Carl von Ossietzky
Universität
Oldenburg

Carl von Ossietzky Universität Oldenburg / 26111 Oldenburg

Faculty of Information Technology
Czech Technical University in Prague

Fakultät II - Informatik,
Wirtschafts- und
Rechtswissenschaften
Department für Informatik

Abteilung
Foundations and Applications of Systems of Cyber-Physical Systems

Prof. Dr. Martin Fränzle

Telefondurchwahl
Tel.: 0441 9722 - 500
Fax: 0441 9722 - 502
fraenzle@informatik.uni-oldenburg.de

Sekretariat
Kathrin Kuper
Tel.: 0441 9722 - 501
kathrin.kuper@uol.de

Oldenburg, den April 19, 2024

Standort
OFFIS e.V.
Escherweg 2
26121 Oldenburg

Postanschrift
26111 Oldenburg

Paketanschrift
Ammerländer Heerstraße 114–118
26129 Oldenburg

Bankverbindung
Landessparkasse zu Oldenburg
IBAN DE46 2805 0100 0001 9881 12
BIC SLZODE22

Steuernummer
6422008701

www.uol.de

Review of the Doctoral Dissertation Thesis “Boolean Satisfiability Modulo Differential Equation Simulation” as submitted by Ing. Tomáš Kolárik in January 2024 to the Faculty of Information Technology, Czech Technical University in Prague, within its Doctoral Study Programme Informatics

Cyber-physical systems, i.e. systems that tightly integrate physical objects with information technology to alter the dynamic behaviour of the physical objects, form the backbone technology of „smart“ environments and are thus at the heart of a current industrial revolution. Central to their function is that they interleave discrete decisions, often safety-critical or mission-critical ones, with piecewise continuous behaviour mediated by the laws of nature pertinent to the physical objects or induced by continuous feedback control. With the analysis of non-linear continuous behaviour on the one hand and of the dynamics of digital electronics and embedded programs on the other hand both being very hard and only partially solved problems, research addressing the even more complex behavioural analysis and planning in the hybrid discrete-continuous domain originating from their combination always is welcome. Tomáš Kolárik’s thesis addresses this issue in an as original as readable way, which is very laudable. It is my pleasure to write a review for this timely and well-written thesis.

Structure and organization of the thesis

The thesis is written in an excellent and, modulo a minor number of typos, flawless English, conveying the author’s ideas clearly through a sequence of 10 clearly defined chapters covering 154 pages. These are accompanied by the usual front matter (abstract, acknowledgements, table of contents, lists of figures, of tables, of algorithms, of mathematical symbols, and of acronyms) plus an extensive bibliography featuring 122 topical entries and the lists of own publications (3 reviewed and presented at pertinent conference, 8 preprints or online resources) as back matter. Chapter 1 provides a general introduction and explains contributions and structure of the thesis, while chapters 2 and 3 provide the general theoretical and methodological background and an overview over existing algorithms, both “competing” algorithms and such used as a basis for or as a component of the own contribution. Chapters 4 to 6 then develop the own technology “SAT Modulo Differential Equation Simulation” (abbreviated to SMDES in the remainder of this review) progressively by first providing an abstract description of the concept, then explaining related tools, and finally describing the own implementation. Chapters 7 and 8 showcase and evaluate the resulting technology by means of verification or falsification case studies from the hybrid-state control domain (chapter 7) as well as via hybrid-state planning problems (railway scheduling, chapter 8). Chapter 9 stands somewhat isolated in that it briefly describes and then evaluates a different solver technology not being based on SMDES, but being similar in spirit by deferring some of the internal mathematical theory reasoning to numerical algorithms. Chapter 10, finally, concludes the thesis with a retrospective discussion of the major contributions and a very brief list of areas for future research.

Relevance of the selected topic and aims of the thesis

As said in the preamble, the general theme of rigorous and exhaustive analysis of hybrid discrete-continuous dynamic behaviour is of high scientific as well as economic relevance and thus absolutely timely. Given our societal reliance on safety-critical cyber-physical systems in virtually every domain, be it health technologies like implanted and autonomously acting medical devices or smart transportation involving automated driving functions or autonomous drones, we are in urgent need of rigorous methods for certifying or constructively establishing their functional correctness and optimality. These two forms of analysis of hybrid discrete-continuous dynamic behaviour, namely verification delivering certificates or synthesis shaping desirable behaviour, have exhaustively been discussed in the scientific community and there is broad, though not universal, consensus that they should be

1. exhaustive, covering all possible behaviours of the generally open systems operating in an only partially known context,
2. mathematically exact, providing verdicts that match a formal mathematical semantics of the objects of discourse,
3. automatic, relieving system designers from tedious and often hardly understood manual or interactive verification efforts, and
4. scalable, covering systems or at least independently analysable subsystems of industrially relevant size.

The combination of these four properties has, however, hitherto remained elusive. While most existing work - with the notable exception of Bjørnar Lutberget’s PhD thesis, which shares many ideas with the one reviewed here, albeit in a considerably more specialized and thus restricted setting - has compromised on items 3 or 4 or both, i.e. on the level of automation of the reasoning or on its scalability beyond academic case studies, Tomáš Kolárik’s work rather compromises on the first two and does so in a logically and technologically well-argued way: observing that most model-based design is fine with using approximate numerical simulation for building digital twins of the system under design, Kolárik opts for replacing costly exact solving of differential equations by approximate numerical solving in the reachability analysis of hybrid discrete-continuous systems, while at the same time preserving exhaustive analysis of the discrete state components and advancing scalability through their strictly symbolic treatment. While such ideas have been contemplated repeatedly, few have been developed with the same logical stringency as Kolárik’s approach, which delivers a logical mostly – for the remaining problematic cases see below – consistent embedding of such numerical solving into satisfiability-modulo-theory (SMT) solving.

The topic of the thesis consequently is without doubt scientifically very relevant and Kolárik's approach highly original, constituting a PhD-worthy contribution to the state of research.

Assessment of the methods used and evaluation of results

Based on an extensive and sound evaluation of the state of the art, Kolárik selects and adopts established and proven methods whenever justified, but combines them in an original way to overcome the central problems he identified in the state of the art concerning scalability and performance. In particular, he builds on satisfiability (SAT) and satisfiability-modulo-theory (SMT) solving as well as numerical integration of ordinary differential equations (ODE), with their seamless integration into a logically consistent and technically beneficial algorithm being his own, mostly unprecedented line of attack. Given the much better scalability of point-based numerical integration of ODE compared to set-based exhaustive reach-set computation for ODE, the resulting algorithm provides substantial performance gains concerning both scalability in the analysable state-space dimension and computational runtimes. It ought to be noted that the scalability increase thereby not only pertains to the dimensionality of the continuous state-space that can be treated, but also to the discrete state component due to its symbolic representation (an analogous encoding of discrete state has already been used by Egger's isat(ODE), yet not by Gao's dReach, with the latter – yet not the former – serving as a point of reference throughout the thesis). These performance gains are theoretically well-argued, and they are also demonstrated rigorously throughout the thesis by thorough benchmarking against Gao's dReach tool, a tool also integrating ODE solving, albeit in the form of verified set-based integration rather than point-based numerical integration, as a theory solver into SMT.

I must admit that I was at first glance sceptical about the scope of Kolárik's approach, as it comes at the price of enforcing point-valued initial-value problems for all ODE solving involved, meaning that one cannot formulate and prove statements about ensembles or families of continuous trajectories, which is the objective of most, if not all other hybrid-system verification tools. That SMDES delivers much more than a single initial-value based numerical simulation – as available in many free or commercial standard tools like Simulink – does provide is not immediately evident; it requires the case study from Chapter 8 to fully understand its immense potential for hybrid-state manoeuvre planning which any single-trajectory simulation – like that in Simulink – naturally lacks. It might have been wise to frontload this case study description to render the actual scope of SMDES immediately obvious.

Once having understood this critical point and having thus correctly located SMDES in the spectrum between trajectory-based numeric simulation and set-based exhaustive verification, where it represents a meet-in-the-middle approach with the associated compromises between scope of its verdicts and computational performance, the experimental cross-validation against the exhaustive analysis by dReach no longer appears as an apples-to-oranges comparison. It then qualifies as a reasoned experiment in trading exhaustiveness (as provided by dReach, yet not by the sampling of initial states necessitated by SMDES) and mathematical exactness (dReach using verified set-based integration while SMDES uses approximate numerical integration) against computational effort (numerical integration providing an orders of magnitude less demanding algorithm) and thus scalability. The corresponding findings are communicated clearly and interpreted thoroughly, though adding some focus on cases where the two methods provide inconsistent verdicts – which certainly must have come up during the work – would have been desirable.

Objections and questions for the defence

While the author tries to establish a rigorous logical interpretation of numeric ODE solutions as arithmetic theories and thereby to the embedding of numeric interpretations of ODE into theory-related logical formulae as well as of numeric ODE solving into the satisfiability solving process of such formulae (“SAT Modulo Differential Equation Simulation”, SMDES), the result does not seem logically completely consistent and contains various pragmatic design decisions that ought have to be argued and discussed.

The major problem concerning logical consistency concerns lacking referential transparency in the semantics of ODE atoms: at least if one admits variable step-size solving (which the current implementation does not, but which is announced as a straightforward and logically completely independent extension), there is extralogical bias in the definition of the result of the numerical solving that destroys referential transparency, giving context-dependent variations in the meaning of ODE atoms. Given the same initial value for x and the same time-span, the ODE dx/dt = f(x) should have a fixed meaning that is the same for all its occurrences. This, however, is not true here: when evaluated in a situation where the underlying SAT solver has also instantiated the propositional atom guarding another ODE dy/dt = g(y) then dx/dt = f(x) will get a different semantics due to the different choices in step widths even when the two ODE are completely independent in that f and g do not depend on common variables. Whether this just constitutes a weakness in the logical foundations or may even impact soundness of the implemented algorithm I cannot judge currently – personally, I have seen many CDCL-based SMT-solving systems stumble across such inconsistencies in the theory solving, as they invalidate assumed invariants about the consistency between learned conflict clauses and subsequent queries to the theory solver, with faulty inferences materializeing across backjumps especially.

Another question of pragmatic rather than logical nature pertains to the choice of only initial-value problems and numerical temporally forward solving, given that the very same mathematical and algorithmic tools do also permit final-value problems and backward solving, which may come handy especially in robotics planning problems, like the railway scheduling or the mutli-agent planning problems discussed in the thesis. The restriction to forward methods therefore needs justification.

Overall evaluation of the thesis

The overall approach of SMDES developed by the author of the thesis is highly original. The design objectives underlying the suggested technology have reasonably been argued from an analysis of the state of the art and have, to the level possible within thesis work, rigorously been proven to have been achieved. With his thesis, Tomáš Kolárik has tangibly contributed to the state of science and engineering in a relevant field. His method selections are consistent with the scientific standards in the field concerning all the dimensions of identifying need for research and detailed research questions, identifying scientific background, developing and implementing his own method, and evaluating its impact. The scientific approach and results and their representation in the well-written thesis are without any doubt PhD-worthy and deserve broader visibility through a series of subsequent scientific publications accompanying the three existing ones.

Recommendation

As the author of this thesis, Tomáš Kolárik proved the ability to conduct independent research and achieve scientific results of outstanding quality. In accordance with § 47(4) of the Higher Education Act, i.e. of law nr. 111/1998 in the form communicated to me, I recommend acceptance of the thesis for presentation, defence, and publication with the aim of obtaining a doctoral degree in Informatics.

Oldenburg, April 19, 2024

(Martin Fränzle)