Review of the Doctoral Thesis:
"Boolean Satisfiability Modulo Differential Equation Simulations"
submitted by
Ing. Tomáš Kolárik

1 Up-to-dateness of the dissertation

The dissertation builds on the framework of Satisfiability Modulo Theories (SMT) which has become one of key technologies for computer-aided verification over the last 15–20 years, and the integration of numerical Ordinary Differential Equations (ODEs) into SMT has not been considered before. The up-to-dateness of the work is supported also by the comprehensive and up-to-date list of references.

2 Formal structure and organization of the dissertation

The thesis consists of ten chapters and 122 references (167 + xvii pages in total).

Chapter 1 is an introduction which summarizes the whole dissertation.

Chapter 2 is a somewhat quick survey of theoretical backgrounds which introduces ODEs, hybrid automata (HA), Boolean satisfiability (SAT), SMT, and bounded model checking (BMC).

Chapter 3 is an algorithmic counterpart of Chapter 2, which surveys numerical solving of ODEs, SAT solving (quick), and SMT solving.

Chapter 4 describes the core framework of the thesis, i.e., SAT modulo differential equation simulations. It introduces "Simulation Semantics" whose time axis is discrete, whose range is the set of floating-point numbers, and which is parameterized w.r.t. a particular numerical ODE solver.

Definition 16 (floating-point interpretation) is a key notion of the author's simulation semantics, where the final value of the interval being considered need not satisfy the invariant (e.g., the height of a bouncing ball on a floor may be negative at the final time-point). The author then introduces the notion of strong satisfiability (Definition 20) to adapt the SMT framework to the simulation semantics.

Chapter 5 is a practical counterpart of Chapter 2, i.e., it is a survey of the tools that implement the theoretical frameworks introduced in Chapter 2. In particular, it introduces (i) relevant tools including SAT modulo ODE solvers, and (ii) simulation tools for hybrid systems.

Chapter 6 is a detailed description of the SAT modulo ODE solver the author has developed based on the semantics described in Chapter 4. It first describes detailed solver algorithms (Sections 6.1 and 6.2). Then it describes the concrete syntax of the input language of the solver (Section 6.3). Finally, it comments on the open-source implementation called "UN/SAT modulo ODEs Not SOT (UN/SOT)" (Section 6.4).

Chapter 7 describes four case studies, of which the full description of the first one, railway scheduling, is left to Chapter 8. Other case studies incorporate comparison with dReal, a tool based on SMT that implements the notion of δ-satisfiability.

Chapter 8 is a detailed description of the railway scheduling problem, its SMT encoding, and experimental results.

Chapter 9 describes another class of problems, i.e., multi-agent path finding. The subject is actively studied in the multi-agent systems community, but the author studies it in a timed setting for collision detection and avoidance (Section 9.5) incorporated into the SMT framework.

Chapter 10 concludes the thesis with a list of future work.

The style of writing is dense, but the concepts and notations introduced are all defined and explained in detail.

3 Completion of the dissertation objectives

The objectives of the dissertation consist of

1. the proposal of the theoretical framework,
2. algorithms for solving the problems defined using the proposed framework,
3. implementation of the algorithms, and
4. evaluation using several non-trivial benchmark problems.

I believe all of them are well achieved. In addition to the thesis itself, an open-source artifact was made publicly available.

4 Assessment of the methods used in the dissertation

The research methods of the dissertation is generally solid.

The thesis starts with a unique motivation and standpoint towards computer-aided verification, namely it bases the theoretical framework directly on numeric ODE solvers (rather than regarding them as an approximation of mathematical ODEs). This is an unusual approach to formal methods, but I think the author managed to convince the readers of the approach he/she proposes. Specifically, the approach handles correctness and optimization *modulo* errors introduced by the numerical method adopted; thus *abstracting* the details and intricacies of

ODE solvers. The claim that optimality of the solutions of planning problems are not crucial in practical settings sounds convincing.

On the other hand, I think the proposed technique could be used to guarantee some mathematical properties (e.g., properties based on mathematical semantics) of the systems without too much extensions, which would make the proposed ideas much more appealing. This would require two things: overestimation of errors introduced by numerical methods, and errors introduced by floating-point computation (the latter of which could be efficiently controlled by specifying rounding modes).

As for the evaluation of the proposed technique, comparison with dReal is done in Chapter 7, which provides some useful information. However, because the objective of dReal is to return a *guaranteed* answer based on a clearly defined notion of correctness, comparison on an equal basis is not straightforward and should be stated carefully. Also, it would be desirable to compare the *results obtained*, not just *execution time*. Are the results the same or different? If they are always the same, it might be because the problem and evaluation settings are not critical enough to highlight the differences of the approaches.

Chapter 8 discusses the railway scheduling problem for comprehensive evaluation of the proposed technique. The problem is interesting and nontrivial, especially in terms of the number of discrete states. From the railway system point of view, on the other hand, I (as a well-informed outsider of railway engineering) wonder if the benchmark instances are realistic or somewhat contrived. For example, real railway systems with heavy traffic come with double/multiple tracks, flat crossing and level junctions. Also, passenger requests (i.e., flow) cannot be neglected in evaluation. I think the SMT framework makes all those factors (much) easier to encode than dedicated tools. Addressing the generality of the framework would be useful in demonstrating the strength of the SMT approach.

Chapter 9 is devoted to another important class of problems, though it handles ODEs in a simplified way. The key technical contribution here in the SMT setting appears to be conflict generalization discussed in Section 9.5.

5 Evaluation of the results and contributions of the dissertation

As I stated in previous sections, I believe that the dissertation is successful in meeting the objectives as defined by the author. The results obtained contribute to the state of the art of the field. In addition, the work yielded a concrete SMT language for describing problems and an open-source tool for solving them.

One concern is that, although the summary of Chapter 9 states “collision detection and avoidance of the agents yields nonlinear constraints which we handle based on simulations (i.e., floating-point computation)" (both in Chapters 1 and 10), the handling of nonlinear constraints is not discussed explicitly in Chapter 9. (I searched the keywords "nonlinear" and "non-linear" throughout Chapter 9).

6 Remark, objections, notes, and questions for the defense

Possible questions for the defense are marked (*).

(General)

1. The description of the dissertation is careful and detailed in general, but is a bit dense; the text could be structured better by using more itemization, subsubsections, displayed math formulas, etc.

(Section 1.1)

1. (*) "classical mathematical semantics of differential equations often does not correspond well to the actual intended semantics" — could you elaborate on this in order to better convince the readers of this claim, possibly using motivating examples?

(Section 4.1)

1. What is a "type" in Definition 9 (page 36)?

2. Stages of functional flows could be related to *hybrid trajectories*, a standard notion in the theory of hybrid systems.

3. (*) Definition 16 (floating-point interpretation) is a key notion of the author's simulation semantics, but I am not sure if "rounding to nearest" is the only legitimate option here. Carefully controlling rounding modes (towards the safe direction) may possibly strengthen the power of the formalism and the tool.

(Section 4.2)

1. The difference between the notations x_i and $x^{[i]}$ was still not clear to me, and x_i was not defined in the List of Mathematical Symbols (page xv).

(Chapter 6)

1. "the result of our solver is `sat` if the formula is (strongly) satisfiable" (page 59) — it would be better to explicitly state that this notion of (strong) satisfiability is parameterized w.r.t. a particular numerical method employed. (This important point was once stated in Definition 16 but it cannot be overemphasized.)

2. Section 6.2.2 (12.5 pages long!) is detailed but not easy to follow. Examples would greatly help!

3. (*) "UN/SAT modulo ODEs Not SOT (UN/SOT)" — what does SOT stand for, and what does the whole phrase mean? (they are not explained anywhere!)

(Chapter 7)

1. (*) As pointed out in Section 4 of this review, it would be desirable to compare the *results obtained*, not just *execution time*. Are the results the same or different? If they are always the same, it might be because the problem and evaluation settings are not critical enough to highlight the differences of the approaches.

(Section 8.3)

1. one-hot ⟶ one-shot?

2. Any reason why the third and the fourth conjuncts of formula (8.7) are written in different styles?

(Section 8.4)

1. Evaluation is a bit too coarse (especially w.r.t. the bnd (bound) parameters which are different from each other by one or more orders of magnitude). More detailed evaluation would make the thesis stronger.

(Chapter 9)

1. Please include more figures for readability, particularly for MAPF problems.

2. (*) What was the effect of conflict generalization (learning)? Experimental results would be very useful.

3. Evaluation parameters (especially δ) look a bit too coarse.

4. (Section 9.6) Over-approximation of conflict intervals is an important idea and could be described in more detail with simple examples.

5. (*) As mentioned in Section 5 of this review, please clarify the handling of non-linearity. (Non-linearity is mentioned only in page 132, line 8 except for Section 9.8 on future plans.)

7 The overall evaluation of the dissertation

As discussed in Sections 3–5, I think the dissertation describes the topic in great detail and in a self-contained manner. Although there are some (rather minor) points to be addressed as listed in Section 6, I hope those points will be clarified in the thesis defense.

8 Statement whether you DO or DO NOT recommend the dissertation for the defense

To conclude, the author of the dissertation proved the ability to conduct research and achieve scientific results. In accordance with par. 47, letter (4) of the Law Nr. 111/1998 (The Higher Education Act) I do recommend the thesis for the presentation and defense with the aim of receiving the Ph.D. degree.

Tokyo, Japan, May 11, 2024

Kazunori Ueda

Prof. Kazunori Ueda, Dr. Eng.
Reviewer of the thesis