Differential Equations via Temporal Logic and Infinitesimals

15-824 Foundations of Cyber-Physical Systems
Evan Cavallo (ecavallo@cs.cmu.edu)

Abstract

We present a temporal logic for reasoning about differential equations which expresses differential behavior via a time domain with a nilpotent (specifically, nilsquare) infinitesimal, an element $\varepsilon > 0$ for which $\varepsilon^2 = 0$. We define a language and axiomatic system as well as a semantics using the ring of dual numbers $\mathbb{R}[\varepsilon]$.

Introduction

One way of modeling systems which evolve through time is via temporal logics, which introduce modal operators expressing in some way "truth at time t." Many variants on this setup have been proposed (see [4, Section 5.2] for a survey), differing both in the form of the logical connectives and in the form of the time domain (for example, discrete vs. continuous time). In this approach, we tailor the time domain for reasoning about differential equations by adding *infinitesimal elements*, positive but smaller than any standard positive real, which can be used to encode properties of derivatives.

There are two main approaches to analysis using infitesimals: *non-standard analysis* (NSA), which uses the *hyperreal numbers* constructed using ultraproducts [10], and *smooth infinitesimal analysis* (SIA), a subfield of *synthetic differential geometry* [6], which uses algebraic and category-theoretic constructions. We take the latter approach. (There is prior work using the former for modeling continuous systems, e.g. [1, 2, 11, 12]; NSA infinitesimals are used with temporal logic in [5], but in order to model instantaneous transitions, not differential behavior.) Smooth infinitesimal analysis is based on the introduction of *nilpotent infinitemisals*, positive points $\epsilon > 0$ such that $\varepsilon^{n+1} = 0$; non-standard analysis, on the other hand, introduces *invertible infinitesimals*, which are smaller than any standard real and have infinitely large inverses. We contend that our approach is semantically cleaner than the SIA approach; although developing a fully general model for SIA on manifolds is complex [7], our restricted use case has a simple semantics, whereas the construction of the hyperreals is highly non-constructive, requiring some form of choice axiom. We also believe that SIA is the more natural setting for differential reasoning, since nilsquare infinitesimals directly encode first-order behavior. (We will not justify this claim further, but see [7, Introduction & Chapter VI] for discussion of the values of the two approaches.)

In §1, we define and motivate the ring of dual numbers, which provides the semantic basis for our theory. In §2 we present the syntax of our logic εTL and its axiom system, and in §3 we give the intepretation of εTL in the intended model. In §4 we describe a decision procedure for the first-order fragment of εTL, assuming completeness. Finally, in §5 we compare εTL to differential dynamic logic [8], in particular its continuous fragment FOD.

1 The Ring of Dual Numbers $\mathbb{R}[\varepsilon]$

Our approach to a theory of differential equations is based on the *ring of dual numbers*, an algebraic construction which adds infinitesimal elements to the real number field $\mathbb{R}$. The ring of dual numbers is defined as $\mathbb{R}[x]/(x^2)$, the quotient of the ring of polynomials with real coefficients by the relation $x^2 = 0$. For readability, we write $\mathbb{R}[\varepsilon]$ for the ring of dual numbers and ε for the element x. This ε is a *nilpotent infinitesimal*, an element so "small" that $\varepsilon^2 = 0$; as we will see, it plays the role of a first-order differential. (Note that $\mathbb{R}[\varepsilon]$ is not a field, as ε has no inverse.) Formally, an element of $\mathbb{R}[\varepsilon]$ is a polynomial $a_0 + a_1\varepsilon + \ldots + a_n\varepsilon^n$. Since $\varepsilon^2 = 0$, however, every element is equal to one of the form $a_0 + a_1\varepsilon$. Addition and multiplication are inherited from their definitions on polynomials:

$$\begin{aligned}(a_1 + b_1\varepsilon) + (a_2 + b_2\varepsilon) &\implies (a_1 + a_2) + (b_1 + b_2)\varepsilon \\ (a_1 + b_2\varepsilon) \cdot (a_2 + b_2\varepsilon) &\implies a_1a_2 + (a_1b_2 + a_2b_1)\varepsilon\end{aligned}$$

An order on $\mathbb{R}[\varepsilon]$ is defined lexicographically, following the intuition that ε is "smaller than any standard real."

$$a_1 + b_1\varepsilon \leq a_2 + b_2\varepsilon \quad \text{iff} \quad \text{either } a_1 < a_2 \text{ or } a_1 = a_2 \text{ and } b_1 \leq b_2$$

For a taste of the motivation behind using $\mathbb{R}[\varepsilon]$, consider a polynomial $P(x) = a_0 + a_1 x + \cdots + a_n x^n$ with coefficients in $\mathbb{R}$. Then

$$P(x+\varepsilon) = a_0 + a_1(x+\varepsilon) + \cdots + a_n(x+\varepsilon)^n$$

For any n, $(x+\varepsilon)^n = x^n + nx^{n-1}\varepsilon + (\cdots)\varepsilon^2 = x^n + nx^{n-1}$. Thus

$$\begin{aligned} P(x+\varepsilon) &= a_0 + a_1(x+\varepsilon) + \cdots + a_n(x + nx^{n-1}\varepsilon) \\ &= P(x) + P'(x)\varepsilon \end{aligned}$$

where $P'(x) = a_1 + \ldots + a_n n x^{n-1}$ is the formal derivative of P at x. Rearranged, $P'(x)\varepsilon = P(x+\varepsilon) - P(x)$. Since an analytical derivative can be defined as

$$f'(x) = \lim_{h \to 0} \frac{f(x+h) - f(x)}{h},$$

we think of ε as a quantity so small that the limiting equation already holds for $h = \varepsilon$; in other words, on the infinitesimal (specifically, first-order differential) scale, f is linear.

2 Syntax and Axiomatics of εTL

Guided by the idea of semantics in the ring of dual numbers, we define a temporal logic εTL and describe its axiomatics. The terms are given as in differential dynamic logic, with the addition of the constant ε.

$$e ::= x \mid A \mid 0 \mid 1 \mid \varepsilon \mid e + e \mid e \cdot e$$

Note also the presence of two varieties of variable, x and A. We use lowercase $x, y, z, \ldots$ (also, t) to denote *differentiable variables* and capital letters $A, B, C, \ldots$ to denote *constant variables.* Importantly, *only terms containing no differentiable variables may be substituted for constant variables.* Any term can be substituted for a differentiable variable. The formulae are those of first-order propositional logic with an order relation, along a universal quantifier for constant variables and a modal operator $\bigcirc_t$.

$$\phi ::= e = e \mid e \leq e \mid \phi \wedge \phi \mid \neg\phi \mid \forall x.\phi \mid \forall A.\phi \mid \bigcirc_e \phi$$

As usual, we write $\vee$ and $\rightarrow$ for the connectives derived from $\wedge$ and $\neg$. The modal operator $\bigcirc_e$ expresses *truth after e time has passed* (where e is evaluated at the current moment). The value t may be negative, so $\bigcirc$ can also be used to express past events. The distinction between constant and differentiable variables comes into play with $\bigcirc$: as their names suggest, constant variables do not change over time, while differentiable variables may change in a differentiable way.

The axioms capturing the first-order properties of $\mathbb{R}[\varepsilon]$ essentially follow those of $\mathbb{R}$, with minor adjustments to accommodate ε. We write $x \approx y$ as shorthand for $x \cdot \varepsilon = y \cdot \varepsilon$ – this expresses that x and y differ by an infinitesimal.

$$\begin{array}{ccc} (x+y)+z = x+(y+z) & \text{"}\leq \text{ is a total order"} & \exists y. x + y = 0 \\ x + y = y + x & x \leq y \rightarrow y + z \leq x + z & x \not\approx 0 \rightarrow \exists y. x \cdot y = 1 \\ x \cdot y = y \cdot x & 0 \leq x \wedge 0 \leq y \rightarrow 0 \leq x \cdot y & x \approx 0 \rightarrow \exists y. x = y \cdot \varepsilon \\ x + 0 = x & 0 < \varepsilon & \exists A. A = x \\ x \cdot 1 = x & & \\ \varepsilon^2 = 0 & & \end{array}$$

In the right column we see that only non-infinitesimal elements are invertible, while all infinitesimals are some multiple of ε. Note that ε is provably smaller than any positive rational: since $(1-\varepsilon)(1+\varepsilon) = 1-\varepsilon^2 = 1$ is positive and $1+\varepsilon \geq 1 \geq 0$, we must have $1-\varepsilon \geq 0$, thus $1-\varepsilon n = (1-\varepsilon)^n \geq 0$ for any $n \in \mathbb{N}$, thus $\varepsilon \leq \frac{1}{n}$. The axiom $\exists A.A = x$ expresses that there is always a constant variable with the current value of a differentiable variable. For compound formulae, we have, in addition to the standard axioms of classical first-order predicate logic, axioms describing the $\bigcirc_t$ modality.

$$\bigcirc_t(\phi \wedge \psi) \leftrightarrow (\bigcirc_t\phi) \wedge (\bigcirc_t\psi) \qquad\qquad (\bigcirc_t\neg\phi) \leftrightarrow \neg \bigcirc_t \phi$$
$$(\bigcirc_t\forall A.\phi) \leftrightarrow \forall A.(\bigcirc_t\phi) \qquad (\bigcirc_{t_1} \bigcirc_{t_2} \phi) \leftrightarrow \exists A.(\bigcirc_{t_1} t_2 = A) \wedge \bigcirc_{t_1+A}\phi$$

Finally, we have axioms describing the relation of $\bigcirc_t$ and ε to differentiation. Here, $\tilde{\exists!}A.\phi(A)$ is shorthand for $\exists A.(\phi(A) \wedge \forall B.\phi(B) \rightarrow B \approx A)$, that is, unique existence up to $\approx$.

$$(\bigcirc_t\phi) \leftrightarrow \phi \quad \text{(if } \phi \text{ contains no differentiable variables)}$$

$$\tilde{\exists!}A.\forall X.\ x = X \rightarrow \forall B.\ \bigcirc_{B\cdot\varepsilon}\,(x = X + A \cdot B \cdot \varepsilon)$$
(Kock-Lawvere axiom)

$$\tilde{\exists!}X_f.\forall x.\ t \geq 0 \wedge x \approx x_i \rightarrow (\forall 0 < A < t.\ \bigcirc_A\,(\forall X.x = X \rightarrow \bigcirc_\varepsilon x = X + e(X))) \rightarrow \bigcirc_t(x \approx X_f)$$
(uniqueness of solutions)

The first of these axioms internalizes that constant variables are in fact constant. The second states that every term has a derivative, unique up to $\approx$, and serves to relate the infinitesimal timesteps $\bigcirc_{B\cdot\varepsilon}$ as B varies. The third is a schema parameterized over terms e (which, as shown explicitly, may mention X, but not X_i or X_f), and gives the uniqueness of solutions (again, up to $\approx$) to differential equations.

3 Semantics of εTL

We now give the intended semantics for εTL. The semantics of terms is defined by an interpretation function $[\![-]\!]_u^{D;C}$, taking terms of εTL to elements of $\mathbb{R}[\varepsilon]$. Here D is an environment mapping differentiable variables x to pairs of functions $D_0(x), D_1(x) : \mathbb{R} \rightarrow \mathbb{R}$ with $D_0(x)$ differentiable, C is an environment mapping constant variables to elements of $\mathbb{R}[\varepsilon]$, and $u \in \mathbb{R}[\varepsilon]$ is a time index. The interpretation function is defined inductively as follows, where $u = u_0 + u_1 \cdot \varepsilon$ for $u_0, u_1 \in \mathbb{R}$.

$$\begin{array}{rcl}
[\![x]\!]_u^{D;C} & = & D_0(x)(u_0) + (D_1(x)(u_0) + (D_0(x))'(u_0) \cdot u_1) \cdot \varepsilon \\
[\![A]\!]_u^{D;C} & = & C(A) \\
[\![0]\!]_u^{D;C} & = & 0 \\
[\![1]\!]_u^{D;C} & = & 1 \\
[\![\varepsilon]\!]_u^{D;C} & = & \varepsilon \\
[\![e_1 + e_2]\!]_u^{D;C} & = & [\![e_1]\!]_u^{D;C} + [\![e_2]\!]_u^{D;C} \\
[\![e_1 \cdot e_2]\!]_u^{D;C} & = & [\![e_1]\!]_u^{D;C} \cdot [\![e_2]\!]_u^{D;C}
\end{array}$$

Observe that for any term e, the mathematical function $u \mapsto [\![e]\!]_u^{D;C}$ can be written in the form as $u \mapsto f_0(u_0) + (f_0'(u_0) \cdot u_1 + f_1(u_0)) \cdot \varepsilon$ for some $f_0, f_1 : \mathbb{R} \rightarrow \mathbb{R}$ with f_0 differentiable. (The factor f_1 is necessary to account for functions such as $u \mapsto u \cdot \varepsilon$, where we take $f_0(u_0) = 0$ and $f_1(u_0) = u_0$.) We are therefore able to substitute the interpretation of a term for a differentiable variable in D, and we have $[\![e]\!]^{D;C} =$

$[\![x]\!]^{D[x\mapsto(f_0,f_1)];C}$ for f_0, f_1 as above (and x fresh). We now define semantic entailment $\models_u^{D;C} \phi$.

$$\begin{array}{lcl}
\models_u^{D;C} (e_1 = e_2) & \iff & [\![e_1]\!]_u^{D;C} = [\![e_2]\!]_u^{D;C} \\
\models_u^{D;C} (e_1 \leq e_2) & \iff & [\![e_1]\!]_u^{D;C} \leq [\![e_2]\!]_u^{D;C} \\
\models_u^{D;C} \phi \wedge \psi & \iff & \models_u^{D;C} \phi \text{ and } \models_u^{D;C} \psi \\
\models_u^{D;C} \neg\phi & \iff & \not\models_u^{D;C} \phi \\
\models_u^{D;C} \forall A.\phi & \iff & \models_u^{D;C[A\mapsto v]} \phi \text{ for all } v \in \mathbb{R}[\varepsilon] \\
\models_u^{D;C} \forall x.\phi & \iff & \models_u^{D[x\mapsto(f_0,f_1)];C} \phi \text{ for all } f_0, f_1 : \mathbb{R} \to \mathbb{R} \text{ with } f_0 \text{ differentiable} \\
\models_u^{D;C} \bigcirc_e\phi & \iff & \models_{u+[\![e]\!]_u^{D;C}}^{D;C} \phi
\end{array}$$

Nearly all of the axioms described in §2 are obviously sound with respect to these semantics. We will only explicitly verify the uniqueness axiom

$$\tilde{\exists!} X_f.\forall x.\ t \geq 0 \wedge x \approx x_i \to (\forall 0 < A < t.\ \bigcirc_A (\forall X.x = X \to \bigcirc_\varepsilon x = X + e(X))) \to \bigcirc_t(x \approx X_f)$$

For any D, C we have a function defined by $(u, a) \mapsto [\![e]\!]_u^{D;C[X\mapsto a]}$. As argued above, this function can be written as $(u, a) \mapsto f_0(u_0, a) + (f_0'(u_0, a) \cdot u_1 + f_1(u_0, a)) \cdot \varepsilon$ for some f_0, f_1 with f_0 differentiable in u_0. The restricted syntax of the term language implies that f_0, f_1 must be polynomial in a. The semantic interpretation of the uniqueness axiom applied at state D, C, u is then

> There exists $a_f \in \mathbb{R}$ such that, for any $g : \mathbb{R} \to \mathbb{R}$ differentiable with $g(u) = a_i$ and $g'(w) = f_0(w, g(w))$ for all $u < v < w$, we have $g(v) = a_f$.

where a_i is the standard part of $[\![x_i]\!]_u^{D;C}$. In other words, there is a unique solution (modulo infinitesimals) to the differential equation f_0 with initial value x_i. Since f_0 is polynomial in a, the Picard-Lindelöf Theorem [13, Theorem 10.VI] applies to validate this axiom. We remark that this axiom is *not* sound if $\approx$ is replaced by $=$ in the statement: the functions $u \mapsto u \cdot \varepsilon$ and $u \mapsto 2 \cdot u \cdot \varepsilon$ are equal at time $u = 0$ and both have derivative 0 everywhere, but are not equal at time $u = 1$. Also, the generalized "differential invariant" axiom

$$\phi \wedge t \geq 0 \wedge (\forall 0 < A < t.\ \bigcirc_A (\phi \to \bigcirc_\varepsilon\phi)) \to \bigcirc_t\phi$$

is unsound even if ϕ is $\approx$-invariant. Consider the formula $\phi(x) = \neg(x \approx 1)$, and suppose x is interpreted by identity function. Then $\phi(x)$ is true at time 0 and preserved along infinitesimal steps, but fails to hold at time 1.

4 Deciding First-Order $\mathbb{R}[\varepsilon]$-Arithmetic

In this section, we describe a decision procedure for the first-order fragment of our axiomatization (FO$\mathbb{R}[\varepsilon]$) in terms of the corresponding procedure for first-order real arithmetic (FO$\mathbb{R}$). (Since $\mathbb{R}[e]$ is not a field, real-closed field methods do not apply directly; while a decision procedure with infinitesimal elements is developed in [3], the infinitesimals in question are *invertible* infinitesimals.) We will use without proof that FO$\mathbb{R}[\varepsilon]$ is a complete theory, which we believe to be true; if it is false, we do not expect that it would be difficult to correct the theory. Of course, once we assume completeness, we already know that a decision procedure exists, but our intent is to observe that any algorithm for deciding FO$\mathbb{R}$ can be adapted absolutely straightforwardly to FO$\mathbb{R}[\varepsilon]$.

We describe a translation $(-)^\uparrow$ from formulae in FO$\mathbb{R}[\varepsilon]$ to formulae in FO$\mathbb{R}$, then argue that ϕ is semantically valid if and only if $\phi^\uparrow$ is. The formula translation $(-)^\uparrow$ takes a formula in n variables (differentiable or constant) to a formula in $2n$ real variables. For the term translation, each term t in FO$\mathbb{R}[\varepsilon]$ produces two terms in FO$\mathbb{R}$, for which we will write $t^\uparrow = (t_0^\uparrow, t_1^\uparrow)$ – the pair is meant to represent the term $t_0^\uparrow + t_1^\uparrow \cdot \varepsilon$. These are defined as follows.

$$\begin{array}{rcl}
x^{\uparrow} & = & (x_0, x_1) \\
A^{\uparrow} & = & (A_0, A_1) \\
0^{\uparrow} & = & (0, 0) \\
1^{\uparrow} & = & (1, 0) \\
\varepsilon^{\uparrow} & = & (0, 1) \\
(t + r)^{\uparrow} & = & (t_0^{\uparrow} + r_0^{\uparrow}, t_1^{\uparrow} + r_1^{\uparrow}) \\
(t \cdot r)^{\uparrow} & = & (t_0^{\uparrow} \cdot r_0^{\uparrow}, t_0^{\uparrow} \cdot r_1^{\uparrow} + t_1^{\uparrow} \cdot r_0^{\uparrow})
\end{array}$$

The translation of formulae is then given by

$$\begin{array}{rcl}
(t = r)^{\uparrow} & = & t_0^{\uparrow} = r_0^{\uparrow} \wedge t_1^{\uparrow} = r_1^{\uparrow} \\
(t \leq r)^{\uparrow} & = & t_0^{\uparrow} < r_0^{\uparrow} \vee (t_0^{\uparrow} = r_0^{\uparrow} \wedge t_1^{\uparrow} \leq r_1^{\uparrow}) \\
(\phi \wedge \psi)^{\uparrow} & = & \phi^{\uparrow} \wedge \psi^{\uparrow} \\
(\neg \phi)^{\uparrow} & = & \neg(\phi^{\uparrow}) \\
(\forall A.\phi)^{\uparrow} & = & \forall A_0.\forall A_1.(\phi^{\uparrow}) \\
(\forall x.\phi)^{\uparrow} & = & \forall x_0.\forall x_1.(\phi^{\uparrow})
\end{array}$$

The translations of $+$, $\cdot$, $=$,and $\leq$ follow exactly their definitions in $\mathbb{R}[\varepsilon]$ in terms of $\mathbb{R}$, so the semantic correctness of the translation is virtually immediate. Thus, ϕ is valid for FO$\mathbb{R}[\varepsilon]$ if $\phi^{\uparrow}$ is valid for FO$\mathbb{R}$, so we can simply apply a real-arithmetic decision procedure to $\phi^{\uparrow}$ to decide ϕ.

5 d$\mathcal{L}$ and εTL

In this section, we describe the relationship between differential dynamic logic [8] and εTL. Since εTL only deals with functions which vary in a differentiable way, its expressive power is limited compared to d$\mathcal{L}$. Since d$\mathcal{L}$ has a sound and complete axiomatization relative to its differential part FOD, the first order fragment with differential modality $[x' = \theta]\phi$ (see [9]), it is possible that the same holds for d$\mathcal{L}$ relative to εTL. However, εTL cannot express piecewise differentiable functions as FOD can, and the author has not investigated the impact of this difference on the proof of completeness in [9].

Nevertheless, we note that certain axioms of d$\mathcal{L}$ concerning the evolution of differential equations are provable in some form in εTL, while others are at least semantically valid. First of all, the modal formula $[x' = \theta(x)]\phi$ can be translated as the formula

$$\forall r \geq 0.(\forall 0 < t < r. \bigcirc_t \, x' \ \mathit{is} \ \theta) \rightarrow \bigcirc_r \phi$$

where $x' \ \mathit{is} \ \theta$ is shorthand for

$$\forall X.x = X \rightarrow \bigcirc_\varepsilon x = X + \theta(X)$$

(As a side note, we observe that various derivative laws are provable in εTL, for example $x_1' \ \mathit{is} \ \theta_1 \wedge x_2' \ \mathit{is} \ \theta_2 \rightarrow (x_1 x_2)' \ \mathit{is} \ x_1\theta_2 + x_2\theta_1$.) We will write $[x' = \theta(x)]\phi$ for the εTL translation as well as the d$\mathcal{L}$ formula. Consider the axiom $[']$ which characterizes the differential modality in FOD:

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

The εTL equivalent would then be

$$\frac{}{[x' = \theta]\phi \leftrightarrow \forall t \geq 0[y(t)/x]\phi} \ (y'(t) = \theta)$$

This formula is not true, as noted in Section 3, because uniqueness of solutions is true only up to $\approx$ and ϕ may not be $\approx$-invariant. However, the axiom is provable from uniqueness of solutions for all ϕ which are provably $\approx$-invariant, which includes the sublanguage of εTL using $\{\approx, \lesssim\}$ but not $\{=, \leq\}$. (Here $e_1 \lesssim e_2$ is of course defined as $e_1 \cdot \varepsilon \leq e_2 \cdot \varepsilon$.)

The uniqueness of solutions axiom is also sufficient to prove at least one instance of a differential invariant (DI) axiom of d$\mathcal{L}$. For example, consider the following example of an invariant rule derivable in d$\mathcal{L}$.

$$\frac{\Gamma \vdash x = y \quad \vdash \theta_1 = \theta_2}{\Gamma \vdash [x' = \theta_1, y' = \theta_2] x = y}$$

Again, although this rule does is unsound in εTL, the adjusted rule

$$\frac{\Gamma \vdash x \approx y \quad \vdash \theta_1 = \theta_2}{\Gamma \vdash [x' = \theta_1, y' = \theta_2] x \approx y}$$

is easily provable from uniqueness of solutions. Other examples of differential invariants, such as

$$\frac{\Gamma \vdash x \leq y \quad \vdash \theta_1 \leq \theta_2}{\Gamma \vdash [x' = \theta_1, y' = \theta_2] x \leq y}$$

do not appear to be derivable in εTL even with $\leq$ replaced by $\lesssim$ in appropriate places. However, they are semantically valid, so additional axioms could be introduced (as they are in d$\mathcal{L}$) for these purposes.

6 Conclusions and Future Work

We have presented a simple logic for reasoning about differential equations using first-order differential elements, which we hope could be used as a foundation for more practical developments. Besides cleaning up loose ends (e.g., the completeness of FO$\mathbb{R}[\varepsilon]$), we mention two potential directions for improvement. The first is to handle hybrid systems, which contain both continuous and discrete elements. As noted in §5, εTL is very limited in its ability to describe discrete transitions, since no variable is capable of discrete evolution. Even if the completeness proof of [9] can be adapted to εTL, constructs for direct reasoning about discrete transitions are necessary for practical purposes. One possibility is to introduce a predicate $\mathrm{Diff}(e)$ which holds when a term is differentiable at the current time, and restrict the Kock-Lawvere and uniqueness of solution axioms appropriately; this would allow for a semantics where functions need not always vary differentiably.

Second is to more critically analyze the choice of axioms for εTL. As we have observed in §5, the uniqueness of solutions axiom allows us to derive some rules of d$\mathcal{L}$, but not all, and the conceptual priority of the axiom and possibility of augmentations or alternatives remains unexplored.

References

[1] Albert Benveniste, Timothy Bourke, Benoît Caillaud, and Marc Pouzet. Non-standard semantics of hybrid systems modelers. *Journal of Computer and Systems Sciences International*, 78(3):877–910, 2012.

[2] Simon Bliudze and Daniel Krob. Modelling of complex systems: Systems as dataflow machines. *Fundamenta Informaticae*, 91(2):251–274, 2009.

[3] Leonardo de Moura and Grant Olney Passmore. Computation in real closed infinitesimal and transcendental extensions of the rationals. In Maria Paola Bonacina, editor, *Automated Deduction – CADE-24: 24th International Conference on Automated Deduction, Lake Placid, NY, USA, June 9-14, 2013. Proceedings*, pages 178–192, Berlin, Heidelberg, 2013. Springer.

[4] Carlo A. Furia, Dino Mandrioli, Angelo Morzenti, and Matteo Rossi. Modeling time in computing: A taxonomy and a comparative survey. *ACM Comput. Surv.*, 42(2):6:1–6:59, March 2010.

[5] Angelo Gargantini, Dino Mandrioli, and Angelo Morzenti. Dealing with zero-time transitions in axiom systems. *Information and Computation*, 150(2):119 – 131, 1999.

[6] Anders Kock. *Synthetic Differential Geometry*. Cambridge University Press, 2nd edition, 2006.

[7] Ieke Moerdijk and Gonzalo E. Reyes. *Models for Smooth Infinitesimal Analysis*. Springer-Verlag, New York, 1991.

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

[9] André Platzer. The complete proof theory of hybrid systems. In *Logic in Computer Science (LICS), 2012 27th Annual IEEE Symposium on*, pages 541–550, June 2012.

[10] Abraham Robinson. *Non-standard Analysis*. Princeton University Press, 1974.

[11] Kohei Suenaga and Ichiro Hasuo. Programming with infinitesimals: A while-language for hybrid system modeling. *Lecture Notes in Computer Science*, 6756:392–403, 2011.

[12] Kohei Suenaga, Hiroyoshi Sekine, and Ichiro Hasuo. Hyperstream processing systems: Nonstandard modeling of continuous-time signals. *Principles of Programming Languages*, 2013.

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