:=], [;], where the ODEs temporarily hidden with $\cdots$. The cut, $\mathbb{R}_{\exp}$ step adds formula $y = z\exp(-2\beta t)$ to the antecedent because the exponential sub-term $\exp(-2\beta t)$ simplifies to 1 when $t = 0$.

$$\begin{array}{rc}
 & y = z\exp(-2\beta t), [\cdots]\ \|x\|_2^2 \le z\exp(-2\beta t) \vdash [\cdots]\ \|x\|_2^2 \le y \\ \hline
\text{cut}, \mathbb{R}_{\exp} & y = \alpha^2\|x\|_2^2, z = \alpha^2\|x\|_2^2, t = 0, [\cdots]\ \|x\|_2^2 \le z\exp(-2\beta t) \vdash [\cdots]\ \|x\|_2^2 \le y \\ \hline
[:=], [;] & [z := \alpha^2\|x\|_2^2; t := 0; \cdots]\ \|x\|_2^2 \le z\exp(-2\beta t) \vdash [y := \alpha^2\|x\|_2^2; \cdots]\ \|x\|_2^2 \le y
\end{array}$$

The derivation then uses axiom DG to add ghost ODEs, where the resulting ODEs (identical in antecedent and succedent) is abbreviated $\alpha_{yt} \equiv x'{=}f(x), y'{=}{-2\beta y}, t' = 1$. The remaining derivation is similar to the "→" direction (details omitted).

$$\begin{array}{rc}
 & * \\ \hline
\text{dbx} & y - z\exp(-2\beta t) = 0 \vdash [\alpha_{yt}]y - z\exp(-2\beta t) = 0 \\ \hline
 & y{=}z\exp(-2\beta t) \vdash [\alpha_{yt}]y{=}z\exp(-2\beta t) \\ \hline
\text{M}['] & y{=}z\exp(-2\beta t) \vdash [\alpha_{yt}\,\&\ \|x\|_2^2 \le z\exp(-2\beta t)]\ \|x\|_2^2 \le y \\ \hline
\text{dC} & y{=}z\exp(-2\beta t), [\alpha_{yt}]\ \|x\|_2^2 \le z\exp(-2\beta t) \vdash [\alpha_{yt}]\ \|x\|_2^2 \le y \\ \hline
\text{DG} & y{=}z\exp(-2\beta t), [\cdots]\ \|x\|_2^2 \le z\exp(-2\beta t) \vdash [\cdots]\ \|x\|_2^2 \le y
\end{array}$$

□

Proof of Lemma 5.11. The proof starts by instantiating the existentially quantified variables α, β in $\mathrm{EStab}(x' = f(x))$ with $\alpha = \frac{k_2}{k_1}$ and decay rate $\beta = k_3$. Since k_1, k_2, k_3 are all positive constants, these choices satisfy $\alpha > 0, \beta > 0$.

$$\exists\text{R}\ \dfrac{\vdash \exists \delta{>}0\, \forall x\, \big(\, \|x\|_2^2 < \delta^2 \to [y := (\frac{k_2}{k_1})^2\|x\|_2^2; x' = f(x), y' = -2k_3 y]\ \|x\|_2^2 \le y\big)}{\vdash \mathrm{EStab}(x' = f(x))}$$

The subsequent cut step adds the premise of rule Lyap_{E} to the antecedents and Skolemizes it with $\exists$L. The resulting antecedent is abbreviated with

$$\textcircled{a} \equiv \forall x\, \big(\|x\|_2^2 \le \gamma^2 \to k_1^2 \|x\|_2^2 \le V \le k_2^2 \|x\|_2^2 \land \dot{V} \le -2k_3 V \big)$$

Then $\exists$R instantiates the succedent with $\delta = \frac{k_1}{k_2}\gamma$ and the sequent is logically simplified. Note $\textcircled{a}$ also implies $k_1 \le k_2$ as $k_1, k_2 > 0$ so $\delta \le \gamma$ and $\|x\|_2^2 < \delta^2$ implies $\|x\|_2^2 < \gamma$ in real arithmetic. The hybrid program $y := (\frac{k_2}{k_1})^2 \|x\|_2^2 ; x' = f(x), y' = -2k_3 y$ is abbreviated with $\cdots$ in the first three steps below.

$$\begin{array}{rc}
 & \gamma > 0, \textcircled{a}, \|x\|_2^2 < (\frac{k_1}{k_2}\gamma)^2 \vdash [y := (\frac{k_2}{k_1})^2 \|x\|_2^2 ; x' = f(x), y' = -2k_3 y]\, \|x\|_2^2 \le y \\ \hline
\text{$\forall$R, $\to$R, $\to$L} & \gamma > 0, \textcircled{a} \vdash \forall x\, \big(\|x\|_2^2 < (\frac{k_1}{k_2}\gamma)^2 \to [\cdots]\, \|x\|_2^2 \le y \big) \\ \hline
\exists\text{R} & \gamma > 0, \textcircled{a} \vdash \exists \delta{>}0\, \forall x\, \big(\|x\|_2^2 < \delta^2 \to [\cdots]\, \|x\|_2^2 \le y \big) \\ \hline
\text{cut, $\exists$L} & \vdash \exists \delta{>}0\, \forall x\, \big(\|x\|_2^2 < \delta^2 \to [\cdots]\, \|x\|_2^2 \le y \big)
\end{array}$$

The discrete assignment $y := (\frac{k_2}{k_1})^2 \|x\|_2^2$ sets the value of variable y to $(\frac{k_2}{k_1})^2 \|x\|_2^2$ initially. It is turned into an equational assumption with the assignment axiom [:=] and [;] as follows [144].

$$\begin{array}{rc}
 & \gamma > 0, \textcircled{a}, \|x\|_2^2 < (\frac{k_1}{k_2}\gamma)^2, y = (\frac{k_2}{k_1})^2 \|x\|_2^2 \vdash [x' = f(x), y' = -2k_3 y]\, \|x\|_2^2 \le y \\ \hline
\text{[:=], [;]} & \gamma > 0, \textcircled{a}, \|x\|_2^2 < (\frac{k_1}{k_2}\gamma)^2 \vdash [y := (\frac{k_2}{k_1})^2 \|x\|_2^2 ; x' = f(x), y' = -2k_3 y]\, \|x\|_2^2 \le y
\end{array}$$

The antecedents are abbreviated $\Gamma \equiv \gamma > 0, \textcircled{a}, \|x\|_2^2 < (\frac{k_1}{k_2}\gamma)^2, y = (\frac{k_2}{k_1})^2 \|x\|_2^2$. The derivation continues with a differential cut dC adding formula $\|x\|_2^2 < \gamma^2$ to the domain constraint. This cut is abbreviated $\textcircled{1}$ and proved further below. The next differential cut dC adds formula $V \le k_1^2 y$ to the domain constraint. This cut is abbreviated $\textcircled{2}$ and also proved further below. The derivation is completed with a dW, $\mathbb{R}$ step with the quantified antecedent $\textcircled{a}$ and the domain constraint, since they imply the chain of inequalities $k_1^2 \|x\|_2^2 \le V \le k_1^2 y$, which implies the succedent (after dW) $\|x\|_2^2 \le y$ by $\mathbb{R}$.

$$\begin{array}{rll}
 & * & \\ \hline
\mathbb{R} & \textcircled{a}, \|x\|_2^2 < \gamma^2, V \le k_1^2 y \vdash \|x\|_2^2 \le y & \\ \hline
\text{dW} & \qquad\qquad \textcircled{2} & \Gamma \vdash [x' = f(x), y' = -2k_3 y \,\&\, \|x\|_2^2 < \gamma^2 \land V \le k_1^2 y]\, \|x\|_2^2 \le y \\ \hline
\text{dC} & \qquad\qquad \textcircled{1} & \Gamma \vdash [x' = f(x), y' = -2k_3 y \,\&\, \|x\|_2^2 < \gamma^2]\, \|x\|_2^2 \le y \\ \hline
\text{dC} & & \Gamma \vdash [x' = f(x), y' = -2k_3 y]\, \|x\|_2^2 \le y
\end{array}$$

Returning to premise $\textcircled{1}$, the derivation uses Enc to assume $\|x\|_2^2 \le \gamma^2$ in the domain constraint, since $\|x\|_2^2 < \gamma^2$ is true initially. Then, a dC, $\text{dI}_{\succcurlyeq}$ step adds formula $V < k_1^2 \gamma^2$ to the domain constraint. This formula is proved true initially by $\mathbb{R}$ with antecedents Γ using the chain of inequalities from $\textcircled{a}$, $V \le k_2^2 \|x\|_2^2 < k_2^2 \big(\frac{k_1}{k_2}\gamma\big)^2 = k_1^2 \gamma^2$. It is proved invariant by $\text{dI}_{\succcurlyeq}$ because domain constraint $\|x\|_2^2 \le \gamma^2$ and quantified antecedent $\textcircled{a}$ proves the chain of inequalities $\dot{V} \le -2k_3 V \le -2k_3 (k_1^2 \|x\|_2^2) \le 0$. A dW, $\mathbb{R}$ step completes the proof because the domain constraint $\|x\|_2^2 \le \gamma^2 \land V < k_1^2 \gamma^2$ and quantified antecedent $\textcircled{a}$ prove the chain of inequalities

$k_1^2 \left\|x\right\|_2^2 \leq V < k_1^2\gamma^2$, which implies the succedent (after dW) $\left\|x\right\|_2^2 < \gamma^2$ by $\mathbb{R}$.

$$
{\scriptstyle \text{Enc}}\;\dfrac{{\scriptstyle \text{dC, dI}_{\succcurlyeq}}\;\dfrac{{\scriptstyle \text{dW, }\mathbb{R}}\;\dfrac{{\scriptstyle \mathbb{R}}\;\dfrac{*}{\text{ⓐ}, \left\|x\right\|_2^2 \leq \gamma^2, V{<}k_1^2\gamma^2 \vdash \left\|x\right\|_2^2 {<} \gamma^2}}{\Gamma \vdash [x' = f(x), y' = -2k_3y \,\&\, \left\|x\right\|_2^2 \leq \gamma^2 \land V{<}k_1^2\gamma^2]\, \left\|x\right\|_2^2 {<} \gamma^2}}{\Gamma \vdash [x' = f(x), y' = -2k_3y \,\&\, \left\|x\right\|_2^2 \leq \gamma^2]\, \left\|x\right\|_2^2 {<} \gamma^2}}{\Gamma \vdash [x' = f(x), y' = -2k_3y]\, \left\|x\right\|_2^2 {<} \gamma^2}
$$

For premise ②, the inequality $k_1^2y - V \geq 0$ is proved invariant using rule $\text{dbx}_{\succcurlyeq}$ with cofactor $g = -2k_3$ as follows.

$$
{\scriptstyle \text{cut, M}[']}\;\dfrac{{\scriptstyle \text{dbx}_{\succcurlyeq}}\;\dfrac{{\scriptstyle \mathbb{R}}\dfrac{*}{\Gamma \vdash k_1^2y - V \geq 0} \qquad {\scriptstyle \mathbb{R}}\dfrac{*}{\text{ⓐ}, \left\|x\right\|_2^2 < \gamma^2 \vdash -2k_1^2k_3y - \dot{V} \geq -2k_3(k_1^2y - V)}}{\Gamma \vdash [x' = f(x), y' = -2k_3y \,\&\, \left\|x\right\|_2^2 < \gamma^2]k_1^2y - V \geq 0}}{\Gamma \vdash [x' = f(x), y' = -2k_3y \,\&\, \left\|x\right\|_2^2 < \gamma^2]V \leq k_1^2y}
$$

The resulting left premise proves by $\mathbb{R}$ because the antecedents ⓐ and $y = (\frac{k_2}{k_1})^2 \left\|x\right\|_2^2$ in Γ prove the chain of inequalities $V \leq k_2^2 \left\|x\right\|_2^2 = k_1^2y$. For the resulting right premise, the Lie derivative of $k_1^2y - V$ from the LHS of the postcondition is $-2k_1^2k_3y - \dot{V}$. With domain constraint $\left\|x\right\|_2^2 < \gamma^2$ and quantified antecedent ⓐ, this derivative provably satisfies the chain of inequalities $-2k_1^2k_3y - \dot{V} \geq -2k_1^2k_3y + 2k_3V = -2k_3(k_1^2y - V)$ by $\mathbb{R}$. □

Proof of Lemma 5.14. The rules are derived in order starting with rule $\text{Lyap}_{>}^{\text{G}}$. First, observe that the first two premises of rule $\text{Lyap}_{>}^{\text{G}}$ imply the premises of $\text{Lyap}_{\geq}$ because if the sign conditions on V and $\dot{V}$ are true globally, then they also hold for any choice of neighborhood of the origin. Thus, the derivation starts with a cut, $\text{Lyap}_{\geq}$ step which proves stability of the origin. Next, the definition of $\text{Attr}^{\text{P}}(x' = f(x), \mathit{true})$ is logically unfolded and axiom SAttr is used to simplify the succedent, together with logical unfolding steps ∀R, →R.

$$
{\scriptstyle \text{cut, Lyap}_{\geq}}\;\dfrac{{\scriptstyle \forall\text{R}, \rightarrow\text{R}}\;\dfrac{{\scriptstyle \text{SAttr}}\;\dfrac{{\scriptstyle \forall\text{R}, \rightarrow\text{R}}\;\dfrac{\varepsilon{>}0 \vdash \langle x' = f(x)\rangle \left\|x\right\|_2^2 < \varepsilon^2}{\vdash \forall \varepsilon{>}0\, \langle x' = f(x)\rangle \left\|x\right\|_2^2 < \varepsilon^2}}{\text{Stab}(x' = f(x)) \vdash \text{Asym}(x' = f(x), \mathit{true})}}{\text{Stab}(x' = f(x)) \vdash \text{Attr}^{\text{P}}(x' = f(x), \mathit{true})}}{\vdash \text{Stab}(x' = f(x)) \land \text{Attr}^{\text{P}}(x' = f(x), \mathit{true})}
$$

The derivation continues with a cut, $\mathbb{R}$, ∃L step, introducing a fresh Skolem variable b which stores the initial value of the Lyapunov function V. Next, a cut adds the box modality formula $[x' = f(x)]\, V \leq b$ to the antecedents. This cut proves by $\text{dI}_{\succcurlyeq}$ because formula $V \leq b$ is true initially and the premises of rule $\text{Lyap}_{>}^{\text{G}}$ prove the formula $\dot{V} \leq 0$.[2]

$$
{\scriptstyle \text{cut, }\mathbb{R}\text{, }\exists\text{L}}\;\dfrac{{\scriptstyle \text{cut, dI}_{\succcurlyeq}}\;\dfrac{\varepsilon{>}0, [x' = f(x)]V \leq b \vdash \langle x' = f(x)\rangle \left\|x\right\|_2^2 < \varepsilon^2}{\varepsilon{>}0, V = b \vdash \langle x' = f(x)\rangle \left\|x\right\|_2^2 < \varepsilon^2}}{\varepsilon{>}0 \vdash \langle x' = f(x)\rangle \left\|x\right\|_2^2 < \varepsilon^2}
$$

[2] When $x = 0$, the premise $f(0) = 0$ implies $\dot{V} = 0$.

The subsequent M['] step strengthens the postcondition $V \leq b$ of the antecedent box modality to $\|x\|_2^2 < \gamma^2$ using the (Skolemized) rightmost premise of rule $\text{Lyap}^{\text{G}}_{>}$. This rightmost premise corresponds to the assumption that the Lyapunov function V is radially unbounded [71, 89].

$$\text{M[']}\,\dfrac{\varepsilon{>}0, [x' = f(x)]\, \|x\|_2^2 < \gamma^2 \vdash \langle x' = f(x)\rangle\, \|x\|_2^2 < \varepsilon^2}{\varepsilon{>}0, [x' = f(x)] V \leq b \vdash \langle x' = f(x)\rangle\, \|x\|_2^2 < \varepsilon^2}$$

Like the derivation of rule $\text{Lyap}_{>}$ in Lemma 5.6, the remaining open premise is an ODE liveness property which is proved using rule SP_c with the choice of compact staging set $S \equiv \varepsilon^2 \leq \|x\|_2^2 \leq \gamma^2$ and $e \equiv V$. The resulting left premise proves with a differential cut dC of the antecedent and dW. The resulting right premise proves using $\mathbb{R}$ from the middle premise of rule $\text{Lyap}^{\text{G}}_{>}$ and antecedents $\varepsilon{>}0$, S.

$$\text{SP}_c\,\dfrac{\text{dC, dW}\,\dfrac{*}{[x' = f(x)]\ \|x\|_2^2 {<} \gamma^2 \vdash [x' = f(x)\,\&\ \|x\|_2^2 \geq \varepsilon^2] S} \qquad \mathbb{R}\,\dfrac{*}{\varepsilon{>}0, S \vdash \dot{V} < 0}}{\varepsilon{>}0, [x' = f(x)]\ \|x\|_2^2 {<} \gamma^2 \vdash \langle x' = f(x)\rangle\, \|x\|_2^2 {<} \varepsilon^2}$$

The derivation of rule $\text{Lyap}^{\text{G}}_{\text{E}}$ is similar to the derivation of rule Lyap_{E} from Lemma 5.11. The proof steps are repeated briefly. The derivation starts by instantiating (using ∃R) the existentially quantified variables in succedent $\text{EStab}^{\text{P}}(x' = f(x), \mathit{true})$ with $\alpha = \frac{k_2}{k_1}$ and decay rate $\beta = k_3$, followed by logical unfolding of the sequent.

$$\begin{array}{ll}
 & y{=}(\frac{k_2}{k_1})^2\, \|x\|_2^2 \vdash [x'{=}f(x), y'{=}-2k_3 y]\ \|x\|_2^2 \leq y \\ \hline
\text{[:=], [;]} & \vdash [y := (\frac{k_2}{k_1})^2\, \|x\|_2^2\, ; x'{=}f(x), y'{=}-2k_3 y]\ \|x\|_2^2 \leq y \\ \hline
\forall\text{R}, \to\text{R} & \vdash \forall x \left(\mathit{true} \to [y := (\frac{k_2}{k_1})^2\, \|x\|_2^2\, ; x'{=}f(x), y'{=}-2k_3 y]\ \|x\|_2^2 \leq y\right) \\ \hline
\exists\text{R} & \vdash \text{EStab}^{\text{P}}(x'{=}f(x), \mathit{true})
\end{array}$$

The proof continues with a differential cut of the formula $V \leq k_1^2 y$, which is proved invariant with its equivalent rephrasing $k_1^2 y - V \geq 0$ and rule dbx$_{\succcurlyeq}$ using cofactor $g = -2k_3$. Similar to Lemma 5.11, the cut formula is true initially because the antecedent $y = (\frac{k_2}{k_1})^2\, \|x\|_2^2$ and the premise of $\text{Lyap}^{\text{G}}_{\text{E}}$ imply the chain of inequalities $V \leq k_2^2\, \|x\|_2^2 = k_1^2 y$. The premise of $\text{Lyap}^{\text{G}}_{\text{E}}$ are also used to show that the Lie derivative of $k_1^2 y - V$ provably satisfies the chain of inequalities $-2k_1^2 k_3 y - \dot{V} \geq -2k_1^2 k_3 y + 2k_3 V = -2k_3(k_1^2 y - V)$. The proof is completed by dW, $\mathbb{R}$ using the premise of rule $\text{Lyap}^{\text{G}}_{\text{E}}$.

$$\begin{array}{ll}
 & * \\ \hline
\mathbb{R} & \vdash k_1^2\, \|x\|_2^2 \leq V \\ \hline
\mathbb{R} & V \leq k_1^2 y \vdash\ \|x\|_2^2 \leq y \\ \hline
\text{dW} & \vdash [x' = f(x), y' = -2k_3 y\,\&\, V \leq k_1^2 y]\ \|x\|_2^2 \leq y \\ \hline
\text{dC} & y = (\frac{k_2}{k_1})^2\, \|x\|_2^2 \vdash [x' = f(x), y' = -2k_3 y]\ \|x\|_2^2 \leq y
\end{array}$$
□

Proof of Corollary 5.16. The two axioms are derived in order, starting with axiom EStabStab. The derivation of axiom EStabStab starts by Skolemizing the existential quantifiers in $\text{EStab}(x' = f(x))$ with ∃L, then Skolemizing the succedent (∀R) and instantiating (∃R) the existentially quantified $\delta{>}0$ in the succedent with $\gamma = \min(\frac{\varepsilon}{\alpha}, \delta)$ (note $\gamma > 0$). The sequent is then simplified, noting that $\gamma \leq \delta$ so that assumption $\|x\|_2^2 < \gamma^2$ proves the implication LHS $\forall x \left(\|x\|_2^2 < \delta^2 \rightarrow \cdots \right)$ in the antecedents. The subformula with discrete assignment to y in the antecedent is abbreviated with $R_y \equiv [y := \alpha^2 \|x\|_2^2 ; x' = f(x), y' = -2\beta y] \ \|x\|_2^2 \leq y$.

$$\begin{array}{r c}
 & \alpha{>}0, \beta{>}0, \delta{>}0, \varepsilon{>}0, R_y, \|x\|_2^2 {<} \gamma^2 \vdash [x'{=}f(x)] \ \|x\|_2^2 {<} \varepsilon^2 \\ \hline
\forall\text{R}, \rightarrow\text{R}, \rightarrow\text{L} & \alpha{>}0, \beta{>}0, \delta{>}0, \varepsilon{>}0, \forall x \left(\|x\|_2^2 {<} \delta^2 {\rightarrow} R_y \right) \vdash \forall x \left(\|x\|_2^2 {<} \gamma^2 {\rightarrow} [x'{=}f(x)] \ \|x\|_2^2 {<} \varepsilon^2 \right) \\ \hline
\exists\text{R} & \alpha{>}0, \beta{>}0, \delta{>}0, \varepsilon{>}0, \forall x \left(\|x\|_2^2 {<} \delta^2 {\rightarrow} R_y \right) \vdash \exists \delta{>}0 \, \forall x \left(\|x\|_2^2 {<} \delta^2 {\rightarrow} \dots \right) \\ \hline
\forall\text{R} & \alpha{>}0, \beta{>}0, \delta{>}0, \forall x \left(\|x\|_2^2 {<} \delta^2 {\rightarrow} R_y \right) \vdash \text{Stab}(x'{=}f(x)) \\ \hline
\exists\text{L} & \text{EStab}(x'{=}f(x)) \vdash \text{Stab}(x'{=}f(x))
\end{array}$$

The discrete assignment in antecedent R_y is unfolded with [:=], [;], similar to the proof of Lemma 5.11, yielding the antecedent $y{=}\alpha^2 \|x\|_2^2$ and abbreviated box modality antecedent $P_y \equiv [x' = f(x), y' = -2\beta y] \ \|x\|_2^2 \leq y$. Axiom DG adds differential ghost $y' = -2\beta y$ to the succedent ODE and the postcondition is strengthened to $y < \varepsilon^2$ by K using assumption P_y. The proof is completed with a dbx$_{\succcurlyeq}$ step with cofactor term -2β because the (rephrased) postcondition $\varepsilon^2 - y > 0$ is proved by $\mathbb{R}$ from the antecedents with the chain of inequalities $y = \alpha^2 \|x\|_2^2 < \alpha^2 \gamma^2 \leq \alpha^2 (\frac{\varepsilon}{\alpha})^2 = \varepsilon^2$. The Lie derivative of $\varepsilon^2 - y$ is $2\beta y$ which provably satisfies the inequality $2\beta y \geq 2\beta y - 2\beta \varepsilon^2 = -2\beta(\varepsilon^2 - y)$ for $\beta > 0$.

$$\begin{array}{r c}
 & * \\ \hline
\text{dbx}_{\succcurlyeq} & \alpha{>}0, \beta{>}0, \delta{>}0, \varepsilon{>}0, y{=}\alpha^2 \|x\|_2^2, \|x\|_2^2 < \gamma^2 \vdash [x'{=}f(x), y'{=}{-2\beta y}] \, \varepsilon^2 - y > 0 \\ \hline
\text{M}['] & \alpha{>}0, \beta{>}0, \delta{>}0, \varepsilon{>}0, y{=}\alpha^2 \|x\|_2^2, \|x\|_2^2 < \gamma^2 \vdash [x'{=}f(x), y'{=}{-2\beta y}] \, y < \varepsilon^2 \\ \hline
\text{K} & \alpha{>}0, \beta{>}0, \delta{>}0, \varepsilon{>}0, y{=}\alpha^2 \|x\|_2^2, P_y, \|x\|_2^2 < \gamma^2 \vdash [x'{=}f(x), y'{=}{-2\beta y}] \ \|x\|_2^2 < \varepsilon^2 \\ \hline
\text{DG} & \alpha{>}0, \beta{>}0, \delta{>}0, \varepsilon{>}0, y{=}\alpha^2 \|x\|_2^2, P_y, \|x\|_2^2 < \gamma^2 \vdash [x'{=}f(x)] \ \|x\|_2^2 < \varepsilon^2 \\ \hline
[:=], [;] & \alpha{>}0, \beta{>}0, \delta{>}0, \varepsilon{>}0, R_y, \|x\|_2^2 < \gamma^2 \vdash [x'{=}f(x)] \ \|x\|_2^2 < \varepsilon^2
\end{array}$$

The derivation of axiom EStabAttr starts by unfolding and Skolemizing the existential quantifiers for the antecedent, with abbreviated ODE $\alpha_{xy} \equiv x' = f(x), y' = -2\beta y$ and subformula $R_y \equiv [y := \alpha^2 \|x\|_2^2 ; \alpha_{xy}] \ \|x\|_2^2 \leq y$ (identically to the preceding derivation for axiom EStabStab). The succedent is logically unfolded and the resulting antecedent P proves the implication LHS in the antecedent $\forall x \, (P \rightarrow R_y)$. Succedent $\text{Asym}(x' = f(x), x = 0)$ is then Skolemized with $\forall$R.[3]

$$\begin{array}{r c}
 & \alpha{>}0, \beta{>}0, R_y \vdash \langle x' = f(x) \rangle [x' = f(x)] \ \|x\|_2^2 < \varepsilon^2 \\ \hline
\forall\text{R} & \alpha{>}0, \beta{>}0, R_y, P \vdash \text{Asym}(x' = f(x), x = 0) \\ \hline
\forall\text{L}, \rightarrow\text{L} & \alpha{>}0, \beta{>}0, \forall x \left(P \rightarrow R_y\right), P \vdash \text{Asym}(x' = f(x), x = 0) \\ \hline
\forall\text{R}, \rightarrow\text{R} & \alpha{>}0, \beta{>}0, \forall x \left(P \rightarrow R_y\right) \vdash \text{Attr}^{\text{P}}(x' = f(x), P) \\ \hline
\exists\text{L} & \text{EStab}^{\text{P}}(x' = f(x), P) \vdash \text{Attr}^{\text{P}}(x' = f(x), P)
\end{array}$$

The derivation continues with axioms DG, DG$_\forall$ to add the linear differential ghost $y' = -2\beta y$ to both ODEs in the succedent. The postcondition of the succedent diamond modality is monotonically strengthened with M$\langle ' \rangle$ and postcondition $(y < \varepsilon^2 \land [\alpha_{xy}] \ \|x\|_2^2 \leq y)$. The two

[3] Unlike earlier proofs , the formula is *not* simplified using a stability assumption because the generalized formula $\text{EStab}^{\text{P}}(x' = f(x), P)$ does not directly imply stability of the origin unless formula P provably contains a neighborhood of the origin.

resulting premises are abbreviated ① and ②; they are shown and proved further below.

$$\text{DG, DG}_\forall\;\dfrac{\text{M}\langle'\rangle\;\dfrac{\text{①} \qquad \text{②}}{\alpha{>}0, \beta{>}0, R_y \vdash \langle \alpha_{xy} \rangle [\alpha_{xy}]\; \|x\|_2^2 < \varepsilon^2}}{\alpha{>}0, \beta{>}0, R_y \vdash \langle x' = f(x) \rangle [x' = f(x)]\; \|x\|_2^2 < \varepsilon^2}$$

From premise ①, a differential cut dC proves formula $y < \varepsilon^2$ invariant for the succedent ODE α_{xy} using rule $\text{dbx}_{\succcurlyeq}$. The subsequent dC adds the postcondition of the antecedent to the domain constraint before dW, $\mathbb{R}$ finish the derivation.

$$\text{dC, dbx}_{\succcurlyeq}\;\dfrac{\text{dC}\;\dfrac{\text{dW}\;\dfrac{\mathbb{R}\;\dfrac{*}{y < \varepsilon^2 \land \|x\|_2^2 \le y \vdash \|x\|_2^2 < \varepsilon^2}}{\vdash [\alpha_{xy} \,\&\, y < \varepsilon^2 \land \|x\|_2^2 \le y]\; \|x\|_2^2 < \varepsilon^2}}{[\alpha_{xy}]\; \|x\|_2^2 \le y \vdash [\alpha_{xy} \,\&\, y < \varepsilon^2]\; \|x\|_2^2 < \varepsilon^2}}{\beta{>}0, y < \varepsilon^2, [\alpha_{xy}]\; \|x\|_2^2 \le y \vdash [\alpha_{xy}]\; \|x\|_2^2 < \varepsilon^2}$$

From premise ②, the antecedent R_y is unfolded with [:=], [;], then K⟨&⟩, D[;] remove the right conjunct of the postcondition because postcondition $[\alpha_{xy}]\; \|x\|_2^2 \le y$ is true after all runs of α_{xy}. Axiom BDG removes the ODEs for x in the succedent because $\|x\|_2^2$ is bounded using antecedent $[\alpha_{xy}]\; \|x\|_2^2 \le y$.

$$\text{[:=], [;]}\;\dfrac{\text{K}\langle\&\rangle\text{, D[;]}\;\dfrac{\text{BDG}\;\dfrac{\beta{>}0 \vdash \langle y' = -2\beta y \rangle y < \varepsilon^2}{\beta{>}0, [\alpha_{xy}]\; \|x\|_2^2 \le y \vdash \langle \alpha_{xy} \rangle y < \varepsilon^2}}{\beta{>}0, [\alpha_{xy}]\; \|x\|_2^2 \le y \vdash \langle \alpha_{xy} \rangle \big(y < \varepsilon^2 \land [\alpha_{xy}]\; \|x\|_2^2 \le y\big)}}{\beta{>}0, R_y \vdash \langle \alpha_{xy} \rangle \big(y < \varepsilon^2 \land [\alpha_{xy}]\; \|x\|_2^2 \le y\big)}$$

The remaining open premise is an ODE liveness property for variable y. Its proof starts by introducing a fresh variable y_0 storing the initial value of y. Then, rule SP_c is used with $S \equiv \big(\varepsilon^2 \le y \le y_0\big)$. The resulting left premise is an invariance property of the ODE which proves using M[′], $\text{dI}_{\succcurlyeq}$, while the right premise proves by $\mathbb{R}$.

$$\text{cut, }\exists\text{L}\;\dfrac{\text{SP}_c\;\dfrac{\text{M}[']\;\dfrac{\text{dI}_{\succcurlyeq}\;\dfrac{*}{\beta{>}0, y = y_0 \vdash [y' = -2\beta y \,\&\, y \ge \varepsilon^2] y \le y_0}}{\beta{>}0, y = y_0 \vdash [y' = -2\beta y \,\&\, y \ge \varepsilon^2] S} \qquad \mathbb{R}\;\dfrac{*}{\beta{>}0, S \vdash -2\beta y < 0}}{\beta{>}0, y = y_0 \vdash \langle y' = -2\beta y \rangle y < \varepsilon^2}}{\beta{>}0 \vdash \langle y' = -2\beta y \rangle y < \varepsilon^2} \qquad \square$$

C.1.2 Proofs for General Stability

This section derives proof rules for general stability and its specialized instances which are introduced and motivated in Section 5.3.

Proof of Lemma 5.17. The derivation of rule GLyap generalizes the ideas behind the derivation of rule $\text{Lyap}_{\ge}$. The derivation starts with an ∀R step, followed by a cut and Skolemization ∃L of the second (bottom) premise of rule GLyap. The resulting assumptions (for Skolem variables γ, δ, W) are abbreviated with: ⓐ $\equiv \forall x\, (\partial(\mathcal{U}_\gamma(R)) \to V \ge W)$, ⓑ $\equiv \forall x\, (\mathcal{U}_\delta(P) \to R \lor V < W)$, and $\textcircled{c} \equiv \forall x \big(R \vee V < W \rightarrow [x' = f(x) \,\&\, \overline{\mathcal{U}_\gamma(R)}](R \vee V < W)\big)$. A subsequent M['] step strengthens the postcondition monotonically since the formula $\mathcal{U}_\gamma(R) \rightarrow \mathcal{U}_\varepsilon(R)$ is provable by $\mathbb{R}$ for $\gamma \leq \varepsilon$.

$$\begin{array}{r c}
 & 0{<}\gamma, 0{<}\delta{\leq}\gamma, \textcircled{a}, \textcircled{b}, \textcircled{c} \vdash \exists \delta{>}0 \,\forall x \,\big(\mathcal{U}_\delta(P) \rightarrow [x' = f(x)]\,\mathcal{U}_\gamma(R)\big) \\ \cline{2-2}
\text{M}['], \mathbb{R} & 0{<}\gamma{\leq}\varepsilon, 0{<}\delta{\leq}\gamma, \textcircled{a}, \textcircled{b}, \textcircled{c} \vdash \exists \delta{>}0 \,\forall x \,\big(\mathcal{U}_\delta(P) \rightarrow [x' = f(x)]\,\mathcal{U}_\varepsilon(R)\big) \\ \cline{2-2}
\text{cut}, \exists\text{L} & \varepsilon > 0 \vdash \exists \delta{>}0 \,\forall x \,\big(\mathcal{U}_\delta(P) \rightarrow [x' = f(x)]\,\mathcal{U}_\varepsilon(R)\big) \\ \cline{2-2}
\forall\text{R} & \vdash \text{Stab}_{\text{R}}^{\text{P}}(x' = f(x), P, R)
\end{array}$$

The existentially quantified δ in the succedent is witnessed with the Skolem variable δ in the antecedents and the sequent is simplified with ∀R, →R, →L, where the implication LHS in $\textcircled{b}$ is proved using antecedent assumption $\mathcal{U}_\delta(P)$.

$$\begin{array}{r c}
 & 0{<}\gamma, 0{<}\delta{\leq}\gamma, \textcircled{a}, \textcircled{c}, \mathcal{U}_\delta(P), R \vee V < W \vdash [x' = f(x)]\,\mathcal{U}_\gamma(R) \\ \cline{2-2}
\forall\text{R}, \rightarrow\text{R}, \rightarrow\text{L} & 0{<}\gamma, 0{<}\delta{\leq}\gamma, \textcircled{a}, \textcircled{b}, \textcircled{c} \vdash \forall x \,\big(\mathcal{U}_\delta(P) \rightarrow [x' = f(x)]\,\mathcal{U}_\gamma(R)\big) \\ \cline{2-2}
\exists\text{R} & 0{<}\gamma, 0{<}\delta{\leq}\gamma, \textcircled{a}, \textcircled{b}, \textcircled{c} \vdash \exists \delta{>}0 \,\forall x \,\big(\mathcal{U}_\delta(P) \rightarrow [x' = f(x)]\,\mathcal{U}_\gamma(R)\big)
\end{array}$$

The derivation continues with rule Enc to assume the closure formula $\overline{\mathcal{U}_\gamma(R)}$ in the domain constraint. This step uses the first premise of rule GLyap, i.e., precondition P implies postcondition R, so that the neighborhood formula $\mathcal{U}_\delta(P)$ provably implies neighborhood formula $\mathcal{U}_\gamma(R)$ initially by $\mathbb{R}$ for $\delta \leq \gamma$. The subsequent dC step uses the antecedent $\textcircled{c}$ to prove the invariance of formula $R \vee V{<}W$ for the ODE $x' = f(x) \,\&\, \overline{\mathcal{U}_\gamma(R)}$ and adds it to the domain constraint. The derivation is completed using dW, $\mathbb{R}$, where the arithmetic step is justified below.

$$\begin{array}{r c}
 & * \\ \cline{2-2}
\mathbb{R} & 0{<}\gamma, \textcircled{a}, \overline{\mathcal{U}_\gamma(R)}, (R \vee V{<}W) \vdash \mathcal{U}_\gamma(R) \\ \cline{2-2}
\text{dW} & 0{<}\gamma, \textcircled{a}, R \vee V{<}W \vdash [x' = f(x) \,\&\, \overline{\mathcal{U}_\gamma(R)} \wedge (R \vee V{<}W)]\,\mathcal{U}_\gamma(R) \\ \cline{2-2}
\text{dC} & 0{<}\gamma, \textcircled{a}, \textcircled{c}, R \vee V{<}W \vdash [x' = f(x) \,\&\, \overline{\mathcal{U}_\gamma(R)}]\,\mathcal{U}_\gamma(R) \\ \cline{2-2}
\text{Enc} & 0{<}\gamma, 0{<}\delta{\leq}\gamma, \textcircled{a}, \textcircled{c}, \mathcal{U}_\delta(P), R \vee V{<}W \vdash [x' = f(x)]\,\mathcal{U}_\gamma(R)
\end{array}$$

To prove arithmetical premise $0{<}\gamma, \textcircled{a}, \overline{\mathcal{U}_\gamma(R)}, R \vee V{<}W \vdash \mathcal{U}_\gamma(R)$, note that the left disjunct R in the antecedent implies its neighborhood formula $\mathcal{U}_\gamma(R)$ in real arithmetic for $\gamma > 0$. Thus, it suffices to justify the premise for the right disjunct, i.e., $\textcircled{a}, \overline{\mathcal{U}_\gamma(R)}, V{<}W \vdash \mathcal{U}_\gamma(R)$. Antecedent $\textcircled{a}$ is instantiated to obtain assumption $\partial(\mathcal{U}_\gamma(R)) \rightarrow V \geq W$. This assumption says that the right disjunct $V < W$ does *not* occur on the boundary of $\partial(\mathcal{U}_\gamma(R))$ which, together with antecedent $\overline{\mathcal{U}_\gamma(R)}$ implies succedent $\mathcal{U}_\gamma(R)$ in real arithmetic (Appendix B.1.3). □

Proof of Corollary 5.21. The three axioms are derived in order.

SetSAttr The two directions of the inner equivalence of SetSAttr are proved separately. The easier "→" direction follows by Skolemizing ε in the succedent with ∀R, choosing the same ε in the antecedent with ∀L, and then by M⟨′⟩ because the resulting postcondition $[x' = f(x)]\,\mathcal{U}_\varepsilon(P)$ monotonically implies postcondition $\mathcal{U}_\varepsilon(P)$ by differential skip DX.

$$\begin{array}{r c}
 & * \\ \cline{2-2}
\text{DX} & [x' = f(x)]\mathcal{U}_\varepsilon(P) \vdash \mathcal{U}_\varepsilon(P) \\ \cline{2-2}
\text{M}\langle ' \rangle & \langle x' = f(x) \rangle [x' = f(x)]\,\mathcal{U}_\varepsilon(P) \vdash \langle x' = f(x) \rangle\,\mathcal{U}_\varepsilon(P) \\ \cline{2-2}
\forall\text{R}, \forall\text{L} & \text{Asym}(x' = f(x), P) \vdash \forall \varepsilon{>}0 \,\langle x' = f(x) \rangle\,\mathcal{U}_\varepsilon(P)
\end{array}$$

The more interesting “←” direction uses the stability assumption. The first step Skolemizes the succedent with ∀R, then the stability antecedent is instantiated with ∀L and Skolemized with ∃L (yielding fresh Skolem variable δ). Using the resulting quantified assumption $\forall x\,\big(\mathcal{U}_\delta(P) \to [x' = f(x)]\,\mathcal{U}_\varepsilon(P)\big)$, the postcondition of the succedent is monotonically strengthened to $\mathcal{U}_\delta(P)$ with M⟨′⟩. The derivation is completed by ∀L instantiating the remaining quantified antecedent with δ.

$$\begin{array}{lr}
 & * \\ \hline
{\scriptstyle \forall\mathrm{L}} & \delta > 0, \forall \varepsilon{>}0\, \langle x' = f(x)\rangle\, \mathcal{U}_\varepsilon(P) \vdash \langle x' = f(x)\rangle\, \mathcal{U}_\delta(P) \\ \hline
{\scriptstyle \mathrm{M}\langle'\rangle} & \delta > 0, \forall x\,\big(\mathcal{U}_\delta(P) \to [x' = f(x)]\,\mathcal{U}_\varepsilon(P)\big), \forall \varepsilon{>}0\, \langle x' = f(x)\rangle\, \mathcal{U}_\varepsilon(P) \vdash \langle x' = f(x)\rangle [x' = f(x)]\,\mathcal{U}_\varepsilon(P) \\ \hline
{\scriptstyle \forall\mathrm{L},\,\exists\mathrm{L}} & \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P), \forall \varepsilon{>}0\, \langle x' = f(x)\rangle\, \mathcal{U}_\varepsilon(P), \varepsilon{>}0 \vdash \langle x' = f(x)\rangle [x' = f(x)]\,\mathcal{U}_\varepsilon(P) \\ \hline
{\scriptstyle \forall\mathrm{R}} & \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P), \forall \varepsilon{>}0\, \langle x' = f(x)\rangle\, \mathcal{U}_\varepsilon(P) \vdash \mathrm{Asym}(x' = f(x), P)
\end{array}$$

SClosure This axiom is derived immediately by equivalently rewriting with arithmetic equivalences because for $\delta > 0$ the open neighborhood formulas $\mathcal{U}_\delta(P)$ and $\mathcal{U}_\delta(\overline{P})$ are provably equivalent in real arithmetic by $\mathbb{R}$ (Appendix B.1.3).

SClosed This axiom is proved by contradiction so its derivation starts by negating the invariance succedent using ⟨·⟩.

$$\begin{array}{lr}
 & \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P), P, \langle x' = f(x)\rangle \neg P \vdash \mathit{false} \\ \hline
{\scriptstyle \langle\cdot\rangle,\,\neg\mathrm{R}} & \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P), P \vdash [x' = f(x)]P \\ \hline
{\scriptstyle \forall\mathrm{R},\,\to\mathrm{R}} & \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P) \vdash \forall x\,\big(P \to [x' = f(x)]P\big)
\end{array}$$

Since formula P characterizes a closed set, every point satisfying $\neg P$ must be contained in some $\varepsilon > 0$ ball in the interior of the open set characterized by $\neg P$. Accordingly, these points are $\varepsilon > 0$ distance away from the set characterized by P and therefore, the formula $\neg P \leftrightarrow \exists \varepsilon{>}0\, \neg\mathcal{U}_\varepsilon(P)$ is provable in real arithmetic (Appendix B.1.3). This equivalence is used to rewrite the postcondition of the diamond modality antecedent before the resulting existentially quantified variable ε is commuted with the diamond modality and Skolemized with B′, ∃L and V to extract the constant assumption $\varepsilon > 0$.

$$\begin{array}{lr}
 & * \\ \hline
{\scriptstyle \langle\cdot\rangle} & [x' = f(x)]\,\mathcal{U}_\varepsilon(P), \langle x' = f(x)\rangle \neg\mathcal{U}_\varepsilon(P) \vdash \mathit{false} \\ \hline
{\scriptstyle \forall\mathrm{L},\,\to\mathrm{L}} & \delta > 0, \forall x\,\big(\mathcal{U}_\delta(P) \to [x' = f(x)]\,\mathcal{U}_\varepsilon(P)\big), P, \langle x' = f(x)\rangle \neg\mathcal{U}_\varepsilon(P) \vdash \mathit{false} \\ \hline
{\scriptstyle \forall\mathrm{L},\,\exists\mathrm{L}} & \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P), P, \varepsilon > 0, \langle x' = f(x)\rangle \neg\mathcal{U}_\varepsilon(P) \vdash \mathit{false} \\ \hline
{\scriptstyle \mathrm{V}} & \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P), P, \langle x' = f(x)\rangle \big(\varepsilon{>}0 \wedge \neg\mathcal{U}_\varepsilon(P)\big) \vdash \mathit{false} \\ \hline
{\scriptstyle \mathrm{B}',\,\exists\mathrm{L}} & \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P), P, \langle x' = f(x)\rangle \exists \varepsilon{>}0\, \neg\mathcal{U}_\varepsilon(P) \vdash \mathit{false} \\ \hline
{\scriptstyle \mathrm{M}\langle'\rangle,\,\mathbb{R}} & \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P), P, \langle x' = f(x)\rangle \neg P \vdash \mathit{false}
\end{array}$$

The stability assumption is instantiated with ε using ∀L and Skolemized with ∃L. Since the formula $P \to \mathcal{U}_\delta(P)$ is provable in real arithmetic for $\delta > 0$, the implication LHS in the antecedents is proved with ∀L, →L. The proof is completed using ⟨·⟩ since the resulting box and diamond modality antecedents are contradictory. □

Proof of Lemma 5.22. Rule $\mathrm{SLyap}^*_{\geq}$ is derived first before $\mathrm{SLyap}_{\geq}$ and $\mathrm{SLyap}_{>}$ are derived using $\mathrm{SLyap}^*_{\geq}$ as a stepping stone further below. The derivation of rule $\mathrm{SLyap}^*_{\geq}$ starts with a GLyap step. The (left) resulting premise $P \to P$ proves trivially and is not shown below. For the (right) resulting premise, the first two conjuncts under the nested quantifiers prove trivially from a cut of the second (bottom) premise of rule $\mathrm{SLyap}^*_{\geq}$. It remains to prove the final conjunct for fresh

Skolem variable W with antecedent abbreviated $\textcircled{a} \equiv \forall x\,(\overline{\mathcal{U}_\gamma(P)} \land \neg P \rightarrow \dot{V} \leq 0)$ from the premise of $\mathrm{SLyap}^*_{\geq}$.

$$\text{GLyap}\,\dfrac{\text{cut}\,\dfrac{\textcircled{a}, P \lor V{<}W \vdash [x' = f(x)\,\&\,\overline{\mathcal{U}_\gamma(P)}](P \lor V{<}W)}{\vdash \forall \varepsilon{>}0\,\exists 0{<}\gamma{\leq}\varepsilon\,\exists W \left(\begin{array}{l} \forall x\,(\partial(\mathcal{U}_\gamma(P)) \rightarrow V \geq W) \land \\ \exists 0{<}\delta{\leq}\gamma\,\forall x\,(\mathcal{U}_\delta(P) \rightarrow P \lor V{<}W) \land \\ \forall x\,\big(P \lor V{<}W \rightarrow [x' = f(x)\,\&\,\overline{\mathcal{U}_\gamma(P)}](P \lor V{<}W)\big) \end{array}\right)}}{\vdash \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P)}$$

The derivation continues using rule DCC to prove that $V < W$ is true along ODE solutions until the invariant P is entered; the first step uses an equivalent propositional rephrasing of $P \lor V < W$ as $\neg P \rightarrow V < W$. The two resulting premises are abbreviated ① and ②. They are shown and proved further below.

$$\text{cut, M}[']\,\dfrac{\text{DCC}\,\dfrac{\textcircled{1} \qquad \textcircled{2}}{\textcircled{a}, \neg P \rightarrow V < W \vdash [x' = f(x)\,\&\,\overline{\mathcal{U}_\gamma(P)}](\neg P \rightarrow V < W)}}{\textcircled{a}, P \lor V{<}W \vdash [x' = f(x)\,\&\,\overline{\mathcal{U}_\gamma(P)}](P \lor V{<}W)}$$

From premise ①, a DX step strengthens the antecedent to $V < W$ using the domain constraint $\neg P$. Rule $\mathrm{dI}_{\succcurlyeq}$ completes the proof because formula $\dot{V} \leq 0$ proves from antecedent $\textcircled{a}$ with domain $\overline{\mathcal{U}_\gamma(P)} \land \neg P$.

$$\text{DX}\,\dfrac{\mathrm{dI}_{\succcurlyeq}\,\dfrac{\dfrac{*}{\textcircled{a}, \overline{\mathcal{U}_\gamma(P)} \land \neg P \vdash \dot{V} \leq 0}}{\textcircled{a}, V < W \vdash [x' = f(x)\,\&\,\overline{\mathcal{U}_\gamma(P)} \land \neg P]V < W}}{\textcircled{a}, \neg P \rightarrow V < W \vdash [x' = f(x)\,\&\,\overline{\mathcal{U}_\gamma(P)} \land \neg P]V < W}$$

From premise ②, a dW step reduces the premise to an invariance question for formula P (since $\neg\neg P$ is propositionally equivalent to P). The DMP step weakens the domain constraint, which proves using the first premise of rule $\mathrm{SLyap}^*_{\geq}$.

$$\text{dW}\,\dfrac{\text{DMP}\,\dfrac{\dfrac{*}{P \vdash [x' = f(x)]P}}{P \vdash [x' = f(x)\,\&\,\overline{\mathcal{U}_\gamma(P)}]P}}{\vdash [x' = f(x)\,\&\,\overline{\mathcal{U}_\gamma(P)}](\neg\neg P \rightarrow [x' = f(x)\,\&\,\overline{\mathcal{U}_\gamma(P)}]\neg\neg P)}$$

Rule $\mathrm{SLyap}_{\geq}$ derives from $\mathrm{SLyap}^*_{\geq}$ because the two rules share the invariance premise on P and the latter two premises of rule $\mathrm{SLyap}_{\geq}$ imply the latter premise of $\mathrm{SLyap}^*_{\geq}$ in real arithmetic when P characterizes a compact set. The variable γ is witnessed by ε in the premise after $\mathrm{SLyap}^*_{\geq}$.

The proof of the arithmetic premise is explained below.

$$\text{SLyap}^*_{\geq}\dfrac{\forall\text{R}, \exists\text{R}\dfrac{\mathbb{R}\dfrac{*}{\varepsilon{>}0 \vdash \left(\begin{array}{l}\exists W \left(\begin{array}{l}\forall x\, (\partial(\mathcal{U}_\varepsilon(P)) \to V \geq W) \land \\ \exists 0{<}\delta{\leq}\varepsilon\, \forall x\, (\mathcal{U}_\delta(P) \land \neg P \to V < W)\end{array}\right) \land \\ \forall x\, (\overline{\mathcal{U}_\varepsilon(P)} \land \neg P \to \dot{V} \leq 0)\end{array}\right)}}{\vdash \forall \varepsilon{>}0\, \exists 0{<}\gamma{\leq}\varepsilon \left(\begin{array}{l}\exists W \left(\begin{array}{l}\forall x\, (\partial(\mathcal{U}_\gamma(P)) \to V \geq W) \land \\ \exists 0{<}\delta{\leq}\gamma\, \forall x\, (\mathcal{U}_\delta(P) \land \neg P \to V < W)\end{array}\right) \land \\ \forall x\, (\overline{\mathcal{U}_\gamma(P)} \land \neg P \to \dot{V} \leq 0)\end{array}\right)}}{\vdash \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P)}$$

The conjunct $\forall x\, (\overline{\mathcal{U}_\varepsilon(P)} \land \neg P \to \dot{V} \leq 0)$ proves logically from the right conjunct of the middle premise of rule SLyap$_\geq$. For the existentially quantified conjunct, $\exists W\, \big(\cdots\big)$, since formula P characterizes a compact set, the boundary $\partial(\mathcal{U}_\varepsilon(P))$ is also compact and therefore the continuous Lyapunov function V must attain its minimum W on that set. This W witnesses the quantifier $\exists W$ and note that $W > 0$ from the left conjunct of the middle premise of rule SLyap$_\geq$ (because $\partial(\mathcal{U}_\varepsilon(P))$ implies $\neg P$ for $\varepsilon > 0$). From the rightmost premise of rule SLyap$_\geq$, the Lyapunov function satisfies $V \leq 0$ for all points $x \in \mathbb{R}^n$ on the boundary characterized by formula ∂P. For each such point y on the boundary, by continuity, there is a radius $\delta > 0$ where points in the open ball $\|x - y\|_2 < \delta$ satisfy $V < W$ because $W > 0$. The union of all such balls over all points on the boundary is an open cover of the compact boundary which therefore has a finite subcover. The minimum radius $\delta > 0$ of balls in this finite subcover witnesses the formula $\exists 0{<}\delta{\leq}\varepsilon\, \forall x\, (\mathcal{U}_\delta(P) \land \neg P \to V{<}W)$, justifying the use of $\mathbb{R}$ (Appendix B.1.3).

Rule SLyap$_>$ is derived from rule SLyap$_\geq$ similar to the derivation of Lyap$_>$ from Lyap$_\geq$. The derivation starts with a cut of the set stability formula $\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x' = f(x), P, P)$ which proves by SLyap$_\geq$ because rules SLyap$_>$ and SLyap$_\geq$ have identical premises except for a strict inequality on $\dot{V}$.

$$\text{cut, SLyap}_{\geq}\dfrac{\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x'{=}f(x), P, P) \vdash \exists\delta{>}0\ \mathrm{Attr}^{\mathrm{P}}_{\mathrm{R}}(x'{=}f(x), \mathcal{U}_\delta(P), P)}{\vdash \mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x'{=}f(x), P, P) \land \exists\delta{>}0\ \mathrm{Attr}^{\mathrm{P}}_{\mathrm{R}}(x'{=}f(x), \mathcal{U}_\delta(P), P)}$$

The stability antecedent is instantiated with $\varepsilon = 1$; the positive constant 1 is chosen arbitrarily to obtain a neighborhood in which solutions are trapped. After Skolemization, this yields an initial disturbance $\delta > 0$ and the antecedent $\forall x\, \big(\mathcal{U}_\delta(P) \to [x' = f(x)]\, \mathcal{U}_1(P)\big)$. The succedent is witnessed with δ, and the resulting sequent is simplified with ∀R, →R, →L. Then, axiom SetSAttr simplifies the succedent using the stability antecedent.

$$\text{cut, }\exists\text{L}\dfrac{\exists\text{R}\dfrac{\forall\text{R}, \to\text{R}, \to\text{L}\dfrac{\text{SetSAttr}\dfrac{\delta{>}0, [x'{=}f(x)]\, \mathcal{U}_1(P), \mathcal{U}_\delta(P) \vdash \forall\varepsilon{>}0\, \langle x'{=}f(x)\rangle\, \mathcal{U}_\varepsilon(P)}{\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x'{=}f(x), P, P), \delta{>}0, [x'{=}f(x)]\, \mathcal{U}_1(P), \mathcal{U}_\delta(P) \vdash \mathrm{Asym}(x'{=}f(x), P)}}{\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x'{=}f(x), P, P), \delta{>}0, \forall x\, \big(\mathcal{U}_\delta(P) \to [x'{=}f(x)]\, \mathcal{U}_1(P)\big) \vdash \mathrm{Attr}^{\mathrm{P}}_{\mathrm{R}}(x'{=}f(x), \mathcal{U}_\delta(P), P)}}{\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x'{=}f(x), P, P), \delta{>}0, \forall x\, \big(\mathcal{U}_\delta(P) \to [x'{=}f(x)]\, \mathcal{U}_1(P)\big) \vdash \exists\delta{>}0\ \mathrm{Attr}^{\mathrm{P}}_{\mathrm{R}}(x'{=}f(x), \mathcal{U}_\delta(P), P)}}{\mathrm{Stab}^{\mathrm{P}}_{\mathrm{R}}(x'{=}f(x), P, P) \vdash \exists\delta{>}0\ \mathrm{Attr}^{\mathrm{P}}_{\mathrm{R}}(x'{=}f(x), \mathcal{U}_\delta(P), P)}$$

The proof of the liveness property in the open premise uses rule SP$_c$ with the choice of compact staging set $S \equiv \overline{\mathcal{U}_1(P)} \land \neg\mathcal{U}_\varepsilon(P)$ and $e \equiv V$. Note that formula S characterizes a compact set because P is compact so the neighborhood $\overline{\mathcal{U}_1(P)}$ is closed, the negation of the

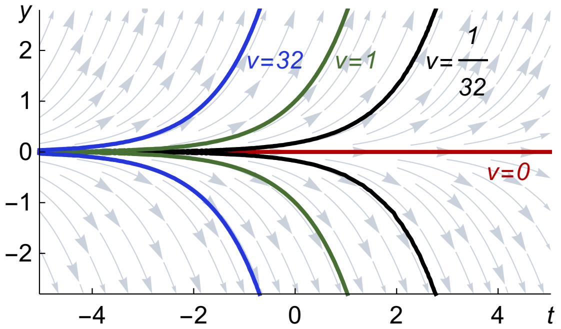

Figure C.1: An illustration of α_c and the Lyapunov function v from Counterexample C.2, with level curves (where $V = k$ for various k) shown in color.

open ε neighborhood is closed $\neg\mathcal{U}_\varepsilon(P)$, and $\overline{\mathcal{U}_1(P)}$ is bounded.

$$\text{∀R}\,\dfrac{\text{SP}_c\,\dfrac{\text{dC, dW}\,\dfrac{*}{[x' = f(x)]\,\mathcal{U}_1(P) \vdash [x' = f(x)\,\&\,\neg\mathcal{U}_\varepsilon(P)]S} \qquad \mathbb{R}\,\dfrac{*}{\varepsilon{>}0, S \vdash \dot{V} < 0}}{\delta{>}0, [x' = f(x)]\,\mathcal{U}_1(P), \mathcal{U}_\delta(P), \varepsilon{>}0 \vdash \langle x' = f(x)\rangle\,\mathcal{U}_\varepsilon(P)}}{\delta{>}0, [x' = f(x)]\,\mathcal{U}_1(P), \mathcal{U}_\delta(P) \vdash \forall\varepsilon{>}0\,\langle x' = f(x)\rangle\,\mathcal{U}_\varepsilon(P)}$$

The left premise proves with a cut dC of the antecedent and dW. The right premise proves by real arithmetic $\mathbb{R}$ using the middle premise of rule SLyap$_>$ because the antecedents imply $\neg P$. □

C.2 Counterexamples

This appendix provides counterexamples for the soundness issues highlighted in Sections 5.3 and 5.4. The first counterexample illustrates the need to assume compactness, i.e., formula P is closed *and bounded* in rule SLyap$_\geq$. The remark after [89, Definition 8.1] suggests that the following variant of SLyap$_\geq$ is sound for formulas P that characterize a closed, invariant set:

$$\text{SLyap}_{\geq}\text{↯}\ \dfrac{P \vdash V = 0 \qquad \neg P \vdash V > 0 \land \dot{V} \leq 0}{\vdash \text{Stab}^{\text{P}}_{\text{R}}(x' = f(x), P, P)}$$

The rule SLyap$_\geq$↯ is unsound (indicated by ↯); indeed, the rule SLyap$_\geq$ from Lemma 5.22 is also unsound if the assumption that formula P characterizes a bounded set is omitted.

Counterexample C.2. Consider the ODE $\alpha_c \equiv y' = y, t' = 1$ and the formula $P \equiv y = 0$ which characterizes a closed invariant set of α_c that is *not bounded.* The Lyapunov function $V = y^2 \exp(-2t)$, satisfies all of the premises of rule SLyap$_\geq$↯ because $V = 0$ when $y = 0$, $V > 0$ for $y \neq 0$, and $\dot{V} = 0$. However, P is not stable for ODE α_c, as can be seen from Fig. C.1. The norm of the right-maximal solution from all initial states that satisfy $y \neq 0$ approach ∞.

This counterexample also illustrates the importance of the boundedness assumption for formula P in Lemma 5.22 for rule SLyap$_\geq$ since all other premises of the rule are satisfied by the above example. ↯

The second counterexample below shows that rule Lyap$_{\geq}$ needs the premise $V(0) = 0$. This premise is unsoundly omitted from the arithmetical conditions in [3, Equation 1].

Counterexample C.3. Consider the ODE $y' = y$ with solution $y(t) = y_0 \exp(t)$ from initial value $y(0) = y_0$. For all perturbed initial states $y_0 \neq 0$, $\|y(t)\|_2$ approaches ∞ as $t \to \infty$ so this ODE is not stable (nor attractive). However, the Lyapunov function $V = 1$ trivially satisfies all of the premises of rule Lyap$_{\geq}$ except the omitted premise $V(0) = 0$. $\lightning$

Appendix D

Appendix: Stability for Switched Systems

D.1 Switched System Models and Stability Proof Rules

This appendix provides proofs for all results that were omitted from Sections 6.2 and 6.3. For ease of reference, this appendix is organized into two sections, corresponding to proofs for Section 6.2 and Section 6.3 respectively. All derived proof rules and axioms used in this appendix are explained earlier in the thesis and in the preceding appendices.

D.1.1 Proofs for Switched Systems as Hybrid Programs

This section contains adequacy proofs for hybrid program models of various switching mechanisms defined in Section 6.2.

Proof of Theorem 6.3. Axiom SAI$_{\text{state}}$ derives immediately from Inv$_{\text{state}}$ by equivalently rewriting its RHS using the derived equivalent characterization of ODE invariants SAI& from Theorem A.11 (page 212). Both directions of axiom Inv$_{\text{state}}$ are derived separately.

"$\leftarrow$" The (easier) "$\leftarrow$" direction uses rule loop to prove that P is a loop invariant of $\alpha_{\texttt{state}}$. The antecedent assumption $\bigwedge_{p\in\mathcal{P}} \forall x\,(P \to [x' = f_p(x)\,\&\,Q_p]P)$ is constant for $\alpha_{\texttt{state}}$, so it is soundly kept across the use of rule loop. The subsequent $[\cup]$, $\land$R step unfolds the nondeterministic choice in $\alpha_{\texttt{state}}$'s loop body, yielding a premise for each ODE in $\mathcal{P}$. These premises are indexed by $p \in \mathcal{P}$ below and they are all proved propositionally from the antecedent assumption.

$$
{\scriptstyle \forall\text{R},\,\to\text{R}}\;\dfrac{{\scriptstyle \text{loop}}\;\dfrac{{\scriptstyle [\cup],\,\land\text{R}}\;\dfrac{{\scriptstyle \land\text{L},\,\forall\text{L},\,\to\text{L}}\;\dfrac{*}{\bigwedge_{p\in\mathcal{P}} \forall x\,(P \to [x' = f_p(x)\,\&\,Q_p]P), P \vdash [x' = f_p(x)\,\&\,Q_p]P \quad (p \in \mathcal{P})}}{\bigwedge_{p\in\mathcal{P}} \forall x\,(P \to [x' = f_p(x)\,\&\,Q_p]P), P \vdash [\bigcup_{p\in\mathcal{P}} x' = f_p(x)\,\&\,Q_p]P}}{\bigwedge_{p\in\mathcal{P}} \forall x\,(P \to [x' = f_p(x)\,\&\,Q_p]P), P \vdash [\alpha_{\texttt{state}}]P}}{\bigwedge_{p\in\mathcal{P}} \forall x\,(P \to [x' = f_p(x)\,\&\,Q_p]P) \vdash \forall x\,(P \to [\alpha_{\texttt{state}}]P)}
$$

"$\rightarrow$" The "$\rightarrow$" direction shows that a run of ODE $x' = f_p(x) \,\&\, Q_p$, $p \in \mathcal{P}$ must also be a run of $\alpha_{\texttt{state}}$, so if formula P is true for all runs of $\alpha_{\texttt{state}}$, it must also be true for all runs of the constituent ODEs. The derivation starts with logical unfolding steps where the resulting premises are indexed by $p \in \mathcal{P}$ below.

$$\begin{array}{r c}
 & [\alpha_{\texttt{state}}]P \vdash [x' = f_p(x) \,\&\, Q_p]P \\ \cline{2-2}
\scriptstyle \forall\text{L}, \rightarrow\text{L} & \forall x\,(P \rightarrow [\alpha_{\texttt{state}}]P), P \vdash [x' = f_p(x) \,\&\, Q_p]P \quad (p \in \mathcal{P}) \\ \cline{2-2}
\scriptstyle \wedge\text{R}, \forall\text{R}, \rightarrow\text{R} & \forall x\,(P \rightarrow [\alpha_{\texttt{state}}]P) \vdash \bigwedge_{p \in \mathcal{P}} \forall x\,(P \rightarrow [x' = f_p(x) \,\&\, Q_p]P)
\end{array}$$

Next, axiom [*] unfolds the loop in the antecedent before axiom $[\cup]$ chooses the branch corresponding to $p \in \mathcal{P}$ in the loop body.

$$\begin{array}{r r}
 & * \\ \cline{2-2}
\scriptstyle [*], \wedge\text{L} & [\alpha_{\texttt{state}}]P \vdash P \\ \cline{2-2}
\scriptstyle \text{M}[\cdot] & [x' = f_p(x) \,\&\, Q_p][\alpha_{\texttt{state}}]P \vdash [x' = f_p(x) \,\&\, Q_p]P \\ \cline{2-2}
\scriptstyle [\cup], \wedge\text{L} & [\bigcup_{p \in \mathcal{P}} x' = f_p(x) \,\&\, Q_p][\alpha_{\texttt{state}}]P \vdash [x' = f_p(x) \,\&\, Q_p]P \\ \cline{2-2}
\scriptstyle [*], \wedge\text{L} & [\alpha_{\texttt{state}}]P \vdash [x' = f_p(x) \,\&\, Q_p]P
\end{array}$$

The derivation is completed using rule $\text{M}[\cdot]$ to monotonically strengthen the postcondition, then unfolding the resulting antecedent with axiom [*]. □

Proof of Proposition 6.4. The proof is similar to Proposition 6.2 but with fresh auxiliary variable u used to control the switching signal. The switched system obeys the guards $G_{p,q}$ along a solution if, for any switch from mode p, q, the system state in mode p satisfies a sequence of guard formulas $G_{p_0,p_1}, G_{p_1,p_2}, \ldots, G_{p_{n-1},p_n}$ with $p = p_0$ and $q = p_n$. Intuitively, this means the system can take a sequence of zero-time jumps along a sequence of modes and guard conditions to switch from mode p initially to mode q at the end. Both directions of the proposition are proved separately for an initial state $\omega \in \mathbb{R}^n$. The initialization program α_i sets u to a choice of mode $p \in \mathcal{P}$ initially but leaves the state variables x unchanged.

"$\Rightarrow$" Suppose $(\omega, \nu) \in [\![\alpha_{\texttt{guard}}]\!]$. By the semantics of dL programs, there is a sequence of states $\omega = \omega_0, \omega_1, \ldots, \omega_n = \nu$ for some $n \geq 0$ and for each $1 \leq i \leq n$, there is a run of the loop body with $(\omega_{i-1}, \omega_i) \in [\![\alpha_u; \alpha_p]\!]$. Unfolding the sequential composition,[1] by the semantics of the controller program α_u, state ω_{i-1} satisfies $u = p$ for some $p \in \mathcal{P}$ and the guard condition $G_{p,q}$ for some $q \in \mathcal{P}$, and there is an intermediate state γ obtained from ω_{i-1} by setting the value of variable u to q. The plant program α_p then runs the selected ODE for mode q from state γ for some time $\zeta_i \geq 0$ to reach state ω_i. Thus, every state ω_i is associated with its chosen mode p_i and time $\zeta_i \geq 0$ for which the ODE for mode p_i is followed in that switching step. The state ω_{i-1} satisfies guard G_{p_{i-1},p_i} for $i \geq 1$.

Similar to the proof of Proposition 6.2, remove from all sequences the chattering indexes $1 \leq i \leq n$ with $\zeta_i = 0$. This yields new sequences $(\tilde{\omega}_0, \tilde{\omega}_1, \ldots, \tilde{\omega}_m)$, $(\tilde{\zeta}_1, \ldots, \tilde{\zeta}_m)$, and $(\tilde{p}_1, \ldots, \tilde{p}_m)$ where $\tilde{\zeta}_i > 0$. Since no continuous evolution occurs when $\zeta_i = 0$, for each resulting pairs of adjacent states $(\tilde{\omega}_{i-1}, \tilde{\omega}_i)$ with $i \geq 1$, the state $\tilde{\omega}_{i-1}$ satisfies the required

[1] In case the controller α_u leaves the mode unchanged with $u := u$, the adjacent continuous evolutions belong to the same mode and can be uniquely concatenated [33, Theorem 1.2]. Thus, assume without loss of generality that the controller always performs a (guarded) switch from mode $p \in \mathcal{P}$ to some mode $q \in \mathcal{P}$.

sequence of guard formulas between mode $\tilde{p}_{i-1}$ and mode $\tilde{p}_i$ by following the (removed) chattering indexes. Consider the switching signal σ with switching times $\tau_i = \sum_{j=1}^{i} \tilde{\zeta}_j$ for $1 \leq i < m$, $\tau_m = \sum_{j=1}^{m} \tilde{\zeta}_j + 1$, and $\tau_i = \tau_{i-1} + 1$ for $i > m$, so $\tau_1 < \tau_2 < \dots$ and $\tau_i \to \infty$. Furthermore, extend the sequence of switching choices with $\tilde{p}_i = \tilde{p}_m$ for $i > m$.[2] By construction using (6.2), σ is well-defined and the solution φ associated with σ from ω reaches ν at time $\sum_{j=1}^{m} \tilde{\zeta}_j$ and it obeys the domains $Q_{\tilde{p}_i}$ until that time.

"$\Leftarrow$" Let σ be a switching signal and $\varphi : [0, \zeta) \to \mathbb{R}^n$ be the associated switched system solution from ω. Suppose that the switched system reaches $\varphi(t)$ for $t \in [0, \zeta)$ while obeying the domains Q_p and guards $G_{p,q}$ for modes $p, q \in \mathcal{P}$. To show $(\omega, \varphi(t)) \in [\![\alpha_{\mathtt{guard}}]\!]$, by the semantics of dL loops, it suffices to construct a sequence of states $\omega = \omega_0, \omega_1, \dots, \omega_n$ for some finite n, with $\omega_n = \varphi(t)$, and $(\omega_{i-1}, \omega_i) \in [\![(\alpha_u; \alpha_p)^*]\!]$ for $1 \leq i \leq n$, because loop $(\alpha_u; \alpha_p)^*$ unfolds to a nested self-loop $((\alpha_u; \alpha_p)^*)^*$.

By (6.2), σ is equivalently defined by a sequence of switching times $\tau_0 < \tau_1 < \tau_2 < \dots$ and a sequence of switching choices $p_1, p_2, \dots$, where $p_i \in \mathcal{P}$. Let τ_n be the first switching time such that $t \leq \tau_n$; the index n exists since $\tau_i \to \infty$. Define the state sequence $\omega_i = \varphi(\tau_i)$ for $0 \leq i < n$ and $\omega_n = \varphi(t)$. Note that $\omega_0 = \omega$ by definition of $\varphi(0)$. It suffices to show $(\omega_{i-1}, \omega_i) \in [\![(\alpha_u; \alpha_p)^*]\!]$ for $1 \leq i \leq n$. By assumption, the switched system takes a sequence of zero-time jumps along a sequence of modes and guard conditions when it switches from mode p_{i-1} to mode p_i. This sequence is simulated by program $(\alpha_u; \alpha_p)^*$ by unfolding the loop, switching to the respective mode(s) by α_u in each iteration and then running the chosen ODE in α_p for 0 time except for the last iteration (mode p_i) where the ODE $x' = f_{p_i}(x)$ is followed continuously until state ω_i is reached. □

Proof of Proposition 5.27. The proof is similar to Proposition 6.4 by reasoning about the effect of the fresh auxiliary variable u used to model the piecewise constant function $u(t)$. The function $u(t)$ is viewed as a switching signal that prescribes switching choices on each interval with associated solution generated according to (6.2). Both directions of the proposition are proved separately for an initial state $\omega \in \mathbb{R}^n$.

"$\Rightarrow$" Suppose $(\omega, \nu) \in [\![\alpha_{\mathtt{piece}}]\!]$. By the semantics of dL programs, there is a sequence of states $\omega = \omega_0, \omega_1, \dots, \omega_n = \nu$ for some $n \geq 0$ and for each $1 \leq i \leq n$, there is a run of the loop body with $(\omega_{i-1}, \omega_i) \in [\![u := *; ?\, \|u\|_\infty \leq \Delta; x' = f(x, u)]\!]$. Unfolding the dL hybrid program semantics, there are real value(s) U_i with $\|U_i\|_\infty \leq \Delta$, such that state ω_i is reached from state ω_{i-1} by following the ODE $x' = f(x, U_i)$ for time $\zeta_i \geq 0$. Define the sequence of switching choices as $p_i = U_i$ for $1 \leq i \leq n$.

Following the proof of Proposition 6.4, remove from all sequences the chattering indexes $1 \leq i \leq n$ with $\zeta_i = 0$. This yields new sequences $(\tilde{\omega}_0, \tilde{\omega}_1, \dots, \tilde{\omega}_m)$, $(\tilde{\zeta}_1, \dots, \tilde{\zeta}_m)$, and $(\tilde{p}_1, \dots, \tilde{p}_m)$ where $\tilde{\zeta}_i > 0$. Consider the associated piecewise constant function $u(t)$ with switching times $\tau_i = \sum_{j=1}^{i} \tilde{\zeta}_j$ for $1 \leq i < m$ and $\tau_i = \tau_{i-1} + 1$ for $i \geq m$, so $\tau_1 < \tau_2 < \dots$ and $\tau_i \to \infty$, together with the extended sequence of switching choices with $\tilde{p}_i = \tilde{p}_m$ for $i > m$. By construction using (6.2), $u(t)$ is a bounded piecewise constant function and the solution φ associated with $u(t)$ from ω reaches ν at time $\sum_{j=1}^{m} \tilde{\zeta}_j$.

[2] The choice of mode switches and guards are irrelevant for times $t > \sum_{j=1}^{m} \tilde{\zeta}_j$ since they only need to be obeyed until time $\sum_{j=1}^{m} \tilde{\zeta}_j$.

“$\Leftarrow$” Let $u(t)$ be a bounded piecewise constant function with bound $\|u(t)\|_\infty \leq \Delta$ and let $\varphi : [0, \zeta) \to \mathbb{R}^n$ be the associated switched system solution from ω. Suppose that the solution reaches $\varphi(t)$ for $t \in [0, \zeta)$. To show $(\omega, \varphi(t)) \in [\![\alpha_{\texttt{piece}}]\!]$, by the semantics of dL loops, it suffices to construct a sequence of states $\omega = \omega_0, \omega_1, \ldots, \omega_n$ for some finite n, with $\omega_n = \varphi(t)$, and $(\omega_{i-1}, \omega_i) \in [\![u := *; ?\, \|u\|_\infty \leq \Delta; x' = f(x, u)]\!]$ for $1 \leq i \leq n$.
By (6.2), $u(t)$ is equivalently defined by a sequence of switching times $\tau_0 < \tau_1 < \tau_2 < \ldots$ and a sequence of switching choices $p_1, p_2, \ldots$, where $u(t) = p_i$ for all $\tau_{i-1} \leq t < \tau_i$ for $1 \leq i$. Let τ_n be the first switching time such that $t \leq \tau_n$; the index n exists since $\tau_i \to \infty$. Define the state sequence $\omega_i = \varphi(\tau_i)$ for $0 \leq i < n$ and $\omega_n = \varphi(t)$. Note that $\omega_0 = \omega$ by definition of $\varphi(0)$. It suffices to show $(\omega_{i-1}, \omega_i) \in [\![u := *; ?\, \|u\|_\infty \leq \Delta; x' = f(x, u)]\!]$ for $1 \leq i \leq n$. By assumption, $\|p_i\|_\infty = \|u(t)\|_\infty \leq \Delta$, so choose value p_i for variable u in $u := *$ which passes the test $?\, \|u\|_\infty \leq \Delta$. By construction of φ, state ω_i is reached from ω_{i-1} by following the solution to ODE $x' = f(x, p_i)$. □

Proof of Proposition 6.5. The proof is similar to Proposition 6.4 but with an additional fresh auxiliary variable τ which tracks the time spent in each mode. Both directions of the proposition are proved separately for an initial state $\omega \in \mathbb{R}^n$.

“$\Rightarrow$” Suppose $(\omega, \nu) \in [\![\alpha_{\texttt{time}}]\!]$. By the semantics of dL programs, there is a sequence of states $\omega = \omega_0, \omega_1, \ldots, \omega_n = \nu$ for some $n \geq 0$ and for each $1 \leq i \leq n$, there is a run of the loop body with $(\omega_{i-1}, \omega_i) \in [\![\alpha_u; \alpha_p]\!]$. Unfolding the sequential composition,[3] by the semantics of the controller program α_u, state ω_{i-1} satisfies $u = p$ for some $p \in \mathcal{P}$ and the minimum dwell time condition $\theta_{p,q} \leq \tau$ for some $q \in \mathcal{P}$, and there is an intermediate state γ obtained from ω_{i-1} by setting the mode variable u to q and resetting the timer τ to 0. The plant program α_p then runs the selected ODE for mode q from state γ for some maximum time $\Theta_q \geq \zeta_i \geq 0$ to reach state ω_i, where $\Theta_q > 0$ by assumption. Thus, every state ω_i is associated with its chosen mode p_i and time $\Theta_{p_i} \geq \zeta_i \geq 0$ for which the ODE for mode p_i is followed in that switching step. Moreover, minimum dwell time $\theta_{p_{i-1}, p_i} \leq \zeta_{i-1}$ must be spent in mode p_{i-1} before switching to mode p_i.
Similar to the proof of Proposition 6.2, remove from all sequences the chattering indexes $1 \leq i \leq n$ with $\zeta_i = 0$. This yields new sequences $(\tilde{\omega}_0, \tilde{\omega}_1, \ldots, \tilde{\omega}_m)$, $(\tilde{\zeta}_1, \ldots, \tilde{\zeta}_m)$, and $(\tilde{p}_1, \ldots, \tilde{p}_m)$ where $\tilde{\zeta}_i > 0$. Consider the switching signal σ with switching times $\tau_i = \sum_{j=1}^{i} \tilde{\zeta}_j$ for $1 \leq i < m$ and $\tau_i = \tau_{i-1} + \Theta_{\tilde{p}_m}$ for $i \geq m$, so $\tau_1 < \tau_2 < \ldots$ and $\tau_i \to \infty$. Furthermore, extend the sequence of switching choices with $\tilde{p}_i = \tilde{p}_m$ for $i > m$. The switching times satisfy the maximum dwell times $\tau_i - \tau_{i-1} \leq \Theta_{\tilde{p}_i}$ and the minimum dwell times $\theta_{\tilde{p}_i, \tilde{p}_{i+1}} \leq \tau_i - \tau_{i-1}$ for $i \geq 1$ until time $\sum_{j=1}^{m} \tilde{\zeta}_j$. By construction using (6.2), σ is well-defined and the solution φ associated with σ from ω reaches ν at time $\sum_{j=1}^{m} \tilde{\zeta}_j$.

“$\Leftarrow$” Let σ be a switching signal and $\varphi : [0, \zeta) \to \mathbb{R}^n$ be the associated switched system solution from ω. Suppose that the switched system reaches $\varphi(t)$ for $t \in [0, \zeta)$ while spending at least time $\tau_{p,q}$ in mode $p \in \mathcal{P}$ before switching to mode $q \in \mathcal{P}$ and spends at most $\Theta_p > 0$ time in mode $p \in \mathcal{P}$. To show $(\omega, \varphi(t)) \in [\![\alpha_{\texttt{time}}]\!]$, by the semantics of dL loops, it suffices

[3] Similar to the proof of Proposition 6.4, assume without loss of generality that the controller always performs a switch since adjacent continuous evolutions can be uniquely concatenated whenever the controller leaves the mode unchanged with $u := u$.

to construct a sequence of states $\omega = \omega_0, \omega_1, \ldots, \omega_n$ for some finite n, with $\omega_n = \varphi(t)$, and $(\omega_{i-1}, \omega_i) \in [\![\alpha_u; \alpha_p]\!]$ for $1 \leq i \leq n$.

By (6.2), σ is equivalently defined by a sequence of switching times $\tau_0 < \tau_1 < \tau_2 < \ldots$ and a sequence of switching choices $p_1, p_2, \ldots$, where $p_i \in \mathcal{P}$. Let τ_n be the first switching time such that $t \leq \tau_n$; the index n exists since $\tau_i \to \infty$. Define the state sequence $\omega_i = \varphi(\tau_i)$ for $0 \leq i < n$ and $\omega_n = \varphi(t)$. Note that $\omega_0 = \omega$ by definition of $\varphi(0)$. It suffices to show $(\omega_{i-1}, \omega_i) \in [\![\alpha_u; \alpha_p]\!]$ for $1 \leq i \leq n$. By assumption, the minimum dwell time constraint in α_u is satisfied when switching from mode p_{i-1} to mode p_i (for $i = 1$, skip with $u := u$). The maximum dwell time constraint on mode p_i means that, by construction of φ, ω_i is reached from ω_{i-1} by following the solution to ODE $x' = f_{p_i}(x), \tau' = 1 \,\&\, \tau \leq \Theta_{p_i}$. □

D.1.2 Proofs for Switched System Stability

This section derives stability proof rules for various switching mechanisms using the loop invariants explained and motivated in Section 6.3. To improve readability in the proofs, formula and premises are often abbreviated, e.g., with ⓐ, ①. To avoid confusion, the scope of these abbreviations always extend to the end of each **paragraph** label, i.e., the abbreviations used in the **Stability** proofs should not be confused with those used in the **Pre-Attractivity** proofs.

Proof of Corollary 6.9. Rule CLF is an instance of rule MLF from Corollary 6.10 where the Lyapunov functions for all modes $p \in \mathcal{P}$ are chosen identically with $V_p = V$. Nevertheless, a full derivation of CLF is given here as it provides the main building blocks used in later derivations. The stability and pre-attractivity conjuncts of $\mathrm{UGpAS}(\alpha_{\mathtt{state}})$ are proved separately with $\wedge$R:

$$
{\scriptstyle \wedge\mathrm{R}}\,\dfrac{\vdash \mathrm{UStab}(\alpha_{\mathtt{state}}) \qquad \vdash \mathrm{UGpAttr}(\alpha_{\mathtt{state}})}{\vdash \mathrm{UGpAS}(\alpha_{\mathtt{state}})}
$$

Stability. The derivation for stability begins by Skolemizing the succedent with $\forall$R, $\to$R, followed by two arithmetic cuts which are justified as follows (recall convention from Appendix B.1.3). For any $\varepsilon > 0$, the Lyapunov function V attains a minimum value on the compact set characterized by $\|x\|_2 = \varepsilon$. From the first (topmost) premise of rule CLF, this minimum is attained away from the origin so it is positive, which proves the first cut of formula $\exists W{>}0\,\text{ⓐ}$ where $\text{ⓐ} \equiv \forall x\, (\|x\|_2 = \varepsilon \to V \geq W)$. After Skolemizing W with $\exists$L, the premise $V(0) = 0$ implies, by continuity of dL term semantics [142], that the sublevel set characterized by $V < W$ with $W > 0$ (see Fig. 6.3) contains a sufficiently small δ ball around the origin (with $\delta \leq \varepsilon$). This proves the second arithmetic cut with the formula $\exists \delta\, (0 < \delta \leq \varepsilon \wedge \text{ⓑ})$ where $\text{ⓑ} \equiv \forall x\, (\|x\|_2 < \delta \to V < W)$. After both cuts, the Skolemized δ from the antecedent is used to witness the succedent δ by $\exists$R.

$$
{\scriptstyle \forall\mathrm{R},\,\to\mathrm{R}}\,\dfrac{{\scriptstyle \mathrm{cut},\,\mathbb{R},\,\exists\mathrm{L}}\,\dfrac{{\scriptstyle \mathrm{cut},\,\mathbb{R},\,\exists\mathrm{L}}\,\dfrac{{\scriptstyle \exists\mathrm{R}}\,\dfrac{\text{ⓐ}, \delta \leq \varepsilon, \text{ⓑ} \vdash \forall x \left(\|x\|_2 < \delta \to [\alpha_{\mathtt{state}}]\, \|x\|_2 < \varepsilon \right)}{\text{ⓐ}, 0 < \delta \leq \varepsilon, \text{ⓑ} \vdash \exists \delta{>}0\, \forall x \left(\|x\|_2 < \delta \to [\alpha_{\mathtt{state}}]\, \|x\|_2 < \varepsilon \right)}}{\varepsilon > 0, W > 0, \text{ⓐ} \vdash \exists \delta{>}0\, \forall x \left(\|x\|_2 < \delta \to [\alpha_{\mathtt{state}}]\, \|x\|_2 < \varepsilon \right)}}{\varepsilon > 0 \vdash \exists \delta{>}0\, \forall x \left(\|x\|_2 < \delta \to [\alpha_{\mathtt{state}}]\, \|x\|_2 < \varepsilon \right)}}{\vdash \mathrm{UStab}(\alpha_{\mathtt{state}})}
$$

The derivation continues from the open premise by Skolemizing the succedent with ∀R, →R and proving the LHS of the implication in ⓑ with ∀L, →L. Then, the loop rule is used with the stability loop invariant $\mathit{Inv}_s \equiv \|x\|_2 < \varepsilon \land V < W$. This results in three premises: ① which shows that the invariant is implied by the initial antecedent assumptions; ② the crucial premise, which shows that the invariant Inv_s is preserved across the loop body of $\alpha_{\mathtt{state}}$; and ③ which shows that the invariant implies the postcondition. These premises are shown and proved further below.

$$\forall\text{R}, \to\text{R}\;\dfrac{\forall\text{L}, \to\text{L}\;\dfrac{\text{loop}\;\dfrac{\textcircled{1} \qquad \textcircled{2} \quad \textcircled{3}}{\textcircled{a}, \delta \le \varepsilon, \|x\|_2 < \delta, V < W \vdash [\alpha_{\mathtt{state}}]\ \|x\|_2 < \varepsilon}}{\textcircled{a}, \delta \le \varepsilon, \textcircled{b}, \|x\|_2 < \delta \vdash [\alpha_{\mathtt{state}}]\ \|x\|_2 < \varepsilon}}{\textcircled{a}, \delta \le \varepsilon, \textcircled{b} \vdash \forall x\, \big(\, \|x\|_2 < \delta \to [\alpha_{\mathtt{state}}]\ \|x\|_2 < \varepsilon\big)}$$

Premise ① proves by $\mathbb{R}$ from the antecedents using the inequalities $\|x\|_2 < \delta$ and $\delta \le \varepsilon$.

$$\mathbb{R}\;\dfrac{*}{\delta \le \varepsilon, \|x\|_2 < \delta, V < W \vdash \mathit{Inv}_s}$$

Premise ③ proves trivially since the postcondition $\|x\|_2 < \varepsilon$ is part of the loop invariant:

$$\mathbb{R}\;\dfrac{*}{\mathit{Inv}_s \vdash \|x\|_2 < \varepsilon}$$

The derivation continues from premise ② by unfolding the loop body of $\alpha_{\mathtt{state}}$ with $[\cup]$, ∧R. This results in one premise for each switching choice $p \in \mathcal{P}$, indexed below by p.

$$[\cup], \land\text{R}\;\dfrac{\textcircled{a}, \mathit{Inv}_s \vdash [x' = f_p(x) \,\&\, Q_p]\mathit{Inv}_s \qquad (p \in \mathcal{P})}{\textcircled{a}, \mathit{Inv}_s \vdash [\bigcup_{p \in \mathcal{P}} x' = f_p(x) \,\&\, Q_p]\mathit{Inv}_s}$$

Each of these $p \in \mathcal{P}$ premises is an ODE invariance question, which completely reduces to an arithmetic question by proof in dL (Chapter 3). The derivation below shows how to directly derive arithmetical conditions on V from these premises. The right conjunct of Inv_s, $V < W$, is added to the domain constraint with a dC step; the cut premise is labeled ④ and proved below. A subsequent dC step adds $\|x\|_2 \ne \varepsilon$ to the domain constraint using the contrapositive of antecedent ⓐ and the derivation is completed with rule Barr since the resulting $\|x\|_2 = \varepsilon$ assumption in its premise contradicts the domain constraint $\|x\|_2 \ne \varepsilon$.

$$\text{dC}\;\dfrac{\text{dC}\;\dfrac{\text{Barr}\;\dfrac{\mathbb{R}\;\dfrac{*}{\|x\|_2 \ne \varepsilon, \|x\|_2 = \varepsilon \vdash \mathit{false}}}{\|x\|_2 < \varepsilon \vdash [x' = f_p(x) \,\&\, Q_p \land V < W \land \|x\|_2 \ne \varepsilon]\ \|x\|_2 < \varepsilon}}{\textcircled{a}, \|x\|_2 < \varepsilon \vdash [x' = f_p(x) \,\&\, Q_p \land V < W]\ \|x\|_2 < \varepsilon \qquad \textcircled{4}}}{\textcircled{a}, \mathit{Inv}_s \vdash [x' = f_p(x) \,\&\, Q_p]\mathit{Inv}_s}$$

The derivation from ④ is completed with a $\text{dI}_{\succcurlyeq}$ step whose resulting arithmetic is implied by the bottom premise of rule CLF.

$$\text{dI}_{\succcurlyeq}\;\dfrac{\mathbb{R}\;\dfrac{*}{Q_p \vdash \mathcal{L}_{f_p}(V) \le 0}}{V < W \vdash [x' = f_p(x) \,\&\, Q_p]V < W}$$

Pre-Attractivity. The derivation for pre-attractivity begins by Skolemizing the succedent δ, ε with ∀R, →R, followed by a series of arithmetic cuts which are justified stepwise. First, the Lyapunov function V is bounded above on the ball characterized by $\|x\|_2 < \delta$, which justifies a cut of the formula $\exists W{>}0\,\textcircled{a}$ with $\textcircled{a} \equiv \forall x \left(\|x\|_2 < \delta \rightarrow V < W \right)$. After Skolemizing the upper bound W, note that the set characterized by formula $V \leq W$ is compact by radial unboundedness (middle premise of rule CLF). Therefore, the set characterized by formula $V \leq W \land \|x\|_2 \geq \varepsilon$ is an intersection of a compact and closed set, which is itself compact. Thus, V attains a minimum U on that set which is positive by the first (topmost) premise. This justifies an arithmetic cut of the formula $\exists U{>}0\,\textcircled{b}$ with $\textcircled{b} \equiv \forall x \, (V \leq W \land \|x\|_2 \geq \varepsilon \rightarrow V \geq U)$, where U is subsequently Skolemized with ∃L. The steps are shown below, with the box modality in $\mathrm{UGpAttr}(\alpha_{\mathtt{state}})$ temporarily hidden with $\ldots$ as it is not relevant for this part of the derivation.

$${\scriptstyle \forall\mathrm{R},\,\rightarrow\mathrm{R}}\;\dfrac{{\scriptstyle \mathrm{cut},\,\mathbb{R},\,\exists\mathrm{L}}\;\dfrac{{\scriptstyle \mathrm{cut},\,\mathbb{R},\,\exists\mathrm{L}}\;\dfrac{\varepsilon > 0, W > 0, \textcircled{a}, U > 0, \textcircled{b} \vdash \exists T{\geq}0 \forall x \left(\|x\|_2 < \delta \rightarrow \ldots \right)}{\varepsilon > 0, W > 0, \textcircled{a} \vdash \exists T{\geq}0 \forall x \left(\|x\|_2 < \delta \rightarrow \ldots \right)}}{\varepsilon > 0 \vdash \exists T{\geq}0 \forall x \left(\|x\|_2 < \delta \rightarrow \ldots \right)}}{\vdash \mathrm{UGpAttr}(\alpha_{\mathtt{state}})}$$

Intuitively (see Fig. 6.3) the next arithmetic steps syntactically determine $T \geq 0$ such that the value of V decreases from W to U along all switching trajectories within time T. Consider the set characterized by formula $Q_p \land U \leq V \leq W$, which is the set of states (before reaching $V < U$) where switching to ODE $x' = f_p(x) \,\&\, Q_p (p \in \mathcal{P})$ is possible. From the third (bottom) premise of rule CLF, $\mathcal{L}_{f_p}(V)$ is negative on the set characterized by the formula $\overline{Q_p} \land U \leq V \leq W$ because conjunct $U \leq V$ bounds the set away from the origin as $U > 0$. Using radial unboundedness again, $V \leq W$ is compact, so the set characterized by $\overline{Q_p} \land U \leq V \leq W$ is an intersection of closed sets and compact sets which is therefore compact. Accordingly, $\mathcal{L}_{f_p}(V)$ attains a maximum value $k_p < 0$ on that set, which justifies the following arithmetic cut, where the bound $k < 0$ is chosen uniformly across all choices of p, e.g., as the maximum over all k_p for $p \in \mathcal{P}$:

$$\exists k < 0 \; \underbrace{\bigwedge_{p \in \mathcal{P}} \forall x \left(\overline{Q_p} \land U \leq V \leq W \rightarrow \mathcal{L}_{f_p}(V) \leq k \right)}_{\textcircled{c}}$$

After Skolemizing k, it suffices to pick $T \geq 0$ for the succedent such that $W + kT \leq U$. Such a T always exists since $k < 0$.

$${\scriptstyle \mathrm{cut},\,\mathbb{R},\,\exists\mathrm{L}}\;\dfrac{{\scriptstyle \exists\mathrm{R}}\;\dfrac{\textcircled{a}, \textcircled{b}, k < 0, \textcircled{c}, W + kT \leq U \vdash \forall x \left(\|x\|_2 < \delta \rightarrow \ldots \right)}{\varepsilon > 0, W > 0, \textcircled{a}, U > 0, \textcircled{b}, k < 0, \textcircled{c} \vdash \exists T{\geq}0 \forall x \left(\|x\|_2 < \delta \rightarrow \ldots \right)}}{\varepsilon > 0, W > 0, \textcircled{a}, U > 0, \textcircled{b} \vdash \exists T{\geq}0 \forall x \left(\|x\|_2 < \delta \rightarrow \ldots \right)}$$

The derivation continues by Skolemizing with ∀R, →R and proving the LHS of the implication in ⓐ with ∀L, →L. The assignment $t := 0$ is unfolded with axioms [;], [:=], then the loop rule is used with the pre-attractivity loop invariant $\mathit{Inv}_a \equiv V < W \land (V \geq U \rightarrow V < W + kt)$. Similar to the stability derivation, this results in three premises, where the crucial premise ② requires showing that Inv_a is preserved across the loop body, while the other premises are labeled ① and

③ (all three premises are shown further below).

$$\text{∀R, →R}\;\dfrac{\text{∀L, →L}\;\dfrac{\text{[;], [:=]}\;\dfrac{\text{loop}\;\dfrac{①\qquad ②\quad ③}{V < W, \text{ⓑ}, k < 0, \text{ⓒ}, W + kT \leq U, t = 0 \vdash [\alpha_{\mathtt{state}}, t' = 1]\ldots}}{V < W, \text{ⓑ}, k < 0, \text{ⓒ}, W + kT \leq U \vdash [t := 0; \alpha_{\mathtt{state}}, t' = 1]\ldots}}{\text{ⓐ}, \text{ⓑ}, k < 0, \text{ⓒ}, W + kT \leq U, \|x\|_2 < \delta \vdash [t := 0; \alpha_{\mathtt{state}}, t' = 1]\ldots}}{\text{ⓐ}, \text{ⓑ}, k < 0, \text{ⓒ}, W + kT \leq U \vdash \forall x \left(\|x\|_2 < \delta \rightarrow \ldots \right)}$$

Premise ① proves by $\mathbb{R}$ from the antecedents, using assumption $t = 0$ to simplify the term $W + kt$ in Inv_a.

$$\mathbb{R}\;\dfrac{*}{V < W, t = 0 \vdash \mathit{Inv}_a}$$

Premise ③ proves by $\mathbb{R}$ from the loop invariant using the following arithmetic argument. Suppose for contradiction that there is a state satisfying the negation of the postcondition, i.e., assume the negation $t \geq T \land \|x\|_2 \geq \varepsilon$. Then, using the left conjunct of Inv_a together with $\|x\|_2 \geq \varepsilon$ to prove the LHS of the implication in ⓑ gives assumption $V \geq U$. The right conjunct of Inv_a then yields the chain of inequalities $V < W + kt \leq W + kT \leq U$, which is a contradiction to assumption $V \geq U$. The steps are outlined below.

$$\mathbb{R}\;\dfrac{\mathbb{R}\;\dfrac{\mathbb{R}\;\dfrac{\mathbb{R}\;\dfrac{*}{V \geq U, k < 0, W + kT \leq U, V < W + kt, t \geq T \vdash \mathit{false}}}{V \geq U, k < 0, W + kT \leq U, \mathit{Inv}_a, t \geq T \vdash \mathit{false}}}{\text{ⓑ}, k < 0, W + kT \leq U, \mathit{Inv}_a, t \geq T, \|x\|_2 \geq \varepsilon \vdash \mathit{false}}}{\text{ⓑ}, k < 0, W + kT \leq U, \mathit{Inv}_a \vdash t \geq T \rightarrow \|x\|_2 < \varepsilon}$$

The proof for premise ② proceeds by unfolding the loop body with $[\cup]$, $\land$R, yielding one premise for each switching choice $p \in \mathcal{P}$. A dC step proves the invariance of the left conjunct $V < W$ of Inv_a with $\mathrm{dI}_{\succcurlyeq}$ (see the stability proof, sublevel sets of V are invariant). The right conjunct of Inv_a is abbreviated $I \equiv V \geq U \rightarrow V < W + kt$ and it is proved below using axiom DCC, which results in premises ④ and ⑤ (shown and proved further below).

$$\text{[∪], ∧R}\;\dfrac{\text{dC, dI}_{\succcurlyeq}\;\dfrac{\text{DCC, ∧R}\;\dfrac{④\qquad ⑤}{\text{ⓒ}, I \vdash [x' = f_p(x), t' = 1 \,\&\, Q_p \land V < W]I}}{\text{ⓒ}, \mathit{Inv}_a \vdash [x' = f_p(x), t' = 1 \,\&\, Q_p]\mathit{Inv}_a \qquad (p \in \mathcal{P})}}{\text{ⓒ}, \mathit{Inv}_a \vdash [\bigcup_{p \in \mathcal{P}} x' = f_p(x), t' = 1 \,\&\, Q_p]\mathit{Inv}_a}$$

From premise ④, the proof is completed with a $\mathrm{dI}_{\succcurlyeq}$ step using the quantified assumption ⓒ because the domain constraint Q implies its closure formula $\overline{Q}$ and the strict inequality $V < W$ implies the nonstrict inequality $V \leq W$ which is needed for the LHS of the nested implication in ⓒ. The Lie derivative of RHS $W + kt$ is k using $t' = 1$.

$$\mathrm{dI}_{\succcurlyeq}\;\dfrac{\mathbb{R}\;\dfrac{*}{\text{ⓒ}, Q_p \land V < W \land V \geq U \vdash \mathcal{L}_{f_p}(V) \leq k}}{\text{ⓒ}, I \vdash [x' = f_p(x), t' = 1 \,\&\, Q_p \land V < W \land V \geq U]V < W + kt}$$

From premise ⑤, the proof is completed with a generalization G step followed by $\mathrm{dI}_{\succcurlyeq}$ to prove the invariance of formula $V < U$ (see the stability proof above, sublevel sets of V are invariant). The ODE in the outer box modality is elided with $\dots$ here.

$$\text{G}, \to\text{R}\;\dfrac{\text{dI}_{\succcurlyeq}\;\dfrac{*}{V < U \vdash [x' = f_p(x), t' = 1 \,\&\, Q_p \land V < W]V < U}}{\vdash [\dots](V < U \to [x' = f_p(x), t' = 1 \,\&\, Q_p \land V < W]V < U)} \qquad \Box$$

Proof of Corollary 6.10. The derivation of rule MLF builds on the ideas of the derivation of rule CLF from Corollary 6.9 so similar proof steps are explained in less detail here. The derivation starts as usual with an $\land$R step for the stability and pre-attractivity conjuncts which are proved separately below.

$$\land\text{R}\;\dfrac{\vdash \mathrm{UStab}(\alpha_{\texttt{state}}) \qquad \vdash \mathrm{UGpAttr}(\alpha_{\texttt{state}})}{\vdash \mathrm{UGpAS}(\alpha_{\texttt{state}})}$$

Stability. The derivation for stability similarly begins with cut and Skolemization steps. The difference compared to the derivation of rule CLF is the cut formulas are now conjunctions over all possible modes $p \in \mathcal{P}$ for the Lyapunov functions V_p. The first cut is $\exists W{>}0\,\text{ⓐ}$ with $\text{ⓐ} \equiv \bigwedge_{p\in\mathcal{P}} \forall x\,(\|x\|_2 = \varepsilon \to V_p \geq W)$, where the upper bound $W > 0$ is chosen to be the maximum of the respective bounds for each V_p on the compact set characterized by $\|x\|_2 = \varepsilon$. After Skolemizing W, the second arithmetic cut is the formula $\exists\delta\,(0 < \delta \leq \varepsilon \land \text{ⓑ})$ with $\text{ⓑ} \equiv \bigwedge_{p\in\mathcal{P}} \forall x\,(\|x\|_2 < \delta \to V_p < W)$. Such a δ exists by continuity for each $V_p, p \in \mathcal{P}$ since $V_p(0) = 0$ from the first (topmost) premise of rule MLF. After both cuts, the Skolemized δ from the antecedent is used to witness the succedent by $\exists$R.

$$\begin{array}{rc}
 & \text{ⓐ}, \delta \leq \varepsilon, \text{ⓑ} \vdash \forall x \left(\|x\|_2 < \delta \to [\alpha_{\texttt{state}}]\ \|x\|_2 < \varepsilon \right) \\ \hline
\exists\text{R} & \text{ⓐ}, 0 < \delta \leq \varepsilon, \text{ⓑ} \vdash \exists\delta{>}0\,\forall x \left(\|x\|_2 < \delta \to [\alpha_{\texttt{state}}]\ \|x\|_2 < \varepsilon \right) \\ \hline
\text{cut}, \mathbb{R}, \exists\text{L} & \varepsilon > 0, W > 0, \text{ⓐ} \vdash \exists\delta{>}0\,\forall x \left(\|x\|_2 < \delta \to [\alpha_{\texttt{state}}]\ \|x\|_2 < \varepsilon \right) \\ \hline
\text{cut}, \mathbb{R}, \exists\text{L} & \varepsilon > 0 \vdash \exists\delta{>}0\,\forall x \left(\|x\|_2 < \delta \to [\alpha_{\texttt{state}}]\ \|x\|_2 < \varepsilon \right) \\ \hline
\forall\text{R}, \to\text{R} & \vdash \mathrm{UStab}(\alpha_{\texttt{state}})
\end{array}$$

The derivation continues with logical simplification steps, Skolemizing the succedent and then proving the LHS of the implications in antecedent ⓑ.

$$\begin{array}{rc}
 & \text{ⓐ}, \delta \leq \varepsilon, \|x\|_2 < \delta, \bigwedge_{p\in\mathcal{P}} V_p < W \vdash [\alpha_{\texttt{state}}]\ \|x\|_2 < \varepsilon \\ \hline
\forall\text{L}, \to\text{L} & \text{ⓐ}, \delta \leq \varepsilon, \text{ⓑ}, \|x\|_2 < \delta \vdash [\alpha_{\texttt{state}}]\ \|x\|_2 < \varepsilon \\ \hline
\forall\text{R}, \to\text{R} & \text{ⓐ}, \delta \leq \varepsilon, \text{ⓑ} \vdash \forall x \left(\|x\|_2 < \delta \to [\alpha_{\texttt{state}}]\ \|x\|_2 < \varepsilon \right)
\end{array}$$

Next, a cut, $\lor$L step case splits on whether the switched system is initially in its domain of definition characterized by formula $Q \equiv \bigvee_{p\in\mathcal{P}} Q_p$. The case where the system is *not* in its domain is labeled ⓪ and the proof for this case is deferred to the end. In case the system is in its domain, the loop rule is used with stability loop invariant $\mathit{Inv}_s \equiv \|x\|_2 < \varepsilon \land \bigvee_{p\in\mathcal{P}} \left(Q_p \land V_p < W\right)$. This yields three premises labeled ①–③ shown and proved further below.

$$\text{cut}, \lor\text{L}\;\dfrac{\text{loop}\;\dfrac{\text{①} \qquad \text{②} \quad \text{③}}{\text{ⓐ}, \delta \leq \varepsilon, \|x\|_2 < \delta, \bigwedge_{p\in\mathcal{P}} V_p < W, Q \vdash [\alpha_{\texttt{state}}]\ \|x\|_2 < \varepsilon} \qquad \text{⓪}}{\text{ⓐ}, \delta \leq \varepsilon, \|x\|_2 < \delta, \bigwedge_{p\in\mathcal{P}} V_p < W \vdash [\alpha_{\texttt{state}}]\ \|x\|_2 < \varepsilon}$$

Premise ① proves by $\mathbb{R}$ from the antecedents using inequalities $\|x\|_2 < \delta$ and $\delta \le \varepsilon$ for the left conjunct and propositionally from antecedents Q and $\bigwedge_{p\in\mathcal{P}} V_p < W$ for the right conjunct.

$$\mathbb{R}\,\dfrac{*}{\delta \le \varepsilon, \|x\|_2 < \delta, \bigwedge_{p\in\mathcal{P}} V_p < W, Q \vdash \mathit{Inv}_s}$$

Premise ③ proves trivially since the postcondition $\|x\|_2 < \varepsilon$ is part of the loop invariant:

$$\mathbb{R}\,\dfrac{*}{\mathit{Inv}_s \vdash \|x\|_2 < \varepsilon}$$

The derivation continues from premise ② by unfolding the loop body of $\alpha_{\texttt{state}}$ with $[\cup]$, $\wedge$R. Premises are indexed by $p \in \mathcal{P}$ in the derivation. The M$[\cdot]$ step propositionally strengthens the postcondition to its constituent disjunct $\|x\|_2 < \varepsilon \wedge V_p < W$ for the chosen mode p. Then, DX assumes domain Q_p in the antecedent and a cut step adds the assumption $\|x\|_2 < \varepsilon \wedge V_p < W$. This cut corresponds to the last (bottom) premise of rule MLF. It is labeled ④ and explained below. The rest of the proof after the cut proceeds identically to the corresponding derivation for rule CLF using the respective conjunct for $p \in \mathcal{P}$ from ⓐ. The steps are omitted here.

$$\begin{array}{ll}
 & * \\ \cline{2-2}
 & \text{ⓐ}, \|x\|_2 < \varepsilon \wedge V_p < W \vdash [x' = f_p(x) \,\&\, Q_p](\|x\|_2 < \varepsilon \wedge V_p < W) \qquad \text{④} \\ \cline{2-2}
\text{cut} & \text{ⓐ}, \mathit{Inv}_s, Q_p \vdash [x' = f_p(x) \,\&\, Q_p](\|x\|_2 < \varepsilon \wedge V_p < W) \\ \cline{2-2}
\text{DX} & \text{ⓐ}, \mathit{Inv}_s \vdash [x' = f_p(x) \,\&\, Q_p](\|x\|_2 < \varepsilon \wedge V_p < W) \\ \cline{2-2}
\text{M}[\cdot] & \text{ⓐ}, \mathit{Inv}_s \vdash [x' = f_p(x) \,\&\, Q_p]\mathit{Inv}_s \qquad (p \in \mathcal{P}) \\ \cline{2-2}
[\cup], \wedge\text{R} & \text{ⓐ}, \mathit{Inv}_s \vdash [\bigcup_{p\in\mathcal{P}} x' = f_p(x) \,\&\, Q_p]\mathit{Inv}_s
\end{array}$$

The cut premise ④ is proved by splitting the disjunction in Inv_s with $\vee$L (indexed by $q \in \mathcal{P}$ below). The disjunct corresponding to mode p proves trivially. For modes $q \neq p$, the derivation yields a compatibility condition for switching from mode q to mode p which is proved using the last (bottom) premise of rule MLF. Note that the rule uses succedent $V_p = V_q$ since a symmetric condition ($V_q \le V_p$) is obtained when the roles of modes $p, q \in \mathcal{P}$ are swapped.

$$\begin{array}{ll}
 & * \\ \cline{2-2}
\mathbb{R} & Q_q, Q_p \vdash V_p \le V_q \\ \cline{2-2}
\mathbb{R} & p \neq q, Q_q, V_q < W, Q_p \vdash V_p < W \qquad (q \in \mathcal{P}) \\ \cline{2-2}
\vee\text{L} & \bigvee_{q\in\mathcal{P}} \big(Q_q \wedge V_q < W\big), Q_p \vdash V_p < W \\ \cline{2-2}
 & \mathit{Inv}_s, Q_p \vdash \|x\|_2 < \varepsilon \wedge V_p < W
\end{array}$$

Returning to premise ⓪, for initial states not in the switched system's domain, i.e., satisfying $\neg Q$, no continuous motion is possible within the model. This is proved using the loop invariant $\mathit{Inv}_s^0 \equiv \|x\|_2 < \varepsilon \wedge \neg Q$. The first and third premise resulting from the loop rule are proved trivially (not shown below). For the remaining premise, $\neg Q$ is preserved (trivially) across the loop body after unfolding it with $[\cup]$, $\wedge$R and using DX to show that the system is unable to switch to the ODE with domain Q_p because $\neg Q$ implies $\neg Q_p$ propositionally.

$$\begin{array}{ll}
 & * \\ \cline{2-2}
\text{DX} & \neg Q \vdash [x' = f_p(x) \,\&\, Q_p]\mathit{Inv}_s^0 \qquad (p \in \mathcal{P}) \\ \cline{2-2}
[\cup], \wedge\text{R} & \mathit{Inv}_a^0 \vdash [\bigcup_{p\in\mathcal{P}} x' = f_p(x) \,\&\, Q_p]\mathit{Inv}_s^0 \\ \cline{2-2}
\text{loop} & \delta \le \varepsilon, \|x\|_2 < \delta, \neg Q \vdash [\alpha_{\texttt{state}}]\ \|x\|_2 < \varepsilon
\end{array}$$

Pre-Attractivity. The derivation for pre-attractivity begins with logical simplification followed by a series of arithmetic cuts. First, the multiple Lyapunov functions $V_p, p \in \mathcal{P}$ are simultaneously bounded above on the ball characterized by $\|x\|_2 < \delta$, with the cut $\exists W{>}0\,\textcircled{a}$ where $\textcircled{a} \equiv \bigwedge_{p\in\mathcal{P}} \forall x \left(\|x\|_2 < \delta \rightarrow V_p < W \right)$. The upper bound W is Skolemized, then the next arithmetic cut uses $\exists U{>}0\,\textcircled{b}$ with $\textcircled{b} \equiv \bigwedge_{p\in\mathcal{P}} \forall x\,(V_p \leq W \wedge \|x\|_2 \geq \varepsilon \rightarrow V_p \geq U)$ (using radial unboundedness of all functions V_p from the second premise of MLF). Then, U is Skolemized with ∃L. The steps are shown below, with the box modality in $\mathrm{UGpAttr}(\alpha_{\mathtt{state}})$ temporarily hidden with $\ldots$ as it is not relevant for this part of the derivation.

$$\begin{array}{l r}
 & \varepsilon > 0, W > 0, \textcircled{a}, U > 0, \textcircled{b} \vdash \exists T{\geq}0 \forall x \left(\|x\|_2 < \delta \rightarrow \ldots \right) \\ \hline
\text{cut}, \mathbb{R}, \exists\text{L} & \varepsilon > 0, W > 0, \textcircled{a} \vdash \exists T{\geq}0 \forall x \left(\|x\|_2 < \delta \rightarrow \ldots \right) \\ \hline
\text{cut}, \mathbb{R}, \exists\text{L} & \varepsilon > 0 \vdash \exists T{\geq}0 \forall x \left(\|x\|_2 < \delta \rightarrow \ldots \right) \\ \hline
\forall\text{R}, \rightarrow\text{R} & \vdash \mathrm{UGpAttr}(\alpha_{\mathtt{state}})
\end{array}$$

Like the derivation of rule CLF from Corollary 6.9, the premises of rule MLF prove that, for each $p \in \mathcal{P}$, the Lie derivatives $\mathcal{L}_{f_p}(V_p)$ are bounded above by some $k_p < 0$ on the compact set characterized by formula $\overline{Q_p} \wedge U \leq V_p \leq W$. This justifies the following arithmetic cut, where the bound $k < 0$ is chosen to be the maximum over all k_p across all switching choices $p \in \mathcal{P}$:

$$\exists k < 0 \; \underbrace{\bigwedge_{p\in\mathcal{P}} \forall x \left(\overline{Q_p} \wedge U \leq V_p \leq W \rightarrow \mathcal{L}_{f_p}(V_p) \leq k \right)}_{\textcircled{c}}$$

The derivation continues similarly to rule CLF, first picking $T > 0$ satisfying $W + kT \leq U$, then Skolemizing and unfolding the succedent propositionally.

$$\begin{array}{l r}
 & \textcircled{a}, \textcircled{b}, k < 0, \textcircled{c}, T > 0, W + kT \leq U, \|x\|_2 < \delta \vdash \ldots \\ \hline
\forall\text{R}, \rightarrow\text{R} & \textcircled{a}, \textcircled{b}, k < 0, \textcircled{c}, T > 0, W + kT \leq U \vdash \forall x \left(\|x\|_2 < \delta \rightarrow \ldots \right) \\ \hline
\exists\text{R} & \varepsilon > 0, W > 0, \textcircled{a}, U > 0, \textcircled{b}, k < 0, \textcircled{c} \vdash \exists T{\geq}0 \ldots \\ \hline
\text{cut}, \mathbb{R}, \exists\text{L} & \varepsilon > 0, W > 0, \textcircled{a}, U > 0, \textcircled{b} \vdash \exists T{\geq}0 \ldots
\end{array}$$

The LHS in antecedent ⓐ is proved and the succedent is further unfolded with [;], [:=]. The antecedents are abbreviated with $\Gamma \equiv \textcircled{b}, k < 0, \textcircled{c}, T > 0, W + kT \leq U$ below. Similar to the stability proof, the derivation continues with a cut, ∨L step that case splits on whether the switched system is initially in its domain of definition $Q \equiv \bigvee_{p\in\mathcal{P}} Q_p$. The case where the system is *not* in its domain is labeled ⓪ and its proof is deferred to the end. In case the system is in domain Q, the loop rule is used with pre-attractivity loop invariant $\mathit{Inv}_a \equiv \bigvee_{p\in\mathcal{P}} \big(Q_p \wedge V_p < W \wedge (V_p \geq U \rightarrow V_p < W + kt)\big)$. This results in three premises ①–③ which are proved below.

$$\begin{array}{l r l}
 & \textcircled{1} \qquad \textcircled{2} \quad \textcircled{3} & \\ \hline
\text{loop} & \Gamma, \bigwedge_{p\in\mathcal{P}} V_p < W, t = 0, Q \vdash [\alpha_{\mathtt{state}}, t' = 1] \ldots & \textcircled{0} \\ \hline
\text{cut}, \vee\text{L} & \Gamma, \bigwedge_{p\in\mathcal{P}} V_p < W, t = 0 \vdash [\alpha_{\mathtt{state}}, t' = 1] \ldots & \\ \hline
[;], [:=] & \Gamma, \bigwedge_{p\in\mathcal{P}} V_p < W \vdash [t := 0; \alpha_{\mathtt{state}}, t' = 1] \ldots & \\ \hline
\forall\text{L}, \rightarrow\text{L} & \Gamma, \textcircled{a}, \|x\|_2 < \delta \vdash [t := 0; \alpha_{\mathtt{state}}, t' = 1] \ldots &
\end{array}$$

Premise ① proves propositionally from the antecedents after simplifying the term $W + kt$ using assumption $t = 0$.

$$\mathbb{R}\,\dfrac{*}{\bigwedge_{p\in\mathcal{P}} V_p < W, t = 0, Q \vdash \mathit{Inv}_a}$$

Premise ③ proves by $\mathbb{R}$ from the loop invariant after using $\vee$L to split the disjuncts of the loop invariant. The disjunct for mode $p \in \mathcal{P}$ is abbreviated $R \equiv V_p < W \wedge (V_p \geq U \rightarrow V_p < W + kt)$. The rest of the arithmetic argument is identical to the corresponding premise for CLF using the conjunct for p in ⓑ (summarized below).

$$\vee\mathrm{L}\,\dfrac{\mathbb{R}\,\dfrac{\mathbb{R}\,\dfrac{\mathbb{R}\,\dfrac{\mathbb{R}\,\dfrac{*}{V_p \geq U, k < 0, W + kT \leq U, V_p < W + kt, t \geq T \vdash \mathit{false}}}{V_p \geq U, k < 0, W + kT \leq U, R, t \geq T \vdash \mathit{false}}}{\text{ⓑ}, k < 0, W + kT \leq U, R, t \geq T, \|x\|_2 \geq \varepsilon \vdash \mathit{false}}}{\text{ⓑ}, k < 0, W + kT \leq U, R \vdash t \geq T \rightarrow \|x\|_2 < \varepsilon}}{\text{ⓑ}, k < 0, W + kT \leq U, \mathit{Inv}_a \vdash t \geq T \rightarrow \|x\|_2 < \varepsilon}$$

The derivation from premise ② proceeds by unfolding the loop body with $[\cup]$, $\wedge$R, DX, yielding one premise for each switching choice $p \in \mathcal{P}$. The M$[\cdot]$ step selects the disjunct R (as defined above for premise ③) in the postcondition corresponding to mode p and the cut adds this disjunct to the antecedents (the cut premise ④ is shown and proved below). The rest of the proof after the cut is omitted here as it is identical to the corresponding derivation for rule CLF using the respective conjunct for mode p in ⓒ.

$$[\cup], \wedge\mathrm{R}, \mathrm{DX}\,\dfrac{\mathrm{M}[\cdot]\,\dfrac{\mathrm{cut}\,\dfrac{\dfrac{*}{\text{④} \qquad \text{ⓒ}, R \vdash [x' = f_p(x), t' = 1 \,\&\, Q_p]R}}{\text{ⓒ}, \mathit{Inv}_a, Q_p \vdash [x' = f_p(x), t' = 1 \,\&\, Q_p]R}}{\text{ⓒ}, \mathit{Inv}_a, Q_p \vdash [x' = f_p(x), t' = 1 \,\&\, Q_p]\mathit{Inv}_a \qquad (p \in \mathcal{P})}}{\text{ⓒ}, \mathit{Inv}_a \vdash [\bigcup_{p\in\mathcal{P}} x' = f_p(x), t' = 1 \,\&\, Q_p]\mathit{Inv}_a}$$

The cut premise ④ is proved by splitting the disjunction in Inv_a with $\vee$L (indexed by $q \in \mathcal{P}$ below). For modes $q \neq p$, the derivation needs a compatibility condition which proves using the last (bottom) premise of rule MLF, similar to the stability proof.

$$\dfrac{\vee\mathrm{L}\,\dfrac{\mathbb{R}\,\dfrac{\mathbb{R}\,\dfrac{*}{Q_q, Q_p \vdash V_p \leq V_q}}{p \neq q, Q_q \wedge V_q < W \wedge (V_q \geq U \rightarrow V_q < W + kt), Q_p \vdash R \quad (q \in \mathcal{P})}}{\bigvee_{q\in\mathcal{P}} \big(Q_q \wedge V_q < W \wedge (V_q \geq U \rightarrow V_q < W + kt)\big), Q_p \vdash R}}{\mathit{Inv}_a, Q_p \vdash R}$$

Returning to premise ⓪, similar to the case for stability, initial states satisfying $\neg Q$ have no continuous motion possible so they are stuck at the initial state (with global clock $t = 0$). This is proved using the loop invariant $\mathit{Inv}_a^0 \equiv t = 0 \wedge \neg Q$. The first and third premise resulting from the loop rule are proved trivially (not shown below). For the remaining premise, $\neg Q$ is preserved (trivially) across the loop body after unfolding it with $[\cup]$, $\wedge$R and using DX to show that the system is unable to switch to the ODE with domain Q_p because $\neg Q$ implies $\neg Q_p$ propositionally.

$$\text{loop}\;\dfrac{[\cup],\wedge\text{R}\;\dfrac{\text{DX}\;\dfrac{*}{\neg Q \vdash [x' = f_p(x), t' = 1 \,\&\, Q_p]\mathit{Inv}_a^0 \qquad (p \in \mathcal{P})}}{\mathit{Inv}_a^0 \vdash [\bigcup_{p\in\mathcal{P}} x' = f_p(x), t' = 1 \,\&\, Q_p]\mathit{Inv}_a^0}}{T > 0, t = 0, \neg Q \vdash [\alpha_{\texttt{state}}, t' = 1](t \geq T \to \|x\|_2 < \varepsilon)} \qquad \square$$

Proof of Corollary 6.11. The derivation of rule MLF_G is similar to the derivation of rule MLF from Corollary 6.10, but adapted to the shape of the guarded state-dependent switching model $\alpha_{\texttt{guard}}$ and its corresponding loop invariants. The derivation starts as usual with an $\wedge$R step for the stability and pre-attractivity conjuncts which are proved separately below.

$$\wedge\text{R}\;\dfrac{\vdash \text{UStab}(\alpha_{\texttt{guard}}) \qquad \vdash \text{UGpAttr}(\alpha_{\texttt{guard}})}{\vdash \text{UGpAS}(\alpha_{\texttt{guard}})}$$

Stability. The derivation for stability proceeds identically to the derivation for rule MLF from Corollary 6.10 until the step before the stability loop invariant is used. These steps are omitted below with $\ldots$ and the resulting premise has antecedent formula abbreviated with ⓐ $\equiv \bigwedge_{p\in\mathcal{P}} \forall x\, (\|x\|_2 = \varepsilon \to V_p \geq W)$.

$$\dfrac{\dfrac{\text{ⓐ}, \delta \leq \varepsilon, \|x\|_2 < \delta, \bigwedge_{p\in\mathcal{P}} V_p < W \vdash [\alpha_{\texttt{guard}}]\; \|x\|_2 < \varepsilon}{\ldots}}{\vdash \text{UStab}(\alpha_{\texttt{guard}})}$$

The derivation continues using the loopT rule with the modified stability loop invariant $\mathit{Inv}_s \equiv \|x\|_2 < \varepsilon \wedge \bigvee_{p\in\mathcal{P}} \big(u = p \wedge V_p < W\big)$. This yields four premises labeled ①–④, shown and proved further below.

$$\text{loopT}\;\dfrac{① \qquad ② \quad ③ \quad ④}{\text{ⓐ}, \delta \leq \varepsilon, \|x\|_2 < \delta, \bigwedge_{p\in\mathcal{P}} V_p < W \vdash [\alpha_{\texttt{guard}}]\; \|x\|_2 < \varepsilon}$$

Premise ① shows that the system state satisfies the invariant Inv_s after running the initialization program $\alpha_i \equiv \bigcup_{p\in\mathcal{P}} u := p$. This is proved by $\mathbb{R}$ after unfolding α_i using $[\cup]$, $[:=]$.

$$[\cup],[:=]\;\dfrac{\mathbb{R}\;\dfrac{*}{\delta \leq \varepsilon, \|x\|_2 < \delta, \bigwedge_{p\in\mathcal{P}} V_p < W, u = p \vdash \mathit{Inv}_s \qquad (p \in \mathcal{P})}}{\delta \leq \varepsilon, \|x\|_2 < \delta, \bigwedge_{p\in\mathcal{P}} V_p < W \vdash [\alpha_i]\mathit{Inv}_s}$$

Premise ④ proves trivially since the postcondition $\|x\|_2 < \varepsilon$ is part of the loop invariant.

$$\mathbb{R}\;\dfrac{*}{\mathit{Inv}_s \vdash \|x\|_2 < \varepsilon}$$

The derivation from premise ② yields *correct-by-construction* arithmetical conditions on the Lyapunov functions from unfolding the switching controller in $\alpha_{\texttt{guard}}$, recall

$$\alpha_u \equiv \bigcup_{p\in\mathcal{P}} \Big(?u = p; (\bigcup_{q\in\mathcal{P}} (?G_{p,q}; u := q) \cup u := u)\Big)$$

Axiom $[\cup]$ unfolds the outer choice $\bigcup_{p\in\mathcal{P}} \big(\,\cdot\,\big)$, yielding one premise for each mode $p \in \mathcal{P}$. Then, axioms $[;]$, $[?]$ add the current mode $u = p$ (before switching) to the assumptions. The cut step propositionally unfolds antecedent loop invariant assumption Inv_s to the corresponding disjunct for $u = p$. The inner choice $\bigcup_{q\in\mathcal{P}}\big(\,\cdot\,\big)$ is unfolded next with axioms [∪], [;], [?], yielding one premise for each possible transition to mode $q \in \mathcal{P}$ guarded by formula $G_{p,q}$; the case with no switching $u := u$ is trivial since the antecedents imply the postcondition (unchanged). The assignment $u := q$ is unfolded with [:=], so the succedent simplifies to the disjunct for $u = q$ in Inv_s. An arithmetic simplification step yieds the bottom premise of rule MLF_G.

$$
\begin{array}{rc}
 & * \\ \cline{2-2}
\mathbb{R} & G_{p,q} \vdash V_q \le V_p \\ \cline{2-2}
\mathbb{R} & V_p < W, G_{p,q} \vdash V_q < W \\ \cline{2-2}
[:=] & \|x\|_2 < \varepsilon, V_p < W, G_{p,q} \vdash [u := q]\mathit{Inv}_s \qquad (q \in \mathcal{P}) \\ \cline{2-2}
[\cup], [;], [?] & \|x\|_2 < \varepsilon, u = p, V_p < W \vdash [\bigcup_{q\in\mathcal{P}} (?G_{p,q}; u := q) \cup u := u]\mathit{Inv}_s \\ \cline{2-2}
\text{cut} & \mathit{Inv}_s, u = p \vdash [\bigcup_{q\in\mathcal{P}} (?G_{p,q}; u := q) \cup u := u]\mathit{Inv}_s \\ \cline{2-2}
[;], [?] & \mathit{Inv}_s \vdash [?u = p; (\bigcup_{q\in\mathcal{P}} (?G_{p,q}; u := q) \cup u := u)]\mathit{Inv}_s \ (p \in \mathcal{P}) \\ \cline{2-2}
[\cup] & \mathit{Inv}_s \vdash [\alpha_u]\mathit{Inv}_s
\end{array}
$$

The derivation from premise ③ unfolds the plant model $\alpha_p \equiv \bigcup_{p\in\mathcal{P}}\big(?u = p; x' = f_p(x,y)\,\&\,Q_p\big)$. The choice $\bigcup_{p\in\mathcal{P}}\big(\,\cdot\,\big)$ is unfolded with axiom [∪], yielding one premise for each mode $p \in \mathcal{P}$. Then, axioms [;], [?] add the mode selected by α_u to the antecedent, where the antecedent loop invariant assumption Inv_s is simplified by cut to the disjunct for $u = p$. Similarly M[·] strengthens the postcondition to the disjunct for $u = p$. The rest of the proof proceeds identically to the corresponding derivation for rule CLF in Corollary 6.9 so it is omitted here.

$$
\begin{array}{rc}
 & * \\ \cline{2-2}
 & \text{ⓐ}, \|x\|_2 < \varepsilon, V_p < W \vdash [x' = f_p(x)\,\&\,Q_p](\|x\|_2 < \varepsilon \land V_p < W) \\ \cline{2-2}
\mathrm{M}[\cdot] & \text{ⓐ}, \|x\|_2 < \varepsilon, V_p < W, u = p \vdash [x' = f_p(x)\,\&\,Q_p]\mathit{Inv}_s \\ \cline{2-2}
\text{cut} & \text{ⓐ}, \mathit{Inv}_s, u = p \vdash [x' = f_p(x)\,\&\,Q_p]\mathit{Inv}_s \\ \cline{2-2}
[;], [?] & \text{ⓐ}, \mathit{Inv}_s \vdash [?u = p; x' = f_p(x,y)\,\&\,Q_p]\mathit{Inv}_s \ (p \in \mathcal{P}) \\ \cline{2-2}
[\cup] & \text{ⓐ}, \mathit{Inv}_s \vdash [\alpha_p]\mathit{Inv}_s
\end{array}
$$

Pre-Attractivity. The derivation for pre-attractivity is also identical to MLF until the step before the pre-attractivity loop invariant is used. These steps are omitted below with … and the resulting premise has antecedent formulas abbreviated with:

$$
\text{ⓑ} \equiv \bigwedge_{p\in\mathcal{P}} \forall x\,(V_p \le W \land \|x\|_2 \ge \varepsilon \to V_p \ge U)
$$

$$
\text{ⓒ} \equiv \bigwedge_{p\in\mathcal{P}} \forall x\,\big(\overline{Q_p} \land U \le V_p \le W \to \mathcal{L}_{f_p}(V_p) \le k\big)
$$

$$
\cfrac{\cfrac{\bigwedge_{p\in\mathcal{P}} V_p < W, \text{ⓑ}, k < 0, \text{ⓒ}, W + kT \le U, t = 0 \vdash [\alpha_{\mathtt{guard}}, t' = 1]\ldots}{\ldots}}{\vdash \mathrm{UGpAttr}(\alpha_{\mathtt{guard}})}
$$

The derivation continues using the loopT rule with pre-attractivity loop invariant $\mathit{Inv}_a \equiv \bigvee_{p\in\mathcal{P}} \big(u = p \wedge V_p < W \wedge (V_p \geq U \rightarrow V_p < W + kt)\big)$. This yields four premises labeled ①–④ which are shown and proved further below.

$$\text{loopT}\ \dfrac{\text{①}\qquad\quad \text{②}\quad \text{③}\quad \text{④}}{\bigwedge_{p\in\mathcal{P}} V_p < W, \text{ⓑ}, k < 0, \text{ⓒ}, W + kT \leq U, t = 0 \vdash [\alpha_{\texttt{guard}}, t' = 1]\ldots}$$

Premise ① proves the invariant Inv_a after unfolding the initialization program α_i using $[\cup]$, $[:=]$.

$$[\cup], [:=]\ \dfrac{\mathbb{R}\ \dfrac{*}{\bigwedge_{p\in\mathcal{P}} V_p < W, t = 0, u = p \vdash \mathit{Inv}_a}}{\bigwedge_{p\in\mathcal{P}} V_p < W, t = 0 \vdash [\alpha_i]\mathit{Inv}_a}$$

Premise ④ is proved by $\mathbb{R}$ after unfolding the disjuncts of the loop invariant with $\vee$L (the arithmetical argument is identical to earlier proofs). The selected disjunct of Inv_a (indexed by p) is abbreviated $R \equiv u = p \wedge V_p < W \wedge (V_p \geq U \rightarrow V_p < W + kt)$.

$$\vee\text{L}\ \dfrac{\mathbb{R}\ \dfrac{*}{\text{ⓑ}, k < 0, W + kT \leq U, R \vdash t \geq T \rightarrow \|x\|_2 < \varepsilon \qquad (p \in \mathcal{P})}}{\text{ⓑ}, k < 0, W + kT \leq U, \mathit{Inv}_a \vdash t \geq T \rightarrow \|x\|_2 < \varepsilon}$$

The derivation from premise ② unfolds α_u using dL's hybrid program axioms similar to the stability proof, and an arithmetic simplification step yields the premises of MLF$_G$ for guarded mode switches from p to q, for $p, q \in \mathcal{P}$.

$$[\cup]\ \dfrac{[;], [?]\ \dfrac{\text{cut}, \vee\text{L}\ \dfrac{[\cup], [;], [?]\ \dfrac{[:=]\ \dfrac{\mathbb{R}\ \dfrac{\mathbb{R}\ \dfrac{*}{G_{p,q} \vdash V_q \leq V_p}}{R, G_{p,q} \vdash V_q < W \wedge (V_q \geq U \rightarrow V_q < W + kt)}}{R, G_{p,q} \vdash [u := q]\mathit{Inv}_a \qquad (q \in \mathcal{P})}}{R \vdash [\bigcup_{q\in\mathcal{P}} (?G_{p,q}; u := q) \cup u := u]\mathit{Inv}_a}}{\mathit{Inv}_a, u = p \vdash [\bigcup_{q\in\mathcal{P}} (?G_{p,q}; u := q) \cup u := u]\mathit{Inv}_a}}{\mathit{Inv}_a \vdash [?u = p; (\bigcup_{q\in\mathcal{P}} (?G_{p,q}; u := q) \cup u := u)]\mathit{Inv}_a \ (p \in \mathcal{P})}}{\mathit{Inv}_a \vdash [\alpha_u]\mathit{Inv}_a}$$

The derivation from premise ③ unfolds the plant model and proceeds identically to the corresponding derivation for rule CLF, with $R \equiv u = p \wedge V_p < W \wedge (V_p \geq U \rightarrow V_p < W + kt)$.

$$[\cup]\ \dfrac{[;], [?]\ \dfrac{\text{cut}\ \dfrac{\text{M}[\cdot]\ \dfrac{\dfrac{*}{\text{ⓒ}, R \vdash [x' = f_p(x), t' = 1 \,\&\, Q_p]R}}{\text{ⓒ}, R \vdash [x' = f_p(x), t' = 1 \,\&\, Q_p]\mathit{Inv}_a}}{\text{ⓒ}, \mathit{Inv}_a, u = p \vdash [x' = f_p(x), t' = 1 \,\&\, Q_p]\mathit{Inv}_a}}{\text{ⓒ}, \mathit{Inv}_a \vdash [?u = p; x' = f_p(x, y), t' = 1 \,\&\, Q_p]\mathit{Inv}_a}}{\text{ⓒ}, \mathit{Inv}_a \vdash [\alpha_p, t' = 1]\mathit{Inv}_a}$$

□

Proof of Corollary 6.12. The derivation of rule MLF$_\tau$ departs more significantly from the derivations of rules CLF, MLF, MLF$_G$. For this proof, $\mathbb{R}_{\exp}$ is used to indicate arithmetic steps that use properties of the real exponential function. Although arithmetic over the exponential function is not known to be decidable, tools are available for answering specialized subsets of such questions [60]. Additional explanation is given below for $\mathbb{R}_{\exp}$ steps that only require elementary properties of the exponential function.

The proof also shows how to derive arithmetic conditions (arising from the time-dependent switching controller) in a correct by construction manner through the hybrid program axioms of dL [142, 144]. Recall from Corollary 6.12 that the modes $p \in \mathcal{P}$ are partitioned into two subsets consisting of the stable $\mathcal{S} = \{p \in \mathcal{P}, \lambda_p > 0\}$ and unstable $\mathcal{U} = \{p \in \mathcal{P}, \lambda_p \leq 0\}$ modes. The derivation starts as usual with an $\wedge$R step for the stability and pre-attractivity conjuncts which are proved separately below.

$$
{\scriptstyle \wedge\mathrm{R}}\ \cfrac{\vdash \mathrm{UStab}(\alpha_{\mathtt{time}}) \qquad \vdash \mathrm{UGpAttr}(\alpha_{\mathtt{time}})}{\vdash \mathrm{UGpAS}(\alpha_{\mathtt{time}})}
$$

Stability. The stability derivation begins with cut and Skolemization steps. The first cut is $\exists W{>}0\, \textcircled{a}$ with the abbreviation $\textcircled{a} \equiv \bigwedge_{p\in\mathcal{P}} \forall x\, (\|x\|_2 = \varepsilon \to V_p \geq W)$, where the upper bound $W > 0$ is chosen to be the maximum of the respective bounds for each V_p on the compact set characterized by $\|x\|_2 = \varepsilon$. After Skolemizing W, the second arithmetic cut is the formula $\exists \delta\, (0 < \delta \leq \varepsilon \wedge \textcircled{b})$, where the conjuncts for $p \in \mathcal{U}$ need the arithmetic fact $\exp(\lambda_p \Theta_p) > 0$.

$$
\textcircled{b} \equiv \bigwedge_{p\in\mathcal{S}} \forall x\, (\|x\|_2 < \delta \to V_p < W) \wedge \bigwedge_{p\in\mathcal{U}} \forall x\, (\|x\|_2 < \delta \to V_p < W \exp(\lambda_p \Theta_p))
$$

Such a δ exists by continuity for each $V_p, p \in \mathcal{P}$, $V_p(0) = 0$ from the premise of rule MLF$_\tau$. After both cuts, the Skolemized δ from the antecedent is used to witness the succedent by $\exists$R.

$$
{\scriptstyle \forall\mathrm{R},\, \to\mathrm{R}}\ \cfrac{{\scriptstyle \mathrm{cut},\, \mathbb{R},\, \exists\mathrm{L}}\ \cfrac{{\scriptstyle \mathrm{cut},\, \mathbb{R}_{\exp},\, \exists\mathrm{L}}\ \cfrac{{\scriptstyle \exists\mathrm{R}}\ \cfrac{\textcircled{a}, \delta \leq \varepsilon, \textcircled{b} \vdash \forall x\, \big(\, \|x\|_2 < \delta \to [\alpha_{\mathtt{time}}]\, \|x\|_2 < \varepsilon \big)}{\textcircled{a}, 0 < \delta \leq \varepsilon, \textcircled{b} \vdash \exists \delta{>}0\, \forall x\, \big(\, \|x\|_2 < \delta \to [\alpha_{\mathtt{time}}]\, \|x\|_2 < \varepsilon \big)}}{\varepsilon > 0, W > 0, \textcircled{a} \vdash \exists \delta{>}0\, \forall x\, \big(\, \|x\|_2 < \delta \to [\alpha_{\mathtt{time}}]\, \|x\|_2 < \varepsilon \big)}}{\varepsilon > 0 \vdash \exists \delta{>}0\, \forall x\, \big(\, \|x\|_2 < \delta \to [\alpha_{\mathtt{time}}]\, \|x\|_2 < \varepsilon \big)}}{\vdash \mathrm{UStab}(\alpha_{\mathtt{time}})}
$$

The derivation continues after both cuts similarly to MLF from Corollary 6.10 by unfolding and proving the LHS of the implications in antecedent $\textcircled{b}$. The resulting assumption on the initial state is abbreviated $B \equiv \bigwedge_{p\in\mathcal{S}} V_p < W \wedge \bigwedge_{p\in\mathcal{U}} V_p < W \exp(\lambda_p \Theta_p)$. Then, the loopT rule is used with the following stability loop invariant *Inv*$_s$, which yields premises ①–④ shown and proved further below:

$$
\mathit{Inv}_s \equiv \tau \geq 0 \wedge \|x\|_2 < \varepsilon\ \wedge \left(\begin{array}{l} \displaystyle\bigvee_{p\in\mathcal{S}} \big(u = p \wedge V_p < W \exp(-\lambda_p \tau)\big) \vee \\ \displaystyle\bigvee_{p\in\mathcal{U}} \big(u = p \wedge V_p < W \exp(-\lambda_p(\tau - \Theta_p)) \wedge \tau \leq \Theta_p\big) \end{array} \right)
$$

$$
{\scriptstyle \forall\mathrm{R},\, \to\mathrm{R}}\ \cfrac{{\scriptstyle \forall\mathrm{L},\, \to\mathrm{L}}\ \cfrac{{\scriptstyle \mathrm{loopT}}\ \cfrac{\textcircled{1} \qquad \textcircled{2} \quad \textcircled{3} \quad \textcircled{4}}{\textcircled{a}, \delta \leq \varepsilon, \|x\|_2 < \delta, B \vdash [\alpha_{\mathtt{time}}]\, \|x\|_2 < \varepsilon}}{\textcircled{a}, \delta \leq \varepsilon, \textcircled{b}, \|x\|_2 < \delta \vdash [\alpha_{\mathtt{time}}]\, \|x\|_2 < \varepsilon}}{\textcircled{a}, \delta \leq \varepsilon, \textcircled{b} \vdash \forall x\, \big(\, \|x\|_2 < \delta \to [\alpha_{\mathtt{time}}]\, \|x\|_2 < \varepsilon \big)}
$$

Premise ① shows that the system state satisfies the invariant *Inv*$_s$ after initialization with program $\alpha_i \equiv \tau := 0; \bigcup_{p\in\mathcal{P}} u := p$. This is proved from B after unfolding α_i using $[\cup]$, $[:=]$ and substituting $\tau = 0$ in the loop invariant (using $\exp(0) = 1$).

$$\scriptstyle{[\cup],[:=]}\dfrac{\scriptstyle{\mathbb{R}_{\exp}}\dfrac{*}{\delta \le \varepsilon, \|x\|_2 < \delta, B, \tau = 0, u = p \vdash \mathit{Inv}_s}}{\delta \le \varepsilon, \|x\|_2 < \delta, B \vdash [\alpha_i]\mathit{Inv}_s}$$

Premise ④ proves trivially since the postcondition $\|x\|_2 < \varepsilon$ is part of the loop invariant.

$$\scriptstyle{\mathbb{R}}\dfrac{*}{\mathit{Inv}_s \vdash \|x\|_2 < \varepsilon}$$

The derivation from premise ② unfolds the switching controller α_u in $\alpha_{\texttt{time}}$ with dL's hybrid program axioms, recall:

$$\alpha_u \equiv \bigcup_{p \in \mathcal{P}} \Big(?u = p; \big(\bigcup_{q \in \mathcal{P}} \big(?\theta_{p,q} \le \tau; \tau := 0; u := q\big) \cup u := u\big)\Big)$$

This unfolding yields four possible shapes of premises (abbreviated as ... and shown immediately below) for a switch from the current mode p to mode q. As usual, the case with no switching $u := u$ is trivial and proves by [;], [$\cup$]. In each (non-trivial) case, the antecedent assumption corresponds to the disjunct of Inv_s for mode p, while the succedent assumption corresponds to the disjunct for mode q with timer τ reset to 0 by the switching controller α_u.

$$\scriptstyle{[\cup]}\dfrac{\scriptstyle{[;],[?]}\dfrac{\scriptstyle{[;],[\cup]}\dfrac{\scriptstyle{[\cup],[;],[?],[:=]}\dfrac{\ldots}{\mathit{Inv}_s, u = p \vdash [\bigcup_{q \in \mathcal{P}} \big(?\theta_{p,q} \le \tau; \tau := 0; u := q\big)]\mathit{Inv}_s}}{\mathit{Inv}_s, u = p \vdash [\bigcup_{q \in \mathcal{P}} \big(?\theta_{p,q} \le \tau; \tau := 0; u := q\big) \cup u := u]\mathit{Inv}_s}}{\mathit{Inv}_s \vdash [?u = p; (\bigcup_{q \in \mathcal{P}} \big(?\theta_{p,q} \le \tau; \tau := 0; u := q\big) \cup u := u)]\mathit{Inv}_s}}{\mathit{Inv}_s \vdash [\alpha_u]\mathit{Inv}_s}$$

The four cases correspond to whether $p \in \mathcal{S}$ or $p \in \mathcal{U}$ and similarly for q, as labeled below.

$$\begin{array}{ll}
\theta_{p,q} \le \tau, V_p < W\exp(-\lambda_p\tau) \vdash V_q < W & (p \in \mathcal{S}, q \in \mathcal{S}) \\
\theta_{p,q} \le \tau, V_p < W\exp(-\lambda_p\tau) \vdash V_q < W\exp(\lambda_q\Theta_q) & (p \in \mathcal{S}, q \in \mathcal{U}) \\
\theta_{p,q} \le \tau, V_p < W\exp(-\lambda_p(\tau - \Theta_p)), \tau \le \Theta_p \vdash V_q < W & (p \in \mathcal{U}, q \in \mathcal{S}) \\
\theta_{p,q} \le \tau, V_p < W\exp(-\lambda_p(\tau - \Theta_p)), \tau \le \Theta_p \vdash V_q < W\exp(\lambda_q\Theta_q) & (p \in \mathcal{U}, q \in \mathcal{U})
\end{array}$$

These premises are correct-by-construction and can be handed to an arithmetic solver directly. They can also be simplified, e.g., for $p \in \mathcal{S}, q \in \mathcal{S}$, the inequalities can be rearranged to eliminate W and τ. The first $\mathbb{R}$ step uses transitivity of $<$ and $\le$, while the second $\mathbb{R}_{\exp}$ step uses monotonicity $\exp(\lambda_p\theta_{p,q}) \le \exp(\lambda_p\tau)$ whenever $\lambda_p > 0$ (since $p \in \mathcal{S}$) and $\theta_{p,q} \le \tau$. Intuitively, the resulting (simplified) premise says that by choosing sufficiently large dwell time $\theta_{p,q}$ (for stable mode p), one can offset an increase in value when switching from V_p to V_q. The resulting arithmetic condition $V_q \le V_p\exp(\lambda_p\theta_{p,q})$ is a *correct-by-construction* premise for rule MLF$_\tau$.

$$\scriptstyle{\mathbb{R}}\dfrac{\scriptstyle{\mathbb{R}_{\exp}}\dfrac{\vdash V_q \le V_p\exp(\lambda_p\theta_{p,q})}{\theta_{p,q} \le \tau \vdash V_q \le V_p\exp(\lambda_p\tau)}}{\theta_{p,q} \le \tau, V_p < W\exp(-\lambda_p\tau) \vdash V_q < W}$$

The derivation from premise ③ unfolds the plant model α_p using dL axioms. There are two possible shapes of the premises resulting from this unfolding, depending if $p \in \mathcal{S}$ or $p \in \mathcal{U}$, these are abbreviated ⑤ and ⑥ respectively. In either case, the derivation shows that the appropriate upper bound on V_p is preserved for the invariant.

$$
\begin{array}{r c}
 & \textcircled{5} \qquad \textcircled{6} \\ \cline{2-2}
{\scriptstyle [;],\,[?]} & \textcircled{a}, \mathit{Inv}_s, u = p \vdash [x' = f_p(x), \tau' = 1 \,\&\, \tau \le \Theta_p]\mathit{Inv}_s \\ \cline{2-2}
{\scriptstyle [;],\,[?]} & \textcircled{a}, \mathit{Inv}_s \vdash [?u = p; x' = f_p(x), \tau' = 1 \,\&\, \tau \le \Theta_p]\mathit{Inv}_s \\ \cline{2-2}
{\scriptstyle [\cup]} & \textcircled{a}, \mathit{Inv}_s \vdash [\alpha_p]\mathit{Inv}_s
\end{array}
$$

For premise ⑤, the proof uses dbx$_{\succcurlyeq}$ with cofactor $-\lambda_p$, where the Lie derivative of subterm $W\exp(-\lambda_p\tau)$ is $(-\lambda_p)W\exp(-\lambda_p\tau)$ from $\tau' = 1$. The resulting premise simplifies to the third premise of rule MLF$_\tau$.

$$
\begin{array}{r c}
 & * \\ \hline
 & \vdash \mathcal{L}_{f_p}(V_p) \le -\lambda_p V_p \\ \hline
 & \vdash \mathcal{L}_{f_p}(V_p) - (-\lambda_p)W\exp(-\lambda_p\tau) \le -\lambda_p(V_p - W\exp(-\lambda_p\tau)) \\ \hline
{\scriptstyle \mathrm{dbx}_{\succcurlyeq}} & V_p - W\exp(-\lambda_p\tau) < 0 \vdash [x' = f_p(x), \tau' = 1 \,\&\, \tau \le \Theta_p]V_p - W\exp(-\lambda_p\tau) < 0 \\ \hline
{\scriptstyle \mathrm{cut},\, \mathrm{M}[\cdot]} & V_p < W\exp(-\lambda_p\tau) \vdash [x' = f_p(x), \tau' = 1 \,\&\, \tau \le \Theta_p]V_p < W\exp(-\lambda_p\tau)
\end{array}
$$

The proof for premise ⑥ also uses dbx$_{\succcurlyeq}$ with cofactor $-\lambda_p$, yielding the third premise of rule MLF$_\tau$ again.

$$
\begin{array}{r c}
 & * \\ \hline
 & \vdash \mathcal{L}_{f_p}(V_p) \le -\lambda_p V_p \\ \hline
{\scriptstyle \mathrm{dbx}_{\succcurlyeq}} & V_p < W\exp(-\lambda_p(\tau - \Theta_p)) \vdash [x' = f_p(x), \tau' = 1 \,\&\, \tau \le \Theta_p]V_p < W\exp(-\lambda_p(\tau - \Theta_p))
\end{array}
$$

Pre-Attractivity. The pre-attractivity proof requires an additional input parameter $\sigma > 0$ for the overall decay factor with $\sigma < \lambda_p$ for $p \in \mathcal{S}$ (σ must also satisfy other arithmetic properties, to be derived in a correct-by-construction manner in the proof). The derivation begins with logical simplification followed by a series of arithmetic cuts. First, the multiple Lyapunov functions $V_p, p \in \mathcal{P}$ are simultaneously bounded above on the ball characterized by $\|x\|_2 < \delta$, with the cut $\exists W{>}0\,\textcircled{a}$ (abbreviated below) where the conjuncts for $p \in \mathcal{U}$ need the arithmetic fact $\exp(\lambda_p\Theta_p) > 0$ (by $\mathbb{R}_{\exp}$).

$$
\textcircled{a} \equiv \bigwedge_{p \in \mathcal{S}} \forall x \,(\|x\|_2 < \delta \rightarrow V_p < W) \land \bigwedge_{p \in \mathcal{U}} \forall x \,(\|x\|_2 < \delta \rightarrow V_p < W\exp(\lambda_p\Theta_p))
$$

The upper bound W is Skolemized, then the next arithmetic cut uses $\exists U{>}0\,\textcircled{b}$ with $\textcircled{b} \equiv \bigwedge_{p \in \mathcal{P}} \forall x \,(V_p \le W \land \|x\|_2 \ge \varepsilon \rightarrow V_p \ge U)$, where U is Skolemized with $\exists$L.

$$
\begin{array}{r r}
 & \varepsilon > 0, W > 0, \textcircled{a}, U > 0, \textcircled{b} \vdash \exists T{\ge}0 \forall x \left(\|x\|_2 < \delta \rightarrow \dots \right) \\ \hline
{\scriptstyle \mathrm{cut},\, \mathbb{R},\, \exists\mathrm{L}} & \varepsilon > 0, W > 0, \textcircled{a} \vdash \exists T{\ge}0 \forall x \left(\|x\|_2 < \delta \rightarrow \dots \right) \\ \hline
{\scriptstyle \mathrm{cut},\, \mathbb{R}_{\exp},\, \exists\mathrm{L}} & \varepsilon > 0 \vdash \exists T{\ge}0 \forall x \left(\|x\|_2 < \delta \rightarrow \dots \right) \\ \hline
{\scriptstyle \forall\mathrm{R},\, \rightarrow\mathrm{R}} & \vdash \mathrm{UGpAttr}(\alpha_{\mathtt{time}})
\end{array}
$$

The derivation continues by picking $T \ge 0$ using the abbreviated formula in the derivation $R \equiv W \le U\exp(\sigma T) \land \bigwedge_{p \in \mathcal{U}} W \le U\exp(\sigma T)\exp(-\sigma\Theta_p)$, such a T exists by $\mathbb{R}_{\exp}$ because

$\sigma > 0$ so the $\exp(\sigma T)$ term on the RHS of each inequality can be chosen arbitrarily large. The quantifiers in the succedent are unfolded and the LHS of the implications in ⓐ are proved. The resulting antecedent (from ⓐ) is abbreviated $B \equiv \bigwedge_{p \in \mathcal{S}} V_p < W \land \bigwedge_{p \in \mathcal{U}} V_p < W \exp(\lambda_p \Theta_p)$. The loopT rule is used with the following pre-attractivity loop invariant Inv_a, which yields premises ①–④ shown and proved further below:

$$
\begin{aligned}
\mathit{Inv}_a \equiv\ & \tau \geq 0 \land t \geq \tau \land \\
& \left(\begin{array}{l}
\displaystyle\bigvee_{p \in \mathcal{S}} \big(u = p \land V_p < W \exp(-\sigma(t-\tau)) \exp(-\lambda_p \tau)\big) \lor \\
\displaystyle\bigvee_{p \in \mathcal{U}} \big(u = p \land V_p < W \exp(-\sigma(t-\tau)) \exp(-\lambda_p(\tau - \Theta_p)) \land \tau \leq \Theta_p\big)
\end{array} \right)
\end{aligned}
$$

$$
\begin{array}{r c}
 & \qquad\qquad \textcircled{1} \qquad \textcircled{2} \quad \textcircled{3} \quad \textcircled{4} \\ \cline{2-2}
\text{loopT} & \textcircled{b}, T \geq 0, R, B, t = 0 \vdash [\alpha_{\mathtt{guard}}, t' = 1] \ldots \\ \cline{2-2}
\forall\text{L}, \to\text{L} & \textcircled{a}, \textcircled{b}, T \geq 0, R, \|x\|_2 < \delta, t = 0 \vdash [\alpha_{\mathtt{guard}}, t' = 1] \ldots \\ \cline{2-2}
[;], [:=] & \textcircled{a}, \textcircled{b}, T \geq 0, R, \|x\|_2 < \delta \vdash [t := 0; \alpha_{\mathtt{guard}}, t' = 1] \ldots \\ \cline{2-2}
\forall\text{R}, \to\text{R} & \textcircled{a}, \textcircled{b}, T \geq 0, R \vdash \forall x \left(\|x\|_2 < \delta \to \ldots \right) \\ \cline{2-2}
\exists\text{R}, \mathbb{R}_{\exp} & \varepsilon > 0, W > 0, \textcircled{a}, U > 0, \textcircled{b} \vdash \exists T {\geq} 0 \forall x \left(\|x\|_2 < \delta \to \ldots \right)
\end{array}
$$

Premise ① is proved from B after unfolding α_i using axioms $[\cup]$, $[:=]$ and substituting $\tau = 0$ and $t = 0$ in the loop invariant (using $\exp(0) = 1$).

$$
\begin{array}{r r}
 & * \\ \cline{2-2}
\mathbb{R}_{\exp} & B, t = 0, \tau = 0, u = p \vdash \mathit{Inv}_a \\ \cline{2-2}
[\cup], [:=] & B, t = 0 \vdash [\alpha_i] \mathit{Inv}_a
\end{array}
$$

Premise ④ is proved by unfolding the loop invariant with $\lor$L. This yields two possible premise shapes, corresponding to $p \in \mathcal{S}$ or $p \in \mathcal{U}$. In both cases, assuming the negation of the succedent proves the corresponding implication LHS in the antecedent assumption ⓑ, which gives $V < U$ as an assumption. The remaining arithmetic argument underlying these premises is proved by $\mathbb{R}_{\exp}$ by contradicting assumption $V < U$ for each case resulting from $\lor$L. The cases are shown and explained further below.

$$
\begin{array}{r c}
 & * \\ \cline{2-2}
\lor\text{L}, \mathbb{R}_{\exp} & \textcircled{b}, R, \mathit{Inv}_a \vdash t \geq T \to \|x\|_2 < \varepsilon
\end{array}
$$

For $p \in \mathcal{S}$, the following sequence of inequalities is used:

$$
\begin{aligned}
V_p &< W \exp(-\sigma(t-\tau)) \exp(-\lambda_p \tau) && \text{(from invariant)} \\
&= W \exp(-\sigma t) \exp(-\tau(\lambda_p - \sigma)) && \\
&\leq W \exp(-\sigma T) \exp(-\tau(\lambda_p - \sigma)) && \text{(from } t \geq T, \sigma > 0\text{)} \\
&\leq U \exp(-\tau(\lambda_p - \sigma)) && \text{(from } R\text{)} \\
&\leq U && \text{(from } \sigma < \lambda_p, \tau \geq 0\text{, contradiction)}
\end{aligned}
$$

For $p \in \mathcal{U}$, the following sequence of inequalities is used (note that $\tau \leq \Theta_p$ is in the invariant Inv_a for $p \in \mathcal{U}$):

$$\begin{aligned}
V_p &< W \exp(-\sigma(t-\tau)) \exp(-\lambda_p(\tau - \Theta_p)) && \text{(from invariant)} \\
&\leq W \exp(-\sigma(t-\tau)) && \text{(from } \tau \leq \Theta_p, \lambda_p \leq 0\text{)} \\
&= W \exp(-\sigma t) \exp(\sigma \tau) && \\
&\leq W \exp(-\sigma t) \exp(\sigma \Theta_p) && \text{(from } \sigma > 0, \tau \leq \Theta_p\text{)} \\
&\leq W \exp(-\sigma T) \exp(\sigma \Theta_p) && \text{(from } t \geq T, \sigma > 0\text{)} \\
&\leq U && \text{(from } R\text{, contradiction)}
\end{aligned}$$

The derivation from premise ② unfolds the switching controller α_u in $\alpha_{\mathtt{time}}$ with dL's hybrid program axioms. Similar to the derivation for the stability conjunct, this unfolding yields four possible shapes of premises (abbreviated as $\ldots$ and shown immediately below) for maintaining the invariant Inv_a after a switch from the current mode p to the next mode q.

$$\begin{array}{rc}
 & \cdots \\ \cline{2-2}
{\scriptstyle [\cup],[;],[?],[:=]} & \mathit{Inv}_a, u = p \vdash [\bigcup_{q \in \mathcal{P}} \big(?\theta_{p,q} \leq \tau; \tau := 0; u := q\big) \cup u := u]\mathit{Inv}_a \\ \cline{2-2}
{\scriptstyle [;],[?]} & \mathit{Inv}_a \vdash [?u = p; (\bigcup_{q \in \mathcal{P}} \big(?\theta_{p,q} \leq \tau; \tau := 0; u := q\big) \cup u := u)]\mathit{Inv}_a \\ \cline{2-2}
{\scriptstyle [\cup]} & \mathit{Inv}_a \vdash [\alpha_u]\mathit{Inv}_a
\end{array}$$

$$t \geq \tau, \theta_{p,q} \leq \tau, V_p < W \exp(-\sigma(t-\tau)) \exp(-\lambda_p \tau) \vdash V_q < W \exp(-\sigma t) \qquad (p \in \mathcal{S}, q \in \mathcal{S})$$

$$t \geq \tau, \theta_{p,q} \leq \tau, V_p < W \exp(-\sigma(t-\tau)) \exp(-\lambda_p \tau) \vdash V_q < W \exp(-\sigma t) \exp(\lambda_q \Theta_q) \qquad (p \in \mathcal{S}, q \in \mathcal{U})$$

$$t \geq \tau, \theta_{p,q} \leq \tau, V_p < W \exp(-\sigma(t-\tau)) \exp(-\lambda_p(\tau - \Theta_p)), \tau \leq \Theta_p \vdash V_q < W \exp(-\sigma t) \qquad (p \in \mathcal{U}, q \in \mathcal{S})$$

$$t \geq \tau, \theta_{p,q} \leq \tau, V_p < W \exp(-\sigma(t-\tau)) \exp(-\lambda_p(\tau - \Theta_p)), \tau \leq \Theta_p \vdash V_q < W \exp(-\sigma t) \exp(\lambda_q \Theta_q) \qquad (p \in \mathcal{U}, q \in \mathcal{U})$$

The derivation from premise ③ unfolds the plant model α_p. resulting in two possible shapes of premises, depending if $p \in \mathcal{S}$ or $p \in \mathcal{U}$, which are abbreviated ⑤ and ⑥ respectively. In either case, the key step is to show that the respective upper bound on V_p is preserved along evolution of the ODE.

$$\begin{array}{rc}
 & \text{⑤} \qquad \text{⑥} \\ \cline{2-2}
{\scriptstyle [;],[?]} & \mathit{Inv}_a, u = p \vdash [x' = f_p(x), \tau' = 1, t' = 1 \,\&\, \tau \leq \Theta_p]\mathit{Inv}_a \\ \cline{2-2}
{\scriptstyle [;],[?]} & \mathit{Inv}_a \vdash [?u = p; x' = f_p(x), \tau' = 1, t' = 1 \,\&\, \tau \leq \Theta_p]\mathit{Inv}_a \\ \cline{2-2}
{\scriptstyle [\cup]} & \mathit{Inv}_a \vdash [\alpha_p]\mathit{Inv}_a
\end{array}$$

For premise ⑤, the proof uses $\mathrm{dbx}_{\succcurlyeq}$ with cofactor $-\lambda_p$, with abbreviated term $P_s = W \exp(-\sigma(t-\tau)) \exp(-\lambda_p \tau)$, noting that the Lie derivative of P_s is $-\lambda_p P_s$. This yields the third premise of rule MLF_τ.

$$\begin{array}{rc}
 & * \\ \cline{2-2}
 & \vdash \mathcal{L}_{f_p}(V_p) \leq -\lambda_p V_p \\ \cline{2-2}
{\scriptstyle \mathrm{dbx}_{\succcurlyeq}} & V_p < P_s \vdash [x' = f_p(x), \tau' = 1, t' = 1 \,\&\, \tau \leq \Theta_p] V_p < P_s
\end{array}$$

The proof for premise ⑥ also uses rule $\text{dbx}_{\succcurlyeq}$ with cofactor $-\lambda_p$, with abbreviated term $P_u = W\exp(-\sigma(t-\tau))\exp(-\lambda_p(\tau-\Theta_p))$, noting that the Lie derivative of P_a is $-\lambda_p P_a$. This yields the third premise of rule MLF_τ.

$$\text{dbx}_{\succcurlyeq}\dfrac{\dfrac{*}{\vdash \mathcal{L}_{f_p}(V_p) \le -\lambda_p V_p}}{V_p < P_u \vdash [x' = f_p(x), \tau' = 1, t' = 1 \,\&\, \tau \le \Theta_p] V_p < P_u} \qquad \square$$

D.2 Counterexamples

The cruise controller automaton from Section 6.5.2 is taken from the suite of examples for the Stabhyli tool [116, 117]. Using the default instructions on a Linux machine, Stabhyli generates a success message with the following output (newlines added for readability):

```
...
SOSSolution( Problem is solved. (accepted); ...
...
### Lyapunov template for mode normal_PI: \
  +V_23*relV^2+V_22*intV^2+V_21*intV*relV \
  +V_20*relV+V_19*intV
### Lyapunov function for mode normal_PI: \
  +572572089848357/144115188075855872*intV*relV \
  +256336575597239/281474976710656*relV^2 \
  +6008302119812893/4611686018427387904*intV^2 \
  +5787253314511645/618970019642690137449562112*relV \
  +5661677770976729/39614081257132168796771975168*intV
...
The hybrid system is stable
```

The generated Lyapunov function candidate V does not satisfy all of the required arithmetical conditions for the normal PI mode [116]. For example, one requirement is that it should be non-negative in the mode invariant $-15 \le relV \le 15 \land -500 \le intV \le 500$. It can be checked that $intV = -\frac{1}{17179869184}, relV = 0$ is a counterexample, with $V = -3.90488 \times 10^{-24} < 0$.

A heuristic approach to resolve this numerical issue is to truncate terms in the candidate V with extremely small coefficients and then check the resulting truncated candidate. This heuristic is applied for the case study in Section 6.5.2, where the KeYmaera X proof succeeded using the truncated candidate together with the rest of the Lyapunov function candidates generated by Stabhyli (for other automaton modes).

More interestingly, it is also possible for Stabhyli to declare that a system is stable because of numerical issues even though the system is unstable. Consider the following unstable system with two modes and no switching allowed between the modes:

- $x' = -x + 1 \,\&\, -\frac{1}{1000000000} \le x \le \frac{1}{1000000000}$ which is unstable at the origin and
- $x' = -x$ which is stable.

Stabhyli always examines stability of the origin of the given hybrid system [116]. Using the default instructions as before, Stabhyli generates a success message with the following output (newlines added for readability):

```
...
SOSSolution( Problem is solved. (accepted); ...
...
### Lyapunov template for mode stable: +V_2*x^2+V_1*x
### Lyapunov function for mode stable: \
  +603702977637151/18889465931478580854784*x^2
### Lyapunov template for mode unstable: +V_4*x^2+V_3*x
### Lyapunov function for mode unstable: \
  +457363293760441/18889465931478580854784*x^2
  -224353181720881/77371252455336267181195264*x
The hybrid system is stable
```