15-424: Foundations of Cyber-Physical Systems

Lecture Notes on Logical Theory & Completeness

André Platzer

Carnegie Mellon University

Lecture 24

1 Introduction

This course has studied a number of logics, first-order logic FOL in Lecture 2, differential dynamic logic d$\mathcal{L}$ [Pla08, Pla10a, Pla12c, Pla12b] in Lecture 3 and Lecture 4 and following, differential temporal dynamic logic dTL [Pla07, Pla10a, Chapter 4] in Lecture 16 and 17, as well as differential game logic dG$\mathcal{L}$ [Pla13] since Lecture 22. There are other logics for cyber-physical systems that have not been included in this course, but share similar principles for further dynamical aspects. Such logics include quantified differential dynamic logic Qd$\mathcal{L}$ for distributed hybrid systems [Pla10b, Pla12a], which are systems that are simultaneously distributed systems and hybrid systems, as well as stochastic differential dynamic logic Sd$\mathcal{L}$ for stochastic hybrid systems [Pla11], which simultaneously involve stochastic dynamics and hybrid dynamics. Logics play a stellar role not just in cyber-physical systems, but also many other contexts. Other important logics include propositional logic, restrictions of first-order logic to certain theories, such as first-order logic of real arithmetic [Tar51], and higher-order logic [And02]. But there are numerous other important and successful logics.

In this lecture, we take a step back and study some common important concepts in logic. This study will necessarily be hopelessly incomplete for lack of time. But it should give you a flavor of important principles and concepts in logic that we have not already run across explicitly in earlier lectures of this course. We will also have the opportunity to apply these more general concepts to cyber-physical systems and learn more about them in the next lecture.

2 Soundness

The most important parts of a logic $\mathcal{L}$ are the following. The logic $\mathcal{L}$ defines what the *syntactically well-formed formulas* are. Every well-formed formula carries meaning, which the *semantics of formulas* in $\mathcal{L}$ defines. The semantics defines a relation $\vDash$ between sets of formulas and formulas, in which $\Phi \vDash \phi$ holds iff ϕ is a semantic consequence of the set of formulas Φ, i.e. ϕ is true (usually written $\nu \models \phi$) in every interpretation ν for which all formulas $\psi \in \Phi$ are true. The most important case for our purposes is the case $\Phi = \emptyset$ of validity, in which case $\vDash \phi$ holds iff ϕ is valid, i.e. true ($\nu \models \phi$) in all interpretations ν of $\mathcal{L}$. An interpretation ν in which ϕ is true (i.e. $\nu \models \phi$) is also called a *model* of ϕ.

For the case of first-order logic FOL, Lecture 2 defined both their syntax and semantics. The syntax and semantics of differential dynamic logic d$\mathcal{L}$ has been defined in Lecture 3 and Lecture 4.

The syntax of a logic $\mathcal{L}$ defines what we can write down that carries meaning. The semantics of a logic $\mathcal{L}$ then defines what the meaning of the syntactic formulas is. The semantics, in particular, defines which formulas express true facts about the world, either in a particular interpretation ν or about the world in general (for valid formulas, which are true regardless of the interpretation). Yet, the semantics is usually highly ineffective, so that it cannot be used directly to find out whether a formula is valid. Just think of formulas in differential dynamic logic that express safety properties of hybrid systems. It would not get us very far if we were to try to establish the truth of such a formula by literally computing the semantics (which includes executing the hybrid system) in every initial state, of which there are uncountably infinitely many.

Instead, logics come with *proof calculi* that can be used to establish validity of logical formulas in the logic $\mathcal{L}$. Those proof calculi comprised *axioms* (Lecture 5) and *proof rules* (Lecture 6 and others), which can be combined to prove or derive logical formulas of the logic $\mathcal{L}$. The proof calculus of the logic $\mathcal{L}$ defines a relation $\vdash$ between sets of formulas and formulas, in which $\Phi \vdash \phi$ holds iff ϕ is provable from the set of formulas Φ. That is, there is a proof of ϕ in the proof calculus of $\mathcal{L}$ that uses only assumptions from Φ. The most important case for our purposes is again $\Phi = \emptyset$, in which case $\vdash \phi$ holds iff ϕ is provable in the proof calculus of $\mathcal{L}$, i.e. there is a proof of ϕ.

Of course, only some formulas of $\mathcal{L}$ are provable, not all of them. The formula $p \wedge \neg p$ should not be provable in any proper logic, because it is inconsistently false and, thus, cannot possibly be valid.

We could have written down any arbitrary axiom, or we could have accidentally had a typo in the axioms. So a crucial question we have to ask (and have asked every time we introduced an axiom in other lectures of this course) is whether the axioms and proof rules are sound. In a nutshell, a proof calculus is sound if all provable formulas are valid.

Theorem 1 (Soundness [Pla08, Pla10a, Pla12b])**.** *The proof calculus of differential dynamic logic is* sound*, i.e.* $\vdash \subseteq \vDash$*, which means that* $\vdash \phi$ *implies* $\vDash \phi$ *for all* dℒ *formulas* ϕ*. That is, all provable* dℒ *formulas are valid.*

The significance of soundness is that, whatever formula we derive by using the dℒ proof rules and axioms, we can rest assured that it is valid, i.e. true in all states. In particular, it does not matter how big and complicated the formula might be, we know that it is valid as long as we have a proof for it. About the axioms, we can easily convince ourselves using a soundness proof why they are valid, and then conclude that all provable formulas are also valid, because they follow from sound axioms by sound proof rules.

Note 2 (Necessity of soundness)**.** *Soundness is a must for otherwise we could not trust our own proofs.*

3 Soundness Challenge for CPS

What good would it do to analyze safety of a CPS using a technique that is as faulty as the original CPS? If an unsound analysis technique says that a CPS is correct, we are, fundamentally, not much better off than without any analysis, because all we can conclude is that we did not find problems, not that there are none.[1] After all, an unsound analysis technique could say "correct", which might turn out to be a lie because the correctness statement itself was not valid.

Note 3 (Challenge of soundness)**.** *In a domain that is as challenging as cyber-physical systems and hybrid systems, it is surprisingly easy for analysis techniques to become unsound due to subtle flaws. Necessary conditions for soundness and the numerical decidability frontier have been identified in the literature [PC07, Col07]. The crux of the matter is that hybrid systems are subject to a numerical analogue of the halting problem of Turing machines [PC07].*

There is a shockingly large number of approaches that, for subtle reasons, are subject to the unsoundness resulting from non-observance of the conditions identified in [PC07, Col07]. Consequently, such approaches need some of the additional assumptions identified in [PC07, Col07] to have a chance to become sound.

[1]Notwithstanding of the fact that unsound analysis techniques can still be very useful in practice, especially if they identify problems in system designs. Yet, we should exercise great care in concluding anything from unsound techniques that have not found a problem. As has been aptly phrased by Dijkstra [Dij70]: "Program testing can be used to show the presence of bugs, but never to show their absence!"

4 First-Order Logic

Even though this course primarily studied extensions of first-order logic by dynamic modalities for hybrid systems instead of pure first-order logic, the sequent proof rules of propositional logic and quantifiers (instantiation and Skolemization) give a suitable proof calculus for first-order logic. And this suitability of the proof calculus for first-order logic is a much stronger statement than soundness.

Soundness is the question whether all provable formulas are valid and is a minimal requirement for proper logics. Completeness studies the converse question whether all valid formulas are provable.

The first-order logic proof calculus can be shown to be both sound and complete, which is a result that originates from Gödel's PhD thesis [Göd30], albeit in a different form.

Theorem 2 (Soundness & completeness of first-order logic)**.** *First-order logic is* sound, *i.e.* $\vdash\ \subseteq\ \vDash$, *which means that* $\vdash \phi$ *implies* $\vDash \phi$ *for all first-order formulas* ϕ *(all provable formulas are valid). First-order logic is* complete, *i.e.* $\vDash\ \subseteq\ \vdash$, *which means that* $\vDash \phi$ *implies* $\vdash \phi$ *for all first-order formulas* ϕ *(all valid formulas are provable). In particular, the provability relation* $\vdash$ *and the validity relation* $\vDash$ *coincide for first-order logic:* $\vdash\ =\ \vDash$*. The same holds in the presence of a set of assumptions* Γ*, i.e.* $\Gamma \vdash \phi$ *iff* $\Gamma \vDash \phi$*.*

This lecture will not set out for a direct proof of this result, because the techniques used for those proofs are interesting but would lead us too far astray. An indirect justification for what makes first-order logic so special that Theorem 2 can hold will be discussed later.

The following central result about compactness of first-order logic is of similar importance. Compactness is involved in most proofs of Theorem 2, but also easily follows from Theorem 2.

Theorem 3 (Compactness of first-order logic)**.** *First-order logic is* compact, *i.e.*

$$\Gamma \vDash A \iff E \vDash A \text{ for some finite } E \subseteq \Gamma \tag{1}$$

Proof. By Theorem 2, $\vdash\ =\ \vDash$. By completeness, semantic compactness theorem (1) is equivalent to the syntactic compactness theorem:

$$\Gamma \vdash A \iff E \vdash A \text{ for some finite } E \subseteq \Gamma \tag{2}$$

Condition (2) is obvious, because provability implies that there is a proof, which can, by definition, only use finitely many assumptions $E \subseteq \Gamma$. □

Compactness is equivalent to the finiteness property, which, for that reason, is usually simply referred to as compactness:

Corollary 4 (Finiteness)**.** *First-order logic satisfies the* finiteness property, *i.e.*

$$\Gamma \text{ has a model} \iff \text{all finite } E \subseteq \Gamma \text{ have a model} \tag{3}$$

Proof. Compactness (Theorem 3) implies the finiteness property. The key observation is that Γ has no model iff $\Gamma \vDash \mathit{false}$, because if Γ has no model, then false holds in all models of Γ of which there are none. Conversely, the only chance for false to hold in all models of Γ is if there are no such models, since false never holds. By Theorem 3,

$$\Gamma \vDash \mathit{false} \iff \exists \text{finite } E \subseteq \Gamma \quad E \vDash \mathit{false}$$

Hence,

$$\Gamma \text{ has a model} \iff \Gamma \nvDash \mathit{false} \iff \forall \text{finite } E \subseteq \Gamma \quad E \nvDash \mathit{false} \iff \text{all finite } E \subseteq \Gamma \text{ have a model}$$

It is worth noting that, conversely, the finiteness property implies compactness.

$$\begin{aligned} \Gamma \vDash A &\iff \Gamma \cup \{\neg A\} \text{ has no model} \\ &\iff \text{some finite } E \subseteq \Gamma \cup \{\neg A\} \text{ has no model} && \text{by finiteness} \\ &\iff E \vDash A \text{ for some finite } E \subseteq \Gamma \end{aligned}$$

The last equivalence uses that we might as well include $\neg A$ in E, because if E has no model then neither does $E \cup \{\neg A\}$. ☐

5 Skolem-Herbrand-Löwenheim Theory

The value of a logical formula is subject to interpretation in the semantics of the logic. In a certain sense maybe the most naïve interpretation of first-order logic interprets all terms as themselves. Such an interpretation I is called *Herbrand model*. It stubbornly interprets a term $f(g(a), h(b))$ in the logic as itself: $[\![f(g(a), h(b))]\!]_I = f(g(a), h(b))$. And likewise for all other ground terms.

That may sound like a surprising and stubborn interpretation. But, even more surprisingly, it is not at all an uninsightful one, at least for first-order logic. So insightful, that it even deserves a name: Herbrand models. Certainly, it is one of the many permitted interpretations.

Definition 5 (Herbrand Model)**.** An interpretation I is called *Herbrand model* if it has the free semantics for ground terms, i.e.:

1. The domain D is the ground terms (i.e. terms without variables) $\mathrm{Trm}^0(\Sigma)$ over Σ
2. $I(f) : D^n \to D; (t_1, \ldots, t_n) \mapsto f(t_1, \ldots, t_n)$ for each function symbol f of arity n

Let Γ be a set of closed universal formulas. $\mathrm{Trm}^0(\Sigma)(\Gamma)$ is the set of all ground term instances of the formulas in Γ, i.e. with (all possible) ground terms in $\mathrm{Trm}^0(\Sigma)$ instantiated for the variables of the universal quantifier prefix.

$$\mathrm{Trm}^0(\Sigma)(\Gamma) = \{\phi(t_1, t_2, \dots, t_n) \;:\; (\forall x_1 \forall x_2 \dots \forall x_n\, \phi(x_1, x_2, \dots, x_n)) \in \Gamma \\ t_1, \dots, t_n \in \mathrm{Trm}^0(\Sigma), \text{ for any } n \in \mathbb{N}\}$$

That is, for any $n \in \mathbb{N}$ and for any formula

$$\forall x_1 \forall x_2 \dots \forall x_n\, \phi(x_1, x_2, \dots, x_n)$$

in Γ and for any ground terms $t_1, \dots, t_n \in \mathrm{Trm}^0(\Sigma)$, the set $\mathrm{Trm}^0(\Sigma)(\Gamma)$ contains the following ground instance of ϕ:

$$\phi(t_1, t_2, \dots, t_n)$$

Theorem 6 (Herbrand [Her30])**.** *Let Γ be a (suitable) set of first-order formulas (i.e. closed* universal formulas *without equality and with signature Σ having at least one constant).*

$$\Gamma \text{ has a model} \iff \Gamma \text{ has a Herbrand model} \\ \iff \text{ground term instances } \mathrm{Trm}^0(\Sigma)(\Gamma) \text{ of } \Gamma \text{ have a model}$$

Using the Herbrand theorem twice gives:

Γ has a model $\iff$ ground term instances $\mathrm{Trm}^0(\Sigma)(\Gamma)$ of Γ have a Herbrand model

Corollary 7. *Validity in first-order logic is semidecidable.*

Proof. For suitable first-order formulas F (i.e. $\neg F$ satisfies the assumptions of Theorem 6), semidecidability follows from the following reductions:

F valid	$\iff$ $\neg F$ unsatisfiable	
	$\iff$ $\mathrm{Trm}^0(\Sigma)(\neg F)$ have no model	by Theorem 6
	$\iff$ some finite subset of $\mathrm{Trm}^0(\Sigma)(\neg F)$ has no Herbrand model	by Corollary 4

Thus, it remains to consider the assumptions in Theorem 6 whether first-order formulas that are not suitable can be turned into formulas that are suitable. First of all, Σ can be assumed without loss of generality to have at least one constant symbol for, otherwise, a constant can be added to Σ without changing validity of F. Furthermore, a formula F is valid iff its universal closure is, where the universal closure of a formula F is obtained by prefixing F with universal quantifiers $\forall x$ for each variable x that occurs free in F. Finally, existential quantifiers in first-order formula $\neg F$ can be removed without affecting satisfiability by Skolemization, which introduces new function symbols much like the quantifier proof rules from Lecture 6 did. □

Note 10 (Limitations of Herbrand models)**.** *Herbrand models are not the cure for everything in first-order logic, because they unwittingly forget about the intimate relationship of the term $2+5$ to the term $5+2$ and, for that matter, to the term $8-1$. All those terms ought to denote the same identical object, but end up denoting different ground terms in Herbrand models. In particular, a Herbrand model would not mind at all if a unary predicate p would hold of $2+5$ but not hold for $5+2$ even though both ought to denote the same object. Thus, Herbrand models are a little weak in arithmetic, but otherwise incredibly powerful.*

Herbrand's theorem has a second form with a close resemblance to the core arguments of quantifier elimination in first order logic of real arithmetic from Lecture 18 and Lecture 19.

Theorem 8 (Herbrand's theorem: Herbrand disjunctions [Her30])**.** *For a quantifier-free formula $\phi(x)$ of a free variable x without equality*

$$\exists x\,\phi(x) \text{ valid} \iff \phi(t_1)\vee\dots\vee\phi(t_n) \text{ valid for some } n\in\mathbb{N} \text{ and ground terms } t_1,\dots,t_n$$

Proof. The proof follows directly from Theorem 6 and Corollary 4:

$$\begin{aligned}
&\exists x\,\phi(x) \text{ valid} && \\
&\iff \neg\exists x\,\phi(x) \text{ unsatisfiable} && \\
&\iff \forall x\,\neg\phi(x) \text{ has no model} && \\
&\iff \mathrm{Trm}^0(\Sigma)(\forall x\,\neg\phi(x)) \text{ has no model} && \text{by Theorem 6}\\
&\iff \{\neg\phi(t)\ :\ t \text{ ground term}\} \text{ has no model} && \text{by definition}\\
&\iff \{\neg\phi(t_1),\dots,\neg\phi(t_n)\} \text{ has no model for some } t_1,\dots,t_n \text{ and some } n && \text{by Corollary 4}\\
&\iff \neg\phi(t_1)\wedge\dots\wedge\neg\phi(t_n) \text{ has no model for some } n \text{ and some } t_1,\dots,t_n && \\
&\iff \phi(t_1)\vee\dots\vee\phi(t_n) \text{ valid for some } n \text{ and some } t_1,\dots,t_n && \square
\end{aligned}$$

Theorem 8 holds for first-order formulas with multiple existential quantifiers. More general forms of the Herbrand theorem hold for arbitrary first-order formulas that are not in the specific form assumed above [Her30].

These more general Herbrand theorems won't be necessary for us, because, for validity purposes, first-order formulas can be turned into the form $\exists x_1\ \dots\exists x_n\,\phi(x_1,\dots,x_n)$ with quantifier-free $\phi(x_1,\dots,x_n)$ by introducing new function symbols for the universal quantifiers using essentially the quantifier proof rules from Lecture 6:[2]

[2] The new function symbols are usually called Skolem functions and the process called Skolemization, because Thoralf Skolem introduced them in the first correct proof of the Skolem-Löwenheim theorem [Sko20]. Strictly speaking, however, Herbrand functions and Herbrandization are the more adequate names, because Jacques Herbrand introduced this dual notion for the first proof of the Herbrand theorem [Her30]. Skolemization and Herbrandization are duals. Skolemization preserves satisfiability while Herbrandization preserves validity.

$$(\forall r)\ \frac{\Gamma \vdash \phi(s(X_1,..,X_n)), \Delta}{\Gamma \vdash \forall x\, \phi(x), \Delta}{}^1 \qquad (\exists l)\ \frac{\Gamma, \phi(s(X_1,..,X_n)) \vdash \Delta}{\Gamma, \exists x\, \phi(x) \vdash \Delta}{}^1$$

[1] s is a new (Skolem-Herbrand) function and $X_1,..,X_n$ are all (existential) free logical variables of $\forall x\, \phi(x)$.

The clou about quantifier rules ∀r,∃l is that they preserve validity. By soundness, if their premiss is valid then so is their conclusion. Yet, in the case of rules ∀r,∃l the converse actually holds as well. If their conclusion is valid then so is their premiss. For rule ∀r, for example, the conclusion says that $\phi(x)$ holds for all values of x in all interpretations where Γ holds and Δ does not. Consequently, in those interpretations, $\phi(s(X_1,..,X_n))$ holds whatever the interpretation of s is, because s is a fresh function symbol, which, thus, does not appear in Γ, Δ.

Lemma 9 (Herbrandization)**.** *With each first-order logic formula ψ, a formula*

$$\exists x_1 \ldots \exists x_n\, \phi(x_1,\ldots,x_n)$$

with quantifier-free $\phi(x_1,\ldots,x_n)$ can be associated effectively that is valid if and only if ψ is. The formula $\exists x_1 \ldots \exists x_n\, \phi(x_1,\ldots,x_n)$ uses additional function symbols that do not occur in ψ.

Theorem 8 enables a second, more straightforward proof of the semidecidability of the validity problem of first-order logic:

Proof of Corollary 7. The semidecision procedure for validity of first-order logic formulas ψ proceeds as follows:

1. Herbrandize ψ to obtain a formula $\exists x_1 \ldots \exists x_n\, \phi(x_1,\ldots,x_n)$ by Lemma 9, which preserves validity.
2. Enumerate all $m \in \mathbb{N}$ and all ground terms t_i^j $(1 \le j \le n, 1 \le i \le m)$, over the new signature.
   a) If the *propositional* formula

   $$\phi(t_1^1,\ldots,t_1^n) \lor \cdots \lor \phi(t_m^1,\ldots,t_m^n)$$

   is valid, then so is $\exists x_1 \ldots \exists x_n\, \phi(x_1,\ldots,x_n)$ and, hence, ψ is valid.

By Theorem 8 and Lemma 9, the procedure terminates for all valid first-order formulas. □

The procedure in this proof will always succeed but it enumerates the ground terms for instantiation rather blindly, which can cause for quite a bit of waiting. Nevertheless, refinements of this idea lead to very successful automated theorem proving techniques for first-order logic known as *instance-based methods* [BT10], which restrict the instantiation to instantiation-on-demand in various ways to make the procedure more goal-directed. There are also many successful automatic theorem proving procedures for first-order logic that are based on different principles, including tableaux and resolution [Fit96].

6 Back to CPS

First-order logic is beautiful, elegant, expressive, and simple. Unfortunately, however, it is not expressive enough for hybrid systems [Pla10a, Pla12b, Pla13]. As soon as we come back to studying hybrid systems, the situation gets more difficult. And that is not by accident, but, instead, a fundamental property of first-order logic and of hybrid systems. Per Lindström characterized first-order logic in a way that limits which properties stronger logics could possess [Lin69]. Hybrid systems themselves are also known not to be semidecidable.

Given that differential dynamic logic talks about properties of hybrid systems, and Turing machines are a special case, undecidability is not surprising. We show a very simple standalone proof of incompleteness by adapting a proof for programs, e.g., [Pla10c].

Theorem 10 (Incompactness)**.** *Differential dynamic logic is not compact.*

Proof. It is easy to see that there is a set of formulas that has no model even though all finite subsets have a model, consider:

$$\{\langle(x := x+1)^*\rangle x > y\} \cup \{\neg(x+n > y) \ : \ n \in \mathbb{N}\} \qquad \square$$

Hence, differential dynamic logic does not have the finiteness property, which is equivalent to compactness (Corollary 4).

Since soundness and completeness imply compactness (see proof of Theorem 3), incompactness implies incompleteness[3], because $\mathsf{d}\mathcal{L}$ is sound. An explicit proof is as follows:

Theorem 11 (Incompleteness [Pla08])**.** *Differential dynamic logic has no effective sound and complete calculus.*

Proof. Suppose there was an effective sound and complete calculus for $\mathsf{d}\mathcal{L}$. Consider a set Γ of formulas that has no model in which all finite subsets have a model, which exists by Theorem 10. Then $\Gamma \vDash 0 > 1$ is valid, thus provable by completeness. But since the proof is effective, it can only use finitely many assumptions $E \subset \Gamma$. Thus $E \vDash 0 > 1$ by soundness. But then the finite set E has no model, which is a contradiction. $\square$

[3] Strictly speaking, incompleteness only follows for effective calculi. *Relative* soundness and completeness can still be proved for $\mathsf{d}\mathcal{L}$ [Pla08, Pla10a, Pla12b], which gives very insightful characterizations of the challenges and complexities of hybrid systems.

Having said these negative (but necessary) results about differential dynamic logic (and, by classical arguments, any other approach for hybrid systems), let's return to the surprisingly amazing positive properties that differential dynamic logic possesses.

For one thing, the basis of differential dynamic logic is the first-order logic of real arithmetic, not arbitrary first-order logic. This enables a particularly pleasant form of Herbrand disjunctions resulting from quantifier elimination in real arithmetic (recall Lecture 18 and Lecture 19).

Definition 12 (Quantifier elimination)**.** A first-order theory admits *quantifier elimination* if, with each formula ϕ, a quantifier-free formula $\mathrm{QE}(\phi)$ can be associated effectively that is equivalent, i.e. $\phi \leftrightarrow \mathrm{QE}(\phi)$ is valid (in that theory).

Theorem 13 (Tarski [Tar51])**.** *The first-order logic of real arithmetic admits quantifier elimination and is, thus, decidable.*

Also recall from Lecture 18 and Lecture 19 that the quantifier-free formula $\mathrm{QE}(\phi)$ is constructed by substitution or virtual substitution from ϕ, with some side constraints on the parameter relations. The quantifier-elimination instantiations are more useful than Theorem 8, because the required terms for instantiation can be computed effectively and the equivalence holds whether or not the original formula ϕ was valid. This makes it possible to use the proof calculus of differential dynamic logic to synthesize constraints on the parameters to make an intended conjecture valid [Pla10a].

Exercises

Exercise 1. The arguments for incompleteness and incompactness of $\mathsf{d}\mathcal{L}$ hardly depend on $\mathsf{d}\mathcal{L}$, but, rather, only on $\mathsf{d}\mathcal{L}$'s ability to characterize natural numbers. Incompleteness and incompactness hold for other logics that characterize natural numbers due to a famous result of Gödel [Göd31]. Both the discrete and the continuous fragment of $\mathsf{d}\mathcal{L}$ can characterize the natural numbers [Pla08].

1. Show that the natural numbers can be characterized in the discrete fragment of $\mathsf{d}\mathcal{L}$, i.e. only using assignments and repetition.
2. Then go on to show that the natural numbers can also be characterized in the continuous fragment of $\mathsf{d}\mathcal{L}$, i.e. using only differential equations.
3. Conclude from this that both the discrete and the continuous fragment of $\mathsf{d}\mathcal{L}$ are not compact, nor is any other logic that can characterize the natural numbers.

References

[And02] Peter B. Andrews. *An Introduction to Mathematical Logic and Type Theory: To Truth Through Proof*. Kluwer, 2nd edition, 2002.

[BT10] Peter Baumgartner and Evgenij Thorstensen. Instance based methods - a brief overview. *KI*, 24(1):35–42, 2010.

[Col07] Pieter Collins. Optimal semicomputable approximations to reachable and invariant sets. *Theory Comput. Syst.*, 41(1):33–48, 2007. doi:10.1007/s00224-006-1338-3.

[DBL12] *Proceedings of the 27th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2012, Dubrovnik, Croatia, June 25–28, 2012*. IEEE, 2012.

[Dij70] Edsger Wybe Dijkstra. Structured programming. In John Buxton and Brian Randell, editors, *Software Engineering Techniques. NATO Software Engineering Conference 1969*. NATO Scientific Committee, 1970.

[Fit96] Melvin Fitting. *First-Order Logic and Automated Theorem Proving*. Springer, New York, 2nd edition, 1996.

[Göd30] Kurt Gödel. Die Vollständigkeit der Axiome des logischen Funktionenkalküls. *Mon. hefte Math. Phys.*, 37:349–360, 1930.

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

[Her30] Jacques Herbrand. Recherches sur la théorie de la démonstration. *Travaux de la Société des Sciences et des Lettres de Varsovie, Class III, Sciences Mathématiques et Physiques*, 33:33–160, 1930.

[Lin69] Per Lindström. On extensions of elementary logic. *Theoria*, 35:1–11, 1969. doi:10.1111/j.1755-2567.1969.tb00356.x.

[PC07] André Platzer and Edmund M. Clarke. The image computation problem in hybrid systems model checking. In Alberto Bemporad, Antonio Bicchi, and Giorgio Buttazzo, editors, *HSCC*, volume 4416 of *LNCS*, pages 473–486. Springer, 2007. doi:10.1007/978-3-540-71493-4_37.

[Pla07] 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.

[Pla08] André Platzer. Differential dynamic logic for hybrid systems. *J. Autom. Reas.*, 41(2):143–189, 2008. doi:10.1007/s10817-008-9103-8.

[Pla10a] André Platzer. *Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics*. Springer, Heidelberg, 2010. doi:10.1007/978-3-642-14509-4.

[Pla10b] 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.

[Pla10c] André Platzer. Theory of dynamic logic. Lecture Notes 15-816 Modal Logic, Carnegie Mellon University, 2010. URL: http://www.cs.cmu.edu/~fp/courses/15816-s10/lectures/25-DLtheo.pdf.

[Pla11] André Platzer. Stochastic differential dynamic logic for stochastic hybrid programs. In Nikolaj Bjørner and Viorica Sofronie-Stokkermans, editors, *CADE*, volume 6803 of *LNCS*, pages 431–445. Springer, 2011. doi:10.1007/978-3-642-22438-6_34.

[Pla12a] André Platzer. A complete axiomatization of quantified differential dynamic logic for distributed hybrid systems. *Logical Methods in Computer Science*, 8(4):1–44, 2012. Special issue for selected papers from CSL'10. doi:10.2168/LMCS-8(4:17)2012.

[Pla12b] André Platzer. The complete proof theory of hybrid systems. In *LICS* [DBL12], pages 541–550. doi:10.1109/LICS.2012.64.

[Pla12c] André Platzer. Logics of dynamical systems. In *LICS* [DBL12], pages 13–24. doi:10.1109/LICS.2012.13.

[Pla13] André Platzer. A complete axiomatization of differential game logic for hybrid games. Technical Report CMU-CS-13-100R, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, January, Revised and extended in July 2013.

[Sko20] Thoralf Skolem. Logisch-kombinatorische Untersuchungen über die Erfüllbarkeit oder Beweisbarkeit mathematischer Sätze nebst einem Theorem über dichte Mengen. *Videnskapsselskapet Skrifter, I. Matematisk-naturvidenskabelig Klasse*, 6:1–36, 1920.

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