

Formal Simulation and Visualisation of Hybrid Programs

An Extension of a Proof-of-Concept Tool

Pedro Mendes

University of Minho, Portugal
pg50685@alunos.uminho.pt

Ricardo Correia

University of Minho, Portugal
pg47607@alunos.uminho.pt

Renato Neves

INESC-TEC & University of Minho, Portugal
nevrenato@di.uminho.pt

José Proença

CISTER, Faculty of Sciences of the University of Porto, Portugal
jose.proenca@fc.up.pt

The design and analysis of systems that combine computational behaviour with physical processes' continuous dynamics – such as movement, velocity, and voltage – is a famous, challenging task. Several theoretical results from programming theory emerged in the last decades to tackle the issue; some of which are the basis of a *proof-of-concept* tool, called Lince, that aids in the analysis of such systems, by presenting simulations of their respective behaviours.

However being a proof-of-concept, the tool is quite limited with respect to usability, and when attempting to apply it to a set of common, concrete problems, involving autonomous driving and others, it either simply cannot simulate them or fails to provide a satisfactory user-experience.

The current work complements the aforementioned theoretical approaches with a more practical perspective, by improving Lince along several dimensions: to name a few, richer syntactic constructs, more operations, more informative plotting systems and errors messages, and a better performance overall. We illustrate our improvements via a variety of examples that involve both autonomous driving and electrical systems.

1 Introduction

Motivation and context. This paper concerns the design and analysis of hybrid systems (*i.e.* those that combine discrete with continuous behaviour) from a programming-oriented perspective. Such a view emerged recently in a series of works [24, 21, 11, 15], and revolves around the idea of importing principles and techniques from programming theory to better handle the behaviour of hybrid systems. In this context programs combine standard program constructs, such as conditionals and while-loops, with certain kinds of differential statement meant to express the dynamics of physical processes, such as time, energy, and motion. Consider the following example of such a program:

$$p' = v, v' = 2 \text{ for } 1 ; p' = v, v' = -2 \text{ for } 1 \quad (1)$$

In a nutshell, it is a sequential composition (;) of two programs where each expresses how the position (p) and velocity (v) of a vehicle evolve over time. The program on the left ($p' = v, v' = 2 \text{ for } 1$) is a differential statement that reads “the vehicle accelerates at the rate of 2m/s^2 for 1 second”. The other program corresponds to a deceleration. Both position and velocity over time are presented in Fig. 1, where we see that the vehicle travels 2 meters and then stops.




Figure 1: Simulation of (1).

Actually there has been a rapid proliferation of such systems, not only in the domain of autonomous driving but also in the medical industry and industrial infrastructures, among others [24, 12, 19, 21]. This spurred extensive research on languages, semantics, and tools for their design and analysis. An example is our work [10, 11] on the semantics of hybrid programs – *i.e.* those that combine program constructs with differential statements, such as in (1) – from which arises a mathematical basis for reasoning about their behaviour, both operationally and denotationally. A proof-of-concept tool, called Lince, also emerged from this: its engine is a previously developed operational semantics [11] that yields trajectories of hybrid programs, just as we saw in Fig. 1. However because our focus was rather theoretical, the tool was not developed with usability in mind, and thus lacks basic features for tackling a broad range of important scenarios. Let us illustrate this problem with a very simple example.

Problem scenario. Suppose that we wish to move a stationary object a distance of $dist$ meters – a basic task in autonomous driving. For simplicity assume that we have access only to the acceleration rates a^{m/s^2} and $-a^{m/s^2}$, where $a > 0$. Our mission can be accomplished by taking the following variation of Eq. (1),

$$p' = v, v' = a \text{ for } t ; p' = v, v' = -a \text{ for } t \quad (2)$$

for a suitable duration t . Then in order to calculate t (*i.e.* the prescribed duration of each differential statement) we simply observe that,

$$dist = \int_0^t v_a(x) dx + \int_0^t v_{-a}(x) dx$$

where $v_a(x) = a \cdot x$ and $v_{-a}(x) = v_a(t) - a \cdot x$ are the velocity functions with respect to the time intervals $[0, t]$ and $[t, 2 \cdot t]$ associated with the program's execution. We now observe, by recalling Fig. 1, that the value $dist$ corresponds to the area of a triangle with basis $2 \cdot t$ and height $v_a(t)$. This geometric shape yields the equations,

$$\begin{cases} dist &= 1/2 \cdot (2 \cdot t) \cdot v_a(t) \text{ (area)} \\ v_a(t) &= a \cdot t \text{ (height)} \end{cases} \implies t = \sqrt{\frac{dist}{a}}$$

Finally observe that if $dist = 3$ and $a = 1$ then $t = \sqrt{3}$. Unfortunately the previous version of Lince does not support square root operations which renders our mission impossible to accomplish.

Contributions and outline. As already alluded to, this paper complements our previous theoretical work on the semantics of hybrid programming [10, 11]. Specifically it improves our proof-of-concept tool Lince so that it can handle a myriad of important scenarios, whilst maintaining both its simplicity and theoretical underpinnings. The improvements were made along different dimensions, and we highlight the most relevant ones next¹.

Extension of basic operations. As illustrated before, the previous version of Lince lacked essential arithmetic operations for handling most basic tasks. Thus as the first main contribution we added standard arithmetic operations, including divisions, trigonometric functions, and square root extractions. Notably the fact that many of these operations are partial required us to extend the operational semantics developed in [11] (the main engine of Lince) with the possibility of failure. The extended semantics is detailed in Section 2 and it is of course the basis of the new engine behind improved Lince.

¹The improved version can be checked online at <http://arcatools.org/lince>.

Extension of numerical methods. Again because our focus in previous work was rather theoretical the previous version of Lince was unable to simulate standard scenarios in hybrid programming. A main reason for this was our method of obtaining solutions of systems of ordinary differential equations (ODEs), which although *exact* lacked in scalability. Precisely for this reason we now integrate a complementary, numerical solver in Lince with the obvious compromise that the solutions obtained for such systems are no longer exact.

The benefits of the extended language (and respective semantics), the numerical solver, and a number of quality-of-life features, are summarised in [Section 3](#) and illustrated with a standard, running example concerning the famous concept of harmonic oscillation.

Extension of visualisation mechanisms. Lince is constituted by two core components: the simulator which, by recurring to the aforementioned operational semantics, parses a received program and presents its output with respect to a *single* time instant. And the visualiser which presents (a sample of) the trajectory over time respective to the program at hand, by querying the operational semantics for a certain sequence of time instants. After trying to properly visualise the behaviour of several types of hybrid program with Lince we identified two major limitations with respect to this architecture. First many real-world problems involve multiple spatial dimensions and thus the described view of trajectories over time is often not the best representation of a hybrid program’s behaviour. Second the user is often interested in observing the overall program behaviour for varying initial conditions, concerning for example position and velocity. We therefore present in [Section 4](#) an improved visualiser for Lince that precisely addresses these two issues. We illustrate it via another classical scenario in autonomous driving, *viz.* manoeuvring around an obstacle.

In [Section 5](#) we illustrate that, whilst keeping its simplicity, Lince can now handle complex central problems in the theory of hybrid systems; we focus specifically on the task of one player pursuing another, *e.g.* a vehicle, a drone, or simply a projectile. Such pursuit games were discussed for example in [\[20, 2, 6, 18\]](#), from an (hybrid-)automata, state-chart, and duration calculus perspective. Here we present a programming-oriented approach. Finally in [Section 6](#) we discuss future work and conclude.

Related work. Several tools for the design and analysis of hybrid systems were already proposed, *e.g.* in the areas of deductive verification [\[24\]](#), model checking [\[3, 8, 4\]](#), simulation [\[16, 9, 15, 11\]](#), and program semantics [\[24, 15, 11\]](#). But only a few are committed to a programming-oriented approach, rooted on formal semantics, and with effective simulation capabilities. The only ones we are aware of are [\[15\]](#) and our own tool Lince [\[11\]](#). Interestingly both cases adopt complementary approaches as well: the former harbours a very powerful concurrent language, particularly well-suited for large-scale distributed systems. The latter, harbouring a sequential while-language, aims at being minimalistic whilst still capturing a broad range of interesting problems on which to study different aspects of (pure) hybrid computation at a suitable abstraction level.

Aside from the obvious pedagogical benefit, our minimalistic approach also allows to capitalise on different programming theories more easily. For example already in [\[11\]](#) we connected our tool to a compositional, denotational semantics – particularly well-suited to study hybrid program equivalence and combinations with other paradigms. An analogous concurrent semantics for [\[15\]](#) would be notoriously more difficult to achieve (*cf.* [\[26, 28\]](#)). Similarly our language is amenable to algebraic reasoning in the style of (weak) Kleene algebras [\[17, 14\]](#) whilst the connection between the latter and concurrent object-oriented programming (as adopted in [\[15\]](#)) is less clear.

2 Lince's Foundations Extended with the Possibility of Failure

We now extend part of Lince's foundations with the possibility of failure. Specifically we present an extension of the language in [11] with partial operations, such as division and square root extraction, and introduce a corresponding operational semantics. As explained in the introduction, such is necessary for extending Lince to 'real-world problems' whilst preserving its merit of having a firm, mathematical basis.

Language. First we postulate a finite set $X = \{x_1, \dots, x_n\}$ of variables and a stock of partial functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$ that contains the usual arithmetic operations. Then we define expressions and boolean conditions via the following BNF grammars,

$$e ::= x \mid f(e, \dots, e) \qquad b ::= e \leq e \mid b \wedge b \mid b \vee b \mid \neg b \mid \text{tt} \mid \text{ff}$$

We omit the explanation of these grammars as they are widely used (see *e.g.* [28, 26]). Next, we qualify as 'linear' those expressions e which aside from the use of variables involve only the operations $+$ and $r \cdot (-)$ for some $r \in \mathbb{R}$. For example the expression $2 \cdot x$ is linear but the expression $x \cdot x$ is not. The concept of linearity is key in the grammar of hybrid programs which we present next.

Programs are built according to the following BNF grammars,

$$\begin{aligned} a &::= x'_1 = \ell_1, \dots, x'_n = \ell_n \text{ for } e \mid x := e \\ p &::= a \mid p; p \mid \text{if } b \text{ then } p \text{ else } p \mid \text{while } b \text{ do } \{ p \} \end{aligned}$$

where the terms ℓ_i ($1 \leq i \leq n$) are linear expressions. We qualify as 'atomic' those hybrid programs that are built according to the first grammar. They can be either classical assignments or *differential* statements as described in the introduction. The linearity constraint is here imposed merely to ensure that the latter kind of statement will always have unique solutions, which renders our semantics more lightweight whilst still being able to treat a broad range of problems (see more details in [11]).

The language of hybrid programs p itself is simply the usual while-language [28, 26] extended with the aforementioned differential statements. It is easy to check that our grammar indeed extends that in the previous version of Lince [11] where *all* expressions involved in the assignments and the durations of differential statements had to be linear. This has of course significant implications in the operational semantics introduced in [11].

Operational semantics. We need a series of preliminaries. First for simplicity we abbreviate differential statements $x'_1 = \ell_1, \dots, x'_n = \ell_n \text{ for } e$ simply to $\vec{x}' = \vec{\ell} \text{ for } e$, where \vec{x}' and $\vec{\ell}$ abbreviate the corresponding vectors of variables $x'_1 \dots x'_n$ and linear expressions $\ell_1 \dots \ell_n$. We call functions of the type $\sigma : X \rightarrow \mathbb{R}$ *environments*; they map variables to the respective valuations. We use the notation $\sigma[\vec{x} \mapsto \vec{v}]$ to denote the environment that maps each x_i in \vec{x} to v_i in \vec{v} and the remaining variables as in σ . Finally we denote by $\phi_\sigma^{\vec{x}' = \vec{\ell}} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ the (unique) solution of a system of differential equations $\vec{x}' = \vec{\ell}$ with σ as the initial condition (recall our previous constraint about linearity which ensures that such solutions indeed exist). When clear from context, we omit both the superscript and subscript in $\phi_\sigma^{\vec{x}' = \vec{\ell}}$. Next, for an expression e the notation $\llbracket e \rrbracket(\sigma)$ denotes the standard (partial) interpretation of expressions [28, 26] according to σ , and analogously for $\llbracket b \rrbracket(\sigma)$ where b is a boolean expression. For example $\llbracket x + 1 \rrbracket(\sigma) = \sigma(x) + 1$ and $\llbracket 1/x \rrbracket(\sigma)$ is undefined if $\sigma(x) = 0$.

We now present an operational semantics for the language. Following traditions in programming theory [22, 28, 26], we present it from two different, complementary perspectives, which gives a much more complete understanding of the language's features. Specifically we present the semantics in two

different styles: one formalises the idea of a machine “running” a hybrid program and describes its *step-by-step evolution*. The other abstracts away from all *intermediate steps* of this machine and is therefore generally more suitable to reason about “input-output behaviours” (although we do not explore such a feature here). Whilst the former style is the basis of Lince’s new version, the latter style is conceptually more intuitive and therefore we present it first. The current section concludes with a proof that both semantics are in fact equivalent. The curious reader can consult for example [28, 26] for a thorough account on the key differences between the small-step and big-step styles in general program semantics.

Our ‘big-step’ operational semantics is given by an ‘input-output’ relation \Downarrow which relates programs p , environments σ , and time instants t to outputs v . The expression $p, \sigma, t \Downarrow v$ can be read as “at time instant t the program p starting from state σ outputs v ”. The relation \Downarrow is built inductively according to the rules in Fig. 2. Specifically the first three rules describe how differential statements are evaluated: first one computes the duration $\llbracket e \rrbracket(\sigma)$ of the differential statement at hand and an error is raised if $\llbracket e \rrbracket(\sigma)$ is undefined; otherwise the output v is the respective modified state (as dictated by the differential statement) paired with one of the flags *stop* or *skip*. Intuitively the flag *stop* indicates that we ‘reached’ the time instant at which the program needs to be evaluated and therefore the evaluation can stop moving forward in time, which fact is reflected in rule (**seq-stop**). The flag *skip* is simply the negation of *stop*. The remaining rules follow analogous principles and therefore we refrain from detailing them – instead we will briefly show how the semantics works via instructive, concrete examples.

Example 2.1. Let us consider the following very simple program,

$$x' = -1 \text{ for } 1 ; x := 1/x$$

which continuously decreases the value of variable x during 1 second and then applies the (discrete) operation $x := 1/x$. Suppose as well that our initial state is the environment σ defined by $x \mapsto 1$. Then by an application of rule (**diff-stop**) one deduces that this program outputs the environment $x \mapsto 1 - t$ at every time instant $t < 1$. On the other hand, by an application of rules (**diff-skip**), (**asg-err**), and (**seq-skip**) one easily deduces that the evaluation of the program fails at every time instant $t \geq 1$, due to a division by 0.

Notably the fact that failure occurs only at the time instants $t \geq 1$ is a fundamental difference with respect to the famous hybrid programming language detailed in [24]. In the *op. cit.* the language was designed in the spirit of Kleene algebra, which in particular forces the previous program to be *indistinguishable* from *e.g.* the program $x := x/0$. Whilst such a feature could be desirable in some verification scenarios it is clearly unnatural in a simulation-based environment such as ours.

Let us continue unravelling prominent features of our semantics with another example. Consider the following hybrid program,

$$\text{while } x \neq 0 \text{ do } \{ x' = -1 \text{ for } x/2 \} ; x := 1/x$$

paired with the environment $x \mapsto 1$. This program is an instance of a so-called Zeno loop: *viz.* the loop involved unfolds *infinitely* many times with the duration of each iteration becoming shorter and shorter (see details *e.g.* in [11]). In this particular case it is straightforward to check that the duration of the i -th iteration is given by $1/2^i$, and thus that the total duration $\sum_{i=1}^{\infty} 1/2^i$ of the loop will be 1. Now, by applying the operational rules in Fig. 2 one can successfully evaluate the program at every time instant $t < 1$ (intuitively because every such t is reached in a *finite* number of iterations). The same is *false* for time instant $t = 1$ since such requires a complete unfolding of the loop which is of course computationally unfeasible. Thus operationally the potential point of failure $x := 1/x$ in the program above never occurs, as the Zeno loop makes it impossible to actually reach this command in the evaluation. These infinite

$$\begin{array}{c}
\text{(diff-skip)} \quad \frac{\llbracket e \rrbracket(\sigma) = t}{\vec{x}' = \vec{\ell} \text{ for } e, \sigma, t \Downarrow \text{skip}, \sigma[\vec{x} \mapsto \phi(t)]} \\
\text{(diff-stop)} \quad \frac{\llbracket e \rrbracket(\sigma) > t}{\vec{x}' = \vec{\ell} \text{ for } e, \sigma, t \Downarrow \text{stop}, \sigma[\vec{x} \mapsto \phi(t)]} \qquad \text{(diff-err)} \quad \frac{\llbracket e \rrbracket(\sigma) \text{ undefined}}{\vec{x}' = \vec{\ell} \text{ for } e, \sigma, t \Downarrow \text{err}} \\
\text{(asg-skip)} \quad \frac{\llbracket e \rrbracket(\sigma) \text{ defined}}{x := e, \sigma, 0 \Downarrow \text{skip}, \sigma[x \mapsto \llbracket e \rrbracket(\sigma)]} \qquad \text{(asg-err)} \quad \frac{\llbracket e \rrbracket(\sigma) \text{ undefined}}{x := e, \sigma, t \Downarrow \text{err}} \\
\text{(seq-skip)} \quad \frac{p, \sigma, t \Downarrow \text{skip}, \tau \quad q, \tau, u \Downarrow v}{p; q, \sigma, t + u \Downarrow v} \\
\text{(seq-stop)} \quad \frac{p, \sigma, t \Downarrow \text{stop}, \tau}{p; q, \sigma, t \Downarrow \text{stop}, \tau} \qquad \text{(seq-err)} \quad \frac{p, \sigma, t \Downarrow \text{err}}{p; q, \sigma, t \Downarrow \text{err}} \\
\text{(if-true)} \quad \frac{\llbracket b \rrbracket(\sigma) = \text{tt} \quad p, \sigma, t \Downarrow v}{\text{if } b \text{ then } p \text{ else } q, \sigma, t \Downarrow v} \\
\text{(if-false)} \quad \frac{\llbracket b \rrbracket(\sigma) = \text{ff} \quad q, \sigma, t \Downarrow v}{\text{if } b \text{ then } p \text{ else } q, \sigma, t \Downarrow v} \qquad \text{(if-err)} \quad \frac{\llbracket b \rrbracket(\sigma) \text{ undefined}}{\text{if } b \text{ then } p \text{ else } q, \sigma, t \Downarrow \text{err}} \\
\text{(wh-true)} \quad \frac{\llbracket b \rrbracket(\sigma) = \text{tt} \quad p; \text{while } b \text{ do } \{ p \}, \sigma, t \Downarrow v}{\text{while } b \text{ do } \{ p \}, \sigma, t \Downarrow v} \\
\text{(wh-false)} \quad \frac{\llbracket b \rrbracket(\sigma) = \text{ff}}{\text{while } b \text{ do } \{ p \}, \sigma, 0 \Downarrow \text{skip}, \sigma} \qquad \text{(wh-err)} \quad \frac{\llbracket b \rrbracket(\sigma) \text{ undefined}}{\text{while } b \text{ do } \{ p \}, \sigma, t \Downarrow \text{err}}
\end{array}$$

Figure 2: Extension of the big-step operational semantics in [11] with the possibility of failure.

behaviours are bounded in Lince by manually setting limits on the total time and on the number of unfoldings of while-loops, adjustable for each program.

Next, the semantics in the aforementioned ‘small-step’ style is given in the form of a relation \rightarrow that is defined inductively according to the rules in Fig. 3. These rules follow an analogous reasoning to the ones in Fig. 2 so we refrain from repeating explanations.

As detailed in Corollary 1 our small-step semantics is deterministic. This is of course a key property in what concerns its implementation and subsequent use in Lince for simulating hybrid programs. The corollary is based on the following theorem.

Theorem 2.1. For every program p , environment σ , and time instant t there is *at most one* applicable reduction rule.

Let \rightarrow^* be the transitive closure of the small-step relation \rightarrow that was previously presented. Intuitively \rightarrow^* represents an evaluation of one or more steps according to the small-step semantics. If $p, \sigma, t \rightarrow^* v$ we call v ‘non-terminal’ whenever it is of the form p', σ', t' for some hybrid program p' , environment σ' , and time instant t' ; we call v ‘terminal’ otherwise.

(asg[→])	$x := e, \sigma, t \rightarrow skip, \sigma[x \mapsto \llbracket e \rrbracket(\sigma)], t$	<i>(if</i> $\llbracket e \rrbracket(\sigma)$ <i>defined)</i>
(asg-err[→])	$x := e, \sigma, t \rightarrow err$	<i>(if</i> $\llbracket e \rrbracket(\sigma)$ <i>undefined)</i>
(diff-stop[→])	$\vec{x}' = \vec{\ell} \text{ for } e, \sigma, t \rightarrow stop, \sigma[\vec{x} \mapsto \phi(t)], 0$	<i>(if</i> $\llbracket e \rrbracket(\sigma) > t$ <i>)</i>
(diff-skip[→])	$\vec{x}' = \vec{\ell} \text{ for } e, \sigma, t \rightarrow skip, \sigma[\vec{x} \mapsto \sigma(t)], t - \llbracket e \rrbracket(\sigma)$	<i>(if</i> $\llbracket e \rrbracket(\sigma) \leq t$ <i>)</i>
(diff-err[→])	$\vec{x}' = \vec{\ell} \text{ for } e, \sigma, t \rightarrow err$	<i>(if</i> $\llbracket e \rrbracket(\sigma)$ <i>undefined)</i>
(if-true[→])	if b then p else $q, \sigma, t \rightarrow p, \sigma, t$	<i>(if</i> $\llbracket b \rrbracket(\sigma) = tt$ <i>)</i>
(if-false[→])	if b then p else $q, \sigma, t \rightarrow q, \sigma, t$	<i>(if</i> $\llbracket b \rrbracket(\sigma) = ff$ <i>)</i>
(if-err[→])	if b then p else $q, \sigma, t \rightarrow err$	<i>(if</i> $\llbracket b \rrbracket(\sigma)$ <i>undefined)</i>
(wh-true[→])	while b do $\{ p \}, \sigma, t \rightarrow p; \text{while } b \text{ do } \{ p \}, \sigma, t$	<i>(if</i> $\llbracket b \rrbracket(\sigma) = tt$ <i>)</i>
(wh-false[→])	while b do $\{ p \}, \sigma, t \rightarrow skip, \sigma, t$	<i>(if</i> $\llbracket b \rrbracket(\sigma) = ff$ <i>)</i>
(wh-err[→])	while b do $\{ p \}, \sigma, t \rightarrow err$	<i>(if</i> $\llbracket b \rrbracket(\sigma)$ <i>undefined)</i>
(seq-stop[→])	$\frac{p, \sigma, t \rightarrow stop, \sigma', t'}{p; q, \sigma, t \rightarrow stop, \sigma', t'}$	(seq-skip[→]) $\frac{p, \sigma, t \rightarrow skip, \sigma', t'}{p; q, \sigma, t \rightarrow q, \sigma', t'}$
(seq-err[→])	$\frac{p, \sigma, t \rightarrow err}{p; q, \sigma, t \rightarrow err}$	(seq[→]) $\frac{p, \sigma, t \rightarrow p', \sigma', t'}{p; q, \sigma, t \rightarrow p'; q, \sigma', t'} \quad (\text{if } p' \neq stop \text{ and } p' \neq skip)$

Figure 3: Extension of the small-step operational semantics in [11] with the possibility of failure.

Corollary 1 (Determinism). Consider a program p , an environment σ , and a time instant t . If $p, \sigma, t \rightarrow^* v$ and $p, \sigma, t \rightarrow^* u$ with both v and u terminal then we have $v = u$.

Proof. Follows by induction on the number of reduction steps and Theorem 2.1. □

Next we will show that the small-step semantics and its big-step counterpart are indeed equivalent. We will use the two following results for this effect.

Lemma 2.1. Given a program p , an environment σ and a time instant t

1. if $p, \sigma, t \rightarrow p', \sigma', t'$ and $p', \sigma', t' \Downarrow skip, \sigma''$ then $p, \sigma, t \Downarrow skip, \sigma''$;
2. if $p, \sigma, t \rightarrow p', \sigma', t'$ and $p', \sigma', t' \Downarrow stop, \sigma''$ then $p, \sigma, t \Downarrow stop, \sigma''$;
3. if $p, \sigma, t \rightarrow p', \sigma', t'$ and $p', \sigma', t' \Downarrow err$ then $p, \sigma, t \Downarrow err$;

Proof. Follows by induction over the rules concerning the small-step relation. □

Proposition 1. For all program p and q , environments σ and σ' , and time instants t, t' and s , if $p, \sigma, t \rightarrow q, \sigma', t'$ then $p, \sigma, t+s \rightarrow q, \sigma', t'+s$; and if $p, \sigma, t \rightarrow skip, \sigma', t'$ then $p, \sigma, t+s \rightarrow skip, \sigma', t'+s$. If $p, \sigma, t \rightarrow err$ then $p, \sigma, t+s \rightarrow err$

Proof. Follows straightforwardly by induction over the rules concerning the small-step relation and the algebraic properties of addition. \square

Theorem 2.2 (Equivalence). The small-step semantics and the big-step semantics are related in the following manner. Given a program p , an environment σ and a time instant t

1. $p, \sigma, t \Downarrow skip, \sigma'$ iff $p, \sigma, t \rightarrow^* skip, \sigma', 0$;
2. $p, \sigma, t \Downarrow stop, \sigma'$ iff $p, \sigma, t \rightarrow^* stop, \sigma', 0$;
3. $p, \sigma, t \Downarrow err$ iff $p, \sigma, t \rightarrow^* err$.

Proof. The right-to-left direction is obtained by induction over the length of the small-step reduction sequence using [Lemma 2.1](#). The left-to-right direction follows by induction over the big-step derivations together with [Proposition 1](#). \square

3 An Improved Simulator for Hybrid Programs

This section summarises several improvements made to Lince’s simulator of hybrid programs since its original publication [11]. These include (1) more expressive assignments and durations in differential statements (by virtue of the results in the preceding section); (2) a more user-friendly program syntax (by means of syntactic sugar); (3) more informative error messages; and (4) a numerical solver of systems of ordinary differential equations. In order to render our summary more lively we complement it with a running example involving an RLC circuit in series with an On-Off source. It is designed to stabilise voltage across the capacitor in the circuit at a specific value.

Running example: RLC circuits and harmonic oscillation. We present in [Fig. 4](#) the simulation of an *RLC circuit in series* (RLCS). This simulation models an electric system composed of a resistor, a capacitor, an inductor, and a power source connected in series. The power source strategically switches on and off, as a way to stabilise voltage across the capacitor at a target value (say, 10V). Such systems are known to yield interesting results that are practically relevant for energy storage voltage control systems, which help to mitigate voltage imbalances that could otherwise damage electronic equipment. More details about such circuits and associated differential equations are available for example in [29, 13]. We present in [Fig. 4](#) two variations of an RLCS circuit: one in which the capacitor voltage is in an underdamped regime – with a resistance r_U of 0.5Ω , a capacitance c of $0.047F$, and an inductance l of $0.047H$ – and one in which the capacitor voltage is in an overdamped regime – with a resistance r_0 of 4Ω and the same values as before for the capacitance and inductance. The general idea of our program is that the associated controller will read the voltage across the capacitor (variable `under` for the underdamped case, `over` for the overdamped one) every 0.01 seconds, and set the voltage at the source either to 0 (off) or 18V (on) depending on the value read.

Improvement’s summary. The program just described is highly problematic for the original version of Lince. This is due to two fundamental reasons related to the ODEs involved: specifically (1) the equations used in the ODEs violate the linearity condition presented in [Section 2](#) (they include variable multiplications); and (2) the original solver of ODEs, mentioned in the introduction, fails to produce




Figure 4: Hybrid program (left) and its plot (right) of two variations of an RLC circuit that tries to maintain the voltage in the capacitor at 10V.

solutions after few iterations, due to the sheer, exponential growth of the involved expressions' size. We detail these issues and others next.

Richer expressions. As illustrated in the introduction and in the previous RLCS example, there are several essential, non-linear operations that are necessary to accommodate if one wishes to employ Lince in the analysis of diverse, common hybrid scenarios. We therefore now permit non-linear expressions outside of ODEs, essentially by using as basis the grammar of hybrid programs that was described in Section 2. Thus expressions outside the ODEs can now include for example the operations: division and multiplication of variables, more complex mathematical functions (such as square root extraction, exponentials, logarithms, minimum/maximum, and (co)sine), and mathematical constants (namely π and Euler's constant).

As for expressions inside ODEs, the linearity constraint is kept but the associated parser is much less rigid. A core feature is that it now tries to convert non-linear expressions into equivalent linear ones via algebraic laws. For example, it converts the expression $x \cdot 5$, which syntactically is not a linear expression, into the linear one $5 \cdot x$ since multiplication is commutative. Most notably, it converts non-linear expressions $x \cdot y$ into scalar multiplications $s \cdot x$ or $s \cdot y$ if it can infer that either x or y is a constant with value s . Such a feature is critical in our RLCS example, where we multiply variables in the respective ODEs.

More informative error messages. Several errors were undetected at an early stage of the simulation process, which resulted in unintelligible error messages in many situations. We thus added and improved the detection and notification of several key errors occurring in typical usages of Lince, including when: (1) a partial function fails (such as in division by 0); (2) a variable is not properly initialised; (3) the number of arguments of a function is incorrect; (4) the solver fails to solve a system of ODEs; and (5) ODEs contain non-linear expressions after de-sugaring. For example, in our RLCS simulation when defining c to be 0 we now obtain the error “Error: the divisor of the division 'rU(c)' is zero.”. In our experience, this more precise detection and notification of errors drastically improved user experience.

Numerical solver. As already mentioned, several hybrid programs such as our RLCS example cannot be properly handled by the (exact) solver of ODEs (*viz.* SageMath [27]) used by Lince. We have therefore implemented an alternative, numerical solver based on the popular fourth-order Runge-Kutta method. At the theoretical level, this only required a small adaptation of the operational semantics presented

in Section 2. Specifically we no longer assume that the solution $\phi_{\sigma}^{\vec{x}=\vec{\ell}}$ associated to a system of ODEs $\vec{x}' = \vec{\ell}$ and an initial condition σ is exact. At the practical level, this allowed us to keep the size of expressions involved in computations highly manageable thus allowing Lince to cover a broader range of examples such as the RLCS.

4 An Improved Visualiser for Hybrid Programs

Many hybrid programs cannot be easily understood by simply plotting values of variables over time. For example, in some cases one may wish to analyse the movement of a vehicle in a 2D plane, or to analyse how its behaviour varies due to changes in its initial position and velocity. This section presents an extension of Lince’s visualisation capabilities in these two directions. In the same spirit of the preceding section, we complement our description with a running example.

Running example: avoiding and manoeuvring around obstacles. The *Automatic Emergency Braking* (AEB) system is an autonomous driving device that after reading its distance to an obstacle and its current velocity, decides whether to decelerate until stopping [1]. Here we present a more advanced version of the AEB that after stopping also manoeuvres around the obstacle – clearly a process involving two or even three spatial dimensions. Such a system is called *Automatic Emergency Braking with an Overtaking Manoeuvre* (AEBOM).

The continuous dynamics of the AEBOM (*i.e.* the differential equations involved) is typically given by Dubins dynamics which essentially describe the object’s orientation over time (an angle) and its effect on the object’s velocity along the different spatial dimensions [25]. We adopt this approach as well. For simplicity we additionally assume that our object is a robot that is able to rotate around itself. The overall process of our AEBOM is thus as follows: move forward until detecting the obstacle and in which case decelerate until stopping; then rotate to the left and move forward a prescribed number of meters (that depends on the obstacle’s size); then rotate right and move forward again a prescribed number of meters; and finally repeat the last step.

Figure 5 depicts the original visualisation of the AEBOM simulation on the left, and a customised 2D visualisation that uses our extension on the right. The respective implementation of the AEBOM, included in Lince online, is not relevant to show at this stage, because our focus is at the moment on describing new visualisation mechanisms and not features concerning code. Observe as well that the plot on the right provides novel insights with respect to the one on the left: whilst in the right it is clear that the robot does not collide with the obstacle and performs the overtaking manoeuvre safely, in the left it is much harder to see that the same occurs. We provide more details about our improved plotting system next.

Higher-dimensional trajectories and beyond. Our new visualisation framework in Lince uses the Plotly JavaScript library to display plots². Among other things, we now support 2D and 3D scatter plots, and include dedicated markers such as the large circles indicating the start and end points of trajectories. When hovering over these markers, extra information is displayed, *e.g.* the respective values, relevant information about the conditionals involved, and potential warnings. We also exploit the animation functionality of Plotly in plots that do not include the time component, by moving a highlighting circle through the trajectories capturing how values vary throughout time. This feature is active by default. To take all these possibilities into account, Lince allows the user to adjust different settings of the plot under

²<http://plotly.com/>




Figure 5: Plot of AEBOM using the traditional plotting system in Lince (left), and a new customised 2D plot (right) relating x with y (the robot's coordinates) and x_1 with y_1 (the obstacle's coordinates).

analysis so that she can obtain the best possible configuration for her needs. We very briefly detail such settings next:

- *Axis*: Allows defining the relationships between variables which will automatically be presented in the respective plots. For example, by setting $[x, y, v]$, if the graph type is scatter, three separate graphs will be generated where the vertical axis represents each of the variables x , y , and v , while the horizontal axis represents time. Choosing which variables to map to the axes is crucial for proper data analysis, allowing direct visual comparisons between different variables over time or with each other.
- *Max Time*: Refers to the duration of the simulation.
- *Max Iterations*: Specifies the maximum number of iterations (in while-loops) that the simulation can perform.
- *Graph Type*: Defines the type of graph to be used for visualising the simulation data, by selecting from the available types ('scatter' or 'scatter3d'). In a nutshell, a scatter plot is a 2D graph used to display the relationship between two variables, with data points plotted in the two-dimensional plane. Scatter3D serves the same purpose but involves three variables, with data points plotted in the three-dimensional space.

The summarised settings are presented in Fig. 6, where the values there listed are the ones used to obtain the plot in Fig. 5 on the right.

Axis	↻
[(x,y),(x1,y1)]	▲
Max Time (Limit of 150)	↻
50	▲
Max Iterations (Limit of 1000)	↻
1000	▲
Graph Type	↻
scatter	▲

Figure 6: Input boxes that allow for the configuration of the visualisation.

Variability of initial conditions. As mentioned before, it is highly relevant take into account how the behaviour of a hybrid program varies due to changes in its initial conditions. In the AEBOM previously described in particular, it is of fundamental importance to understand how the robot manoeuvres around

an obstacle with respect to different initial positions and velocities – for it is unrealistic to expect that it moves with well-known, exact conditions. A similar, more general discussion can be consulted in [25].

In order to address this aspect we extended Lince in two steps: first its syntax now allows the listing of different initial conditions at the same time. Such is illustrated in Fig. 7 on the left, with a snippet of code used to specify initial values with respect to our robot in the AEBOM example. The latter’s initial position (x,y) for example, can now be either $(0,0)$, $(2,0)$, or $(4,0)$; and similarly we have different initial velocities (v_x) towards the obstacle, 4, 8, and 12m/s . Second Lince now pre-processes such listings in the code and derives all possible combinations of initial conditions, which of course yields several hybrid programs at once (in the standard syntax). These data is then fed into Lince’s visualiser which presents multiple simulations overlapped in the same plot. Such is seen in Fig. 7 on the right, again with our AEBOM example, where we see that our robot behaves in the same way under different initial conditions.




Figure 7: Visualisation of multiple simulations overlapped concerning the AEBOM.

5 Lince at Work: a Showcase of the Overall List of Improvements

This section illustrates the overall list of improvements made to Lince (as described in the preceding sections) working together in the design and analysis of a complex hybrid scenario – specifically we focus on a multi-dimensional pursuit game between two players (for example two drones) [20, 2, 6, 18]. Our illustration focuses mainly on two aspects: (1) Lince’s capability to simulate such scenarios, with optimally configured 3D plotting systems; and (2) the time that Lince takes to simulate increasingly larger systems, to provide insights over limitations of the current implementation.

Pursuit Games. Pursuit games are a captivating class of problems involving multiple agents, where at least one them (the pursuer) aims to capture or reach another (the evader) [20, 2, 6, 18, 25]. Such games are extensively studied across various disciplines, including mathematics, game theory, robotics, and computer science, due to their practical and theoretical significance. Indeed they model a wide range of real-world situations, from military and security operations to animal behaviour and industrial applications.

In this section we explore a specific 3D pursuit game, where we perceive the pursuer as a drone that attempts to capture another one in the three-dimensional space. This scenario is particularly challenging, due to the additional complexity introduced by the third dimension which requires a higher level of planning and coordination between the drones’ movements. In order to model this problem we base our game’s continuous dynamics on Dubins dynamics [25], *i.e.* as in Section 4 but now in three dimensions.

Our overarching strategy for the pursuer is to simply point its orientation to the evader’s position at every iteration in a certain while-loop. Of course there are other options, such as that of (variations of)

Dubins paths [25, 5], but our version already suffices to properly illustrate Lince at work. Technically our approach utilises the angular velocity tensor to perform 3D infinitesimal rotations [7]. Additionally we use the cross product between the projection of the relative velocity vector and the relative position vector in each plane to determine the orientation of rotation among the three axes. We do not show here the coding details of all these processes, since this is unnecessary for our illustration. However the interested reader can consult details about these in [5, 7], and the complete code of our program is included in the examples available in Lince online.

We now show the simulation of our game in Lince across different scenarios. In the first case, the pursuer starts from the position $(300, 300, 600)$ with a velocity of $(-20, -10, 0)m/s$, while the evader begins at the position $(600, 600, 500)$ with a velocity of $(10, 0, 10)m/s$. The pursuer's angular velocity along each axis is $(1/20)*2*\pi(\text{rad})/s$ (20 seconds to complete a full rotation); and for the evader $(1/40)*2*\pi(\text{rad})/s$ (40 seconds to complete a full rotation). The pursuer is allowed to actuate every 0.1s, and it wins the game if it reaches a distance of less than one meter with respect to the evader. Finally, for simplicity we assume a pre-defined set of movements for the latter player. Using these parameters, we simulated the corresponding program in Lince and generated a 3D scatter plot of the positional variables for both the pursuer and the evader, resulting in the graphical representation shown in the Fig. 8 after 73 seconds.




Figure 8: Two views of the same plot, where a pursuer (blue) captures an evader (orange).

We can see that the decision strategy for the pursuer adopted in this hybrid program successfully guided it to the evader, resulting in a capture at the position $(691.26, 441.92, 561.12)$ after 27.7 seconds. However if we change the initial velocity of the evader to a higher value, such as $(20, 0, 9)m/s$, we no longer can visualise the capture of the evader within the limits used for this simulation (Fig. 9). Indeed, Lince supports the customisation of bounds both on the maximal time and on the number of times loops are unfolded, to avoid infinite computations. In this case, using larger bounds would allow the pursuer to capture the evader in the plot.

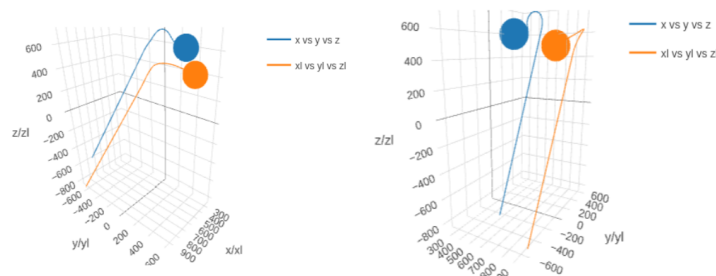


Figure 9: Similar plot to the one in Fig. 8, but using different initial velocities while keeping the same bounds on the size of the plot; this leaves out the point of the capture.

Finally by taking advantage of the variability results presented in Section 4 we very briefly study the




Figure 10: Two simulations (left and middle) of a pursuit game using different initial velocities $((1/40)*2*\pi \text{ rad/s})$ and $((1/100)*2*\pi \text{ rad/s})$, respectively); the right plot depicts both simulations overlaid.

effects of using different velocities in this pursuit game. Specifically we adjust the angular velocity of the pursuer along each axis to be either $(1/40)*2*\pi \text{ rad/s}$ or $(1/100)*2*\pi \text{ rad/s}$, whilst keeping all other aspects. The resulting graphical representation (after 220 seconds) is shown in Fig. 10. From the plots we observe that the pursuer successfully captures the evader when the angular velocity is $(1/40)*2*\pi \text{ rad/s}$ at the position $(692.07, 415.62, 464.63)$ in 34.8 seconds (left plot). However with an angular velocity of $(1/100)*2*\pi \text{ rad/s}$, the pursuer does not capture the evader in this time frame (middle plot). These simulations showcase Lince’s ability to model and simulate complex scenarios, thus providing valuable insights into a system’s behavior.

A brief overview of Lince’s time performance. As shown in the previous example, Lince still has a few limitations concerning performance. In order to give the reader a more concrete idea of them we provide next an overview of how Lince fares performance-wise against the examples presented in this paper. First we need to give further context on how Lince operates.

The first main observation is that now that Lince is equipped with an effective numerical solver (recall Section 3) it can operate in two starkly different ways: one analytical with exact methods that rely on SageMath’s framework [27], the other numerical, based on progressively closer approximations as described in Section 3. Both operation modes have significant differences performance-wise: most notably the former is obviously slower and gives timeouts much more frequently than the latter (recall our RLCS example in Section 3). Interestingly the bottleneck hinges not only on the employment of a precise solver, but also on the fact that:

1. this solver is external to Lince, specifically our tool needs to interact with a server, with all the usual delays that this implies;
2. along the evaluation of a hybrid program, Lince needs to simplify resulting expressions over and over to make them tractable (due to them being symbolic and not numerical).

We saw first-hand in Section 3 how all these extra tasks running behind the curtains inhibit Lince to simulate programs such as the RCLS circuit. The numerical solver, on the other hand, avoid these problems, but at the cost of less precision which may have deep implications if one wishes to have full guarantees that a simulation is correct, particularly if the system at hand is chaotic [23]. Needless to say, to find methods that have the virtues of both approaches is a very interesting challenge.

Table 1 lists several execution times of Lince against different variations of the examples presented in the paper. More specifically, each row represents one of our three examples with varying sampling times and total number of iterations. The example AEB is a variation of AEBOM, where the vehicle stops instead of performing an overtaking manoeuvre. All these examples are fully available in our improved Lince online.

Table 1: An overview of Lince’s time performance with respect to the examples discussed in this paper. We consider different sampling times, number of iterations, and both exact and approximate methods.

	Sampling Time	N\checkmark of Iterations	Time Symb-Server	Time Symb-Total	Time Numerical-Total
RLCS	0.01s	1000	-	-	11.46s
	0.1s	1000	-	-	10.98s
	1s	150	-	-	1.14s
AEB	0.01s	184	23.56s	23.70s	0.41s
	0.1s	19	13.04s	13.08s	0.18s
	1s	2	11.90s	11.97s	0.14s
AEBOM	0.01s	1000	-	-	8.85s
	0.1s	128	-	-	0.62s
	1s	21	-	-	0.35s
Pursuit Games	0.01s	1000	-	-	66.60s
	0.1s	322	-	-	18.26s
	1s	150	-	-	7.85s

We used a Linux laptop with a Intel quad-core i5 processor and 16GB RAM running both the server and the client. The columns *Sampling time* and *N \checkmark of Iterations* refer respectively to the rate at which computational tasks need to be performed and the total number of times the while-loop in the program involved is unfolded. The column *Time Symb-Total* presents the time since a new program is loaded, before parsing, until the plot is displayed in the browser. The column *Time Symb-Server* measures only the time taken since the launch of a dedicated process running SageMath until it is terminated at the end of a trajectory. The column *Time Numerical-Total* measures the time taken since a program is loaded until its plot is displayed, computed using numerical approximations. Some observations over the values on [Table 1](#) follow below.

- Most examples, except for AEB, reach a timeout (set in our server) when using the symbolic analysis, marked in the table with “-”. The feasibility of AEB is mainly due to the smaller number of required calls to the symbolic engine.
- In the AEB example we observe that, when using exact methods, around 99% of the total time was spent by the computations at the server.
- The numerical mechanisms in the AEB example yield simulations significantly faster than in the exact counterpart.
- The total time taken to numerically simulate the RLCS and AEBOM examples are shorter than in the Pursuit Games example. This is because these two examples involve fewer computations and the Pursuit Games use a 3D scatter plot, which is more computationally intensive than the 2D scatter plot.
- Larger sampling times imply reduced times in generating both the exact and numerical plots, due to the decreased number of computational operations. Consequently, it takes longer to simulate controllers with higher precision that actuate on physical processes such as movement, velocity,

and time. However, many critical systems, e.g., in the context of autonomous driving and other embedded systems, may require such a high precision.

6 Conclusion and Future work

We presented an improved version of Lince, which can now handle a broader class of hybrid programs and aims overall at improving user experience. As previously discussed, this required an extension with the possibility of failure of the operational semantics introduced in [11], the implementation of an efficient numerical solver, and more informative error messages, among other things.

We believe that our work opens up several research paths that we would like to explore next. For example, thanks to the numerical solver it is now straightforward to extend our language with non-linear differential equations, which widens even more the range of programs that Lince can currently tackle. Another interesting research path is the addition of probabilistic constructs to Lince, such as measure sampling. We conjecture that this could be handled easily in Lince via a random-number generator and part of the implemented variability mechanisms that were presented in Section 4.

Yet another interesting research line is to connect Lince to the theorem prover for hybrid programs KeYmaera X [25] – specifically the connection would consist of a suitable encoding from programs written in Lince to programs written in KeYmaera X. Such would establish a workflow in which the engineer first *analyses* a given hybrid program via simulation mechanisms (provided by Lince) and subsequently *proves* properties about this program (e.g. correctness) via KeYmaera X.

Acknowledgments. This work is financed by National Funds through FCT - Fundação para a Ciência e a Tecnologia, I.P. (Portuguese Foundation for Science and Technology) within the project IBEX, with reference 10.54499/PTDC/CCI-COM/4280/2021. This work is also partially supported by National Funds through FCT/MCTES, within the CISTER Unit (UIDP/UIDB/04234/2020); and by the EU/Next Generation, within the Recovery and Resilience Plan, within project Route 25 (TRB/2022/00061 – C645463824-00000063).

References

- [1] Proctor Acura: *Technology Guide: What is an Automatic Braking System?* <https://www.proctoracura.com/automatic-braking-system-guide>.
- [2] T. Anderson, R. de Lemos, J. S. Fitzgerald & A. Saeed (1993): *On formal support for industrial-scale requirements analysis*. In Robert L. Grossman, Anil Nerode, Anders P. Ravn & Hans Rischel, editors: *Hybrid Systems*, Springer Berlin Heidelberg, pp. 426–451, doi:[10.1007/3-540-57318-6_39](https://doi.org/10.1007/3-540-57318-6_39).
- [3] Paolo Ballarini, Hilal Djafri, Marie Duflot, Serge Haddad & Nihal Pekergin (2011): *COSMOS: A Statistical Model Checker for the Hybrid Automata Stochastic Logic*. In: *Eighth International Conference on Quantitative Evaluation of Systems, QEST 2011*, IEEE Computer Society, pp. 143–144, doi:[10.1109/QEST.2011.24](https://doi.org/10.1109/QEST.2011.24).
- [4] Davide Bresolin, Luca Geretti, Tiziano Villa & Pieter Collins (2015): *An Introduction to the Verification of Hybrid Systems Using Ariadne*, pp. 339–346. *Lecture Notes in Control and Information Sciences*, Springer, doi:[10.1007/978-3-319-10407-2_39](https://doi.org/10.1007/978-3-319-10407-2_39).
- [5] Xuan-Nam Bui, J.-D. Boissonnat, P. Soueres & J.-P. Laumond (1994): *Shortest path synthesis for Dubins non-holonomic robot*. In: *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pp. 2–7 vol.1, doi:[10.1109/ROBOT.1994.351019](https://doi.org/10.1109/ROBOT.1994.351019).
- [6] Zhou Chaochen, Anders P. Ravn & Michael R. Hansen (1992): *An Extended Duration Calculus for Hybrid Real-Time Systems*. In Robert L. Grossman, Anil Nerode, Anders P. Ravn & Hans Rischel, editors: *Hybrid Systems, Lecture Notes in Computer Science 736*, Springer, pp. 36–59, doi:[10.1007/3-540-57318-6_23](https://doi.org/10.1007/3-540-57318-6_23).

- [7] Garanin Dmitry (2008): *Rotational motion of rigid bodies*. https://www.lehman.edu/faculty/dgaranin/Mechanics/Mechanis_of_rigid_bodies.pdf.
- [8] Goran Frehse, Colas Le Guernic, Alexandre Donzé, Scott Cotton, Rajarshi Ray, Olivier Lebeltel, Rodolfo Ripado, Antoine Girard, Thao Dang & Oded Maler (2011): *SpaceEx: Scalable Verification of Hybrid Systems*. In Ganesh Gopalakrishnan & Shaz Qadeer, editors: *Computer Aided Verification*, Springer Berlin Heidelberg, pp. 379–395, doi:[10.1007/978-3-642-22110-1_30](https://doi.org/10.1007/978-3-642-22110-1_30).
- [9] Peter Fritzson (2014): *Principles of object-oriented modeling and simulation with Modelica 3.3: a cyber-physical approach*. John Wiley & Sons, doi:[10.1002/9781118989166](https://doi.org/10.1002/9781118989166).
- [10] Sergey Goncharov & Renato Neves (2019): *An Adequate While-Language for Hybrid Computation*. In Ekaterina Komendantskaya, editor: *Proceedings of the 21st International Symposium on Principles and Practice of Programming Languages, PDP 2019*, ACM, pp. 11:1–11:15, doi:[10.1145/3354166.3354176](https://doi.org/10.1145/3354166.3354176).
- [11] Sergey Goncharov, Renato Neves & José Proença (2020): *Implementing Hybrid Semantics: From Functional to Imperative*. In Violet Ka I Pun, Volker Stolz & Adenilso Simão, editors: *Theoretical Aspects of Computing - ICTAC 2020 - 17th International Colloquium, Macau, China, November 30 - December 4, 2020, Proceedings, Lecture Notes in Computer Science 12545*, Springer, pp. 262–282, doi:[10.1007/978-3-030-64276-1_14](https://doi.org/10.1007/978-3-030-64276-1_14).
- [12] Volkan Gunes, Steffen Peter, Tony Givargis & Frank Vahid (2014): *A Survey on Concepts, Applications, and Challenges in Cyber-Physical Systems*. *Transactions on Internet and Information Systems* 8(12), pp. 4242–4268, doi:[10.3837/TIIS.2014.12.001](https://doi.org/10.3837/TIIS.2014.12.001).
- [13] Ahammodullah Hasan, Md Abdul Halim & MA Meia (2019): *Application of linear differential equation in an analysis transient and steady response for second order RLC closed series circuit*. *American Journal of Circuits, Systems and Signal Processing* 5(1), pp. 1–8.
- [14] Peter Höfner (2009): *Algebraic calculi for hybrid systems*. Ph.D. thesis, University of Augsburg. Available at <http://opus.bibliothek.uni-augsburg.de/volltexte/2010/1481/>.
- [15] Eduard Kamburjan, Stefan Mitsch & Reiner Hähnle (2022): *A Hybrid Programming Language for Formal Modeling and Verification of Hybrid Systems*. *Leibniz Transactions on Embedded Systems* 8(2), pp. 04:1–04:34, doi:[10.4230/LITES.8.2.4](https://doi.org/10.4230/LITES.8.2.4).
- [16] Harold Klee (2007): *Simulation of Dynamic Systems with MATLAB and Simulink*. CRC Press, Inc., USA.
- [17] Dexter Kozen (1997): *Kleene algebra with tests* 19(3), p. 427443. doi:[10.1145/256167.256195](https://doi.org/10.1145/256167.256195).
- [18] Tomas Krilavicius (2006): *Hybrid Techniques for Hybrid Systems*. Ph.D. thesis, University of Twente, Enschede, Netherlands. Available at <http://eprints.eemcs.utwente.nl/9609/>.
- [19] Edward A. Lee & Sanjit A. Seshia (2016): *Introduction to embedded systems: A cyber-physical systems approach*. MIT Press.
- [20] Zohar Manna & Amir Pnueli (1992): *Verifying Hybrid Systems*. In Robert L. Grossman, Anil Nerode, Anders P. Ravn & Hans Rischel, editors: *Hybrid Systems, Lecture Notes in Computer Science 736*, Springer, pp. 4–35, doi:[10.1007/3-540-57318-6_22](https://doi.org/10.1007/3-540-57318-6_22).
- [21] Renato Neves (2018): *Hybrid programs*. Ph.D. thesis, University of Minho. Available at <https://repositorium.sdum.uminho.pt/handle/1822/56808>.
- [22] E-R Olderog (1992): *Nets, terms and formulas: three views of concurrent processes and their relationship*. Cambridge University Press.
- [23] Lawrence Perko (2013): *Differential equations and dynamical systems*. 7, Springer Science & Business Media, doi:[10.1007/978-1-4613-0003-8](https://doi.org/10.1007/978-1-4613-0003-8).
- [24] André Platzer (2010): *Logical Analysis of Hybrid Systems - Proving Theorems for Complex Dynamics*. Springer, doi:[10.1007/978-3-642-14509-4](https://doi.org/10.1007/978-3-642-14509-4).
- [25] André Platzer (2018): *Logical Foundations of Cyber-Physical Systems*. Springer, doi:[10.1007/978-3-319-63588-0](https://doi.org/10.1007/978-3-319-63588-0).
- [26] John C Reynolds (1998): *Theories of programming languages*. Cambridge University Press, doi:[10.1017/CBO9780511626364](https://doi.org/10.1017/CBO9780511626364).

- [27] W.A. Stein et al. (2015): *Sage Mathematics Software (Version 6.4.1)*. The Sage Development Team. <http://www.sagemath.org>.
- [28] Glynn Winskel (1993): *The formal semantics of programming languages - an introduction*. Foundation of computing series, MIT Press, doi:[10.7551/mitpress/3054.001.0001](https://doi.org/10.7551/mitpress/3054.001.0001).
- [29] Yue Zhang & Anurag Srivastava (2021): *Voltage Control Strategy for Energy Storage System in Sustainable Distribution System Operation*. *Energies* 14(4), doi:[10.3390/en14040832](https://doi.org/10.3390/en14040832).